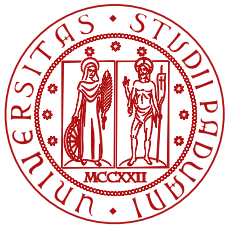

Università degli Studi di Padova

Dipartimento di Ingegneria dell'Informazione

Master's Degree in Electronic Engineering

Embedded system for in-situ measurement of sediment inertial dynamics in marine environments

Candidate

Marco Migliorini

N. 2087051

Supervisor

Prof. Alessandro Pozzebon

Academic Year
2023/2024

"O da la testa o da la coa,
tutti quanti fa la soa".
— Ai veri maestri

Contents

Abstract

In recent years, natural hazards such as coastal erosion, flooding, and landslides have become common events due to climate change and global warming. Understanding sediment transport is crucial for addressing these extreme events as well as to achieve an in-depth comprehension of physical processes underlying sediment transport. However, existing technologies are expensive, require frequent human-led testing, and lack long-term acquisition capabilities, making extremely complex their deployment in more challenging environments. Moreover, the unique obstacles presented by underwater (UW) settings, such as communication limitations, continuous thermal and physical stress, and the need for ultra-low power design due to off-grid operation, mean that many sediment movement parameters in aquatic environments remain completely unmonitored.

This work proposes a novel ultra-low power IoT (Internet of Things) system to perform sediment inertial dynamics measurements and tracking in adverse marine environments. Starting from this prototype, we aim to study a new generation of technologies for the in-situ sediment transport monitoring extendable to various aquatic settings. Furthermore, this project faces critical problems for ultra-low power IoT applications in harsh aquatic scenarios. By overcoming these challenges, we will deliver a significant technological advancement for the entire field of ultra-low power IoT applications operating in critical conditions.

The thesis explores the design and implementation of the initial prototype, with a deep focus on the firmware and hardware optimization for an ultra low power design.

1

Introduction

In recent years, our planet has undergone continuous transformation due to climate change. As a consequence, extreme phenomena such as coastal erosion, flooding and landslides are becoming increasingly frequent, threatening both economic activities and human safety.
Indeed, coastal and fluvial zones sustain a significant regional population and are vital for providing essential ecosystem services that support crucial socioeconomic activities. However, these zones face pressing challenges, with erosion being a prominent concern. In particular, coastal erosion not only compromises coastal stability but also leads to occasional loss of natural habitats due to the consistent shoreline retreat in specific regions. Short term adverse effects of coastal erosion are often attributed to human activities. Nonetheless, the influence of hydrodynamic processes, intensified by extreme events, is gaining importance. This is due to the combined impact of tides, waves, currents, and other factors, resulting in accelerated shoreline retreat and significant losses of infrastructure and beach areas.

Given the strict dependence of these events on sediment transport processes both along river courses and coastal areas, studying sediment transport becomes crucial for a better understanding of fluvial and coastal morphodynamics. By studying new technologies to monitor sediment dynamics, we can gain valuable insights into the interactions between natural forces and human activities, ultimately leading to more sustainable management of vulnerable coastal and fluvial zones.

1.1 Problem definition

Nowadays many measurement techniques are used to obtain information from the environment. Modern measurement technologies have significantly advanced our understanding of coastal and riverine environments, enabling detailed analysis at both macro and micro scales.
Satellite technologies, for instance, have proven effective for long-term, large-scale monitoring of fluvial and coastal areas, offering valuable insights into environmental changes over time [Ouillon et al. (2004), Haihong (2011), and Yangdong et al. (2016)]. Similarly, autonomous underwater vehicles (AUVs) and unmanned aerial vehicles (UAVs) have been deployed to collect in-situ data, providing detailed samples of specific locations [Gaël (2016), Jaud et al. (2016), and Pucino et al. (2021)]. However, despite these advancements, significant limitations remain when it comes to performing in-situ measurements of punctual sediment transport dynamics. Satellite technologies, while powerful for large-scale analysis, are limited in their ability to conduct detailed, localized studies of sediment movements. On the other hand, AUVs and UAVs, although capable of providing on-site data, require substantial human intervention for deployment, limiting their scalability and operational efficiency.
Even more robust methodologies, such as extensive bathymetric surveys, sediment sampling across large areas, and the deployment of multiple specialized instruments [Fitri et al. (2019)], require intensive human involvement. This not only increases costs but also limits the scalability and frequency of data collection, making it challenging to monitor sediment transport in a dynamic and continuous manner for prolonged periods.

Consequently, the technologies currently available to monitor sediment transport mechanisms and their environmental impact are either limited or non-existent in providing the continuous, detailed, and scalable data required for effective analysis. Thus, the development of new methodologies capable of acquiring such information directly in the environment would revolutionize the study of these geophysical processes.

However, these advancements must also address the typical challenges of deploying embedded sensors in underwater environments. These systems must withstand continual physical and mechanical stress, overcome communication

limitations due to UW operation, and guarantee an ultra-low power design due to their off-grid nature. The extreme conditions of aquatic settings present various adversities that need to be carefully addressed to open new perspectives in the field of embedded systems designed for harsh environments.

1.2 Aim of the work

The aim of this thesis is to study, design and implement a new technology for the in-situ measurement of sediment transport dynamics in marine ecosystems. By developing this prototype, we aim to lay the foundation for a novel set of embedded systems capable of accessing and monitoring phenomena that currently remain unobserved.
For instance, this work has potential applications in the monitoring of:

- Coarsed grained sediment inertial dynamics in riverine and marine settings;
- Sediment tracking along long stretches of rivers;
- Longshore sand transport mechanisms;
- Amount of sediment moved during extreme events;
- Amount of sediments moved from a river to a coast.

Building on existing techniques and combining them with new measurement methodologies, this study will conduct an in-depth analysis of hardware components and firmware optimization specifically tailored for continuous underwater operation over extended periods.
Although tested in a laboratory setting using readily available shields and microcontroller (MCU) boards, this prototype optimizes all operational phases, laying the foundation for future deployment in real marine environments.

This work will have two significant impacts. Firs, it will overcome an existing gap in the literature by providing an ad-hoc-designed technology for in-situ sediment dynamics measurements and tracking in harsh aquatic settings. This technology will not only improve existing measuring methods but also enable the monitoring of previously unmonitored phenomena related to sediment dynamics. This will significantly affect future works on the research of effective

solutions to contrast natural hazards like coastal erosion, land loss, floods and other extreme events which are becoming day-by-day more common as a consequence of global warming and sea levels rise. Second, the robust IoT architecture, implemented for long-term UW operation, will offer significant technological advancements for the entire field of low-power IoT applications deployed in critical environments.

1.3 Working principle of the proposal

To overcome these challenges, this work explores the development and testing phases of an embedded system based on the following logic: the core of the system will be a microcontroller that processes data from an Inertial Measurement Unit (IMU). The IMU measures sediment inertial movements, triggering the system to log data on a Micro SD card when movement is detected. This enables continuous monitoring of sediment inertial dynamics during active transport periods.
To ensure periodic data retrieval, necessary since measurements are logged into a local memorization device, the system exploits a Low Power Wide Area Network (LoRaWAN) communication infrastructure. Once a day, the device sends a status message through LoRaWAN. If the message is acknowledged, the system activates a GPS (Global Positioning System) module to attempt to track the sediment's location, which will facilitate retrieval of the smart pebble directly in the coastal area. While GPS can actively track superficial sediments, future works will consider LoRaWAN triangulation to spot buried pebbles. As a backup solution, for the sediments being UW, where radiofrequency transmissions are not feasible, each smart pebble will include an RFID (Radiofrequency Identification) tag, enabling efficient identification and retrieval directly on the field. The working principle is summarized in the flowchart reported in Figure 1.1.

The rest of the thesis is structured as follows: Chapter 2 covers the state of the art, Chapter 3 introduces the system's architecture, Chapter 4 describes the firmware optimization techniques, Chapter 5 presents the overall power consumption results, and Chapter 6 provides conclusions and potential future design steps.

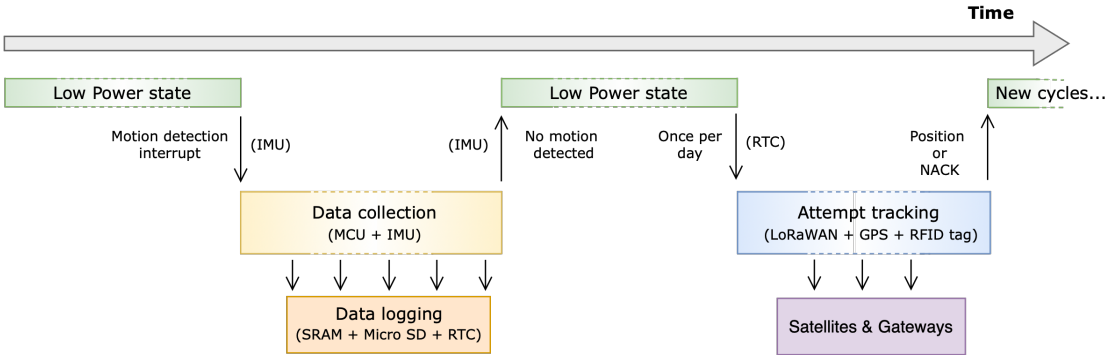

Figure 1.1: General working logic of the system and key operational steps.

2

State of the Art

In this chapter we will first introduce typical design methods and optimization techniques used to ensure efficient battery management in energy constrained environments. Following this, we will explore existing works in the field of sediment transport dynamics measurement, covering technologies such as RFID systems, distributed sensors architectures deployed in sandy beaches and the use of smart spheres in underwater environments.

2.1 Energy efficiency techniques

One of the main objectives of this research is to develop a sensor node with reduced power consumption. Extending battery's lifespan would result in the capability to operate over longer periods, reducing the necessity for frequent battery replacements. Additionally, extending the overall operating time allows for the collection of larger datasets, which is crucial for analyzing sediment movement patterns. However this project involves energy-hungry components, such as a transmission module and memorization devices, that pose several constraints to the overall operation. Therefore this makes absolutely necessary to use at least one energy-saving policy.
In general, the easiest solution would be to provide the system with an energy harvesting source [Peruzzi et al. (2020)]. Finding the correct balance between activation periods and energy available from the harvesting medium can guarantee very long lifespans as widely proved in previous works [Migliorini et al. (2023) and Mondal et al. (2017)]. Even other energy sources suitable for UW

scenarios have already been investigated as demonstrated by [Olumide et al. (2023)] focusing on microbial fuel cell, solar irradiance, acoustic and optical signals, near-field wireless power transfer based on inductive and capacitive coupling and hydro-kinetic energy from river and ocean waves, tides and vibrations. While all these solutions are able to continuously power the monitoring system, they are complex to be integrated in our context where compact size and high dynamism is required. For example, even a small turbine or compact piezoelectric generator might not be suitable. This is because our embedded system is designed to constantly move with the water current. The energy required to operate a battery management system (BMS) could overcome the advantages of energy harvesting in this case.

One promising approach to enhance energy efficiency is to dynamically adjust the node's operating time. Research by [Alippi et al. (2009)] explores various energy optimization techniques, including adaptive sensing, which is particularly relevant to this study. By aligning sensor activity with real-world conditions, adaptive sensing can significantly reduce data collection and processing. This method complements duty-cycling, which involves controlling sensor on/off periods. The microcontroller can orchestrate these functions, managing sensor power and implementing data acquisition strategies.
This dynamic technique can be combined with predictive models to further enhance the energy efficiency of the system. Indeed, starting from a set of data preliminary acquired, once the model is accomplished, it can be used to trigger the system activation only when relevant phenomena are expected rather than a periodic basis as it is done by standard duty cycling techniques. However this latter part will be the target of future works related to the one described here.

Other works are strictly related to the energy-saving aspect, proposing event driven architectures [Sundaresan et al. (2009)]. In this research work, the authors optimize the performance of an embedded system deployed under energy-constrained conditions in the context of event-detection applications. Two adaptive event-driven techniques are exposed, highlighting both their simple implementation directly at system level as well as their efficiency in both cutting down wasteful energy consumption and in improving the detection performance in applications characterised by long periods of quiescence followed by bursts of activity.

Briefly, this approach can be adapted to our system by turning on each sensor only when it detects an occurrence of the monitored event and then to turn it off once the transmission phase is completed. Of course the necessary components must be carefully selected given that to generate trigger signals such components must be optimized to remain constantly on with low energy requirement.

Some inspiration can also be drawn from the work proposed in [Raghunathan et al. (2006)] which, among the various optimizations, suggests the possibility of structuring the system around an heterogeneous multiprocessor sensor node. This approach aligns well with the event-driven architectures discussed earlier. In low-power designs, the typical workload can be divided into two main phases: low workload and high workload, each with corresponding proportional power consumption.
A practical solution would be to equip the embedded system with two microcontrollers or microprocessors that share the same hardware resources, such as the power supply, regulators, and sensors. Each processor would manage one of the two working phases. Specifically, during the lower phase, the more energy efficient MCU would operate keeping off the other controller, and vice-versa during the high workload phase. As will be discussed later in this thesis, one processor would handle event detection and trigger generation, while the other would act as the main controller during the high workload phase.

When designing low-power embedded systems for harsh aquatic environments, power management is essential. Optimizing the microcontroller, sensors, conditioning circuits, and wireless transceiver is crucial. Given the limited battery power available for sensor nodes, it is necessary to identify and reduce the energy consumption of each component during sensing and data processing. The subsequent chapter presents an optimized architecture based on these optimization principles.

2.2 Technologies for sediment tracking

While optimizing individual node components is crucial for energy efficiency, a comprehensive understanding of the broader context is equally important. To this end, the following discussion shifts focus to existing methods for monitoring sediment movement dynamics. By examining these approaches, we can identify

opportunities to integrate energy-efficient components and develop innovative solutions.
Researchers have explored various methods for directly measuring sediment transport dynamics in situ. Despite the importance of understanding these events, logistical challenges have limited research efforts. For example, tracking the movement of sand grains over extended periods remains still particularly difficult even with nowadays technologies. One of the biggest limitation, besides the challenges of tracking smaller grain sediments (such as sand), is that sediments can move both in and out of water, with overall displacements of up to hundreds of meters. This requires highly precise tracking equipment to retrieve the "smart sediment" from the environment.

If tracking smaller particles is still an hard task, innovative techniques have been invented in the years to follow the movements of coarse-grained sediments. Radiofrequency identification (RFID) is one of the most promising technologies in the field [Benelli et al. (2009), Benelli et al. (2013), and Pozzebon et al. (2014)]. RFID is a popular way to quickly identify things. It is used in many different areas, from letting people in to tracking products in factories and stores. In these studies the authors present a smart solution to overcome the limitations existing while using radiofrequency technologies in the underwater scenario. By exploiting the lowest operative frequency available in RFID spectrum, they effectively provide a new way to investigate the motion of sediments in grained beaches. RFID tags are directly placed inside the sediments, allowing for their retrieval and recognition after being deployed for weeks in a coastal area. By recording both the initial and recovery positions, authors were also able to assess the overall displacement.
The same solution has been widely tested by [Bertoni et al. (2010)], where 100 smart pebbles were released onto the beach and 74 of them were retrieved after two months. By using the RFID technology the authors were first able to determine the overall movement trajectory during this time span, second by retrieving the sediment they were able access the erosion effects of the environment on the tested samples.

More complex testing phases are described by [Benelli et al. (2011)], where the RFID technology has been tested by using new kinds of transponders and a mixture of smart pebbles with different materials composition. The track-

ing capabilities of smaller size sediments (3x3x1 cm wide) were also tested, demonstrating the efficiency to use the RFID technology in tracking typologies of pebbles on beaches composed by different materials, both under and outside water. All these results were widely confirmed also in other studies [Bertoni et al. (2012b), Benelli et al. (2012), and Bertoni et al. (2012a)]. RFID technology was also tested during short term experiences [Grottoli et al. (2015) and Grottoli et al. (2019)], in which it proved to be effective to understand how the size and shape of pebbles can affect their transport under low energy conditions.

One of the most recent studies on RFID technology was conducted by [Diet et al. (2016)], where a prototype for low-frequency RFID glasstag detection was presented. This research combined MATLAB simulations with laboratory testing to address a major issue that limits the functionality of RFID technologies in underwater settings: the limited detection range of the smart pebble. This limitation is primarily due to the reduced size of the glasstag, which typically restricts the mutual coupling between the tag and the detector. In their work the authors proposed a strong geometrical modification of the reader loop structure, opting for this approach over standard techniques like increasing the current or using a wider single-loop antenna to expand the detection area.

While research on using of RFID technology to monitor the movement of coarse sediments is ongoing, few works have focused on analyzing the dynamics of finer sediments, due to the logistical challenges typical of the field. However, also in this direction most of the existing studies are still based on a statistical approach, underlying once more an existing lack of technologies capable to acquire such information directly on the field.
An interesting work is described by [Pozzebon et al. (2018b)] where an innovative methodology is presented for direct measurements of aeolian sand transport on coastal dunes and beaches. The proposed solution is based on a a sensing structure capable to orient itself according to wind direction and directly calculate the overall amount of sediments transported thanks to on board electronics. Moreover this solution integrates a ZigBee radio module, which enables to perform all the measurements remotely without the need for human presence on the field. The reduced coasts, together with a simple structure and the low power consumption make the proposed research ideal for studying sand transport mechanisms over extended periods. In particular, the reduced power

consumption and the ease of installation are qualities we aim to achieve also in our work. These characteristics would allow to deploy multiple measuring devices simultaneously, thereby enhancing the quality of the analysis on the sediment transport dynamics.

Another curious work proposed by the same author [Pozzebon et al. (2018a)] presents a similar topology based on a distributed sensor network, for measuring sand level change in a sandy beach setting. The sensing technique was implemented by measuring dune height variation levels through an array of Light-Dependent Resistors (LDRs). The overall design underwent careful components selection (including counters, MUXs and MOSFETs) to obtain low cost and low energy requirement. The final results showed that very long life spans, potentially up to years, were achievable. This is a crucial aspect, considering that these systems find their more typical application in a scenario lacking other power sources.
In contrast, this research highlighted a common trade off for low power applications. To achieve extremely low power consumption it is often required to reduce the overall processing time by eliminating the processing unit from the setup. To address this challenge, in their work the authors modified the network structure. Two nodes were activated simultaneously for data collection and transmission to a central coordinator. The coordinator was then in charge to process the received data and deliver the results to a web application, allowing the user to see them in quasi real time. Each pair of nodes was activated with a delay of two minutes from the previous one, avoiding in this way possible interference due to data overlaps at the coordinator level.
Of course this setup is not suitable for all the fields of application. In our study, due to the required dynamism and the specific operating environment, such solution cannot be replicated. Achieving the correct balance between power consumption and on board processing time must be carefully analyzed. This will be one of the main topics in possible future works, where innovative, ad-hoc designed optimization techniques will be extensively studied to upgrade the prototype presented here.

Monitoring hard to detect phenomena often requires innovative and smart solutions. This is the case described by [Pozzebon et al. (2023a)], where the usability of grids of ultrasonic sensors for estimating the volume of granular materials

is presented. This solution introduced a new methodology to monitor finer sediments transport mechanisms, presenting both an hardware infrastructure and a mathematical model that can be easily implemented on a low-cost microcontroller to calculate the material volume from sensor readings. In such experimental setup, different sensor layouts were tested, suggesting that the sensor disposition highly affects the system performances.
Once again this research underlines the necessity to invent new setups to access unmonitored phenomena. In particular, the use of ultrasonic sensors is something we can consider to be included in our setups to monitor suspended particles.

2.3 Challenges with underwater tracking

Almost all the sensor nodes present a transmission technology, enabling real-time data and alarm signal delivery. However, due to the high attenuation introduced by the medium, the reliability of radio technologies is compromised in underwater scenarios. Some studies have explored this issue, focusing on LoRaWAN technology.
For instance, the propagation of electromagnetic waves in UW settings is widely described by [Cappelli et al. (2022)], focusing on LoRa (Long Range) technology's ability to communicate with a gateway outside water. Theoretical analysis and field tests highlight the correlation between penetration depth and water salinity, demonstrating that in fresh water this technology is reliable up to 120 cm from the water's surface.

Another interesting work is presented by [Pozzebon et al. (2023b)], where the transmission performance of the LoRa technology is investigated in terms of mean and standard deviation of received signal strength indicator (RSSI) and signal to noise ratio (SNR), and packet loss (PL) percentage. Laboratory tests demonstrated that the transmission is feasible only with the receiver submerged by a few cm of water, limiting the use of this technology to almost superficial monitoring when deployed in salty water.

These findings significantly impact all the underwater applications requiring both long transmission range and low power consumption due to the environmental constraints. To address this challenge, a more complicated infrastructure,

such as the one presented by [Campagnaro et al. (2022)], may be necessary. In this setup, underwater sensor nodes equipped with acoustic transceivers, communicate with emerged buoys, which can transmit to a base station by means of a standard low-power wide are network (LPWAN) technology.
However, the dynamic nature of our system, which needs to follow the water flow, introduces many complications, such as establishing reliable underwater communication with an emerged buoy. Considering our application scenario, LoRaWAN could be used to setup an alarm system. Although it may not be 100% reliable, several optimization techniques will be presented later on to maximize the reliability for the conditions we are operating in. For instance, by attempting transmissions during low tide, we can increase the likelihood of the system being above water, thus improving communication success.

The same issues that affect LoRaWAN infrastructures limit also the use of a GPS in UW settings. As a result, tracking in UW environments is usually achieved by means of alternative techniques, such as relying on acoustic waves. For instance DVLs (Doppler Velocity Logs) [Lu et al. (2022)] exploits Doppler effects of the acoustic waves to calculate the speed of the vehicle relative to the seabed by measuring the frequency variation between the emitted wave and the reflected wave. Typically four transducers are combined, measuring the velocity in different directions. By integrating these velocity measurements over time, the position of the vehicle can be estimated relative to a starting point. Another commonly used solution involves Long Baseline (LBL), Short Baseline (SBL) and Ultra-Short Baseline (USBL) positioning systems [Yoerger et al. (2007)]. These techniques rely on triangulation based on acoustic waves to determine the position. To further enhance performance, acoustic techniques are usually combined with other methods. For example, an inertial measurement unit (IMU) can be combined with acoustic beacons, where the IMU's drift error is calibrated using these acoustic signals [Wu et al. (2024)].
Despite their reliability, above-mentioned solutions are usually complex and not space efficient, making them unsuitable for our application, where all the electronics should be housed within relatively small pebbles.

Over the past five years, underwater backscatter localization has emerged as a promising technology for ultra-low-power, battery-less, and scalable underwater positioning, as discussed by [Ghaffarivardavagh et al. (2020)]. However,

this approach is still in its early stages due to limited 1-D positioning and short communication range.

Understanding the complexities of underwater tracking will require future dedicated research. In this study, we will rely on a combination of GPS and RFID technology for operation in a longshore coastal environment. The specific challenges and solutions associated with this approach will be explored in detail in subsequent chapters.

2.4 Comparison with existing methods

Eventually, this chapter is concluded by highlighting some common aspects as well as some improvements in comparison with some related works belonging to the literature. While RFID technology offers a valid solution for sediment tracking over extended periods, it lacks the capability to monitor precise dynamics in a continuous way, thereby limiting our understanding of the phenomena.

Let us start by citing [Sear et al. (2002)] which show a novel system capable of tracking multiple particles simultaneously in both littoral and fluvial environments. This system can provide finer spatial resolution, exploiting magnetic fields generated by loops of wire placed within the beach or riverbed. Basically the proposed solution calculates the electro-magnetic field at any position and uses this information to estimate the location of sediment particles that record the field strength. To enable position recording, the smart pebbles are equipped with orthogonally-mounted receiving coil and non-volatile EPROM memory. Although this research provides accurate information, it faces limitations in terms of scalability in a real world setting, as it constantly requires magnetic field generation, limiting its application to a smaller area.
More advanced solutions are proposed by [Spazzapan et al. (2004) and Abeywardana et al. (2012)]. These studies rely entirely on on-board electronics, facilitating deployment in real-world scenarios. The first study introduces a novel approach for acquiring sediment inertial dynamics directly in situ through an embedded system using an MCU, an IMU, and a memory device. The latter study takes this a step further by employing a similar architecture but introduces interesting post-processing phases. Starting with the inertial measurements acquired in the field, it enables the approximate reconstruction of the trajectory followed by the

smart pebble. While the architecture of these studies offers many advantages over older solutions, they both lack power optimization and a recovery system for retrieving the sediment from the environment to download the stored measurements. This limitation restricts their analysis to short observation periods and controlled environments where human supervision is possible.

Other interesting works based on similar architectures are proposed by [Akeila et al. (2010) and Alhusban et al. (June 2020)]. These studies focus on sensor error models to identify the necessary calibration procedures. Although they describe essential procedures for ensuring reliable sensor measurements, all the existing systems lack an architecture capable to operate over extended periods. Furthermore, as previously mentioned, none of these systems are equipped with a retrieval technology capable to recover the smart pebble from the environment. This limitation confines their use to scenarios where human supervision is possible, highlighting an existing gap in the literature.

In the subsequent chapter a novel architecture is described, with a focus on the hardware components needed to move toward a miniaturized design, which is necessary for reducing power consumption.

3

System Architecture

This thesis is focused on developing a solution for the measurement of the sediment inertial dynamics over extended periods. Consequently, optimizing power consumption will be the primary objective. In addition to this main goal, although readily available boards were used, exploring the mounted hardware components has been considered as the secondary goal. This preliminary studying phase is essential for understanding the functioning of each component, and will help guide future design steps toward a PCB design to enable on the field testing while minimizing power consumption.
The low power consumption requirement implicitly indicates that the node's deployment position cannot expect accessibility to the electrical grid. Therefore, it is likely that the node will be powered by batteries or energy-harvesting technologies. However, the use of renewable energy sources cannot be easily achieved in our application context, for this reason, the architecture proposed in this study has been customized on battery-powered embedded systems. In particular, power is expected to be provided by standard 4.2 V lithium batteries whose capacity is 3200 mAh.

Since the system discussed here requires many energy-hungry components, including the external memory, the LoRaWAN module, and the GPS module, the design activity will be focused not only on the structure of the embedded system but also on implementing custom optimization techniques.
In addition, careful estimation of the processing power is required. A powerful MCU unit is necessary if extensive on-board processing phases are expected, as

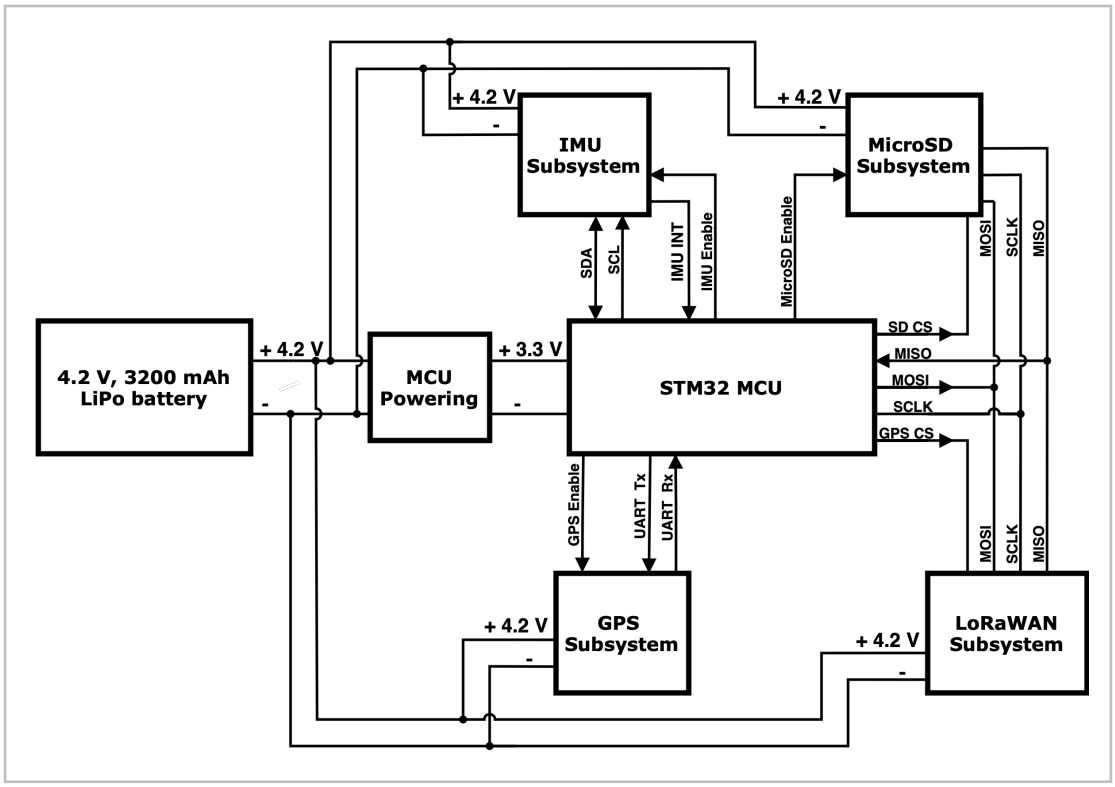

Figure 3.1: Block diagram of the overall embedded system.

it provides better performance at the expense of higher energy. Conversely, if the system's operation is limited to data collection, simpler devices with lower power consumption are more suitable. Therefore, it is necessary to select the MCU that offers the optimal trade-off between complexity and power consumption. Eventually, sensors have to be chosen looking for the best compromise between accuracy of the measured data and the required sampling frequency, and so the required energy to perform such task. Thus, by privileging sensors with event-driven capabilities, we would provide a way to efficiently control power hungry components that may become otherwise the crucial factor for energy consumption.

The proposed embedded system was designed using this approach, beginning with the selection of basic components: as an initial prototype, a single board microcontroller (STM32L476xx by ST microelectronics) was used, while for tracking and data transmission a standard GPS shield and LoRa module were employed. Finally, a low-power Inertial Measurement Unit (IMU) and a basic micro SD module were used for the measurement and logging phase.

After selecting the individual components, the design concentrated on integrat-

ing them into a single board, which included extra devices necessary for power management. Figure 3.1 depicts the final node architecture, which includes powering circuits and enable pins that deactivate various subsystems to execute duty-cycling rules. In particular, the MCU actively controls those switching subsystems and enable pins.
The diagram is divided into five main blocks: the measuring and data logging section, composed of an IMU and a micro SD, and the retrieval section which includes a LoRaWAN module and a GPS. The MCU orchestrates the entire system. The structure of each single block as well as its functioning will be described in detail in the following subsections.

3.1 Inertial measurement unit

The primary objective of this system is to measure the sediment inertial dynamics in an UW scenario, therefore, the most suitable sensor for this purpose is an IMU. To correctly choose the sensor we evaluated the main requirements for our application:

- Low power consumption is essential for the IMU. Indeed, as we will widely explain in the subsequent chapters, we are aiming to an event-driven architecture, meaning that the entire device will be triggered exploiting movement detection.
- The IMU should utilize its internal digital motion processor (DMP) for efficient operation. This would allow to use a duty-cycling approach in the IMU subsystem, reducing power consumption of the sensor, while keeping everything else completely off, minimizing power requirements.
- The sensor should come with a compact size and a standard communication protocol to facilitate the integration within the system.

Following all these considerations, the choice for the best solution fell on an MPU-6050 from InvenSense (Borregas Ave, Sunnyvale, CA, U.S.A.)[1].
The MPU-6050 integrates a 6-axis motion tracking device that combines a 3-axis

[1]https://invensense.tdk.com/wp-content/uploads/2015/02/MPU-6000-Datasheet1.pdf (last accessed: August 29, 2024)

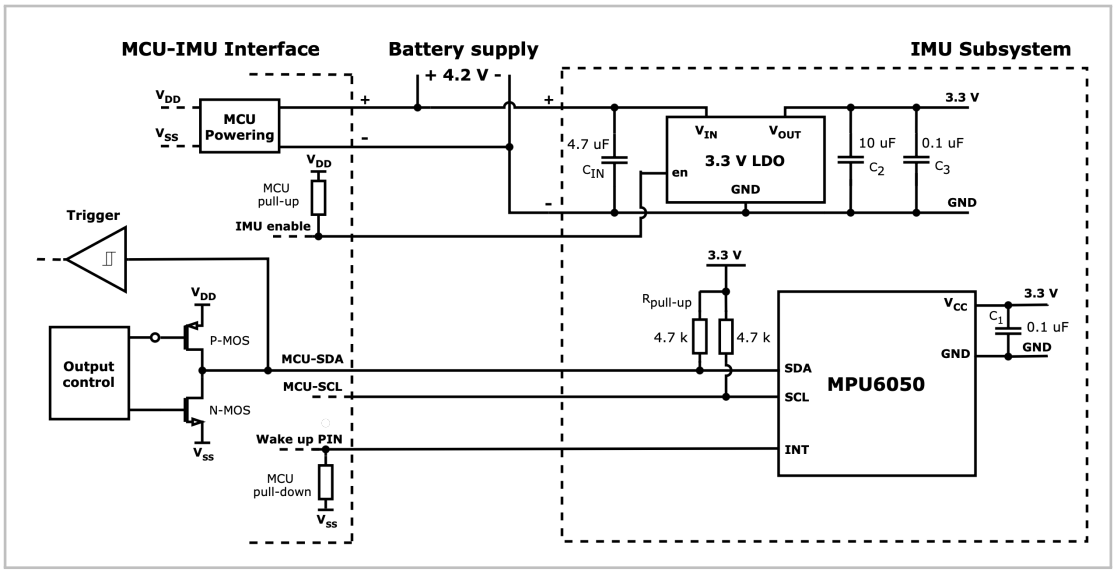

Figure 3.2: Schematic of the IMU and its interface with the MCU. The diagram illustrates the connections for interrupt generation, IMU enable/disable, and I2C communication, as well as internal configurations for pull-up and pull-down resistors.

MEMS gyroscope, 3-axis MEMS accelerometer, and a DMP all in a small 4x4x0.9 mm package. Briefly, micro electromechanical systems (MEMS) are capable to measure acceleration and angular velocity by measuring change in capacitance induced by movements and rotations. MEMS wafers are directly integrated with CMOS technology at the wafer level, resulting in power consumption reduced to just a few milliamperes, as highlighted in Chapter 5.

The IMU subsystem, together with the communication interface with the MCU, is represented in Figure 3.2. As we can see the shield can be directly connected to the battery supply. Indeed an internal low dropout voltage regulator (LDO)[2] adapts the battery voltage to the operating range of the IMU. As suggested in the datasheet a low equivalent series resistance (ESR) ceramic capacitor (C_{IN}) is placed at the input side, while at the output we find a combination of a larger tantalum capacitor (C_2) and a smaller ceramic one (C_3) to filter residual noise and enhance overall stability of the regulated voltage.

[2]https : //win.adrirobot.it/datasheet/integrati/pdf/regolatore_ld3985_4A2D.pdf (last accessed: August 29, 2024)

In addition to the power line, four additional pins are required: one for interrupt generation, one to enable/disable the IMU, and two for communication with the MCU via the I2C protocol.
It is important to carefully analyze how the MCU pins are internally configured, as this can eliminate the need for additional components, enhancing design efficiency. STM devices allow for internal customization, as highlighted in the schematic representation above. Specifically, an internal controller enables the activation/deactivation of internal pull-up and pull-down resistors. For output setups, you can choose between a push-pull configuration (both NMOS and PMOS activated) and an open-drain configuration (only NMOS enabled). In our case, a pull-up keeps the IMU on until the MCU actively turns it off, while a pull-down is used on the wake-up line to prevent the MCU from unintentionally waking up during standby mode. Concerning the SDA (serial data) and SCL (serial clock) lines, they were configured as open-drain given that pull-up resistors were already present at the IMU side. In this case, the selected values are standard ones, ensuring correct operation both when the NMOS transistor switches on, driving the line to a digital low, and when it turns off, allowing the resistor to pull the line to a high state. A quick verification can be done using equations 3.1, 3.2.

$$R_{pull,min} = \frac{V_{CC} - V_{OL}}{I_{OL}} = \frac{3.3 - 0.6}{10}\,\text{k}\Omega = 0.27\,\text{k}\Omega \tag{3.1}$$

$$R_{pull,max} = \frac{t_{r,max}}{0.8473 \cdot C_{par}} = \frac{300}{0.8473 \cdot 12}\,\text{k}\Omega \approx 17.7\,\text{k}\Omega \tag{3.2}$$

The first equation uses the fact that when $V_{OUT} = 0.6\,\text{V}$ the internal MOSFET can sink 10 mA, as specified in the MCU datasheet. The second equation is based on the I2C fast mode (400 kHz), which permits a maximum rise time of 300 ns. The parasitic capacitance of the bus is specified as less than 12 pF in the datasheet.
Regarding the internal pull-down on the interrupt line, its value is around 40 kΩ, making it a weak pull-down resistor that can be easily driven by the MPU internal components.

All the components described so far are capable of ensuring low power consumption while maintaining the required performances. The accelerometer, with its 16-bit ADC, has a maximum sensitivity of 16384 LSB/g with a full-scale

of ±2 g. The gyroscope provides a sensitivity of 131 LSB/(°/s) with a full scale ±250 °/s. The accelerometer can operate at an output data rate of up to 1 kHz, while the gyroscope supports up to 8 kHz, making both sensors sufficiently accurate for our application.
Further optimization on the software side will be treated extensively later on.

3.2 Micro SD memory

Our system will be deployed in an underwater environment, where continuous data transmission is not affordable. Furthermore, we aim to continually collect large amounts of data, making impossible to solely rely on the MCU's internal memory. Given these constraints, the decision was made to include a memorization device into our system. Both EEPROM (Electrically Erasable Programmable Read-Only Memory) and micro SD were evaluated, taking into consideration many factors.
EEPROMs typically have memory capacities ranging from a few kilobytes to a few megabytes. While suitable for applications requiring small amounts of data storage, such as storing configuration settings or small logs, EEPROMs capacity is limited for applications like ours. The acquired data volume requires a storage solution capable of handling several megabytes of data, which is well beyond the capability of even the largest available EEPROMs.
Conversely, micro SD cards are suitable for situations where a lot of data needs to be stored since they can provide storage capabilities ranging from a few gigabytes to terabytes. We chose a 32GB micro SD card as it had sufficient capacity to log sensor data for extended periods of time without requiring frequent data transfers. Micro SD cards offer a simple extensibility solution. Indeed, it is easy to upgrade to a higher-capacity card, so future requirements can be met by the system without major hardware changes. On the other hand, scaling up EEPROM storage, would require to add more EEPROM chips complicating the circuit design and increasing power consumption.
Micro SD cards also provide the benefits of simplicity of use and flexibility in data transfer. Data on a micro SD card is easily transferable to other devices, such as computers, for analysis. This plug-and-play feature simplifies data retrieval, especially in field circumstances where specialized tools may not be available. EEPROM, on the other hand, requires more complicated data retrieval methods, which often need custom interfaces and software.

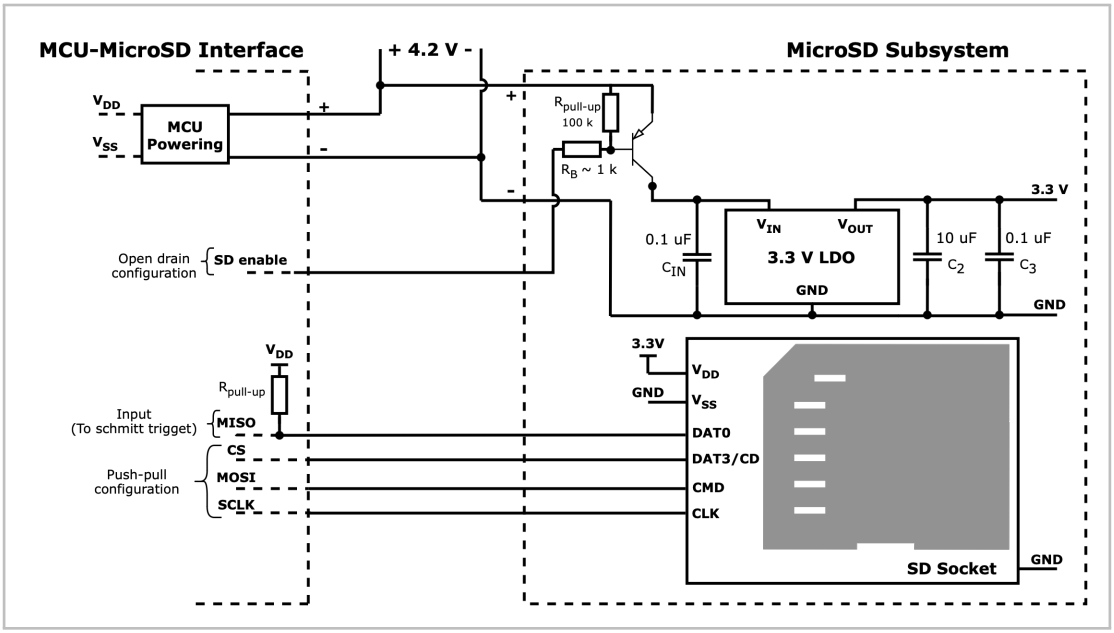

Figure 3.3: Schematic of the Micro SD and its interface with the MCU.

While EEPROMs have the benefit of consuming less power during read/write operations for smaller data chunks, the overall power efficiency of a micro SD card is evident when considering the system's total data storage requirements. The capacity to store vast volumes of data on a single, small memory card eliminates the need for multiple read/write cycles, saving energy on storage processes.

Figure 3.3 reports the connection of the micro SD module to the MCU. The MCU controls the Micro SD, therefore communication between the two components is managed via the Serial Peripheral Interface (SPI) bus. These devices interact synchronously in full-duplex mode, with the MCU acting as the master. Four signals are used: the slave uses MISO (Master Input Slave Output) to transfer data to the master, while the master uses MOSI (Master Output Slave Input). The master uses CS (Chip Select) to select the slave to communicate with, and SCK (Serial Clock) to synchronize communication. The SPI lines are actively driven using a push-pull configuration, meaning that no additional pull-up or pull-down resistors are required, except for the MISO line, which is pulled up during start-up.

The remaining pins are dedicated to the supply line and the enable signal. An LDO with the associated by-pass capacitors are used to provide stable 3.3 V to the Micro SD socket. In this case the shield did not have an exposed enable

pin for the LDO, forcing the use of an external pnp-type BJT (Bipolar Junction Transistor)[3] to manage the Micro SD activation. The sizing of the switch was done by considering that the Micro SD can require up to 100 mA during writing operations. The base resistor and the pull-up resistor were selected using the following procedure:

$$I_{base} = I_{MCU} \geq \frac{I_{c,max}}{h_{fe}} = \frac{100}{40}\,\text{mA} = 2.5\,\text{mA} \tag{3.3}$$

where we assumed that the minimum DC current gain is equals to $h_{fe} = 40$, which combined to the base-emitter voltage $V_{be,on} = 0.7\,\text{V}$ leads to a resistance value:

$$R_{base} \leq \frac{V_{bat} - V_{be,on}}{I_{base}} = \frac{4.2 - 0.7}{2.5}\,\text{k}\Omega = 1.4\,\text{k}\Omega \tag{3.4}$$

The final value selected was around $1\,\text{k}\Omega$, combined with a weak pull-up of $100\,\text{k}\Omega$ to keep the Micro SD off when the MCU is not actively driving the pin.

One important aspect to keep in mind is that the Micro SD socket behaves like a very low ESR capacitor. When the system switches on, this capacitor, initially discharged, generates a current spike know as "in-rush-current". Due to the low ESR, this current spike can reach up to hundreds of milliamps, causing a significant voltage drop on the supply line that may lead to unstable behavior in other connected components. In our case the effect was managed by using an LDO with internal current limiter, which prevents inrush current spike propagating to the input side, and so, to the main supply line. This limitation guarantees proper operation with minor effects, such as a slower turn on time of the Micro SD, which can be easily compensated for in the firmware.
Interesting optimization techniques for the Micro SD module are presented in the next chapters.

[3]https://www.st.com/resource/en/datasheet/cd00001225.pdf (last accessed: August 29, 2024)

3.3 Microcontroller unit

The IMU sensor and the Micro SD require the presence of a microcontroller unit to operate properly; this component must be chosen based on the best compromise of cost, power consumption, and efficiency. Low power microcontrollers often absorb few mA of current, making their influence on node power consumption important when no power-hungry transmission components are used.
As a result, the first prototype of the node was built using an STM32-L476RG Nucleo development board[4]. This platform provided a plug-and-play mechanism for integrating the surrounding components. The Nucleo board was also chosen because it has a sufficient number of analog and digital pins, allowing direct connection of various sensors. Even if we employed a development board rather than a single-chip MCU, the total consumption measurements were unaffected. In fact, an on-board jumper gives direct access to the current absorbed by the MCU chip and the associated voltage regulator, excluding all of the surrounding components included on any development board, such as additional LEDs, regulators and input buttons.
We decided to move from basic platforms, such as the Arduino UNO, to more advanced ones since STM-side development allows to achieve excellent performance with reduced power consumption. As it will be better discussed in subsequent sections the high clock frequency (up to 80 MHz), the large flash memory (up to 1 MB), together with extremely low power consumption during sleep mode enable us to scale the processing tasks directly on board. This allows the system to obtain information on-board, enhancing the device's overall operation.

The MCU schematic is reported in Figure 3.4. As previously mentioned, even if the device was prototyped using a Nucleo board, we are reporting only the few components necessary for the operation of the single-chip MCU that were accessed through the dedicated jumper.
An external 500 mA low quiescent current and low noise LD39050 voltage reg-

[4]https://www.st.com/en/evaluation-tools/nucleo-l476rg.html (last accessed: August 29, 2024)

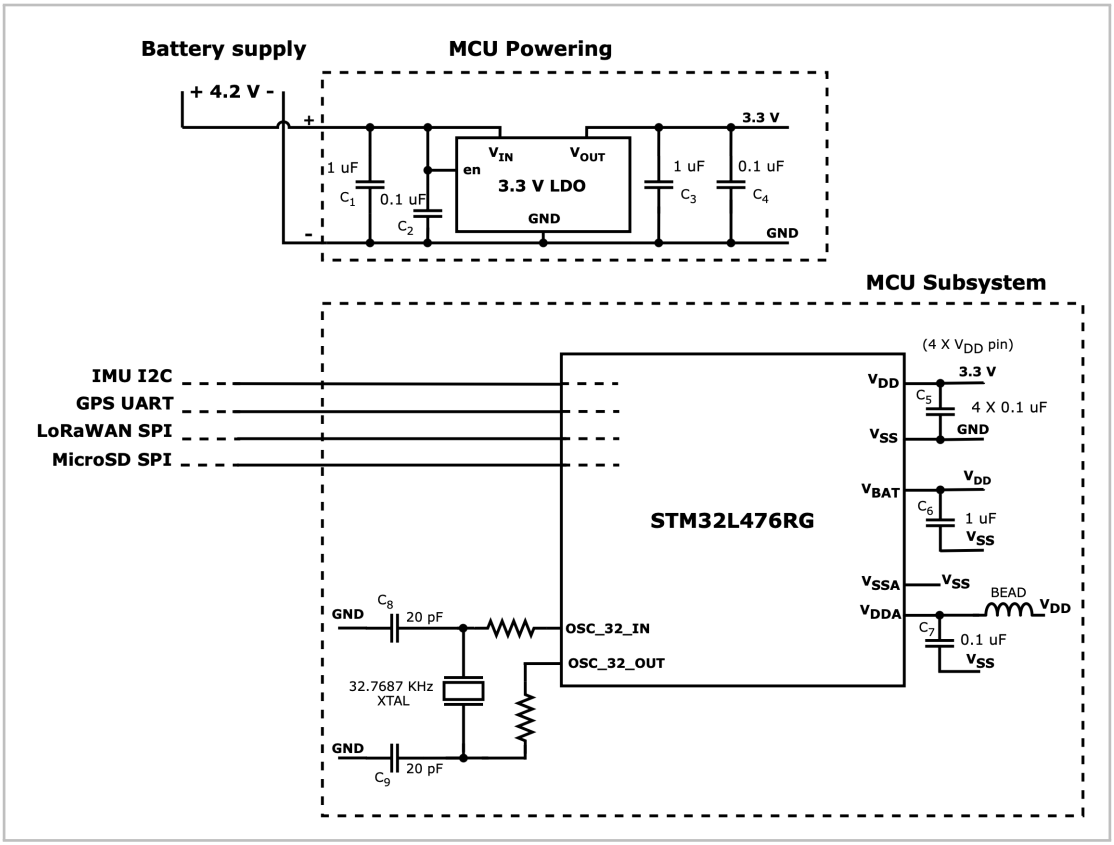

Figure 3.4: Schematic of the MCU and its interface with the surrounding components. The diagram highlights the use of a dedicated LDO and external quartz crystal.

ulator[5] is used to adapt the battery voltage to the 3.3 V needed for the MCU operation. In this case, it is extremely important that the regulator guarantees very low quiescent current (20 μA typical at no load) to avoid impacting the system's performance when operating in a low power state. In addition to standard decoupling capacitors placed near the power pins, a small ferrite bead is used to isolate analog supply (V_{DDA}) from high frequency noise coming from digital lines (V_{DD}). A ferrite bead opposes changes in current, but instead of storing energy in a magnetic field like an inductor, it uses resistive behavior to limit current variation, creating a high impedance path for high-frequency noise.
Another component worth mentioning is the external quartz crystal. While STM MCU chips typically generate a high-speed internal clock (HSI), which is then distributed to various components and peripherals using internal PLLs (Phase-

[5]https://www.st.com/resource/en/datasheet/ld39050.pdf (last accessed: August 29, 2024)

Locked Loops) and eventually reduced to the 32 kHz needed to maintain the reference in the internal RTC (Real-Time Clock) for timestamps, we opted for an external quartz crystal. This is because, in our application, we need to maintain the time reference even when the system enters standby mode, which does not allow the powering of all the internal circuitry required to provide the RTC's clock internally.

In addition to managing the IMU and the Micro SD memory during the measurement phase, the MCU is the core of our system during the retrieval phase when the GPS and LoRaWAN module cooperate to fix the system's position. These two remaining subsystems are better described in the subsequent subsections.

3.4 GPS module

While the Micro SD card and the IMU allow for the acquisition and logging of inertial dynamics measurements, a reliable methodology for retrieving the smart pebble from the field is still needed. Our design will include passive RFID tags, as discussed in Chapter 2. However, this study also aims to enhance existing methodologies for recovering the system from the environment.

As a first prototype, we have integrated a standard positioning system based on GPS into our architecture. Although GPS is not completely reliable in underwater settings, a careful analysis of the operational conditions can highlight opportunities for radiofrequency technologies to be effective also in these harsh environments. For example, in a longshore scenario where low tide exposes up to ten meters of the seabed with a clear sky view at certain times of the day, attempting a GPS position fix during these moments can significantly increase the likelihood of accurately mapping the pebble's location.
Furthermore, we are currently using LoRaWAN technology to establish an alarm system. This means that the smart sediment is not continuously tracked; instead, our goal is to enhance the chances of recovering it from the environment by combining GPS with existing RFID methods. This approach will allow us to analyze the data logged during its working period.

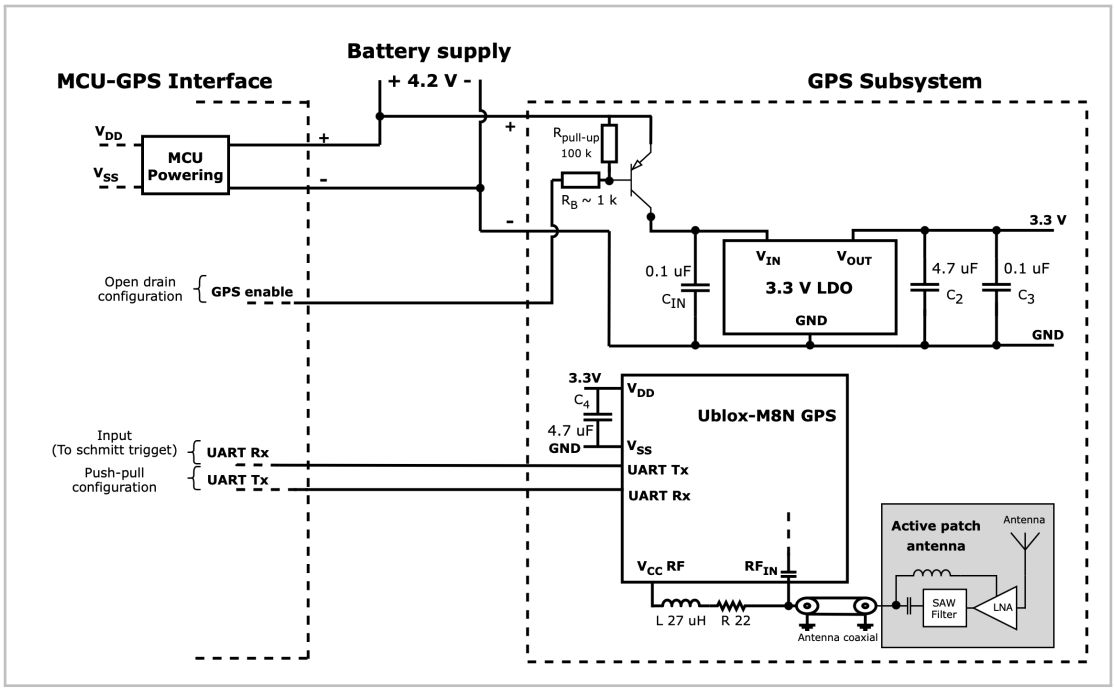

Figure 3.5: Schematic of the GPS and its interface with the MCU.

To this we started with a standard GPS shield from Ublox[6]. The complete schematic together with the interface with the MCU is reported in Figure 3.5.
Once again a dedicated LDO is used, along with a pnp-type BJT switch sized following the procedure reported in equations 3.3 and 3.4. We chose to hard switching the entire system given that the internally managed low power routines were not compliant with our operating conditions, as it will be better expanded in the next chapter.
In addition to the supply lines, the GPS comes with an active patch antenna with an internal LNA (Low Noise Amplifier). The antenna was linked to the GPS chip through an inductor needed to isolate the antenna supply from the high frequency signal, and a small 22 Ω resistor to limit the current flowing on the RF branch in the event of antenna short circuits.

The interface with the MCU is simpler compared to previous components. It relies on a straightforward UART protocol, requiring only two pins to establish a complete full-duplex communication.

[6]https : //content.ublox.com/sites/default/files/NEO − M8 − FW3_DataSheet_UBX − 15031086.pdf (last accessed: August 29, 2024)

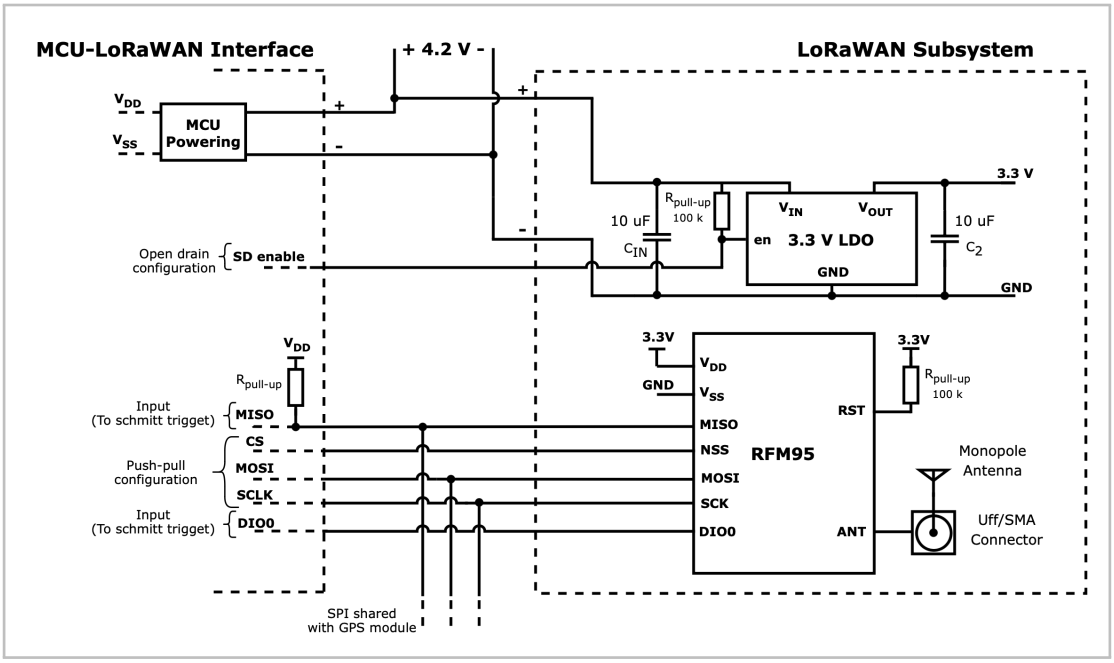

Figure 3.6: Schematic of the LoRaWAN shield and its interface with the MCU.

3.5 LoRaWAN module

This chapter concludes by introducing the final component of our setup.
Many wireless communication technologies are nowadays available, including Local Area Networks (LANs) such as Bluetooth, ZigBee, and WiFi; Wide Area Networks (WANs) like SigFox and LoRa; and Global Area Networks (GANs) such as 2G, 3G, and 4G, which have already proven effective in many low-power applications.
However, after analyzing the constraints imposed by the target environment, Local Area Networks were immediately excluded due to their limited communication range. While Global Area Networks are reliable, they involve higher costs compared to other solutions, both because they are SIM-based, requiring subscription payments, and because GPRS/UMTS modules are more expensive than other standard technologies.
For our initial study, we decided to rely on LoRaWAN technology, a cost-effective solution that provides long communication range with reduced power consumption. The same constraints that affect GPS operation in an UW setting also limit LoRaWAN performance. However, as previously mentioned, our goal is not to implement continuous tracking, which would consume too much energy and

limit the device's operation to just a few hours. Instead, we are developing a mechanism that tracks the sediment when the likelihood of achieving a position fix is maximized. This approach should assist in retrieving those sediments that cannot be located using simple RFID tags. While RFID technology has proven effective in recovering sediment on the exposed parts of the beach, by combining GPS signals with LoRaWAN, we aim to detect smart pebbles that typically lie on the seabed but emerge during low tide.

A readily available RFM95 shield from Adafruit[7] was employed. The interface with the MCU is reported in Figure 3.6. No additional components were used in this case, given that the internal LDO provides access to the enable pin to manage the component's duty cycle. The communication happens through an SPI protocol, that is shared with the GPS SPI lines.

[7]https://cdn-shop.adafruit.com/product-files/3076/RFM69HCW-V1.1.pdf (last accessed: August 29, 2024)

4

Energy-Saving Policy

This Chapter covers two relevant themes for the correct operation of our system: the description of the energy-saving policy and the implemented optimization methods adopted to extend the lifetime of the 3200 mAh lithium battery supplying the device.
The overall operation of the system is divided into two main phases. The first phase uses the IMU and the micro SD to actively collect data from the environment. The second phase attempts to track the device's position by using the LoRaWAN and GPS subsystems. Besides this strategy, the overall battery capacity has been partitioned. The first part, corresponding to 90% of the initial capacity, is used to actively collect data, while once per day the GPS and LoRaWAN subsystems are activated to attempt a position fix. Conversely, in the second portion, corresponding to the remaining 10%, the system solely relies on the transmission section, deactivating the IMU and the micro SD to extend the battery-life duration. Indeed, the battery level is constantly monitored to stop data acquisition when it falls within certain threshold. This will allow the preservation of a portion for the retrieval phase.
Both the operating phases are better expanded in the subsequent sections.

4.1 Phase 1: measurement and data acquisition

To facilitate the explanation of the energy-saving policy, let us introduce a first flow chart reported in Figure 4.1. The flow chart qualitatively reports the operating state, during two different cycles, of the main components employed in

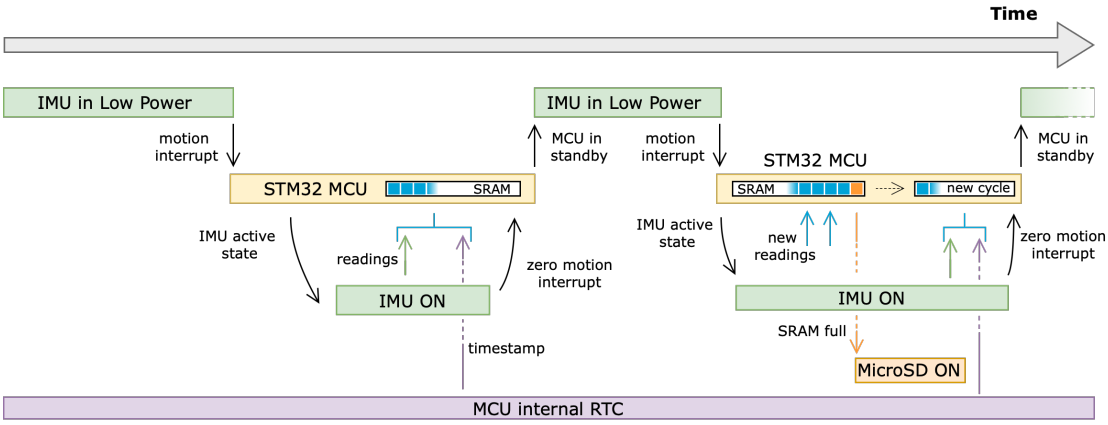

Figure 4.1: Flow chart of the first operating mode of the device.

this first phase: the IMU sensor as a green bar, the MCU as a yellow bar, the Micro SD as an orange bar, and the MCU's internal RTC as a purple bar. The time line in not linear, and the arrows labels try to give an intuitive way of the overall working principle, summarized in the following points.

Stage 1. Initially, all components are disconnected from the supply line using switches or enable pins, except for the MCU, which operates in standby mode, and the IMU sensor, that is programmed to operate in a cyclic way. In this state, both the DMP and the gyroscope are kept off, while the accelerometer axes are periodically switched on and off at a frequency of 1.25 Hz. The internal digital high-pass filter is configured in a "hold" state, so the IMU output corresponds to the difference between the current sample and the reference value stored in the filter when the sensor was at rest (i.e., before entering sleep mode). This configuration allows the DMP to operate in a motion-detection mode, triggering a signal only when the overall output exceeds a predefined threshold for a specified duration.

Stage 2. When a trigger event occurs, the MCU takes control of the system. The IMU is reprogrammed to operate continuously, reading both the accelerometer and gyroscope data, while the internal DMP is used to trigger a "zero-motion" interrupt when no consistent movement is detected for a specified time. At this stage, the samples are not directly written to the micro SD. As we will better discuss at the end of this section, writing operations are expensive and require careful analysis to avoid corrupting the overall device's operation. Therefore,

the incoming samples, along with the timestamps generated by the MCU's internal RTC, are initially stored in a dedicated SRAM region. When movement is no longer detected, the MCU reverts the IMU to the configurations illustrated in Stage 1, and it enters standby mode, awaiting a new interrupt signal.

Stage 3. Stages 1 and 2 are repeatedly cycled. Once the MCU's internal SRAM is full, the micro SD is powered on, and the entire buffer is transferred in a single operation. After this transfer, the micro SD is immediately powered off, and the SRAM is reinitialized for the next measurement cycle.

Optimizing write operations is a critical aspect during the data acquisition phase. Micro SD cards support both reading and writing, and in small embedded systems they are typically handled using the FatFs[1], which operates with the FAT (File Allocation Table) file system to manage files and directories on storage devices. MMC/SDC memories are organized into sectors, meaning that while reading a chunk of data, the entire sector containing that data must be loaded before accessing the specific bytes of interest. This procedure is facilitated by an intermediate I/O buffer. As shown in Figure 4.2, aligning read and write operations with sector boundaries avoids to load data into the intermediate I/O buffer, thus speeding up the whole procedure. Furthermore, reading multiple sectors at a time is recommended, as this approach allows the sector's FAT information to be loaded only once, further enhancing efficiency.
Another important aspect concerns the writing operations effectiveness. Similarly to reading operations, writing data requires erasing an entire memory block even when storing a small amount of data. This behavior can easily corrupt the energy efficiency of our system given that even simple writing operations involve current peaks up to 100 mA. One possible solution to address this problem consists in delaying writing operations until enough measurements are collected to fill an entire data block in the Micro SD memory. As shown in Figure 4.3, the busy time required to erase a block before writing new data remains constant, as long as the data transferred stays within the memory's block boundaries. While writing larger data blocks involves additional time for data transfer via the SPI bus, the overall throughput is improved.

[1]http://elm-chan.org/fsw/ff/ (last accessed: August 29, 2024)

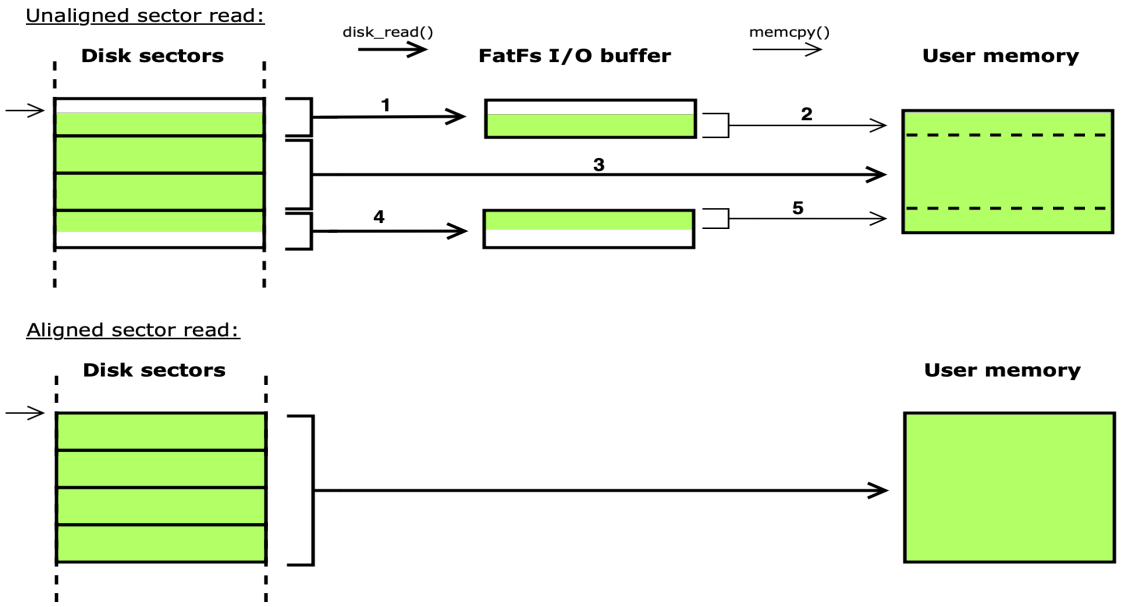

Figure 4.2: Impact of sector alignment on reading and writing operations.

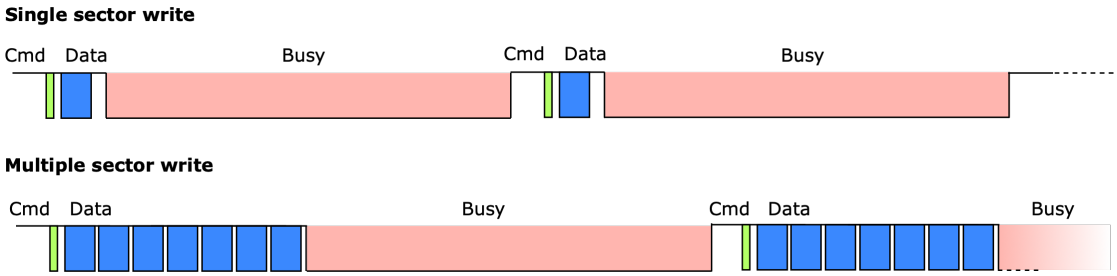

Figure 4.3: Comparison between single-sector and multi-sector writing operations.

This procedure was implemented exploiting the MCU's internal SRAM. Indeed, STM32-L4 series allows to allocate a dedicated RAM region that is continuously powered even when the device enters a low-power state. This makes possible to preserve all the measurements across multiple operating cycles, significantly reducing the need to activate the Micro SD. Additionally, among all the possible power down routines available on a STM microcontrollers, we opted for standby mode. As shown in Figure 4.4 this operating mode allows to power down almost all the internal components except for few dedicated interrupt pins, the LSE (low speed external) oscillator which maintains the time reference on the internal RTC, and the allocated SRAM region. This is made possible by a low-power regulator that reduces power consumption to just a few microamperes during low-power routines, as we will demonstrate in the next chapter.

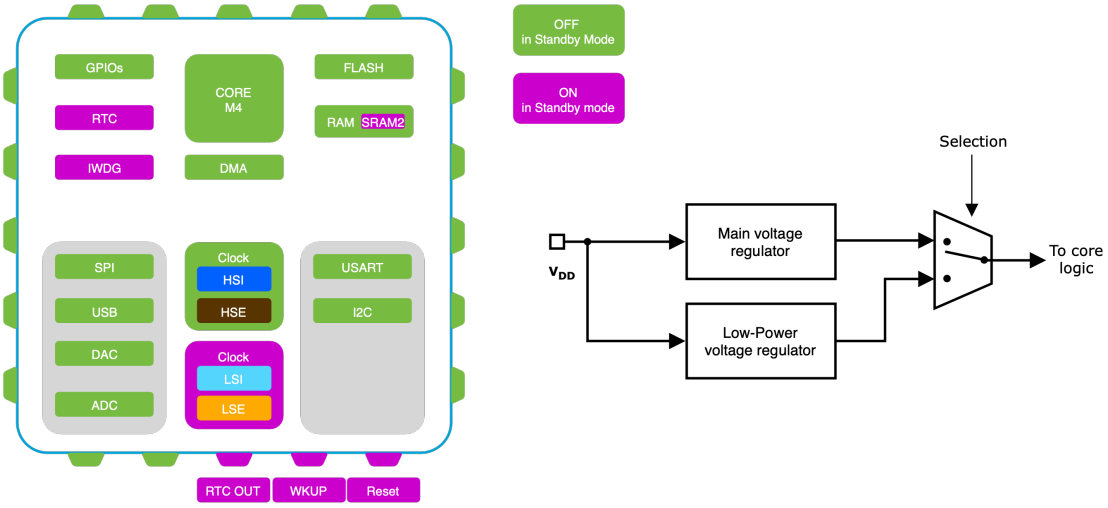

Figure 4.4: Diagram of MCU internal components with standby mode power status.

4.2 Phase 2: tracking operation

In our application, achieving real-time transmissions is not feasible due to all the limitations imposed by the operating conditions. As a consequence, all the measurements are logged into an onboard memorization device. To prevent data corruption, it is essential to ensure the capability to track and retrieve the device from the environment.

To facilitate retrieval operations if GPS fails to establish a position fix, our setup will include passive RFID tags. We can rely on recovery methods similar to those presented in Chapter 2, where it was proved that relying solely on RFID technology can recover more than 70 % of the samples after two months of operation in a coastal area. The aim of this study is not only to develop an innovative methodology for continually measuring sediment inertial dynamics but also to improve the likelihood of retrieving samples from the environment and to estimate sediment position across the operating period.

To solve this problem, we have developed a new method that helps track the smart pebble's position at least once a day. The MCU's internal clock is used to activate the transmission section (LoRaWAN module and GPS) daily. The goal is to provide a rough idea of the sediment's movement over time. Although the exact position may not be recorded every day due to challenges with UW transmissions, we can reconstruct the sediment's path by combining IMU data with the positions we manage to record. This will allow us to reconstruct

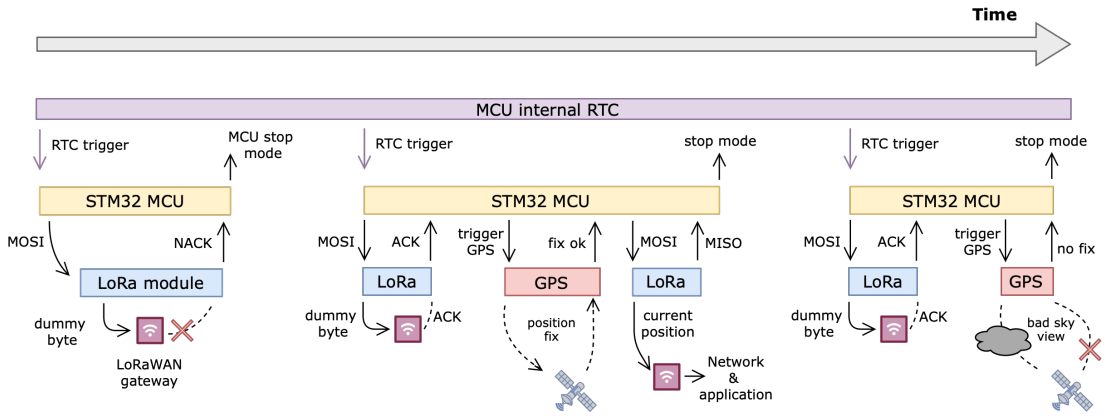

Figure 4.5: Flow chart of the second operating mode of the device.

the sediments overall movement during its operation period by means of post-processing phases.
Additionally, LoRaWAN transmissions enable remote access to the smart pebble's position, facilitating retrieval operations.

The use of radiofrequency technologies in UW settings adds additional challenges for our application. These components are usually power consuming, requiring dedicated optimization routines to prevent premature battery drain. The strategy implemented is better described in the following points, which expand the flow chart shown in Figure 4.5.

Scenario 1. The MCU's internal RTC triggers the activation of LoRaWAN module at specific instants designed to maximize the chances of a GPS lock. This mechanism is explained in more detail at the end of this section.
Our setup is configured for ABP (Activation By Personalization) and operates as class-A device. Upon activation, the LoRaWAN module sends a dummy byte and opens two receive windows to wait for an ACK (Acknowledgement). This initial phase determines whether a reliable connection can be established and whether it is worth activating the GPS module. If no response is received within the dedicated windows, the device reverts its operation to the acquisition mode and retries a position fix the following day.
This situation corresponds to the first scenario illustrated in the flowchart above. In this case, while we may miss the device's position, this approach helps avoid activating power-hungry components like the GPS if the system is submerged.

Scenario 2. If an ACK is received, the MCU activates the GPS. Given the resource-constrained operating conditions, we opted to use the GPS in its full performance mode rather than relying on internal low-power routines. This is because, to have any benefit under low-power routines, the Ephemeris and Almanac datasets, which are necessary for maintaining the current position, need to be hourly updated. However, in our operating conditions the risk of failing a position fix is high. This failure would involve restarting the whole procedure, consequently losing all the benefits of the internally-managed routines.
By evaluating the RSSI (Received Signal Strength Indicator) of the received ACK, we can assess the quality of the communication link. If the signal strength is satisfactory, it indicates that the system is likely to be above water's surface, increasing the chances of the GPS successfully completing a position fix. This information will then be available to the user on the network server, as illustrated in the second scenario of the flowchart above.

Scenario 3. Altought LoRaWAN provides an effective way to avoid GPS activation while submerged, there may still be situations where an ACK is received under poor sky view conditions, preventing the correct operation of the GPS, as shown in the third scenario in the flow chart. To overcome this problem, and avoid unnecessarily long activation times, which would lead to higher power consumption, we implemented an optimized method to dynamically assess GPS signal quality. This mechanism, better explained at the end of the section, allows to quickly shut down the system, preventing excessive power waste in situations where position fix is likely to be missed.

As mentioned earlier, LoRaWAN activation should not happen randomly. The transmission section of the device can be optimized for activating when the chances of having a GPS lock are maximized. This analysis can be done before deploying the system, exploiting TLE (Two-Line-Element) data, which are regularly updated by various organizations and provide orbital information for satellites (SVs).
In our application we relied on Python's Skyfield[2], a Python library that is widely used for high-precision astronomy calculations, including tracking and

[2]https://rhodesmill.org/skyfield/ (last accessed: August 29, 2024)

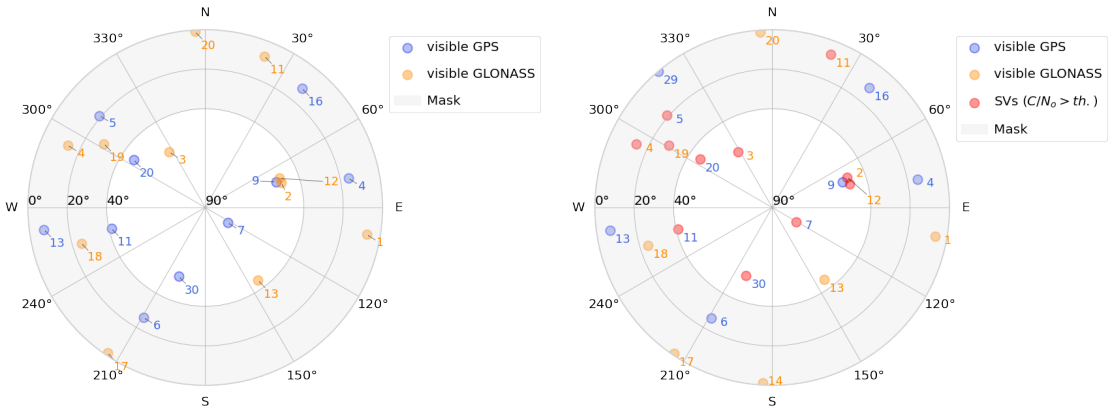

(a) **(b)**

Figure 4.6: GNSS satellite visibility for: (a) one week old TLE-based prediction, (b) on the field measurements using u-blox GPS.

predicting satellite trajectories. Skyfield's orbit prediction capability enabled us to calculate the precise position of GNSS satellites up to two weeks in the future. This data can be stored on the MCU's internal memory, and the system can activate the GPS receiver prioritizing the windows with the highest satellite visibility.

In our setup we considered only time windows when more than 7 satellites were visible with an altitude angle greater than 40°. This conservative approach was used because, while generally four satellites are enough to triangulate the device's position, the harsh operating conditions require aiming for a higher number of satellites and a higher altitude angle. Despite this, we were able to obtain at least a couple of valid windows per day, with an average visibility period of around seven minutes. This duration was used to setup a timeout; if the GPS cannot lock a position within seven minutes, it reverts to standby mode. One test using one week old TLE data is reported in Figure 4.6. On the left side, we report predictions made with Python, while on the right, we can observe satellite positions obtained from the GPS module.

A minimal displacement between the predicted and observed positions indicates that the predictions are reliable for up to a couple of weeks. Additionally, we evaluated the C/N_0 (Carrier-to-Noise Density Ratio), which represents the ratio of the power of the GPS signal (carrier) to the power of the background noise (noise density), and allows to determine the strength and reliability of the GPS signal. As shown in Figure 4.6(b), most of the satellites within the mask exceed the C/N_0 threshold equals to 30 dB-Hz, as recommended in the datasheet.

We also considered the quality index provided by the GPS, which ranges from 1 to 7, with higher values indicating better signal quality and positioning accuracy. By assessing both the C/N_0 and the quality index, the MCU determines whether it is feasible to attempt a position fix. In particular, according to the datasheet, a GPS lock should be achieved in less than 40 seconds under clean sky conditions. Therefore, we set an initial timeout of 30 seconds. During this time, we expect to lock at least a couple of satellites with good signal quality; otherwise, the GPS reverts to standby mode, as a failed lock within this initial window, possibly due to external factors like weather, suggests that a position fix is unlikely to be completed.

This solution minimizes overall GPS activation time by setting a timeout, speeding up the acquisition process by attempting a fix only during high visibility windows, and analyzing the quality of both LoRaWAN and GPS signals to determine if a position fix can be completed.

In addition, as mentioned earlier, the MCU continually monitors the battery status. When the residual capacity drops below a certain threshold, for instance 10 % of the initial full charge, the operation of the device is reverted from the acquisition state to the retrieval phase. During this phase, the Micro SD and the IMU sensor are powered down, and only the GPS and LoRaWAN module remain active, focusing solely on the retrieval operation.

5

Power consumption analysis

This Chapter presents an analysis of the power consumption of the overall setup, optimized with the implemented strategies. First, the current profiles of the single components embedded in the system are presented, highlighting their nominal absorption during different working phases, as well as the time spent in each phase according to the predefined policy. The contribution of passive components, capacitors and switches is considered negligible in the overall power consumption of the system.
The acquisition was done by using the Power Profiler Kit II from Nordic Semiconductor[1]. The device was configured to operate as an ammeter and was placed in series with the DUT (Device Under Test).

Starting with the IMU, the corresponding profiles during transients and active mode are respectively reported in Figures 5.1 and 5.2. In the first figure we can clearly distinguish the instant in which the interrupt is generated. After a 50 ms delay, during which the IMU is reprogrammed for full performance operation, all the accelerometer and gyroscope axes are activated and the sensor starts the measurement mode, that lasts until consistent movement is no longer detected. At that time, as shown on the right figure, a zero motion interrupt is triggered, and after another 50 ms delay the IMU reverts its operation to the cyclic mode, ready to trigger a new interrupt.

[1]https://www.nordicsemi.com/Products/Development-hardware/Power-Profiler-Kit-2 (last accessed: August 29, 2024)

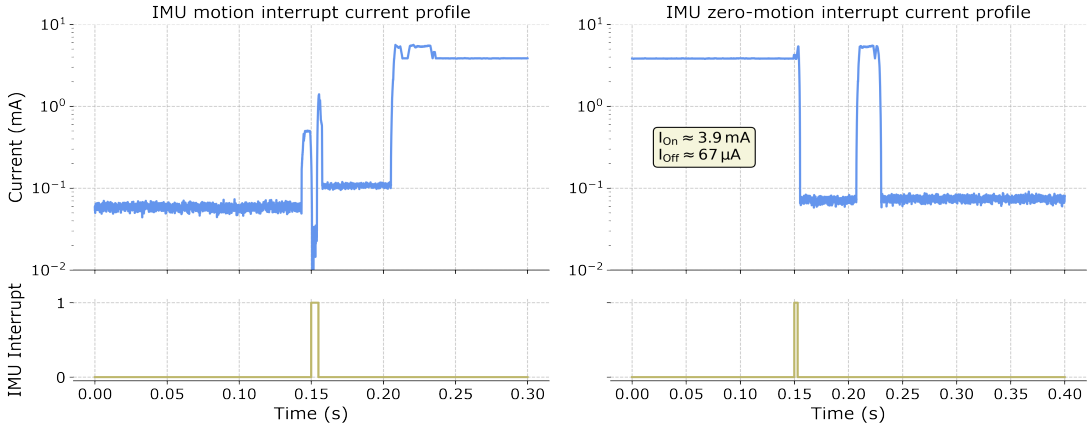

Figure 5.1: Current consumption profile of the IMU during transitions from low power to active state (left) and from active state to cyclic operation (right).

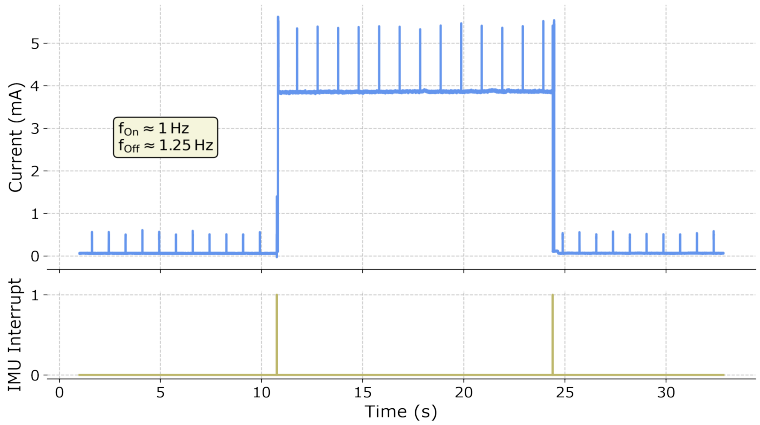

Figure 5.2: Current consumption profile of the IMU during the active state.

The second figure reports a complete operating cycle. As shown, the update frequency during low-power mode was set to 1.25 Hz, which provided a good trade-off between power consumption and the capability to detect coarse movements. For the same reason, during active phases the update frequency was set to 1 Hz. Of course, future works will take into consideration the development of more advanced algorithms capable to dynamically adjust the sampling frequency according to the measured intensity of the movement.

For the Micro SD, we report four different profiles, each corresponding to a different data block length. This allowed us to validate the optimization technique introduced in the previous Chapter.
Specifically, we tested data chunks of 1 byte, 512 bytes, 4096 bytes and 16384

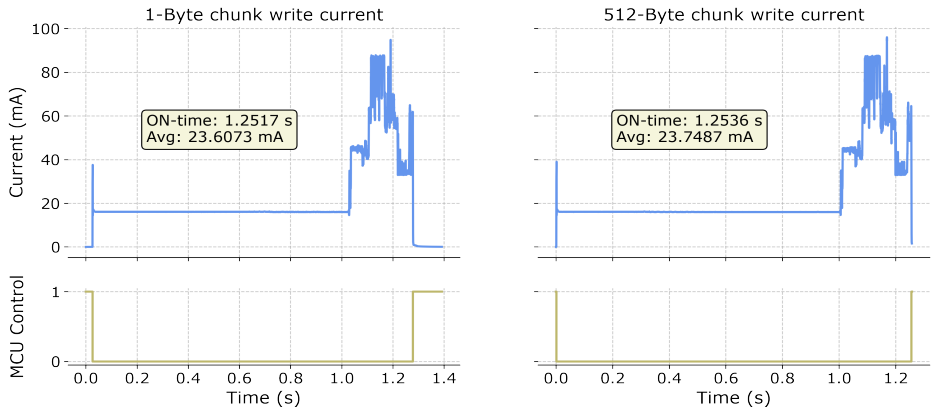

Figure 5.3: Current consumption profile of the Micro SD: 1 byte long data (left), 512 bytes long data (right).

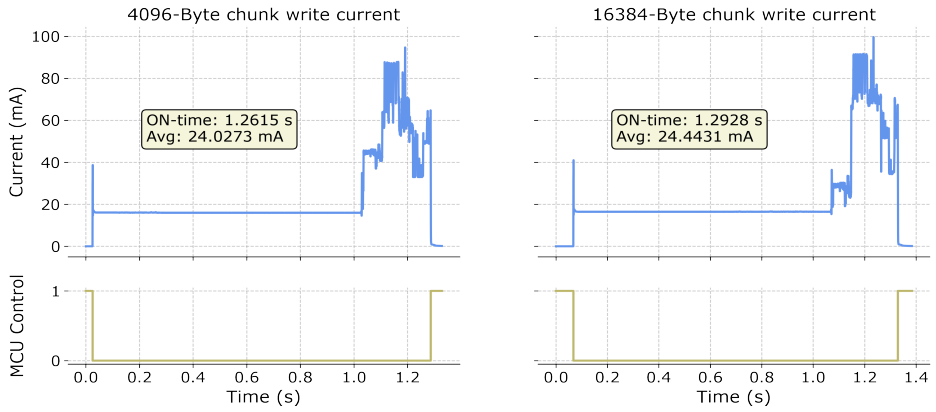

Figure 5.4: Current consumption profile of the IMU while switching from: low power to active state (left), active state to cyclic operation (right).

bytes, corresponding to the current profiles shown in Figures 5.3 and 5.4. As we can see, the overall shape remains consistent despite the increase in the data writing size. The only noticeable difference is a slight increase in writing time which moves from 1.2517 s, for a single-byte transmission to 1.2928 s, for 16 -kilobyte writing operation. However, this discrepancy aligns with the time required to transfer data using a 5 MHz SPI bus. Another key aspect to notice is that all the profiles present a flat region lasting approximately 1 s immediately after the switching instant. This behavior is controlled by the firmware, which imposes a 1 s delay to stabilize the Micro SD after powering on the device.
The activation signal, shown at the bottom of the plot, is inverted because the p-type switch conducts when a low logic level is applied at the input.

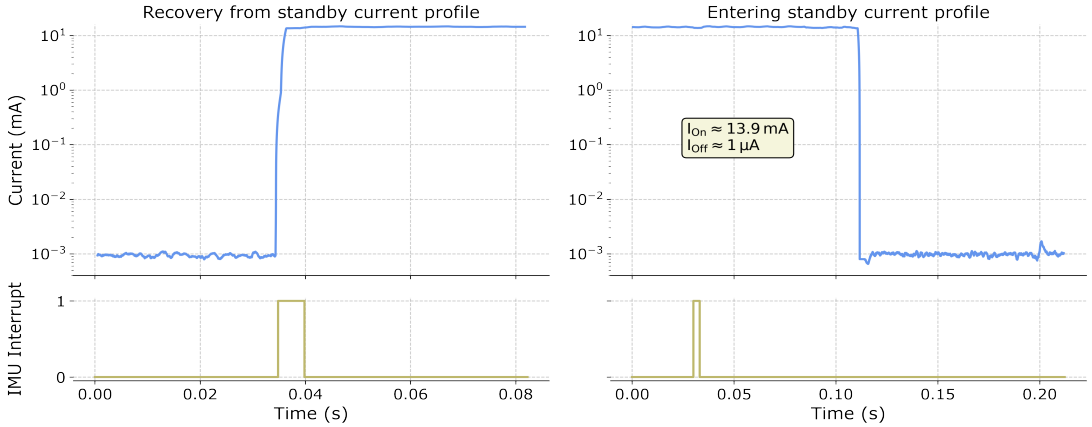

Figure 5.5: Current consumption profile of the MCU when switching from standby mode to active state (left) and from active state to low power mode (right).

Focusing on the MCU, the current profile is reported in Figure 5.5. Both the active current, at 13.9 mA, and the standby current, around 1 μA, align with the nominal values reported in the datasheet. The standby current is slightly higher because we are retaining both an SRAM region and RTC settings during power-down cycles. Another important consideration is that, to obtain a complete overview of the power consumption of the entire MCU subsystem during power-down states, the current drawn by the dedicated LDO, approximately 20 μA, must also be accounted for.
Furthermore, we can notice that during the measurement phase, i.e. until the device doesn't switch to pure tracking mode (observed when battery residual capacity falls below 10 %), the operation of the MCU is driven by interrupt signals generated by the IMU.

The components discussed so far respond to external events with periodic behavior, characterized by repeating the consumption patterns we have analyzed. On the other hand, the transmission section of the system, which includes the GPS and LoRaWAN modules, behaves differently. Indeed, depending on environmental conditions, the device may or may not successfully complete a transmission. To provide a general idea of how our setup reacts under different conditions, we analyzed the GPS in two scenarios. First, we considered a hypothetical poor visibility condition by placing the GPS indoors, where only a small portion of the sky was visible through a window.

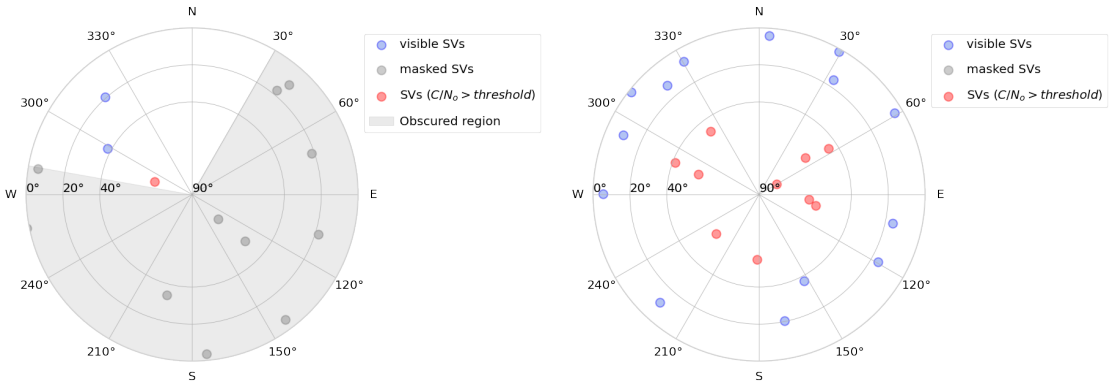

(a) (b)

Figure 5.6: GPS satellite visibility under two conditions: (left) limited sky visibility with only one satellite detected; (right) clear sky visibility with multiple satellites detected.

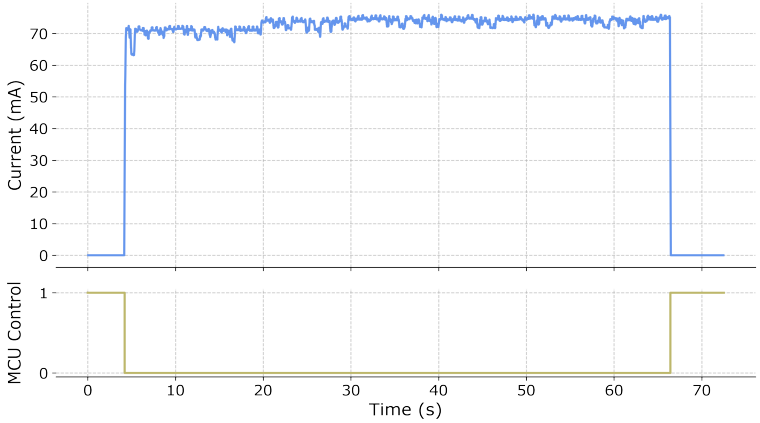

Figure 5.7: GPS current consumption in full performance mode: activation lasts 60 s with poor visibility and 40 s with clear visibility.

The corresponding satellite map is shown in Figure 5.6(a), and the observed consumption profile is reported in Figure 5.7. As we can see, under these conditions, only one satellite was visible with satisfactory signal strength, resulting in a current drawn of 70 mA for approximately one minute. After this time the MCU states that it was not worth to attempt a position fix, given that not enough satellites were visible with strong signals. A second test was performed with a clear sky view. The corresponding map is shown in Figure 5.6(b), while the consumption profile maintains the same shape, but with a shorter duration, decreasing to 40 s. This is fully consistent with the expectations, as we previously noted, since the GPS is being used in its full performance mode in this setup.

Table 5.1: States, variables and their values for LoRaWAN unacknowledged transmission.

Description	Duration		Current absorption	
	Variable	**Value (ms)**	**Variable**	**Value (mA)**
wake up	T_{wu}	168.2	I_{wu}	22.12
radio preparation	T_{pre}	83.8	I_{pre}	13.33
transmission	T_{tx}	1155.1	I_{tx}	83.04
wait 1st window	T_{w1w}	983.3	I_{w1w}	27.05
1st receive window	T_{rx1w}	262.0	I_{1w}	38.16
wait 2nd window	T_{w2w}	738.0	I_{w2w}	27.17
2nd receive window	T_{rx2w}	33.0	I_{2w}	35.08
radio off	T_{off}	170.0	I_{off}	13.29

Concerning the LoRaWAN module, taking into account all the scenarios and possible combinations goes beyond the aim of this work. A wide description of all the parameters involved in a LoRaWAN infrastructure is proposed by [Casals et al. (2017)], where the consumption profiles of various LoRaWAN modules are presented. Based on their work and focusing on the Semtech-SX1276 chip mounted in our system, we were able to derive the behavior described in Table 5.1. We assumed spreading factor (SF) 12, transmission power set to default value of 14 dBm, default settings for both the receive windows (data rate DR0, so SF 12), and confirmed uplink messages with no retransmission mechanisms upon failure.

The Table reports the consumption required for the operation of both the MCU and the LoRaWAN module in the case of an unacknowledged single byte transmission. The variable T_{tx} has been estimated using an Air Time calculator[2]. Its value increases to $T_{tx} = 2138.1$ ms when transmitting the 32 bytes containing the time stamp and the position latitude and longitude.

Acknowledged transmissions show minimal differences, so we focused on unacknowledged ones. Indeed, these represent a worst-case operating condition with higher power requirements, as both receiving windows must remain open until their timeout.

[2]https://www.thethingsnetwork.org/airtime-calculator (last accessed: August 29, 2024)

Table 5.2: Power consumption and activation times during different operating modes of all components analyzed in the system.

Description	Duration		Current absorption	
	Variable	**Value (s)**	**Variable**	**Value (mA)**
IMU (active)	$T_{\text{IMU, run}}$	†	$I_{\text{IMU, run}}$	3.90
IMU (pwr)	$T_{\text{IMU, pwr}}$	△	$I_{\text{IMU, pwr}}$	$67 \cdot 10^{-3}$
Micro SD	$T_{\text{SD, on}}$	1.29	$I_{\text{SD, on}}$	24.44
MCU (active)	$T_{\text{MCU, on}}$	†	$I_{\text{MCU, on}}$	13.90
MCU (standby)	$T_{\text{MCU, pwr}}$	△	$I_{\text{MCU, pwr}}$	$21 \cdot 10^{-3}$
GPS (good view)	$T_{\text{GPS, good}}$	40.00	$I_{\text{GPS, good}}$	71.33
GPS (bad view)	$T_{\text{GPS, bad}}$	60.00	$I_{\text{GPS, bad}}$	71.33
GPS (max timeout)	$T_{\text{GPS, max}}$	$7 \cdot 60$	$I_{\text{GPS, max}}$	71.33
LoRaWAN (dummy)	$T_{\text{tx, dummy}}$	3.59	$I_{\text{tx, dummy}}$	44.76
LoRaWAN (packet)	$T_{\text{tx, pos}}$	4.58	$I_{\text{tx, pos}}$	52.98

For a complete overview, the results of the power consumption analysis are summarized in Table 5.2.
The values marked by asterisks † and △ correspond to operating modes dependent on external conditions. Specifically, when the IMU operates in active mode, the MCU will also be active. Conversely, when the IMU enters power-down mode, the MCU reverts to standby.
The last two rows, dedicated to LoRaWAN operation, highlight the unacknowledged transmission of a dummy byte, used to trigger GPS activation, and the transmission of an entire packet containing the measured position and timestamp. As previously observed, the analysis is limited to unacknowledged transmissions, as they represent a worst case scenario. Additionally, even in the case of acknowledged transmissions, the consumption profile would still be dominated by the transmission peak, leading to negligible variations.

At this point, these results can be used to estimate the battery life of our setup. For this estimation, we will consider a worst case scenario, in which the device is continuously moving and never entering low-power state. Additionally, for the transmission section, we will assume that the first dummy byte sent by the LoRaWAN module is acknowledged, forcing GPS activation. However, we will also suppose poor visibility conditions, resulting in the GPS remaining active

for the maximum time before stating that a position lock is unachievable, after which it reverts to low-power state. The analysis is based on hourly consumption to facilitate battery life estimation. Starting with the operation of the MCU and the IMU, assuming continuous movement, their combined consumption is:

$$c_{\text{mcu, imu}} = \frac{(I_{\text{IMU, run}} + I_{\text{MCU, on}}) \cdot 3600\,\text{s}}{3600\,\text{seconds per hour}} \approx 17.8\,\text{mAh} \tag{5.1}$$

To account for the consumption of the Micro SD, we first estimate its activation frequency. Each timestamp logged into the SD occupies 64 bytes, including years, months, days, hours, minutes, seconds, and the 6 axes measurements provided by the IMU. Considering that a 32768-bytes SRAM region was reserved for the measurements, and the sampling frequency is 1 Hz, the Micro SD activation period is:

$$T_{\text{SD, act}} = \frac{32768\,\text{bytes}}{64\,\text{bytes/s}} = 512\,\text{s} \tag{5.2}$$

This corresponds to approximately $N_{\text{SD, on}} = 7$ activations per hour.
The Micro SD consumption is derived as:

$$c_{\text{sd}} = I_{\text{SD, on}} \cdot \frac{T_{\text{SD, on}}}{3600\,\text{seconds per hour}} \cdot N_{\text{SD, on}} \approx 0.061\,\text{mAh} \tag{5.3}$$

The daily consumption of these three components can be estimated as follows:

$$c_{\text{measure, daily}} = (c_{\text{mcu, imu}} + c_{\text{sd}}) \cdot 24 = 428.66\,\text{mAh per day} \tag{5.4}$$

This allows for a direct comparison with the daily activation of the GPS and LoRaWAN section. The energy consumption for these two last components, considering the worst-case scenario of acknowledging the first dummy packet and a GPS failure (only the last timeout blocks the GPS functioning) lasting seven minutes, is given by:

$$\begin{aligned} c_{\text{GPS, failure}} &= \frac{I_{\text{tx, dummy}} \cdot T_{\text{tx, dummy}} + (I_{\text{GPS, max}} + I_{\text{MCU, on}}) \cdot T_{\text{GPS, max}}}{3600\,\text{seconds per hour}} \\ &\approx 9.99\,\text{mAh} \end{aligned} \tag{5.5}$$

that corresponds also to the daily consumption, as these two components are activated only once per day.

Considering the overall capacity (c_{bat}) of the 3200 mAh battery, and assuming that 10 % of the battery is reserved solely for the retrieval operation (i.e. only the localization procedure is performed), the expected life time (L) is equals to:

$$L = \frac{c_{\text{bat}} \cdot 0.9}{c_{\text{GPS, failure}} + c_{\text{measure, daily}}} \approx 6.57 \,\text{days} \tag{5.6}$$

Thus, in a worst case scenario the device is expected to operate at least for almost one week. However, in these conditions, the energy consumption is dominated by the MCU, as the implemented duty cycling techniques have no effect when this component is continually powered.
On the other hand, thanks to the optimization strategies, we are able to limit the consumption of power-hungry components, particularly the ones involved in the transmission section. Additionally, the current drawn during low-power operation is more than two orders of magnitude smaller than the one required during the device's active-mode. This guarantees a negligible contribution during power-down phases, which is essential to extend the system's battery life during low-intensity events. For instance, if the active time halves, the lifetime doubles to 2 weeks. Conversely, one month operation is achievable if the device actively moves 25 % of the initial estimation (i.e. less than 6 hours per day).

Furthermore, by focusing on the 10 % of the battery capacity dedicated for the retrieval operation, and assuming to attempt a position lock two times per day to facilitate on the field recovery phases, the device can operate for more than two weeks. This provides a wide enough window during which the system can be retrieved from the coastal area relying on the GPS positioning. After this period, the retrieval operation must rely on the backup solution, which uses RFID tags.

This analysis confirms the promising nature of this preliminary work. Significant progress is achieved compared to earlier studies, which only guaranteed a few hours of operation. Moreover, further improvements may be adopted. For instance, the above calculations suggest that lowering the clock frequency of the microcontroller could be beneficial. Another possibility could be the employment of more efficient hardware components. However, this necessitates additional studies and improvement that will be part of future works.

6

Conclusions and future works

The aim of this thesis was to present an energy efficient embedded system for the monitoring of sediment inertial dynamics in a continuous and punctual way. In particular, the solution described in this work focuses on the realization of a system capable to perform these measurements directly in-situ over extended operating periods. Readily available boards were used, focusing on their preliminary integration by implementing dedicated optimization techniques, capable to guarantee a battery life long enough to assess phenomena that up to now remain completely unmonitored.

Regarding the power optimization, standard duty-cycling techniques were combined with event-driven approaches capable to guarantee active operation only when consistent movement is detected. The IMU was programmed to read the accelerometer axes at low frequency intervals. This, combined with the optimization of the MCU, guarantees months of operation in the case the smart pebble remains relatively static. Data logging was guaranteed by using an on board Micro SD, with power consumption minimized thought SRAM-based data buffering, which maximizes writing throughput. This operation was possible by efficiently managing the MCU firmware operation.
For the device retrieval, a passive RFID tag enables manually scanning of the beach to recover the sediment. Additionally, an infrastructure based on a GPS and a LoRaWAN network attempts to track the device's position once per day. Despite not being tested in a real world scenario, this transmission system was optimized to attempt position locks only during high visibility windows, min-

imizing activation time. As a result, daily power consumption is reduced to more than one order of magnitude lower than the MCU's continuous operation energy use.
In addition to facilitating on-field recovery, this tracking methodology enables the reconstruction of overall sediment movement patterns, offering valuable additional information.

The developed system guarantees a lifetime of approximately one week under worst case operating conditions. This performance is expected to improve with the actual implementation of the system. Direct long-run field tests were not carried out, as the design requires a PCB implementation, which will enable the deployment into real smart pebbles or by using customized plastic casings to insulate the system from the underwater environment. Anyway, even under worst-case operating conditions, the system guarantees a longer operational lifetime compared to similar solutions presented in the literature.

Additional work will be done in the future to increase the system's efficiency. Specifically, implementing more software-based energy-efficient optimizations is one option that will be explored. Adaptive techniques that adjust the sampling rate based on the intensity of movement and transmission policies that activate the localization section only when a considerable displacement from the prior position is detected are a couple of examples.
In conclusion, the efficiency of the embedded system was demonstrated through its development and laboratory testing for the purpose of validating the measured data. In the near future, it is expected that the entire architecture will be implemented in a real-world setting.

List of Figures

List of Tables

Bibliography

Abeywardana, D. K., A. P. Hu, and N. Kularatna (2012). "IPT charged wireless sensor module for river sedimentation detection". In: *2012 IEEE Sensors Applications Symposium Proceedings*, pp. 1–5.

Akeila, E., Z. Salcic, and A. Swain (2010). "Smart Pebble for Monitoring Riverbed Sediment Transport". In: *IEEE Sensors Journal* 10.11, pp. 1705–1717.

Alhusban, Z. and M. Valyrakis (June 2020). "Studying sediment transport dynamics by using the Smart sphere". In: *International Journal of Energetica*.

Alippi, C., G. Anastasi, M. Di Francesco, and M. Roveri (2009). "Energy management in wireless sensor networks with energy-hungry sensors". In: *IEEE Instrumentation & Measurement Magazine* 12.2, pp. 16–23.

Benelli, G., E. Panzardi, A. Pozzebon, D. Bertoni, and G. Sarti (2011). "An analysis on the use of LF RFID for the tracking of different typologies of pebbles on beaches". In: *2011 IEEE International Conference on RFID-Technologies and Applications*, pp. 426–431.

Benelli, G. and A. Pozzebon (2013). "RFID Under Water: Technical Issues and Applications". In: *Radio Frequency Identification*. Ed. by Mamun Bin Ibne Reaz. Rijeka: IntechOpen. Chap. 18.

Benelli, G., A. Pozzebon, D. Bertoni, and G. Sarti (2012). "An RFID-Based Toolbox for the Study of Under- and Outside-Water Movement of Pebbles on Coarse-Grained Beaches". In: *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 5.5, pp. 1474–1482.

Benelli, G., A. Pozzebon, G. Raguseo, D. Bertoni, and G. Sarti (2009). "An RFID Based System for the Underwater Tracking of Pebbles on Artificial Coarse Beaches". In: *2009 Third International Conference on Sensor Technologies and Applications*, pp. 294–299.

Bertoni, D., G. Sarti, G. Benelli, et al. (2010). "Radio Frequency Identification (RFID) technology applied to the definition of underwater and subaerial coarse sediment movement". In: *Sedimentary Geology* 228.3, pp. 140–150. ISSN: 0037-0738.

Bertoni, D., G. Sarti, G. Benelli, and A. Pozzebon (2012a). "In situ abrasion of marked pebbles on two coarse-clastic beaches (Marina di Pisa, Italy)". In: *Italian Journal of Geosciences* 131.2, pp. 205–214. ISSN: 2038-1719.

Bertoni, D., G. Sarti, G. Benelli, A. Pozzebon, and G. Raguseo (2012b). "Transport trajectories of smart pebbles on an artificial coarse-grained beach at Marina di Pisa (Italy): Implications for beach morphodynamics". In: *Marine Geology* 291-294, pp. 227–235.

Campagnaro, F., N. Toffolo, A. Pozzebon, R. Francescon, A. Barausse, L. Airoldi, and M. Zorzi (2022). "A Network Infrastructure for Monitoring Coastal Environments and Study Climate Changes in Marine Systems". In: *OCEANS 2022, Hampton Roads*, pp. 1–8.

Cappelli, I., A. Fort, M. Mugnaini, S. Parrino, and A. Pozzebon (2022). "Underwater to above water LoRaWAN networking: Theoretical analysis and field tests". In: *Measurement* 196, p. 111140. ISSN: 0263-2241.

Casals, L., B. Mir, R. Vidal, and C. Gomez (2017). "Modeling the Energy Performance of LoRaWAN". In: *Sensors* 17.10. ISSN: 1424-8220.

Diet, A., Y. Le Bihan, C. Conessa, F. Alves, M. Grzeskowiak, M. Benamara, G. Lissorgues, M. Biancheri-Astier, and A. Pozzebon (2016). "LF RFID chequered loop antenna for pebbles on the beach detection". In: *2016 46th European Microwave Conference (EuMC)*, pp. 41–44.

Fitri, A., R. Hashim, S. Abolfathi, and K. N. Abdul Maulud (2019). "Dynamics of Sediment Transport and Erosion-Deposition Patterns in the Locality of a Detached Low-Crested Breakwater on a Cohesive Coast". In: *Water* 11.8, p. 1721.

Gaël, Many (2016). "Particle assemblage characterization in the Rhone River ROFI". In: *Journal of Marine Systems* 157, pp. 39–51. ISSN: 0924-7963.

Ghaffarivardavagh, R., S. S. Afzal, O. Rodriguez, and F. Adib (2020). "Underwater Backscatter Localization: Toward a Battery-Free Underwater GPS". In:

Proceedings of the 19th ACM Workshop on Hot Topics in Networks. HotNets '20. New York, NY, USA: Association for Computing Machinery, pp. 125–131. ISBN: 9781450381451.

Grottoli, E., D. Bertoni, A. Pozzebon, and P. Ciavola (2019). "Influence of particle shape on pebble transport in a mixed sand and gravel beach during low energy conditions: Implications for nourishment projects". In: *Ocean & Coastal Management* 169, pp. 171–181. ISSN: 0964-5691.

Grottoli, E., B. Duccio, P. Ciavola, and A. Pozzebon (2015). "Short term displacements of marked pebbles in the swash zone: Focus on particle shape and size". In: *Marine Geology* 367, pp. 143–158. ISSN: 0025-3227.

Haihong, Zhao (2011). "A study of sediment transport in a shallow estuary using MODIS imagery and particle tracking simulation". In: *International Journal of Remote Sensing* 32.21, pp. 6653–6671.

Jaud, M., F. Grasso, N. Le Dantec, R. Verney, C. Delacourt, J. Ammann, J. Deloffre, and P. Grandjean (2016). "Potential of UAVs for Monitoring Mudflat Morphodynamics (Application to the Seine Estuary, France)". In: *ISPRS International Journal of Geo-Information* 5.4.

Lu, D., S. Song, J. Wang, Y. Cai, and H. Shen (2022). "Review on the development of SINS/DVL underwater integrated navigation technology". In: *Control Theory Appl* 39, pp. 1159–1170.

Migliorini, M. and A. Pozzebon (2023). "Thermoelectric Generators (TEG) for the Powering of Energy-Hungry LoRaWAN-Based Sensor Nodes in Industrial Applications". In: *2023 IEEE International Workshop on Metrology for Industry 4.0 & IoT (MetroInd4.0&IoT)*, pp. 24–29.

Mondal, S. and R. Paily (2017). "Efficient Solar Power Management System for Self-Powered IoT Node". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 64.9, pp. 2359–2369.

Olumide, A., O. O. Thomas, and D. Karim (2023). "Energy harvesting techniques for sustainable underwater wireless communication networks: A review". In: *e-Prime - Advances in Electrical Engineering, Electronics and Energy* 5, p. 100265. ISSN: 2772-6711.

Ouillon, S., P. Douillet, and S. Andréfouët (2004). "Coupling satellite data with in situ measurements and numerical modeling to study fine suspended-sediment transport: a study for the lagoon of New Caledonia". In: *Coral Reefs* 23.1, pp. 109–122.

Peruzzi, G. and A. Pozzebon (2020). "A Review of Energy Harvesting Techniques for Low Power Wide Area Networks (LPWANs)". In: *Energies* 13.13, p. 3433.

Pozzebon, A., A. Andreadis, D. Bertoni, and C. Bove (2018a). "A Wireless Sensor Network Framework for Real-Time Monitoring of Height and Volume Variations on Sandy Beaches and Dunes". In: *ISPRS International Journal of Geo-Information* 7.4. ISSN: 2220-9964.

Pozzebon, A., M. Benini, C. Bocci, A. Fort, S. Parrino, and F. Rapallo (2023a). "Grid-layout ultrasonic LoRaWAN-based sensor networks for the measurement of the volume of granular materials". In: *Measurement* 220, p. 113404. ISSN: 0263-2241.

Pozzebon, A. and D. Bertoni (2014). "A wireless waterproof RFID reader for marine sediment localization and tracking". In: *2014 IEEE RFID Technology and Applications Conference (RFID-TA)*, pp. 187–192.

Pozzebon, A., I. Cappelli, F. Campagnaro, R. Francescon, and M. Zorzi (2023b). "LoRaWAN Transmissions in Salt Water for Superficial Marine Sensor Networking: Laboratory and Field Tests". In: *Sensors* 23.10. ISSN: 1424-8220.

Pozzebon, A., I. Cappelli, A. Mecocci, D. Bertoni, G. Sarti, and F. Alquini (2018b). "A Wireless Sensor Network for the Real-Time Remote Measurement of Aeolian Sand Transport on Sandy Beaches and Dunes". In: *Sensors* 18.3. ISSN: 1424-8220.

Pucino, N., D. M. Kennedy, R. C. Carvalho, and et al. (2021). "Citizen science for monitoring seasonal-scale beach erosion and behaviour with aerial drones". In: *Scientific Reports* 11, p. 3935.

Raghunathan, V., S. Ganeriwal, and M. Srivastava (2006). "Emerging techniques for long lived wireless sensor networks". In: *IEEE Communications Magazine* 44.4, pp. 108–114.

Sear, D., M. Lee, M. Collins, and P. Carling (2002). "The Intelligent Pebble: A New Technology for Tracking Particle Movements in Fluvial and Littoral En-

vironments". In: *In Proceedings of Erosion and Sediment Transport Measurement: Technological and Methodological Advances Workshop*. Oslo, Norway: Association for Computing Machinery, pp. 19–21.

Spazzapan, M., J. Petrovi, and M. Miko (2004). "New Tracer for Monitoring Dynamics of Sediment Transport in Turbulent Flows". In: *ActaHydrotech*, pp. 135–138.

Sundaresan, S., I Koren, Z Koren, and C. M. Krishna (2009). "Event-driven adaptive duty-cycling in sensor networks". In: *International Journal of Sensor Networks* 6.2, pp. 89–100.

Wu, P., W. Nie, and Y. Liu (2024). "Improving the underwater navigation performance of an IMU with acoustic long baseline calibration". In: *Satell Navig*.

Yangdong, Li and Li Xiaofeng (2016). "Remote sensing observations and numerical studies of a super typhoon-induced suspended sediment concentration variation in the East China Sea". In: *Ocean Modelling* 104, pp. 187–202. ISSN: 1463-5003.

Yoerger, D. R., M. Jakuba, A. M. Bradley, and B. Bingham (2007). "Techniques for deep sea near bottom survey using an autonomous underwater vehicle". In: *Robotics Research: Results of the 12th International Symposium ISRR*. Springer, pp. 416–429.