UNIMORE

UNIVERSITÀ DEGLI STUDI DI
MODENA E REGGIO EMILIA

Dipartimento di Ingegneria "Enzo Ferrari"

Master's Degree in
Electronics Engineering

Development of a TDC-based Ultrasonic Gas Flow Sensor

Candidate: Andrea Circelli

Supervisor: Prof. Giorgio Matteo Vitetta

Co-supervisor: Ing. Felice Di Iorio

Academic Year 2024/2025

Abstract

The transition toward modern smart gas metering is being modelled by the Measuring Instruments Directive (MID). This thesis details the development of a TDC-based (Time-to-Digital Converter) ultrasonic gas flow sensor, starting from an in-depth comparative analysis of current commercial state-of-the-art solutions.
Central to this work is a custom hardware architecture featuring an optimized single-path, V-shaped conduit with a 65° incidence angle. This specific geometry is designed to match a balance: maximizing Differential Time-of-Flight (DToF) sensitivity while avoiding lateral beam deviation and signal attenuation. By operating at a transducer frequency of 500 kHz, combined with a Split Burst excitation technique and a Programmable Gain Amplifier (PGA), the system achieves picosecond-level temporal resolution and robust immunity to Cycle Slip phenomena without heavy computational loads.
To further optimize metrological stability, an adaptive filter was developed for signal processing that merges the zero-flow stability of a Kalman filter with the rapid step-response of an Improved Weighted Recursive Filter (IWRF).
Experimental validation on a certified test bench demonstrates that the prototype is comparable with the current commercial standards. Thanks to mechanical and hardware precision, a simple two-parameter linear calibration was sufficient to confine the percentage error between -1.04% and +1.16%, safely within the strict Maximum Permissible Error (MPE) limits required by MID Class 1.5.
Finally, the optimized hardware-firmware co-design achieves an ultra-low power consumption of approximately $20\mu A$, confirming its suitability for battery-powered smart meters.

Contents

List of Figures

List of Tables

1. Introduction

Gas flow measurement is applied in a wide range of sectors, from residential natural gas distribution to industrial process control, to critical safety and energy monitoring applications. In recent years, technological advances and new national rules have pushed a progressive shift from traditional mechanical meters to static solutions (with no moving parts) that offer superior reliability, virtually eliminating wearing and ensuring long-term stable metrological performance.

With the widespread implementation of smart metering systems, the requirements placed on a gas meter have become extremely stringent, not only accuracy and long-term stability are expected, but also extremely low power consumption, self-diagnostic capability, remote communication, and intrinsic robustness to variations in temperature, in pressure, and in the chemical composition of the gas.

When looking at the range of available technologies, classification is typically based on the physical principle used for detection, the main types include volumetric flowmeters, differential pressure meters, vortex flowmeters, thermo-mass flowmeters, and ultrasonic flowmeters. Each approach brings specific advantages and limitations regarding accuracy, dynamic measurement range, pressure drop induced in the system, and overall construction complexity. For example, industry requires robustness for high flow rates, while the domestic sector focuses on compactness and energy efficiency.

An aspect that cannot be neglected, especially when measuring for billing and trade purposes, is compliance with international regulatory requirements. In Europe, the reference point is the MID (Measuring Instruments Directive), which defines the accuracy classes, the maximum permissible error (MPE) limits, and the repeatability criteria that the instrument must guarantee throughout its operational lifetime. Compliance with these standards is not only a legal requirement, but also a guarantee of economic fairness between supplier and end consumer.

1.1 MID Directive

The international regulatory requirements for gas meters in the EU is Directive 2014/32/EU, commonly known as MID (Measuring Instruments Directive).
For gas meters and volume conversion devices, the specific requirements are detailed in Annex IV (MI-002) of the directive. The regulation defines the flow sensor requirements, which concern not only measurement accuracy within the defined accuracy classes (Class 1.0 and 1.5 for residential and commercial use), but also long-term stability, resistance to external influences (electromagnetic, climatic, and mechanical), and protection toward software and hardware tampering [12].
The document defines the critical parameter that determines the accuracy of the instrument:

Table 1.1: MID parameter definition

Minimum flow rate ($Q_{\min}$)	The minimum flow rate at which the gas meter provides indications that satisfy the requirements regarding maximum permissible error.
Maximum flow rate ($Q_{\max}$)	The maximum flow rate at which the gas meter provides indications that satisfy the requirements regarding maximum permissible error.
Transitional flow rate (Q_t)	The transitional flow rate is the flow rate value occurring between the maximum and minimum flow rates, where the flow rate range is divided into two zones, the "upper zone" and the "lower zone". Each zone has a characteristic maximum permissible error.
Overload flow rate (Q_r)	The overload flow rate is the highest flow rate at which the meter can operate for a short period of time without deteriorating.
Base conditions	The specific conditions to which the measured quantity of fluid is converted.

The normative, for the devices intended to be used for use residential, commercial and light industry, define two classes:

- Class 1.5 that includes domestic and residential meters (typically sizes G4 and G6).
- Class 1.0 that includes larger meters, used in industrial environmental or for large commercial applications (typically G10 and above)

The manufacturer must also specify the instrument's nominal operating conditions, taking into account the following elements:

- The gas flow range shall at least satisfy the following conditions:

Table 1.2: Gas flow ranges

Classe	Q_{max}/Q_{min}	Q_{max}/Q_t	Q_r/Q_{max}
1,5	≥ 150	≥ 10	1,2
1,0	≥ 20	≥ 5	1,2

- The temperature range of the gas, with a minimum range of 40 °C.
- The instrument shall be designed for the range of gasses and the range of delivery pressures in the country of destination. In particular, the manufacturer must indicate the family or group to which the gas belongs and the maximum operating pressure.
- A minimum thermal range of 50 °C with respect to the climatic environment.

The sensor must ensure compliance with the Maximum Permissible Errors (MPE) defined as the deviation between the measured value of the test instrument (Vmis) and the reference volume considered "true" (Vref), provided by a certified calibration bench or a master meter.
The analytical formula for calculating the percentage error ($E_{\%}$) is as follows:

$$E_{\%} = \frac{V_{mis} - V_{rif}}{V_{rif}} \times 100 \tag{1.1}$$

This error must remain within a tolerance band that varies depending on the operating range and the instrument's accuracy class.
The error curve is divided into two operating zones based on the flow rate Q (Table 1.3):

- The zone between Q_{min} and Q_t (excluded) where the fluid dynamic conditions are more unstable and the signal-to-noise ratio of the sensor is lower.
- The Nominal zone, between Q_t and Q_{max} where there is the normal operating in which maximum precision is required.

Table 1.3: Maximum Permitted Errors

Classe	**1,5**	**1,0**
$Q_{min} \leq Q < Q_t$	3 %	2 %
$Q_t \leq Q \leq Q_{max}$	1,5 %	1 %

Finally, the legislation also defines the characteristics related to electromagnetic immunity, durability, and power supply (which guaranties use in the event of a blackout if connected to the grid or guaranties a minimum lifetime of 5 years in the case of battery powered) [12].

1.2 Technology Comparison

Among the predominant technologies for gas flow measurement, two differ in importance and diffusion: thermal dispersion mass meters (thermomassive) and ultrasonic meters [10, 17].

The operating principle of thermomassive meters is based on the heat transfer laws of a moving fluid. The typical sensor architecture involves the use of two thermometric resistors, the first measures the actual temperature of the fluid acting as a reference, while the second is electrically heated in order to maintain a constant temperature difference compared to the first. When the gas flows through the pipe, the fluid molecules remove a specific amount of heat from the heated sensor by forced thermal convection, the electrical current necessary to compensate for this dispersion and keep the thermal jump constant is directly proportional to the mass flow rate of the gas. This principle exploits an intrinsic property of the gas, such as its specific heat at constant pressure (c_p) which determines the capacity of the molecules to absorb heat [1, 6].

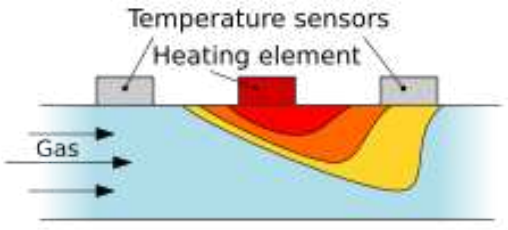

Figure 1.1: Thermomassive technology

As for ultrasonic sensors, there are two main variants: Doppler effect meters and transit time (ToF) meters. Since Doppler meters require the presence of particles or bubbles to reflect the acoustic signal, in the natural gas sector ToF meters are used exclusively, in which a pair of piezoelectric transducers (or multiple in a multi-path

configuration) are mounted with a certain angle θ which respects to the flow axis. The transducers alternately emit acoustic pulses upstream and downstream, the acoustic wave that travels in the same direction of the flow is accelerated, while the one traveling in the opposite direction is slowed down. Through a rigorous physical-mathematical model, it is possible to derive the gas velocity (v) as a function of the upstream (t_{up}) and downstream (t_{down}) ToFs, the acoustic path length (L) and the incidence angle (θ) [5, 6], as will be subsequently discussed in paragraph 2.1.

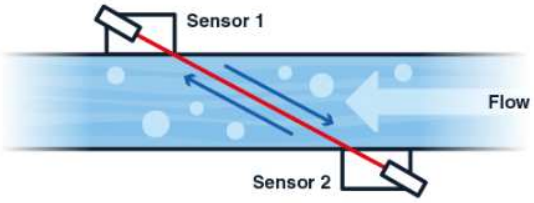

Figure 1.2: Ultrasonic technology

A fundamental divergence between the two technologies regards the nature of the variable directly extrapolated from the sensor. The thermomassive is a direct mass measurement instrument, this provides the advantage of not requiring compensation calculations for temperature or pressure fluctuations in order to obtain the transferred mass value. On the contrary, the ultrasonic ToF calculates the average gas velocity and then derives the instantaneous volumetric flow rate at the actual operating conditions, since the gas is a highly compressible fluid, for billing purposes or for standardized gas measurement (Sm^3), the meter must be coupled with temperature and pressure sensors, to convert the measured volume into the volume in Base conditions [1, 17].

The two technologies are designed to satisfy different fluid dynamic requirements and precision needs. Thermomassive meters offer an intrinsic high sensitivity at low flow rates, they offer an extremely wide measurement range (turndown ratio) up to even 1000:1, but their relative accuracy is generally positioned on average values (often between 1% and 2% of the reading). On the contrary, ToF ultrasonic meters reach a higher degree of precision, ensuring uncertainties that can reach between $\pm 0.5\%$ and $\pm 0.1\%$. However, the critical limit of the ultrasonic technology is detecting extremely reduced ToF differences at low flow rates since the difference between t_{up} and t_{down} goes into the order of picoseconds. Consequently, while guaranteeing very high precision, ultrasonic systems require a more complex signal

processing compared to the thermal counterpart [1, 2].

1.2.1 Robustness to Gas Composition Variations and Hydrogen Blending

The resistance of the two meters to the chemical fluctuations of the transported gas highlights the most important objective disparity. Modern challenges, such as the blending of hydrogen within natural gas distribution networks for decarbonization purposes, profoundly alter density, specific heat, and thermal conductivity.
Ultrasonic technology is mathematically independent from the speed of sound in the medium, consequently, abrupt variations in the gas mixture do not directly affect the determination of the flow velocity. Field studies demonstrate that the metrological errors of ultrasonic instruments remain almost unchanged also in the presence of H2NG mixtures with hydrogen up to 23% by volume, once this threshold is exceeded, the sensor can no longer process the signal because hydrogen has a speed of sound approximately 3.7 times higher than that of air or methane, consequently, as hydrogen is injected, the speed of sound of the mixture increases until it exceeds the maximum limit of hardware design. Furthermore, an advantage of ultrasounds is the ability to inversely calculate the speed of sound in real time, making it possible to deduce gas alterations or the presence of contaminations.
In the case of thermomassive measurement, since the principle is crucially based on the thermal absorption of the fluid, the original calibration of the instrument is strongly dependent to a chemically specific gas (its specific heat (c_p) and its thermal conductivity (λ)). In the presence of hydrogen, which has a thermal conductivity three times higher than air or methane, the heat transfer is massively modified, experimental tests show that the measurement maintains metrological conformity only if the hydrogen fraction does not exceed 2%, with mixtures containing 23% of hydrogen, the alteration of the thermophysical profile brings the weighted mean error (WME) up to 15.8% in the worst cases. Although modern sensors and predictive algorithms attempt to mitigate these discrepancies, the use of thermal meters in networks with a highly unstable or heterogeneous chemical composition remains problematic if not strictly calibrated for the specific mixture [13].

2. Fundamental Principles

The ultrasonic sensors differs from mechanical technology, that depends on kinetic energy of the flux, or from thermomassive technology, that depends on thermal property of the gas, the ultrasonic method is based on a Time of Flight (Transit-Time), it is derived from the variation of velocity that an ultrasonic wave ensures where across a flux in moving. The wave that travels in the flux direction (downstream) is accelerated by the flux velocity, while the wave that travels in the opposite way of the flux (upstream) is slowed down.
In this chapter we will see a physical-mathematical model that relates the Time-of-flight (ToF) with the average flux velocity.

2.1 Transit-Time Measure

The function of the sensor is based on the Differential Time-of-Flight (D_{ToF}) method. The physical concept is that the propagation velocity of an acoustic wave in a moving fluid is the vector sum of the speed of sound in the static medium (c) and the component of the fluid velocity (v) projected along the direction of propagation.
In an ultrasonic system, two transducers (A and B) alternatively works as emitter and receiver, the ultrasonic wave traveling in the direction of flow (downstream) is accelerated by the gas velocity, while the wave traveling in the opposite direction (upstream) is slowed down.

2.1.1 Time of Flight derivation

As mentioned above the propagation velocity is the vector sum of the sound of velocity in a static medium (c) and the component of the fluid velocity (v) along the path, we can write:

$$c_{up} = c - v\cos\theta \tag{2.1}$$

$$c_{down} = c + v\cos\theta \tag{2.2}$$

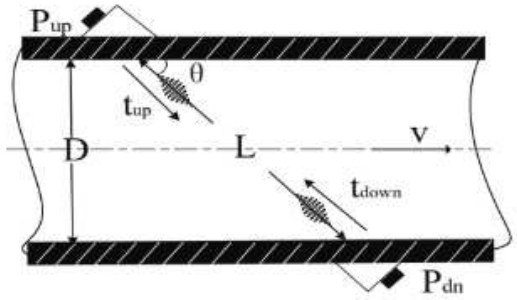

Figure 2.1: Schematic diagram of the principle of time-difference method

As a consequence, the absolute value of ToF is given by the equations 2.3 and 2.4, where L is the length of the ultrasonic path, c is the velocity of sound in the medium, v is the flow velocity, and θ is the angle of incidence:

$$t_{up} = \frac{L}{c - v\cos\theta} \tag{2.3}$$

$$t_{down} = \frac{L}{c + v\cos\theta} \tag{2.4}$$

2.1.2 Flux velocity derivation

One of the main advantages of the D_{ToF} technique is the deletion of the dependence on the speed of sound c, making the measurement robust to changes in gas composition. Combining equations 2.3 and 2.4 we can define the gas velocity v as a function of the measured times:

$$v = \frac{L}{2 \cdot \cos(\theta)} \cdot \left(\frac{1}{t_{down}} - \frac{1}{t_{up}}\right) \tag{2.5}$$

The equation 2.5 shows that the velocity is independent from the sound velocity (c) of the medium and so this implies that the velocity is not directly influenced by the gas composition.

2.1.3 Derivation of volumetric flow rate

The velocity (v) obtained in the equation 2.5 represent the average velocity along the ultrasonic path, but, the flux in the pipe can be laminar or turbulent modifying the velocity profile making the velocity along the ultrasonic path different from the average velocity along the overall tube, for this reason an hydrodynamics correction

factor[1] (k or K_p) must be applied. To obtain the instant volumetric flow rate (Q) is obtained by multiply the velocity for the transversal section (S) of the tube:

$$Q = k \cdot S \cdot v \tag{2.6}$$

2.2 Thermodynamic Dependence

While the equation 2.5 deletes the sound velocity (c) dependence in the calculation of the velocity flux (v), the tacking of c is fundamental. From the ToF measurements the system can derive the sound velocity c_{means} in the medium [14]:

$$c_{means} = \frac{L}{2} \cdot (\frac{1}{t_{up}} + \frac{1}{t_{down}}) \tag{2.7}$$

For an ideal gas, the theoretical sound velocity depends on temperature (T) and molar composition following the Newton-Laplace equation:

$$c_{theo} = \sqrt{\frac{\gamma \cdot R \cdot T}{M}} \tag{2.8}$$

From the comparison between c_{means} and c_{theo} is possible to know if there are contamination in the fluid watching if the measured sound velocity significantly differs from the theoretical one for the gas at a given temperature [16].

Another parameter that affects the measured value is the pressure; in fact, the gas is a compressible fluid and so the quantity of matter (in number of moles) and so the energy content present in a cubic meter of gas varies, in order to obtain a measure that represents the effective mass of gas transported it is necessary to normalize the measured volume to standardized reference conditions [12].
In order to obtain a base volume (V_b) from the measured volume (V_m) is necessary convert the parameter following the relation:

$$V_b = V_m \cdot \frac{P_{mis}}{P_b} \cdot \frac{T_b}{T_{mis}} \cdot \frac{Z_b}{Z_{mis}} \tag{2.9}$$

where P_{mis} and T_{mis} are the Pressure and Temperature measured to the operating condition, P_b and T_b are the base pressure and the base temperature (1013.25hPa and 288.15K), Z_{mis} and Z_b are the gas compressibility factors at measured and base conditions. In a domestic gasmeter the preassures are low (few mbars) and so the

[1]The correction factor k depends by the Reynolds number (Re) and from the flux profile, typically k=0.75 for laminar regime and k=0.95 for turbulent regime

impact of Z_b/Z_{mis} can be assumed equal to one ($Z \approx 1$), but becomes essential in industrial applications [11].

3. Mechanical design

The accuracy, the sensitivity, and the energy consumption of an ultrasonic gas flow meter based on time-of-flight (ToF) depends, in critical way, on the mechanical propagation and spatial position of the ultrasonic transducers. The topology choice defines the interaction between the ultrasonic wave and the fluid, determining the main performances of the system before that the signal is elaborated [14].

3.1 Single-path and Multi-path

The first fundamental classification in ultrasonic flow sensor design concerns the number of transducers used to calculate the ToF.
In a multi-path system, multiple pairs of ultrasonic transducers are used to measure gas velocity along different directions of the pipe cross-section. Therefore, the multi-path design offers high accuracy in compensating for flow noise, making it useful for industrial applications. However, they have higher hardware costs, greater mathematical complexity, and higher power consumption, making them an oversized solution for low-power devices [20].
On the other hand, the single-path topology uses a single pair of transducers, the ultrasonic wave travels through the pipe along a single defined path, however, optimizind of the duct shape and wave reflection geometry becomes crucial to have high accuracy and low noise. You can see the two system paths in the figure 3.1.

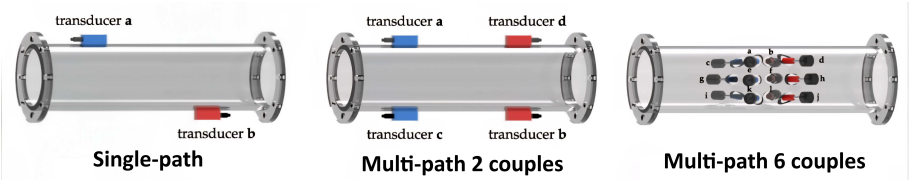

Figure 3.1: Single and Multi path

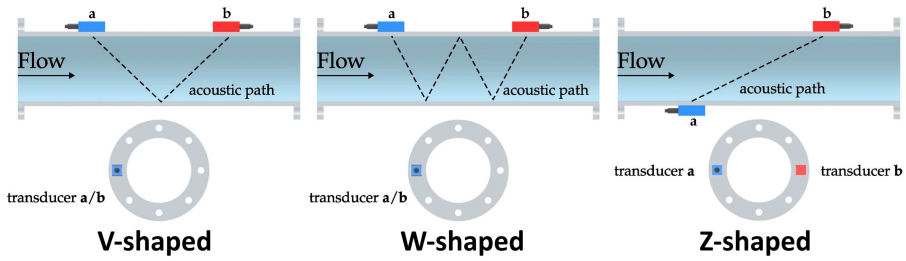

Figure 3.2: V W and Z shaped

3.2 Clamp-on

In the single-path system there are several way to arrange the transducers, which can lead to different configuration, some examples are the V-shaped, the Z-shaped and the W-shaped. An overview of this mounting is shown in figure 3.2 [9].

- **Z-shaped**: the transducers are mounted on the opposite side of the tube and the ultrasonic wave across the fluid in direct line, the wave moves on a distance equal to the distance of the transducers.
- **V-shaped**: both transducers are placed on the same side, the ultrasonic wave is emitted across the fluid and bounces one time on opposite face and go on the other transducer, this allows the waves to move on a distance higher than the physical distance of the transducers.
- **W-shaped**: the ultrasonic wave suffers multiple reflection (typically 3) on the face of the tube before to go to the receiver; it is used is a very thin tube to maximize the path length.

3.3 Topology Selection

An analysis of experimental tests shows that the design based on a Z-shaped has an uncertainty higher risk than the V-shaped, particularly at high flows, due to a higher deviation of the ultrasonic wave and on negative incidence angular as you can see in figure 3.3 where different V-shaped and Z-shaped are compared [14]. However, the W-shaped theoretically maximizes the Δt increasing the length but every reflection gives a loss of energy [21].

The V-shape offers a series of advantages, in the first place, this architecture has an intrinsic capacity in noise compensation due to fluid dynamics, in fact, the velocity profile in a pipe is not uniform due to turbolence. Since the V-shape has a double crossing of the tube section, the ultrasonic wave propagates along specular

paths, allowing a self-compensation of transversal deviations or disturbances. In the second place, in parity of diameter, the V-shape doubles the length of the acustic path (L) with respect to the Z-shape. A higher L means a higher interaction between ultrasonic wave and the flow of gas, increasing proportionality the difference between measurable ToF increasing the accuracy at low flux, so it increases the Signal-to-Noise Ratio (SNR) [14].

The mathematical core of mechanical sizing is the balance (trade-off) between the incidence angle (θ), the acoustic length (L) and trasversal section of the tube (S). From these three parameter depends the mechanical uncertainty of the measurement.

Modifying the transversal section (S) returns different fluid velocities (v) for the same volumetric flow rate (Q), since v=Q/S. If S is reduced, v increases, which improves the resolution at low flow rates. However, if v becomes too high, the ultrasonic beam suffers a larger lateral deviation (it is physically shifted downstream by the moving gas). This deviation is a parasitic effect that causes asymmetric signal attenuation and different ADC signal levels between the upstream and downstream directions, increasing the measurement uncertainty.

Modifying the acoustic length (L) it have a different value of D_{ToF} ($D_{\text{ToF}} \propto \frac{L \cdot v \cdot \cos\theta}{c^2}$, where c is the sound velocity), an higher L means an higher DToF and it increase the sensitivity, increase the resolution at low flux but increases the signal attenuation and increase also the deviation of the beam.

Modifying the incidence angle (θ) also change the D_{ToF} ($D_{\text{ToF}} \propto v \cos\theta$), if θ increase $\cos\theta$ decrease, and so D_{ToF} decrease and it have a lower sensitivity but it have a more compact route and lower sensitivity to lateral deviation of the beam [14].

Table 3.1: Geometrical and performance parameters of the analyzed designs in [14]

Design	L(mm)	W(mm)	H(mm)	S(mm^2)	θ(°)	Min.D_{ToF}(ns)	Max.v(m/s)
V1	73.8	9.9	19.2	190.55	45	51.71	10.50
V2	61.1	8.3	23.6	194.75	65	25.04	10.27
V3	58.8	8.5	22.3	189.49	67	22.90	10.55
V4	54.7	10.2	19.1	195.30	63	24.00	10.24
V5	63.4	8.1	23.7	191.13	65	26.48	10.46
V6	58.6	9.1	22.5	203.39	65	23.00	9.83
V7	56.6	8.1	20.7	166.32	65	27.17	12.02
Z1	54.8	5.0	36.0	180.21	60	28.72	11.10
Z2	50.1	6.3	30.6	192.98	58	25.96	10.36

Another mechanical choice that affects the uncertainty is due to the separation plates for flow stabilization purposes, the figure 3.4 shows the uncertainty achieved

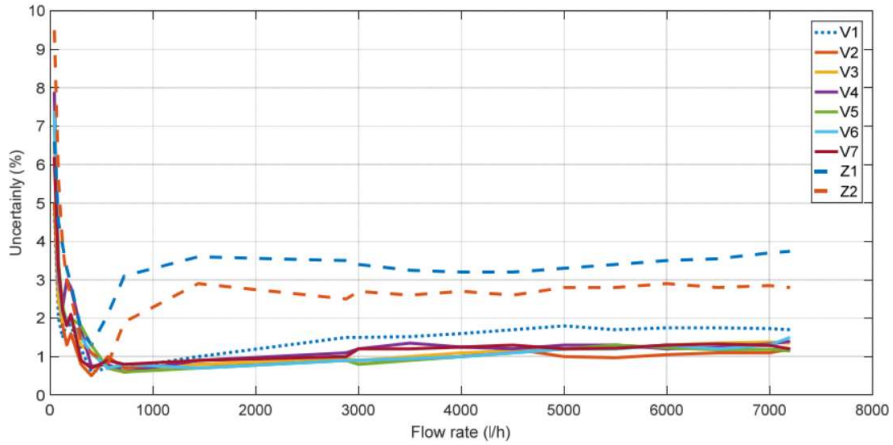

Figure 3.3: Uncertainty by shape analyzed in [14]

with a V design with and without lamination (V2 that has L=61.1mm, W = 8.3mm, H=23.6mm, S=194.75, and $\theta = 65°$).

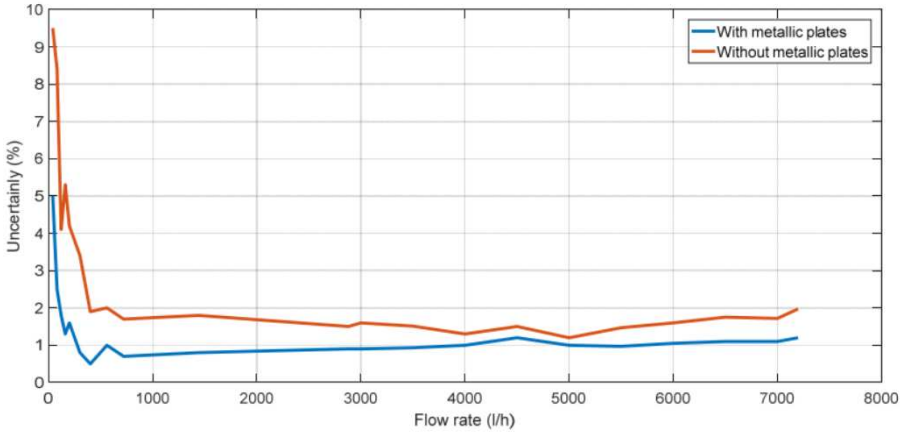

Figure 3.4: Uncertainty achieved with and without lamination analyzed in [14]

3.4 Transducer Frequency Selection

The frequency choice for transducers is typically in the 200 kHz - 500 kHz range; this frequency band represents the optimal trade-off between signal attenuation and timing precision for gas applications. Higher frequencies (in the MHz range, usually adopted for liquid meters) suffer from excessive acoustic attenuation in gases. On the contrary, lower frequencies exhibit a slower slew rate at the zero-crossing point increasing the timing uncertainty in the case of TDC (Time-to-Digital) architecture as we will see in chapter 4.1.

In gases, the ultrasonic attenuation increases quadrically with frequency ($\alpha_{att} \propto f^2$), experimental data shows that in air the attenuation is more or less 42.7 dB/m @ 200kHz and 44.8 dB/m @ 500kHz, and this difference amplifies in a real gas media(e.g. methane), as a consequence a 200kHz system has a greater received signal amplitude, ensuring in lower gain in amplifies stage (ensuring also lower power consumption) and with a higher SNR. Another advantage of 200kHz regards the zero-flow drift (800ps @ 200kHz vs 1200ps @ 500kHz), allowing to 200kHz to have a lower threshold of minimum detectable flux (i.e 1.3L/h @ 200kHz vs 2.1L/h @ 500kHz), providing an higher accuracy at low flux without using complex compensation algorithms [14].
On the other hand, a 500kHz signal have a period of only $2\mu s$, which respects $5\mu s$ of 200kHz, this entails into a significantly steeper signal slope (slew rate) in the point in which the signal across the zero, this higher slope allow the detecting of Zero-Cross Detection higher robust to the SNR, allowing to TDC to obtain a time resolution of picoseconds.

4. ToF Detection Algorithms

The signal transmitted by the ultrasonic wave, after propagating through the fluid medium, reaches the receiving transducer with a very low amplitude (often in the order of tens of millivolts) and high noise interference, which is why ToF detection algorithms play a fundamental role in defining the final quality of the measurement [8].

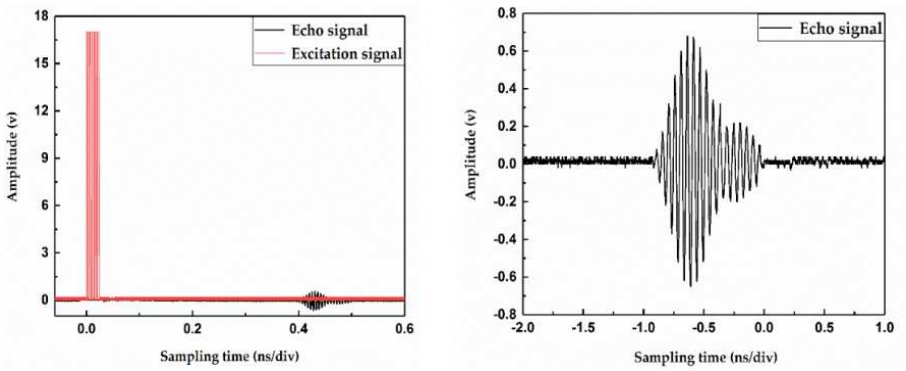

Figure 4.1: Transmitted and received signal

The main two methods to extract ToF in a transit time-based gas flow sensor are based on:

- time-to-digital (TDC), based on the threshold and zero crossing detection of the received signal
- analog-to-digital (ADC), based on the capture of the whole received waveform

4.1 Time-to-Digital Converter

The simplest way to deduce a ToF measure by a TDC is by the threshold detection. The measurements beings when the signal is transmitted and ends when the received echo signal first exceeds the threshold point.
However, according to acoustic energy formula, the ultrasonic sound energy density

is:

$$E = \frac{\rho \cdot c \cdot \omega^2 \cdot A^2}{2} \tag{4.1}$$

where E represents the sound energy density, ρ is the medium density, c is the ultrasonic speed, ω is the ultrasonic frequency, and A is the ultrasonic amplitude. Variuos factor such as impurities, uneven distribution of velocity fields, and changes in temperature and pressure can cause variation in ultrasonic energy attenuation resulting in a decrease in amplitude of the echo signal [15, 23].

If the amplitude variations is relatively small could lead in a phenomenon called Amplitude-dependent time walk. Mathematically, the ultrasonic signal received near the first wavefront can be approximated as a sine wave ($V(t) = A \cdot sin(2\pi ft)$, where A is the maximum amplitude and is dependent on the gas attenuation and f is the resonant frequency of the transducer that is a constant). The rate at which the signal grows over time, called Slew Rate (SR), is defined by the first derivative of the voltage with respect to time and so it have its maximum value when the slope across the zero ($\cos(2\pi ft) = 1$) demonstrating that the slope of the signal is directly proportional to its amplitude (since f is constant):

$$SR_{max} = A \cdot 2\pi f \tag{4.2}$$

To free time measurement from amplitude fluctuations the Zero-Crossing Detection technique can be applied. With this technique the time measurement ends when the wave equation is equal to zero ($A \cdot sin(2\pi ft) = 0$), so when the sin argument is an integer multiple of π. The time instant in which the wave crosses zero depends only on f of the transducer and the wave phase, resulting mathematically independent of SR variations and therefore independent of amplitude (A), a graphical representation of the threshold and zero-crossing can be seen in Figure 4.2.

Instead, if the amplitude attenuation variation due by gas condition is relatively high, the waves could not reach the threshold amplitude with the same peak, leading in a critical issue called Cycle Slip, as you can see in figure 4.3. In this case the system will be affected from ad error equal to an entire period ($T = 1/f$) totally variates the flux measure, to solve this issue a Dynamic Threshold is applied, figure 4.4 [25].

The main advantage of TDC is the low energy consumption because it is not a continuous sampling, but the circuit wakes up only in the microseconds required for signal emission and reception. Another two advantages are the low computational load, allowing to MCU to do only simple opreations staying in sleep mode for most

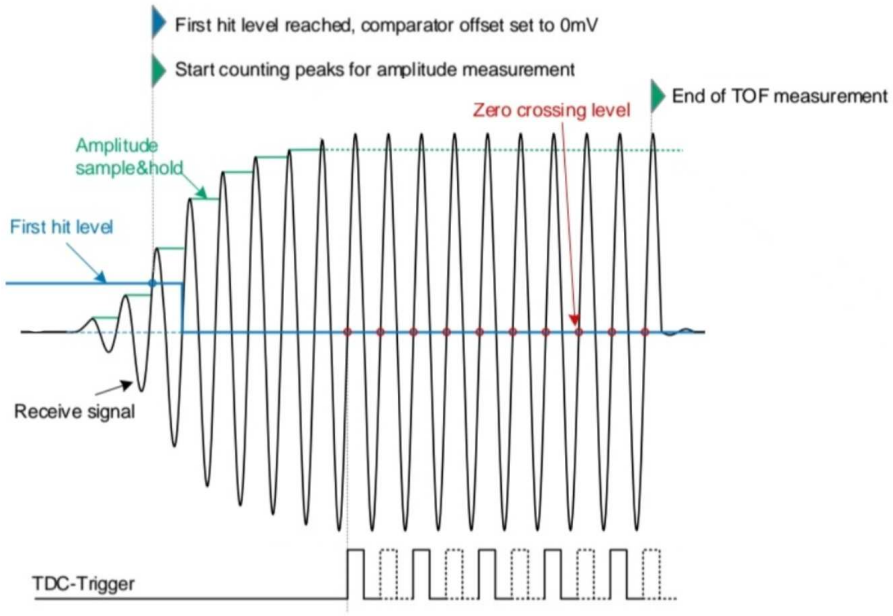

Figure 4.2: TDC method

of the time (the latency between receiving the signal and obtaining the data is practically zero, allowing the system to return to "deep sleep" almost instantly), and the costs, due to a simple architecture. Moreover, modern TDC chips allow to measure the time intervals with a resolution in the order of picoseconds, to obtain this result with an ADC it is mandatory to use very high sampling frequency.

4.1.1 Dynamic Threshold Adjustment Method

This approach is divided into two sequential phases: the definition of a reference value via a weighted recursive filter (which will be explained in the chapter 5.2.2) and the subsequent feedback adjustment of the threshold voltage [25].

Once a reliable filtered value $y[n]$ is obtained, the system performs a feedback check by comparing the latter with a group of N historical measurements x_i. The goal is to understand whether the current voltage threshold is locking on the correct wave peak, or whether, due to signal attenuation or amplification, it is detecting an earlier or later peak.

The algorithm calculates two counters:

- R: The number of times the measured value is greater than the filtered value ($x_i > y$), indicating that the threshold captured a peak late with respect to the target.

- F: The number of times the measured value is less than or equal to the filtered

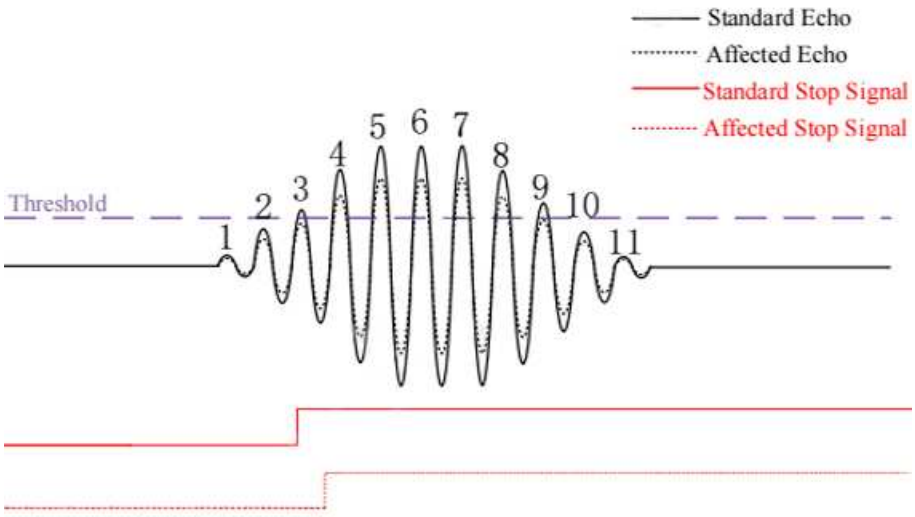

Figure 4.3: Cycle Slip phenomena

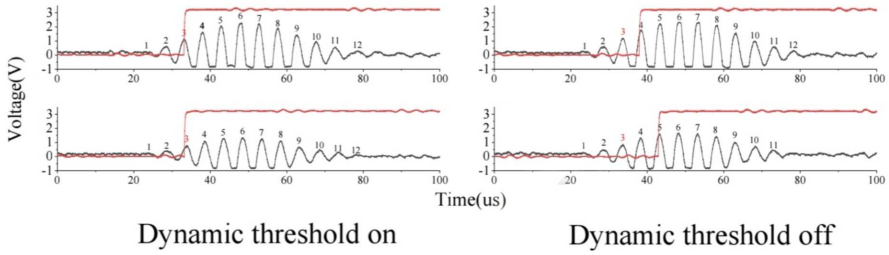

Figure 4.4: Dynamic threshold

value, indicating early detection.

By defining the error as $E = R - F$, the threshold update follows these decision rules:

1. If $|E| \leq N/2$, the error distribution is balanced. This indicates that the threshold is correctly centered on the characteristic peak and is kept unchanged.

2. If $E > N/2$, the system is predominantly detecting peaks after the target. This is an indicator of a decrease in signal amplitude (e.g. due to an increase in flow rate), making it necessary to lower the threshold level.

3. If $E < -N/2$, the system is capturing earlier peaks. The amplitude of the target is increased and the threshold must be raised to avoid early false triggers.

To prevent instability at very low flows or in the presence of impurities, the entire dynamic compensation algorithm is activated only when the flow rate exceeds a minimum safety value Q_s [25].

4.2 Analog-to-Digital Converter

The ADC digitizes the entire received acoustic waveform. To obtain a high time resolution, the signal is typically oversampled at frequencies range from 4 MHz to 16 MHz. Having a complete digital representation of the signal allows at the application to have versatile mathematical algorithms, enabling two different approach analysis: the extraction of the absolute ToF or the direct calculation of their difference (Δt) [24].
First, the absolute ToF (t_{up} and t_{down}) can be calculated by applying digital threshold or envelope detection algorithms to the sampled array. This software-based approach mimics the logic of a TDC but offers greater immunity to noise, providing the absolute temporal values often required by the flow velocity equations.
Second, to obtain an extremely precise time difference, it calculate the Δt directly by applying a cross-correlation algorithm. This algorithm calculates the cross-correlation function $R_{12}(\tau)$ between the downstream acquired segment ($x_1(t)$) and the upstream acquired segment ($x_2(t)$):

$$R_{12}(\tau) = \int_{-T/2}^{T/2} x_1(t)\, x_2(t+\tau)\, dt \tag{4.3}$$

Where the time delay τ is equal to (Δt), the Cross-Correlation function reach its maximum value, as shown in Figure 4.6 [14, 18].

Comparative studies [14] demonstrate that the ADC-based method offers higher accuracy and greater robustness against noise than TDC, especially at high flow rates, since it uses the global properties of the wave and not a single point. However, it involves significantly higher power consumption and requires computational resources not compatible with ultra-low power microcontrollers (MCU) [25].
However, similarly to TDC systems, if the signals suffers an high deformation travelling across the gas the cross-correlation algorithms can be confused between two near peaks resulting in a Cycle Slip, a way to prevent this situation is using the Hilbert Transform (HT).
From a mathematical point of view, the application of this transform to a real signal $x(t)$ is defined by the following integral:

$$\tilde{x}(t) = \mathcal{H}[x(t)] = \frac{1}{\pi} \int_{-\infty}^{\infty} \frac{x(u)}{t-u} du \tag{4.4}$$

In practice, this operation is equivalent to calculating the convolution between the starting signal and the function $1/(\pi t)$. The result of this transformation represents

the imaginary component which, once added to the original real signal, defines the so-called analytical signal:

$$x_a(t) = x(t) + j\tilde{x}(t) \tag{4.5}$$

The main advantage of this formulation is the possibility of deriving the envelope of the signal $x(t)$ by simply calculating the modulus of the analytical signal just described:

$$|x_a(t)| = \sqrt{x^2(t) + \tilde{x}^2(t)} \tag{4.6}$$

Once the envelopes of the signals under consideration have been extracted (e.g. an emitted signal $x(t)$ and a received one $y(t)$), a cross-correlation can be performed to locate the peak position and estimate the ToF. A more often advantageous algorithmic alternative is to apply cross-correlation by using the Hilbert Transform of the signal $y(t)$, obtaining a function (CCFHT) expressible as:

$$R_{x\tilde{y}}(\tau) = CCF[x(t)\tilde{y}(t+\tau)] \tag{4.7}$$

Due to the limitations imposed by the sampling rate f_s, an estimate based only on finding the maximum value of $R_{x\tilde{y}}(\tau)$ provides a raw resolution (the ToF being proportional to n/f_s, with n representing the identified sample). To overcome this problem and obtain a sub-sample resolution (necessary for calculating the DToF), zero-crossing detection (zero-crossing) is used.

The application advantage of this technique lies in the possibility of using a simple linear interpolation on the CCFHT data around the zero-crossing point to correct the peak position. Compared to the heavier parabolic or Gaussian interpolations used traditionally, linear interpolation provides a much lighter, lower-power implementation while maintaining the desired "subsample" accuracy [14].

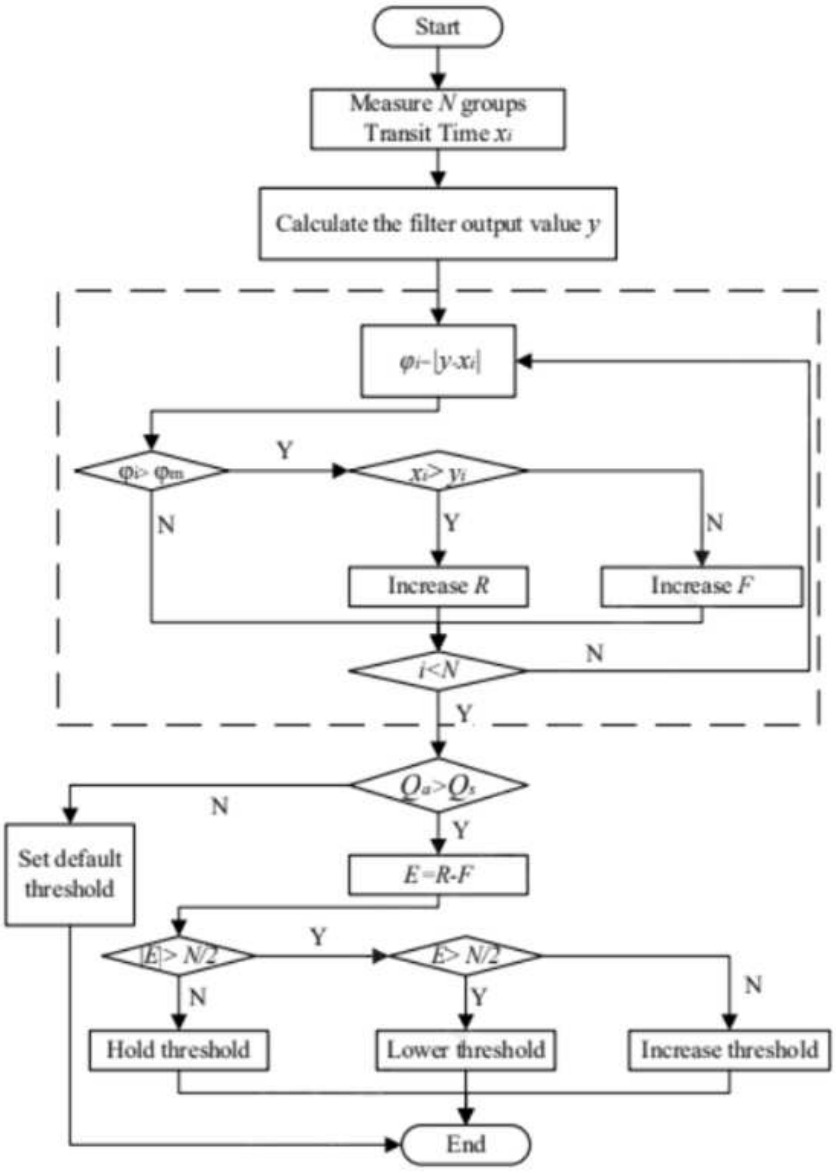

Figure 4.5: Feedback-Type Dynamic Threshold Method Flowchart in [25]

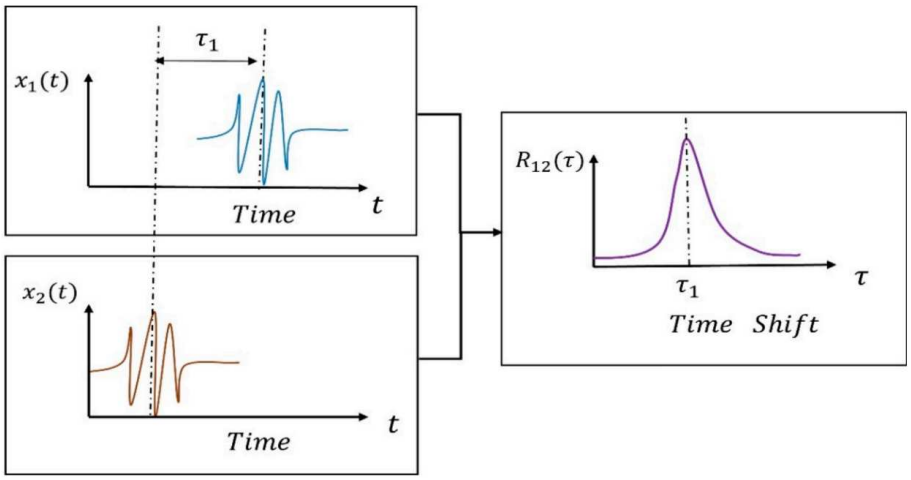

Figure 4.6: Schematic diagram of the principle of Cross-Correlation

5. Post-Processing data

The raw numerical value ToF obtained by TDC (Time-to-Digital Converter) or ADC (Analog-to-Digital Converter) is typically affected by error due to physical phenomena and ambient noise that affect the ultrasonic signal and by temporal quantization. The errors can be classified into three categories:

1. Impulsive noise (outliers): extreme values caused by electromagnetic interference, air bubbles, impurities, cycle skipping or false triggering phenomena. These data must be identified and fixed prior to any future calculations.

2. White noise (jitter): stochastic fluctuations around the real value, due to thermal noise in the electronics and the temporal quantization. This noise is particularly critical at low flow rates [22].

3. Flow dynamics: real changes in flow rate that should be over-filtered. Too aggressive filtering would introduce a delay in the system response, compromising accuracy [8].

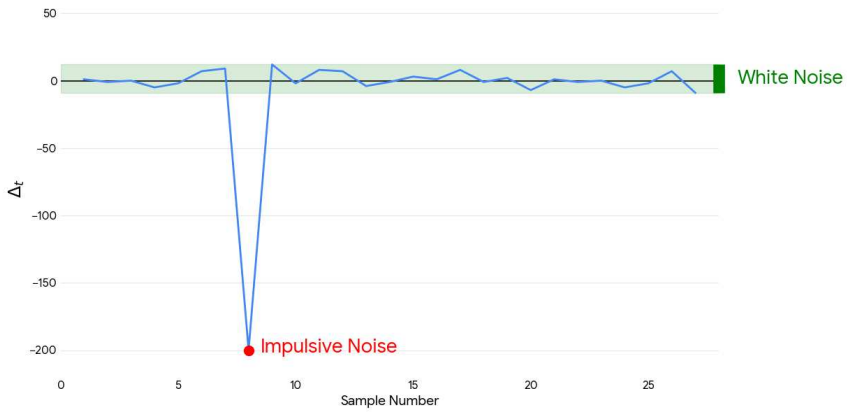

Figure 5.1: Time Difference under static

5.1 Pre-Filtering

In order to mitigate the error due to Impulsive noise the data can be sorted, filtered and error compensated.
Although numerous sorting algorithms are available (e.g. bubble sort, insertion sort), the Shell sort is particularly suitable for ultrasonic flowmeters due to its low computational complexity, which is highly advantageous given that the typical sample size is small (usually n = 10).
After sorting the data will be analyzed in order to remove peaks due to impulsive noise, to do this we can use the Peters methods. The Peters formulas is used to calculate the standard deviation (s):

$$s = \sqrt{\frac{\pi}{2}} \sum_{i=1}^{n} \frac{|x_i - \bar{x}|}{\sqrt{n(n-1)}} \tag{5.1}$$

where x_i represents the i-th sampled data point, $\bar{x}$ is the average of the measurements, and n is the number of measurements. For a number of measurements $n \leq 10$, a new measurement x_{new} is considered valid if it satisfies the condition:

$$|x_{new} - \bar{x}| \leq 2s \tag{5.2}$$

If the measurement falls within this tolerance it is considered effective, otherwise it is classified as a random error and must be mitigated, typically it is re-established with an average value compared to adjacent values [22].

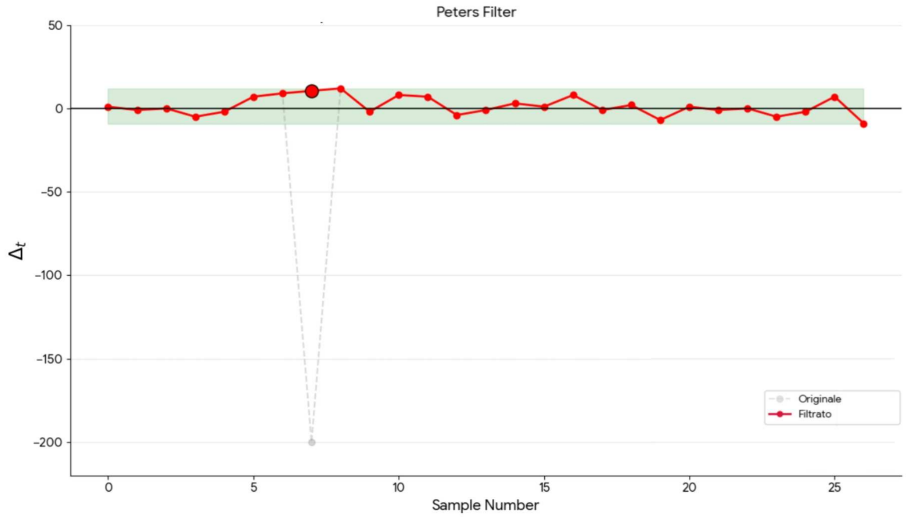

Figure 5.2: Peters Filter

5.2 Post-Filtering

The White Noise effect is removed in post-processing with classical, low computational, digital filter like Movine Average, Improved Weighted Recursive, Kalman or alwo with a combination of these like in Kalman and Arithmetic Average Filer, how you see in this chapter every filtering algorithm have its pro and cons.

5.2.1 Moving Average Filter

The Moving Average is the most simple filtering that can be applied. The algorithms works on a time windows of n samples and the output value is simply the arithmetic average of last n samples:

$$y_k = \frac{1}{n} \sum_{i=0}^{n-1} x_{k-i} \tag{5.3}$$

where x_{k-i} is the sample at $(k - i)$ instant and n is the window size. The effect of this technique depends on window size, an high n means an higher filtration resulting in a smoother curve, however, if the window size is too large, the filter introduces a significant time delay proportional to $(n-1)/2$ and excessively smooths the data, compromising the system's capacity to track rapid flow variations and abrupt changes in the real trend [25].

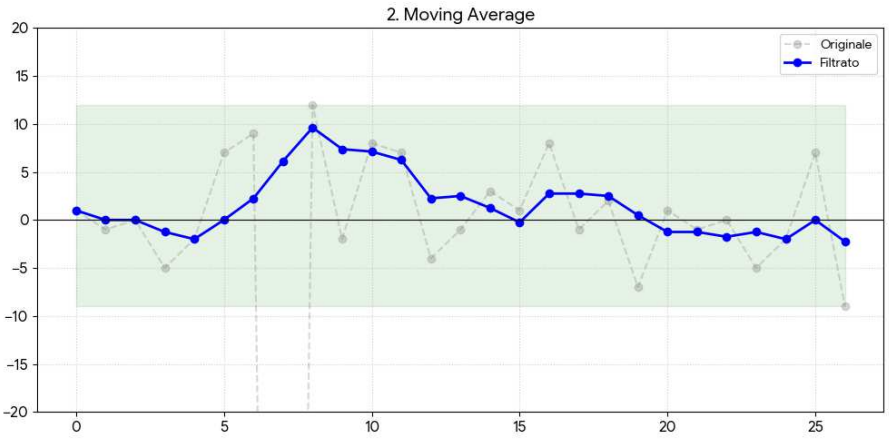

Figure 5.3: Moving Average Filter in rest condition

5.2.2 Improved Weighted Recursive Filter

Another typical approach, to minimize the impact of noise, is the Standard Recursive Filter (or Exponential Moving Average), which calculates the current filtered value

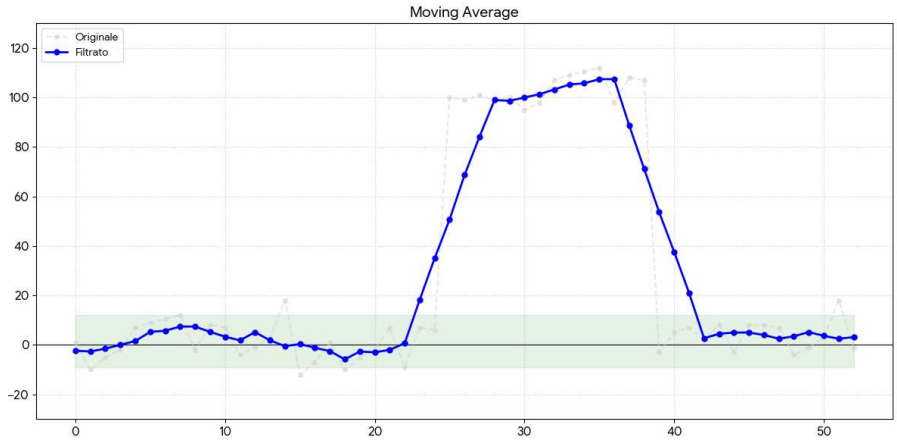

Figure 5.4: Moving Average Filter in flux variation condition

y_n by combining the raw measured data x_n with the previous filtered output y_{n-1}:

$$y_n = \alpha \cdot x_n + (1 - \alpha) \cdot y_{n-1} \tag{5.4}$$

where α is a fixed smoothing factor between 0 and 1.
While this filter, much like the simple Moving Average, is computationally light and effective in steady-state conditions, it presents significant disvantage during dynamic operations because it assigns fixed weights suffering of the same problem of Moving Average Filter, so it suffers from a "smearing" effect and introduces a noticeable lag [8]. To overcome the lag and the inadequate dynamic response of classical filters, the Improved Weighted Recursive (IWR) Filter can be applied. Unlike standard methods, this algorithm dynamically assigns a specific weight (W_i) to each data based on the mathematical and physical characteristics of the transit time [25]. The weight of a single measurement is evaluated through two diagnostic parameters:

- Consistency (λ_i): The absolute difference between the variation of the upstream transit time (t_{up_dif}) and the downstream transit time (t_{down_dif}), In an ideal fluid, these two variations should be almost equal.

- Deviation (φ_i): The absolute difference between the current raw data and the previous filtered output, representing the temporal stability of the measurement.

Based on these parameters, the Dynamic Reliability Coefficient (or raw weight η_i) of the individual measurement sample is calculated as:

$$\eta_i = \frac{\lambda_m - \lambda_i}{\lambda_m} \cdot \frac{\varphi_m - \varphi_i}{\varphi_m} \tag{5.5}$$

where λ_m and φ_m are the maximum allowable threshold values for consistency and deviation, respectively. If a measurement is affected by strong interference or dynamic anomalies, its error parameters will approach the maximum limits, and its assigned weight η_i will tend to zero, effectively neutralizing its impact on the final average.

Considering a window of N samples, the raw weights are then normalized to ensure their sum equals 1:

$$W_i = \frac{\eta_i}{\sum_{i=0}^{N-1} \eta_i} \tag{5.6}$$

Finally, the filtered output y_n is calculated as the normalized weighted average of the acquisitions:

$$y_n = \sum_{i=0}^{N-1} W_i \cdot x_{n-i} \tag{5.7}$$

This methodology guarantees excellent zero-point stability while maintaining high responsiveness to actual flow changes, as reliable measurements are heavily weighted, whereas physically incoherent data are dynamically excluded [25].

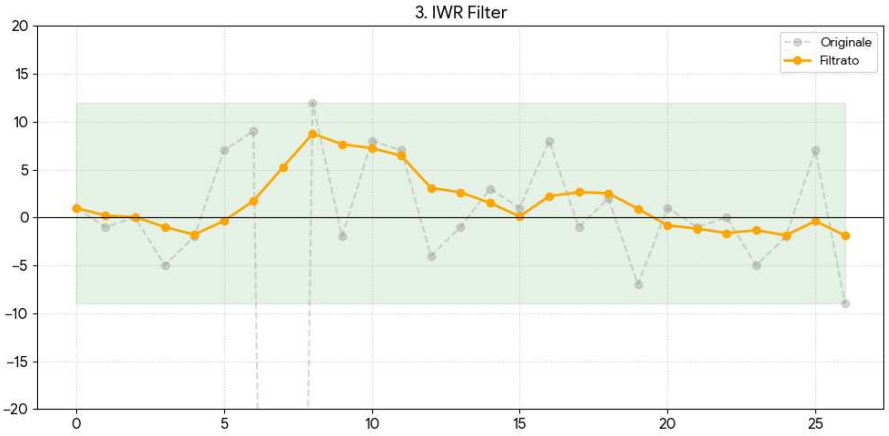

Figure 5.5: Improved Weighted Recursive Filter in rest condition

5.2.3 Kalman Filtering Algorithm

The Kalman Filter is an optimal recursive estimator that minimizes the mean square error, making it highly effective at suppressing white noise and improving the sys-

Figure 5.6: Improved Weighted Recursive Filter in flux variation condition

tem's zero-point stability [8].
The flow dynamics measurement system can be represented by the following discrete state-space models:

$$x(k) = Ax(k-1) + Bu(k-1) + w(k-1) \tag{5.8}$$

$$z(k) = Cx(k) + v(k) \tag{5.9}$$

where $x(k)$ is the state variable; the state transition matrix A linearly maps the state from step $k-1$ to step k at the current time; matrix B represents the gain of the optional control input $u(k-1)$, however, due to the fact that the state of the flow process has no control inputs, u is assumed equal to zero; the random signal $w(k-1)$ is the process excitation noise with covariance matrix Q (which represents the model uncertainty); the observation matrix C maps the true state space into the observed space, and $v(k)$ represents the measurement noise with covariance matrix R (which represents the sensor measurement error) [8].

Statistically, the noise vectors $w(k)$ and $v(k)$ are modeled as mutually independent, zero-mean, white-noise processes [19]. Their covariance matrices, Q and R, are formally defined by their expected values (E):

$$E[w(k)w(j)^T] = \begin{cases} Q(k) & fork = j \\ 0 & fork \neq j \end{cases} \tag{5.10}$$

$$E[v(k)v(j)^T] = \begin{cases} R(k) & fork = j \\ 0 & fork \neq j \end{cases} \tag{5.11}$$

In a practical implementation, Q and R act as the fundamental tuning parameters of the Kalman Filter. The matrix R is related to the intrinsic noise of the ultrasonic hardware (e.g. the TDC quantization error), while Q represents the confidence in the mathematical model tracking the flow dynamics. Tuning these parameters (e.g. setting $R = 10^{-5}$ and $Q = 10^{-1}$) allows the designer to balance the filter's behavior: Un rapporto Q/R più elevato implica che il sistema segua più da vicino le misurazioni del sensore, aumentando la reattività, mentre un rapporto più basso si basa maggiormente sul modello, migliorando l'effetto di levigatura contro forti jitter del segnale [8].

The Kalman filtering algorithm can be divided into two main phases:

1. **Prediction:** in which the filter evolves the state based exclusively on the mathematical model of the system.

   - The prior state estimate $\hat{x}(k|k-1)$ is calculated by applying the transition matrix A to the previous state, representing the theoretical trajectory of the system:
   $$\hat{x}(k|k-1) = A(k-1) \cdot \hat{x}(k-1|k-1) \tag{5.12}$$

   - The prior error covariance estimate is calculated as:
   $$P(k|k-1) = A(k-1)P(k-1|k-1)A^T(k-1) + Q(k-1) \tag{5.13}$$

2. **Correction:** in which the filter compares the prediction with the actual measurement to update and adapt the system state.

   - The calculation of the Kalman Gain $K(k)$ matrix, which defines the optimal weighting between the model prediction and the new measurement:
   $$K(k) = P(k|k-1)C^T(k) \left[C(k)P(k|k-1)C^T(k) + R(k)\right]^{-1} \tag{5.14}$$

   - The posterior state estimate is obtained by the sum of the prediction and a corrective parameter based on the measurement residual:
   $$\hat{x}(k|k) = \hat{x}(k|k-1) + K(k)[z(k) - C(k)\hat{x}(k|k-1)] \tag{5.15}$$

   - The posterior error covariance is updated using the identity matrix I:
   $$P(k|k) = [I - K(k)C(k)]P(k|k-1) \tag{5.16}$$

where $P(k|k)$ is the covariance of the estimation error, $K(k)$ is the filter gain, $R(k)$ is the observation noise covariance matrix, and $Q(k-1)$ is the process excitation noise covariance matrix [8, 19].

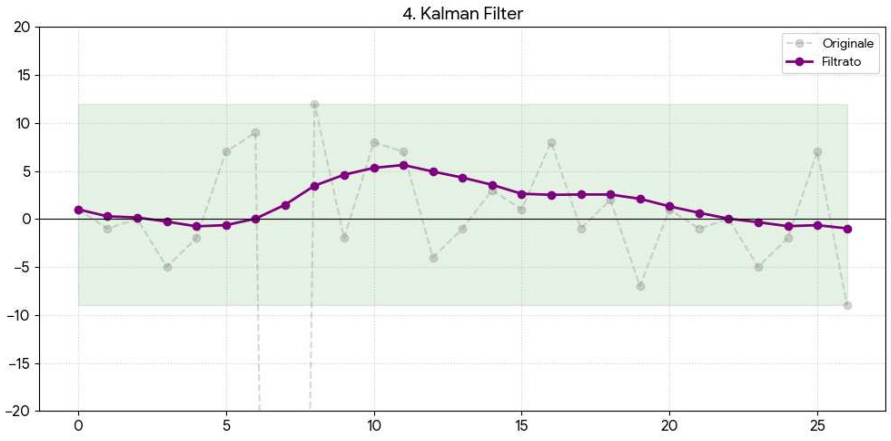

Figure 5.7: Kalman Filter in rest condition

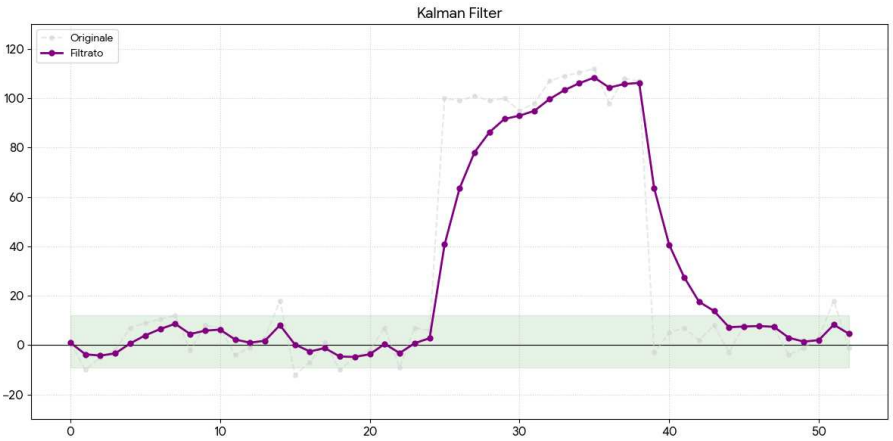

Figure 5.8: Kalman Filter in flux variation condition

6. State of the art

In this chapter, the state of the art of today's sensors will be shown by executing a comparative study between three different ultrasonic gas flow sensors currently used in modern smart meters.
Before entering in the test configuration, it is important to note that the three analyzed devices (which, for confidentiality and corporate compliance reasons, will be identified anonymously as Sensor A, Sensor B, and Sensor C) cannot be purchased on the consumer market by private individuals, but are designed and calibrated exclusively for Business-to-Business (B2B) dynamics. The absence of availability on the consumer market is justified by the critical nature of the application, given that gas measurement for billing purposes requires the entire gas meter system to be certified according to the European MID (Measuring Instruments Directive) regulations. Consequently, these sensors are supplied only to industrial partners capable of guaranteeing an integration that respects the manufacturing tolerances and strict metrological requirements.

6.1 Commercial Solutions

The three examined sensors belong to the G4 category, this means that they are certified to operate up to $6m^3/h$, all three are in a Single-path configuration, and all three have two internal separation plates to obtain laminar flow.

6.1.1 Sensor A

Sensor A uses a V-shaped topology but since the exact geometric details are covered by industrial secrecy, the internal dimensions and topology have been estimated empirically: the acoustic path (L) has an approximate length of $55mm$, with a cross-section (S) of about $286mm^2$ and an incidence angle (θ) of about $60°$. Analyzing the transducer excitation via oscilloscope revealed a frequency of about 500kHz.

The module is designed to measure air and CH4 up to a full scale of $7200l/h$, operating in an environmental temperature range from $-25°$ to $+55°$ and in a pressure range from 1.0 to 3.5kPa. The intrinsic pressure drop of the tube is $\leq 200Pa$ (measured at 25° with a flow of $6000L/h$), while the error on the thermal measurement is $\pm 0.675°$.
As for energy management, in Normal mode it gives a flow measurement every 2 seconds (0.5Hz) with a consumption of about $26.8\mu A$, but if the measured flow is zero for over 30 seconds the device enters in Low-power mode, reducing the absorption to about $10.2\mu A$. For quick bench testing purposes, an Inspection Mode can also be activated, which lowers the sampling period to 125ms.
Regarding the communication interface with the microcontroller, a standard UART at 9600bps is available. The data frame is delivered in HEX format and returns the instantaneous volumetric flow rate value (in L/h), the measured times of flight (T_{up} and T_{down}) and every 30 seconds also a value of temperature (the temperature reading causes an increase in consumption of about $1.4\mu A$). The protocol also handles alignment through handshaking mechanisms (ACK/NAK) to ensure maximum robustness of the serial communication.

6.1.2 Sensor B

Sensor B is characterized by a different mechanical design, estimated as W-shaped, while regarding the mechanical dimensions they have been estimated as: a path length (L) of about 80mm, a cross-section of about $330mm^2$ and an incidence angle (θ) of about 30°, with a transducer frequency of 200kHz.
The module is designed to measure Real Gas and air up to $6000l/h$, therefore it is characterized by a lower full scale compared to sensor A. In this case, the pressure range is higher attesting to a value of 30kPa, while as for the temperature, an operating range from $-25°$ to $+55°$ is guaranteed, therefore equal to sensor A. Regarding the intrinsic pressure drop of the tube, we obtain a lower value compared to sensor A, which stands at $\leq 140Pa$, the same applies to the temperature measurement error which in this case drops to $\pm 0.607°$.
Also in this case we have two operating modes: the Normal mode which returns a measurement output every 2 seconds registering an absorption of $36\mu A$, and the Low-power mode with a consumption of $25\mu A$, therefore in this case we find higher consumptions compared to sensor A. In this case there are two modes to facilitate bench testing, one called Pulsation mode, which returns an output every 2 seconds based on various samplings performed every 125ms, bringing the instantaneous consumption to $197\mu A$, and a second called Detection mode, which returns an output

every 125ms maintaining the same sampling frequency (8Hz), but in this case the consumptions rise to $317\mu A$.
As for the communication interface, also in this case we have a UART protocol at 9600bps, with HEX packets. Through specific commands the system can return not only the flow rate Q (in L/h) and the ToFs T_{up} and T_{down} but also some status flags capable of signaling any anomalies. Information regarding the temperature can be updated once a second or once every 30 seconds based on user requests, but every temperature reading leads to a consumption of about $0.4\mu A$. Also in this case, a handshaking management is provided to guarantee the robustness of the serial communication.

6.1.3 Sensor C

Sensor C, like sensor A, employs a V-shaped topology, with internal dimensions (estimated empirically) of $66mm$ regarding the path length (L), a cross-section (S) of about $264mm^2$ and an incidence angle (θ) of about 45°, with 200kHz transducers. The module is designed to measure Air, CH4 or even natural gas mixtures up to a flow rate of 6000l/h, operating in an environmental temperature range from $-10°$ to $+40°$ (therefore narrower compared to A and B) and in a pressure range up to 20kPa. The accuracy on the thermal measurement is $\pm 0.5°$, while regarding the intrinsic pressure drop of the tube, no data is provided.
This sensor is characterized by two operating modes: the Normal mode generating an output once per second with a consumption of $40\mu A$, and a Low-power mode in which the output is provided once every 2 seconds where the absorption drops to $25\mu A$. In this case the device, in addition to being equipped with the temperature sensor, is also equipped with a pressure sensor, however, the energy consumptions of the two measurements are unrevealed.
Compared to sensors A and B, the UART interface works at a higher baud rate, equal to 38400bps. The logical interface is more advanced compared to A and B, in fact, in this case it is possible to obtain, in addition to the instantaneous flow, also the calculated volume both in the raw state (working conditions) and thermally compensated (base conditions), the calculated speed of sound (with a resolution of 0.1m/s), the gas concentration, and various error flags. It also allows the selective enabling via software of the temperature and pressure compensation. It is important to point out, however, that sensor C is the only one not able to know the reverse flow value, but in case of reverse flow it turns on an error flag and this can be a serious problem in some cases.

6.2 Test Methodology and Metrics

To guarantee a direct and fair comparison of the three commercial sensors, and subsequently also the sensor developed during this thesis project, it was decided to test all the devices using the same certified test bench in order to eliminate systematic uncertainties deriving from variations in baseline conditions, such as fluctuations in temperature or atmospheric pressure.

The testing was structured by defining a duty cycle that simulates the typical operational dynamics of a domestic user (class 1.5, for G4 nominal sizes). The chosen profile includes a setpoint ramp first increasing and then decreasing, articulated in the following volumetric flow rate values: 0, 0.04, 0.6, 1.2, 3.6, 6, 3.6, 1.2, 0.6, 0.04, 0 m^3/h. This profile was designed to cover the entire dynamic range required by the MID directive, including the defined minimum flow rate (Q_{min}) of 0.04 m^3/h, the transitional flow rate (Q_t) of 1.2 m^3/h, and the maximum flow rate (Q_{max}) of 6 m^3/h.

The data collected from the bench will be analyzed and compared based on the following metrics:

- **Offset**: The Offset represents the residual flow value read by the sensor when the actual flow rate on the bench is strictly equal to $0m^3/h$. It is a vital parameter (also known as Zero-flow drift), as it determines the minimum detection threshold; a fluctuating or high offset risks counting "phantom flow" when the user is inactive.

- **Accuracy**: For each setpoint of the cycle, the relative percentage error of the measurement will be calculated by applying the formula provided by the MID directive (formula 1.1):

$$E_{\%} = \frac{V_{mis} - V_{rif}}{V_{rif}} \times 100$$

  in order to verify if the sensors respect the limits of the Maximum Permissible Error (MPE), as defined in chapter 1.1.

- **Range and Turndown**: Going to evaluate the maximum flow value read by the sensor falling within a 1.5% error as defined by the MID directive. From here it is also possible to obtain the Turndown value as the ratio between Q_{max} and Q_{min}, which as defined by the MID directive must be ≥ 150 for a class 1.5 device.

- **Linearity and Sensibility**: The flow values imposed by the bench (X axis) and the values measured by the sensors (Y axis) will be plotted on a scatter graph to draw a regression line. The slope of this line will provide the sensibility of the system for the different flow rates. Linearity will instead be quantified by evaluating the maximum deviation and the variance of the individual acquired points compared to the ideal linear model. High linearity enormously simplifies the work of the software compensation algorithms.

- **Hysteresis**: Going to compare the reading obtained at the setpoints during the increment phase with the reading during the flow decrement phase. This comparison will highlight any "fluid dynamic memory" phenomena or settling delays in the digital filters of the sensors.

- **Repeatability**: To guarantee the economic fairness required by the market, a sensor must provide identical readings under equal conditions. Repeatability will be tested by calculating the statistical dispersion of the measurements, repeating the acquisition for the same setpoint n consecutive times.

6.3 Comparison

Based on the test methodology and the metrics defined in Section 6.2, the empirical data acquired from the test bench have been plotted and analyzed. Figure 6.1 illustrates the transient response of the three commercial sensors (A, B and C) subjected to the standardized duty cycle that covers the dynamic range of the nominal size of G4, from 0 to 6 m^3/h.

Analyzing the offset and the Zero-flow drift (Figure 6.2), Sensor A demonstrates an almost ideal behavior with a residual reading of about 0.001 m^3/h and a negligible standard deviation. Sensor B, on the other hand, shows a negative offset (about -0.013 m^3/h at the beginning), which would require a zero calibration procedure to avoid metrological errors. Sensor C presents an unnatural flattening exactly at 0.0 m^3/h with zero variance at the beginning of the test; this behavior strongly suggests the implementation of a very aggressive software deadband to artificially mask the intrinsic white noise at zero flow.

Regarding the transient response, Sensors A and B strictly track the setpoint ramp over the entire dynamic range. On the contrary, Sensor C reveals serious problems at both ends of the spectrum. At high flow rates, it completely fails to reach the maximum setpoint ($Q_{max} = 6$ m^3/h), erroneously flattening out around

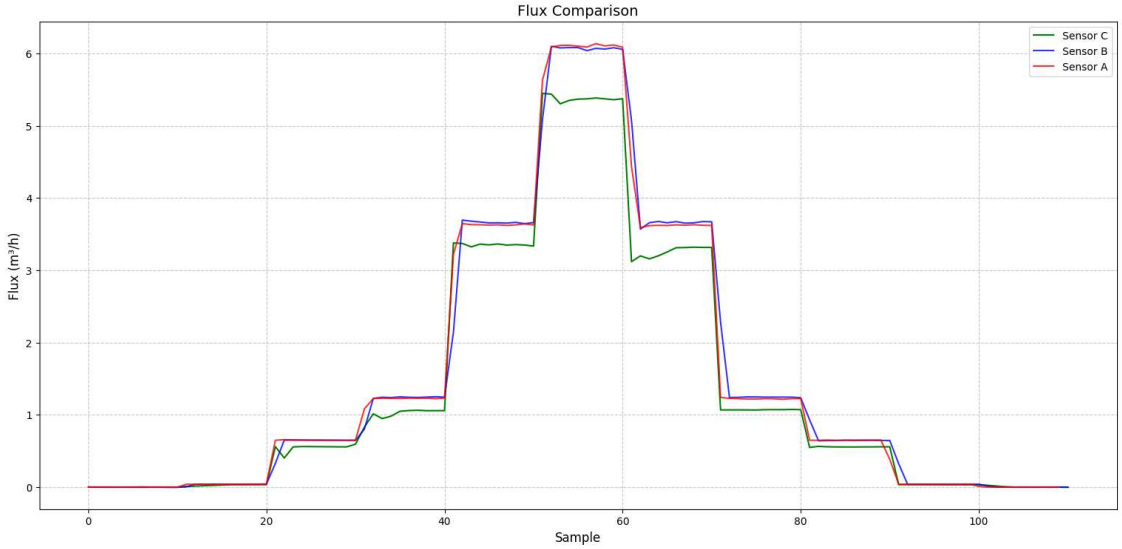

Figure 6.1: Flow comparison of sensors A, B, C in dynamic test

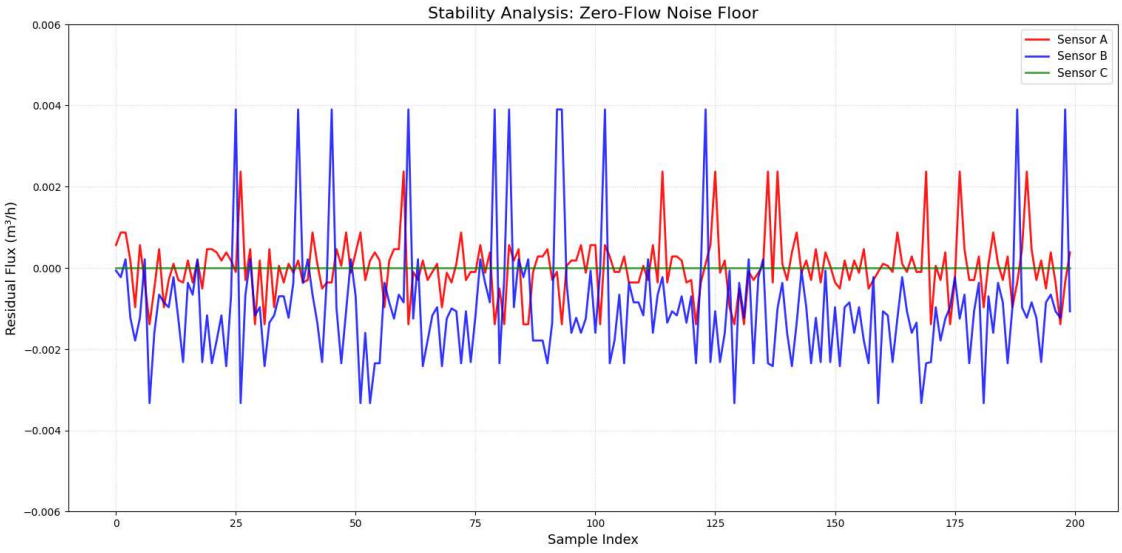

Figure 6.2: Flow comparison of sensors A, B, C at zero flow

5.368 m^3/h. Furthermore, as shown more clearly in Figure 6.3, Sensor C is remarkably slow in detecting small flow variations. This delayed response is symptomatic of a too aggressive post-processing filter, probably adopted to compensate for a poor raw Signal-to-Noise Ratio (SNR). As theoretically analyzed in Chapter 5, standard filtering techniques like the Moving Average introduce a significant time delay proportional to the window size, excessively smoothing the data and compromising the system's ability to track rapid dynamic variations.

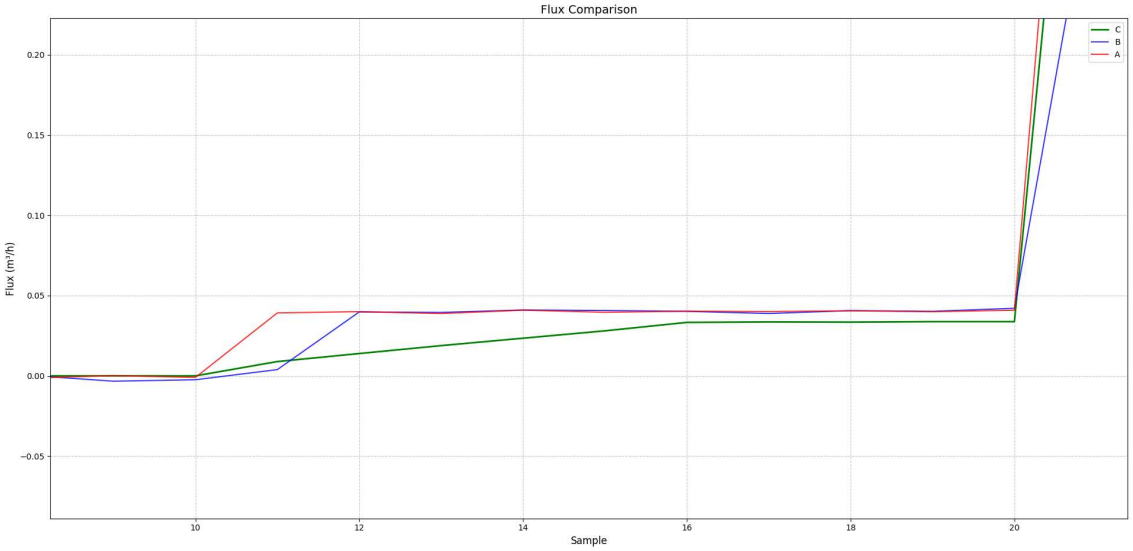

Figure 6.3: Flow comparison of sensors A, B, C in response to small variations

The evaluation of hysteresis and repeatability highlights further discrepancies. Comparing the ascending and descending steps (for example, at the 3.6 m^3/h setpoint), Sensor A maintains a sharp, consistent and square-wave-like profile in both directions, keeping the white noise well contained. This indicates the probable use of an advanced adaptive filtering algorithm, such as the Improved Weighted Recursive or Kalman filters discussed in Chapter 5, capable of dynamically assigning weights to mitigate noise without introducing severe phase delays. In contrast, Sensor C suffers from a massive "fluid dynamic memory" during the descending ramp: its measurement collapses, showing violent and irregular fluctuations with a noise standard deviation that exceeds 0.06 m^3/h. Sensor B also shows a slight increase in ripple during the descent phase compared to the ascent, although not as high as Sensor C.

A quantitative analysis was conducted by evaluating the Percentage Error and the Linearity and Sensitivity parameters, in order to rigorously verify compliance

with the limits imposed by the MID directive for Class 1.5 devices.

Flow Rate (m^3/h)	**A Error (%)**	**B Error (%)**	**C Error (%)**
0.04	-0.75	-0.50	-16.50
0.60	-1.67	2.20	-7.00
1.20	1.67	2.38	-10.83
3.60	1.33	4.11	-4.20
6.00	1.92	3.90	-10.53

Table 6.1: Comparison of the Percentage Error for sensors A, B and C at different flow rate setpoints.

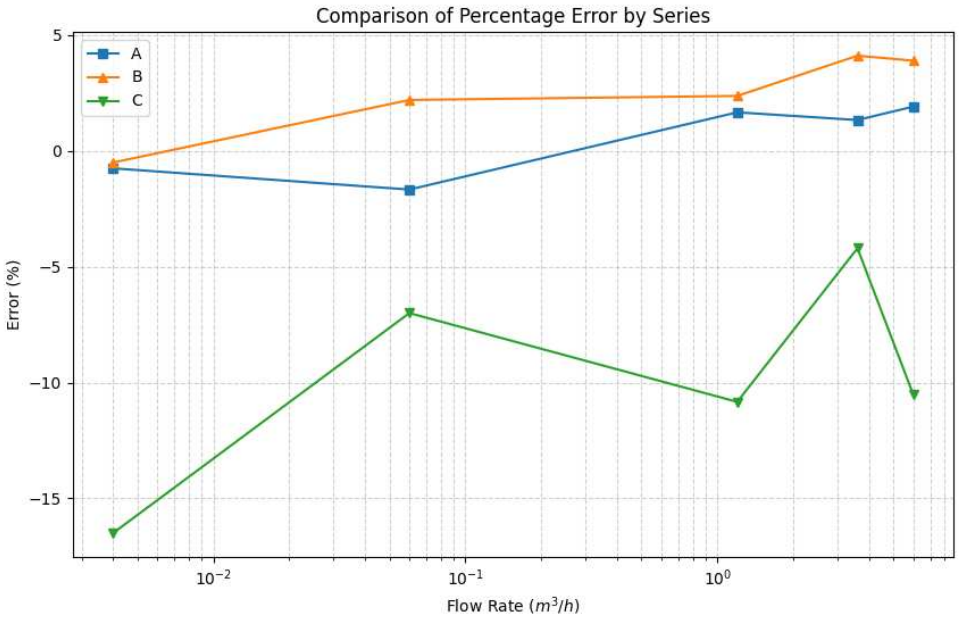

Figure 6.4: Comparison of the Percentage Error per series relative to the volumetric flow rate.

As theoretically defined by Equation 1.1, the accuracy analysis (Percentage Error) allows evaluating compliance with the Maximum Permissible Error (MPE) limits. The regulation tolerates an error of ±3% in the lower zone (between Q_{min} and Q_t) and a more restrictive limit of ±1.5% in the nominal operating zone (between Q_t and Q_{max}). The values calculated on the test bench, summarized in Table 6.1 and plotted in Figure 6.4, confirm the clear performance differences.
Sensor A is clearly the most accurate. At low flow rates the error remains well within the ±3% limit (registering −0.75%). In the nominal zone, the sensor only slightly touches the restrictive ±1.5% limit, with a peak of 1.92% at the maximum flow rate (6 m^3/h). Sensor B shows good accuracy at Q_{min}, but its performance

rapidly degrades as the flow increases, systematically falling outside the MID limits (error of 4.11% at 3.6 m^3/h). Finally, Sensor C critically fails the metrological requirements across the entire spectrum, with an underestimation peak of -16.5% at the minimum flow rate and settling around a very severe -10% at high flow rates.

These accuracy data are directly reflected in the calculation of Range and Turndown (Q_{max}/Q_{min}). The MID directive imposes a turndown ratio ≥ 150 for class 1.5 meters, since the valid turndown is calculated considering exclusively the dynamic range in which the sensor respects the MPE limits, the three devices show very different compliances. Sensor A is the only device to approach the regulatory target: strictly respecting the limit at Q_{min} (0.04 m^3/h) and only marginally exceeding it at Q_{max} (6 m^3/h), it guarantees an effective turndown of 150. Sensor B, despite starting with an accurate measurement at minimum flow, exits the $\pm 1.5\%$ tolerance limits already starting from the transition flow Q_t (1.2 m^3/h with an error of 2.38%), prematurely halting its validatable maximum flow rate and reducing its useful turndown to a value below 30, totally insufficient for homologation. Sensor C, resoundingly failing the MPE requirements right from the minimum flow, does not possess an operating range compliant with the directive, effectively making its turndown ratio null and unclassifiable.
It is important to underline, however, that the MID directive does not evaluate the single isolated sensor, but rather certifies the entire measurement system (the gas meter as a whole). Consequently, Sensors B and C could theoretically fall within the regulatory parameters following a rigorous calibration procedure of the entire system before being placed on the market. It is crucial to consider, however, that the correction of such marked non-linearities and deviations requires the implementation of heavy software calibration curves (such as extensive look-up tables). This compensatory approach inevitably entails a greater computational load for the host microcontroller, translating into an increase in overall energy consumption.

The data extracted via linear regression (Figure 6.5 and Table 6.2) mathematically quantify the described anomalies. Sensor A and Sensor B demonstrate excellent linearity ($R^2 = 0.99999$) with negligible maximum deviations ($\sim 0.014\ m^3/h$). However, Sensor A presents an ideal sensitivity of 1.0183, while Sensor B shows a higher slope (1.0402), translating into a systematic overestimation at high flow rates. Sensor C confirms itself as the least performing solution: the sensitivity drops to 0.9068 and the maximum deviation touches 0.1654 m^3/h (a value ten times higher than the other devices), making the output difficult to compensate via software.

In conclusion, Sensor A represents the reference "gold standard" among the

Sensor	**Sensitivity (Slope)**	**Max Deviation** (m^3/h)	R^2
Sensor A	1.0183	0.0145	0.99999
Sensor B	1.0402	0.0148	0.99999
Sensor C	0.9068	0.1654	0.99829

Table 6.2: Linearity and sensitivity parameters

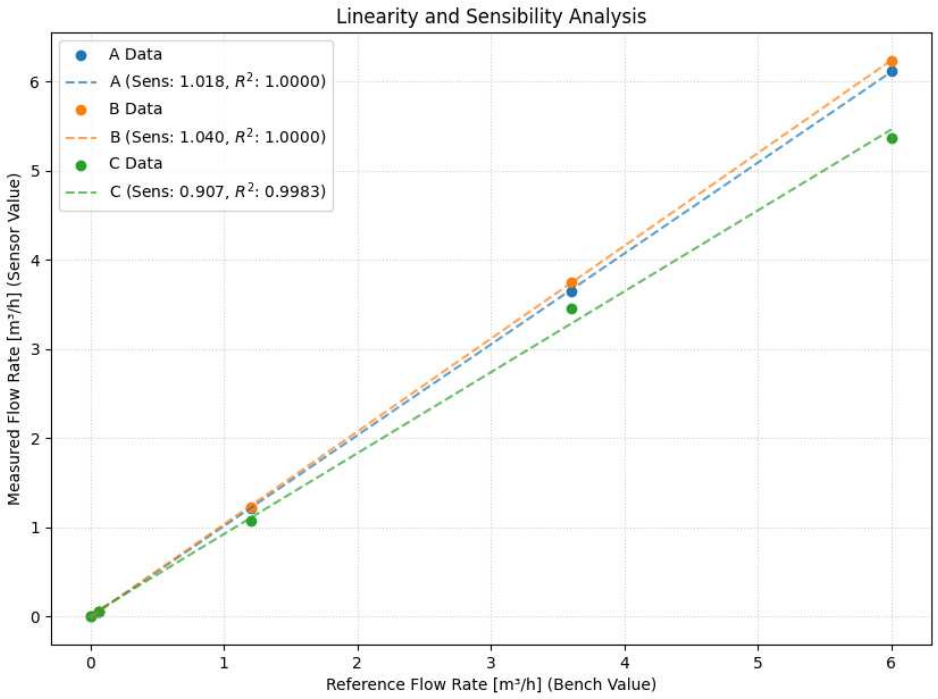

Figure 6.5: Analysis of linearity and sensitivity: comparison between the reference flow rates (bench) and the measured values.

tested state-of-the-art devices, balancing excellent zero stability, high tracking linearity and low noise. These empirical results clearly quantify the problems of the current state of the art and provide an ideal target for the validation of the architecture developed in this thesis.

7. Realization

Following the studies relative to the system and the state of the art, we proceeded towards the realization of our own Gas Flow Sensor.
SISTEMARE

7.1 Architecture Design

The fundamental first step in the realization of the hardware architecture of the ultrasonic flow sensor is the choice of the Time-of-Flight detection technology. As we have seen in chapter 4, the two main methodologies for ToF extraction are the Time-to-Digital Converter (TDC) and the Analog-to-Digital Converter (ADC). Wanting to make a quick recall, the TDC method stands out because it does not digitize and process the entire received acoustic waveform, but simply measures the time interval between the emission of the signal and the instant in which the received echo exceeds a specific amplitude threshold and crosses zero (Zero-Crossing technique). The ADC technology is theoretically capable of providing greater noise immunity, but this approach requires superior computational resources and energy budget. From the tests carried out on the state of the art, discussed in chapter 6, it was noted how an ADC method can fall within the energy consumptions required by the MID directive only provided that algorithms capable of using the lowest possible computational cost are applied and using microcontrollers designed and optimized for this architecture. However, this compromise leads to metrological results slightly inferior compared to those obtainable with the TDC technology, while failing to reduce consumptions in an equal way. For this reason, it was decided to opt for a TDC architecture, in order to obtain the best possible metrological result at the lowest computational and, consequently, energy cost.
In a TDC architecture, the hardware is active only for the time necessary to send and receive the acoustic pulse, furthermore, since there are no continuous data arrays to process, the calculation latency is almost zero and the microcontroller can return almost instantaneously to a deep sleep state. To counterbalance the intrinsic

vulnerability of the TDC to amplitude variations of the acoustic wave in the gas (the Cycle Slip phenomenon), in addition to the integration of a Split Burst algorithm, it was also decided to use a programmable gain amplifier (PGA) that dynamically adapts the amplitude of the input signal.
Finally, the adoption of the TDC technology led to the choice of an operating frequency of the transducers of 500kHz in order to maximize the detection precision. At this frequency, in fact, a steeper Slew Rate is obtained compared to 200kHz, thus optimizing the Zero-Crossing Detection and making the detection extremely more robust against the signal-to-noise ratio (SNR). Furthermore, exploiting the high slope of the signal, the TDC chip is able to compensate for the greater amplitude attenuation suffered by the acoustic wave in the gas, reaching a very high timing stability and a metrological resolution in the order of picoseconds.

The initial validation phases of this architecture were conducted through a developer board equipped with an STM32U5 and an evaluation board equipped with the TDC chip, to then conclude in the design of a custom PCB aimed at integration into the prototype.

7.2 Mechanical Design

Parallel to the development of the electronic architecture, we proceeded with the mechanical design of the measurement pipe. As seen in chapter 3, the choice of topology defines the physical integration between the ultrasonic wave and the fluid, determining the basic metrological performances even before signal processing. Among the available single-path configurations, the V-shaped topology was selected, excluding the Z-shaped and W-shaped architectures. On one hand, the Z configuration exposes the system to high uncertainty at high flow rates, due to the strong lateral deviation of the ultrasonic beam dragged by the gas. On the other hand, the W configuration, while theoretically maximizing the time of flight difference (D_{ToF}), entails an excessive dissipation of acoustic energy at each reflection on the walls. The V-shaped structure offers two advantages: the double transverse crossing of the pipe section provides an intrinsic self-compensation of the noise caused by turbulences and by the asymmetry of the gas velocity profile, and that, with the same pipe diameter, the acoustic path length (L) doubles compared to the Z configuration, increasing the interaction with the fluid and improving the Signal-to-Noise Ratio (SNR) at minimum flow rates (Q_{min}).

The sizing of the measurement pipe was guided by the analysis of the empirical results obtained from the experimental tests by José R. García Oya shown in paragraph 3.3 and by the experimental data obtained during the state of the art shown in chapter 6. Based on this, the sensor was engineered with a cross-section S of $310mm^2$, a path length L of $53.2mm$ and an incidence angle θ of $65°$. Each of these parameters responds to a specific physical need:

- Since the gas velocity is $v = Q/S$, excessively reducing the section would increase the fluid velocity, improving the resolution at low flows, but it would lead to a parasitic deviation of the ultrasonic beam at high flow rates. A section of $310mm^2$ guarantees that the lateral deviation remains contained, avoiding the asymmetrical attenuation of the signal between upstream and downstream measurements.

- The path length of $53.2mm$ represents the ideal compromise to guarantee a sufficiently large D_{ToF} (proportional to L) without incurring an excessive attenuation of the acoustic signal. Considering the choice to operate at 500kHz (which suffers a greater attenuation compared to 200kHz), an $L = 53.2mm$ allows the wave to arrive at the receiver with enough energy so as not to require either a high transmission voltage or a high input gain that would lead to higher energy consumptions.

- It was decided to fix the incidence angle θ at $65°$ because it is the value that returned absolutely the best results in the tests performed by José R. García Oya in [14], the models characterized by such an angle demonstrated the lowest and most stable uncertainty curve along the entire range of flow rates. Although smaller angles would mathematically generate a greater D_{ToF} (since $D_{ToF} \propto \cos(\theta)$), tests demonstrated that a more acute angle makes the ultrasonic beam excessively vulnerable to lateral deviation caused by gas dragging at high speeds, decreasing the accuracy of the system.

Finally, the pipe was equipped with three separation plates, as tests confirmed their crucial role in the laminar stabilization of the flow.

In order to quickly validate these geometric choices without incurring the high costs and long realization times of an industrial mold, the physical structure of the sensor was manufactured using rapid prototyping techniques in 3D printing. However, from a fluid dynamic point of view, it is fundamental to consider the impact of this production process, which, unlike an injection molding process that guarantees perfectly smooth walls, generates the object by successive layer deposition,

introducing an intrinsic surface roughness along the pipe walls, risking triggering parasitic micro-turbulences that hinder the hold of a purely laminar flow.

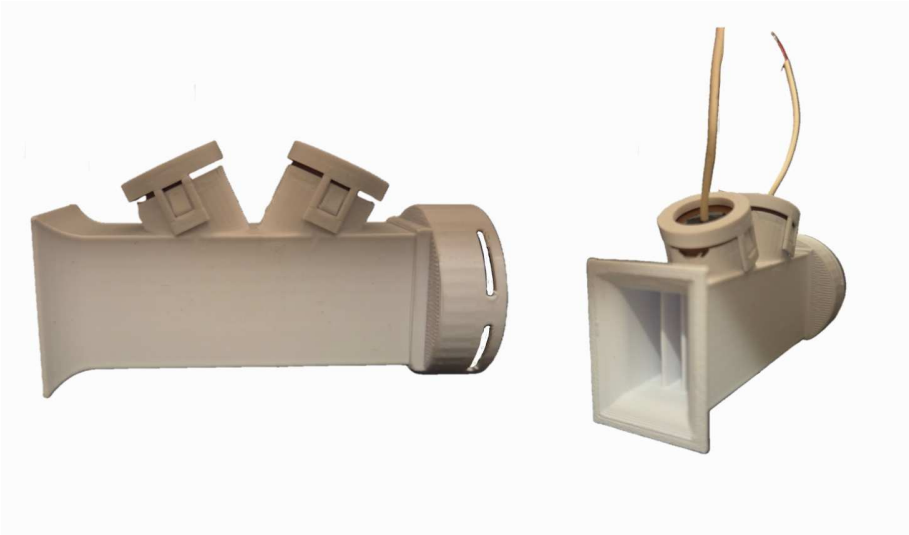

Figure 7.1: 3D Printed Sensor

7.3 Transducer Matching

Transducer Matching represents a crucial phase to guarantee maximum energy efficiency and minimum measurement uncertainty. In this work, in order to optimize the acoustic response of the device and guarantee an accurate echo detection, the joint adoption of two strategies became necessary: the insertion of a damping resistor in series with the transducer and the implementation of a frequency sweeping algorithm for dynamic calibration.

A piezoelectric transducer can be electrically seen as a resonant RLC system, the fundamental parameter that describes its frequency selectivity and the decay time of the oscillations is the Quality Factor (Q). In TDC-based ToF applications, a high Q represents a critical limit, because when the front-end sends the excitation burst, the piezoelectric crystal stores elastic energy which, at the end of the burst, causes an oscillation due to mechanical and electrical inertia. This phenomenon, known as ringing, protracts over time, lengthening the overall duration of the emitted ultrasonic pulse. To mitigate this effect, a series resistor was inserted into the excitation circuit in order to increase the energy dissipation by Joule effect in the driving branch. In a series RLC resonant circuit, the relationship linking the quality

factor to the resistance is described by the following equation:

$$Q = \frac{1}{R}\sqrt{\frac{L}{C}} \tag{7.1}$$

By increasing R, the Q factor decreases proportionally, although this entails a reduction in efficiency by reducing the amplitude of the transmitted peak, the result is an extremely short and defined acoustic pulse [7]. Furthermore, the addition of this resistor also has a hardware protection task by going to limit the peak currents and protecting the output stages of the TDC chip. Finally, it is important to specify that during the echo reception phase, the internal architecture of the TDC switches the excitation pin to ground, consequently, the damping resistor will find itself connected in parallel to the receiving transducer, acting as passive damping on the input wave, absorbing the parasitic oscillations of the receiving crystal and improving the Zero-Crossing of the signal.

By lowering the Q factor, the impedance and frequency response curve of the transducer flatten out, the bandwidth increases, but the resonance peak becomes less pronounced. Added to this dynamic are the non-idealities of the transducers, for a component with a nominal frequency of 500kHz and a typical tolerance of $\pm 5\%$, the actual resonance frequency can vary between 475kHz and 525kHz, not counting the further drifts caused by temperature (a meter must guarantee a temperature range going from $-25°$ to $+55°$) or by the acoustic coupling with the pipe. In a system damped by the resistor, blindly exciting the piezoelectric at a nominal 500kHz risks making the system work outside the efficiency peak, for this reason, a frequency sweeping algorithm was implemented. During the calibration phase, the microcontroller commands the TDC to generate a series of excitation bursts by progressively varying the carrier frequency within a predefined range (for example $\pm 10\%$), and, for each tested frequency, the system acquires the peak-to-peak amplitude of the return echo and identifies the frequency that produces the echo with the maximum amplitude and the highest SNR. However, due to the intrinsic quantization constraints of the chip's hardware divider, the frequency sweeping algorithm does not perform a linear analog scan, but a set of discrete frequencies derived from the 8MHz clock; around 500kHz, the physically generable values are: 444.4kHz, 470.5kHz, 500kHz, 533.33kHz and 571.4kHz.

Precisely in this scenario, the flattening of the resonance curve induced by the damping resistor (i.e. the increase in the transducer's bandwidth) becomes a systemic advantage because the real mechanical resonance frequency of the crystal will almost certainly fall at an intermediate point between two of these discrete hardware

steps, a high-Q transducer (characterized by a very narrow resonance curve) would suffer a dramatic collapse in efficiency. In contrast, a low-Q system presents a much more homogeneous and flat amplitude response around the natural frequency. This characteristic physically compensates for the chip's inability to generate the exact frequency, minimizing the mismatch error induced by the signal quantization.

7.4 TDC settings

As analyzed in chapter 4, one of the main challenges in ultrasonic measurements is the resolution of the Cycle Slip. To mitigate this problem, during the design phase it was chosen to discard the classic uniform burst excitation in favor of the advanced Split Burst mode. In this mode, the sending of the wave packet is divided into two parts: a first part of 10 pulses, followed by a phase jump, and concluded by another train of 8 pulses. Physically, while the mechanical oscillation of the transducer jumps from the initial to the final one, the imposed phase jump generates a marked deviation in the acoustic wave periods, this anomaly acts as a temporal marker within the ultrasonic echo, making the recognition of the wave's position certain regardless of its absolute amplitude.
The use of this technique allowed not having to apply the Dynamic Threshold algorithm (seen in paragraph 4.1.1), allowing a saving of computational and, consequently, energy resources. However, although the Split Burst allows the temporal marker to be uniquely identified, the reception architecture guarantees maximum reliability if the amplitude of the received acoustic wave settles in an optimal range (ideally around $300mV$). Since the wave amplitude is variable as a function of the flow, gas composition, and temperature, a software-variable Programmable Gain Amplifier was applied in order to keep the echo amplitude within the optimal range so as to guarantee the stable switching of the comparator.

The front-end is capable of sending excitation pulses at various voltage levels thanks to a system with a series of capacitors called a Charge Pump. Having chosen an ultrasonic beam length (L) of less than $55mm$, together with the use of the PGA, allowed maintaining a hardware firing voltage at the minimum settable, thus managing to limit the peak consumption of the system. However, the emission of these high-frequency pulses can generate ground returns and electromagnetic interference. To avoid false triggers caused by this crosstalk, a total hardware blanking of the comparator (Noise Mask Windows) was implemented for the first $30\mu s$ fol-

lowing the firing, making the system physically deaf during the energy discharge. In addition to this initial block, a fixed delay window (Release Delay Windows) mathematically calibrated on the pipe geometry was set, in this way, the comparator input is unblocked to actively analyze the signal only upon expiration of this time, i.e. in the narrow theoretical surroundings of the arrival of the first useful acoustic wave.

Once the delay window expires, the comparator waits for the ultrasonic wave to exceed a specific threshold, called the First Hit Level (FHL). Once this threshold is exceeded, the system was configured to acquire the zero-crossing times of all consecutive waves starting from the second useful wave, dividing the ToFs into three main blocks: the wave times before the phase anomaly, the wave times during the anomaly, and the final wave times to confirm the trend. The processing of all zero-crossing times allows white noise and quantization error to be drastically reduced, providing an extremely solid temporal data point. Furthermore, the firmware analyzes the period differences between the three blocks, locates the phase jump and, if it detects a Cycle Slip, mathematically realigns the measurement by subtracting or adding a nominal wave period to the ToF. This design choice guarantees exceptional metrological stability and relieves the microcontroller of the computational burden required by the Dynamic Threshold, containing the energy consumptions of the system.

Confirming and validating the architectural choices just described, the acquisitions captured via oscilloscope allow physically visualizing the behavior of the analog front-end during the measurement phase.
Observing the overview of the acquisition cycle (as shown in Figure 7.2), it is possible to appreciate the entire measurement time window. The yellow trace represents the transmitted signal (excitation burst) sent to the transducer, while the green trace represents the signal received by the other transducer.

Focusing attention on the detail of the transmission pulse (Figure 7.3), the structure and effectiveness of the Split Burst technique emerge clearly on the yellow trace. The acquisition highlights with precision the initial pulse train, the sharp phase jump imposed by the TDC chip, and the subsequent final pulse train. This intentional anomaly acts precisely as a temporal marker, marking the acoustic wave and guaranteeing the system a unique reference to prevent the Cycle Slip phenomenon. At the end of the excitation burst, the mechanical and electrical damping transient (ringing) is also visible, the duration of which is properly mitigated by the use of the damping resistor inserted in the driving branch.

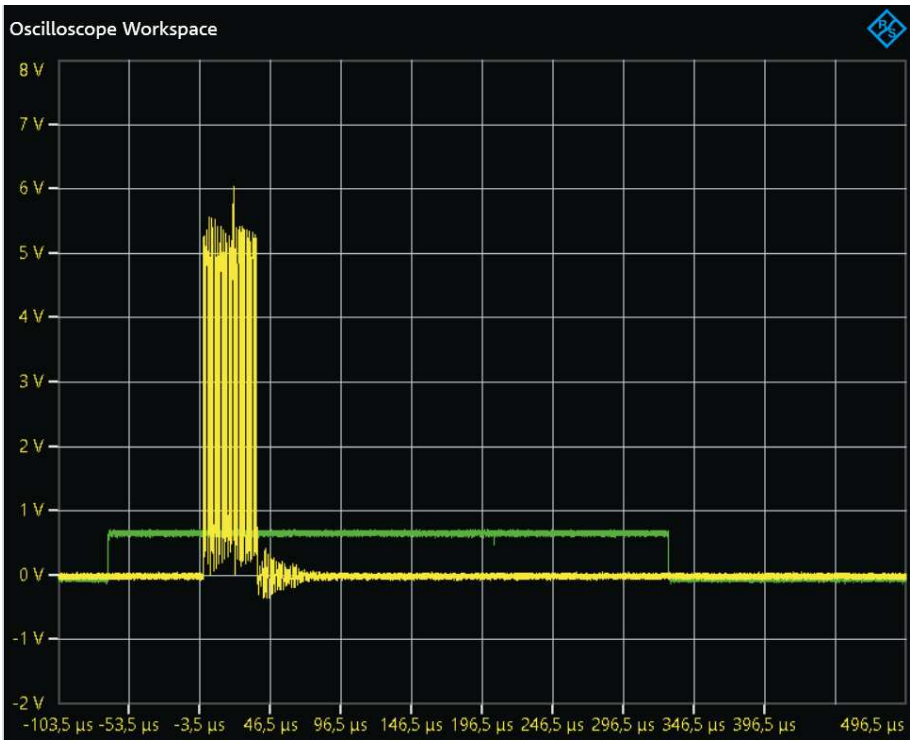

Figure 7.2: Overview of the acquisition cycle: transmission (yellow) and reception (green).

Figure 7.3: Detail of the transmission signal: evidence of the Split Burst technique and the phase jump.

Shifting the analysis to the return echo (Figure 7.4), the arrival of the acoustic wave on the receiving transducer is observed on the green trace. The amplitude of this signal, although attenuated by passing through the gaseous medium, is kept within an optimal range thanks to the intervention of the Programmable Gain Amplifier (PGA). This signal conditioning is fundamental to guarantee stable switching of the comparator once the First Hit Level (FHL) is exceeded, allowing the correct acquisition of the Zero-Crossing instants from which the firmware will locate the phase jump.

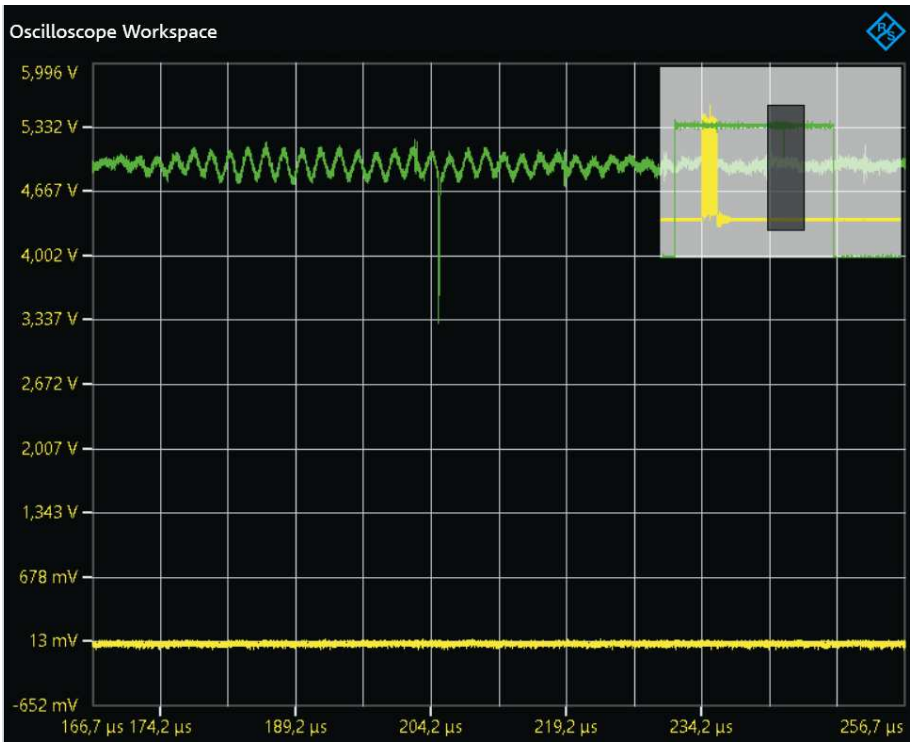

Figure 7.4: Detail of the return echo and the signal conditioned by the PGA.

Finally, to expand the analysis, Figure 7.5 illustrates the macroscopic behavior of the system over a wider time scale. As discussed in the operating principle (Chapter 2), flow measurement via the DToF method requires the alternating acquisition of flight times in both directions, upstream and downstream. On the oscilloscope trace, it is clearly observed how the system inverts the roles of the two transducers between one measurement cycle and the next. In the first event (on the left), highlighted by a predominant peak on the green trace, one transducer acts as an emitter, in the subsequent burst, the peak appears specularly on the yellow trace, confirming that the transmission and reception roles have perfectly inverted. This mechanism is managed entirely via hardware by the TDC chip, guaranteeing the symmetrical acquisition necessary for the differential calculation.

Figure 7.5: Macroscopic behavior: alternation of the transducers in the transmission and reception phases in the upstream and downstream cycles.

7.5 Filtering

The raw data extracted from the TDC chip, although characterized by a very high temporal resolution, are intrinsically affected by environmental and instrumental disturbances. As theoretically analyzed in Chapter 5, the measurement environment inside a gas pipe is subject to turbulences, micro-bubbles of impurities, and thermal noise from the electronics. Furthermore, the stringent MID regulations and the constraint of operating with an extremely reduced energy budget dictate that the host microcontroller (STM32) performs filtering operations in the shortest possible time, in order to rapidly return to a deep sleep state.

For these reasons, the software architecture of the sensor was structured by implementing a sequential processing pipeline divided into two main stages: a first Pre-Filtering stage, dedicated to the rejection of impulsive disturbances (outliers) on individual measurement bursts, and a second Post-Filtering stage, aimed at reducing white noise and smoothly tracking fluid dynamics.

7.5.1 Pre-Filtering

During a single measurement cycle (composed of one upstream and one downstream firing), the TDC does not return a single temporal data point, but acquires an array of n samples (typically $n = 10$) corresponding to the Zero-Crossing instants of the waves within the temporal marker (Split Burst). Before extracting the average time of flight, it is fundamental to purge this array of any anomalous measurements due to false hardware triggers or isolated Cycle Slip phenomena not natively compensated by the front-end.

The first operation performed by the firmware is sorting the array of samples. Among the various available algorithms, it was chosen to implement the Shell Sort. This engineering choice derives from the need to minimize CPU clock cycles: since the array size is very small ($n \leq 10$), the Shell Sort offers an ideal compromise between code simplicity and execution speed, avoiding the memory overhead required by more complex algorithms like Quick Sort or Merge Sort.

Once the array is sorted, the data is passed to a Peters filter for statistical validation. The firmware calculates the arithmetic mean $\bar{x}$ and the standard deviation s of the samples (as described by Equation 5.1). Subsequently, it applies a stringent tolerance window equal to $\pm 2s$ and each single acquired ToF is tested against this band: if the value deviates from the mean by more than two standard deviations, it is classified as impulsive noise.

From the firmware implementation point of view, if a data point is flagged as an outlier, brutally discarding it would reduce the number of useful samples, altering the statistical weight of the measurement. Therefore, the software replaces it by recalculating it through a linear interpolation based on adjacent valid samples. The application of the Peters filter allows identifying and effectively neutralizing anomalous peaks that would otherwise massively shift the calculation of the Δt.

7.5.2 Post-Filtering

The data output from the Pre-Filtering block represents the validated Time-of-Flight for a single sampling instant. However, the analysis over a wider time window reveals the presence of white noise (quantitative jitter). As evident from the acquisition of the raw flow (Figure 7.6), the management of this stochastic component presents a twofold challenge linked to the gas dynamics.

On one hand, under zero-flow conditions and at minimum flow rates (Q_{min}), the noise is particularly critical in percentage terms: since the difference in time of flight ($DToF$) is in the order of a few picoseconds, even a minimal fluctuation

risks overwhelming the useful signal, destroying the accuracy of the measurement. On the other hand, as the flow rate increases up to the maximum limit (Q_{max}), the absolute amplitude of the dispersion drastically increases. As clearly visible in the figure, the thickness of the raw trace widens considerably at the higher flow steps (e.g. at $6m^3/h$). This phenomenon is attributable to the strong fluid dynamic turbulences triggered by the high speed and to the attenuation of the useful signal due to the lateral deviation of the ultrasonic beam.

To mitigate this stochastic component of the signal along the entire operating range, without losing reactivity to transients, three different algorithmic approaches were coded and validated on the microcontroller.

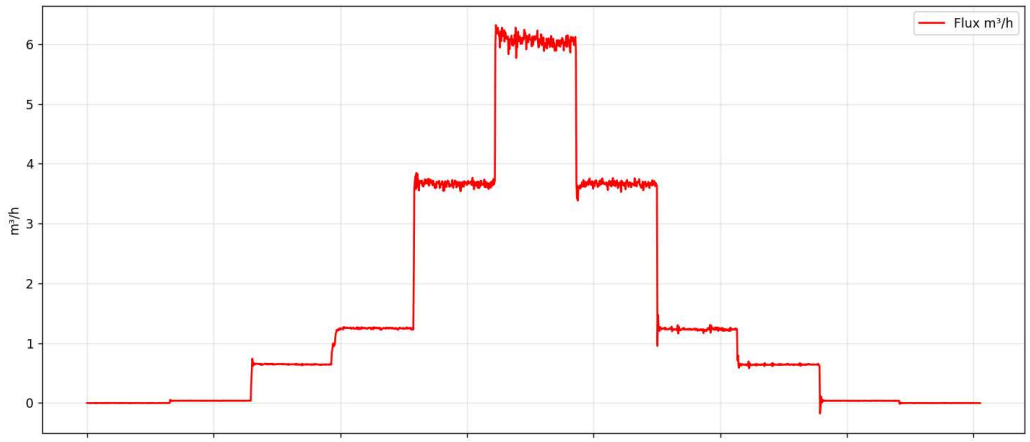

Figure 7.6: Time course of raw data at various flow rate steps.

The first tested algorithm was the Moving Average Filter. Its firmware implementation is trivial and requires very few resources: it uses a circular buffer to keep the last N samples in memory and returns their arithmetic mean. In stationary rest conditions, the filter significantly reduces the standard deviation of the noise. However, the analysis of the transient response (Figure 7.8) highlighted its structural limit: to obtain high noise suppression, a large buffer is necessary, which inevitably introduces a severe temporal phase delay (lag). During sudden flow variations, the filter "drags" the measurement, making the sensor extremely slow in detecting flow gradients, an unacceptable defect for the restrictive MID specifications.

To overcome the phase delay limit, we proceeded to implement the Improved Weighted Recursive (IWR) Filter. At the firmware level, complexity increases since the microcontroller must iteratively calculate the Consistency (λ_i) and Deviation (φ_i) parameters to evaluate the dynamic reliability of the single data point. When the system detects a coherent variation between upstream and downstream (an unequivocal sign of a real fluid dynamic change), it assigns a very high weight W_i to the new sample. Conversely, if the variation is asymmetrical or incoherent, it

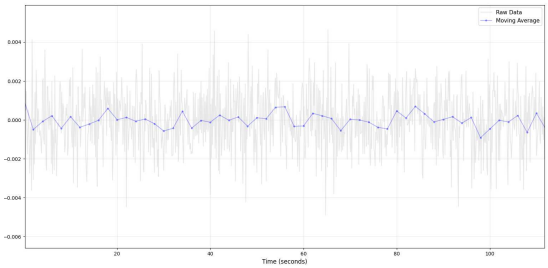

Figure 7.7: Moving Average in stationary conditions (*zero-flow*).

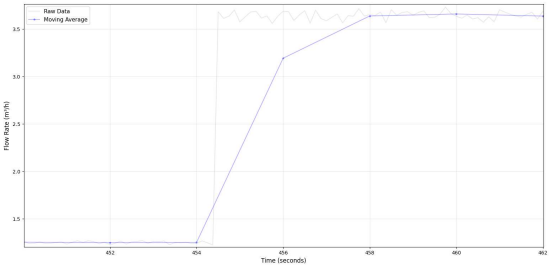

Figure 7.8: Moving Average during a flow variation

is downgraded to noise. The tests performed (Figures 7.9 and 7.10) confirm the drastic improvement: in stable conditions the filter behaves almost similarly to the moving average, while in the presence of a flow step, the algorithm reacts almost instantaneously, faithfully following the real flow rate setpoint.

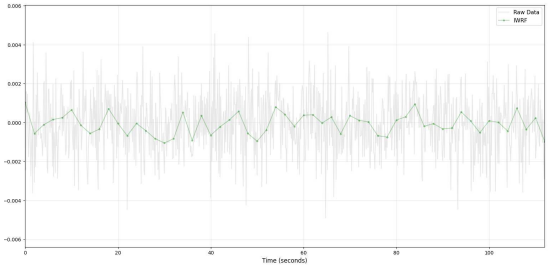

Figure 7.9: IWR Filter in stationary conditions

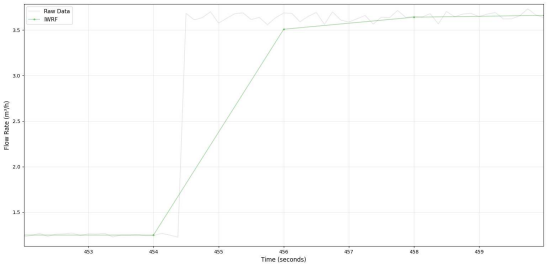

Figure 7.10: IWR Filter during a flow variation

Finally, to seek absolute optimization of metrological stability, a one-dimensional Kalman Filter was coded. This filter, although computationally more burdensome for the STM32 CPU due to the continuous matrix operations (Prediction and Correction), offers the advantage of estimating the optimal state by minimizing the mean squared error. The filter tuning was performed by acting on the covariance matrices Q (process noise) and R (TDC measurement noise). By setting a Q/R ratio experimentally calibrated on the tolerances of the ultrasonic chip, excellent behavior was obtained in zero-flow, paying however a greater delay during transients (Figures 7.11 and 7.12). The Kalman filter ignores strong transient jitters by relying on the internal mathematical model, but rapidly updates the gain $K(k)$ as soon as the readings confirm a coherent shift in the fluid state.

At the conclusion of the analysis on the individual algorithms, a test campaign was conducted to simultaneously compare the performance of the filters over the entire dynamic range. While the hardware front-end and the Pre-Filtering block acquire and validate raw data with a sampling period of 125 ms (equal to a frequency

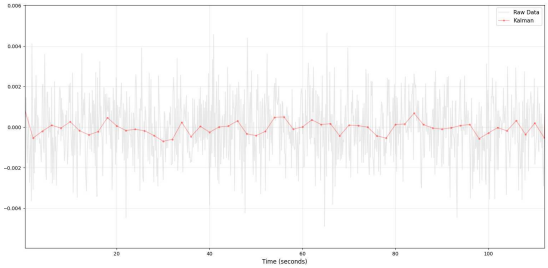

Figure 7.11: Kalman Filter in stationary conditions

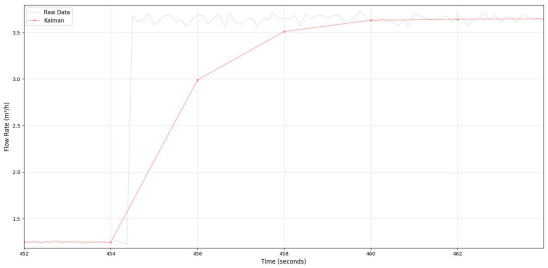

Figure 7.12: Kalman Filter in flow variation

of 8 Hz), the Post-Filtering block aggregates these samples to return the final fluid dynamic data point every 2 seconds.

This engineering choice responds to a twofold need. On one hand, it provides the algorithms with a sufficiently populated statistical window (exactly 16 raw samples for each useful output) to drastically reduce the variance of stochastic noise. On the other hand, it aligns the prototype with the operational specifications of the commercial meters analyzed in Chapter 6 (which natively operate at 0.5 Hz), allowing the duty cycle of the microcontroller to be optimized and energy consumptions to be reduced.

To critically evaluate the behavior of the three standard logics (Moving Average, IWRF and Kalman), the related outputs decimated at 2 seconds were compared in the two operational scenarios of zero-flow and rapid flow rate variation (Figure 7.13, Figure 7.14).

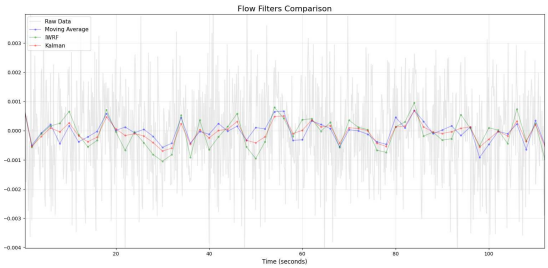

Figure 7.13: Comparison of the 3 standard filters in stationary zero-flow condition

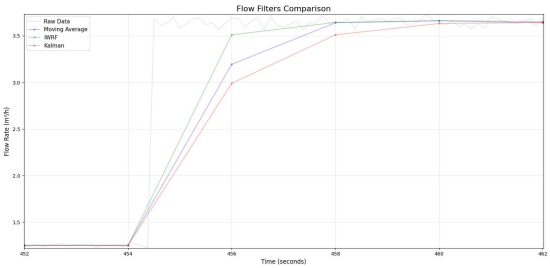

Figure 7.14: Comparison of the 3 standard filters in response to a rapid flow transient

Analyzing the detailed graphs, the characteristic behavior of the various algorithms clearly emerges. The Moving Average confirms its structural limit: while partially reducing noise, it introduces a phase delay in the steepest flow gradients, "dragging" the measurement. The IWRF, instead, is extremely reactive in promptly locking onto the new step setpoints, but pays a slight price in terms of signal flatness in stationary conditions. Conversely, the Kalman filter proves unsurpassable in

flattening white noise in constant flow stretches (guaranteeing exceptional zero stability); however, its intrinsic predictive inertia generates a "rounding" of the signal during sudden flow variations, slowing down the sensor's response to the step.

To obtain absolute maximization of metrological stability, overcoming the intrinsic limits of the individual methods just highlighted, the final architecture of the firmware was ultimately equipped with a dynamic solution that merges the IWRF's response logic to variations with the one-dimensional statistical robustness of the Kalman. Under stable conditions, the system entrusts the filtering to the Kalman algorithm, which zeros out the stochastic jitter. As soon as a coherent and sudden flow rate variation is detected, the hybrid logic bypasses the inertia of the Kalman's predictive model, allowing the measurement to "lock onto" the new setpoint with the immediate promptness of the IWRF. The results of this hybrid filter are visible in Figure 7.15, and in Figure 7.16.

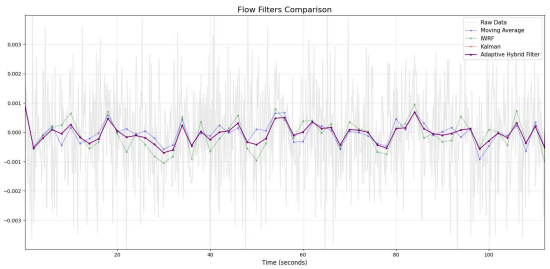

Figure 7.15: Comparison of the filters in stationary zero-flow condition

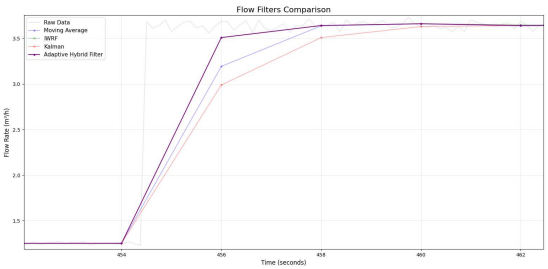

Figure 7.16: Comparison of the filters in response to a rapid flow transient

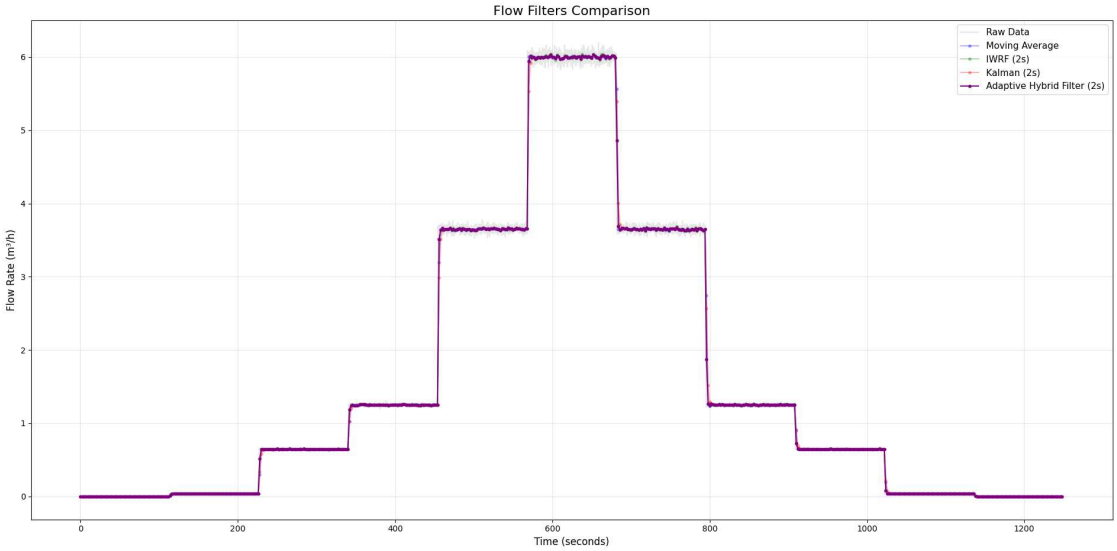

Figure 7.17: Comparison over the entire flow profile

As evident from analyzing the entire dynamic range (Figure 7.17), the final result

of the Adaptive Hybrid Filter (purple trace) is an almost ideal stepped profile.

7.6 Temperature and Pressure Compensation

As analyzed in chapter 2.2, extracting the fluid velocity provides a measurement of the gas volume at current working conditions (V_{mis}). Since gas is a highly compressible fluid, for billing purposes and to comply with the metrological requirements of the MID directive, it is strictly necessary to convert the measured volume into the volume at standard base conditions (V_b). Such thermodynamic compensation depends linearly on the operating temperature and pressure, in accordance with the ideal gas equation (Eq. 2.9).

$$V_b = V_m \cdot \frac{P_{mis}}{P_b} \cdot \frac{T_b}{T_{mis}} \cdot \frac{Z_b}{Z_{mis}}$$

Furthermore, accurate knowledge of the temperature is fundamental to estimate the correct theoretical speed of sound (c_{theo}) and validate the goodness of the ultrasonic measurement.

To acquire these environmental parameters, the system was equipped with a high-precision combined piezoresistive MEMS sensor, the LPS22DF model by STMicroelectronics, this component integrates both a barometer and a temperature sensor in a single chip. From a metrological point of view, this chip provides absolute pressure and temperature readings that are factory pre-calibrated and immune to external electromagnetic interference, guaranteeing a resolution of 0.00024 hPa and a relative accuracy of ± 0.01 hPa. Given such performance, the sensor maintains a current absorption of just $1.7\,\mu$A (sampling at 1 Hz), to then drop to only $0.9\,\mu$A in power-down mode. This extremely low absorption represents an indispensable requirement to meet the energy budget of a battery-powered smart meter.

At the architectural level, the LPS22DF was connected directly to the host microcontroller exploiting the I^2C serial communication bus; this allows the LPS22DF to work in a decoupled way with respect to the host, which will only have to deal with acquiring the temperature reading when necessary. Since the thermodynamic drifts of the gas mixture inside the pipe present time constants significantly higher compared to the flow dynamics (pressure and temperature variations occur very slowly), the acquisitions on the I^2C bus are activated with an extremely low frequency, typically once every 30 seconds in order to keep system consumptions to a minimum while still obtaining a sufficiently updated compensation for volumetric compensations.

7.7 Printed Circuit Board

Following the preliminary logical and circuit validation performed via evaluation board, we proceeded with the design of a custom PCB (Printed Circuit Board). The routing design required particular attention primarily on the signal tracks dedicated to the upstream and downstream transducers; these were traced maintaining strict geometric symmetry and a balancing of parasitic capacitances. This circuit precaution was necessary because any asymmetries in the electrical path would introduce differential propagation delays, which would add to the actual acoustic time of flight, compromising the calculation of D_{ToF} and the subsequent flow estimation. To guarantee maximum rejection of electromagnetic interference (EMI) and minimize background noise, the board was engineered by maximizing the extension of the ground plane (GND) on both layers. Specifically, the PCB was divided into functional macro-areas to physically separate them so as not to create disturbances between the digital domain and the ultrasonic domain. To protect the TDC chip, an extensive network of decoupling capacitors was added, strategically positioned in the immediate vicinity of the power supply pins.

The board was also equipped with a linear voltage regulator (XC6206P) in order to provide a stabilized, filtered and constant power supply to the rest of the components, protecting them from any power supply voltage ripples. As discussed in paragraph 2.2, a MEMS sensor (LPS22DF) was integrated, connected via I^2C, to acquire local temperature and pressure data, necessary for the volumetric conversion algorithm. As for the communication interface, the PCB hosts a serial communication bus (UART) directed towards the microcontroller (STM32) in order to dynamically vary the various settings of the firmware and the TDC and extract the various system output data.

Finally, the system debugging and calibration phase was supported by the insertion of numerous Test Points (TP) along the critical signal lines. These access points allow monitoring the analog waveforms and the various timing flags directly from the oscilloscope, allowing physically validating the firmware behavior and testing any architectural modifications in the field.

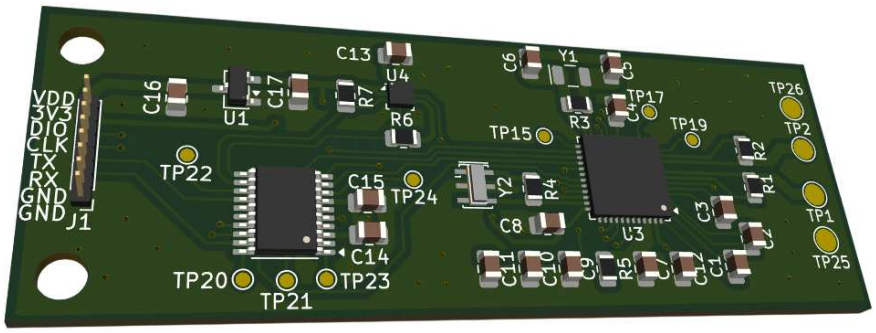

Figure 7.18: PCB

8. Results

In this chapter, the experimental results obtained from testing the developed ultrasonic flow sensor prototype are presented and analyzed. In the first section (8.1), the direct comparison between the raw (uncalibrated) measurements of the prototype and the industrial state of the art will be illustrated, evaluating key parameters such as offset, linearity, sensitivity, and drawing a projection of metrological compliance with the strict standards of the European MID directive. Subsequently (Section 8.2), the calibration techniques applicable to the system aimed at further refining the percentage error curve will be explored. Finally, Section 8.3 will be dedicated to the analysis of the power budget and global energy consumption, confirming the absolute suitability of the developed architecture for integration into modern battery-powered smart gas meters.

8.1 Comparison with the State of the Art

In order to rigorously evaluate the performance of the developed ultrasonic flow sensor, the prototype was tested on the same certified test bench used for the state-of-the-art analysis. To guarantee a direct and impartial comparison with the B2B commercial devices (Sensors A, B and C), the same standardized duty cycle was used, which covers the entire dynamic range required by the European MID directive for G4 size meters (from 0 to 6 m^3/h).

As highlighted in the previous chapters, the MID directive does not evaluate the single isolated sensor, but rather the entire finished and calibrated measurement system. However, in this experimental phase, it was decided to test the prototype without the application of fluid dynamic calibration curves, in order to be able to have a direct comparison with the other tested sensors. The objective is to demonstrate that the design choices guarantee a structural metrological stability requiring a minimum computational load for calibration, safeguarding the restriction of very low energy consumption.

8.1.1 Raw Performance Analysis

The raw acquisition of the developed prototype, in zero-flow conditions, shows an average offset of -0.0012 m^3/h with a standard deviation of just 0.00025 m^3/h. This result is excellent, it confirms that the implementation of the hybrid filter effectively suppresses thermal and quantization noise, positioning itself at the same level as the commercial "gold standard" Sensor A (which registers a residual reading of 0.0001 m^3/h) and Sensor B (which starts from -0.0013 m^3/h with greater deviation) without resorting to the aggressive software deadbands used by Sensor C to artificially mask the noise.

Analyzing the transient response to the flow rate steps, the prototype demonstrates to closely track the setpoints imposed by the bench along the entire dynamic range. Unlike Sensor C, which fails the measurement at high flow rates by collapsing to 5.368 m^3/h, and unlike Sensor B which shows ripple and fluctuations during descent, the realized sensor does not suffer from "fluid dynamic memory", reacting instantly and cleanly both in the ascending ramp and in the descending ramp.

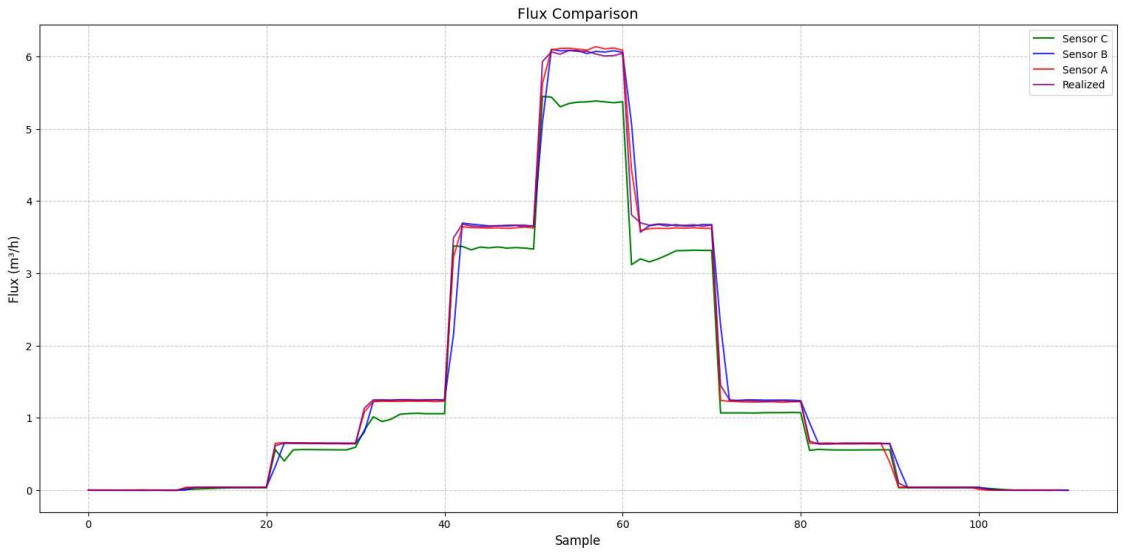

Figure 8.1: Comparison of the response of the sensors

Shifting attention to Linearity and Sensitivity, we confront the most critical hardware validation parameter, as it attests to the intrinsic quality of the mechanical design. As can be seen from the graph (Figure 8.2) and Table 8.1, the raw performance of the prototype is very good. The coefficient of determination (R^2) stands at 0.9999, effectively equaling the linear perfection of sensors A and B ($R^2 = 1.0000$) and surpassing Sensor C ($R^2 = 0.9983$). The sensitivity (Slope) equal to 1.001 also demonstrates that the structural gain of the system is almost ideally 1:1 compared

to the actual flow rate imposed by the bench.

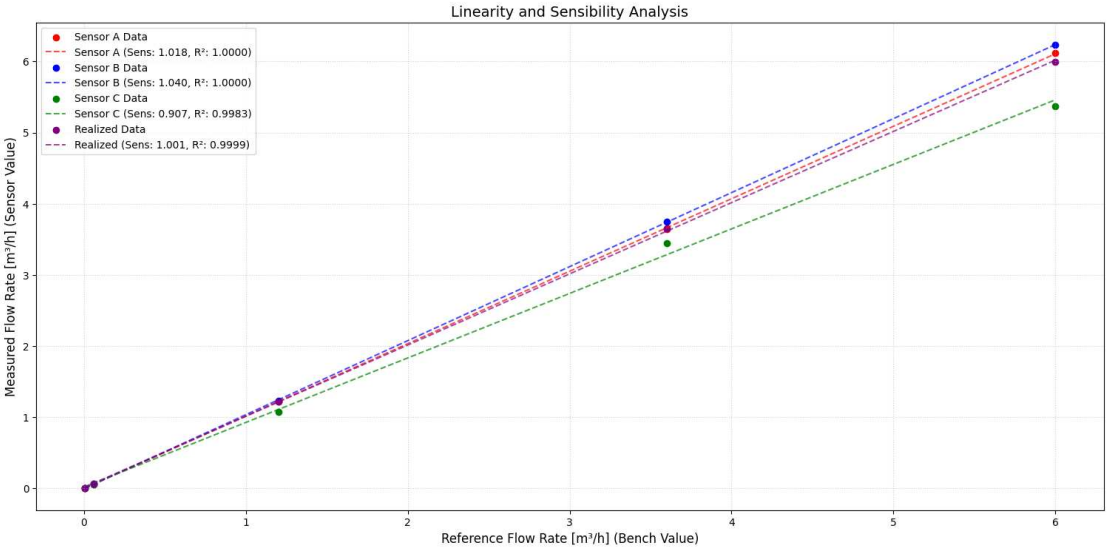

Figure 8.2: Raw linearity analysis: comparison between the bench reference flow rates and the measured values

Sensor	**Sensitivity (Slope)**	R^2
Sensor A	1.018	1.0000
Sensor B	1.040	1.0000
Sensor C	0.907	0.9983
Realized	1.001	0.9999

Table 8.1: Comparison of linearity and sensitivity parameters

8.1.2 Error Analysis and MID Compliance Projection

To validate the compliance of the system to the regulatory parameters, the Percentage Error (E%) of the uncalibrated prototype was calculated against the strict MID Class 1.5 tolerance bands, namely ±3% in the low zone from Q_{min} to Q_t, and ±1.5% in the nominal zone from Q_t to Q_{max}.

The plotted data (Figure 8.3) provide the final and most important design validation. Despite being a device without a calibration curve, the raw percentage error already touches MID compliance at high flow rates, recording an exact 0.00% deviation at 6 m^3/h and touching the regulatory limits with the tolerance band in the rest of the curve. The deviations found (in particular the error recorded around -7.50% at the minimum flow rate of $Q_{min} = 0.04\ m^3/h$) are the pure and physiological consequence of the natural offset and the lack of a trivial software compensation factor at the origin.

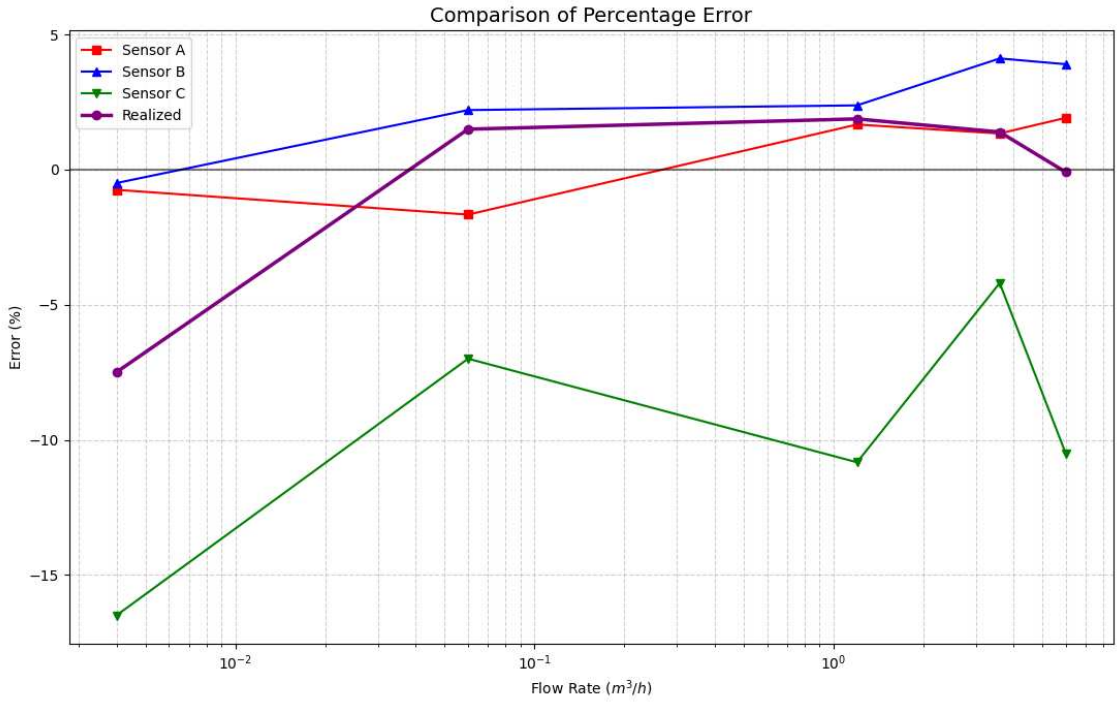

Figure 8.3: Comparison of the Percentage Error relative to the volumetric flow rate

Unlike less performing commercial sensors (like Sensor C, whose strong fluid dynamic distortions and sensitivity collapse to 0.907 force the use of complex polynomial interpolations with high computational absorption), the error structure of the developed prototype is strictly linked to a linear equation with $R^2 = 0.9999$. This means that, during the industrialization phase, to translate the entire curve within the narrow MID Class 1.5 band, it will be sufficient to apply a very light two-point correction via firmware (such as a simple offset adjustment or linear gain).

8.2 Calibration

In this section, two different compensatory strategies applicable to the system will be explored and compared: a two-parameter linear calibration (Offset and Multiplier), made possible by the marked intrinsic linearity of the hardware, and a more complex multipoint calibration based on polynomial regression, necessary if nonlinear trends need to be corrected or extreme accuracy requirements are required.

8.2.1 Two-Parameter Linear Calibration

The simplest and most efficient method is the linear calibration, which models the correction according to the equation of a straight line:

$$Q_{corr} = (Q_{mis} + \text{Offset}) \times \text{Multiplier} \tag{8.1}$$

This strategy uses only two parameters to translate and rotate the reading curve. The addition of a constant Offset has a percentage impact inversely proportional to the flow rate, it intervenes massively on minimum flows, compensating for the severe physiological underestimates of the sensor close to zero and allowing the starting threshold (Q_{start}) to be calibrated precisely. Although a positive offset generates a high percentage error below the minimum approved flow rate (Q_{min}), this deviation is irrelevant for certification purposes, since in that zone the meter must only detect the passage of gas (e.g. for leak detection) without strict accuracy constraints. The Multiplier (or Gain), instead, acts as a span correction, proportionally compressing or expanding the values to counteract the sensor's tendency to overestimate or underestimate in medium-high regimes.

In the present case, the empirical optimization led to the identification of an ideal pair of parameters: an offset of +3.7 L/h and a multiplier of 0.99. As visible in Table 8.2, this simple calibration allowed flattening the error curve, confining the entire excursion of the sensor to an excellent range between -1.04% and +1.16% over the entire operating spectrum, falling within the MPEs required by the MID directive, which we recall being ±3% for flow values between Q_{min} (0.04 m^3/h for category G4) and Q_t (1.2 m^3/h for category G4), and ±1.5% for values between Q_t and Q_{max} (6 m^3/h for category G4).

Real Flow [l/h]	Q_{mis} [l/h]	**Raw Error**	Q_{corr} [l/h]	**Calibrated Error**
40	37.00	-7.5%	40.29	**+0.73%**
600	608.96	+1.49%	606.53	**+1.09%**
1200	1222.49	+1.87%	1213.93	**+1.16%**
3600	3650.00	+1.39%	3617.16	**+0.48%**
6000	5994.03	-0.1%	5937.75	**-1.04%**

Table 8.2: Results of the linear calibration

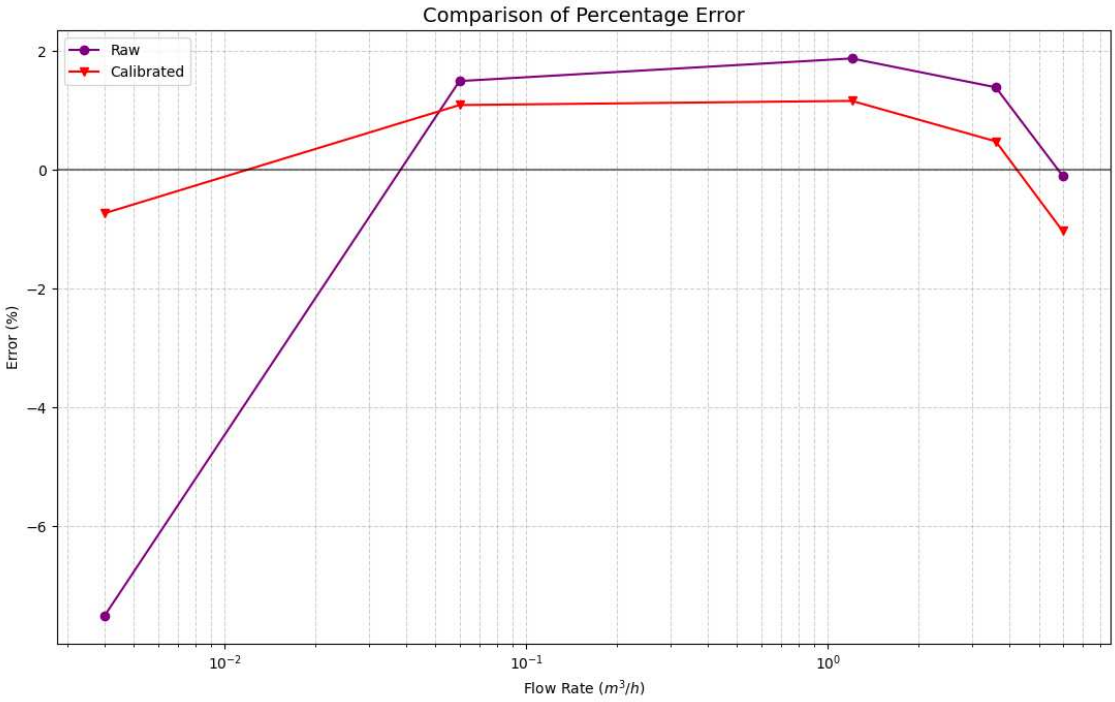

Figure 8.4: Comparison between the raw percentage error and the error after linear calibration

8.2.2 Multipoint Calibration and Polynomial Regression

Should the required accuracy exceed the physical limits of two-parameter linearization, it is necessary to employ to a multipoint calibration. The most sophisticated algorithm to obtain a single continuous curve, free from logical discontinuities, is the use of a polynomial of degree $N-1$, where N represents the number of calibration points.

The fundamental rule of this approach is that as the sampling points increase, the curve's ability to accommodate the sensor's distortions increases, reducing the overall error. However, this increase in precision brings with it an increase in computational cost. Using, for example, 5 calibration points, the system calculates a 4th-degree polynomial:

$$Q_{corr} = a \cdot Q_{mis}^4 + b \cdot Q_{mis}^3 + c \cdot Q_{mis}^2 + d \cdot Q_{mis} + e \tag{8.2}$$

To mitigate the heavy computational cost of the polynomial and make the algorithm usable even on embedded architectures, Horner's Method can be used. This mathematical technique decomposes the polynomial by factoring out the independent variable Q_{mis}, completely eliminating the need to use slow power functions and

reducing the equation to a simple sequence of nested multiplications and additions:

$$Q_{corr} = ((((a \cdot Q_{mis}) + b) \cdot Q_{mis} + c) \cdot Q_{mis} + d) \cdot Q_{mis} + e \tag{8.3}$$

The peculiar mathematical characteristic of a polynomial generated on 5 points is that the calculated curve will pass exactly over those coordinates. Consequently, in correspondence with the selected calibration flow rates, the measurement error can be mathematically equal to zero, however, it is essential to analyze the behavior of the curve outside these exact nodes. First, although the curve ensures a smooth transition between the calibrated points, the error does not remain strictly at zero. If the physical behavior of the sensor is very irregular, the high-degree polynomial will tend to oscillate physiologically between one point and another, the higher the degree of the polynomial, the greater the risk of instability in the intermediate values. Secondly, if the sensor detects a flow rate lower than the first calibration point (e.g. flow rates close to zero) or higher than the last (beyond 6 m^3/h), the polynomial equation tends to diverge violently towards infinity. In these so-called "out-of-scale" zones, the error can explode to unacceptable values. It is therefore mandatory to implement safety limits (clamping) at the firmware level, forcing the algorithm to use a fixed correction factor or to freeze the curve as soon as the tested safety margins are exceeded.

8.3 Power Consumption

Energy efficiency represents a fundamental technical requirement for modern smart gas meters, which must guarantee years of battery autonomy as prescribed by the European MID directive. The total energy consumption of the system is given by the sum of the contributions of the three main active components: the ultrasonic front-end (TDC), the thermodynamic sensor (LPS22DF), and the microcontroller (i.e. STM32U031F4Px).

The choice to adopt a TDC-based architecture proved to be a winner from an energy point of view. The hardware of this chip is awakened exclusively for the microseconds necessary for the emission and reception of the ultrasonic burst. From the empirical measurements acquired via the Otii Battery Profiler (Figure 8.5), the TDC chip registers an average absorption (AVG) of just 11.0 μA over a standard acquisition window. During the active phase of acoustic excitation, physiological current peaks (MAX) reaching 571 μA are observed; however, given the extreme brevity of these transients, their impact on the long-term average consumption is minimal.

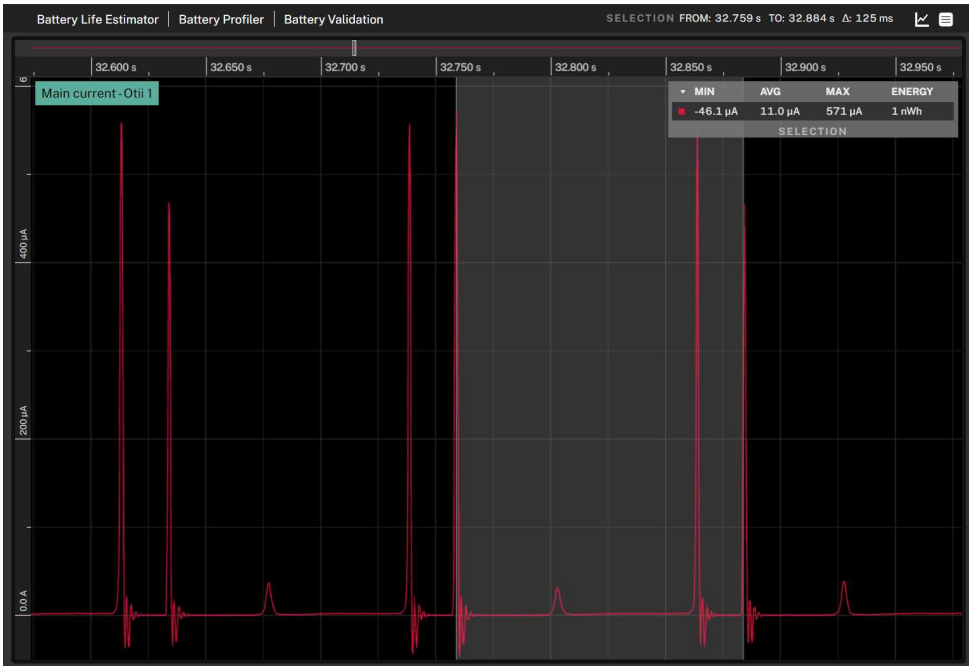

Figure 8.5: Absorption profile of the TDC chip

For volumetric thermodynamic compensation, the system relies on the LPS22DF combined MEMS pressure and temperature sensor. This component is characterized by intrinsically low consumptions: it absorbs just 1.7 μA when sampling at 1 Hz, and drops to 0.9 μA in power-down mode. Given that the thermodynamic phenomena of the gas and temperature variations are slow processes, acquisitions via the I^2C bus occur with a reduced frequency (once every 30 seconds), allowing the sensor to spend almost all of its time in power-down and stabilizing its overall average contribution around 0.9 μA [3].

The task of managing and processing the measurements can be entrusted, for example, to the ultra-low-power STM32U031F4Px microcontroller. The use of TDC technology guarantees near-zero computational latency, allowing the microcontroller to almost instantly return to a deep sleep state between one acquisition and the next. According to the official specifications, in active mode (Run mode with LDO regulator) the STM32U031F4Px registers a dynamic absorption of only 52 μA/MHz. For the vast majority of the time, having an output duty cycle set to 0.5 Hz, the device enters a state of deep energy saving, called Stop Mode, which keeps the contents of the SRAM and registers intact, limiting consumption to only 0.515 μA without RTC, or 0.63 μA with active RTC. Furthermore, the extremely fast wake-up time of just 4 μs from Stop mode minimizes any energy waste during the transition between the stand-by and active calculation phases [4].

Excluding the brief dynamic operating time of the microcontroller, the resting

baseline established by the TDC block and the LPS22DF sensor stands at about 11.9 μA, adding the static currents of the microcontroller ($\sim 0.6\mu A$) and averaging the short wake-up peaks of the STM32 ($\sim 6\mu A$), the complete system maintains a global operating consumption that hovers around 20 μA.

This power budget positions the realized prototype in a condition of absolute metrological and energy advantage compared to the commercial solutions examined in Chapter 6:

- The **Sensor A** (primary reference) consumes on average 26.8 μA working at 0.5 Hz, managing to drop to 10.2 μA only by partially disabling itself after 30 consecutive seconds of absence of flow.
- **Sensor B** consumes on average 36 μA in normal mode at 0.5 Hz.
- **Sensor C** is the most expensive, absorbing 40 μA in Normal mode at 1 Hz and dropping to 25 μA in Low-power mode at 0.5 Hz.

9. Future Developments

The raw performance of the prototype already satisfies, at an embryonic level, the rigorous parameters imposed by the European MID Directive. However, to evolve this prototype from an academic validation phase to an "industrial-ready" device ready for the smart meter market, it is possible to outline a series of future developments aimed at software, functional, and certification optimization.

9.1 Implementation of an Industrial Communication Firmware (HEX Stack)

Although the current system successfully acquires, filters, and processes the fluid dynamic data, the fundamental step for commercial integration is the development of a standardized communication stack. The developed PCB hardware is already physically predisposed with a direct UART serial bus to the STM32 microcontroller. The future objective will be to align this interface with the standards of the B2B devices analyzed in the state of the art (such as Sensors A, B, and C), implementing a transmission strictly based on packets in HEX format.

Specifically, the implementation of the following is foreseen:

- Handshaking protocols (ACK/NAK) to guarantee the robustness of the transmission against electromagnetic interferences.
- A dynamic management for switching between a Normal mode (with measurements every 2 seconds) and a Low-power mode (which partially deactivates the system in case of prolonged absence of flow), in order to further optimize the energy budget.
- Implementing diagnostic error flags to promptly signal the onset of critical phenomena such as reverse flow, hardware drifts, or the presence of bubbles and impurities in the pipe.

9.2 Gas Composition Estimation and "Hydrogen Blending"

One of the most current metrological challenges, linked to the decarbonization of distribution networks, is the introduction of natural gas and hydrogen mixtures (Hydrogen Blending).

As analyzed in Chapter 1, unlike thermomassive meters, ultrasonic technology enjoys the mathematical advantage of being independent of the speed of sound in the calculation of the flow velocity. This allows the sensor not to suffer direct metrological drifts in the presence of hydrogen, tolerating mixtures with H_2 concentrations up to 23% by volume.

A software development of extreme industrial interest will consist in exploiting the native ability of the sensor to inversely derive the real speed of sound of the measured medium (c_{means}) and compare it with the theoretical speed of sound (c_{theo}), allowing real-time gas profiling capabilities. The device will become able to estimate the exact percentage of hydrogen injected into the network or to detect any anomalies or contaminations, allowing the volumetric meter to analyze the gas quality, without requiring the addition of any hardware component.

9.3 Evaluation of Long-Term Stability and Mechanical Transition

The ultimate goal of a smart metering device is the compliance with strict durability requirements, requiring a guaranteed battery autonomy from 5 to 10 years and a metrological incorruptibility over time.

Future developments in this perspective will have to focus on Stress Tests and Long-Term Stability, subjecting the system to prolonged test cycles in climatic chambers to simulate the entire operating thermal range (from $-25°C$ to $+55°C$) for an extended period.

This will allow evaluating empirically the possible zero-flow drift caused by the aging of the piezoelectric transducers and validating the hold over time of the very low current absorption (about $20\mu A$).

10. Conclusions

The present thesis work has illustrated the development, design and metrological validation of an ultrasonic gas flow sensor based on a TDC (Time-to-Digital Converter) architecture. The primary objective was to realize a prototype capable of satisfying the rigorous requirements of the European MID (Measuring Instruments Directive) for class 1.5 meters.

From a mechanical point of view, a single-path architecture with a V-shaped geometry was engineered to exploit the self-compensation of fluid dynamic disturbances. The measuring tube was optimized with an incidence angle of 65°, an acoustic path length of 53.2 mm and the adoption of separation plates for the laminar stabilization of the flow. The adoption of TDC technology, operating with 500 kHz transducers, proved to be the winning choice to minimize computational latencies and reduce energy consumption.

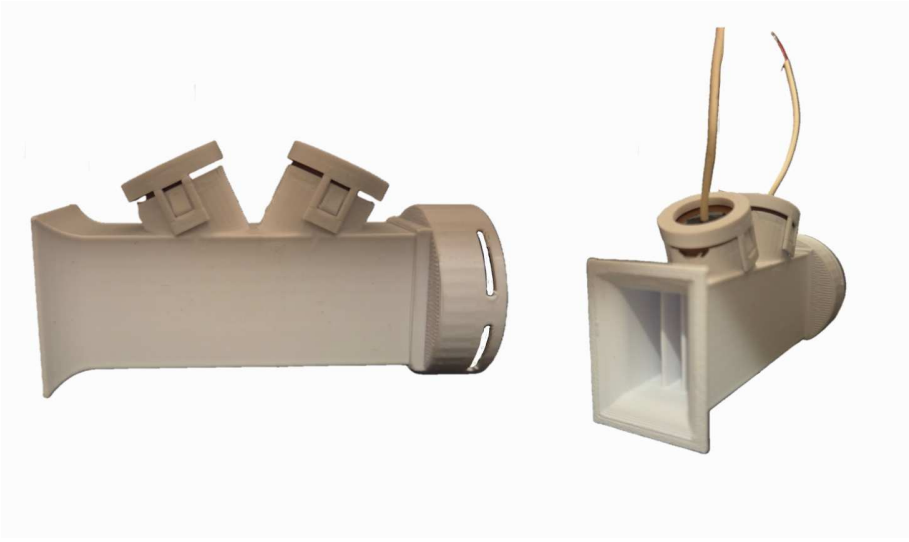

Figure 10.1: 3D Printed Sensor

Through the use of the Split Burst excitation technique, which introduces a phase shift as a time marker, and the integration of a programmable gain amplifier (PGA), the hardware system proved to be intrinsically immune to the Cycle Slip

phenomenon and highly robust against signal amplitude variations.

A crucial milestone of the project was the development of the firmware architecture for filtering stochastic and impulsive noise. For pre-filtering, a combination of Shell Sort and Peters filter was adopted for the isolation of outliers. For post-filtering, overcoming the heavy phase delays intrinsic to classic moving average filters, an Adaptive Hybrid Filter was developed that dynamically mix the consistent zero-flow stability of the (one-dimensional) Kalman filter with the immediate reactivity of the IWR (Improved Weighted Recursive) filter. In this way, the microcontroller is able to follow flow variations without delays, guaranteeing a response profile capable of perfectly following the fluid variations.

The experimental tests, performed on a certified calibration bench and based on a MID duty cycle from 0 to 6 m^3/h, demonstrated that the prototype in comparabte to the performance of the commercial state of the art (B2B Sensors A, B and C). In the raw measurements, the sensor exhibited a negligible offset (-0.0012 m^3/h) and an exceptional linearity, quantified in a coefficient of determination R^2 equal to 0.9999 and a sensitivity (slope) of 1.001.

This extraordinary mechanical and hardware linearity avoided the heavy use of heavy multipoint polynomial regressions, it was in fact sufficient to apply a very light two-parameter linear calibration (an Offset of +3.7 l/h and a Multiplier of 0.99) to translate and rotate the error curve. After this compensation, the system error was confined to a narrow interval between -1.04% and +1.16%, hitting with an excellent safety margin the strict limits of the Maximum Permissible Error (MPE) dictated by the MID class 1.5.

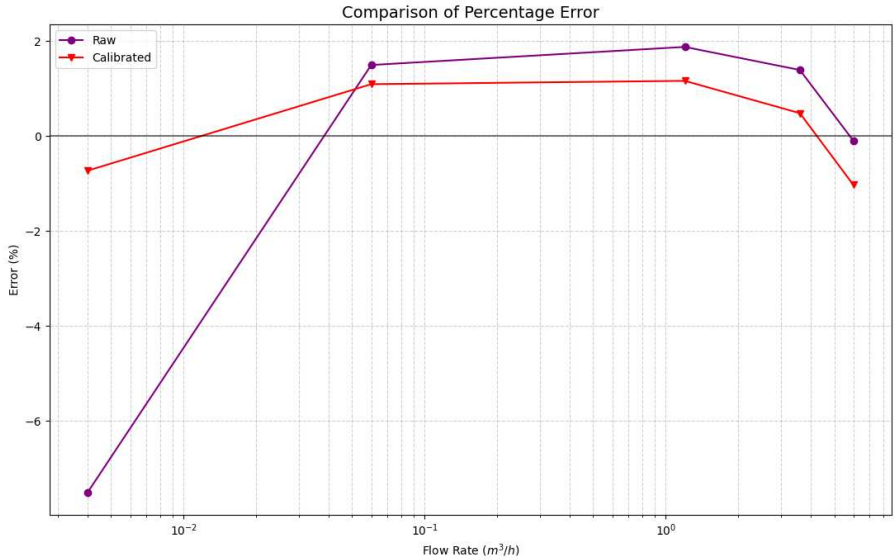

Figure 10.2: Comparison between the raw error curve and the calibrated one

Finally, the most decisive result for the industrial applicability of the thesis concerns the energy budget, in fact, the global average consumption of the running system was estimated at about 20 μA. This value positions the prototype in a condition of absolute metrological and energy advantage compared to all the commercial solutions analyzed in the state of the art, which register variable absorptions between 26.8 μA and 40 μA.

In conclusion, the developed prototype concretely demonstrates how a co-design between mechanical architecture, ultra-low-power TDC electronics and adaptive filtering can produce a highly reliable measuring instrument, ready to support the technological evolution towards modern smart metering networks.

Bibliography

[1] Affidabilità, versatilità e accuratezza con i misuratori massici a dispersione termica. 2020.

[2] Natural Gas Flow Meter Selection Guide. 2022.

[3] Lps22df datasheet: Low-power and high-precision mems nano pressure sensor: 260-1260 hpa absolute digital output barometer. 2023.

[4] Stm32u031x4/6/8 datasheet: Ultra-low-power arm cortex-m0+ 32-bit mcu, up to 64-kbyte flash memory, 12-kbyte sram. 2024.

[5] Aurélien Barra. Misuratore di portata Doppler vs misuratore di portata a tempo di transito: confronto tra misuratori di portata a ultrasuoni per prestazioni ottimali.

[6] Aurélien Barra. Come funzionano i flussimetri a ultrasuoni? Scoprire la tecnologia dei flussimetri a ultrasuoni. 2024.

[7] Ching-Mei Chen and Bhaskar Choubey. Ultrasound Transducer Quality Factor Control Using Coupled External Electrical Resonator. 2018.

[8] Jianfeng Chen, Kai Zhang, Leiyang Wang, and Mingyue Yang. Design of a High Precision Ultrasonic Gas Flowmeter. *Sensors*, 2020.

[9] Guilherme Siqueira de Aquino, Ramon Silva Martins, Marcio Ferreira Martins, and Rogério Ramos. An Overview of Computational Fluid Dynamics as a Tool to Support Ultrasonic Flow Measurements. *Metrology*, 2025.

[10] M. Dell'Isola, P. Vigo, G. Ficco, and L. Celenza. Analisi e caratterizzazione metrologica dei sistemi di misura delle reti termiche distribuite. 2015.

[11] A. Ebrahimi-Moghadam, B. Mohseni-Gharyehsafa, A. Jabari moghadam, M. Farzaneh-Gord, Y. Bijarzehi, and V. Okati. Prediction of Natural Gas Z-Factor by Using ANN Method. 2018.

[12] European Parliament and Council of the European Union. Directive 2014/32/EU of the European Parliament and of the Council of 26 February 2014 on the harmonisation of the laws of the Member States relating to the making available on the market of measuring instruments (recast). *Official Journal of the European Union*, 2014.

[13] Giorgio Ficco, Marco Dell'Isola, Giorgio Graditi, Giulia Monteleone, Paola Gislon, Pawel Kulaga, and Jacek Jaworski. Reliability of Domestic Gas Flow Sensors with Hydrogen Admixtures. *Sensors*, 2024.

[14] José R. García Oya, Alejandro Sainz Rojas, Daniel Narbona Miguel, Ramón González Carvajal, and Fernando Muñoz Chavero. Low-Power Transit Time-Based Gas Flow Sensor with Accuracy Optimization. *Sensors*, 2022.

[15] Claas Hartmann, Christoph Haugwitz, Gianni Allevato, Matthias Rutsch, Jan Hinrichs, Johannes Brötz, Dieter Bothe, Peter Pelz, and Mario Kupnik. Ray-tracing simulation of sound drift effect for multi-path ultrasonic high-velocity gas flow metering. *IEEE*, 2021.

[16] Luis Iglesias Hernandez, Priyadarshini Shanmugam, Jean-François Michaud, Daniel Alquier, Dominique Certon, and Isabelle Dufour. Gas discrimination by simultaneous sound velocity and attenuation measurements using uncoated capacitive micromachined ultrasonic transducers. *Scientific Reports*, 2022.

[17] Jiangsu Hongguang Instrument Factory Co., Ltd. Misuratore di portata massico termico vs misuratore di portata magnetico: Le principali differenze spiegate. 2024.

[18] Lingkai Kong, Liang Zhang, Hulin Guo, Ning Zhao, and Xinhu Xu. Time Delay Study of Ultrasonic Gas Flowmeters Based on VMD-Hilbert Spectrum and Cross-Correlation. *Sensors*, 2024.

[19] Yaoying Lin. Signal Processing and Experimental Technology in Ultrasonic Flow Measurement. 2004.

[20] Rui Ren, Hongliang Wang, Xiaolei Sun, and He Quan. Design and Implementation of an Ultrasonic Flowmeter Based on the Cross-Correlation Method. *Sensors*, 2022.

[21] Merima Toromanović, Jasmina Ibrahimpašić, Ifet Šišić, and Samira Hotić. Elements of Ultrasonic Flowmeter Installation, Model Omni TDI 200H with Test Work Results. *ISSN*, 2022.

[22] Rang-ding Wang, Qiang Liu, Chen-tou Du, and Ling Yao. A Signal Preprocessing Algorithm Applied for Ultrasonic Flow-meter. *TELKOMNIKA*, 2013.

[23] Nannan Zhao, Liang Hu, Xiangxiang Peng, Zehua Fang, Wenyu Chen, and Xin Fu. A method combining measurement tool and numerical simulation for calculating acoustic signals of ultrasonic flowmeter. *IEEE*, 2019.

[24] Hongliang Zhou, Yanchu Liu, and Yunxiao Wu. Transit Time Determination Based on Similarity-Symmetry Method in Multipath Ultrasonic Gas Flowmeter. *Metrology*, 2025.

[25] Junyang Zhou, Peng Wang, Ruibing Wang, Zhenyu Niu, and Kai Zhang. Signal Processing Method of Ultrasonic Gas Flowmeter Based on Transit-Time Mathematical Characteristics. *Sensors*, 2025.