



POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

CMOS Capacitive Dust Detection System for Space Applications Characterized with Lunar Simulant

TESI DI LAUREA MAGISTRALE IN
ELECTRONICS ENGINEERING - INGEGNERIA ELETTRONICA

Author: **Marco La Barbera**

Student ID: 243244

Advisor: Prof. Marco Sampietro

Co-advisors: Francesco Zanetto

Academic Year: 2025-2026

Abstract

In space, where physical access is extremely limited or even entirely impractical, satellites, spacecraft, and large scientific instruments, such as telescopes, are difficult to repair once deployed. Given the high costs of development and launch, it is important to ensure both reliability and cleanliness, anticipating or monitoring potential failures before they become critical. One particularly relevant aspect concerns long life missions or vibrations generated during launch, which can release dust particles that may then deposit on satellite instrumentation, compromising its functionality. For this reason, an effective solution is represented by the integration of a capacitive sensor implemented in CMOS technology.

This work presents the current development of a system designed to support sensor integration, with specific attention to EMC, power, and dimensional constraints. The proposed architecture is based on a modular two-board system (Carrier and Motherboard) that integrates a high-precision 16-bit ADC and an ultra-low-power microcontroller. The system supports autonomous data logging and wireless communication via Bluetooth Low Energy (BLE). A specific firmware was developed to manage signal acquisition, data processing, and system communication.

Finally, the performance of a previous sensor generation was validated through an experimental campaign involving lunar dust simulant. The results demonstrated a high correlation between the electrical signals and the physical particle size, achieving a reliable detection accuracy and allowing for a real-time reconstruction of the dust grain distribution.

Keywords: electronic dust sensor, impedance dust sensor, impedance dust counter, wireless system, space application, lunar simulant, CMOS

Abstract in lingua italiana

Nello spazio, dove l'accesso fisico è estremamente limitato o del tutto impraticabile, satelliti, veicoli spaziali e grandi strumenti scientifici, come i telescopi, risultano difficili da riparare una volta operativi. Dati gli elevati costi di sviluppo e lancio, è fondamentale garantire sia l'affidabilità che la pulizia, anticipando o monitorando potenziali guasti prima che diventino critici. Un aspetto particolarmente rilevante riguarda le missioni di lunga durata o le vibrazioni generate durante il lancio, che possono rilasciare particelle di polvere in grado di depositarsi sulla strumentazione di bordo, compromettendone la funzionalità. Per questo motivo, un'efficace soluzione è rappresentata dall'integrazione di un sensore capacitivo realizzato in tecnologia CMOS.

Questo lavoro presenta lo sviluppo di un sistema progettato per supportare l'integrazione del sensore, con specifica attenzione ai vincoli di compatibilità elettromagnetica (EMC), alimentazione e dimensioni. L'architettura proposta si basa su un sistema modulare a due schede (Carrier e Motherboard) che integra un ADC a 16 bit ad alta precisione e un microcontrollore a bassissimo consumo. Il sistema supporta il data logging autonomo e la comunicazione wireless tramite Bluetooth Low Energy (BLE). È stato inoltre sviluppato un firmware specifico per gestire l'acquisizione dei segnali, l'elaborazione dei dati e la comunicazione di sistema.

Infine, le prestazioni di una generazione precedente del sensore sono state validate attraverso una campagna sperimentale con simulante di polvere lunare. I risultati hanno dimostrato un'elevata correlazione tra i segnali elettrici e la dimensione fisica delle particelle, raggiungendo un'accuratezza di rilevamento affidabile e consentendo la ricostruzione in tempo reale della distribuzione granulometrica della polvere.

Parole chiave: sensore di polvere elettronico, sensore di polvere a impedenza, contatore di particelle di polvere, sistema senza fili, applicazioni spaziali, simulante lunare, CMOS

Contents

Abstract	i
Abstract in lingua italiana	iii
Contents	v
1 Introduction	1
1.1 Dust detection	1
1.2 Standard detection methods	2
1.3 Capacitive sensing	2
2 Capacitive Sensor	5
2.1 Working principle	5
2.2 Signal Acquisition Chain Architecture	8
2.2.1 Readout Front End	9
2.2.2 Lock-in Amplifier: demodulation and LPF	12
2.2.3 DCC: Digital to Capacitance Converter	14
2.2.4 Ring Oscillator	17
3 Boards and Hardware design	23
3.1 Carrier Board	25
3.1.1 Power Signals	26
3.1.2 Digital Signals	28
3.1.3 ADC	30
3.1.4 Stability and Antialiasing Filter	32
3.1.5 Signal Handling, waveform and timing	34
3.2 Motherboard	37
3.2.1 Bluetooth PCB design	38
3.2.2 User LED and visual feedback	41

3.2.3	BOOT0 Switch and System Reset	42
3.3	Power Management System	43
3.3.1	External Power Supply	43
3.3.2	USB-C power supply	44
3.3.3	Power multiplexer	44
3.3.4	Filtering stage	45
3.3.5	Power regulation stage: CP2102 and LT3045	46
3.4	Connection within Carrier and Motherboard	50
4	Firmware and GUI	53
4.1	Firmware Architecture and Implementation	53
4.1.1	Data Acquisition and Bluetooth Transmission	54
4.1.2	Local Data Logging and microSD Management	57
4.1.3	Particle Counting Algorithm	58
4.1.4	CP2102N USB-to-UART Bridge Configuration	61
4.2	Graphical User Interface	63
4.2.1	Connection and Data Logging	63
4.2.2	Real-Time Visualization and Hardware Status	64
4.2.3	Advanced Dashboard and Material Configuration	66
4.2.4	Offline SD Card Analysis	67
5	Experiments and Results	69
5.1	Old sensor layout	69
5.2	Old sensor measurements	71
5.2.1	Image Registration and post-processing	72
5.2.2	Comparison with acquired sensor data	74
5.2.3	Fitting method to predict particle size	78
5.2.4	Specific and particular cases	82
6	Conclusions and future developments	85
	Bibliography	87
	List of Figures	89
	List of Tables	93

1 | Introduction

1.1. Dust detection

In the harsh environment of space, where physical access is extremely limited or entirely impractical, satellites, spacecraft, and large-scale scientific instruments, such as space telescopes, are nearly impossible to repair once deployed. Given the significant costs associated with the development, testing, and launch of these assets, ensuring long-term reliability and cleanliness is a primary concern. A critical threat to mission success is the presence of micro-particles and dust.

During the launch phase, the extreme mechanical vibrations generated by the carrier rocket can cause the release of dust particles from the internal surfaces of the satellite. In a microgravity environment, particles can eventually deposit on sensitive instrumentation. The consequences of such deposition are severe: the obscuration of mirrors and lenses in optical payloads drastically reduces the quality of scientific data, while dust accumulation on solar arrays leads to significant power degradation. Furthermore, the abrasive and electrostatic nature of lunar regolith poses a constant threat to moving mechanical parts and electronic integrity. Consequently, *in-situ* dust monitoring has evolved from a secondary requirement to a fundamental necessity for anticipating potential failures before they become mission-critical [1].

The categories of spaceborne dust and particulate matter (PM) comprise a wide variety of solid particles found in conditions ranging from atmospheric suspension to the vacuum of space, including deposits on the lunar regolith or Martian soil. The origins of such particles is attributed to both natural and human-derived mechanisms. In the natural context, this involves the electrostatic lofting of extraterrestrial regolith, a phenomenon propelled by solar winds and surface electrical fields [2], as well as the presence of microscopic aerosols in the air. On the other hand, anthropogenic sources are primarily linked to the structural degradation of spacecraft and satellites; specifically, the rigorous conditions of extreme thermal oscillations and persistent UV exposure cause the erosion and flaking of external coatings and polymeric materials [3].

1.2. Standard detection methods

Historically, spaceborne dust monitoring has relied on a few established technologies. However, these traditional methods often struggle to meet the increasingly stringent requirements for miniaturization and integration in modern space missions. The most common detection techniques include:

- **Gravimetric Method:** This technique involves an air pump and a series of cascading filters with specific porosities designed to capture particles of targeted sizes. While it allows for precise size discrimination, the filters must be manually removed and weighed to determine concentration [4]. The primary downsides of this approach are its bulkiness, high power consumption, the inability to provide real-time data, and its general impracticality for space-based platforms.
- **Light Scattering Methods:** These systems utilize a laser beam pointed transversely through a small air channel where particles flow. Photodetectors sense the intensity and angle of the light scattered by passing particles to determine their size in real time [5]. Although capable of detecting particles as small as $0.5\mu m$, these sensors remain unsuitable for space applications due to their significant physical footprint, power requirements, and high susceptibility to inaccuracies caused by microscopic vibrations.
- **MEMS Resonators:** To overcome the bulkiness of traditional methods, Micro-Electro-Mechanical Systems (MEMS) can be employed as mass-sensitive resonators. The resonant frequency of these devices shifts depending on the mass of the deposited particulate matter, with the rate of frequency change corresponding to the particle concentration [6]. While these devices offer easy integration for space-based applications, they are limited by their inability to provide detailed information regarding the granulometry (size distribution) of the sampled dust.

The common drawback of these solutions is their high impact on the satellite's resources in terms of size, power and reliability. As the industry shifts towards smaller platforms like CubeSats, the need for a more compact, low-power and reliable detection technology has become critical.

1.3. Capacitive sensing

Capacitive sensing represents one of the most promising alternatives to overcome the limitations of traditional methods. The detection principle is based on measuring the

variation in capacitance between two electrodes when a dust particle enters their electric field. Due to the high dielectric constant of typical space dust materials, such as silicates, even particles as small as a few micrometers can induce detectable capacitance changes in the attofarad (10^{-18} F) or femtofarad (10^{-15} F) range.




Figure 1.1: Capacitive variation as function of time.

The primary advantage of this approach is its compatibility with CMOS technology. This allows for the integration of both the sensing electrodes and the readout electronics on a single silicon chip, drastically reducing noise and enabling unprecedented miniaturization. However, transforming a CMOS sensor into a fully operational space instrument requires a robust hardware and firmware ecosystem capable of high-precision acquisition, offset management, and autonomous communication.

This work focuses on the development of such a system. An architecture based on a modular two-board design (Carrier and Motherboard) was developed to support the integration of a capacitive CMOS sensor. The system incorporates a high-precision acquisition chain and a low-power microcontroller for real-time data processing. Finally, the performance of the system was validated through an experimental campaign at ESA-ESTEC laboratories using lunar dust simulant, demonstrating the reliability of the proposed solution for future space missions.

2 | Capacitive Sensor

The sensor, developed by the I3N Laboratory at Politecnico di Milano [7], is designed to measure typical particle dimensions found in space applications, such as lunar regolith. Its interdigitated structure permits the measurement of particles ranging from $5\mu m$ to $100\mu m$, which is the specific target range for this device, within a $12mm^2$ active sensing area. The sensor is composed of 32 independent channels, each of one with an $1880\mu m \times 202\mu m \simeq 0.37976mm^2$ active area, where dust deposition can occur.

2.1. Working principle

When two electrodes are placed side-by-side on a flat surface, the electric field between them forms a circular pattern. This field is usually most intense at the inner edges where the electrodes are closest to each other. The shape and strength of this field change significantly whenever a particle is nearby. This effect becomes much stronger the closer the particle gets to the electrodes. Because of this, particles that land right in the small gap between the electrodes cause the largest change in capacitance. Essentially, these two side-by-side electrodes work like a capacitor, and its value depends on the type of material located in the space between them. The expected capacitive variation per micrometer for this implementation is given by the following formula 2.1:

$$C'_{el} = 2\epsilon_0 \frac{\epsilon_{ox} + \epsilon_{air}}{2\pi} \ln \left[\left(1 + \frac{2W}{G} \right) + \sqrt{\left(1 + \frac{2W}{G} \right)^2 - 1} \right] \quad (2.1)$$

where G represents the gap between the electrodes, W is their width, $\epsilon_{ox}(\simeq 3.9)$ and $\epsilon_{air}(\simeq 3.9)$ are the relative permeabilities of the oxide layer and of the air between the electrodes.

The chosen dimension for this application for the electrodes are $W = 4\mu m$ and a gap $G = 4\mu m$. With these sizes, the capacitance per micrometer is about $24.3 \text{ aF}/\mu m$. Given that the electrodes are $1860\mu m$ long, the total capacitance (C_{el}) comes out to roughly

45.3 fF. The change in capacitance depends on the size of the particle compared to the total "sensing volume" around the electrodes. It is possible to visualize this volume as a cylinder with a radius of $6\mu m$ (calculated as $W + G/2$) and a length of $1860\mu m$. Since the sensor is surrounded by both air and oxide, we only calculate the signal (ΔC) based on the top half of this cylinder, adjusting for the different materials. From this data, a theoretical estimation of the signal produced by a small PDMS sphere (a material with $\epsilon_r \approx 2$) with radius R and diameter D as depicted in Figure 2.1 was computed using the following formula:

$$\Delta C_{el} \approx C_{el} \cdot \frac{V_{Sphere}}{V_{total}} \cdot \frac{\Delta \epsilon_r}{\epsilon_r} = 45.3 \text{ fF} \cdot \frac{1}{1 + \epsilon_{ox}} \frac{\frac{4}{3}\pi R^3}{\pi(W + G/2)^2 \cdot \frac{L}{2}} \cdot \frac{\Delta \epsilon_r}{\epsilon_r} \approx 6 \text{ aF} \quad (2.2)$$




Figure 2.1: Coplanar electrodes architecture and electric field visualization.

One key detail is that while the total capacitance is tied to the length of the electrodes, while the actual change in capacitance (ΔC_{el}) caused by a particle stays the same regardless of how long the electrodes are.

The sensor architecture consists of multiple channels with interlocking "fingers", known as an interdigitated design. Each channels presents a differential sensing design to cancel out background noise by comparing two separate signals. This approach keeps the capacitive baseline as low as possible, limited only by minor physical differences between the electrodes. The differential structure also helps to filter out interference from the electri-

cal square-wave signals used to power the sensor. The environmental variations such as humidity and temperature are detected as common-mode effects; since they affect both signals equally, the system can ignore them and focus strictly on particle detection.

The electrodes are energized by two opposite square-wave signals that are 180 degrees out of phase (when one is high, the other is low) as it can be seen in Figure 2.2. To prevent signal loss when large particles land across adjacent sections, the design alternates the electrodes belonging to different reading channels. In terms of dimensions, each electrode block has a width of $84\mu m$ and includes five reading fingers (depicted in purple) and six stimulation fingers (in orange). By combining 64 of these blocks, the total active sensing area forms a 32 channels rectangle of approximately $1860\mu m \times 84\mu m \times 64$, resulting in a total area of about $10mm^2$.




Figure 2.2: Electrodes interdigitate structure.

In the Figure below 2.5, it is possible to see the layout of the fingers of the first half of a channel, followed by the first half of the successive channel.




Figure 2.3: Real layout of first half of two consecutive channels.

2.2. Signal Acquisition Chain Architecture

The integrated circuit architecture (Figure 2.5) is designed to transform the minute capacitive variations detected by the electrodes into a robust analog signal suitable for ADC digital sampling. This signal processing chain is organized into four fundamental stages that work to maximize the signal-to-noise ratio. The process begins with the Analog Front-End, which serves as the primary interface to the physical sensor; its role is to convert the charge variations into a voltage signal while providing the necessary transimpedance gain and initial noise suppression.

Following the initial amplification, the signal enters the Mixer stage, where synchronous demodulation occurs to shift the information from the high-frequency excitation carrier down to the baseband. The chain continues with a Low-Pass Filter (LPF), which is responsible for removing high-frequency residuals and modulation artifacts. The resulting output is a clean, conditioned signal that accurately represents the particle detection events, ready to be processed by the system's microcontroller. The Digital to Capacitance Converter (DCC) module, which manages and compensates for offsets present in the electrodes, ensures that the demodulation remains accurate even in the presence of parasitic components or saturation.




Figure 2.4: Sensor signal acquisition chain overview.

2.2.1. Readout Front End

The readout front-end (Figure 2.5) acts as the primary interface between the capacitive sensing electrodes and the subsequent processing stages. Its fundamental role is to detect the small capacitance variations (ΔC) produced by particles depositing over the electrodes, that induces a difference in the current between C_{UP} and C_{DOWN} , and convert them into a proportional voltage signal.

The front-end architecture is typically divided into two main stages to ensure high sensitivity and low noise. The first stage is the Charge Amplifier, which integrates the input current onto a feedback capacitor to produce an initial voltage variation. This is then followed by a 40x gain stage to further increase the initial voltage variation. The primary function of the amplifiers in these stages is to boost the weak sensor signal to a level that can be accurately processed by the mixer and filters, effectively improving the overall Signal-to-Noise Ratio (SNR) of the system.

The implementation of pseudo-resistor within the first feedback loop sets the integral behavior inside the 200kHz signal band. These are implemented using MOSFET transistors biased to operate in the sub-threshold region, where they behave like resistors with extremely high values, often reaching the Giga- Ω range. By combining with the feedback capacitance, they introduce a very low-frequency pole. Furthermore, these pseudoresistors provide a stable DC path to bias the input nodes of the amplifiers, preventing them from saturating due to leakage currents or accumulated charge. Pseudo-resistors can be biased with external voltages in order to tune its value.




Figure 2.5: Front end stage with pseudo resistors configuration.

Figure 2.6: Pseudoresistors bias network implementation at transistor level.

These pads allow to modify the internal bias of the mirroring stages, thereby directly influencing the pseudo-resistance value and the resulting low-frequency cut-off of the sensor (described in section 3.1).

The V_{Ra} pad provides access to the primary branch of the current generation circuit. Under default conditions, if this pad is left floating, the bias current is defined by the internal resistance series $R_1 + R_2$ and the 3.3 V power supply. However, in the PCB, is possible to manually override this current as follows:

- **Voltage Control:** By applying an external voltage to V_{Ra} , the voltage drop across R_2 is redefined. This change alters the current flowing through the first branch of the mirror.
- **Resistance Tuning:** Modifying the voltage at V_{Ra} , the current will go as $i = \frac{V_{Ra}}{R_2 + 1/g_m}$, which subsequently enters in the mirroring stages and modifies the pseudo-resistance value. With the maximum value of V_{Ra} it is possible to obtain the maximum value for the pseudo-resistors.

If needed, in order to bypass the initial mirroring stages, the I_{Rb} pad is connected to the third branch of the bias network:

- **Current Injection:** A small external current can be directly injected into this node to supplement or replace the mirrored current from the previous stages.
- **Noise Mitigation:** The design includes an internal capacitor C_1 (10 pF) to filter out any low-frequency noise that might be introduced through the external probe or connection to the I_{Rb} pad.
- **Protection:** A series resistor R_3 (6.15 k Ω) is integrated between the pad and the internal node to protect the sensitive CMOS circuitry from direct electrical contact and potential ESD events.

The following table describes the values of the bias network circuitry in Figure 2.6:

Parameter	Value
k	3
R_1	200 k Ω
R_2	50 k Ω
R_3	6.15 k Ω
C_1	10 pF

Table 2.1: Implemented values for the tunable bias network components.

2.2.2. Lock-in Amplifier: demodulation and LPF

Mixer

The Mixer represents the second critical stage in the signal processing chain, acting as a synchronous demodulator. Its primary function is to translate the high-frequency modulated signal received from the front-end back down to the baseband, where the particle information can be more easily extracted.

The combination of the mixer, and the low-pass filter constitutes a Lock-in Amplifier (LIA) architecture. By "locking in" to the frequency of the selected square wave, the system effectively ignores any interference or noise that does not occur at that exact frequency and phase.

The operation of this stage relies on the multiplication of the amplified sensor signal by a reference square-wave signal. By performing this multiplication, the mixer effectively performs a frequency down-conversion: the information carried at the excitation frequency is shifted toward DC, while out-of-band noise and interference are shifted to higher frequencies.

The operation of the mixer within the lock-in chain can be described mathematically. Considering the input signal $V_{in}(t)$ coming from the front-end and a reference square wave $V_{ref}(t)$ provided by the timing controller:

$$V_{in}(t) = A_{sig} \sin(\omega_0 t + \phi) + n(t) \quad (2.5)$$

where A_{sig} is the amplitude of the signal, ω_0 is the excitation frequency, ϕ is the phase shift, and $n(t)$ represents the total noise. The mixer performs a multiplication between the input and the reference signal:

$$V_{mix}(t) = V_{in}(t) \cdot V_{ref}(t) \quad (2.6)$$

For simplicity, considering a sinusoidal reference $V_{ref}(t) = A_{ref} \sin(\omega_0 t)$, the result of the multiplication becomes:

$$V_{mix}(t) = \frac{1}{2} A_{sig} A_{ref} \cos(\phi) - \frac{1}{2} A_{sig} A_{ref} \cos(2\omega_0 t + \phi) + V_{ref}(t) \cdot n(t) \quad (2.7)$$

This operation produces two distinct components:

- **A DC component:** $\frac{1}{2}A_{sig}A_{ref}\cos(\phi)$, which is directly proportional to the signal amplitude A_{sig} .
- **A high-frequency component:** at double the excitation frequency ($2\omega_0$), which carries no useful information for the baseband measurement.

The Low-Pass Filter (LPF) is then used to suppress the $2\omega_0$ term and the modulated noise $n(t)$, leaving only the clean DC value as the final output. This process ensures that the capacitive variations caused by particles are detected with high precision, as any noise at frequencies different from ω_0 is shifted far away from the DC band during demodulation.

Gm-C Low Pass Filter

The final stage of the signal conditioning chain is a Low-Pass Filter (LPF) implemented using a G_m -C architecture. This block is placed immediately after the mixer to extract the baseband signal from the demodulated output and to suppress high-frequency components and switching residuals.

The G_m -C topology is chosen for its suitability in integrated CMOS technology, as it allows for the implementation of low cut-off frequencies without requiring large, area-consuming passive resistors. The filter operates based on the principle of transconductance (G_m): a transconductance amplifier converts the input voltage into a current, which then charges a capacitor (C), creating a pole in the frequency response.

The primary purpose of the G_m -C filter is to define the final bandwidth of the sensing system, ensuring that only the low-frequency variations associated with particle detection are passed to the output. It is possible to tune also this structure. Since the cut-off frequency is directly proportional to the transconductance G_m ($f_c \propto G_m$), the filter's behavior can be adjusted by varying the bias current of the amplifiers.

To allow for post-fabrication calibration, the G_m -C bias network is accessible through two dedicated external pads:

- $G_m C_a$: Used to provide an external voltage that defines the current flowing into the G_m -C bias network.
- $G_m C_b$: Used to directly inject or monitor the bias current that determines the transconductance value.




Figure 2.7: Gm-C filters bias network implementation at transistor level.

By interacting with these pads, it is possible to shift the filter's cut-off frequency to optimize the trade-off between noise suppression and the temporal resolution required to detect a greater amount of particles per second.

The final implementation of the bias network was configured to operate at the lowest cutoff frequency observed in simulation, namely 49 Hz, thus setting the settling time for a deposition at around 20ms.

In order to obtain these values, the sizing of the components was chosen as seen in the following

Parameter	Value
k	7
R_1	900 k Ω
R_2	50 k Ω
R_3	6.15 k Ω
C_1	10 pF
C_2	2 pF

Table 2.2: Implemented value for k , resistances and capacitances.

2.2.3. DCC: Digital to Capacitance Converter

The Digital-to-Capacitance Converter (DCC) is a module within the readout chain, designed to ensure capacitive offset cancellation. In a differential capacitive sensing system, unavoidable physical asymmetries between electrodes or layout parasitics can generate a

significant baseline signal. If left uncorrected, this offset could lead to the saturation of the high-gain amplifiers in the front-end, making it impossible to detect the minute variations produced by particles. It can be also possible that there is a large deposition of dust particles on only one of the two electrodes of the channel, thus generating a immediate large ΔC_{el} .

Architecture and Operation

The DCC architecture is based on a bank of programmable capacitors integrated directly at the input nodes of the analog front-end. The system operates as a digitally balanced capacitive bridge: by injecting a specific amount of compensation capacitance, the initial mismatch can be canceled to "zero" the sensor's baseline.

The initial structure of the DCC is an ADC that basically consists of four comparators, that confront the output voltage signal of the low pass filters with two on-chip generated threshold voltages, V_{TL} and V_{TH} . Whenever one of these two threshold is crossed, the comparator switches its output and the digital word is processed by the following DCC logic block. The control logic relies on a 5-bit counter that allows for the selection of different capacitance combinations, offering precise resolution for calibration. As specified in the design, the total implemented capacitance for the DCC is approximately 14 fF, a value optimized to cover the expected mismatch range without degrading overall sensitivity.

The clock signal required to drive the 5-bit counter is produced by a frequency divider stage composed of D flip-flops (Figure 2.8).




Figure 2.8: D flip-flop based frequency divider.

This module processes the primary square-wave signal (CK_{in}) generated by the ring oscillator, scaling its frequency down by a factor of 2^{13} . This significant frequency reduction is necessary to provide the system with an adequate time window for signal stabilization. Specifically, after each adjustment of the counter, the output of the G_m -C filter must be allowed to reach its steady-state value, a process that requires a duration of approximately 5τ :

$$5\tau = 5 \cdot \frac{1}{2\pi f_{G_m} C} \quad (2.8)$$

Once this settling time has elapsed, the output voltages are evaluated against the comparator thresholds according to the following logic:

- **Balanced State:** If the channel has reached a balanced condition where the output voltage falls within the range defined by the two thresholds, the comparator output is updated. Consequently, the counter is stopped, and the calibration loop is opened.
- **Unbalanced State:** If the channel remains unbalanced, the comparators do not trigger the stop condition. Instead, the subsequent rising edge of the counter clock increments or decrements the digital word by one bit. This adjustment modifies the compensation current injected by the DCC, and the balancing status is re-evaluated after another 5τ interval.

Furthermore, the digital word derived from the four comparators is utilized to generate the SIGN bit, as illustrated in Figure 2.11.




Figure 2.9: Digital logic of the DCC.

External Control and Tuning Signals

The user can interact with the DCC calibration process through a series of dedicated PADS, which allow for both automatic and manual compensation management:

- $RESET_{DCC}$: This signal allows for the resetting of the internal DCC counter, returning the compensation to its minimum value to start a new calibration procedure.
- $CK_{DCC,SEL}$: A digital bit used to select between the internal clock (generated by the Ring Oscillator) and an external clock for incrementing the calibration counter.
- $CK_{DCC,PAD}$: A dedicated PAD for inputting a manual clock signal, useful during testing phases to finely adjust the compensation level.
- OUT_D : Provides digital feedback when the DCC calibration process is active or completed, allowing the external microcontroller (STM32WBA) to monitor the balancing status.

The integration of this block allows the chip to maintain a high dynamic range, ensuring that the full amplification is dedicated exclusively to the useful signals originating from dust debris.

2.2.4. Ring Oscillator

The square wave signals for electrode stimulation and mixer synchronization are generated on-chip. The architecture utilizes an internal 12.8 MHz ring oscillator followed by a D flip-flop frequency divider chain.

The operating frequency can be selected by driving two external digital pads A_0 and A_1 . This allows the system to toggle between three predefined frequencies (50 kHz, 200 kHz, and 400 kHz) and an external clock mode for debugging purposes. If these pads are left floating or unbonded, the internal logic sets the default state to $A_0 = 1$ and $A_1 = 0$ selecting the default 200kHz clock frequency.

The configuration logic for the PCB implementation is summarized in the following table:

A_1	A_0	Operating Mode / Frequency
0	0	$f_{CK1} = 400 \text{ kHz}$
0	1	$f_{CK2} = 200 \text{ kHz}$ (Default)
1	0	$f_{CK3} = 50 \text{ kHz}$
1	1	External Clock (Debug Mode)

Table 2.3: Frequency selection logic through digital pads A_0 and A_1 .




Figure 2.10: Overview circuit of a single channel, presenting all the basic elements: Electrodes, Front End, Mixer, DCC, Gm-C filter and Ring Oscillator.




Figure 2.11: Layout of the ASIC without power MET4 layer & PAD description.

PAD	Signal
$V_{ref,el}$	+3.3 V charging the electrodes
GND_{el}	0 V decharging the electrodes
GND	0 V
$TV+$	+1.65 V of the + nodes of the amplifiers
V_{mid}	+1.65 V power supply
V_{TH}	+2.65 V for the comparators
V_{TL}	+0.65 V for the comparators
GmC_a	V to define I flowing in G_mC bias network
GmC_b	I to define I flowing in G_mC bias network
$R_{f,a}$	I to define I flowing in R_{pseudo} bias network
$R_{f,b}$	V to define I flowing in R_{pseudo} bias network
$vss_{an}; vdd_{an}$	Analog power supplies (0 V / +3.3 V)
$OUT+$	Output positive pin
$OUT-$	Output negative pin
$vss_{buff}; vdd_{buff}$	Buffer power supplies (0 V / +3.3 V)
$RESET_{lettura}$	Reset for the output reading MUX
$RESET_{DCC}$	Reset the DCC counter
V_{midCK}	+1.65 V disturbed by the CK
$vss_{dig}; vdd_{dig}$	Digital power supplies (0 V / +3.3 V)
HF_{B1}	BIT_{dig} to control the signal frequency
HF_{B0}	BIT_{dig} to control the signal frequency
SEL	BIT_{dig} to select automatic/manual channel reading
$S1$	BIT_{dig} for manual channel reading
$S2$	BIT_{dig} for manual channel reading
$S3$	BIT_{dig} for manual channel reading
$S4$	BIT_{dig} for manual channel reading
$S0_{CK}$	BIT_{dig} for manual channel reading or reading CK
$CK_{DCC,SEL}$	BIT_{dig} to select internal/external DCC $CK_{COUNTER}$
$CK_{DCC,PAD}$	External CK for DCC counter
OUT_D	BIT_{dig} when DCC is activated
CK_{PAD}	Manual signal frequency
$2CK_{PAD}$	Manual signal at double frequency

Table 2.4: Description of the CMOS sensor PAD signals.

3 | Boards and Hardware design

This chapter details the hardware architecture designed to interface with, control, and configure the sensor. It outlines the key design choices and specifies the electronic components selected for the PCBs.

Since the sensors will eventually be completely covered in dust and require replacement, the system was conceived as a two-board design. The idea for the development of the hardware was inspired by a previous work done in the I3N Laboratory [8] for the sensor of the previous family. In Figure 3.1 it is possible to see a simplified schematic of the system. The carrier board serves as the dedicated support for the sensor, which is soldered via bonding wires. This board is then mounted onto the motherboard, which hosts the microcontroller, that is responsible for managing the entire system architecture, handling everything from the acquisition of input and output sensor signals to the management of user communication protocols. This two-boards structural separation, allows the sensor to be easily swapped multiple times, significantly lowering replacement costs and avoiding the risk of damaging the bonding contacts.




Figure 3.1: Simplified schematic of the System.

The sensor is designed for in-situ measurements, it can be deployed in hard-to-reach locations, such as a satellite's photovoltaic panels. Since the system is engineered for remote operation far from a main computing center, it supports wireless communication, specifically via a Bluetooth Low Energy (BLE) connection. To ensure data integrity during long-term experiments, the architecture integrates a micro SD card slot for local storage, allowing the system to operate autonomously even when a wireless link is unavailable. Nevertheless, a USB-C interface is also available for programming, data extraction, and real-time communication.




Figure 3.2: Motherboard connected to the computer and two unmounted Carrier boards.

The following sections provide a detailed analysis of these two components: the carrier board, focusing on the sensor integration and analog-to-digital conversion, and the motherboard, describing the processing unit and the communication interfaces.

3.1. Carrier Board

To ensure a high degree of modularity and facilitate future integration with diverse flight hardware or updated microcontroller platforms, the sensor board relies exclusively on a digital communication interface. This architectural choice is supported by the fact that the board independently generates all its required analog signals and internal voltage references, such as the $+3.3V_{sensor}$ line provided by the high-precision LT3045 LDO. Consequently, the module remains fully compatible with any external digital system operating at a maximum logic level of 3.3V. The physical implementation is realized on a six-layers, highly compact $20\text{mm} \times 20\text{mm}$ printed circuit board. Despite the high density of interconnects required for the sensor, the design utilizes a six-layer copper stack-up, with routing specifically optimized across three of these layers. All the pads of the sensor have been routed with a precise trace, but the "optional" pads (cited in table 2.4) have been left unsoldered, using missing resistor spots that can be used in the case they are needed. A key feature of the layout is the strategic segregation of components: all auxiliary hardware, that is mainly composed by the ADC, LDO and the connector, is populated on the bottom layer (figure 3.3b). This asymmetrical component distribution leaves a clear area (3.3a) and simplifies the critical soldering and bonding processes necessary for the sensor's assembly and its management through the measurements.




Figure 3.3: Carrier board design details.

The upper section of the PCB layout is explicitly designed to support rigorous testing and debugging activities. Surrounding the central sensor footprint, a series of exposed test pads has been routed to the top layer (figure 3.4). These circular pads provide direct

physical access to critical signal nodes that would otherwise be inaccessible once the main sensor package is mounted.

Specifically, these test points allow for the real-time monitoring of power supply integrity ($VDD_3.3V$, GND), reference levels ($V_{REF_3.3V}$), and key analog voltages such as TV_+ , ADC analog output OUT_N , OUT_P , threshold settings (VT_{high} , VT_{low}), but also digital signals from the microcontroller as channel selection pins ($S0_{ck}$, $S1$, $S2$, $S3$, $S4$) and digital sensor output OUT_D .



Figure 3.4: Detail of the top layer and test pad locations.

3.1.1. Power Signals

The power distribution network of the carrier board, which originates from the primary power management system detailed in section 3.3, is engineered to provide highly stable and low-noise voltage rails, which are essential for the correct and precise functioning of the sensor, and the high-resolution analog-to-digital conversions performed by the ADC:

- **VDD_3.3V:** This primary supply line serves as the fundamental power rail for the entire functional circuitry of the carrier board. It is designed to provide a stable

and continuous voltage level to all active components, originating as a dedicated, pre-filtered voltage source derived directly from the main power regulation stage located on the motherboard.

- **VREF_3.3V:** Specifically dedicated to power-up the sensor at +3.3V, this rail is derived from the VDD_3.3V source and is conditioned through a dedicated low-pass π filter (Figure 3.5). The filtering stage is composed of two $10\mu\text{F}$ capacitors and a 2500Ω @ 100MHz ferrite bead, designed to suppress high-frequency noise and electromagnetic interference.





Figure 3.5: π filter, featuring a ferrite bead and decoupling capacitors for EMI suppression.

- **TV+:** This +1.65V reference signal is generated by a high-precision LT3045 voltage regulator to establish the operating point for the positive nodes of the amplifiers. It serves as the critical reference for the pseudo-resistor stages and the GmC filter bias networks.
- **VT_{high} and VT_{low}:** These voltages define the operational thresholds for the Digital to Capacitance Converter (DCC) comparators, set at +2.65V and +0.65V respectively. Whenever the channel output voltage crosses these limits, the DCC logic toggles its output to activate the balancing mechanism. While the standard board utilizes internal voltages (with resistors R6 to R11 left unpopulated), external thresholds can be implemented by soldering specific resistor values to change the comparator input voltages as it can be seen in Figure 3.3b.
- **V_{Ra}, G_mC_a, V_{Rb}, G_mC_b:** These bias voltages play a critical role in the analog front-end, as they are responsible for establishing and maintaining the precise current flow required within the Gm-C filter stages and the associated pseudo-resistor bias networks. A detailed technical explanation of the operational principles of these specific sub-circuits has been provided in Section 2.2. In the current default hardware configuration, the sensor is designed to operate exclusively through its integrated internal bias networks; this is due to the fact that the external resistor pads, specifically labeled from R17 to R22 on the PCB layout, are not populated.


3.1.2. Digital Signals

The digital interface facilitates the synchronization between the microcontroller and the CMOS sensor, managing channel addressing and data acquisition:


- **SPI3:** This bus handles the communication with the ADC (better described in section 3.1.3) , where the **SPI3_MISO** line transmits the converted data values every 8ms. The bus is shared with the motherboard's microSD slot; timing is managed such that a positive edge on **SPI3_CS** initiates acquisition, while a negative edge signals the ADC to release converted data, all synchronized by a 16MHz clock on the **SPI3_SCK** line.
- **SEL:** This digital bit is used to determine the channel selection logic, allowing the user to switch between the automatic internal channel iteration of the sensor and the manual channel selection programmed within the STM32 microcontroller firmware.
- **S0ck, S1, S2, S3, S4:** These signals provide the binary addressing for the 32 available sensing channels. When the **SEL** signal is high, **S0ck** functions as a PWM input with a period of 8ms and a 50% duty cycle, driven by the microcontroller's TIM2 hardware timer.




Signal S_0




Signal S_1



Signal S_2



Signal S_3



Signal S_4

Figure 3.6: Oscilloscope traces of the sensor signals from S_0 to S_4 , showing the timing and channel switching logic.

- **DCC_{SEL}**: This control bit selects whether the 5-bit DCC counter is fed by an internal clock or an external clock to control the calibration behavior.
- **DCC_{Counter}**: This line provides the PWM signal that establishes the clock frequency for the DCC counter. Due to pin-multiplexing constraints, this pin is shared with the USART1_RX peripheral, that required firmware management to handle both functional behaviors.
- **OUT_D**: This digital output notifies the microcontroller when the DCC has detected an out-of-range voltage and has successfully brought the signal back within the defined thresholds.
- **RESET**: The system utilizes two reset signals: one able to clear the DCC counter and the other one to return the multiplexer to channel 0. This is crucial for synchronizing the starting point whenever the system toggles from manual to automatic channel iteration modes.




Figure 3.7: Oscilloscope capture of the short reset pulse (approximately 450 ns). This signal is generated by the firmware to trigger the transition from manual channel selection, managed by the STM32, to the automatic internal iteration mode of the sensor.

- **CKpad, 2CKpad**: These signals can be utilized to provide external clocking to the sensor; however, they are not routed to the main Motherboard.
- **Hfb1, Hfb2**: These digital control bits are used to configure the internal operating frequency of the sensor. Specifically, they allow to select the signal frequency among the predefined values of 50 kHz, 200 kHz, and 400 kHz, or to decide to use an external clock source, although the latter is currently not routed on the board.

3.1.3. ADC

The digitization of the conditioned analog signal is a crucial task of the dust detection system. To meet the requirements imposed by the physical characteristics of the sensor and the expected signal magnitude, the Texas Instruments ADS8867 [9] was selected as the designated Analog-to-Digital Converter (ADC) (Figure 3.8).




Figure 3.8: ADC ads8867 schematic functioning from datasheet.

The choice of this specific component is motivated by analysis of the system's dynamic range and bandwidth constraints, which will be detailed in this sections. The ADS8867 is a 16-bit, successive approximation register (SAR) ADC offering high-precision performance with true differential inputs.

The fundamental detection limit for this specific application is dictated by the minimum capacitive variation that the system must resolve. Based on simulation models, the target threshold is established at $\Delta C_{el} \approx 10$ aF, which corresponds to the signal perturbation induced by a dust particle with a diameter of $5\mu m$ and a relative permittivity of $\epsilon_r \approx 2$. The absolute lower limit of the measurement capability is defined by the system's electronic noise floor. The total equivalent input noise has been calculated to correspond to a root-mean-square capacitance difference of $\Delta C_{el,rms} \approx 236$ zF_{rms}.

On the opposite end of the dynamic range, the maximum detectable signal—defined as the largest particle the system can measure before non-linearity occurs—is constrained by the saturation limits of the voltage source, assuming the DC Current Cancellation (DCC) circuit is not active. By relating the buffer output voltage to the input capacitance difference, it is possible to derive the system's full-scale range. Given the differential sensing topology utilized in this design, the maximum allowable voltage swing on a single branch is limited to 1.65V relative to the common-mode voltage. Consequently, this voltage constraint translates to a maximum measurable capacitance of $\Delta C_{max} \approx 1.95$ fF.

With the full-scale range and the noise floor clearly defined, it is possible to analytically determine the minimum resolution required for the Analog-to-Digital Converter (ADC). The required number of bits is derived from the dynamic range ratio as follows:

$$N_{bits} = \log_2 \left(\frac{\Delta C_{max}}{\Delta C_{rms}} \right) = \log_2 \left(\frac{1.95 \text{ fF}}{236 \text{ zF}} \right) \approx 13.01 \quad (3.1)$$

The analytical calculation yields a theoretical requirement of 13.01 bits, implying that the absolute minimum integer resolution necessary to avoid information loss is $N_{bits} = 14$. However, to ensure a robust safety margin and to minimize quantization errors, a 16-bit architecture was selected for this implementation. It is important to note that the chosen ADC allocates 1 bit for the sign; consequently, the effective resolution available for the positive signal swing is 15 bits.

This 15-bit allocation defines the vertical resolution of the system, determining the smallest detectable change in voltage, also known as the Least Significant Bit (*LSB*). Given a Full Scale Range (*FSR*) of $3.3V$, the voltage resolution can be derived as follows:

$$LSB_V = \frac{FSR}{2^{15}} = \frac{3.3V}{32768} \approx 100.7\mu V \quad (3.2)$$

This value indicates that the system is capable of resolving differential voltage variations as small as approximately $100\mu V$. To contextualize this precision with a practical example, it is possible to consider the minimum detectable particle of interest, which corresponds to a capacitive variation of $\Delta C \approx 10aF$. Using the function below 3.3, obtained by simulation during sensor design:

$$V_{OUT_D} \approx V_{ref} \cdot \frac{\Delta C}{C_{eq}} = \frac{3.3 \text{ V}}{3.9 \text{ fF}} \cdot \Delta C \quad (3.3)$$

The $10aF$ capacitance induces a voltage signal of approximately $8.46mV$. By dividing this analog value by the digital step size, we obtain:

$$\text{Digital Output} = \frac{8.46mV}{100.7\mu V} \approx 84 \text{ LSB} \quad (3.4)$$

A signal magnitude of 84 LSBs confirms that even the smallest detectable particle generates a distinct digital signature significantly above the single-bit quantization floor, ensuring high-fidelity reconstruction.

Regarding the interface, the ADC performs the conversion on a fully differential signal, where the positive input ($AINP$) is connected to V_{OUT+} and the negative input ($AINN$) is tied to V_{OUT-} . This topology allows the device to effectively reject common-mode noise present on the sensing lines. The digital communication between the host microcontroller and the ADS8867 is established through a Serial Peripheral Interface (SPI) protocol—which will be analyzed in depth in the firmware chapter—specifically optimized to operate using a 3-wire configuration to minimize the digital I/O count.

3.1.4. Stability and Antialiasing Filter

A comprehensive technical evaluation of the topology of the buffer responsible for outputting the channel signal was performed to ensure full compatibility with the stringent requirements of the ADS8867. The ADS8867 datasheet explicitly suggests the utilization of input amplifiers—which in this design correspond to the output buffer of the sensor, as illustrated in Figure 3.9—that provide a sufficiently high small-signal bandwidth. This parameter is a critical factor for effectively reducing the closed-loop output impedance. Such a characteristic allows the driver to effectively manage the low-cutoff RC anti-aliasing filter and to properly settle the sampling transients, often referred to as kickback noise, that are generated by the ADC during the acquisition phase.




Figure 3.9: Output buffers of the multiplexer inside the DUST sensor.

Furthermore, the substantial available bandwidth ensures that total harmonic distortion remains minimized, even when operating at higher input frequencies. Finally, the stability of the entire input driver stage was rigorously verified by cross-referencing the amplifiers' bandwidth with the theoretical constraints established in Equation 3.5.

$$\text{Unity - Gain Bandwidth} \geq 4 \times \frac{1}{2\pi \times R_{FLT} \times C_{FLT}} \quad (3.5)$$

The output buffers integrated within the dust sensor chip are characterized by an output resistance value of $R_{FLT} = 50\Omega$ each (Figure 3.9). In order to determine the most appropriate capacitance value, extensive parametric simulations were conducted prior to selecting the component that yields the optimal performance fit.




Figure 3.10: GBWP and Phase Margin of C_{FLT} capacitance from $1pF$ to $10\mu F$.

To investigate the circuit stability in greater detail, the Phase Margin was plotted as a function of varying capacitance values (Figure 3.11). This analysis allows for a clear visualization of how the stability margins evolve across the capacitive range.




Figure 3.11: Phase Margin for C_{FLT} from $1pF$ to $10\mu F$.

Based on the simulation results, the best fit for this specific application was determined to be $C_{FLT} = 10nF$, which was subsequently selected for the final design. The Unity

Gain Bandwidth of the output buffer, computed via Cadence simulation, yielded a value of 10MHz; this confirms the correct selection of the components, as the resulting phase margin is maintained above 80°. Furthermore, with this selected parameters, Equation 3.5 dictates a minimum required value of 2.12MHz. Since the simulated bandwidth significantly exceeds this threshold, it ensures perfect operating conditions for the ADC.




Figure 3.12: GBWP and Phase Margin of $C_{FLT} = 10nF$ capacitance.

The anti-aliasing filter is designed as a low-pass RC network, for which the 3-dB bandwidth has been also optimized based on the specific requirements of this application. Although a single dust signal typically settles within 16ms and is already pre-filtered internally within the sensor, additional precautions are necessary. To mitigate any potential residual noise, the filter bandwidth should be kept conservatively low to band-limit the wideband noise fed into the ADC input, thereby significantly increasing the overall signal-to-noise ratio (SNR) of the system.

$$f_{-3dB} = \frac{1}{2\pi \times (R_{FLT} + R_{FLT}) \times C_{FLT}} \quad (3.6)$$

Consequently, the cut-off frequency is set to 100kHz to ensure absolute stability. This implementation discussed above, confirms that the system is a robust and oscillation-free signal conditioning chain.

3.1.5. Signal Handling, waveform and timing

Effective communication between the microcontroller and the ADC depends on strict understanding of the signal functions and adherence to the timing requirements specified in

the device datasheet. By examining the operational waveforms, it is possible to understand the exact sequence of the conversion cycle: from the moment the signal is sampled to the final transmission of the digital data.




Figure 3.13: ADC ads8867 waveform from datasheet.

The ADC is operated in *3-Wire $\overline{CS}$ Mode*.

- **DVDD:** Is the power supply. Decoupled with a $1\mu\text{F}$ capacitor as suggested in the datasheet.
- **DIN:** The level of this signal at the start of a conversion selects the mode of operation. In this case, its state is always HIGH, as it is connected directly with VDD. This configuration permits CONVST to be the $\overline{CS}$ signal, allowing to use the *3-Wire $\overline{CS}$ Mode*.
- **CONVST:** This signal functions as a $\overline{CS}$ pin.
- **DOUT:** It is the Serial data output, compatible with SPI communication.
- **SCLK:** It is the Clock input, data output are synchronized with this clock.

The process is initiated by a rising edge on CONVST, which triggers the internal sample-and-hold circuit and places the DOUT pin into a high-impedance state. The conversion is then performed using the internal clock. The minimum conversion time t_{conv_min} is 500ns and t_{conv_max} is 8800ns, for this reason, inside the firmware the CONVST signal returns low after at least $10\mu\text{s}$. Upon completion of the conversion, the device automatically enters the acquisition phase and switches to a low-power mode to save energy. To read the results, CONVST is pulled low. This falling edge brings DOUT out of its high-impedance state and immediately presents the Most Significant Bit (MSB). The subsequent bits are shifted out on each SCLK falling edge. While the data is valid on both edges of the clock, capturing it on the falling edge allows for a faster overall reading rate. The communication

concludes, and DOUT returns to tri-state, after the 16th SCLK cycle or when CONVST is driven high again. The ADC shares the SPI3 bus with the microSD card. After the communication concludes, DOUT is in tri-state, and it is possible to communicate with the microSD peripheral.

3.2. Motherboard

The motherboard serves as the central processing and power distribution hub for the entire system, implemented on a six-layer PCB to ensure optimal signal integrity and easier routing. At the core of the board is the STM32WBA5MMGH6 module, an ultra-low-power microcontroller operating at 100MHz, which provides the necessary computational overhead while featuring an integrated, certified antenna for low-energy Bluetooth communication. The physical layout is optimized for modularity, hosting the primary Board-to-Board (BTB) connector, which allows for a compact interface with the sensor-carrier board.




Figure 3.14: Top and Bottom side of the motherboard that host control electronics and power system.

A power management system, discussed in detail in 3.3, is integrated into the motherboard to handle multiple operational scenarios. The system utilizes a Power MUX based on the TPS2117 to intelligently switch between the USB-C input and an external power supply connector, the latter enabling long-duration experimental runs without relying on a host computer. Reliability is further enhanced by an ESD protection stage, utilizing HSP053-4M5 diode arrays on critical data lines, including the Micro SD card interface (visible in Figure 3.14b) used for high-capacity event storage and data logging.

Development and debugging are facilitated by a firmware deployment infrastructure. Users can upload firmware via the main USB port or a high-performance Tag-Connect interface for the ST-LINK programmer. The board includes tactile controls for the development phase, specifically a BOOT0 switch to enter the bootloader and a dedicated

RESET button with an RC debouncing circuit to restart the system. Visual feedback is provided by an RGB user programmable LED.

To ensure peak operational performance, the design prioritizes electromagnetic compatibility (EMC) and high power integrity, as these factors are fundamental for minimizing noise interference throughout the system. These challenges, as already stated, were addressed by implementing a six-layer PCB architecture for the motherboard. This multi-layer stack-up provides the necessary routing density to facilitate complex interconnects. By leveraging this configuration, the layout achieves strategic isolation between analog, digital, and power domains, utilizing dedicated planes to provide effective shielding. This architectural choice significantly mitigates crosstalk and prevents signal degradation.

3.2.1. Bluetooth PCB design

The microcontroller has been chosen to be placed in the top-right corner of the motherboard. This design choice takes care of the goal of keeping the board as tiny as possible, the mechanical constraint of the carrier mount on top, and the design rules given from STMicroelectronics [10] (Figure 3.15).




Figure 3.15: Design rules from ST, chosen corner identified as "Good".

The STM32WBA5MMG uses an internal ceramic antenna with a slot-type topology. This antenna is very small and relies on a ground plane in the PCB to radiate effectively. The slot antenna type is considered an LC resonator. The capacitance (C) is defined by the ceramic antenna component, and the inductance (L) is proportional to the length of the current path around the antenna clearance on the PCB. The resonance frequency is inversely proportional to LC. Increasing either C or L inadvertently decreases the frequency. The layout must be designed with particular care to avoid performance loss caused by a shift in LC resonance.

A black area to the left of the sensor, indicates a clearance constraint in which no ground planes and in general no copper has to be placed. A larger clearance area results in a longer current path and reduces the resonance frequency. If the clearance area is too small, the frequency increases. A reference layout (Figure 3.16) has been implemented to ensure the correct functioning of the resonator.




Figure 3.16: Design clearance based on STMicroelectronics official technical specifications.

The parameter with which the power had been analyzed, give an indication of the power transferred between the two terminals of the communication system. S_{YX} gives the power that is received by Y after transmission from X. If X transmits 0 dB and Y receives -10 dB then $S_{YX} = -10$ dB. The negative value indicates a loss. Systems can be compared by looking at the S value. The closer the (negative) value is to 0, the lower the loss.

$$S_{YX} = \frac{\text{Transmitted}}{\text{Incident}} \quad (3.7)$$




Figure 3.17: Basic correspondences between watt and dBm.

$$P = 1mW \times 10^{x/10} \quad (3.8)$$

In the standard operating scenario, the microcontroller is configured to transmit at a nominal power of -0.3dBm that, accordingly with equation 3.8, corresponds to a power consumption of approximately 0.93mW for the Bluetooth signal establishment. The propagation model assumes an initial signal loss starting from -35 dB. The system maintains a stable connection up to a maximum distance of 100m (measured with 0.5m precision with Google Earth tool, Figure 3.18), at which point the signal reaches a threshold of -95dBm. Beyond this limit, the signal-to-noise ratio (SNR) drops significantly, and the transmitted packets become indistinguishable from environmental background noise.



Figure 3.18: Distance measured using Google Earth tool.

The system's flexibility is guaranteed by the programmable output power features of the microcontroller. By modifying the `CFG_TX_POWER` register, typically set to 0x19 for standard operations, it is possible to modulate the TX output stages. This allows for a trade-off between power efficiency and transmission range, adapting the device to either low-power proximity applications or high-coverage scenarios.

- **Ultra Low-Power Configuration (`CFG_TX_POWER = 0x00`):** This setting is optimized for short-range efficiency, achieving a maximum reliable distance of 2.5m. The generated output power is reduced to -20.4dBm, equivalent to 9.1 μ W. In this mode, the signal experiences a steep initial attenuation of -60dB before reaching the receiver sensitivity limit of -95dBm.
- **High-Power Configuration (`CFG_TX_POWER = 0x23`):** Designed for long-range communication, this mode extends the total coverage to 160m. The

output power is boosted to +10dBm, which translates to a significant 10mW of power. The link budget in this case is much more robust, with an initial path loss of only -18dBm before reaching the final threshold of -95dBm at the maximum range.




Figure 3.19: Grounding EMC strategy for the System.

3.2.2. User LED and visual feedback

The user interface is equipped with an integrated RGB LED that provides immediate diagnostic feedback regarding the board's operational status. The three color channels of the LED are driven by the microcontroller through three independent Pulse Width Modulation (PWM) outputs.



Figure 3.20: Green LED: Standalone operational mode.



Figure 3.21: Blue LED: Active and stable BLE connection.

When the board is powered and functioning correctly in standalone mode, the LED is

illuminated in green, indicating that the internal logic is stable but a Bluetooth connection it is not established. This state is visually confirmed in Figure 3.20, which shows the device during its initial idle phase. Once the Bluetooth Low Energy (BLE) pairing procedure is successfully completed and the data stream begins transmitting to the host computer, the LED transitions to a blue light, as illustrated in Figure 3.21.

In contrast to these normal operating conditions, a red light is used to warn the user of critical system errors. This error state is triggered specifically by a hardware over-current condition, which may point to a potential sensor malfunction, or in cases where the SD card data logging has been enabled via the GUI but the memory module fails to respond or initialize correctly. By integrating these visual cues, the system ensures that the user can monitor the health of the acquisition process and the connectivity status at a glance, even when away from the workstation monitor.

3.2.3. BOOT0 Switch and System Reset

The board features a dedicated hardware interface to manage the STM32 microcontroller's boot modes and firmware update cycles. This interface consists of a two-position BOOT0 switch and a manual Reset button. To initiate a manual firmware upload, the BOOT0 switch must be moved from its default Ground (GND) position to the +3.3V rail. This configuration forces the microcontroller into its internal Bootloader mode upon the next hardware trigger. Once the switch is set, pressing the Reset button activates this mode, allowing the host computer to recognize the device and upload the new firmware via the USB interface. During the flashing process, the activity is visually confirmed by the rapid toggling of the dedicated TX and RX diagnostic LEDs.

Following the successful completion of the firmware update, the BOOT0 switch must be returned to the GND position to ensure the microcontroller boots from the internal Flash memory. A subsequent press of the Reset button is required to restart the system under the newly loaded firmware. The successful initialization of the user application is then signaled by the RGB status LED turning green, indicating that the board is ready for operational use.

3.3. Power Management System

The power management system is based on a cascaded architecture (Figure 3.22) designed to provide stable and low-noise supply rails for both digital and sensitive analog components through all the system. The strategy involves a primary high-efficiency conversion followed by secondary precision regulation stages.




Figure 3.22: Power Management System.

The system can be powered either with an external power supply (battery, generic power source) or via a USB-C connector. To select which power source to use between them, the power management relies on a power multiplexer.

3.3.1. External Power Supply

The external power supply is intended to be powered with +5V. To avoid any problems, this is clearly indicated on the board silkscreen as well. This input is protected with an ESDA7P60-1U1M Zener diode [11], preventing overvoltage from damaging the board by creating a direct path to ground. The TPS2521 [12] is an integrated electronic fuse solution designed to ensure robust power management and safe power delivery in low-voltage systems. Unlike traditional passive fuses, which rely on thermal melting and require manual replacement after a fault, the eFuse is a silicon-based protection device that actively monitors electrical parameters and provides high-speed response to faults. The operational range is characterized by an input voltage range of +2.7V to +5.7V. When the input voltage exceeds a user-defined threshold, pin (OVCSEL), selectable at either 3.8V (if shorted to ground) or 5.7V (if left open), the device reacts within approximately 5 μ s to clamp the output voltage to a safe level. In this case, the pin is left open, and any voltages above +5.7V will be brought back to the chosen threshold value. If the part stays in a clamping state for an extended period of time, there is higher power dissipation inside the part, which can eventually lead to thermal shut-down (TSD). After the part shuts down due to a TSD fault, it automatically attempts to restart after a cooling period

and a fixed recovery delay (t_{RST}) of 110 ms. This feature is critical for maintaining high system uptime in remote or automated environments where manual intervention is not feasible. The Absolute Maximum Rating value is +28V; this value represents the survival limit of the internal silicon structures. The maximum current given to the load can be controlled via an external resistor connected to a dedicated pin (ILM). The e-Fuse has been programmed to limit the maximum drawn current to 500mA with a $6.8k\Omega$ resistor.

3.3.2. USB-C power supply

The USB-C connector has been placed to use a +5V source of a generic computer connected to the board. This design incorporates a USBLC6-2P6 [13], from STMicroelectronics, in order to provide robust Electrostatic Discharge (ESD) protection for the $D + / D -$ differential pair. The "very-low capacitance" characteristic of the USBLC6-2P6, features a typical parasitic capacitance of only 3.5pF. This specification is vital for preserving the performance of the USB 2.0 High-Speed protocol used in this case. Operating at a data rate of 480 Mbps, the rise and fall times of the digital pulses are extremely short; employing protection diodes with higher capacitance would introduce a parasitic low-pass filtering effect, rounding off the signal edges.

3.3.3. Power multiplexer

Battery operation becomes possible through the external power supply option, thereby greatly improving the system's portability. Connecting a USB cable to a computer for interfacing purposes while external power remains active would, without adequate protection, create a parallel connection between the two sources, potentially causing irreversible system damage. A power multiplexer (TPS2117DRLR [14]) manufactured by Texas Instruments has been integrated to address this concern. Independent switching between primary and secondary supplies is guaranteed by this component. The external power supply serves as the primary source in this setup, with the USB's +5V functioning as the secondary option. The switch function between VIN1 and VIN2 is guaranteed by the PR1 and Mode pins: when MODE is tied high, PR1 determines the channel selected. To configure VIN1 as the priority supply, Mode is connected to VIN1 and a proper threshold through a resistor divider from VIN1 to PR1 is set. When PR1 is pulled above an internal voltage $V_{REF} = 1V$ generated inside the device (Figure 3.23), the voltage on VIN1 is used to power the output, and when it is pulled below V_{REF} , VIN2 is used to power the output.




Figure 3.23: Design configuration, V_{REF} and internal details of TPS2117.

3.3.4. Filtering stage

The section dedicated to noise filtering in the power supply network utilizes an RLC low-pass filter, strategically positioned downstream of the Power MUX. This stage receives the $+5V_{MUX}$ voltage at its input and has the fundamental task of cleaning the signal from high-frequency interference before it is distributed as the $+5V_{Filtered}$.

The circuit architecture of the filter involves the use of an L1 inductor with a value of $10\mu H$, which features a maximum parasitic series resistance of $900m\Omega$. A 2Ω resistor, R9, is placed in series with this component, contributing to the definition of the filter's damping. The shunt branch to ground consists of a C7 capacitor with a capacitance of $10\mu F$, characterized by an equivalent series resistance (ESR) of approximately $10m\Omega$.




Figure 3.24: RLC filter schematic.

By analyzing the frequency response provided by the component characterization (figure 3.25, it is observed that the filter has been sized to work with a cutoff frequency f_c located at approximately $15.9kHz$ using the formula 3.9:

$$f_c = \frac{1}{2\pi\sqrt{L \cdot C}} \quad (3.9)$$

As highlighted by the magnitude plot, the system ensures a progressive attenuation of electromagnetic noise, reaching a reduction of -10dB at high frequencies around 100kHz . This configuration ensures that the $+5V_{Filtered}$ supply distributed throughout the system is effectively decoupled from high-frequency power supply noise (typical values of PC power supply source noise are around 100kHz), thereby significantly improving the board's overall SNR.




Figure 3.25: RLC transfer function.

3.3.5. Power regulation stage: CP2102 and LT3045

The final stage of the power management system relies on two distinct regulation circuits that distribute power to different sections of the board, ensuring signal integrity and noise isolation. A dedicated LT3045 Low Dropout Regulator (LDO) has been selected to convert the $+5\text{V}$ rail to $+3.3\text{V}$. This component is specifically focused on providing a low-noise, clean power supply to the carrier board, ensuring that only the sensitive sensing circuitry is powered by this ultra-low noise source.

The CP2102N, previously introduced in Section 3.2, features an integrated voltage regulator capable of converting an input voltage (VREGIN) to a stable $+3.3\text{V}$ output (V_{DD}). To ensure proper decoupling from the noisy USB host supply, the component has been configured in the *Self-Powered* mode, as recommended by the datasheet specifications.

In this configuration, the device is powered by the local board supply rather than the

USB cable. As illustrated in Figure 3.26, the local $+5V_{\text{Filtered}}$ rail is connected directly to the V_{REGIN} pin. However, even in self-powered mode, the device must detect the presence of the USB host. This is achieved by connecting the USB $+5V_{\text{USB}}$ signal to the V_{BUS} pin. This function ensure that when USB is disconnected, the CP2102 does not look for a connection, lowering the current consumption from $9.5mA$ to $195\mu A$.




Figure 3.26: CP2102N Power Configuration (Self-Powered Mode).

According to the specific electrical requirements for self-powered operation, a voltage divider is required on the V_{BUS} pin. This divider, composed of a $22.1k\Omega$ and a $47.5k\Omega$ resistor, serves two purposes:

1. It ensures the voltage on the V_{BUS} pin does not exceed the absolute maximum rating defined as $V_{IO} + 2.5V$.
2. It allows the device to reliably detect when it is connected to a bus while preventing high leakage current when the device is not powered.

It should be noted that a discrepancy exists between the final schematic design and the currently manufactured prototype. Due to a legacy configuration carried over from a previous test board, the physical PCB was inadvertently routed in a *Bus-Powered* configuration (where the device draws power directly from the USB host). However, the implementation depicted in Figure 3.26 represents the corrected, final design intended for the production revision, ensuring robust self-powered operation.

For what concern the second power path, the selection of the LT3045 is critical for the

performance of the capacitive sensing stage. Since the detection of zepto-farad variations relies heavily on a stable voltage reference, any ripple or switching noise present on the supply line could be coupled into the measurement, covering the dust particle signals and degrading the dynamic range. The LT3045 features an architecture that minimizes spot noise to $0.8\mu V_{RMS}$, essentially isolating the carrier board from upstream perturbations originating from the USB +5V supply.

The schematic implementation, highlights the unique configuration of this device. Unlike traditional Low Dropout Regulators (LDOs) that rely on a resistor divider network connected to a voltage reference, the LT3045 employs a precision current source reference. The device sources a constant $100\mu A$ current out of the **SET** pin. Consequently, the output voltage is determined strictly by Ohm's law across the set resistor (R_{SET}), providing unity-gain operation and constant noise performance regardless of the output voltage level.

In this design, a resistor value of $R_{SET} = 33.2k\Omega$ (component R15) has been selected. The resulting output voltage is calculated as follows:

$$V_{OUT} = 100\mu A \cdot R_{SET} = 100\mu A \cdot 33.2k\Omega \approx 3.32V \quad (3.10)$$

To further enhance the spectral purity of the supply, a capacitor ($C_{SET} = 4.7\mu F$) is placed in parallel with the setting resistor. This component filters the thermal noise generated by the reference current source and improves the transient response. The stability of the regulator is ensured by a $10\mu F$ ceramic capacitor on the output rail ($+3.3V_{Sensor}$).

The LT3045 provides a programmable current limit feature to protect the downstream sensing circuitry from short-circuit conditions or excessive power dissipation. This limit is established by connecting a resistor between the **ILIM** pin and ground. The relationship between the programming resistor (R_{ILIM}) and the maximum output current (I_{LIMIT}) is governed by the following equation:

$$I_{LIMIT} \approx \frac{150mA \cdot k\Omega}{R_{ILIM}} \quad (3.11)$$

For this implementation, a resistor $R16 = 3k\Omega$ was selected. Substituting this value into the equation yields:

$$I_{LIMIT} = \frac{150mA \cdot k\Omega}{3k\Omega} = 50mA \quad (3.12)$$

This $50mA$ limit is conservatively chosen to supply the necessary current for the carrier board operations while ensuring rapid shutdown in the event of a fault, thus preserving the integrity of the sensor. In fact, the complete system, consisting of the motherboard and carrier board, excluding the sensor, consumes $15mA$ when powered via USB with an active Bluetooth connection. This consumption can be further reduced to $5mA$ by enabling the self-powered configuration and a correct routing of the CP2102N.

3.4. Connection within Carrier and Motherboard

The electrical and mechanical coupling between the carrier board and the motherboard is established through the pair of Receptacle Molex 5024262630 and Plug Molex 5024302630, a high-density SlimStack connectors belonging to the 0.40mm pitch series [15]. This specific component was selected for its very tiny dimensions and ultra-low profile, offering a very low stacking height, which is a critical requirement for maintaining the compact vertical envelope of the final device.

The connector's design features a dual-contact structure (Figure 3.27), with the plug on top (soldered on the bottom-side of the carrier) and the receptacle on bottom (soldered on top-side of the motherboard), that ensures high reliability and a distinct tactile "click" during mating, providing immediate feedback on the correctness of the connection.




Figure 3.27: Mechanical coupling of the Board-to-Board (BTB) interfaces.

This type of interconnect is particularly susceptible to mechanical stress during the mating and unmating phases. Any excessive tilt or "peeling" motion during disconnection can lead to permanent deformation or damage of the pins or damage to the plastic housing. Therefore, Molex suggests a careful positioning strategy, that is mandatory to avoid structural failure (Figure 3.28).




Figure 3.28: BTB correct usage on how to remove carriers from the motherboard.

The following Figure 3.29 and table 3.1 detail the pinout mapping for the connector, defining the electrical interface between the carrier board and the motherboard. The pinout refers to the motherboard, top-view. This configuration allows all the signals discussed in section 3.1.1 and 3.1.2 to become available to the motherboard.




Figure 3.29: BTB connector pinout.

Pin number	Signal	I/O
1	OUT_P	AO
2	OUT_N	AO
3	DCC_Counter	O
4	OUT_D	O
5	CK_pad	I
6	2CK_pad	I
7	DCC_SEL	I
8	RESET_DCC	I
9	S4	I
10	S3	I
11	S2	I
12	S1	I
13	SEL	I

Pin number	Signal	I/O
14	SPI3_SCK	I
15	SPI3_MISO	O
16	SPI3_CS_DT	I
17	GND	A
18	GND	A
19	RESET_ch_read	I
20	GND	A
21	Hfb1	I
22	Hfb2	I
23	+3.3V_Sensor	A
24	GND	A
25	GND	A
26	S0_ck	I

Table 3.1: Pinout of the 26-pin board-to-board (BTB) connector.

4 | Firmware and GUI

The firmware for the system is developed in C and deployed on the STM32WBA5MMGH6 microcontroller, that features an ARM Cortex-M33 core and integrated Bluetooth Low Energy (BLE) capabilities. The software architecture is tailored to handle precise real-time data acquisition, digital signal processing, and wireless transmission, ensuring zero data loss during the continuous monitoring of the 32 channels. A dedicated Graphical User Interface (GUI) provides the user with the necessary tools to interact with the hardware. Furthermore, the system is designed such that it can maintain its standard operational workflow and routine execution entirely uninterrupted even if it is not connected via USB or Bluetooth. It is possible to continue storing data into a microSD peripheral for long-term experiments, allowing a successive access to the information previously recorded, assuming that the storage peripheral is inserted.

Under the experimental conditions currently established, the system demonstrates the capability to detect an average of 20 particles per channel every second. Given the 32-channel architecture of the sensor, this individual performance translates into an aggregate detection capacity of 640 particles per second across the entire device. It is important to highlight, that the operational environments for which the sensor is specifically designed exhibit particle rates that are substantially lower than these thresholds, ensuring reliable measurements conditions.

4.1. Firmware Architecture and Implementation

The system's core performance is driven by a high-speed clock configuration operating at 96 MHz. This frequency provides the necessary computational headroom to manage multiple concurrent tasks without performance degradation. For external connectivity and development, the microcontroller is configured with dedicated UART communication ports and a Serial Wire Debug (SWD) interface via the `SWDIO` pin, ensuring both robust data exchange and efficient real-time debugging capabilities. To optimize the management of the interfaces and all other data-intensive peripherals, the system utilizes Direct Memory Access (DMA) controllers whenever feasible. The implementation of DMA is crucial

as it offloads data transfer tasks from the CPU, allowing the processor to focus on high-level logic and signal processing while data is moved autonomously between peripherals and memory.

A critical aspect of the hardware design involves the management of the SPI3 communication bus, which is shared between the sensor and the microSD card platform. To ensure seamless operation and prevent any potential bus contention or signal interference, the system utilizes a Chip Select (CS) mechanism. This logic enables the microcontroller to activate the required peripheral only when strictly necessary, effectively isolating the devices during their respective communication cycles.

Thanks to the high operating frequency of 96 MHz, the processor can execute its routines with extreme efficiency. Processing tasks—such as simultaneous data acquisition from the sensor’s ADC and the subsequent writing operations to the microSD card—are completed in very short intervals. This high execution speed ensures that the system maintains a significant amount of idle time, preventing bottlenecks even during intensive data-logging operations and ensuring that the two peripherals do not interfere with each other’s performance.

4.1.1. Data Acquisition and Bluetooth Transmission

The continuous data stream generated by the 32 sensor channels is processed and subsequently packed into a customized Bluetooth Low Energy (BLE) frame structure.

Whenever the user initiates a data acquisition sequence via Bluetooth or USB through the User Interface, the LPTIM1 (a low-power timer) begins its counting operation. The LPTIM1 peripheral is clocked at a frequency of 96 MHz; its counter is managed through a prescaler that divides the input frequency by two, utilizing an Auto-Reload Register (ARR) value of 11999. This specific hardware configuration allows the system to achieve the precise timing parameters illustrated in the following equations:

$$f_{LPTIM1} = \frac{f_{clk}}{\text{Prescaler} \cdot (ARR + 1)} = \frac{96 \text{ MHz}}{2 \cdot (11999 + 1)} = 4000 \text{ Hz} \quad (4.1)$$

$$T_{LPTIM1} = \frac{1}{f_{LPTIM1}} = \frac{1}{4000 \text{ Hz}} = 250 \text{ } \mu\text{s} \quad (4.2)$$

Upon reaching its terminal count, the LPTIM1 generates a hardware interrupt and enters its dedicated callback routine. Within this function, the system executes the channel selection logic and initiates the Analog-to-Digital Conversion (ADC) acquisition cycle.

The addressing of the 32 sensor channels is implemented through direct register-level manipulation of the microcontroller's GPIO peripherals. To select a specific channel, the system modifies five control bits (S0, S1, S2, S3, and S4) via the Digital Output interface. This is achieved by performing a direct write operation to the `GPIOA_BSRR` (Bit Set/Reset Register), a method that ensures atomic bit manipulation for optimal timing and reliability.

The selection logic is executed through the following instruction:

```
GPIOA->BSRR = bsrrA[channel & 0x1F];
```

In this instruction, a bitwise mask of `0x1F` (corresponding to the 5 least significant bits) is applied to the `channel` variable. This masking operation guarantees that only the required 5 bits are addressed within the `bsrrA` lookup table, ensuring the correct signal routing for the dust sensor as defined in the hardware register specifications of the reference manual [16]. The `bsrrA` look up table is built in order to start from the electrode number 1 and continue in order until electrode number 32.

The ADC acquisition is triggered by asserting the `GPIOB_PIN13`, which is connected to the ADC `CONVST` (Conversion Start) signal, as previously discussed in Section 3.1.3. Following the initiation of the conversion, a worst-case delay of $t_{\text{conv-max}} = 8800$ ns (highlighted in the ADC datasheet) must be observed to allow for the completion of the sampling process. To manage this interval, an auxiliary timer, `TIM1`, is configured to count for a duration of $10\ \mu\text{s}$; after this period, the timer deasserts the `CONVST` signal and initiates the SPI3 communication sequence.

Finally, when the 12MHz SPI3 peripheral returns an interrupt, the system enters its callback and invokes the `DUST_Process` function, that is responsible for detecting deposition events. If the entire sensor has been read, meaning that the 32-data vector has been completely filled following a full "scan" of the sensor, the system informs the Bluetooth stack via the `UTIL_SEQ_SetTask` utility that a packet is ready for transmission.

While a complete sensor scan occurs approximately every 8 ms, a specific transmission strategy is implemented to avoid overstressing the microcontroller and to comply with the maximum Bluetooth packet rate of 7.5 ms. In fact, data packets are sent on a decimated basis (every other "scan" is transmitted, effectively skipping one frame), resulting in a final transmission interval of 16 ms for each packet containing the sensor condition and deposition data. However, it should be noted that all particle presence calculations are performed using the full set of data provided by the ADC without any skipping. This decimation does not affect the end-user's visual perception, as the physical sensor requires

a settling time of at least 17 ms following a particle deposition event.

The wireless communication is handled by a dedicated application layer (`ble_sensor_app.c`), which utilizes a custom GATT service. The payload is aggregated into a data structure and sent as a notification via the `BLE_SENSOR_MYDATA` characteristic. The frame length is composed by a frame of 100 bytes, ensuring that all 32 channel outputs and their respective cumulative particle counts are transmitted efficiently within a single connection interval.




Figure 4.1: Compact visual representation of the 100-byte BLE data frame structure.

$$2 \times 8\text{bit synch frame} + 32 \times (8\text{bit particle count} + 16\text{bit ADC}) + \backslash\text{n} + \backslash\text{r} = 100\text{bytes}$$

Byte Index	HEX	Description
[0]	0xAA	Frame Sync 1
[1]	0x55	Frame Sync 2
[2]	0xXX	Channel 1 - Cumulative Particle Count
[3]	0xXX	Channel 1 - ADC Value (MSB)
[4]	0xXX	Channel 1 - ADC Value (LSB)
...	...	...
$[2 + (k \times 3)]$	0xXX	Channel $k + 1$ - Cumulative Particle Count
$[3 + (k \times 3)]$	0xXX	Channel $k + 1$ - ADC Value (MSB)
$[4 + (k \times 3)]$	0xXX	Channel $k + 1$ - ADC Value (LSB)
...	...	...
[95]	0xXX	Channel 32 - Cumulative Particle Count
[96]	0xXX	Channel 32 - ADC Value (MSB)
[97]	0xXX	Channel 32 - ADC Value (LSB)
[98]	0x0D	Carriage Return $\backslash\text{n}$
[99]	0x0A	Line Feed $\backslash\text{r}$

Table 4.1: Detailed explanation of the 100-byte BLE data packet transmitted by the firmware.

4.1.2. Local Data Logging and microSD Management

To allow for long-term experiments, the system features a local data logging unit. The non-volatile storage is managed using a microSD card, physically interfaced with the microcontroller through the SPI3 peripheral. At the software level, the file system operations are handled by the robust FatFs generic FAT/exFAT file system module [17].

Writing data to a flash-based memory card is an inherently non-deterministic process. SD cards occasionally perform internal wear-leveling or block-erase operations, which can introduce write latencies of several milliseconds. If a write operation were triggered synchronously during each 8 *ms* reading cycle, it would inevitably stall the microcontroller, causing the analog front-end to drop critical samples and corrupting the particle detection algorithm.

To decouple the strict real-time acquisition from the slow and unpredictable storage operations, a RAM-based buffering strategy is implemented. The firmware allocates a large circular buffer of 32 *KB* (RAM_BUFFER_SIZE). The continuous data stream from the sensor is pushed into this buffer at low latency. A background process continuously monitors the buffer occupancy; when the unwritten data exceeds a 16 *KB* threshold (WRITE_THRESHOLD), a bulk write operation is triggered. This approach maximizes the SPI3 throughput by writing large contiguous chunks of data and completely shields the high-priority signal processing interrupts from any SD card latency.

To further optimize storage space and write times, the data frame saved to the microSD is strictly limited to the essential analog measurements. Unlike the BLE packet, which includes the cumulative particle count for real-time visualization, the binary log file stores only the 16-bit ADC values for the 32 channels. This results in a highly compact, 68-byte binary frame. The structure of the logged data frame is depicted in Figure 4.2.




Figure 4.2: Compact visual representation of the 68-byte binary data frame structure logged to the microSD card.

From a temporal perspective, a simple calculation demonstrates the efficiency of this logging strategy. The system requires approximately 32 hours and 40 minutes to accumulate

1 GB of data, which translates to a daily storage footprint of just 734.4 MB. Consequently, by equipping the board just with a standard commercial microSD card, it is technically feasible to sustain continuous, uninterrupted data acquisition for a single experiment over several months.

4.1.3. Particle Counting Algorithm

The core logic of the firmware is governed by a Finite State Machine (FSM) designed to ensure reliable particle counting inside the call of the `DUST_Process` function. Since a deposited particle permanently alters the local capacitance, the algorithm must detect a persistent "step" in the signal rather than a transient pulse.

To achieve this, the raw signal first undergoes a Moving Average (MA) filter, with an averaging buffer size decided by the user, in order to suppress high-frequency noise. During the initial system power-up, the channel remains in a Warm-up phase for a fixed number of samples, allowing the software baseline to track and align with the initial hardware settling state. Once the warm-up is complete, the FSM transitions through the following states:

- **DUST_STATE_MONITORING:** During the monitoring phase, the algorithm continuously computes the absolute difference between the filtered incoming signal and the established baseline, allowing it to detect both positive and negative signal variations. As long as the signal remains stable and within the noise margin, the software baseline is continuously and slowly updated to track the signal's resting value. This adaptive tracking mechanism effectively compensates for slow environmental changes, such as thermal drifts. Conversely, if a sudden, sharp variation strictly exceeds the pre-calibrated, material-specific threshold, the FSM flags a potential physical impact. At this point, the system immediately transitions to the subsequent validation state to verify the event and filter out transient noise.
- **DUST_STATE_CONFIRMING:** To effectively reject fast electrical spikes, the algorithm requires the step variation to persist for five consecutive samples. Given the 8 ms sampling period between consecutive readings of the same channel—which accounts for the scanning cycle across all the other 31 channels—a particle impact requires exactly 40 ms to be formally validated. If the signal drops back below the detection threshold during this validation window, the event is immediately discarded as high-frequency noise. If, however, the step maintains its new level for the entire duration of the confirmation window, the system officially recognizes a valid particle deposition. As a result, the cumulative particle counter is successfully

incremented, and the finite state machine advances to the latching phase in order to properly settle the baseline.

- **DUST_STATE_FOUND**: The physical deposition of a particle creates a rising or falling edge that takes a few samples to fully settle. In this state, a "recover time" (1 sample) is enforced. The baseline is forced to rapidly track the filtered signal, completely absorbing the step. Once the the deposition event has been saved, the baseline is perfectly aligned with the new voltage level, and the FSM returns to the **MONITORING** state, ready to detect the next particle.

To detect subsequent particles, the algorithm calculates the absolute difference (Δ) between the newly filtered signal and the tracked baseline. If this Δ exceeds the calibrated material threshold, a particle is counted, and the baseline is immediately updated to the new signal level, effectively absorbing the step and resetting the reference for future events.

This structure guarantees that the digital signal processing is highly robust against both high-frequency noise (mitigated by the moving average) and low-frequency baseline drifts or settling times (handled by the warm-up state and the dynamic baseline tracking). This algorithm operates independently for each of the 32 channels. The use of a cumulative counter array (`particle_counters`) ensures that the host application (PC GUI) does not miss any particle events even if a BLE packet is dropped during transmission.




Figure 4.3: Waveform of the FSM for a particular channel.

The complete logic integrating the FSM and the particle detection is formalized in Algorithm 4.1. This process is repeated for each channel, with all state variables and counters stored in arrays using the monitored channel number as the index.

Algorithm 4.1 DUST_Process Channel FSM Algorithm

```

1: filtered = Update_Moving_Average(raw_sample)
2:
3: if warmup_cnt < WARMUP_LIMIT then
4:   increment warmup_cnt
5:   baseline = filtered {Track initial transient}
6: end if
7:
8: State ← MONITORING
9:  $\Delta = |\textit{filtered} - \textit{baseline}|$ 
10:
11: if State == MONITORING then
12:   if  $\Delta > \textit{Threshold}$  then
13:     over_cnt ← over_cnt + 1
14:     if over_cnt ≥ SPIKE_FILTER_LIMIT then
15:       State ← CONFIRMING
16:       over_cnt ← 0
17:     end if
18:   else
19:     over_cnt ← 0
20:     baseline ← Update_Slow_Drift(baseline, filtered) {Noise compensation}
21:   end if
22: else if State == CONFIRMING then
23:   if  $\Delta > \textit{Threshold}$  then
24:     over_cnt ← over_cnt + 1
25:     if over_cnt ≥ 4 then
26:       particle_count ← particle_count + 1 {Particle Confirmed (40 ms)}
27:       State ← FOUND
28:     end if
29:   else
30:     State ← MONITORING {False alarm, return to monitoring}
31:     over_cnt ← 0
32:   end if
33: else if State == FOUND then
34:   baseline ← filtered {Force baseline to absorb the step}
35:   event_buf ++ {Update particle deposition event counter}
36:   State ← MONITORING {Signal settled, ready for next deposition}
37: end if

```

4.1.4. CP2102N USB-to-UART Bridge Configuration

The communication interface between the hardware and the computer is managed by the CP2102N USB-to-UART bridge. By utilizing the *Simplicity Studio* software, the CP2102N can be easily configured by updating its internal non-volatile memory; this allows for the modification of crucial hardware parameters, such as the power mode. To optimize the overall energy efficiency of the system, the device is explicitly configured to operate in self-powered mode. This power configuration ensures that the bridge relies on the local board supply rather than continuously drawing current from the USB bus. As a result, when the USB cable is disconnected or the host connection enters a suspended state, the CP2102N significantly reduces its power consumption, according to the component’s datasheet, the supply current drops from $9.5mA$ to approximately $195\mu A$ during USB Suspend mode, thereby preventing unnecessary energy waste during inactive periods.




Figure 4.4: CP2102N Configuration window, with a detail in *Self Powered* Mode.

In addition to the power management optimizations, the device has been configured to provide immediate visual feedback during serial data exchanges. The General Purpose Input/Output (GPIO) pins of the CP2102N are programmed to operate in their alternate functions as Transmit Toggle (TXT) and Receive Toggle (RXT). These specific pins are physically routed to dedicated indicator LEDs on the PCB. By exploiting this hardware-level feature, the LEDs automatically illuminate and blink whenever a data transfer occurs between the microcontroller and the computer, such as during firmware flashing procedures or real-time data logging. This design choice implements a reliable diagnostic interface that allows the user to instantly monitor the communication status

without requiring any additional software overhead from the main microcontroller.




Figure 4.5: Port Configuration Window, here it is possible to set how the GPIOs behave: GPIO2 and GPIO3 here are set as TX and RX Toggle for motherboard LEDs.

Although most modern operating systems natively include drivers for recognizing standard USB peripherals, the operation of the CP2102N bridge requires the specific Virtual COM Port (VCP) drivers provided by the manufacturer. If these are not already pre-installed or integrated into the host operating system, it is necessary to manually download and install the dedicated software package. This step is very easy but strictly required, as the VCP drivers provide the fundamental abstraction layer needed to translate USB protocol packets into conventional serial communication. Only following the correct installation of these drivers can the operating system map the hardware device as a virtual COM port, thereby allowing application software on the host computer to interface seamlessly and bidirectionally with the custom board.

4.2. Graphical User Interface

To effectively monitor, record, and analyze the data stream generated by the 32-channel capacitive sensor, a custom Graphical User Interface (GUI) was developed. The software is written in Python and leverages the PyQt5 framework for the application architecture and PyQtGraph for high-performance, real-time data rendering. The application is divided into specific tabs, each handling a dedicated aspect of the system: Connection, Global real-time visualization, Channels real-time visualization, Dashboard, and offline file analysis.

4.2.1. Connection and Data Logging

The system communicates with the host computer primarily via Bluetooth Low Energy (BLE). The connection logic is handled by an asynchronous event loop utilizing the `bleak` library, ensuring non-blocking operations during continuous data reception. Through the Connection tab, the user can scan for the specific BLE UUID associated with the hardware, establish a secure link, and initiate the data stream. Additionally, this section allows the user to enable raw data logging, which records the incoming 100-byte BLE packets directly to serial monitor.




Figure 4.6: The Connection tab of the GUI, displaying the BLE device discovery and the serial port selection interface.

4.2.2. Real-Time Visualization and Hardware Status

The core feature of the software is the real-time visual representation of the sensor's status. The Visual tab implements a dynamic histogram (bar chart) updated every time a new 100-byte BLE frame is received and decoded. This interface instantly displays the cumulative particle counts in LSB or the reconstructed diameter size of the particle for all 32 channels simultaneously.




Figure 4.7: Live histogram visualization of the sensor deposited particle sizes, allowing instant evaluation of the particle counts and signal distribution.

The real-time visualization extends beyond hardware diagnostics to actively characterize the deposited dust. During experimental validation with a lunar dust simulant, the electrical step amplitudes, quantified in Least Significant Bits (LSB), were correlated with the simulant's physical volumetric distribution. This statistical fitting procedure established a specific power law relating the measured capacitance variations to the actual particle diameters. The Graphical User Interface directly implements this empirical power law to instantly convert the raw 16-bit ADC data into equivalent particle dimensions in micrometers. To effectively monitor the ongoing experiment, the central graph dynamically plots the particles' falling trend over time. Alongside this visual tracking, the dashboard features a cumulative particle counter and a real-time concentration metric expressed in Parts Per Million (PPM), offering immediate physical feedback on the dust accumulation.



Figure 4.8: Visualization of the initial floating state of the electrodes before the analog front-end biasing is fully stabilized.

For a more granular control over the experiment, the Advanced tab provides a comprehensive numerical dashboard. It displays the exact ADC value, the calculated voltage, and the particle count for each specific channel in a grid format. Once the warm-up transient is complete (barely 1 second after acquisition starts), the data are correctly printed in the monitors. Given the 3.3 V reference of the analog section, the stable baseline settles exactly at 0V and 0 LSB, with $\pm 1LSB$ of noise.




Figure 4.9: Stable analog baseline detection. The channels change their voltage when stimulated.

4.2.3. Advanced Dashboard and Material Configuration

Furthermore, the Settings dashboard provides comprehensive control over both the hardware behavior and the software interface. Beyond dynamically adjusting the detection thresholds, users can select specific physical targets (e.g., Lunar Dust, Martian Dust) whose parameters are loaded directly from a structured JSON configuration file (`materials.json`).




Figure 4.10: The Dashboard tab, allowing threshold adjustments based on the target material.

This ensures high flexibility depending on the simulant used during the experimental campaigns. At the hardware configuration level, the interface allows the user to modify the sensor's internal clock and define the channel scanning strategy. Specifically, the "rolling" of the channels can be set to Automatic, where the sensor's internal logic autonomously sequences through the 32 inputs, or Manual, where the microcontroller firmware explicitly drives the channel addressing during the acquisition cycle. On the software side, the

dashboard includes tuning options for the data monitor's print frequency, regulating how often the incoming data is visually refreshed on the screen to optimize both readability and CPU load. Finally, the user can toggle the overall GUI appearance between light and dark themes, ensuring optimal visual comfort under various laboratory lighting conditions.

4.2.4. Offline SD Card Analysis

To analyze the high-frequency data logged locally on the hardware or saved as text in the computer, the GUI includes an Analysis tab dedicated to the microSD binary files and generic text file. This module bypasses the BLE connection and parses the compact 68-byte SD card frames directly from the uploaded `.bin` file or a local `.txt` file. The software reconstructs the historical timeline of the experiment, plotting the ADC signal variations over time. This offline tool is essential for post-processing the data and high-speed transients associated with individual particle impacts that can be studied in detail.




Figure 4.11: The Analysis tab, utilized for parsing and plotting the 68-byte binary frames recorded offline by the hardware's SD card data logger.

5 | Experiments and Results

The validation and characterization of dust sensors require controlled experimental conditions that can accurately reproduce the target operational environment. To achieve this, a measurement campaign was carried out at ESA’s ESTEC facility on the previous generation sensor, developed in the previous years in the I3N Laboratory [18], where specialized equipment enabled a complete characterization that had not been previously performed.

The experimental setup employed a Dust Generation Unit (DGU) [19], a specialized machine designed to release controlled quantities of dust simulant through precise pressure regulation. A small amount of dust was loaded into the DGU, which then released the particles in a controlled manner, enabling systematic characterization of the sensor’s response. Initial tuning of the system was performed using Teflon powder, followed by the actual measurement campaign conducted with lunar dust simulant. This characterization campaign provided comprehensive experimental data and valuable insights that guided the development of the new firmware for the updated electronics and sensor architecture presented in this thesis.

5.1. Old sensor layout

The previous generation sensor was designed to detect particles with a nominal size of $1\mu m$. As shown in Figure 5.1, the monolithic CMOS chip measures $2.5mm \times 2.4mm$ and integrates all the necessary analog front-end circuitry alongside the sensing electrodes. The chip architecture can be divided into distinct functional regions, organized in horizontal bands from top to bottom.

The Digital-to-Capacitance Converter (DCC) section is located at the top of the chip, this 32-channel digital section, similarly to the new generation sensor, implements the mismatch compensation system, allowing fine-tuning of each channel’s baseline through programmable capacitance addition.

The central region contains the 32 sensing channels, each composed of 13 coplanar electrodes (7 stimulated and 6 reading) arranged vertically. The electrodes feature precise

dimensions: width $W = 1\mu m$, gap $G = 1\mu m$ between adjacent fingers, and length $L = 500\mu m$. Each channel occupies an area of $70\mu m \times 500\mu m$, resulting in a total active sensing area of approximately $1.12mm^2$. The vertical arrangement of the 32 channels maximizes the probability of particle interception on the chip surface.




Figure 5.1: Picture of the old CMOS sensor with the main system blocks highlighted.

Below the electrode region, 32 parallel charge preamplifiers with their associated OP2 amplification stages ($40\times$ gain) convert the capacitance variations into measurable voltage signals. Finally, the Lock-in and Gm-C Low-pass filters in the bottom of the sensor, implement the demodulation circuits and the low-pass filters with adjustable cut-off frequency, enabling the extraction of the DC signal proportional to the deposited particle size.

Each of the 32 channels is configured in a differential structure, where the channel is divided into two sensing capacitors (C_{UP} and C_{DW}), each with a nominal capacitance of 180 fF. The electrodes in each half are energized with counter-phase square-wave signals at $f_{clk} = 200kHz$ to improve common-mode noise rejection. Unlike the updated sensor

design where the two sections of each channel are separated by half of the subsequent channel, the previous generation maintains a simpler adjacent electrode arrangement within each differential pair, with both C_{UP} and C_{DW} electrodes grouped together in each $70\mu m$ -wide channel strip.

5.2. Old sensor measurements

The experimental validation of the 32-channel sensor was performed by depositing micro-particles of lunar simulant on the active surface using the Dust Generator Unit. Figure 5.2 shows the optical microscope images of the sensor area (the electrodes in light gray) before and after a deposition. The used lunar simulant comprises a particle size distribution ranging from $2\mu m$ to $100\mu m$. The distribution follows a lognormal relationship, characterized by a high concentration of fine sub-micron particles and a significantly lower frequency of larger particles.




Figure 5.2: Images of the first 16 channels before and after deposition.

Figure 5.2b illustrates the heterogeneous size distribution after the deposition across the sensor's active area. The particle population shows exactly the expected distribution pattern, with a higher concentration of small particles in the sub-micrometer to few-micrometer range, progressively decreasing in frequency as particle size increases, culminating in the occasional deposition of significantly large particles or agglomerates.

The optical microscopy images were acquired using a Leica microscope at $40\times$ magnification, providing sufficient resolution to resolve individual particles down to approximately $1\mu m$ in diameter. The comparison between the pre-deposition baseline image (Figure 5.2a) and the post-deposition image (Figure 5.2b) enables direct visual confirmation of

newly deposited particles and validation of the electrical signal data recorded by the sensor.

To enhance the detection and quantification of particle deposition events, particularly for smaller particles that are challenging to identify through manual visual inspection alone, an automated image processing algorithm has been developed and applied to the microscopy data. This algorithm, detailed in the following section, performs subtraction between the pre-deposition baseline image (Figure 5.2a) and post-deposition images (Figure 5.2b), effectively highlighting only the newly deposited particles.

5.2.1. Image Registration and post-processing

In order to study the quality of the sensor acquired data, it is first important to see what is the difference between the sensor before being exposed and after. To accomplish this task, an imaging process script has been written in Python.

The procedure is designed to isolate morphological changes between two states: the initial condition (before exposure) and the final condition (after exposure).

Since the sensor is physically removed from the microscope between the pre-deposition and post-deposition imaging sessions to perform the particle exposure in the DGU, the resulting images are not spatially aligned. Therefore, an image registration step is mandatory to correct for translational and rotational offsets before performing differential analysis. The "after exposure" image is aligned to the "before exposure" reference using an Enhanced Correlation Coefficient (ECC) algorithm. This transformation ensures sub-pixel correspondence between the two frames. A global brightness normalization is then applied to equalize the mean intensity levels, preventing false positives caused by fluctuations in ambient lighting or sensor exposure.

Once aligned, the two images are compared using an absolute difference operator. The resulting difference image D is calculated pixel-wise as:

$$D(x, y) = |I_{after}(x, y) - I_{before}(x, y)| \quad (5.1)$$

where I_{before} and I_{after} represent the intensity values at coordinates (x, y) . In this representation, static background features are suppressed, while new particles appear as high-intensity regions.

To distinguish actual particles from electronic sensor noise, a dynamic statistical threshold T is implemented. Rather than using a fixed value, the threshold is defined as a function

of the noise distribution within the difference image:

$$T = \mu + k \cdot \sigma \quad (5.2)$$

In this formula, μ represents the mean intensity of the difference image, σ is its standard deviation, and k is a sensitivity multiplier. Any pixel with an intensity exceeding T is classified as part of a deposited particle, highlighting it in red.

Finally, the result is what can be seen in the image below (figure 5.3):




Figure 5.3: The filter applied in Python highlights in red the new particles.

The newly deposited particles are then highlighted with yellow circles across the 32 sensing channels (figure 5.4). Each particle is manually annotated with its measured signal amplitude in LSB units from the acquired sensor data, showing the wide range of particle sizes detected during the deposition event. The majority of detected particles fall within the 50-300 LSB range, corresponding to particles in the few-micrometer diameter range, with occasional larger particles (>1000 LSB). Each sensor output value (in LSB) was manually correlated with a specific particle visible in the microscopy image through careful spatial analysis. This assignment is based on the reasonable assumption that the electrical signal magnitude directly correlates with particle dimensions, as supported by the theoretical sensor response model.




Figure 5.4: Highlighting of newly detected particles with respective LSB dimension annotations.

5.2.2. Comparison with acquired sensor data

For each detected particle, two values were recorded: the digital output from the sensor interface (expressed in LSB) and the physical diameter measured optically (expressed in μm). As previously described, optical measurements were performed post-deposition using a Leica optical microscope at $40\times$ magnification. A calibrated pixel-to-micrometer conversion was applied during post-processing to ensure accurate real dimensional measurements. The raw sensor data were instead analyzed using a custom Python script, which enables the visualization and processing of the acquisition data stream across all 32 channels, allowing the extraction of the relevant deposition events and their corresponding signal amplitudes.

The sensor output signal is proportional to both the contact area of the particle over the electrode surface and the height to which the particle perturbs the electric field generated between the coplanar electrodes. This electric field, however, decays rapidly with height and is sensitive only within approximately $1\ \mu\text{m}$ above the electrode surface. Consequently, for particles whose height exceeds this threshold, only the volume contained within this sensitive region contributes to the measured capacitance variation $\Delta C'$. The material extending beyond $1\ \mu\text{m}$ above the surface does not interact with the fringing electric field and therefore does not contribute to the sensor output. This effect introduces a clear non-linearity in the signal-to-size relationship for larger particles.

Taking for instance Channel 2, it is possible to see what was acquired through the sensor by seeing the graphs below (figure 5.5).




Figure 5.5: Acquisition on Channel 2 revealed 5 different particles.

Starting with the detection shown in Figure 5.5a, it is possible to observe a signal intensity of 160 LSB, which has been associated with a particle diameter of $12.1\mu m$. Subsequently, a significantly stronger signal of 2650 LSB was recorded (Figure 5.5b); this high value is consistent with the presence of a relatively large particle, clearly visible above the electrode, measuring $17.6\mu m$ in diameter. Following these events, a series of smaller particle acquisitions occurred. Figure 5.5c depicts two distinct detections of 50 LSB and 30 LSB, corresponding to diameters of $4.9\mu m$ and $1.9\mu m$, respectively. The final detection in this sequence (Figure 5.5d) shows a 40 LSB signal associated with a particle size of $1.5\mu m$. It is particularly interesting to note the behavior of these last two detections (30 LSB and 40 LSB): both particles landed on the very edge of the sensing electrode. Consequently, their full physical volume did not entirely contribute to the capacitive change and the resulting LSB count. For this reason, the optical measurement in micrometers was strictly limited to the portion of the particle body actually positioned inside the active electrode area, ensuring a consistent comparison between the partial overlap and the sensor output. It is important to note that in the data processing pipeline, an absolute value function ($|x|$)

is applied to the raw signal to quantify the magnitude of the impact event, which correlates to particle size. With this rectification, the physical distinction between the electrodes is no more observable in the time-domain signal as opposing "jumps". Specifically, a jump appearing as a "drop" is not related with a deposition on the negative electrode.

The statistical analysis of the collected data reveals the nature of the particle population. Figure 5.6 displays the histogram of the particle sizes measured optically all over the electrodes. The distribution follows a Lognormal distribution, which is exactly what was expected, as the distribution declared in the dust simulant is characterized by a high concentration of fine sub-micron particles with a peak at around $5\mu m$ and a significantly lower frequency of larger particles. The calculated mean diameter is approximately $5.69\mu m$, with a high concentration of small particles and a long tail towards larger dimensions as expected.




Figure 5.6: Particles measured with optical inspection ($6px = 1\mu m$).

Similarly, the distribution of the sensor digital output (LSB), shown in Figure 5.7, exhibits a clear Lognormal behavior, which is in excellent agreement with the Lognormal trend observed in the optical measurements. This consistency confirms that the sensor response faithfully reproduces the underlying particle size distribution. The average sensor response is recorded at 334.8LSB.




Figure 5.7: Particles revealed by the sensor in LSB.

The similarity between the two distributions suggests a consistent correlation between the physical size of the particle and the electrical signal generated by the sensor.

The optical analysis identified a total of 131 physical particles. The sensor system successfully resolved 122 *matches*, while recording 9 *misses* and 12 *false positives*. Applying the Generalized Jaccard metric [20] (equation 5.3), this performance results in an overall detection accuracy of 85.3%.

$$\text{Accuracy} = \frac{\sum_{i=1}^N \min(S_i, R_i)}{\sum_{i=1}^N \max(S_i, R_i)} = \frac{\text{Matches}}{\text{Matches} + \text{Misses} + \text{False Pos}} = \frac{122}{122 + 9 + 12} = 0.853 \quad (5.3)$$

The adoption of the Generalized Jaccard Index (Min-Max Accuracy) over a simple count ratio ($\frac{N_{\text{sensor}}}{N_{\text{real}}}$) is crucial for a conservative performance assessment. A simple ratio would allow opposing errors to cancel each other out; for example, a sensor missing 50 true particles while generating 50 false positives would misleadingly yield a 100% count accuracy despite failing to represent the reality. Alternatively, the Jaccard metric prevents

this compensation by defining the accuracy as the ratio between the *intersection* (the "secure" matches common to both methods, 122 in this case) and the *union* (the total envelope of events including all errors, 134 for this measurement). This approach ensures symmetric penalization.

The observed discrepancies were physically justified by experimental dynamics: *misses* were primarily caused by residual airborne particles observed under microscope settled on the sensor during its retrieval from the Dust Generation Unit (post-acquisition) or, to a lesser extent, by coincidence events masking simultaneous impacts on the same electrode. On the other hand, *false positives* were attributed to particle displacement, where correctly detected particles shifted away from the electrode area during handling prior to the optical inspection, or with minor probability, to sporadic electronic noise.

5.2.3. Fitting method to predict particle size

To establish the sensitivity law of the device for lunar simulant, the optical measurements were plotted against the sensor digital values. The resulting scatter plot is presented in Figure 5.8. A spread in the data points is observed, where a single particle size can correspond to different capacitance variations. This can be attributed to the irregular shape of the particles and their landing geometry: the sensor responds only to the contact area between the particle and the electrode surface, and for only $1\mu m$ in height. For instance, if a particle deposits with an irregular or "inverted teardrop" shape, the optical measurement records the equivalent spherical diameter as seen from above, which may be significantly larger than the actual contact area with the electrode. Consequently, particles with similar optical diameters can produce different electrical signals depending on their orientation and contact geometry upon deposition.

In order to mathematically characterize the behavior of the sensor, three distinct regression models were evaluated and compared: a Linear model, a Logarithmic model, and a Power Law model.

As illustrated in the comparative graph presented in Figure 5.8, visual inspection reveals significant limitations in the first two approaches. Specifically, both the simple Linear fit (represented by the dashed red line) and the Logarithmic fit (depicted as the solid blue line) fail to accurately capture the trend of the measurements. This deviation becomes particularly critical in the lower sensitivity range, specifically between 0 and 500 LSB. This region is of great importance, as it contains the vast majority of the particle events recorded during the experimental session. Consequently, the Power Law model (highlighted by the solid green line) emerges as the superior choice, providing the most robust

and accurate approximation of the experimental dataset.




Figure 5.8: Fitting functions that describe sensor behavior.

The observed non-linear relationship clearly indicates that the sensitivity of the sensor is not constant; rather, it progressively decreases as the particle size increases. This behavior suggests the presence of a saturation effect that becomes particularly evident for larger particles. It is possible to see a detailed inspection with Figure 5.9 that, particularly in the region characterized by a high density of data points, reveals the limitations of standard regression approaches. While both linear and logarithmic fitting models fail to capture the trend, the power law model demonstrates a superior ability to fit the experimental data. Consequently, the calibration equation is defined as:

$$d = a \cdot \text{LSB}^b \quad (5.4)$$

where LSB represents the raw digital output of the sensor and d denotes the estimated physical diameter of the particle.




Figure 5.9: Fitting function detail in the $0\mu m$ to $10\mu m$ range.

The calibration equation value that has been extrapolated from this fitting curve is:

$$d = 0.264 \cdot LSB^{0.553} \quad (5.5)$$

Once the calibration parameters a and b were determined, the raw dataset was converted into a vector of estimated diameters, that can provide a real-time measure of falling particle directly in the software. To validate the quality of this reconstruction, it was necessary to compare the distribution of the estimated values against the distribution of the real optical measurements. Even if in this experiment we are managing the same number of samples, both datasets were normalized into Probability Density Functions (PDFs). In this representation, the area under the histogram is normalized to 1, allowing for a shape-based comparison that is independent of the total number of events recorded.

The fidelity of the sensor was then quantified using two specific statistical metrics. The first metric is the Overlap Index, which is calculated by integrating the intersection area between the real PDF and the estimated PDF. This index, calculated with equation 5.6,

ranges from 0 to 1 (or 0% to 100%), where a higher value indicates that the sensor is correctly reproducing the statistical profile of the dust.

$$\text{Overlap Index} = \int \min(P_{real}(x), P_{sensor}(x)) dx \quad (5.6)$$

To illustrate this with a simplified numerical example, let's consider a distribution divided into two bins. If the real distribution exhibits probability densities of 0.4 and 0.6 for the first and second bin respectively, while the sensor estimates densities of 0.3 and 0.7, the algorithm calculates the intersection for each interval. It takes the minimum value for the first bin (0.3) and the minimum for the second bin (0.6). The sum of these intersections yields an Overlapping Index of 0.9, indicating a 90% fidelity in the reconstruction of the particle population profile.

The second metric is the Root Mean Square Error (RMSE), which measures the average magnitude of the error between the estimated diameter and the true diameter for each data pair. To ensure the analysis focused on the region of interest for fine dust, a threshold filter was applied to exclude large outliers from the statistical computation, ensuring that the performance metrics reflect the sensor's behavior in its optimal operating range.




Figure 5.10: Reconstructed histogram with diameter estimation.

Figure 5.10 shows the reconstructed histogram of particle dimension displayed to the user in real time during the measurement (in blue), compared against the optical particle size distribution obtained from post-deposition microscopy analysis (in green). The analysis shows that, filtering out particle bigger than $70\ \mu m$, the the overlap index is in the order of the 84%, and the RMSE is $4.64\ \mu m$.

5.2.4. Specific and particular cases

There are some particular cases in which particle detection can show non-standard signal characteristics. Specifically, when a particle deposits at the boundary between the positive and negative electrodes, it simultaneously interacts with both electrode sets, resulting in a transient pulse rather than a permanent step change in the output signal. This occurs because the particle's dielectric field influences both C_{UP} and C_{DW} electrodes equally, creating a spike that is then rejected by the differential architecture, but still can produce a detectable transient response.

Two representative examples of this behavior can be observed in Figure 5.4 on channels 4 and 11. In this figure, particles that produced pulse-type signals are circled and marked with an “s” (for spike), distinguishing them from the typical step-response particles. The pulse width and amplitude depend on the particle's exact landing position relative to the electrode gap and its contact geometry with both electrode surfaces.




Figure 5.11: Detected particles between positive and negative electrodes.

Channel 4 presents a spike with an amplitude of approximately 100 LSB. Based on the power law relationship established in Equation 5.5, the effective contact area of the particle bridging both the positive and negative electrodes corresponds to an equivalent diameter on the order of $2\ \mu m$. The optical microscopy image confirms that the particle is indeed

positioned at the electrode boundary, with portions of its surface in contact with both electrode sets, explaining the observed pulse behavior rather than a sustained step change.

Similarly, the particle deposited on channel 11 shows a smaller contact area with the electrodes and was detected as a 30 LSB transient spike. Applying the same calibration relationship, this corresponds to an equivalent particle dimension on the order of $1.2\ \mu\text{m}$. The reduced signal amplitude reflects both the smaller particle size and the potentially less favorable contact geometry at the electrode interface.

6 | Conclusions and future developments

This thesis work successfully aimed at the design and development of a comprehensive hardware and firmware ecosystem specifically engineered to host and manage the new generation of CMOS capacitive sensors developed at the I3N laboratory. The proposed system has proven to be fully ready for implementation, demonstrating its capability to handle all the critical control signals required for the correct operation of the sensor. By integrating a modular architecture consisting of a Motherboard and a specialized Carrier Board, the system provides a stable and low-noise readout environment, which is fundamental for the precise analysis of micrometric dust particles. The results achieved throughout this work indicate that the developed platform is capable of reconstructing the deposition rate with high reliability and, most importantly, estimating the particle size distribution with significant accuracy.

The experimental activities conducted at the ESA-ESTEC laboratories have confirmed the validity of this technological path. By employing a previous generation of the sensor family, it was possible to fully characterize the detection performance using lunar dust simulant. The analytical results demonstrated a high correlation between the electrical signal magnitude (expressed in LSB) and the physical particle size, allowing for the definition of a Power Law regression model to reconstruct the diameter of the grains. The validation process, performed by comparing electrical signals with optical microscope inspections, achieved a detection accuracy of 85.3% according to the Jaccard index, with an Overlap Index of 84.4%. This campaign enabled real-time tracking of dust deposition and validated the sensing principle in a relevant environment. From a software perspective, the project delivers an intuitive, cross-platform interface that allows the end-user to manage the system on any operating system. This software platform simplifies the interaction with the hardware by providing clear explanations of settings and parameters directly within the GUI, ensuring that the system is accessible even for complex testing scenarios.

Building upon these successful results, future developments will focus on the characterization of the latest sensor generation, which has been specifically optimized to match the particle size requirements defined by ESA, particularly regarding lunar regolith. The newly developed hardware system, characterized by ultra-low noise levels and a compact form factor, provides the necessary versatility to perform these tests in a wide variety of experimental setups. The solid guidelines and regression models established during the characterization of the previous sensor family will serve as a fundamental benchmark. Future efforts will involve calibrating the acquisition software for this new family of sensors, fine-tuning the response to deposition events to achieve even higher precision in real-time diameter estimation and mass distribution analysis, targeting particles in the range of 5 μm to 100 μm .

Bibliography

- [1] J. Schultz. "Cleanroom tech key to success in space". <https://www.nasa.gov/organizations/osi/setmo/cleanroom-tech-key-to-success-in-space/>.
- [2] F. L. Johansson M. Sperl Y. C. Wang, F. Cipriani and M. Adachi. "Electrodynamic dust shield efficiency characterisation under uv in vacuum for lunar application". *Advances in Space Research*, vol. 74, no. 11, pp. 6194–6204, 2024. <https://www.sciencedirect.com/science/article/pii/S0273117724008019>.
- [3] S. K. R. Miller and B. A. Banks. "Degradation of spacecraft materials in the space environment". *NASA Glenn Research Center, Cleveland, OH, United States, Tech. Rep., January 2010, public Use Permitted; Glenn Research Center report*. <https://ntrs.nasa.gov/citations/20110016109>.
- [4] P. V. Wells and R. H. Gerke. "An oscillation method for measuring the size of ultramicroscopic particles.1". *Journal of the American Chemical Society*, vol. 41, no. 3, pp. 312–329, 1919. <https://pubs.acs.org/doi/abs/10.1021/ja01460a004>.
- [5] J. K. Agarwal and G. J. Sem. "Continuous flow, single-particle-counting condensation nucleus counter". *Journal of Aerosol Science*, vol. 11, no. 4, pp. 343–357, 1980. <https://www.sciencedirect.com/science/article/abs/pii/0021850280900427>.
- [6] I. Paprotny, F. Doering, P. Solomon, R. White, and L. Gundel. "Microfabricated air-microfluidic sensor for personal monitoring of airborne particulate matter: Design, fabrication, and experimental results". *Sensors and Actuators A: Physical*, 201:506–516, 10 2013.
- [7] Tiziano Fabrizio D'Amico, Francesco Zanetto, and Marco Sampietro. "CMOS 12 mm² active area dust detector for spaceborne systems based on capacitive sensing". Master's thesis, 2025.
- [8] Florin Ciobanu, Francesco Zanetto, and Marco Sampietro. "Fully electronic system for real-time monitoring of dust particle deposition for space applications". Master's thesis, 2024.

- [9] Texas Instruments. *ADS8867 16-Bit, 100-kSPS, 1-Ch SAR ADC with True-Differential Input, SPI Interface and Daisy-Chain*, 2013. <https://www.ti.com/product/ADS8867>.
- [10] STMicroelectronics. *AN6305, PCB design guidelines for STM32WBA5MMG*, 2025. <https://www.st.com/en/microcontrollers-microprocessors/stm32wba5mmg.htmldocumentation>.
- [11] STMicroelectronics. *ESDA7P60-1U1M, 60 A - 8/20 us, 5 V TVS in QFN-2L*, 2025. <https://www.st.com/en/protections-and-emi-filters/esda7p60-1u1m.html>.
- [12] Texas Instruments. *TPS2521xx, 2.7-V-5.7-V, 4-A, 31-m True Reverse Current Blocking eFuse with Input Reverse Polarity Protection*, 2022. <https://www.ti.com/product/TPS2521/part-details/TPS25210ARPWR>.
- [13] STMicroelectronics. *USBL6-2, Very low capacitance ESD protection*, 2021. <https://www.st.com/en/protections-and-emi-filters/usblc6-2.html>.
- [14] Texas Instruments. *TPS2117, 1.6-V to 5.5-V, 20-mΩ, 4-A low IQ power multiplexer with manual and priority switchover*, 2023. <https://www.ti.com/product/TPS2117>.
- [15] Molex. *Molex SlimStack Board-to-Board connector*, 2005. <https://www.molex.com/en-us/products/part-detail/5024262630>.
- [16] STMicroelectronics. *RM0493 Reference manual, stm32wba5xxx*, 2025. <https://www.st.com/en/microcontrollers-microprocessors/stm32wba5mmg.htmldocumentation>.
- [17] elm-chan. *FatFs Module Application Note*. <https://elm-chan.org/fsw/ff/doc/appnote.htmllicense>.
- [18] Pietro Ciccarella. *"High Sensitivity CMOS Dedicated Electronics for Innovative Impedance- based Sensors"*. PhD thesis, Politecnico di Milano, 2016.
- [19] TOPAS. *Dust Generation Unit (DGU)*. <https://www.topas-gmbh.de/en/products/generation/product/sag-410>.
- [20] Paul Jaccard. "Étude comparative de la distribution florale dans une portion des alpes et des jura". *Bulletin de la Société vaudoise des sciences naturelles*, 37(142):547–579, 1901.

List of Figures

1.1	Capacitive variation as function of time.	3
2.1	Coplanar electrodes architecture and electric field visualization.	6
2.2	Electrodes interdigitate structure.	7
2.3	Real layout of first half of two consecutive channels.	7
2.4	Sensor signal acquisition chain overview.	8
2.5	Front end stage with pseudo resistors configuration.	9
2.6	Pseudoresistors bias network implementation at transistor level.	10
2.7	Gm-C filters bias network implementation at transistor level.	14
2.8	D flip-flop based frequency divider.	15
2.9	Digital logic of the DCC.	16
2.10	Overview circuit of a single channel, presenting all the basic elements: Elec- trodes, Front End, Mixer, DCC, Gm-C filter and Ring Oscillator.	18
2.11	Layout of the ASIC without power MET4 layer & PAD description.	19
3.1	Simplified schematic of the System.	23
3.2	Motherboard connected to the computer and two unmounted Carrier boards.	24
3.3	Carrier board design details.	25
3.4	Detail of the top layer and test pad locations.	26
3.5	π filter, featuring a ferrite bead and decoupling capacitors for EMI sup- pression.	27
3.6	Oscilloscope traces of the sensor signals from S_0 to S_4 , showing the timing and channel switching logic.	28
3.7	Oscilloscope capture of the short reset pulse (approximately 450 ns). This signal is generated by the firmware to trigger the transition from man- ual channel selection, managed by the STM32, to the automatic internal iteration mode of the sensor.	29
3.8	ADC ads8867 schematic functioning from datasheet.	30
3.9	Output buffers of the multiplexer inside the DUST sensor.	32
3.10	GBWP and Phase Margin of C_{FLT} capacitance from 1pF to $10\mu\text{F}$	33

3.11	Phase Margin for C_{FLT} from $1pF$ to $10\mu F$	33
3.12	GBWP and Phase Margin of $C_{FLT} = 10nF$ capacitance.	34
3.13	ADC ads8867 waveform from datasheet.	35
3.14	Top and Bottom side of the motherboard that host control electronics and power system.	37
3.15	Design rules from ST, chosen corner identified as "Good".	38
3.16	Design clearance based on STMicroelectronics official technical specifications.	39
3.17	Basic correspondences between watt and dBm.	39
3.18	Distance measured using Google Earth tool.	40
3.19	Grounding EMC strategy for the System.	41
3.20	Green LED: Standalone operational mode.	41
3.21	Blue LED: Active and stable BLE connection.	41
3.22	Power Management System.	43
3.23	Design configuration, V_{REF} and internal details of TPS2117.	45
3.24	RLC filter schematic.	45
3.25	RLC transfer function.	46
3.26	CP2102N Power Configuration (Self-Powered Mode).	47
3.27	Mechanical coupling of the Board-to-Board (BTB) interfaces.	50
3.28	BTB correct usage on how to remove carriers from the motherboard.	50
3.29	BTB connector pinout.	51
4.1	Compact visual representation of the 100-byte BLE data frame structure.	56
4.2	Compact visual representation of the 68-byte binary data frame structure logged to the microSD card.	57
4.3	Waveform of the FSM for a particular channel.	59
4.4	CP2102N Configuration window, with a detail in <i>Self Powered</i> Mode.	61
4.5	Port Configuration Window, here it is possible to set how the GPIOs behave: GPIO2 and GPIO3 here are set as TX and RX Toggle for motherboard LEDs.	62
4.6	The Connection tab of the GUI, displaying the BLE device discovery and the serial port selection interface.	63
4.7	Live histogram visualization of the sensor deposited particle sizes, allowing instant evaluation of the particle counts and signal distribution.	64
4.8	Visualization of the initial floating state of the electrodes before the analog front-end biasing is fully stabilized.	65
4.9	Stable analog baseline detection. The channels change their voltage when stimulated.	65

4.10	The Dashboard tab, allowing threshold adjustments based on the target material.	66
4.11	The Analysis tab, utilized for parsing and plotting the 68-byte binary frames recorded offline by the hardware's SD card data logger.	67
5.1	Picture of the old CMOS sensor with the main system blocks highlighted. .	70
5.2	Images of the first 16 channels before and after deposition.	71
5.3	The filter applied in Python highlights in red the new particles.	73
5.4	Highlighting of newly detected particles with respective LSB dimension annotations.	74
5.5	Acquisition on Channel 2 revealed 5 different particles.	75
5.6	Particles measured with optical inspection ($6px = 1\mu m$).	76
5.7	Particles revealed by the sensor in LSB.	77
5.8	Fitting functions that describe sensor behavior.	79
5.9	Fitting function detail in the $0\mu m$ to $10\mu m$ range.	80
5.10	Reconstructed histogram with diameter estimation.	81
5.11	Detected particles between positive and negative electrodes.	82

List of Tables

2.1	Implemented values for the tunable bias network components.	11
2.2	Implemented value for k , resistances and capacitances.	14
2.3	Frequency selection logic through digital pads A_0 and A_1	17
2.4	Description of the CMOS sensor PAD signals.	20
3.1	Pinout of the 26-pin board-to-board (BTB) connector.	51
4.1	Detailed explanation of the 100-byte BLE data packet transmitted by the firmware.	56