



POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Self-corrected dual-comb spectroscopy in the near-infrared by 1-GHz solid state lasers

TESI DI LAUREA MAGISTRALE IN
ENGINEERING PHYSICS - INGEGNERIA FISICA

Author: **Serena Iacone**

Student ID: 273621

Advisor: Prof. Alessio Gambetta

Co-advisors: Dr. Francesco Canella, Dr Gianluca Galzerano

Academic Year: 2025-26

Abstract

Optical frequency comb (OFC) technology has revolutionized modern laser spectroscopy, merging broad spectral coverage with the high accuracy and precision of atomic frequency standards. Among comb-based techniques, **Dual-Comb Spectroscopy (DCS)** stands out as a powerful tool for high-speed, calibration-free, and moving-part-free Fourier transform interferometry. This technique has transformative applications in a range of fields, including long-range remote sensing, human breath analysis, LiDAR (Light Detection and Ranging), and hyperspectral imaging.

However, the performance of DCS is fundamentally limited by the tight **mutual coherence** requirements between the two optical combs, where tiny relative phase and timing fluctuations (**jitter**) prevent **coherent averaging**, degrading the signal-to-noise ratio (SNR). Therefore, several methods were implemented to correct these fluctuations, such as hardware comb stabilisation and active feed-forward loops.

Nevertheless, all these techniques require complex and bulky equipment, resulting in an expensive and not very scalable setup.

This thesis presents a different approach: a Self-Corrected Dual-Comb Spectrometer that can bypass stringent hardware feedback loops via advanced numerical post-processing. By utilizing the optical interferogram itself to extract phase and frequency noise, an inverse-time-warping and **Hilbert-transform**-based phase correction algorithm is implemented to correct timing fluctuations *a posteriori*, restoring full mutual coherence and enabling long-term coherent averaging.

In the final part of this work, an even lighter and more flexible approach is explored: a real-time self-correction implemented directly on dedicated hardware. To pave the way toward this kind of solution, the analytical self-correction algorithm is investigated for future hardware acceleration, evaluating its translation into a hardware-friendly **Finite Impulse Response (FIR) filter** framework. Preliminary steps are taken utilizing an **AMD RFSoc 4x2** development board, establishing its development workflow within the **Vivado** design suite and testing basic test-bench programs. The board's initial acquisition capabilities are characterized at a high level using synthesized waveforms from a generator, while the feasibility of the real-time pipeline is validated through Python-based

simulations. This foundational exploration maps out the core architecture and identifies the primary implementation challenges, laying the groundwork for future on-chip **real-time phase correction**.

Keywords: Dual-Comb Spectroscopy (DCS), Optical Frequency Combs (OFC), Mutual Coherence, Phase Jitter, Coherent Averaging, Hilbert Transform, Finite Impulse Response (FIR) Filters, AMD RFSoc, Vivado, Real-Time Phase Correction

Abstract in lingua italiana

La tecnologia dei **pettini di frequenza ottica (OFC)** ha rivoluzionato la moderna spettroscopia laser, unendo un'ampia copertura spettrale con l'elevata accuratezza e precisione degli standard di frequenza atomici. Tra le tecniche basate sui pettini, la **spettroscopia a doppio pettine (DCS)** si distingue come un potente strumento per l'interferometria a trasformata di Fourier ad alta velocità, senza calibrazione e priva di parti in movimento. Questa tecnica trova applicazioni rivoluzionarie in molteplici settori, tra cui il telerilevamento a lungo raggio, l'analisi del respiro umano, il LiDAR (Light Detection and Ranging) e l'hyperspectral imaging. Tuttavia le prestazioni della DCS sono fundamentalmente limitate dai rigidi requisiti di **mutua coerenza** tra i due pettini ottici, dove piccole fluttuazioni relative di fase e di tempo (**jitter**) impediscono l'operazione di **media coerente**, degradando il rapporto segnale-rumore (SNR). Pertanto, sono stati implementati diversi metodi per correggere queste fluttuazioni, come la stabilizzazione hardware dei pettini e active feed-forward loops. Tuttavia, tutte queste tecniche richiedono apparecchiature complesse e ingombranti, risultando in un setup costoso e poco scalabile.

Questa tesi presenta un approccio differente: uno spettrometro a doppio pettine auto-corretto (self-correction) in grado di aggirare i rigidi feedback loops hardware tramite un'elaborazione numerica avanzata in post-processing. Utilizzando l'interferogramma ottico stesso per estrarre il rumore di fase e di frequenza, viene implementato un algoritmo di correzione di fase basato sull'inverse-time-warping e sulla **trasformata di Hilbert** per correggere le fluttuazioni temporali *a posteriori*, ripristinando la mutua coerenza e consentendo medie coerenti a lungo termine.

Nell'ultima parte di questo lavoro, viene esplorato un approccio ancora più leggero e flessibile, una self-correction in tempo reale implementata direttamente su un hardware dedicato. Per spianare la strada verso questa soluzione, l'algoritmo analitico di autocorrezione viene studiato in ottica di una futura accelerazione hardware, valutandone la conversione nell'ambito di **filtri a risposta impulsiva finita (FIR)** adatti all'hardware. I passi preliminari vengono mossi utilizzando una scheda di sviluppo **AMD RFSoc 4x2**, definendo il flusso di lavoro all'interno della suite di progettazione **Vivado** e testando programmi di prova di base. Le capacità iniziali di acquisizione della scheda sono carat-

terizzate ad alto livello utilizzando forme d'onda sintetizzate da un generatore, mentre la fattibilità della pipeline in tempo reale è validata attraverso simulazioni in Python. Questa esplorazione fondamentale delinea l'architettura di base e identifica le principali sfide di implementazione, gettando le basi per una futura **correzione di fase in tempo reale** direttamente sul chip.

Parole chiave: Spettroscopia a Doppio Pettine (DCS), Pettini di Frequenza Ottica (OFC), Mutua Coerenza, Jitter di Fase, Media Coerente, Trasformata di Hilbert, Filtri a Risposta Impulsiva Finita (FIR), AMD RFSoc, Vivado, Correzione di Fase in Tempo Reale

Contents

Abstract	i
Abstract in lingua italiana	iii
Contents	v
Introduction	1
1 Principles of Optical Frequency Comb Spectroscopy	3
1.1 Introduction to Laser Spectroscopy	3
1.2 Optical Frequency Combs	5
1.2.1 Mode-locking	5
1.2.2 Carrier-Envelope Phase and the Comb Equation	7
1.3 Chirped Pulse Amplification (CPA) and Supercontinuum (SC) Generation	8
1.3.1 Chirped Pulse Amplification: Managing Peak Power	9
1.3.2 Supercontinuum Generation: Spectral Broadening	11
1.4 Frequency Comb Stabilization	12
1.4.1 Phase-Locked Loop (PLL)	13
1.4.2 The $f - 2f$ Interferometer: Measuring the Carrier-Envelope Offset .	14
1.5 Introduction to Detection Techniques	16
2 Dual-Comb Spectroscopy	19
2.1 Dual Comb Spectroscopy: the principle	19
2.1.1 Multi-heterodyne detection	19
2.1.2 The Time-Domain Signal: The Interferogram	21
2.1.3 Down-conversion and the Nyquist-Shannon sampling theorem . . .	22
2.2 Impact of the frequency noise on dual-comb interferograms	24
2.2.1 Origins of Instabilities	24
2.2.2 Coherent Averaging and SNR Constraints	26

2.2.3	Techniques for Fluctuation Compensation in Dual-Comb Systems	27
2.3	Self-Correction algorithms	29
2.3.1	Post-Processing Self-Correction: Algorithmic Framework	29
2.4	Towards real-time self-correction	30
2.4.1	Discrete Hilbert Transform	31
2.4.2	From SoC to RFSoc	33
2.4.3	FPGA : The Computational Engine of the RFSoc	33
3	Dual Comb Spectroscopy Setup	37
3.1	Experimental layout	37
3.2	Absolute Calibration and Detection	40
3.3	Data Acquisition Systems (DAQ)	43
3.3.1	High-Definition Digital Oscilloscope	43
3.3.2	AMD Xilinx RFSoc 4x2 Platform	44
4	Experimental measurements and results: Post processing Data Analysis	47
4.1	The Self-Correction Framework	49
4.2	Absolute Frequency Calibration	54
4.3	Further averaging	57
4.4	Final Spectrum obtained: normalization	59
5	Experimental measurements and results: towards Real-Time Analysis	63
5.1	Python simulation of the real-time algorithm	64
5.2	External Signal Generation and RFSoc Acquisition Setup	66
5.3	Vivado Design Suite and the Base Overlay	69
6	Conclusions and future developments	75
	Bibliography	77
A	Appendix A	81
A.1	Fundamentals of Molecular Laser Spectroscopy	81
A.1.1	Molecular Energy Levels and Roto-Vibrational Transitions	81
B	Appendix B	83
B.1	Analytical Signal Representation and Hilbert Transform	83

C	Appendix C	87
C.1	Code Implementation and Software Framework	87
C.1.1	Script 1: The Self-Correction Framework	87
C.1.2	Script 2: Multi-Trace Averaging	92
C.1.3	Script 3: Absolute Frequency Calibration	95
C.1.4	Script 4: Baseline Extraction and Spectrum Normalization	96
C.1.5	Python PYNQ Online Acquisition Script	97
	List of Figures	101
	List of Tables	107
	Acknowledgements	109

Introduction

Optical spectroscopy represents one of the most versatile and powerful methods for the non-intrusive investigation of matter. By analyzing the absorption patterns that arise from the interaction between radiation and molecular energy levels, it is possible to identify chemical species and quantify their presence with high sensitivity.

In this context, traditional techniques like Fourier Transform Infrared (FTIR) spectroscopy have long provided a reliable standard for broadband measurements. However, these methods are often constrained by the physical movement of mechanical components and the use of incoherent light sources, which introduce critical mechanical misalignments and limit the speed of data acquisition, the ultimate precision and the sensitivity of the measurement.

The development of the Optical Frequency Comb (OFC) has bypassed these limitations [18]. By emitting a train of ultrashort pulses that corresponds, in the frequency domain, to a dense grid of evenly spaced laser lines, the OFC provides a perfectly defined spectral reference. Unlike thermal sources, a frequency comb combines an extremely broad emission bandwidth with the high brightness and spatial coherence of a laser, allowing for the simultaneous excitation of a vast number of molecular transitions. This unique structure turns the laser into a powerful tool for high-resolution sensing, particularly when targeting the mid-infrared region where molecules exhibit significantly stronger roto-vibrational absorption bands.

Building upon the single-comb architecture, Dual-Comb Spectroscopy (DCS) has emerged as a transformative tool [4] for high-speed molecular fingerprinting. By employing two combs with slightly different repetition rates, DCS enables a process of multi-heterodyne detection: the optical absorption information is down-converted directly into the radio-frequency (RF) domain. This enables the recording of entire broadband spectra, eliminating the need for any moving parts and allowing for acquisition times several orders of magnitude shorter than conventional spectrometers.

Thanks to these distinct advantages, DCS can be seen as a transformative tool across various high-impact fields. In environmental monitoring, it enables long-range remote sensing [21], allowing for the simultaneous tracking of multiple greenhouse gases with

unprecedented temporal resolution [20]. In the biomedical field, the combination of high selectivity and rapid acquisition opens up new frontiers for non-invasive medical diagnostics through highly sensitive breath analysis, where complex mixtures of volatile organic compounds can be detected at trace levels as potential disease biomarkers [12]. The unique coherence and agility of frequency combs have also revolutionized optical imaging, paving the way for next-generation high-speed LIDAR systems and ultra-fast Optical Coherence Tomography (OCT) capable of capturing dynamic processes without any mechanical beam-steering. Furthermore, while historically focused on the infrared molecular fingerprint region, the application of dual-comb architectures is rapidly expanding into the visible spectral range, unlocking new possibilities [14]. The primary challenge in DCS, however, is maintaining the phase coherence between the two sources. Even in systems where the combs are individually stabilized, residual timing and phase jitters can rapidly randomize the relative phase, causing a "smearing" of the absorption features and preventing the coherent averaging of the signal. Addressing this instability is the central theme of this thesis. The present work explores the shift from complex hardware-based feedback loops toward computational stabilization. This approach is first addressed through the development of a post-processing self-correction algorithm in the MATLAB environment, where the phase noise is extracted directly from the recorded interferograms using Hilbert-transform-based techniques. The study then moves toward the next frontier of this technology: real-time acquisition and self-correction. By using the AMD RFSoc (Radio Frequency System on Chip) platform and the Vivado design suite, the thesis investigates how these analytical correction methods can be translated into digital filters and hardware architectures. This transition from offline analysis to an integrated, real-time one represents a critical step toward the development of compact and robust dual-comb spectrometers.

1 | Principles of Optical Frequency Comb Spectroscopy

This chapter provides the theoretical and technical foundation for understanding Optical Frequency Combs (OFCs) and their central role in modern laser spectroscopy. The discussion begins with the fundamental principles of molecular absorption and the limitations of traditional spectroscopic techniques, which set the stage for the revolutionary impact of frequency combs. The core of the chapter focuses on the generation of OFCs through mode-locking, deriving the fundamental comb equation. This equation establishes a direct link between the optical domain and measurable radio-frequency (RF) signals, mapping hundred-terahertz optical oscillations down to the RF domain via the repetition rate (f_{rep}) and the carrier-envelope offset frequency (f_{ceo}). To adapt these sources for high-precision sensing, the text explores the essential stages of Chirped Pulse Amplification (CPA) and Supercontinuum (SC) generation, which provide the necessary peak power and spectral broadening. Finally, the chapter details the critical stabilization techniques, including Phase-Locked Loop (PLL) architectures and $f - 2f$ self-referencing interferometry, which transform the frequency comb into a stable "optical ruler" capable of achieving hertz-level resolution and high accuracy. This framework enables an ultimate relative precision on the order of $\frac{\Delta\nu_{\text{line}}}{\nu_0} \approx 10^{-17}$ (assuming a central optical frequency ν_0 of 100 THz).

These principles form the basis for the advanced dual-comb sensing architectures that will be discussed in detail in the subsequent chapters.

1.1. Introduction to Laser Spectroscopy

Molecular spectroscopy is a powerful tool for the investigation of matter, allowing for the identification and quantification of chemical species by studying their interaction with electromagnetic radiation. Each molecule is characterized by a unique set of energetic transitions (rotational, vibrational, and electronic) that result in a "spectral fingerprint". The fundamental principle governing these interactions is the Beer-Lambert law. When a beam of light with incident intensity $I_i(\nu)$ propagates through a gas sample of length

L , the transmitted intensity $I_t(\nu)$ is attenuated according to the exponential relation:

$$I_t(\nu) = I_i(\nu)e^{-\alpha(\nu)L} \quad (1.1)$$

where ν is the optical frequency, $\alpha(\nu)$ represents the frequency-dependent absorption coefficient, and L is the absorption length. According to this exponential relation, the frequency-dependent absorption coefficient $\alpha(\nu)$ can be directly extracted by inverting the formula:

$$\alpha(\nu) = -\frac{1}{L} \ln \left(\frac{I_t(\nu)}{I_i(\nu)} \right) \quad (1.2)$$

This expression demonstrates that the extracted absorption coefficient $\alpha(\nu)$, which carries the fundamental information regarding the sample's concentration and molecular cross-section, is inversely proportional to the interaction path length L for a given attenuation ratio [3].

Traditionally, while dispersion spectroscopy represents the standard for broadband molecular sensing in the visible and ultraviolet ranges, Fourier Transform Infrared (FTIR) spectroscopy has been the gold standard for broadband molecular spectroscopy in the Infrared, enabling the collection of extensive molecular databases over the last four decades. However, conventional FTIR often relies on incoherent and weak light sources, which limits its sensitivity and requires long acquisition times, often limited by the mechanical movement of a mirror in the setup.




Figure 1.1: Principle of molecular spectroscopy using an optical frequency comb: the source spectrum is modified by the absorption lines of the sample. Adapted from [18].

The invention of the Optical Frequency Comb (OFC), recognized with the Nobel Prize in Physics in 2005 [25], has revolutionized this field by providing a link between measurable radio frequencies and ultrafast optical frequencies. An OFC acts as a "high-precision

ruler" in the frequency domain, offering a unique combination of broad spectral bandwidth and the high brightness and coherence of a laser. As illustrated in Fig. 1.1, a frequency comb can simultaneously excite a vast number of molecular transitions, allowing for the recording of complex spectra with unprecedented accuracy and significantly shorter measurement times compared to traditional Fourier transform spectrometers. This technology has not only improved time-keeping and fundamental metrology but has also become an indispensable tool for molecular fingerprinting, overcoming the speed and resolution trade-offs of traditional spectroscopic methods [11].

1.2. Optical Frequency Combs

1.2.1. Mode-locking

The most established method to generate an Optical Frequency Comb (OFC) relies on the mode-locking technique within a laser cavity. In a standard laser, multiple longitudinal modes oscillate simultaneously but with random phases, leading to a continuous-wave output with erratic intensity fluctuations. Mode-locking is the process of establishing a strict and constant phase relationship between all these oscillating modes. When this synchronization occurs, the modes interfere constructively at regular intervals, concentrating the laser's energy into a single, high-peak-power pulse that circulates inside the resonator.

The immediate physical result of this regime is a train of pulses emitted at a regular interval T_R , which corresponds to the time the pulse takes to complete one round-trip in the cavity, as illustrated in Fig. 1.2.




Figure 1.2: Time-domain representation of a pulse train. The electric field $E(t)$ and time t are expressed in arbitrary units (a.u.). Adapted from [18].

To provide a rigorous mathematical foundation for this phenomenon, we consider a laser resonator of length L . The electromagnetic field within the cavity must satisfy the fol-

lowing boundary conditions:

$$E(0) = 0 \quad E(L) = 0 \quad (1.3)$$

Recalling the associated wavevector formula $k_n = 2\pi\nu_n/c_n$ and applying the resonator boundary conditions, we find:

$$\sin(k_n L) = 0, \quad k_n L = n\pi, \quad n \in \mathbb{Z} \quad (1.4)$$

Therefore, only a set of specific electromagnetic frequencies ν_n are allowed in the resonator:

$$\nu_n = n \frac{c_n}{2L} \quad (1.5)$$

$$\nu_n = \frac{n}{T_R}, \quad T_R = \frac{2L}{c_n} \quad (1.6)$$

where c_n is the speed of light in the medium and T_R is the Round Trip Time. For a single mode, the electric field as a function of time can be written as:

$$E_n(t) = A_n \cos(2\pi\nu_n t + \phi_n) = \frac{A_n}{2} e^{i(2\pi\nu_n t + \phi_n)} + \frac{A_n}{2} e^{-i(2\pi\nu_n t + \phi_n)} \quad (1.7)$$

In the spectral domain, this corresponds to two Dirac delta functions centered at $\pm\nu_n$. Therefore if all the modes are taken into consideration, a frequency comb is obtained (the resulting picture indeed resembles a comb). In spectral domain only the modes amplified by the gain medium will oscillate. If $\tilde{G}(\nu)$ is the Fourier transform of the gain line centered around the frequency ν_0 and $g(t)$ is the gain envelope, the overall emission of the laser will be:

$$\tilde{E}(\nu) = \tilde{G}(\nu) \sum_{n=-\infty}^{+\infty} [\delta(\nu - \nu_n) \cdot e^{i\phi_n}] \quad (1.8)$$

In the mode-locking regime, the simplest condition is $\phi_n = 0$, meaning the phases are no longer random. The electric field in the time domain is then obtained through the inverse Fourier transform:

$$E(t) = \mathcal{F}^{-1}\{\tilde{G}(\nu)\} * \mathcal{F}^{-1}\left\{\sum_{n=-\infty}^{+\infty} \delta\left(\nu - \frac{n}{T_R}\right)\right\} \quad (1.9)$$

Applying the standard Fourier analysis result where a periodic impulse train in the frequency domain transforms into a periodic impulse train in the time domain, we arrive at

the expression for the pulse train:

$$E(t) = T_R \sum_{n=-\infty}^{+\infty} g(t - nT_R) \cos(2\pi\nu_0(t - nT_R)) \quad (1.10)$$

This result rigorously demonstrates that the mode-locking regime produces a periodic train of pulses in the temporal domain, which serves as the physical basis for the frequency comb structure.

1.2.2. Carrier-Envelope Phase and the Comb Equation

To provide a more accurate description of the frequency comb, the dispersion within the laser cavity must be taken into account. In a real medium, the phase velocity and the group velocity are different. Consequently, the wave vector $k(\nu)$ can be rewritten using a first-order Taylor expansion around the central frequency ν_0 :

$$k(\nu) = k(\nu_0) + \frac{2\pi(\nu - \nu_0)}{c_g(\nu_0)} = \frac{2\pi\nu_0}{c_n(\nu_0)} + \frac{2\pi(\nu - \nu_0)}{c_g(\nu_0)} \quad (1.11)$$

where c_n is the phase velocity (the speed of the carrier wave in the medium) and c_g is the group velocity (the speed of the pulse envelope in the medium). By recalling the resonator boundary conditions and multiplying both sides of the equation by the round-trip length $2L$, we obtain:

$$\nu_0 \frac{2L}{c_n} + \nu_n \frac{2L}{c_g} - \nu_0 \frac{2L}{c_g} = n \quad (1.12)$$

We can now define $\tau_{ph} = 2L/c_n$ as the phase delay and $\tau_g = T_R = 2L/c_g$ as the group delay accumulated by the pulse after a single round-trip. Using these definitions, the relationship between the carrier and the envelope is governed by the carrier-envelope phase shift $\Delta\phi_{CEP}$:

$$\Delta\phi_{CEP} = 2\pi\nu_0 (\tau_g - \tau_{ph}) \quad (1.13)$$

This phase shift represents the "slip" of the carrier wave relative to the envelope peak for each pulse in the train. By solving the previous resonance equation for ν_n , we find:

$$\nu_n = \frac{n}{T_R} + \frac{\Delta\phi_{CEP}}{2\pi T_R} \quad (1.14)$$

To express this in the final standard form, we introduce the repetition rate $f_{rep} = 1/T_R$ and the carrier-envelope offset frequency $f_{ceo} = \Delta\phi_{CEP}/(2\pi T_R)$. This leads directly to

the fundamental comb equation:

$$\nu_n = f_{ceo} + n f_{rep} \quad (1.15)$$

This crucial result demonstrates that every optical frequency ν_n in the comb is strictly determined by two radio-frequency parameters, f_{rep} and f_{ceo} , effectively linking the optical domain to measurable electronic signals.




Figure 1.3: Overview of an optical frequency comb. The top panel shows the pulse train in the time domain, while the bottom panel illustrates the corresponding teeth in the frequency domain, shifted by the offset frequency f_{ceo} . Adapted from [18].

1.3. Chirped Pulse Amplification (CPA) and Supercontinuum (SC) Generation

In the previous section, the theoretical nature of the optical frequency comb was defined, and its generation via a mode-locked laser was described. However, the frequency comb produced directly by a laser oscillator cannot be used for high-precision spectroscopy without further modification. Typically, the pulses emerging from the cavity possess low energy (on the order of picojoules to nanojoules, pJ–nJ), and their spectral bandwidth is strictly limited to tens of nanometers by the gain medium of the laser.

To overcome these practical limitations, the pulse train must undergo two fundamental transformations. First, the pulse energy must be increased without inducing catastrophic damage to the optical components; this is achieved through Chirped Pulse Amplification (CPA), a technique that manages peak power by stretching and subsequently compressing the pulses. Second, the narrow laser spectrum must be dramatically broadened to cover a wide range of frequencies, often exceeding an octave. This process, known as Supercontinuum (SC) generation, is essential both for broadband molecular sensing and for the experimental stabilization of the comb parameters.

In the following sections, the physical principles and the experimental stages of these two fundamental processes are discussed.

1.3.1. Chirped Pulse Amplification: Managing Peak Power

The pulse train generated by a mode-locked laser typically consists of pulses with durations in the femtosecond range. While these pulses are characterized by high stability and coherence, their energy at the oscillator output is often too low for many practical applications. To bridge this requirement, an amplification stage is necessary.

However, amplifying ultrashort pulses directly presents a significant physical challenge: because the energy is concentrated in such a brief time interval, the resulting peak power is extremely high. If these pulses were injected directly into an optical amplifier, the intensity would reach levels capable of inducing catastrophic effects:

- **Optical Damage:** The peak intensity can exceed the damage threshold of the optical fibers or the gain medium, leading to the permanent destruction of the components.
- **Non-linear Distortions:** Even below the damage threshold, high intensities trigger unwanted non-linear effects, such as Self-Phase Modulation (SPM), which uncontrollably distort the pulse's temporal shape and spectral phase.




Figure 1.4: Schematic representation of the Chirped Pulse Amplification (CPA) technique, illustrating the stretching, amplification, and compression stages of the pulse. Adapted from (Schimpf, 2010).

This is why the chirped pulse amplification (CPA) technique is necessary. The standard implementation of a CPA system follows a three-step sequence:

1. **Stretching (DCF - Dispersion Compensating Fiber):** The pulses are first injected into a dispersive delay line, such as a Dispersion Compensating Fiber (DCF). Due to the high chromatic positive dispersion, different spectral components travel at different velocities, inducing a positive chirp where the lower frequencies ("red") travel faster than the higher ones ("blue"). As a result, the pulse is stretched in time, effectively lowering its peak power by several orders of magnitude.
2. **Amplification (EDFA - Erbium-Doped Fiber Amplifier):** The stretched, low-intensity pulse is then fed into an amplifier, typically an Erbium-Doped Fiber Amplifier (EDFA) for systems operating in the 1550 nm region. Since the pulse's energy is spread over a much longer time interval, the medium can safely boost the power without reaching the non-linear threshold or damaging the fiber core.
3. **Compression (Grating Compressor):** Once the desired energy level is reached, the pulse must be returned to its ultrashort duration. This is achieved using a Grating Compressor, which provides the opposite dispersion (negative dispersion) to the stretcher. By forcing the spectral components that were previously "ahead" to take a shorter path, the compressor realigns them in time, "squeezing" the pulse back to its original femtosecond width.

The result of this process is a high-energy, ultrashort pulse that maintains the coherence of the original source while reaching the peak intensities necessary for subsequent stages.

1.3.2. Supercontinuum Generation: Spectral Broadening

Supercontinuum (SC) generation is the fundamental process that transforms a narrow-band frequency comb into an ultra-broadband source. While the pulses emerging from the CPA stage possess high energy, their spectral width is typically limited by the gain medium of the laser. Usually, it is necessary to achieve a bandwidth that spans an entire octave (meaning that the highest frequency component of the spectrum is at least twice the lowest frequency). This broad spectral coverage is essential for two primary reasons: it enables broadband measurements across a vast range of molecular transitions and, crucially, it provides the spectral overlap required for the $f - 2f$ self-referencing technique, that will be discussed in more detail in the next paragraph (reference al paragraph dolo). In modern setups, this is achieved using Highly Non-Linear Fibers (HNLF) or integrated nanophotonic waveguides, which are specifically engineered to enhance non-linear interactions.




Figure 1.5: Example of supercontinuum generation through numerically simulated spectral broadening in a standard telecom fiber. The simulation shows the evolution of 1 ps pulses with 50 nJ energy as they propagate through the non-linear medium. Adapted from (RP Photonics).

The initial stage of spectral broadening is driven by the Optical Kerr Effect. In this regime, the refractive index of the medium is no longer constant but becomes dependent

on the instantaneous intensity of the pulse, following the relation:

$$n(I) = n_0 + n_2 I(t) \quad (1.16)$$

where n_0 is the linear refractive index and n_2 is the non-linear index (Kerr coefficient). Because the laser pulse has a specific temporal profile (typically Gaussian), its center experiences a higher refractive index compared to its edges, causing it to travel more slowly. This intensity-dependent phase shift, known as Self-Phase Modulation (SPM), generates new frequency components, effectively stretching the spectrum.

However, to achieve extreme broadening exceeding an octave, SPM alone is not sufficient. At this stage, the propagation enters a more complex regime governed by the Generalized Non-Linear Schrödinger Equation (GNLS). The evolution of the pulse envelope $A(z, t)$ along the propagation coordinate z is mathematically described by:

$$\frac{\partial A}{\partial z} + \frac{\alpha}{2} A - \sum_{n \geq 2} \frac{i^{n+1} \beta_n}{n!} \frac{\partial^n A}{\partial t^n} = i\gamma \left(|A|^2 A + \frac{i}{2\pi\nu_0} \frac{\partial}{\partial t} (|A|^2 A) - T_R A \frac{\partial |A|^2}{\partial t} \right) \quad (1.17)$$

where α is the fiber attenuation, β_n represents the higher-order dispersion terms, γ is the non-linear coefficient and T_R is the Raman scattering contribution. Under this framework, the dynamics are dictated by the interplay between non-linear effects and the medium's Group Velocity Dispersion (GVD). In the anomalous dispersion regime ($\beta_2 < 0$), where lower frequencies travel slower than higher ones, the GVD can effectively compensate for the chirping induced by the Kerr effect. This balance leads to the formation of optical solitons: stable pulses that maintain their shape over long distances while undergoing further non-linear transformations. These solitons, through interactions with higher-order dispersion and Raman scattering within the medium, trigger additional frequency shifts that fill the spectral gaps and extend the comb into the mid-infrared and visible regions. The final result is a coherent frequency comb that provides the necessary spectral coverage for high-precision molecular sensing and the subsequent stabilization of the comb's parameters.

1.4. Frequency Comb Stabilization

The theoretical model of an optical frequency comb defines the frequency of each tooth through the relation $f_n = n f_{rep} + f_{ceo}$. In an ideal scenario, the repetition rate (f_{rep}) and the carrier-envelope offset frequency (f_{ceo}) would remain perfectly constant, providing a rigid and stable grid of frequencies. However, in any real experimental setup, these parameters are extremely sensitive to environmental perturbations. Fluctuations in tem-

perature, mechanical vibrations, or acoustic noise induce continuous changes in the cavity length and the refractive index of the gain medium, causing the comb's frequencies and phases to drift and jitter.

To achieve high-precision measurements and absolute frequency accuracy, these two degrees of freedom must be actively stabilized. This stabilization is fundamental for two primary reasons:

- **Enhancement of Signal-to-Noise Ratio (SNR):** A stable comb allows for longer integration times and coherent averaging of the signal, which is essential to distinguish weak absorption features from the background noise.
- **Prevention of Spectral Line Broadening:** Without active control, the fluctuations of f_{rep} and f_{ceo} cause a temporal "smearing" of the comb teeth during the detection process. This instability leads to an artificial broadening of the measured absorption lines, which limits the effective resolution of the spectrometer and prevents the system from resolving narrow molecular transitions.

The most traditional and robust approach to achieve this stability is through a phase-locking architecture, where the laser's internal parameters are continuously adjusted to match a stable external reference [10].

1.4.1. Phase-Locked Loop (PLL)

The most established method for stabilizing an optical frequency comb is the implementation of a Phase-Locked Loop (PLL). This approach treats the instabilities of the laser as dynamic errors that must be corrected continuously by physically adjusting the laser parameters.

The architecture of a full feedback loop is based on a classic closed-loop control system. The goal is to "lock" the optical frequencies of the comb to a stable external reference (usually a narrow optical reference laser). As illustrated in Fig.1.6, the process can be decomposed into three main blocks:

- **Error Detection:** A portion of the laser output is diverted from the main beam to measure the necessary parameters. In particular to stabilize f_{rep} , a fast photodiode detects the pulse repetition rate. While to stabilize f_{CEO} , an $f - 2f$ interferometer (discussed in Section 2.2.2) is typically used to reveal the offset frequency.
- **Comparison and Control:** The measured frequencies are compared to a stable reference. The difference between the measured value and the reference generates an error signal. This signal is then processed by an electronic controller, typically a

Proportional-Integral-Derivative (PID) controller, to determine the precise adjustment needed to maintain the laser's stability.

- **Actuation (Physical Correction):** The controller sends a correction voltage to physical actuators inside the laser cavity. These typically include:
 - **Piezoelectric Transducers:** Which change the cavity length to adjust f_{rep} .
 - **Pump Power Modulation:** Modulating the pump current is possible to tune f_{CEO} .




Figure 1.6: Schematic diagram of the complete dual-parameter active feedback loop for frequency comb stabilization. The error signals for both f_{rep} and f_{ceo} are processed by electronic controllers to generate real-time correction voltages. These voltages drive physical actuators within the cavity for cavity length adjustments (f_{rep}) and a pump power modulation system for carrier-envelope offset control (f_{ceo}).

An example of this technique is the work of Coddington, Swann, and Newbury [23], where a robust hardware-based feedback loop was employed to achieve very high stability. This rigorous synchronization resulted in relative subhertz linewidths.

1.4.2. The $f-2f$ Interferometer: Measuring the Carrier-Envelope Offset

Stabilizing the optical frequency comb requires the simultaneous measurement and control of two fundamental parameters: the repetition rate f_{rep} and the carrier-envelope offset frequency f_{ceo} . The repetition rate f_{rep} can be easily monitored by sending a portion of

the laser output onto a fast photodiode. Conversely, measuring f_{ceo} is significantly more complex, as it originates from the pulse-to-pulse carrier-envelope phase shift $\Delta\phi_{CEP}$. To isolate and measure this parameter, an $f - 2f$ self-referencing interferometer is required. The core principle of an $f - 2f$ interferometer relies on non-linear optics, specifically the Second Harmonic Generation (SHG) process. As illustrated in the block diagram in Fig. 1.7, the initial laser beam is emerging from the supercontinuum generation stage, therefore possessing an octave-spanning spectrum. This broad spectrum simultaneously contains a low-frequency region (long wavelengths, λ_{longer}) and a high-frequency region (short wavelengths, $\lambda_{shorter}$).




Figure 1.7: Schematic block diagram of the $f - 2f$ interferometer setup using a dichroic mirror to split the octave-spanning spectrum into longer-wavelength and shorter-wavelength arms.

To implement the interferometer, the beam is directed onto a dichroic mirror, which splits the spectrum into two distinct paths based on wavelength:

- The long-wavelength components (λ_{longer}) are routed into the first arm, The corresponding frequency can be written as follows:

$$\nu_m = f_{ceo} + m f_{rep} \quad (1.18)$$

This beam passes through a non-linear crystal to undergo Second Harmonic Generation. The frequency of the newly generated second-harmonic (SH) mode is doubled, yielding:

$$\nu_{SHG} = 2\nu_1 = 2(f_{ceo} + m f_{rep}) = 2f_{ceo} + 2m f_{rep} \quad (1.19)$$

- The short-wavelength components ($\lambda_{shorter}$) are directed into the second arm, bypassing the crystal. This path retains the original fundamental frequency of a comb

with index n :

$$\nu_n = f_{ceo} + n f_{rep} \quad (1.20)$$

Subsequently, a beam splitter is employed to spatially recombine the frequency-doubled beam ($2\nu_m$) and the other beam (ν_n).

As displayed in the spectral representation in Fig. 1.8, there is a region where these two waves cover the same high-frequency spectral region ($\lambda_{shorter}$), they overlap on a fast photodetector to perform a heterodyne measurement. The detector registers an intensity modulation oscillating at the beat frequency:

$$\nu_{beat} = 2\nu_m - \nu_n = (2f_{ceo} + 2mf_{rep}) - (f_{ceo} + nf_{rep}) = f_{ceo} + (2m - n)f_{rep} \quad (1.21)$$

As a result of this heterodyne beating, when the condition $2m = n$ is satisfied, the repetition rate f_{rep} cancels out perfectly, successfully isolating f_{ceo} as a clean radiofrequency signal.




Figure 1.8: Spectral representation of the $f - 2f$ self-referencing technique for the two modes where the condition $2m = n$ is satisfied. The fundamental comb (red) contains modes ν_1 and ν_2 . The mode ν_1 is frequency-doubled via SHG (SH-OFC, blue). The heterodyne beating between $2\nu_1$ and the native fundamental mode ν_2 isolates the carrier-envelope offset frequency f_{ceo} . Adapted from Keller, U. (2021). Optical Frequency Comb from Modelocked Lasers. In: Ultrafast Lasers. Graduate Texts in Physics. Springer.

1.5. Introduction to Detection Techniques

The previous sections of this chapter delineated the fundamental theory of Optical Frequency Combs (OFCs), their stabilization mechanisms and their revolutionary role as

high-precision "frequency rulers." However, to fully exploit the potential of these sources in molecular spectroscopy, a detection system capable of resolving individual comb lines and extracting the encoded absorption information is required. Historically the two primary techniques used for broadband spectroscopy with frequency combs were Fourier Transform Infrared Spectroscopy (FTIR) and Virtually Imaged Phased Array (VIPA) spectroscopy.

Originally developed for incoherent light sources, the FTIR technique is based on the Michelson interferometer architecture but was subsequently adapted for use with frequency combs to drastically improve the signal-to-noise ratio. The operating principle relies on splitting the light beam into two arms: one of fixed length and one modulated by a moving mirror. The recombination of these beams generates an interferogram in time. Through a Fourier transform, the optical spectrum of the source can be reconstructed. Although integration with OFCs allows for accurate frequency calibration, FTIR remains constrained by mechanics: the spectral resolution depends on the maximum displacement of the mirror. To resolve individual lines of a comb at high speeds, the mirror would need to travel kilometric distances in extremely short timeframes. This physical limit restricts acquisition speed to the millisecond range and makes the system highly sensitive to environmental vibrations.

In parallel, VIPA technology was specifically developed to handle the discrete structure




Figure 1.9: Scheme of a FTIR Spectrometer.

of frequency combs through a spatial dispersion strategy. The core of the system is an etalon, light undergoes multiple internal reflections and emerges with very high angular dispersion. To prevent overlapping of different spectral orders, the VIPA is typically combined with a secondary dispersive element (such as a standard diffraction grating) oriented perpendicularly. This configuration generates a two-dimensional spectral map where each individual comb line is spatially separated and projected onto a 2D sensor,

such as a CCD or CMOS camera. While this technique allows for simultaneous and fast measurements, it requires a bulky optical setup that is complex to align. Furthermore, the final resolution is limited by the physical pixel size of the camera and optical aberrations, often preventing the system from reaching the theoretical resolution offered by the comb itself.

Both techniques, while effective, rely on external components—moving mirrors or com-




Figure 1.10: Scheme of a VIPA Spectrometer.

plex dispersive optics that impose a rigid trade-off between acquisition speed, resolution, and system robustness. The Dual-Comb Spectroscopy (DCS) technique represents the overcoming of these limitations. By employing two combs with slightly different repetition frequencies, DCS realizes a static version of the Michelson interferometer. In this scheme, the mechanical scanning of the mirror is replaced by the multi-heterodyne beating between the two laser sources. The slight offset in their repetition rates causes the pulses to periodically "slide" over each other in the time domain, effectively acting as a virtual scanning mirror. This allows the mapping of optical frequencies directly into the radio-frequency (RF) domain, eliminating the need for moving parts or complex dispersive apparatus, and combining the advantages of multiplex spectroscopy with unprecedented speed and compactness. Because this multi-heterodyne approach represents the core methodology of this thesis, its physical and mathematical framework, experimental architectures, and self-correction schemes will be investigated in the next chapters.

2 | Dual-Comb Spectroscopy

2.1. Dual Comb Spectroscopy: the principle

This section introduces the fundamental operating principles of Dual-Comb Spectroscopy (DCS), focusing on its ability to perform high-resolution optical measurements without moving parts. First, the mechanism of multi-heterodyne detection is detailed. Then, the resulting time-domain signal (the interferogram) is analyzed to show how optical information is encoded into periodic bursts. Finally, the discussion addresses the coherent down-conversion of the spectrum into the radio-frequency domain, outlining the strict constraints imposed by the Nyquist-Shannon sampling theorem to avoid spectral aliasing.

2.1.1. Multi-heterodyne detection

Dual Comb Spectroscopy is the main topic of this work and is a technique that exploits two optical frequency combs to reconstruct the absorption features of molecular samples [3].




Figure 2.1: Dual Comb Spectroscopy general scheme.

Consider a configuration in which one comb (comb 1) is propagated through the sample and is then combined with the second comb (comb 2) before reaching the detector. The optical signal from comb 1 undergoes amplitude and phase modifications due to the molec-

ular absorption features of the sample. However, when using just one optical frequency comb, the information it encodes cannot be directly detected by a photodiode, as the oscillations of the optical carrier occur at frequencies far beyond the detector bandwidth. This is where the second comb plays a crucial role. The two combs are nearly identical, except for a small difference in their repetition frequencies, $\Delta f_{rep} = f_{rep,1} - f_{rep,2}$, which lies in the radio-frequency range. When the two combs are combined, they interfere and generate a series of beatings at frequencies on the order of Δf_{rep} . These beatings can be detected by a photodiode, effectively mapping the optical spectrum into the radio-frequency domain. In this way, the absorption features of the sample are retrieved from the measured signal.

Now the formal description [3] of the interaction between the two combs can be considered. The electric field of a pulse train can be written in complex exponential form as the superposition of the n monochromatic plane waves that form it. The electric field of the k -th frequency comb is:

$$E_k(t) = \sum_n A_{n,k} e^{-i2\pi f_{n,k}t} + \text{c.c.} \quad (2.1)$$

where $A_{n,k}$ and $f_{n,k}$ are the amplitude and frequency of the n -th mode, respectively. The mode frequencies are given by:

$$f_{n,k} = f_{0,k} + n f_{rep,k} \quad (2.2)$$

where $f_{0,k}$ is the carrier-envelope offset frequency of the k -th comb.

In the specific configuration of this work, only comb 1 passes through the sample, undergoing amplitude attenuation α_n and phase shift ϕ_n . Comb 2, instead, acts as a local oscillator. The electric fields at the detector are:

$$E_1(t) = \sum_n A_{n,1} e^{-\alpha_n} e^{-i\phi_n} e^{-i2\pi f_{n,1}t} + \text{c.c.} \quad (2.3)$$

$$E_2(t) = \sum_n A_{n,2} e^{-i2\pi f_{n,2}t} + \text{c.c.} \quad (2.4)$$

The total electric field is the superposition $E(t) = E_1(t) + E_2(t)$. The photodiode measures the intensity $I(t) \propto E(t)E^*(t)$, giving rise to the interferogram. Considering only the cross-terms (heterodyne beatings) between modes of the same order n , one obtains:

$$I(t) \propto \sum_n \left[A_{n,1}^2 e^{-2\alpha_n} + A_{n,2}^2 + 2A_{n,1}A_{n,2} e^{-\alpha_n} \cos \left(2\pi(f_{n,1} - f_{n,2})t + \phi_n \right) \right] \quad (2.5)$$

The frequency difference between corresponding modes is:

$$\begin{aligned}
 f_{n,1} - f_{n,2} &= (f_{0,1} + n f_{\text{rep},1}) - (f_{0,2} + n f_{\text{rep},2}) \\
 &= (f_{0,1} - f_{0,2}) + n(f_{\text{rep},1} - f_{\text{rep},2}) \\
 &= \Delta f_0 + n \Delta f_{\text{rep}}
 \end{aligned} \tag{2.6}$$

where $\Delta f_0 = f_{0,1} - f_{0,2}$. Thus, the relevant term of the interferogram can be expressed as:

$$I(t) \propto \sum_n 2A_{n,1}A_{n,2}e^{-\alpha_n} \cos(2\pi(\Delta f_0 + n\Delta f_{\text{rep}})t + \phi_n) \tag{2.7}$$

This term contains the spectroscopic information mapped at radio frequencies:

$$f_{\text{RF},n} = \Delta f_0 + n\Delta f_{\text{rep}} \tag{2.8}$$

2.1.2. The Time-Domain Signal: The Interferogram

In the time domain, the resulting signal $I(t)$ is a periodic sequence of high-intensity interference packets, commonly referred to as *bursts*. Each burst occurs when the pulses from the two combs perfectly overlap at the detector.




Figure 2.2: Interferograms (bursts) observed at the photodiode, resulting from the interference between pulses from the two lasers.

Since the two pulse trains have slightly different repetition periods ($T_{\text{rep},1}$ and $T_{\text{rep},2}$), they create a new temporal overlap every $T_r = 1/\Delta f_{\text{rep}}$.

By computing the Fourier Transform of this train of bursts, the radio-frequency comb spectrum is retrieved. The peaks at frequencies $f_{\text{RF},n}$ are modulated in amplitude and phase by the interaction with the sample, allowing for the reconstruction of its spectro-

scopic features. Knowing the values of $f_{0,k}$, $f_{rep,k}$, and Δf_{rep} , it is then possible to map these RF frequencies back to the absolute optical frequency scale.

The cross terms involving modes of different order give rise to signals at higher frequencies. Since they do not provide additional spectroscopic information, they are typically removed by means of low-pass filtering, which retains only the interference between modes of the same order [3]. The beating phenomenon of the two comb sources is detailed more clearly in Fig. 2.3.




Figure 2.3: Spectrum of the two optical frequency combs (three modes are shown for simplicity). The corresponding beatings spectrum is reported below.

Beatings arising from the interference between modes of the same order are highlighted in red.

2.1.3. Down-conversion and the Nyquist-Shannon sampling theorem

The defining feature of Dual-Comb Spectroscopy is its ability to perform a coherent down-conversion of the optical frequency information into the radio-frequency (RF) domain.

This process is governed by a scaling factor a , which defines the relationship between the optical and RF frequency scales:

$$a = \frac{\Delta f_{rep}}{f_{rep,1}} \quad (2.9)$$

where $f_{rep,1}$ is the repetition rate of the probing comb and Δf_{rep} is the detuning between the two sources.

However, this mapping is strictly constrained by the Nyquist-Shannon sampling theorem [22]. In a DCS system, the sampling rate is effectively defined by the repetition frequency f_{rep} of the pulses. To ensure an unambiguous mapping and avoid spectral aliasing, meaning the folding and overlapping of spectral replicas, the entire down-converted RF spectrum must be contained within the first Nyquist zone, meaning its bandwidth must not exceed $f_{rep,1}/2$.

In practical experimental conditions, if $\Delta\nu$ is the optical bandwidth to be mapped, the RF spectrum can be seen as a collection of N_{bw} modes, where $N_{bw} = \Delta\nu/f_{rep,1}$. To avoid aliasing, the total span of these modes in the RF domain ($N_{bw} \cdot \Delta f_{rep}$) must satisfy:

$$N_{bw} \Delta f_{rep} \leq \frac{f_{rep,1}}{2} \quad (2.10)$$

By substituting the expression for N_{bw} , we derive the fundamental design constraint for the difference in repetition rates:

$$\frac{\Delta\nu}{f_{rep,1}} \Delta f_{rep} \leq \frac{f_{rep,1}}{2} \implies \Delta f_{rep} \leq \frac{f_{rep,1}^2}{2\Delta\nu} \quad (2.11)$$

This inequality defines the maximum allowable detuning $\Delta f_{rep,max}$. If the difference in repetition rates exceeds this threshold, the RF comb "overflows" the $f_{rep}/2$ limit and aliases back, making it difficult to correctly retrieve the original optical information [3]. The dependence of this constraint on the repetition rate highlights the advantages of using high- f_{rep} sources [8]. Specifically, increasing the repetition rate benefits the system in two primary ways:

1. **Expansion of the Nyquist Zone:** Since the first Nyquist zone extends up to $f_{rep}/2$, a higher repetition frequency provides a broader RF bandwidth. This allows for the accommodation of larger optical spans without incurring aliasing.
2. **Relaxation of the Detuning Constraint:** As shown in Equation 2.11, the maximum allowable detuning $\Delta f_{rep,max}$ scales quadratically with $f_{rep,1}$. Doubling the repetition rate quadruples the limit for Δf_{rep} , significantly easing the experimental challenge of maintaining a small but stable difference in repetition rates while still

maximizing the acquisition speed.

In the next section, it will be discussed how to properly manage these sampling constraints and address the further complications that arise when moving to a more realistic description of the system.

2.2. Impact of the frequency noise on dual-comb interferograms

This section explores the role of mutual coherence in Dual-Comb Spectroscopy, analyzing how system instabilities impact experimental measurements. First, the physical origins of timing and phase jitter are discussed, illustrating their effects on the interferogram. Next, coherent averaging is addressed, highlighting the signal-to-noise ratio (SNR) constraints imposed by these fluctuations. Finally, an overview of traditional hardware stabilization techniques is provided.

2.2.1. Origins of Instabilities

The practical implementation of Dual-Comb Spectroscopy (DCS) fundamentally depends on the ability to maintain mutual coherence between the two optical frequency combs. As seen in the previous Section 1.4, the repetition frequency f_{rep} and the carrier-envelope offset frequency f_{CEO} of a frequency comb are not perfectly stable constants in a real-world scenario. These instabilities are problematic in single combs and become even more significant in dual comb spectroscopy. Here, these instabilities affect the combination of the two lasers, with the resulting effects visible in the output interferogram (burst) train. These instabilities manifest as due to two distinct effects: timing jitter and phase jitter. Timing jitter originates from fluctuations in the repetition rate f_{rep} . In a dual-comb system, any residual instability in the difference between the repetition rates, Δf_{rep} , translates directly into an uncertainty in the repetition period of the interferograms, $T_r = 1/\Delta f_{rep}$. This can be visualized as a random variation in the time delay between consecutive bursts. Instead of appearing at perfectly equidistant intervals, the bursts arrive with a slight temporal displacement. As shown in Fig. (2.4), this timing jitter causes a distortion of the temporal mapping, making it difficult to define a stable time axis for the measurement.

Phase jitter, on the other hand, has a much more significant impact on the spectroscopic signal and arises from fluctuations in the carrier-envelope offset frequency f_{ceo} . While timing jitter affects the arrival time of the pulse envelope, phase jitter governs the stability




Figure 2.4: Visual representation of timing jitter in a dual-comb interferogram train. The lower trace illustrates the ideal, unperturbed scenario where bursts occur at perfectly equidistant intervals with a constant repetition period $T_r = 1/\Delta f_{rep}$. The upper trace demonstrates the effect of timing jitter, where fluctuations in Δf_{rep} lead to random variations in the arrival time of consecutive bursts, causing a dynamic distortion of the temporal mapping axis.

of the carrier oscillations under that envelope. To isolate and clarify this phenomenon, an ideal harmonic radio-frequency (RF) comb can be considered, where the offset frequency is fixed to zero ($f_0 = 0$). In this scenario, any misalignment between the carrier crests and the envelope center does not stem from a static frequency offset. Instead, as a consequence of the residual phase noise in a free-running dual-comb system, the relative phase between the two lasers continuously drifts over time. As illustrated in Fig. (2.5A), this instability causes the carrier waves of consecutive interferograms to randomly "slide" or shift beneath the envelope from one burst to the next. Conversely, when active phase stabilization is successfully achieved, as depicted in Fig. (2.5B), the relative phase relation is rigidly locked, freezing the interference fringes in a fixed position across the entire burst train.




Figure 2.5: Visual representation of phase jitter in dual-comb spectroscopy bursts, assuming an underlying harmonic RF comb configuration ($f_0 = 0$). (A) In the presence of phase jitter, the carrier oscillations randomly shift beneath the pink envelope. (B) In the phase-stabilized scenario, the phase relation is rigidly locked, ensuring that the fringes remain perfectly stationary and identical from burst to burst.

The combination of these two instabilities leads to a rapid decay of the mutual coherence between the two lasers. Without a synchronization mechanism, the phase relationship between the optical carrier and the down-converted RF signal is lost. This lack of stability represents the primary obstacle to achieving high-precision measurements, as it prevents the use of coherent averaging to improve the signal-to-noise ratio, a challenge that will be addressed in the following section.

2.2.2. Coherent Averaging and SNR Constraints

The necessity of addressing the instabilities described in the previous section is deeply rooted in the requirement for high-precision measurements. In Dual-Comb Spectroscopy, as in many spectroscopic techniques, the raw signal is often buried under various sources of noise, such as detector thermal noise or laser intensity fluctuations. To improve the quality of the data and resolve weak molecular transitions, it is fundamental to increase the Signal-to-Noise Ratio (SNR) through a process known as coherent averaging.

Unlike incoherent averaging, which sums the power spectra of the signals, coherent averaging is performed by summing multiple successive interferograms (bursts) directly in the time domain. The power of this technique lies in the different statistical evolution of a deterministic signal compared to stochastic noise. When N successive interferograms are summed, the amplitude of the signal adds up linearly with N . In contrast, the noise is an uncorrelated random process; its fluctuations tend to cancel each other out over time,

and its cumulative amplitude (root mean square) grows only with the square root of N . This leads to the fundamental scaling law:

$$SNR \propto \frac{N}{\sqrt{N}} = \sqrt{N} \quad (2.12)$$

This relationship implies that the noise-to-signal ratio decreases as $1/\sqrt{N}$, allowing the detection of signals that would be completely invisible in a single-shot acquisition.

However, the effectiveness of this averaging is strictly limited by the mutual coherence time of the two lasers. As discussed in Sec.2.2.1, timing and phase jitter cause the pulses to shift and the carrier oscillations to "slide" from one burst to the next. If the integration time exceeds the period over which the lasers remain in phase, the summation becomes counterproductive. Instead of adding constructively, the offset bursts will interfere destructively, effectively washing out the spectroscopic information and leaving only a blurred noise floor. Therefore, implementing strategies to compensate for jitter is not just an optimization, but a strict prerequisite for high-resolution dual-comb measurements [4].

2.2.3. Techniques for Fluctuation Compensation in Dual-Comb Systems

As established in Section 2.2, the mitigation of timing and phase jitter is a requirement for performing coherent averaging and achieving high-precision results. Over the years, several methodologies have been implemented to synchronize the two combs. The most traditional approach to ensure mutual coherence is the implementation of a dual phase-locked loop (PLL) architecture, one for each comb. As discussed in Section 1.4 regarding the stabilization of a single comb, this method involves locking both f_{rep} and f_{CEO} of each laser to a stable external reference (such as an ultra-stable Continuous Wave (CW) laser [6]). While this method provides a very high degree of stability, it comes with significant drawbacks: it requires extremely fast and expensive electronics, it is highly sensitive to the surrounding environment, and the resulting experimental setup is bulky and difficult to transport.

A natural evolution toward more flexible systems is represented by feed-forward stabilization techniques [9]. These schemes allow the lasers to operate in a free-running state (or loosely stabilized), reducing the complexity of the locking electronics. However, to compensate for frequency fluctuations, they still require an independent and continuous measurement of the f_{CEO} and f_{rep} drifts for both sources.

Once these errors are measured, the system applies the necessary corrections via a "stabilization block" without closing a physical feedback loop back to the laser cavity. This

approach avoids the bandwidth limitations and potential instabilities connected to the presence of a loop, constrained by the finite response time of physical actuators. The stabilization block can be implemented in different ways:

- **Optical Feed-Forward:** Employs external modulators, such as Acousto-Optic (AOM) or Electro-Optic (EOM) modulators, to shift the frequency of the comb teeth in the optical domain.
- **Analog Electronic Feed-Forward:** Utilizes analog electronic circuitry (mixers, filters, and VCOs) to process the RF beatings and apply phase or frequency corrections directly to the signals before final acquisition [13].
- **Digital Feed-Forward:** Employs digital signal processors (DSP) or high-speed Field Programmable Gate Arrays (FPGA) to calculate and apply corrections to the digitized signals, offering high flexibility and precision.




Figure 2.6: Example of an optical feed-forward stabilization scheme using an Acousto-Optic Modulator (AOM). In this configuration, the relative fluctuations of f_{rep} and f_{CEO} are measured via auxiliary beat-notes, typically requiring gratings to isolate spectral portions and high-speed detectors. These error signals are processed by a control unit that drives an external AOM to shift the frequency of the slave comb in real-time, ensuring mutual coherence without affecting the laser cavity parameters [19].

A common requirement of all the methods described above is the need for auxiliary hardware to monitor the fluctuations of each comb separately. This typically involves f - $2f$ interferometers for f_{CEO} detection and high-speed photodetectors for f_{rep} monitoring, sig-

nificantly increasing the cost and footprint of the spectrometer.

In recent years, research has shifted toward leaner architectures, where the burden of stabilization is moved entirely to the software domain. The primary advantage of these solutions lies in the fact that the need for separate, independent measurements of individual laser parameters is completely eliminated. Instead, the raw, free-running interferograms are recorded directly, as the dual-comb signal itself inherently contains all the information required for its own stabilization. This information can be extracted numerically through sophisticated algorithms, which then automatically apply the necessary corrections to the signal [7]. This software-based approach, in all its various aspects, represents the core of the work presented in this thesis and will be extensively discussed in the following chapters.

2.3. Self-Correction algorithms

As previously mentioned, using a software-based stabilization [7] removes the need for independent physical measurements of f_{rep} and f_{ceo} , since the correction can be made directly by using information extracted from the interferogram itself. In this section, we delve into the details of the self-correction algorithm developed in this work. The approach described in this section is a post-processing algorithm, meaning that the data is first fully acquired and subsequently analyzed offline. The primary objective of this algorithm is to stabilize and "clean" the dual-comb signal by removing what were previously defined as timing and phase jitters (Section 2.2).

By shifting the stabilization to the software domain, the interferogram becomes its own reference, allowing for high-precision results without the complexity of active hardware locking. This self-referencing capability is possible as long as the phase jitter between consecutive bursts remains below 2π ; under this condition, the interferogram itself can be used as a reliable measure of the system's phase and frequency noise, enabling an unambiguous reconstruction of the phase evolution.

2.3.1. Post-Processing Self-Correction: Algorithmic Framework

To understand the effectiveness of the proposed approach, it is necessary to examine the underlying mathematical and physical logic of the correction framework. The algorithm is designed to treat timing and phase fluctuations as distinct phenomena, addressing each sequentially. The foundational step of the entire framework requires converting the experimental real-valued interferogram into its complex analytical signal via the Hilbert transform (B). This mathematical tool is essential as it grants direct access to both the

instantaneous amplitude (the envelope) and the instantaneous phase of the dual-comb signal, providing the necessary metrics to address both forms of noise.

First, the algorithm compensates for the frequency jitter by realigning the temporal arrival of the bursts onto a perfectly linear and equidistant time grid. By exploiting the extracted envelope, the exact positions of the temporal peaks are identified, effectively compensating for fluctuations in the repetition rate difference (Δf_{rep}). It is important to note that while this temporal realignment is a necessary preprocessing step, timing jitter typically hinders frequency stability much less than phase jitter, and this first stage does not alter the underlying carrier phase.

Once the temporal grid is uniform, the phase fluctuations can be isolated and corrected. The phase evolution is tracked across the burst train by sampling the phase values at the specific maxima of the interferograms. These discrete values are then used to construct a continuous phase correction vector, $\Phi(t)$, representing the time-dependent phase drift of the system. By multiplying the analytical signal by the inverse of this phase function, the algorithm achieves straightforward phase cancellation and stabilization:

$$s_{corr}(t) = s(t) \cdot e^{-i\Phi(t)} \quad (2.13)$$

where $s_{corr}(t)$ and $s(t)$ represent the corrected and original analytical signals, respectively. This procedure effectively stabilizes the carrier phase (associated with f_{ceo} fluctuations), ensuring the long-term mutual coherence required for high-resolution spectroscopy and resulting in significantly narrower comb lines.

2.4. Towards real-time self-correction

It has already been discussed in Section 2.2.3 that increasing the signal-to-noise ratio by means of hardware-based methods requires active feedback on the lasers via several complex optical and electronic components. It has therefore been pointed out that post-processing correction of the fluctuations in repetition rates and carrier-envelope offset frequencies is a leaner and less expensive alternative. This software-centric option implies fewer physical components, is universally applicable to different laser sources, and simplifies future commercialization.

However, traditional post-processing techniques also present distinct disadvantages. These self-correction algorithms operate offline on large, pre-recorded dual-comb interferogram datasets to perform coherent averaging. While computationally powerful, this approach is extremely time-consuming. Furthermore, the maximum achievable averaging time is strictly constrained by the physical memory capacity of the DAQ (Data Acquisition) sys-

tem and the computer exploited for the post-processing analysis.

The ideal solution to overcome these problems lies in a Real-Time phase correction framework. By integrating the correction directly into a Digital Signal Processing Hardware, such as a Radio Frequency System-on-Chip (RFSoc) featuring an embedded FPGA (Field Programmable Gate Array), the stabilization can be executed "on the fly". In this configuration, the algorithm processes the incoming data stream using high-speed, localized sliding-window mechanisms rather than digesting the entire experimental dataset at once. This approach fundamentally alters the data handling strategy, requiring a discrete implementation of the Hilbert transform tailored for hardware execution (as detailed in Sect. 2.4.1). By processing the data sequentially, the system drastically reduces its memory footprint and bypasses the need for massive data buffering or post-acquisition storage. This real-time mitigation of timing and phase jitters effectively removes the memory constraints imposed by standard PC architectures, enabling continuous acquisition, instantaneous data reduction, and indefinitely long averaging times [7].

2.4.1. Discrete Hilbert Transform

In the context of real-time stabilization for Dual Comb Spectroscopy (DCS), the transition from post-processing algorithms to live hardware implementation such as Field Programmable Gate Arrays (FPGAs) is necessary to handle instantly the large amount of data and with perfect timing. The hardware implementation necessitates a shift from continuous-time to discrete-time processing [15]. While the ideal Hilbert transform is defined in the continuous domain, an FPGA operates on sampled data at a fixed clock frequency f_s . Therefore, the analytical operator must be discretized to be compatible with digital signal processing (DSP) architectures.

Starting from the continuous-time definition derived in sect.B, the ideal Hilbert transfer function is given by:

$$H(\Omega) = -j\text{sgn}(\Omega) \quad (2.14)$$

where $\Omega = 2\pi f$ represents the continuous angular frequency. To adapt this to a discrete system, we consider a continuous exponential signal $x_c(t) = e^{j\Omega t}$ sampled at discrete intervals $t = nT_s$, where $T_s = 1/f_s$ is the sampling period. The resulting discrete sequence is:

$$x_D[n] = e^{j\Omega nT_s} = e^{j\omega n} \quad (2.15)$$

In this formulation, $\omega = \Omega T_s = 2\pi f/f_s$ represents the normalized discrete angular frequency. Due to the inherent periodicity of discrete-time systems, where $e^{j(\omega+2k\pi)n} = e^{j\omega n}$ for any integer $k \in \mathbb{Z}$, the discrete Hilbert transfer function $H_D(e^{j\omega})$ must be defined

within the fundamental interval $[-\pi, \pi]$ as follows:

$$H_D(e^{j\omega}) = \begin{cases} -j & 0 < \omega < \pi \\ +j & -\pi < \omega < 0 \\ 0 & \omega = 0, \pi \end{cases} \quad (2.16)$$

To find the impulse response $\hat{x}_D[n]$ of the discrete Hilbert transformer, we apply the Inverse Discrete-Time Fourier Transform (IDTFT):

$$\hat{x}_D[n] = \frac{1}{2\pi} \int_{-\pi}^{\pi} H_D(e^{j\omega}) e^{j\omega n} d\omega \quad (2.17)$$

By substituting the values of $H_D(e^{j\omega})$ and evaluating the integral over the two symmetric halves of the spectrum, the following result is obtained:

$$\hat{x}_D[n] = \frac{1}{2\pi} \left[\int_{-\pi}^0 j e^{j\omega n} d\omega + \int_0^{\pi} -j e^{j\omega n} d\omega \right] = \frac{1 - \cos(\pi n)}{\pi n}, \quad n \neq 0 \quad (2.18)$$

Evaluating the cosine term for different values of n , the final expression for the discrete impulse response is:

$$\hat{x}_D[n] = \begin{cases} \frac{2}{\pi n} & \text{if } n \text{ is odd} \\ 0 & \text{if } n \text{ is even} \\ 0 & n = 0 \end{cases} \quad (2.19)$$




Figure 2.7: Impulse response $\hat{x}_D[n]$ of the discrete Hilbert transformer.

This result is of particular importance for FPGA-based real-time processing. In the digital domain, the Hilbert Transform is no longer an ideal, instantaneous phase-shifting operator but is implemented as a Finite Impulse Response (FIR) filter. From a theoretical standpoint, applying the Hilbert Transform to a digitized interferogram $x[n]$ is equivalent to performing a discrete convolution between the input signal and the discrete impulse response $\hat{x}_D[n]$:

$$\hat{x}[n] = \sum_{k=-M}^M \hat{x}_D[k] \cdot x[n-k] \quad (2.20)$$

where $2M + 1$ represents the number of coefficients, or taps, of the filter. This discrete representation is what allows the FPGA to generate the analytical signal $z[n] = x[n] + i\hat{x}[n]$ in real-time. By substituting the ideal operator with a FIR structure, the processing unit can perform the phase extraction as the data is streamed, providing the necessary foundation for the real-time self-correction algorithm.

2.4.2. From SoC to RFSoc

An hardware implementation of the real-time self correction algorithm requires a Digital Signal Processing Hardware. Specifically is important to define what are a System-On-Chip (SoC) and a Radio Frequency System-On-Chip.

A System-on-Chip (SoC) is an integrated circuit that consolidates all the necessary components of a computer, such as the central processing unit (CPU), memory, and input/output ports, onto a single silicon substrate. This integration significantly reduces power consumption and communication latency compared to traditional multi-chip designs.

A major milestone in this field, pioneered by Xilinx, is the Radio Frequency System-on-Chip (RFSoc). This platform extends the SoC concept by embedding high-performance Analog-to-Digital Converters (ADCs) and Digital-to-Analog Converters (DACs) directly into the die, alongside the processing cores and programmable logic. For Dual-Comb Spectroscopy, this is transformative: it enables Direct RF Sampling, allowing analog interferograms to be digitized and fed to the processing unit almost instantaneously. By eliminating external chips and complex interfaces, the RFSoc provides the path necessary to implement a real-time stabilization for DCS.

2.4.3. FPGA : The Computational Engine of the RFSoc

The computational engine of the stabilization system is the Field Programmable Gate Array (FPGA), a versatile type of integrated circuit, directly embedded within the RF-

SoC chip, that represents a middle ground between software-based processing and fixed hardware [5]. The term “*field programmable*” refers to the device’s fundamental ability to be configured and reconfigured by the user after the chip has been manufactured. This flexibility allows the researcher to adapt the hardware to specific or experimental needs without physically altering the device and without needing each time a different instrument.

Architecturally, an FPGA consists of an array of programmable logic blocks and flexible interconnects. By using a Hardware Description Language (HDL) like VHDL or Verilog, the user can define how these blocks are linked together, effectively “drawing” a custom digital circuit directly onto the silicon. This reconfigurable nature is achieved through two main types of internal blocks:

- **Configurable Logic Blocks (CLB):** These are used to create the general structure of the system, acting as the building blocks for data routing and the overall control logic.
- **Digital Signal Processing (DSP) Blocks:** These are specialized sections of the chip optimized for high-performance calculations. They are critical for the stabilization task because they allow complex mathematical operations, such as the Discrete Hilbert Transform, to be performed with massive parallelism.




Figure 2.8: Modern FPGA architecture showing different basic components (Adapted from [2])

Unlike a standard processor that executes instructions one after another (sequentially), an FPGA can perform multiple operations at the same time. This hardware-level execution ensures a deterministic timing and a very high speed, which are essential to correct the laser signal instantly, as the data arrives, which is precisely what is needed for a real-time correction algorithm.

3 | Dual Comb Spectroscopy Setup

Having established the theoretical foundations, this chapter details the experimental implementation developed in this work. First, the overall architecture of the optical setup is outlined, identifying the key optoelectronic components and their integration. Subsequently, a specific focus is dedicated to the data acquisition system, which represents a central asset of this thesis. This section details the performance and roles of both the high-definition digital oscilloscope and the AMD Xilinx RFSoc 4x2 FPGA-based board. Fig. 3.1 illustrates the general experimental scheme. The complete optical path can be divided into three main functional blocks: the optical frequency comb sources, the pulse amplification and spectral broadening stage, and the interferometric detection and CW calibration module. In the following Sections the experimental layout will be described in further detail.

3.1. Experimental layout

The core of the optical setup consists of two free-running mode-locked Erbium solid-state frequency combs (manufactured by Menhir Photonics) emitting short pulse trains (impulse duration around 150 fs) centered at a wavelength of $\lambda_c = 1.55\text{ }\mu\text{m}$. Both laser oscillators operate at a $f_{\text{rep}} \approx 1\text{ GHz}$. To enable multi-heterodyne down-conversion into the radio-frequency (RF) domain, the two combs are slightly detuned to establish a repetition rate difference:

$$\Delta f_{\text{rep}} \approx 0 \div 50\text{ kHz} \quad (3.1)$$

This detuning determines the periodicity of the time-domain interferograms which therefore is around $T_{\text{rep}} = 1/\Delta f_{\text{rep}} \geq 20\text{ }\mu\text{s}$, thereby enabling extremely rapid update rates for the retrieved spectra.

The fiber-coupled outputs of the oscillators are managed through an active safety interlock system that continuously monitors the pulse intensity via a dedicated photodiode. If the input power deviates from a safe operating window (with the upper limit to the generated power set to 40 mW), this system safeguards the following stages by automatically disabling the pump currents of the downstream NuPhoton Erbium-Doped Fiber Amplifier

(EDFA). This mechanism directly protects the active gain fibers inside the amplifier.




Figure 3.1: Scheme of the dual-comb spectroscopy setup. The pulses from two independent 1-GHz Erbium solid-state lasers are temporally stretched by polarization-maintaining dispersion-compensating fibers (PM-DCF) and amplified via Erbium-Doped Fiber Amplifiers (EDFA). After re-compression through diffraction gratings, supercontinuum (SC) generation is performed. The broadened output of the first comb interrogates the reference gas cell, whereas the second comb acts as a local oscillator. To provide optical frequency calibration, a CW laser is combined alongside the combs via a beam splitter, subsequently the signal goes through a narrow-band filter and is detected by a photodiode and saved via an oscilloscope. The multi-heterodyne beating signal goes instead through a $4f$ filter and is then captured by a fast photodiode and recorded using a digital oscilloscope.




Figure 3.2: Portion of the real setup dedicated to chirped pulse amplification.



Figure 3.3: Commercial supercontinuum generation module (Cosmo by Octave Photonics) used in the experimental setup.

To achieve the high peak intensities required for efficient spectral broadening, the optical pulses must undergo amplification. A Chirped Pulse Amplification (CPA) scheme (1.3) is implemented. The pulses are initially directed through approximately 15 meters of

Polarization-Maintaining Dispersion-Compensating Fiber (PM-DCF). Then they are subsequently amplified within a NuPhoton Erbium-Doped Fiber Amplifier (EDFA), boosting the average power to several hundreds of milliwatts. Following amplification, the pulses are re-compressed in space using a pair of transmission diffraction gratings. The geometric parameters of the grating compressor (specifically a relative distance of ≈ 60 mm and an orientation angle of $\approx 47.5^\circ$) are precisely tuned to compensate for the total accumulated GVD across the system (Fig.3.2). This optimization yields transform-limited pulses with durations below 300 fs (specifically $\tau_1 \approx 255$ fs for the first comb, and $\tau_2 \approx 257$ fs for the second, previously characterized by autocorrelation measurements) and a beam waist of 1.3 - 1.7 mm (this was characterized by an IR camera). The highly intense, compressed pulses are then coupled into the supercontinuum (SC) generation module (Cosmo by Octave Photonics 3.3). This module relies on a tantala (Ta_2O_5) nanophotonic waveguide. The strong confinement and extreme optical non-linearity within the waveguide broadens the discrete laser spectrum, expanding it across the entire near-infrared molecular window.




Figure 3.4: Experimental optical spectrum of the generated dual-comb supercontinuum for different EDFA drive currents. The emission spans over a full octave, covering a broad spectral window from approximately 1000 nm up to 2400 nm (corresponding to a wavenumber range from roughly 4160 cm^{-1} to 10000 cm^{-1}). This broadening completely covers the target molecular absorption bands, including the acetylene (C_2H_2) transitions centered around $1.55 \text{ }\mu\text{m}$ ($\approx 6450 \text{ cm}^{-1}$).

3.2. Absolute Calibration and Detection

Once the broad supercontinuum spectrum is generated for both combs, the beams are routed toward the final stage of the optical setup. For the spectroscopic validations in this work, a 1.5m meter long cell containing $^{13}\text{CH}_4$ (methane-13) and C_2H_2 (acetylene) is used. The cell operates at a total internal pressure of 150 mbar and a room temperature of 296.15 K (partial pressure shares of 64.3% for $^{13}\text{CH}_4$ and 35.7% for C_2H_2). In this work

the focus will be on acetylene that, within the near-infrared window, provides a regular series of roto-vibrational lines spanning from 6450 cm^{-1} to 6600 cm^{-1} . As the broadband pulses pass through the medium, the gas imprints its specific molecular absorption profile onto the spectral envelope of the comb.

The layout adopts a standard configuration: the broadened output of the first comb (f_{rep})




Figure 3.5: Photograph of the initial section of the experimental setup, showing the two laser sources and the beam entering the molecular gas absorption cell.

is collimated and transmitted directly through the reference gas cell, whereas the output of the second comb bypasses the gas cell entirely and is subsequently recombined with the first via a beam splitter (BS). This superimposes the two pulse trains in the time domain to generate the interferograms. The two combined combs then propagate through a 4f filter, which is crucial for selecting the desired spatial frequencies and managing aliasing. This first signal is then detected by a high-speed extended Indium Gallium Arsenide (InGaAs) photodiode (sensitive in the $1 - 2.5\text{ }\mu\text{m}$ spectral range) and saved by a data acquisition system (DAQ), in this case an oscilloscope. This initial signal constitutes the dual comb trace, which contains all the spectroscopic information. Nevertheless, to establish the absolute metrological frequency axis, a second trace needs to be saved. For this reason, a fraction of the combined optical power is co-propagated with an ultra-stable

Continuous-Wave (CW) laser (TeraXion-LXM) centered at an anchor frequency of:

$$\nu_{CW} = 192234.9 \text{ GHz} \quad \lambda_{CW} \approx 1559.51 \text{ nm}$$

The exact, real-time frequency of this CW reference is continuously monitored by a high-precision wavemeter. The CW laser light and the comb signal are combined through a 90:10 fiber coupler. By selecting the output port that delivers 90 % of the comb power and 10 % of the CW power, the setup successfully balances the significant intensity mismatch between the high-power continuous laser and the individual comb lines, maximizing the visibility of the heterodyne beat note and preventing photodiode saturation. Finally, the combined beams are filtered by a narrow-band ($1557 \pm 3,5nm$) optical filter and then focused onto a high-speed extended Indium Gallium Arsenide (InGaAs) photodiode. The photodetector possesses an electronic bandwidth wide enough to capture the full multi-heterodyne RF beat spectrum. The output of the photodiode yields the raw time-domain interferogram, which contains the down-converted optical absorption information along with the calibration beat notes. This trace will be essential to recover the optical spectrum. The analog electrical signal is then routed directly to a different channel of a Data Acquisition (DAQ) system for digitization(3.3).




Figure 3.6: View of the final section of the dual-comb spectrometer layout, capturing the path from the exit of the gas cell down to the high-speed photodetector

3.3. Data Acquisition Systems (DAQ)

The analog electrical signals generated by the high-speed photodiodes contain the multi-heterodyne RF spectra, which encapsulate both the molecular absorption profile from the gas cell and the optical frequency calibration beat notes from the CW reference laser. Digitizing these distinct raw time-domain interferograms simultaneously and with high fidelity is a critical step in the Dual-Comb Spectroscopy architecture. In the framework of this thesis, the data acquisition strategy relies on two different instrumentation approaches, serving distinct roles within the experimental implementation.

On one hand, a high-definition digital oscilloscope is employed as the primary, benchmark DAQ system. By dedicating separate input channels to the spectroscopy and calibration signals, it provides the necessary resolution, sampling rate, and memory depth to capture and store waveforms for subsequent offline processing. On the other hand, an innovative hardware solution based on the AMD Xilinx RFSoc 4x2 board represents a newer alternative that will be integrated into the setup in the future. While the oscilloscope represents the standard tool for current measurements, the RFSoc platform constitutes a central asset of this work. A key objective of this thesis is indeed to investigate and take the first development steps toward a multi-channel, fully embedded, FPGA-based DAQ system capable of performing real-time data acquisition and on-board digital signal processing.

In the following Subsections, the technical specifications and specific operational roles of both platforms are detailed.

3.3.1. High-Definition Digital Oscilloscope

For standard experimental acquisitions and spectroscopic validations, the multi-heterodyne radio-frequency signals generated by the photodiodes are routed to a Teledyne LeCroy HDO6054 high-definition digital oscilloscope [24]. This instrument features an analog bandwidth of 500 MHz and a maximum sample rate of 2.5 GSps which are perfect to capture the full down-converted RF beat spectrum.



Figure 3.7: The Teledyne LeCroy HDO6054 high-definition digital oscilloscope used as the benchmark data acquisition system for offline multi-heterodyne signal processing.

Crucially, the oscilloscope operates with a native 12-bit vertical resolution. Compared to standard 8-bit digitizers, this high-definition architecture provides a significantly higher dynamic range and minimizes quantization noise. In Dual-Comb Spectroscopy, an elevated vertical resolution is fundamental to resolve extremely shallow molecular absorption depths and to preserve a high Signal-to-Noise Ratio (SNR) during data retrieval.

The instrument captures the raw, synchronous time-domain waveforms over long acquisition windows, exploiting its deep onboard memory. The stored time-domain interferograms are then exported for offline processing.

3.3.2. AMD Xilinx RFSoc 4x2 Platform

To transition from offline analysis toward an embedded, real-time data acquisition paradigm, this work investigates the AMD Xilinx RFSoc 4x2 development board. While the Teledyne LeCroy oscilloscope serves as the high-fidelity benchmark for data gathering in this thesis, the RFSoc 4x2 Platform represents the hardware baseline intended for the future integration of on-board digital signal processing pipelines.

The board is built around an AMD Xilinx Zynq UltraScale+ ZU48DR Gen 3 RFSoc device [5]. This single-chip architecture drastically minimizes latency and power consumption compared to traditional designs. In older systems, the data converters (ADC/DAC) and the processing units (FPGA) were separate chips placed on the board, requiring complex, high-speed connection lines to communicate. By contrast, the RFSoc integrates the high-speed data converters, the FPGA Programmable Logic (PL), and the multi-core ARM Processing System (PS) onto the exact same silicon die, removing the need for

external connection interfaces and ensuring nearly instantaneous data transfer. [5].




Figure 3.8: The AMD Xilinx Zynq UltraScale+ RFSoc 4x2 development board configured with a Gen 3 ZU48DR device.

For signal acquisition, the board does not use all the channels available on the integrated ZU48DR chip. Instead, it provides access to a specific selection of them through physical, high-quality SMA screw connectors. Specifically, the platform features:

- **Four RF Analog-to-Digital Converters (RF-ADCs):** These channels offer a vertical resolution of 14 bits, support a high-speed sample rate of 5 GSPS (Giga-Samples Per Second) and allow direct RF sampling.
- **Two RF Digital-to-Analog Converters (RF-DACs):** Operating alongside the inputs with a 14-bit resolution, these channels support a sample rate of 9.85 GSPS and manage transmitter signal generation.

From a software and control perspective, the RFSoc 4x2 relies on the open-source PYNQ framework, which activates a Linux-based subsystem on the internal ARM processor cores [1]. This ecosystem allows the researcher to control and interact with custom hardware overlays implemented inside the FPGA PL via Python scripts running inside interactive Jupyter Notebooks [5]. In this context, a hardware overlay is fundamentally a pre-compiled bitstream file that, when loaded, physically reconfigures the internal digital circuits and routing logic of the PL. Instead of writing low-level hardware description languages, this high-level software abstraction provides a seamless pipeline for fast prototyping, run-time parameter tuning, and data visualization, paving the way for real-time

digital filtering and windowing implementations as explored in Chapter 5.

4 | Experimental measurements and results: Post processing Data Analysis

This chapter represents the first of two connected sections dedicated to the discussion of the experimental results. While the subsequent chapter focuses on the transition toward embedded hardware architectures, the present one discusses the post-processing data analysis. This approach relies on a separation between data acquisition and signal processing. The digital oscilloscope acts as a recorder, subsequently the stored data files are transferred to a computer and processed offline. This strategy avoids the necessity of complex analog feedback loops or physical phase-locked electronics to stabilize the laser cavities during the measurement. All physical instabilities are systematically corrected inside the digital domain via software algorithms.

The primary objective of this introductory section is to map out the logical workflow of the developed scripts (appendix C). Rather than presenting the algorithms purely from an abstract mathematical perspective, each script will be illustrated in parallel with its practical effects on actual laboratory data. To achieve this, the progressive stages are shown using the intermediate results extracted from a single reference experimental measurement. In our specific experimental configuration, the raw data generated by the oscilloscope is stored as text files (ASCII format) structured in two main columns: the first registers the discrete time axis, while the second records the raw voltage amplitude of the interferometric signal captured by a fast photodiode. An acquisition sampling rate of 1.25 GSps is used and a record length of 20 ms is set due to physical memory constraints. Indeed a single file encompasses millions of points, occupying around hundreds of *Mbytes*. There are many reasons why this happens, such as the wide Nyquist window ($f_{rep}/2 = 500MHz$), the need to record many (typically >100) long interferograms ($T_{rep} = 1/\Delta f_{rep} \geq 20 \mu s$), and the required high vertical resolution.

Another crucial aspect worth mentioning is the optical bandwidth that can be captured in the near-infrared region, which, for broadband configuration, spans approximately 230 nm.

This broad spectral range is the result of an optimization between different parameters, according to Eq. 2.11. The maximum optical bandwidth ($\Delta\nu_{\max}$) without aliasing is a trade-off between the repetition rate f_{rep} and the detuning Δf_{rep} . Furthermore it should be noted that Δf_{rep} cannot be arbitrarily diminished, since the burst period ($T = 1/\Delta f_{\text{rep}}$) would scale toward infinity, compromising the acquisition.

Now the different scripts written to perform the Self-Correction analysis are briefly described. The processing framework is structured as series of separate scripts, where each program takes the output generated by the previous step as its input. This sequence can be summarized as follows:

- **Script 1: The Self-Correction Framework.** As the entry point, this script takes as input the raw ASCII data file and subtracts the continuous background offset. It localizes the temporal positions of the interferogram bursts via Hilbert envelope detection and applies a digital correction scheme to cancel out relative timing and phase fluctuations, concluding with a coherent averaging of the stabilized bursts in the time domain.
- **Script 2: Multi-Trace Averaging.** To overcome the physical memory limitations of a single oscilloscope trace and increase the signal-to-noise ratio, this optional script aggregates and averages a successive series of pre-corrected files.
- **Script 3: Baseline Extraction and Spectrum Normalization.** The final module of the pipeline addresses the non-flat intensity profile imposed by the super-continuum generation. It implements an algorithm to extract the upper envelope (baseline) of the calibrated spectrum without being distorted by the molecular absorption lines, and performs a point-by-point division to yield the final, normalized transmittance spectrum.
- **Script 4: Absolute Frequency Calibration.** Grounded on a comprehensive multi-heterodyne sampling theory, this script processes further parameters extracted from the oscilloscope (and eventually adjusted by the beating detection routine) along with the absolute frequency of the CW reference measured by a wavemeter. It applies the non-linear scaling relations to map the radio-frequency axis into the optical frequency scale.

It is worth noting that the execution order of these scripts is flexible, with the exception of Program 1, which must strictly be executed at the beginning of the pipeline. While the sequence described above represents an ideal workflow, it is by no means the only viable approach. In fact, in the following sections, where these individual processing steps will be shown in detail, an alternative sequence will be adopted, consisting of script 1, fol-

lowed by absolute frequency calibration (script 4), multi-trace averaging (script 2), and ultimately baseline normalization (script 3).

Having established this overview, the following sections will delve deeper into the technical and algorithmic details of each individual script, illustrating their impact on the experimental data.

4.1. The Self-Correction Framework

As discussed before when operating under free-running conditions, the relative instabilities of the two laser sources introduce severe temporal and phase fluctuations from burst to burst. Specifically, variations in the repetition-rate difference (Δf_{rep}) manifest as timing jitter, altering the arrival times of consecutive interferograms. Concurrently, fluctuations in the carrier-envelope offset frequencies cause a random phase drift of the optical carrier beneath the temporal envelope. Without compensation, these instabilities rapidly degrade the mutual coherence between the combs, causing successive waveforms to interfere destructively and effectively washing out the radio-frequency (RF) comb structure during long integration windows.

The self-correction algorithm (code reported in Appendix C.1.1) is the one implemented to extract and compensate for these errors directly from the raw multiheterodyne signal, restoring full phase coherence.

The digital pipeline begins with the preliminary processing of the raw multiheterodyne voltage trace recorded by the digital oscilloscope over an acquisition window (typically 20 ms at a digitization rate of 1.25 GS/s). To handle the huge dataset consisting of millions of points, the arrays are allocated within a parallel graphics processing architecture to expedite the vector operations. A preliminary baseline subtraction is then applied to remove any continuous-current (DC) offset.

To establish an absolute phase and timing reference for each individual multiheterodyne event, the exact temporal centers of successive bursts must be identified. This is achieved by isolating the analytical signal representation of the waveform via an Hilbert transform, generating the complex-valued vector:

$$s_a(t) = s(t) + i\mathcal{H}\{s(t)\} \quad (4.1)$$

The time-dependent amplitude envelope is then extracted as the absolute magnitude of this analytical representation, $A(t) = |s_a(t)|$. A peak-detection routine is executed on the envelope to localize the maximum of each burst. Furthermore, to handle highly irregular or noisy profile envelopes, the algorithm can recalculate the discrete peak coordinates

by calculating a localized center-of-mass (centroid) within a narrow symmetric window around each maximum.

The first correction stage addresses the relative timing jitter. Because the physical interval between consecutive peaks varies over the acquisition window due to Δf_{rep} fluctuations, the discrete time axis is mathematically remapped. The algorithm performs a numerical time-warping interpolation that maps the original non-equidistant burst coordinates onto a perfectly linear, homogeneous temporal grid where the individual burst centers are forced to be rigorously equidistant. This operation effectively stabilizes the macroscopic repetition period of the multiheterodyne train, realigning the temporal arrival of the pulse envelopes across the entire trace.

The second and more critical stage compensates for the carrier-envelope phase jitter.




Figure 4.1: Experimental profiles of the tracked instabilities across the 20 ms acquisition window. (Top panel) Relative timing jitter.(Bottom panel) Cumulative carrier-envelope phase jitter $\Phi(t)$.

Although the envelopes are perfectly equidistant after the time-warping interpolation, the underlying optical carrier oscillations still exhibit random phase offsets from burst to burst. The absolute phase evolution is retrieved by evaluating the argument of the complex analytic signal at each localized peak coordinate, tracking the cumulative carrier phase drift across the acquisition window. The discrete phase profile is then continuously interpolated over the entire time axis to generate a smooth phase correction vector, $\Phi(t)$. Phase stabilization is ultimately achieved by multiplying the time-warped analytical signal by the inverse exponential of this continuous phase function:

$$s_{\text{corr}}(t) = s_a(t) \cdot e^{-i\Phi(t)} \quad (4.2)$$

By taking the real part of $s_{\text{corr}}(t)$, the algorithm yields a fully stabilized interferogram train. This phase cancellation procedure eliminates the relative carrier-envelope offset




Figure 4.2: Continuous carrier-envelope phase correction vector $\Phi(t)$ obtained after linear interpolation across the entire 20 ms window.

fluctuations, locking the underlying carrier oscillations to the envelope symmetry.




Figure 4.3: Comparison of the multi-heterodyne interferogram train over a localized temporal window. (Top panel) The original uncorrected waveform. (Bottom panel) The fully corrected interferogram waveform highlights the homogeneity and uniform periodicity achieved after the computational time and phase corrections.

Following correction, individual bursts are isolated using a symmetric temporal window centered on each peak. The extracted segments are synchronously accumulated to generate a coherently averaged time-domain burst with an improved signal-to-noise ratio. Concurrently, the Fourier transform of each single segment is computed and averaged to yield the mean spectral distribution.




Figure 4.4: Zoomed view of a single burts before and after the timing and phase jitters corrections. It is evident how the carrier envelope phase was canceled.




Figure 4.5: Overlay of individual multi-heterodyne interferogram bursts. Due to the action of the correction all sequential bursts are successfully forced into complete mutual coherence, showing excellent constructive alignment.




Figure 4.6: Comparison demonstrating the impact of coherent time-domain averaging on the interferogram tails. (Top panel) Overlay of the un-averaged individual segmented traces; far from the main burst center, the molecular signature is completely masked by the high noise floor. (Bottom panel) Resulting coherently averaged single burst profile; the random noise contributions are actively suppressed, revealing the molecular absorption lines of the sample.




Figure 4.7: Multi-heterodyne frequency comb structure before (blue) and after (red) digital computational correction. In the uncorrected spectrum (blue), the individual longitudinal modes are severely broadened. Following the self-correction processing (red), the teeth collapse back into sharp discrete spectral components.

This procedure completely restores mutual coherence across the entire record length, transforming the initially broadened, unstable multiheterodyne signal into a stable se-

quence of identical phase-aligned bursts suitable for coherent accumulation.

4.2. Absolute Frequency Calibration

The radio-frequency (RF) spectrum generated through the dual-comb measurement provides relative frequency information. To reconstruct the absolute optical frequencies and map the down-converted RF features onto a high-fidelity optical scale, an external, well-characterized frequency reference is required. In this experimental configuration, this is achieved by optically mixing a continuous-wave (CW) laser of known absolute frequency, ν_{CW} (measured simultaneously via a high-accuracy wavemeter), with both frequency combs.

The optical heterodyne mixing between the CW reference and its nearest tooth of Comb 1 and Comb 2 generates two distinct radio-frequency beat notes, denoted as $f_{\text{CW},1}$ and $f_{\text{CW},2}$, respectively. The relative difference between these two auxiliary beatings defines a fundamental reference parameter:

$$\Delta f_b = (\pm f_{\text{CW},1}) - (\pm f_{\text{CW},2}) \quad (4.3)$$

where the signs of the individual beat notes depend strictly on the relative spectral positioning of the CW laser with respect to the local comb teeth.

Now recalling that only comb 1 is passing through the sample and that the scaling factor is to be considered, under standard baseband mapping conditions, the absolute optical frequency ν_{abs} corresponding to any measured radio-frequency component f_{RF} can be reconstructed as:

$$\nu_{\text{abs}} = \pm(f_{\text{RF}} - \Delta f_b) \frac{f_{\text{rep},1}}{\Delta f_{\text{rep}}} + \bar{\nu}_{\text{CW}} \quad (4.4)$$

where $f_{\text{rep},1}$ represents the repetition rate of Comb 1, Δf_{rep} is the repetition rate difference between the two sources, and $\bar{\nu}_{\text{CW}} = \nu_{\text{CW}} \pm f_{\text{CW},1}$ defines the absolute optical frequency of the comb tooth acting as the local reference. To uniquely resolve the sign ambiguities in Equation (4.4), both the global orientation of the RF spectrum and the localized position of the CW laser must be experimentally determined. The overall sign of the term $(f_{\text{RF}} - \Delta f_b)$ depends on whether the multi-heterodyne process yields a direct or an inverted representation of the optical spectrum. This mapping condition is dictated by whether the point where two teeth from opposite combs perfectly overlap, occurs at an optical frequency lower or higher than the spectral region of interest. Experimentally, this ambiguity is solved by actively tuning the repetition rate of the fastest comb ($f_{\text{rep},1}$) and observing the direction of the spectral shift on the digital grid; a shift toward higher

radio-frequencies indicates a direct mapping (positive sign), whereas a shift toward lower frequencies reveals an inverted representation (negative sign).




Figure 4.8: Direct and Inverted RF spectrum

Concurrently, the individual signs of $f_{CW,1}$ and $f_{CW,2}$ within the Δf_b term are governed by the CW laser relative position. If the CW laser frequency is lower than the nearest teeth of both combs, any increase in the comb repetition rates will physically move the teeth away from the reference, shifting the RF beatings toward higher frequencies (assigned as positive signs). Conversely, if the CW laser is situated at higher frequencies than the adjacent teeth, a rightward shift of the combs moves the teeth closer to the reference, driving the beat notes toward lower frequencies (assigned as negative signs). Finally, if the CW reference lies directly in the gap between the adjacent teeth of the two combs, a global frequency shift will move one tooth closer and the other farther away, yielding opposite signs for the two beat notes. These dynamics are systematically validated in the laboratory by monitoring the real-time displacement of the beat notes on the oscilloscope.

	Sign of $f_{CW,1}$	Sign of $f_{CW,2}$
Case 1	+	+
Case 2	−	−
Case 3a	−	+
Case 3b	+	−

Table 4.1: Summary of sign assignments for the CW beatings corresponding to the scenarios illustrated in Fig.4.9.




Figure 4.9: Possible positions of the CW laser with respect to the different combs

While this calibration model operates effectively when the multi-heterodyne signal is completely contained within the baseband, the wide optical bandwidth imposed by the supercontinuum generation module typically spans across multiple Nyquist zones. As a consequence of the sampling theorem applied to discrete periodic structures, the high-frequency optical components undergo a periodic folding process back into the baseband spectral window ($0 \leq f_{\text{RF}} \leq f_{\text{rep}}/2$). To preserve absolute accuracy across the entire wideband spectrum, the calibration framework is generalized by incorporating the specific Nyquist zone index, n_{fsr} , and the Nyquist frequency limit, $f_{\text{Ny}} = f_{\text{rep},1}/2$.

When the spectrum in the baseband is direct, the absolute optical frequency axis (into absolute wavenumber units cm^{-1}) is reconstructed as follows:

$$\nu_{\text{abs}} = \frac{\bar{\nu}_{\text{CW}} + M \cdot [f_{\text{RF}} \cdot (-1)^{n_{\text{fsr}}} - |\Delta f_b| + (n_{\text{fsr}} + 0.5 \cdot (1 + (-1)^{n_{\text{fsr}}+1})) \cdot f_{\text{Ny}}]}{c} \quad (4.5)$$

Conversely, if the spectrum is inverted, the formula needs to be slightly modified, resulting in:

$$\nu_{\text{abs}} = \frac{\bar{\nu}_{\text{CW}} + M \cdot [f_{\text{RF}} \cdot (-1)^{n_{\text{fsr}}+1} + |\Delta f_b| + (n_{\text{fsr}} - 0.5 \cdot (1 + (-1)^{n_{\text{fsr}}+1})) \cdot f_{\text{Ny}}]}{c} \quad (4.6)$$

In both expressions, $\bar{\nu}_{\text{CW}}$ represents the absolute optical frequency of the reference comb tooth interrogating the sample, which is determined relative to the continuous-wave laser reference via $\bar{\nu}_{\text{CW}} = \nu_{\text{CW}} \pm f_{\text{CW},1}$. The speed of light is indicated by c while the scaling factor $M = f_{\text{rep},1}/\Delta f_{\text{rep}}$ acts as the multi-heterodyne magnification coefficient that links the radio-frequency domain to the optical domain. By applying these generalized algorithms, the processing pipeline (code reported in Appendix C.1.3) successfully unwraps the periodically replicated multi-heterodyne tracks, transforming the raw, folded oscilloscope frequencies into a high-fidelity optical frequency scale.

4.3. Further averaging

While the self-correction framework successfully restores mutual coherence within a single continuous acquisition, in our experimental conditions, a single 20 ms record length represents a physical ceiling that limits the achievable signal-to-noise ratio. To overcome this hardware constraint and push the detection sensitivity further, multiple independent sequential acquisitions can be aggregated.

However, directly averaging these separate files is prevented by long-term environmental instabilities. Over the time separating consecutive measurements, the free-running laser cavities experience slow thermal and mechanical drifts. These fluctuations cause a macroscopic frequency shift of the entire multi-heterodyne structure, meaning that the discrete lines would misalign from trace to trace. If a simple linear average were performed without correction, these shifts would artificially broaden the lines, ruining the spectral resolution. To circumvent this issue, the multi-trace processing pipeline (code reported in Appendix C.1.2) implements a digital alignment routine. Each acquired file contains two parallel columns of information that share the exact same physical drift: the spectroscopic signal and the down-converted radio-frequency comb structure. To realign these datasets to a common grid, the algorithm first applies a localized smoothing filter to suppress high-frequency noise and then computes the numerical first derivative of the spectrum.




Figure 4.10: Overlap of the calibrated optical spectra after the cross-correlation alignment processing.

This mathematical differentiation effectively removes broadband baseline variations and accentuates the sharp transitions of the absorption lines. Using the first acquired file as a static reference, the algorithm performs a statistical cross-correlation analysis on the spectral derivatives of all subsequent traces. This mathematical tool acts as a digital ruler, determining the integer pixel shift required to realign each specific dataset with the reference. Once this frequency offset is calculated, a rigid linear translation is applied simultaneously to both the spectrum and the comb structure vectors, ensuring they remain perfectly synchronized.




Figure 4.11: Final averaged and corrected optical spectrum of the sample (cell H). By linearly averaging the previously aligned datasets, the result is a drastically improved signal-to-noise ratio (SNR).

To guarantee the high fidelity of the final data, a strict quality control check is executed immediately after the alignment. The algorithm calculates the Pearson correlation coefficient to measure the exact overlap between the reference amplitude and the newly shifted target array. If a trace exhibits a similarity score below a chosen threshold (usually over 0.95), it implies that an anomalous disturbance occurred and that script is automatically discarded. Finally, all validated and realigned traces are linearly averaged, dramatically improving the signal-to-noise ratio.

4.4. Final Spectrum obtained: normalization

Once the post-processing pipeline has successfully executed the self-correction routine, performed the coherent time-domain averaging of individual bursts, and eventually applied the multi-trace stage alongside the absolute frequency calibration, the experimental data is fully projected onto the definitive optical wavenumber grid (cm^{-1}). At this stage, the absolute frequency positions of the comb modes are completely locked, but the raw experimental intensity spectrum $I(\nu)$ remains structurally superimposed on a non-uniform background. This envelope reflects the asymmetric spectral profile of the supercontinuum source, which distorts the relative amplitudes of the molecular transitions. To isolate the transmission of the sample, a terminal baseline-flattening routine is executed within the

MATLAB framework (code reported in Appendix C.1.4) .




Figure 4.12: Acetylene spectrum (blue) with the calculated non-flat laser baseline (red).

To eliminate the non-uniform supercontinuum envelope, the algorithm filters the intensity profile using native MATLAB data-smoothing functions. This filtering process suppresses localized noise and effectively bypasses the deep, narrow molecular absorption valleys, computationally tracing the unattenuated envelope of the light source. A smooth, regularized baseline function, $I_0(\nu)$, is then generated to precisely model the reference background profile without tracking any sharp absorption line. Finally, the true normalized transmittance spectrum $T(\nu)$ is extracted simply by dividing the raw calibrated intensity spectrum by the computed baseline profile:

$$T(\nu) = \frac{I(\nu)}{I_0(\nu)} \quad (4.7)$$




Figure 4.13: Final normalized transmittance spectrum of the acetylene sample, showing a unity transmittance.




Figure 4.14: Normalized experimental spectrum of acetylene overlaid with the HITRAN database simulation.

This numerical division successfully flattens the background and bounds the final molecular transmittance strictly between 0 and 1 (or 0% and 100%); providing a clean spectroscopic trace ready for analysis.

The regular absorption dips of the Acetylene vibrational band are clearly resolved with high contrast across a broad spectral window. To rigorously assess the quantitative accuracy of the post-processing software, the normalized experimental trace can be directly

superimposed on a forward-model simulation retrieved from the HITRAN atmospheric database under identical physical boundary conditions (50 mbar pressure, 15 cm interaction path length, and 296 K equilibrium temperature). The excellent overlay between the experimental line shapes and the database references confirms that the offline self-correction and normalization algorithms achieve instrument-limited resolution and absolute amplitude fidelity without requiring active hardware stabilization loops.

5 | Experimental measurements and results: towards Real-Time Analysis

As already discussed in Section 2.4, shifting the self-correction algorithm from an offline post-processing analysis to a real-time, hardware-integrated platform represents an important step to achieve immediate frequency and phase-jitters compensations during data acquisition. To evaluate the technical feasibility of this hardware implementation, this chapter focuses on a preliminary study and a reverse-engineering analysis of the recent open-source framework proposed by Eber et al. [7], who successfully demonstrated a streaming self-corrected dual-comb spectrometer using the exact same hardware architecture available in our laboratory: the AMD RFSoc 4x2 board.

Previously (3.3.2), it has been seen that the AMD RFSoc 4x2 is a system-on-chip platform that natively combines high-speed radio-frequency analog-to-digital converters (ADCs) with programmable logic (FPGA) and an ARM processing system. A key feature of this board is its integration with the PYNQ framework, an open-source project from AMD that allows developers to exploit the hardware acceleration of the FPGA using the Python language. Through PYNQ, the hardware architecture designed via the Xilinx Vivado Design Suite [2] is compiled into a specific deployment package, consisting of a bitstream file (.bit) and a hardware handoff file (.hwh), which is loaded as a hardware *overlay* directly from a Python environment running on the board's embedded Linux system. A detailed description of the Vivado design environment [1], the concept of Intellectual Property (IP) blocks, and the specific architecture of the base overlay used in this work will be described in Section 4.4.

The main purpose of this chapter is to detail the preliminary steps undertaken to understand and test this architecture. First, the reference Python-based simulation algorithm from the literature repository [17] was thoroughly analyzed to map the digital signal processing (DSP) workflow required for future hardware deployment. Second, initial experimental benchmarks were carried out by exploiting the manufacturer's default RF

acquisition overlay via PYNQ, using an external waveform generator to validate the signal integrity and the ADC performance. Finally, the internal block design of the utilized overlay (base.bit) was analyzed within the Vivado Design Suite environment to assess the complexity of the pre-existing logic and identify the integration points for future self-correction implementations.

5.1. Python simulation of the real-time algorithm

To bridge the gap between the theoretical self-correction framework and its implementation on the AMD RFSoc platform, a test of the algorithm is performed in a Python environment. The core of this analysis consisted of understanding and executing a simulation script designed to emulate the digital signal processing (DSP) operations that will then be executed on the the FPGA hardware. This code models a streaming-like data flow to verify how the algorithm tracks and compensates for phase instabilities in real-time.

To gain a deep understanding of the hardware-oriented workflow, the analysis was focused on the specific simulation model designed by Eber et al. [7] to replicate the FPGA operations. This architectural model simulates a phase correction framework tailored for a fixed detuning of $20 \text{ KHz} \pm 5\%$. Choosing a fixed Δf_{rep} algorithm simplifies the mathematical structure.

To accurately model the behavior of the real-time self-correction system, the software architecture is structured around a fully parallelized digital signal processing pipeline. The processing framework does not use a standard sequential execution. It is divided into independent blocks that work at the same time. These blocks communicate dynamically through standardized First-In, First-Out (FIFO) data queues. This approach is essential to simulate what really happens on an FPGA, where different logic sectors process incoming samples simultaneously in a continuous streaming regime.

The workflow of this parallelized simulation can be mapped into four main logical stages:

1. **Data Ingestion and Scaling:** The first stage emulates the physical Analog-to-Digital Converter (ADC). It continuously streams the raw experimental dual-comb data into the processing pipeline.
2. **Discrete Hilbert Transform (FIR Filtering):** To extract the envelope and phase profile, as seen before, the signal must be converted into its complex analytic form. The architecture achieves this by passing the real input stream through a localized, windowed Finite Impulse Response (FIR) filter consisting of 101 symmetric coefficients. This sliding-window filter is the implementation of the discrete Hilbert transform discussed in section 2.4.1.

3. **Interferogram Localization and Triggering:** The complex analytical stream is fed into a circular buffer where an amplitude threshold-based trigger logic is continuously executed. When a burst from a dual-comb interferogram crosses the predefined activation threshold, the system flags the event and isolates a specific, localized temporal window surrounding the optical peak. This windowed subset is then routed to the measurement stage, while the main stream continues to flow uninterrupted.
4. **Barycentric Parameter Estimation:** The isolated interferogram window is analyzed to extract its physical properties. The algorithm calculates the precise temporal center of the burst by computing the center-of-mass of its energy distribution (centroid, this is the same procedure that was also implemented in the post-processing algorithm 4.4). Simultaneously, a localized Fast Fourier Transform (FFT) is performed on the centered burst to evaluate its mean frequency via a weighted barycentric calculation. By storing these coordinates and comparing them with the values obtained from the historical reference of the previous burst, the system quantifies the exact fluctuations induced by the laser phase and timing jitters.
5. **Phase and time jitters correction:** Once the phase and timing deviations are estimated, the parameters are instantly forwarded to the final execution stage of the pipeline, which performs the physical stabilization of the delayed data stream retrieved from a large circular buffer consisting of 250,000 samples. This final stage is responsible for two distinct hardware-emulated tasks: on-the-fly phase compensation and timing jitter mitigation.

The phase correction is executed via a digital phase accumulator that updates its tracking slope at every clock cycle, dynamically multiplying the data samples by a complex exponential correction vector ($e^{-i\Phi(t)}$) to cancel the optical phase noise.

Concurrently, a resampling stage dynamically adjusts the time spacing between subsequent dual-comb bursts. By stretching or compressing the timeline based on the fluctuations measured from the previous burst, the algorithm ensures that all processed interferograms are aligned at perfectly regular intervals, successfully neutralizing the timing jitter.

The output stream synthesized by this multi-threaded stage is then collected to evaluate the resulting spectroscopic performance via Fast Fourier Transforms and plots of different frequency windows. In particular, here the complete time-domain sequence containing the sequence of bursts before and after the algorithmic execution is reported in Fig. 5.1 and Fig. 5.2, respectively.




Figure 5.1: Uncompensated raw signal.




Figure 5.2: Corrected output stream.

Figure 5.3: Comparison of the complete dual-comb interferogram time-domain sequences before and after the algorithmic execution, plotted against the number of samples. The analysis was conducted using the reference acetylene dataset provided in the authors' official GitHub repository [17]

As can be observed by comparing the two plots, the uncompensated raw signal (Fig. 5.1) spans over 3.3×10^6 samples. Upon activation of the self-correction pipeline, the processed output stream (Fig. 5.2) features a reduced data distribution that caps at approximately 1.4×10^6 samples. This variation in the sample axis directly confirms the physical operation of the resampling stage, which dynamically re-aligns and compresses the time-domain data stream according to the synchronized grid. This comparison confirms that the python simulation works as intended, proving that the algorithm is ready to be implemented directly into the RFSoc 4x2 firmware.

5.2. External Signal Generation and RFSoc Acquisition Setup

The primary objective of this stage is to test the basic signal acquisition capabilities of the RFSoc 4x2 board. To ensure a controlled evaluation environment, the initial testing phase relies on the simplest possible waveform, specifically pure sinusoids synthesized by an external waveform generator (hewlett packard 33120A hp). To process these signals, the setup utilizes a Vivado overlay (base.bit) designed for standard Analog-to-Digital (ADC) and Digital-to-Analog (DAC) acquisition tasks, which is provided directly by the board manufacturers. Within this specific section, the scope is strictly limited to up-

loading this pre-compiled overlay onto the board and operating it via the PYNQ web interface. A detailed architectural analysis of how this base overlay is structured and configured within AMD Vivado will be comprehensively discussed in section 5.3.



Figure 5.4: Laboratory experimental setup with the external Hewlett Packard 33120A waveform generator connected via SMA cable to the RFSoc 4x2 receiver channel input.

To manage the hardware resources a dedicated control script (Appendix C.1.5) was executed within the Jupyter Notebook environment using the `xrfdc` Python package. The program initializes the `base.bit` overlay and configures the board’s clocking network, establishing a hardware sampling frequency (f_s) of 2457.6 MHz. The software configures the internal RF Data Converter (RF DC) hierarchy by accessing the specific receiver channel (`channel[3]`) mapped to the physical ADC port.

The data capture mechanism is managed directly through high-level functions provided by the overlay environment. Initially, the script reads a short data segment from the selected receiver channel to act as a probe. The program evaluates the peak amplitude of this probe against a predefined threshold to detect if a valid input signal is present. Once the signal level exceeds this threshold, the acquisition is triggered, and the incoming waveform is digitized by the ADC and transferred into the board’s memory buffers, where it becomes available for the Python environment to collect and process. These data are then retrieved by the host PC as complex-valued arrays (I and Q channels) for visualization and spectral analysis.




Figure 5.5: Acquisition with low sample size ($N_{\text{samples}} = 256$, $f_s = 2457.6$ MHz) for a 100 MHz single-tone input. The short observation time window captures fewer harmonic cycles in the time domain and results in broad, less defined frequency peaks in the FFT.




Figure 5.6: Acquisition with increased sample size ($N_{\text{samples}} = 30000$, $f_s = 2457.6$ MHz) for a 100 MHz single-tone input. The longer time axis allows the system to capture a dense train of oscillations, drastically narrowing the symmetric FFT peaks into ultra-sharp lines at ± 100 MHz.

To demonstrate the correct operation of the entire acquisition chain, the captured data were plotted in both the time and frequency domains. Two distinct experimental runs are reported below for comparison. In both cases, the external waveform generator was configured to synthesize a pure single-tone sinusoid at an exact frequency of 100 MHz. By comparing the two configurations, it can be observed how the precision of the frequency identification drastically changes as a function of the total number of captured samples. When a low number of samples is acquired, the peak in the FFT spectrum appears broad and less defined, limiting the accuracy of the frequency estimation. Conversely, by increasing the number of samples, the frequency peak narrows into a sharp and precise spectral line centered at the expected value.

In the time-domain plots, both configurations successfully trace the continuous sinusoidal oscillation of the signal. This behavior directly confirms that the board correctly samples and buffers the physical waveform.

5.3. Vivado Design Suite and the Base Overlay

To understand how the hardware works, it is necessary to study the development environment. AMD Xilinx Vivado is the software used to design and program the logic inside the FPGA. Unlike standard software programming, FPGA design defines physical hardware circuits. This is done using Hardware Description Languages (HDL) like Verilog or VHDL. In Vivado, a graphical tool called *Block Design* is used to simplify the process. Inside the Block Design, different blocks called Intellectual Property (IP) cores can be managed: they can be connected and arranged in space to generate the desired digital circuit. IPs are pre-made functional blocks: such as filters, memory controllers, interfaces or anything that performs some kind of operation. Users can import pre-made IP blocks ready to be used, for example, in this specific case, the IP block containing the physical processor system of the RFSoc 4x2 board was provided by the manufacturer. However, if a custom function is needed, a developer can write Verilog code, package it as a new IP and place it into the graphical workspace to connect it with the other components. After finishing the connections, it is necessary to run the *Validate Design* tool. This tool automatically checks the block design for errors or wrong connections before the compilation. When the design is validated, Vivado compiles the project through synthesis and implementation to generate the final output files. The two most important files for the PYNQ framework are:

- **The Bitstream file (.bit):** This file contains the binary configuration data that physically programs the hardware logic cells inside the FPGA.

- **The Hardware Handoff file (.hwh):** This file contains the architectural description of the design. It tells the Python environment the names, types, and memory addresses of the hardware blocks inside the FPGA. It is the file thanks to which we can communicate with the board via the PYNQ environment.

Once these files are transferred to the Linux system on the RFSoc 4x2 board, the overlay can be loaded dynamically using Python code. For example, the base overlay is initialized with the following command:

```
from pynq.overlays.base import BaseOverlay
base = BaseOverlay("base.bit")
```

To better understand how the Vivado Design Suite works, the block design of the base overlay was analyzed. This is the same overlay used for the signal acquisitions described in Section 5.2. However, before analyzing the specific block design of the manufacturer's base overlay, it is useful to clarify how the internal structure of the RFSoc chip is organized. While the development board includes many secondary components, the core of the system is divided into three main functional areas that work together.




Figure 5.7: Diagram of the three major functional components inside the Zynq RFSoc device architecture.

As shown in Fig. 5.7, the internal structure consists of:

- **The Processing System (PS):** Positioned on the left, this area contains the multi-

core ARM processors. It acts as the software environment where the embedded Linux OS, the PYNQ framework, and the Jupyter Notebook web application run.

- **The Programmable Logic (PL):** Positioned in the center, this is the flexible FPGA fabric. It is where custom hardware algorithms are physically implemented.
- **The RF Data Converters (RF DCs):** Positioned on the right, these are the high-speed analog-to-digital (ADC) and digital-to-analog (DAC) hardware blocks. They are organized into independent units called *tiles*. Because these converters are integrated directly onto the same silicon die next to the PL, they can stream data to the user logic with virtually zero delay.

With this hardware structure in mind, we can now analyze the complete block design configuration of the base overlay (Fig. 5.8). To understand how data flows, it is useful to analyze the main hardware blocks visible in the design.




Figure 5.8: Full block design of the manufacturer base overlay in AMD Vivado.

The design is divided into the following main components:

- **Zynq UltraScale+ MPSoC IP Core:** This is the central brain of the design, shown in detail in Fig. 5.9. It represents the physical processing system (PS) of the board, which runs the embedded Linux OS and the PYNQ framework. This block communicates directly with the Python environment. It is connected to all the other IP cores. These other IPs are the components that will physically form and configure the programmable logic (PL) section when the bitstream file is loaded.
- **RF Data Converter (RFDC) IP Core:** This block, included in the radio sub-




Figure 5.9: Detail of the Zynq UltraScale+ MPSoC processing system block.

system shown in Fig. 5.10, manages the high-speed radio frequency analog-to-digital converters (ADCs) and digital-to-analog converters (DACs).

- **Clocking Wizard and Synchronization:** Clock management is critical for data acquisition. The block design uses a Clocking Wizard IP to generate stable internal clock signals at specific frequencies. In the PYNQ environment, a dedicated software function is used to synchronize these internal hardware clocks with the external sampling systems to avoid timing mismatches.
- **Processor System Reset IP Blocks:** The design includes multiple reset blocks. These blocks handle the reset signals for the different clock domains in the system. They make sure that all the hardware logic cells and external interfaces start or restart in a synchronized state.
- **AXI Interconnect and Data Streams:** The blue internal lines in the diagram represent the communication network inside the chip. They use a standard communication ruleset called AXI. These lines are like a highway system where different components send data to each other. There are two main types of data traffic:
 - **AXI-Stream (High-Speed Data):** This path is fast and continuous. It is used for the real-time data coming from the ADCs. Because the data stream never stops, this connection has no addresses or delays; it just moves the digitized signal samples as fast as possible for processing.
 - **AXI-Lite (Control Signals):** This path is used for small, occasional messages, like sending setup commands or changing configurations. Here every message has a specific memory address and it is used by the Python environment to turn things on or off or to read eventual errors.

The blocks called AXI Interconnect act like routers on this highway. They manage all these lines, making sure that the data flows correctly between the FPGA logic, the processor system, and the external DDR (Double Data Rate) memory (the

high-speed RAM of the board).




Figure 5.10: Detailed view of the radio subsystem block containing the RF Data Converter configuration.

In conclusion, the analysis of this block design demonstrates how all the described components work together to form a complete hardware architecture. This integrated system allows the board to handle the high-speed acquisition, digital processing, and generation of radio frequency (RF) signals, directly connecting the external analog environment with the Python software framework.

6 | Conclusions and future developments

In this thesis work, a computational framework to overcome the mutual coherence limits in Dual-Comb Spectroscopy (DCS) was successfully investigated, providing a digital alternative to complex hardware stabilization. While fully phase-locked architectures offer superior mutual coherence, they require complex electronic feedback loops and strictly stabilized reference lasers, significantly increasing experimental complexity and cost. The results analyzed in this work demonstrate that a free-running configuration, when combined with software corrections, represents an advantageous trade-off. It drastically reduces hardware constraints while still delivering the high spectral resolution and mutual coherence necessary to retrieve clear molecular absorption features. This was achieved first by validating the self-correction pipeline offline: using experimental data from a Teledyne LeCroy oscilloscope, an algorithm based on inverse-time-warping and the Hilbert transform, extracted frequency and phase jitter directly from the interferogram, restoring full mutual coherence *a posteriori*. Subsequently, to pave the way toward a real-time correction, the algorithm was translated into a hardware-friendly FIR filter framework, the AMD RFSoc 4x2 was characterized through PYNQ and Vivado Design Suite, validating the data streaming between the RF Data Converters, the Programmable Logic, and the Processing System.

Moving forward, the work completed in this thesis opens up several paths for future technical improvements. At the software level, the immediate next step involves optimizing the current MATLAB post-processing scripts. Eventually merging the scripts into a single, automated software organized with a primary master file and modular sub-functions. This would enhance execution speed and user accessibility for daily spectroscopic measurements. However, the most crucial upgrade concerns the data processing itself. Now the corrections are performed post-processing, on pre-acquired data, which limits the measurement speed. To achieve live monitoring, future developments will focus on the acquisition of a dual comb signal with the RFSoc 4x2 board and the real-time implementation of the algorithm directly on the board itself. Following successful hardware

acceleration workflows found in literature, such as the work by Ossiander et al.[7]. In particular the creation of custom IP cores within Vivado to implement the corrections currently performed via MATLAB will be fundamental. These blocks will be integrated directly inside the hardware block design, allowing the board to process the data streams autonomously, replicating the exact data flow validated through the Python-based simulations analyzed in Section 5.1.

From a laser technology perspective, a particularly promising direction involves extending these dual-comb sources toward the mid-infrared (Mid-IR) spectral region. This would allow the system to target fundamental molecular transitions with strong absorption cross-sections, enhancing detection sensitivity for trace molecules of interest.

To sum up, these advancements align with the broader scientific effort to transition dual-comb spectroscopy from specialised laboratories towards more practical applications. It is important to remember that, in the long term, the ultimate goal is to develop compact instruments suitable for real-world scenarios, such as environmental monitoring, LiDAR and human breath analysis.

Bibliography

- [1] AMD Xilinx. Vivado Design Suite. <https://www.xilinx.com/products/design-tools/vivado.html>, 2026. Accessed: June 8, 2026.
- [2] R. Bajaj and S. A. Fahmy. Mapping for maximum performance on FPGA DSP blocks. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 35(4):573–585, 2015. doi: 10.1109/TCAD.2015.2474363.
- [3] B. Bernhardt. *Dual Comb Spectroscopy*. PhD thesis, Ludwig-Maximilians-Universität München, München, Germany, June 2011.
- [4] I. Coddington, W. C. Swann, and N. R. Newbury. Coherent dual-comb spectroscopy at high signal-to-noise ratio. *Physical Review A*, 82(4):043817, 2010. doi: 10.1103/PhysRevA.82.043817.
- [5] L. H. Crockett, D. Northcote, and R. W. Stewart, editors. *Software Defined Radio with Zynq UltraScale+ RFSoc*. Strathclyde Academic Media, first edition, 2023. URL <https://www.RFSocbook.com>.
- [6] M. Deschênes, P. Giaccari, and J. Genest. Optical referencing technique with cw lasers as intermediate oscillators for continuous full-delay-range frequency-comb interferometry. *Optics Express*, 18(22):23358–23370, 2010. doi: 10.1364/OE.18.023358.
- [7] A. Eber, C. Gruber, M. Schultze, B. Bernhardt, and M. Ossiander. Streaming self-corrected dual-comb spectrometer. *Optics Express*, 33(17):35314, 2025.
- [8] A. Eber, M. Pal, L. Fürst, E. Hruska, C. Gruber, M. Ossiander, and B. Bernhardt. Nir/vis dual-comb spectroscopy comparing high and low repetition rate regimes. *Laser & Photonics Reviews*, 2026. doi: 10.1002/lpor.202501010. Accepted manuscript.
- [9] T. Eber, J. Fürst, M. Ossiander, J. Schötz, T. W. Hänsch, and B. Bernhardt. Phase-locked feed-forward stabilization for dual-comb spectroscopy. *Ultrafast Science*, 2022: 9831742, 2022. doi: 10.34133/2022/9831742.
- [10] R. W. Fox, S. A. Diddams, A. Bartels, and L. Hollberg. Optical frequency measure-

- ments with the global positioning system: tests with an iodine-stabilized he–ne laser. *Applied Optics*, 44(1):113–120, 2005. doi: 10.1364/AO.44.000113.
- [11] D. Gatti, T. Sala, M. Marangoni, G. Galzerano, and L. Gianfrani. Precision molecular spectroscopy with frequency combs. In *Encyclopedia of Analytical Chemistry*. Wiley, Politecnico di Milano, Italy, 2014.
- [12] N. Hoghooghi, P. Chang, S. Egbert, et al. Ghz repetition rate mid-infrared frequency comb spectroscopy of fast chemical reactions. *Optica*, 11(6):876–882, 2024. doi: 10.1364/OPTICA.521655.
- [13] T. Ideguchi, A. Poisson, G. Guelachvili, N. Picqué, and T. W. Hänsch. Adaptive real-time dual-comb spectroscopy. *Nature Communications*, 5:4375, 2014. doi: 10.1038/ncomms4375.
- [14] C. Mashburn, K. F. Chang, M. J. Wahl, M. Walsh, D. I. Herman, M. Heyrich, T.-H. Wu, N. Hoghooghi, R. Sekine, L. Ledezma, E. Jerris, A. Marandi, J. Genest, and S. A. Diddams. Visible dual-comb spectroscopy across more than 100 thz with lithium niobate nanophotonic waveguides. *arXiv preprint arXiv:2602.07239*, 2026. doi: 10.48550/arXiv.2602.07239.
- [15] A. V. Oppenheim and R. W. Schaffer. *Discrete-Time Signal Processing*. Prentice Hall, Upper Saddle River, NJ, 3rd edition, 2010.
- [16] A. V. Oppenheim, A. S. Willsky, and S. N. Hamid. *Signals and Systems*. Prentice Hall, Upper Saddle River, NJ, 2nd edition, 1997.
- [17] M. Ossiander. Real-time fpga phase correction for dual-comb spectroscopy. https://github.com/marcus-o/real_time_rfsoc_phase_correction, 2023. Accessed: June 2026.
- [18] N. Picqué and T. W. Hänsch. Frequency comb spectroscopy. *Nature Photonics*, 13(3):146–157, 2019. doi: 10.1038/s41566-018-0347-5.
- [19] I. Pupeza, T. Eidam, J. Rauschenberger, B. Bernhardt, A. Ozawa, E. Fill, F. Krausz, T. W. Hänsch, J. Limpert, and R. Holzwarth. Power scaling of a high-repetition-rate enhancement cavity. *Optics Letters*, 35(12):2052–2054, 2010. doi: 10.1364/OL.35.002052.
- [20] G. B. Rieker, F. R. Giorgetta, W. C. Swann, J. a. Kofler, A. M. Zolot, L. C. Sinclair, E. Baumann, C. Cromer, G. Petron, C. Sweeney, P. P. Tans, I. Coddington, and N. R. Newbury. Frequency-comb-based remote sensing of greenhouse gases over kilometer air paths. *Optica*, 1(5):290–298, 2014. doi: 10.1364/OPTICA.1.000290.

- [21] A. Schliesser, M. Brehm, F. Keilmann, and D. W. van der Weide. Frequency-comb infrared spectrometer for rapid, remote chemical sensing. *Optics Express*, 13(22): 9029–9038, 2005. doi: 10.1364/OPEX.13.009029.
- [22] C. E. Shannon. Communication in the presence of noise. *Proceedings of the IEEE*, 86(2):447–457, 1998. doi: 10.1109/JPROC.1998.659497. Classic Paper Reprint.
- [23] W. C. Swann, I. Coddington, and N. R. Newbury. Fiber-laser frequency combs with subhertz relative linewidths. *Nature Photonics*, 1(1):48–51, 2006. doi: 10.1038/nphoton.2006.40.
- [24] *HDO6000 High Definition Oscilloscopes Operator’s Manual*. Teledyne LeCroy, Inc., 2013. Revision B, November 2013.
- [25] The Nobel Foundation. The nobel prize in physics 2005. <https://www.nobelprize.org/prizes/physics/2005/summary/>, 2005. Accessed: June 2026.

A | Appendix A

A.1. Fundamentals of Molecular Laser Spectroscopy

Laser spectroscopy is based on the interaction between electromagnetic radiation and matter. When a laser beam propagates through a medium, its intensity is attenuated as a result of absorption processes that are characteristic of the molecular species present in the sample. This macroscopic behavior is described by the Beer-Lambert law:

$$I(\nu, L) = I_0(\nu) \exp(-\alpha(\nu)L) \quad (\text{A.1})$$

where $I_0(\nu)$ represents the incident intensity at frequency ν , L is the optical path length, and $\alpha(\nu)$ is the absorption coefficient. This coefficient is the key parameter in spectroscopic analysis, as it contains the unique “spectral fingerprint” of the molecule, allowing for both its unambiguous identification and its characterization. At a microscopic level, this absorption phenomenon is a resonant process linked to the transition of the molecule between different energy states. In the infrared (IR) spectral region, these transitions primarily involve roto-vibrational energy levels. In the following sections, the molecular origin of these roto-vibrational transitions will be explained. Furthermore, the absorption spectrum of acetylene will be presented, as it represents a standard reference in this field and has been investigated also in this work. Finally, the limitations of conventional spectroscopic techniques will be discussed briefly, highlighting why Dual-Comb Spectroscopy can become a fundamental resource.

A.1.1. Molecular Energy Levels and Roto-Vibrational Transitions

Molecular spectroscopy is based on the quantization of energy states. The total internal energy of a molecule, E_{tot} , can be expressed as the sum of distinct contributions:

$$E_{tot} = E_{tr} + E_{el} + E_{vib} + E_{rot} \quad (\text{A.2})$$

where E_{tr} is the translational kinetic energy of the molecule, E_{el} represents the electronic energy, E_{vib} the vibrational energy of the nuclei, and E_{rot} the rotational energy of the molecule as a whole. The translational component E_{tr} is not considered further in this context. Indeed translational motion in free space is not quantized, so it does not involve discrete energy levels; consequently, it is not useful for the acquisition of a spectrum that is characteristic of the molecule. Conversely, the other three components are quantized according to the laws of quantum mechanics.

Transitions between these quantized levels occur when the energy of the photon (photons in some cases) matches the energy difference between the two quantized states. Namely when the so-called resonance condition is satisfied:

$$h\nu = E_f - E_i \quad (\text{A.3})$$

where $h\nu$ is the photon's energy (ν is the frequency of the radiation and h is Planck's constant) and E_f and E_i are the final and initial energy states, respectively. Since each molecular species possesses a unique set of energy levels, the resulting absorption spectrum acts as a characteristic signature.

For most molecules, the energy gaps between electronic states are relatively large, typically falling in the ultraviolet or visible regions. In contrast, the transitions between vibrational and rotational levels involve much lower energies. Specifically, vibrational transitions typically correspond to the infrared (IR) region of the electromagnetic spectrum, while rotational transitions, which require even less energy, appear as a fine structure superimposed on the vibrational peaks.

This leads to the formation of roto-vibrational bands, governed by specific selection rules ($\Delta v = \pm 1$ for vibration and $\Delta J = \pm 1$ for rotation). These rules dictate the distribution of the absorption lines into two main groups: the P-branch and the R-branch. Indeed a vibrational transition is almost always accompanied by a simultaneous change in the rotational state. If the transition involves a decrease in the rotational quantum number ($\Delta J = -1$), the energy required is slightly lower than the pure vibrational gap, resulting in absorption lines at lower frequencies, the P-branch. On the other hand, if the rotational state increases ($\Delta J = +1$), the photon must provide additional energy with respect to the pure vibrational transition, leading to the formation of the R-branch at higher frequencies. From what it has been discussed so far, it is clear that IR spectroscopy is one of the primary tools for probing these internal degrees of freedom of molecules, providing a direct access to the molecular structure.

B | Appendix B

B.1. Analytical Signal Representation and Hilbert Transform

The Hilbert transform $\hat{x}(t)$ of a continuous-time real signal $x(t)$, a fundamental tool in statistical communication and signal theory [16], is defined in the time domain as the convolution of the signal with the impulse response $1/\pi t$:

$$\hat{x}(t) = \mathcal{H}\{x(t)\} = x(t) * \frac{1}{\pi t} = \frac{1}{\pi} \int_{-\infty}^{+\infty} \frac{x(\alpha)}{t - \alpha} d\alpha \quad (\text{B.1})$$

This integral produces a version of $x(t)$ where all its harmonic components are shifted in phase. To better understand this operation, we examine its behavior in the frequency domain. By applying the Fourier Transform to $\hat{x}(t)$, we obtain the frequency domain representation $X_h(f)$:

$$X_h(f) = X(f) \cdot H(f) = X(f) \cdot [-j \operatorname{sgn}(f)] \quad (\text{B.2})$$

where $H(f) = -j \operatorname{sgn}(f)$ represents the transfer function of the Hilbert transformer. This function can be explicitly written as:

$$H(f) = \begin{cases} -j = e^{-j\pi/2} & f > 0 \\ 0 & f = 0 \\ +j = e^{+j\pi/2} & f < 0 \end{cases} \quad (\text{B.3})$$

From this expression, it is evident that the Hilbert Transform acts as an ideal phase shifter : it shifts the phase of every positive frequency component by -90° and every negative frequency component by $+90^\circ$.

Now that the mathematical definition has been revealed, it is important to understand why the Hilbert transform is such a vital tool for reconstructing analytic signals. This transform produces a complementary analytic signal that reveals amplitude and phase

information. Indeed the phase-shifting property of the Hilbert transform makes it ideal for constructing signals with only positive frequencies, known as analytic signals.

Consider a real signal $x(t)$, once its Hilbert transform $\hat{x}(t)$ is computed, the complex analytic signal $z(t)$ can be defined as:

$$z(t) = x(t) + j\hat{x}(t) \quad (\text{B.4})$$

This can also be rewritten in polar form:

$$z(t) = A(t)e^{j\phi(t)} \quad (\text{B.5})$$

Clearly, this complex signal containing only positive frequencies simplifies the analysis process considerably, making the extraction of the phase and envelope much easier. Indeed the computation will be just:

$$A(t) = |z(t)| \quad (\text{B.6})$$

$$\phi(t) = \arg(z(t)) \quad (\text{B.7})$$

To visualize better what is happening, it is useful to examine $z(t)$ in the frequency domain. Applying the Fourier Transform:

$$Z(f) = \mathcal{F}\{x(t) + j\hat{x}(t)\} = X(f) + j[-j\text{sgn}(f)X(f)] \quad (\text{B.8})$$

Simplifying the expression, we obtain:

$$Z(f) = X(f)[1 + \text{sgn}(f)] = \begin{cases} 2X(f) & f > 0 \\ X(0) & f = 0 \\ 0 & f < 0 \end{cases} \quad (\text{B.9})$$

This derivation shows that the analytic signal $z(t)$ is characterized by the total absence of negative frequency components. This property is illustrated in Figure B.1, where the spectrum of the real signal is compared with the analytic one.




Figure B.1: Spectra of the Real and the Analytic Signals.

C | Appendix C

C.1. Code Implementation and Software Framework

This appendix presents the Scripts developed and implemented throughout this thesis. To maintain clarity and focus on the algorithmic logic, the routines dedicated exclusively to data plotting, graphical user interface and file saving have been omitted. The appendix is divided into the MATLAB offline post-processing scripts and the Python-based (PYNQ) data acquisition program.

C.1.1. Script 1: The Self-Correction Framework

This script is the beginning of the post processing analysis. It performs the frequency and phase jitter corrections and the coherent avergaing. The output is the corrected trace in the RF frequencies.

```

1      % Pre-processing and Baseline Conditioning
2      tic
3      bgnd = sum(s_gpu)/length(s_gpu); % Compute the mean value to define the background
        offset
4      s_gpu = s_gpu - bgnd; % Perform background subtraction from the interferogram
        signal
5      clear bgnd
6
7      if interpolazione
8          t_interp=(interp1(t_full,t_full,linspace(t_full(1),t_full(end),N_interp*length(
                t_full))))'; % Define an upsampled time scale by increasing the number of
                points by an N_interp factor
9          s_gpu = spline(t_full, s_gpu, t_interp); % Generate the upsampled interferogram
                via cubic spline interpolation onto the new time scale
10         t_full = t_interp; % Update the reference time scale array
11         dt = t_full(2)-t_full(1); % Recompute the updated sampling time interval
                between the first two samples
12     end

```

```

13 tempo_pre_proc = toc;
14 disp(['Tempo impiegato per pre_processing dell interferogramma ', num2str(
    tempo_pre_proc), ' s']);
15
16 % Envelope Extraction and Burst Detection
17 tic
18 if envelope_smoothing
19     s_abs = smooth(abs(hilbert(s_gpu)),N_env_smooth); % Compute the smoothed
        envelope using the magnitude of the Hilbert transform
20 else
21     s_abs = abs(hilbert(s_gpu)); % Compute the raw envelope using the magnitude of
        the Hilbert transform
22 end
23
24 [~, locs] = findpeaks(s_abs, 'MinPeakProminence', prominence, 'MinPeakDistance'
    ,0.8*burst_size_estimate); % Extract peak locations of the envelope suppressing
        the magnitude output (~)
25 locs = gather(locs); % Transfer peak location indices from the GPU memory to the
        CPU host
26 s_abs = gather(s_abs); % Transfer the calculated envelope signal from the GPU
        memory to the CPU host
27 delta_f_rep_stim = gather((mean(diff(t_full(locs))))^-1); % Estimate the mean
        repetition frequency in Hz and download to CPU
28 T_rep = delta_f_rep_stim^-1; % Determine the mean repetition period between
        consecutive envelope peaks
29 % min_peak = min(pks);
30 tempo_burst_detection = toc;
31 disp(['Tempo impiegato per burst-detection con Hilbert-transform ', num2str(
    tempo_burst_detection), ' s']);
32
33 % === Array Allocation for Pre-Correction Visualizations and Export ===
34 N = numel(s_gpu); % Count the total number of sampled data elements
35 win = -burst_interval:burst_interval; % Define the index window size for plotting
36 % validPeaks = locs > burst_interval & locs <= N-burst_interval;
37 % locs_v = locs(validPeaks);
38 intervals = locs(2:end-1) + win; % Apply the defined index window symmetrically
        around the detected envelope peaks
39 intervals = sort(intervals(:)); % Sort and linearize the visualization interval
        indices
40 s_abs_sampled = s_abs(locs); % Extract the envelope peak intensities at the
        detected locations
41 t_sampled = t_full(locs); % Extract the time instances corresponding to the
        envelope peak locations
42

```

```

43     s_abs_plot = s_abs(intervals); % Isolate envelope intensities within the defined
        plotting intervals
44     s_raw = gather(s_gpu(intervals)); % Retrieve uncorrected interferogram intensities
        within the plot intervals to the CPU
45     t_raw = t_full(intervals); % Define time instances associated with the uncorrected
        plot intervals
46
47     % Pre-Correction Jitter Metrics Computation
48     Tr_vect = diff(t_full(locs)); % Compute individual repetition periods between
        successive envelope peaks
49     Delta_fr_vect = Tr_vect.^-1; % Calculate local repetition frequencies between
        successive envelope peaks
50     Delta_fr_jitter_vect = diff(Delta_fr_vect); % Compute local frequency variations
        between successive peaks (frequency jitter)
51     Tr_jitter_vect = diff(Tr_vect); % Compute local timing variations between
        successive peaks (timing jitter)
52     t_delta = t_full(locs(1:end-1));
53     t_jitter = t_full(locs(1:end-2));
54     Tr_avg = ones(size(t_delta))*sum(Tr_vect)/length(Tr_vect); % Generate an array
        filled with the mean repetition period value
55     delta_fr_avg = ones(size(t_delta))*sum(Delta_fr_vect)/length(Delta_fr_vect); %
        Generate an array filled with the mean repetition frequency value
56
57     % Time-Domain Frequency Jitter Correction (Time-Warping Algorithm)
58     tic
59     if auto_interferogram_length
60         N_peaks_tot = length(locs); % Calculate the total number of detected envelope
            peaks
61         pos_cut = gpuArray(locs(1):locs(N_peaks_tot)); % Isolate and allocate indices
            between the first and last envelope peaks on the GPU
62         n_BM = ceil((length(pos_cut)-length(locs))/(length(locs)-1)); % Calculate mean
            number of cells between consecutive peaks (BM = Between Maxima)
63         if balanced
64             for k = 1:N_peaks_tot
65                 idx_burst_min = locs(k)-floor(n_BM*0.005);
66                 idx_burst_max = locs(k)+floor(n_BM*0.005);
67                 centroid(k) = floor(sum((s_abs((idx_burst_min:idx_burst_max).')).^2.*((
                    idx_burst_min:idx_burst_max).'))./(sum((s_abs((idx_burst_min:
                    idx_burst_max).')).^2))));
68             end
69             locs = centroid';
70             pos_cut = gpuArray(locs(1):locs(N_peaks_tot)); % Re-isolate and allocate
                updated centroid indices between peaks on the GPU
71         end

```

```

72     L_exact = n_BM*(length(locs)-1)+length(locs); % Compute the ideal total number
        of data points including the peak positions
73     locs_cut = 1:n_BM+1:L_exact; % Establish new equidistant indices for the
        envelope peaks to linearize spacing
74     dt = (t_full(pos_cut(end))-t_full(pos_cut(1)))/(L_exact-1); % Calculate the
        newly uniform sampling time interval based on ideal length
75 else
76     pos_cut = gpuArray(locs(1):locs(N_peaks_tot)); % Isolate and allocate indices
        between the first and last envelope peaks on the GPU
77     n_BM = L_burst-1; % Set the mean number of cells between successive peaks
        according to user-defined L_burst
78     if balanced
79         for k = 1:N_peaks_tot
80             idx_burst_min = locs(k)-floor(n_BM*0.005);
81             idx_burst_max = locs(k)+floor(n_BM*0.005);
82             centroid(k) = floor(sum((s_abs((idx_burst_min:idx_burst_max).')).^2.*((
                idx_burst_min:idx_burst_max).'))./(sum((s_abs((idx_burst_min:
                idx_burst_max).')).^2))));
83         end
84         locs = centroid';
85         pos_cut = gpuArray(locs(1):locs(N_peaks_tot)); % Re-isolate and allocate
            updated centroid indices between peaks on the GPU
86     end
87     L_exact = (N_peaks_tot-1)*L_burst+1; % Compute the ideal total number of data
        points based on user-specified burst length
88     locs_cut = 1:L_burst:L_exact; % Establish new equidistant indices for the
        envelope peaks based on fixed L_burst
89     dt = dt*length(t_full)/L_exact; % Scale the sampling time interval to account
        for the total ideal document length
90 end
91
92 if freq_corr
93     t_cut = dt*((1:L_exact)'-1) + t_full(pos_cut(1)); % Generate a uniform,
        truncated time grid with constant spacing dt
94     p_locs = gpuArray(linspace(1, N_peaks_tot, N_peaks_tot)); % Initialize an index
        vector for the envelope peaks on the GPU
95     p_locs_interp = (linspace(1, N_peaks_tot, L_exact))'; % Define an upsampled
        linear index grid corresponding to uniform peak spacing
96     locs_interp = gpuArray((spline(p_locs, locs(1:N_peaks_tot), p_locs_interp))); %
        Compute interpolated peak trajectories onto the uniform grid via cubic
        splines on the GPU
97     s_analytic_corr = (spline(pos_cut, hilbert(s_gpu(pos_cut)), locs_interp)); %
        Resample the analytical signal via cubic spline interpolation to eliminate
        frequency jitter

```



```

98     s_gpu = gather(s_gpu); % Retrieve the raw uncorrected interferogram data back
        to the CPU memory
99     locs_interp; % Deallocate the interpolation array from GPU memory to free
        resources
100 else % Bypass frequency jitter correction algorithms entirely
101     t_cut = t_full(pos_cut);
102     s_analytic_corr = hilbert(s_gpu(pos_cut));
103     locs_cut = locs-locs(1)+1;
104 end
105 clear t_full;
106 tempo_f_jitter = toc;
107 disp(['Tempo impiegato per correzione Jitter di frequenza ', num2str(tempo_f_jitter
    ), ' s']);
108
109 % Phase Jitter Detection and Phase-Unwrapping Correction
110 tic
111 phase_pks = unwrap(angle(s_analytic_corr(locs_cut))); % Evaluate and unwrap the
        instantaneous phase at each linearized burst maximum
112 if phase_corr_spline
113     phase_vector = spline(t_cut(locs_cut), phase_pks, t_cut); % Track intra-burst
        phase evolution using a cubic spline interpolation
114 else
115     phase_vector = interp1(t_cut(locs_cut), phase_pks, t_cut); % Track intra-burst
        phase evolution using a linear 1D interpolation
116 end
117
118 if phase_correction
119     s_corr_tot = real(s_analytic_corr .* exp(-1i * phase_vector)); % Demodulate the
        analytic signal by its instantaneous phase to cancel phase jitter
120 else
121     s_corr_tot = real(s_analytic_corr); % Maintain the real part of the frequency-
        corrected signal bypassing phase demodulation
122 end
123
124 %Burst Averaging
125 for k = 2:length(locs_cut)-1
126     idx_center = locs_cut(k); % Isolate the index corresponding to the current
        burst maximum
127     idx_range = idx_center - burst_left_length : idx_center + burst_right_length; %
        Define the index range for window slicing around the center
128     burst(:,k-1) = s_corr_tot(idx_range); % Extract the windowed segment and
        populate the columns of the burst matrix
129     burst_avg = burst_avg + burst(:,k-1); % Coherently superimpose individual burst
        waveforms to increase SNR

```

```

130     end
131     tempo_f_jitter = toc;
132     burst_avg = burst_avg/(length(locs_cut)-2); % Compute the mathematical mean of all
           accumulated burst profiles

```

C.1.2. Script 2: Multi-Trace Averaging

This script aggregates and averages a successive series of pre-corrected files to increase the final Signal-to-Noise Ratio (SNR).

```

1  for k = 1:2
2      amps = results(k).amps; % Extract the amplitudes of the current dataset
3      freqs = results(k).freqs; % Extract the frequencies of the current dataset
4
5      local_indices_to_exclude = results(k).indices_to_exclude; % Locally define the
           structure element containing the file indices to be excluded
6
7      amp_ref = amps(:,1); % Set the amplitudes of the first file as the reference
           amplitudes
8      amp_ref_smooth = smooth(amp_ref, 5); % Apply smoothing to prevent noise
           amplification during differentiation
9      diff_ref = diff(amp_ref_smooth); % Compute the derivative of the reference
           amplitudes
10     freq_ref = freqs(:,1); % Set the frequencies of the first file as the reference
           frequencies
11     n_punti = size(amps, 1); % Calculate the total number of rows (data points)
12
13     for i = 2:N_files
14         if ismember(file_numeric_ids(i), local_indices_to_exclude)
15             continue;
16         end % Skip the file if it is already included in the indices_to_exclude
           variable
17
18         current_amp = amps(:,i); % Amplitude vector of the current file
19         current_amp_smooth = smooth(current_amp, 5); % Apply smoothing to the current
           spectrum
20         diff_curr = diff(current_amp_smooth); % Compute the derivative of the current
           spectrum
21         [c, lags] = xcorr(diff_ref - mean(diff_ref), diff_curr - mean(diff_curr), '
           coeff'); % Compute the cross-correlation between the two derivatives,
           normalizing the result
22         [~, max_idx] = max(c); % Locate the index of the maximum correlation value
23         delta_pos = lags(max_idx); % Identify the shift corresponding to the maximum
           correlation peak

```

```

24
25     % Apply the calculated shift to a temporary amplitude variable
26     temp_amp = current_amp; % Save the current file's spectrum amplitude into a
        temporary variable
27     if delta_pos > 0 % Apply a rightward shift to the current spectrum
28         temp_amp(delta_pos+1:n_punti) = current_amp(1:n_punti-delta_pos);
29         temp_amp(1:delta_pos) = current_amp(1);
30     elseif delta_pos < 0 % Apply a leftward shift to the current spectrum
31         abs_lag = abs(delta_pos);
32         temp_amp(1:n_punti-abs_lag) = current_amp(abs_lag+1:n_punti);
33         temp_amp(n_punti-abs_lag+1:n_punti) = current_amp(n_punti);
34     end
35
36     R = corrcoef(amp_ref, temp_amp); % Calculate the spectral overlap as a
        correlation coefficient
37     overlap = R(1,2); % Extract the overlap value from the correlation matrix
38
39     if overlap < overlap_threshold % Check if the overlap is below the predefined
        threshold
40         fprintf('\n[!] File %d REJECTED: Overlap (%.4f) below threshold',
            file_numeric_ids(i), overlap);
41         local_indices_to_exclude = unique([local_indices_to_exclude,
            file_numeric_ids(i)]); % Append the rejected file index to the exclusion
            vector for the averaging stage
42     else % The file is valid and meets quality criteria
43         % Permanently save the aligned amplitude spectrum of the current file
44         amps(:,i) = temp_amp;
45         temp_freq = freqs(:,i);
46
47         % Apply the corresponding shift to the frequency axis
48         if delta_pos > 0 % Apply a rightward shift to the current spectrum
49             freqs(delta_pos+1:n_punti, i) = temp_freq(1:n_punti-delta_pos);
50             freqs(1:delta_pos, i) = temp_freq(1);
51         elseif delta_pos < 0 % Apply a leftward shift to the current spectrum
52             abs_lag = abs(delta_pos);
53             freqs(1:n_punti-abs_lag, i) = temp_freq(abs_lag+1:n_punti);
54             freqs(n_punti-abs_lag+1:n_punti, i) = temp_freq(n_punti);
55         end
56         fprintf('\nFile %d: Aligned (Shift: %d, Overlap: %.4f)', file_numeric_ids(i)
            ), delta_pos, overlap);
57     end
58 end
59

```

```

60     results(k).amps = amps; % Save the aligned amplitudes in the current dataset
        structure
61     results(k).freqs = freqs; % Save the aligned frequencies in the current dataset
        structure
62     results(k).indices_to_exclude = local_indices_to_exclude; % Update the structure
        element with the excluded file indices, incorporating those rejected due to
        insufficient overlap
63     valid_mask = ~ismember(file_numeric_ids, results(k).indices_to_exclude); % Compare
        file IDs against the excluded files list
64     valid_indices = find(valid_mask); % Define the array indices of the accepted files
65     range_ok = find(diff(freq_ref) ~= 0); % Discard edge-padding cells at the beginning
        and end of the column
66
67     if plot_postcorrection
68         figure(k+2)
69         hold on
70         grid on
71         colors = lines(N_files);
72         for i = valid_indices
73             plot(freq_ref(range_ok), amps(range_ok, i), ...
74                 'Color', colors(i,:), ...
75                 'DisplayName', ['ID: ' num2str(file_numeric_ids(i))]); % Plot the
                superposition of aligned spectra (post-correction) that are part of
                the accepted list
76         end
77         title(['Post-correction shifted spectra - ' char(results(k).name)]);
78         xlim(results(k).freq_lim);
79         legend('show', 'Location', 'bestoutside');
80         hold off
81     end
82
83     results(k).amp_avg_ok = mean(amps(range_ok, valid_indices), 2); % Compute the
        mathematical mean of the corrected spectra, restricting it to the valid range
84     results(k).freq_ref_ok = freq_ref(range_ok); % Define the reference frequencies
        within the valid range
85     results(k).range_ok = range_ok; % Save the indices of the valid range
86     results(k).valid_indices = valid_indices; % Save the indices of the accepted files
87
88     figure(k+4)
89     plot(results(k).freq_ref_ok, results(k).amp_avg_ok); % Plot the final averaged
        spectrum computed across all valid files and corrected via cross-correlation
        shift
90     title(['Corrected averaged spectrum - ' char(results(k).name)]);
91     xlabel('Frequency');

```

```

92     ylabel('Amplitude');
93     grid on
94     xlim(sort(results(k).freq_lim));
95 end

```

C.1.3. Script 3: Absolute Frequency Calibration

The aim of this script is to map the radio-frequency axis into an absolute optical frequency scale.

```

1     % Compute the absolute RF beatnote difference between the two optical combs
2     delta_f_b = abs(fb_Asx * sign_fb_Asx - fb_Obx * sign_fb_Obx);
3
4     % Evaluate absolute optical frequency anchored to the CW transfer laser
5     nu_Asx = nu_idfix + sign_fb_Asx * fb_Asx + mode_jump * fr_Asx;
6
7     f_Ny = fr_Asx / 2;           % Define the Nyquist frequency limit
8     M = fr_Asx / delta_f_rep;    % Determine the spectral interleaving scaling factor
9     sign_n_fsr0 = sign_spectrum * (-1)^n_fsr; % Track the spectral orientation inside
    the Nyquist window
10
11    % Optical Frequency Axis Mapping and Conversion to Wavenumbers (cm^-1)
12    for k = 1:numel(spectra)
13        name_spectrum = spectra{k};
14        f_RF_k = data(i).(name_spectrum).x(:); % Extract the raw RF frequency grid
15        spettro_k = data(i).(name_spectrum).y(:); % Extract the raw spectral
    intensities
16
17        % Map the RF axes to absolute optical frequencies and convert to cm^-1 (c =
    29979245800 cm/s)
18        if sign_n_fsr0 == +1 % Direct spectral orientation mapping
19            nu_abs_k = (nu_Asx + M * ...
20                (f_RF_k * (-1)^n_fsr - delta_f_b + ...
21                (n_fsr + 0.5 * (1 + (-1)^(n_fsr + 1))) * f_Ny)) / 29979245800;
22        elseif sign_n_fsr0 == -1 % Inverted spectral orientation mapping
23            nu_abs_k = (nu_Asx + M * ...
24                (f_RF_k * (-1)^(n_fsr + 1) + delta_f_b + ...
25                (n_fsr - 0.5 * (1 + (-1)^(n_fsr + 1))) * f_Ny)) / 29979245800;
26        end
27
28        % Store calibrated data points
29        data(i).(name_spectrum).nu_abs = nu_abs_k;
30        data(i).(name_spectrum).calibrated_values = [nu_abs_k, spettro_k];
31    end

```

C.1.4. Script 4: Baseline Extraction and Spectrum Normalization

This final module normalizes the corrected and calibrated spectrum, leading to data ready to be studied and compared.

```

1  for k = 1:numel(data)
2      spectra(k).name = data(k).name;
3      spectra(k).freqs = data(k).values(:,1); % Copy spectrum frequencies into the
        spectra structure
4      spectra(k).amps = data(k).values(:,2); % Copy spectrum amplitudes into the spectra
        structure
5
6      % Invert spectrum frequencies if the spectral orientation is reversed
7      if mean(diff(spectra(k).freqs)) < 0
8          spectra(k).freqs = flip(spectra(k).freqs);
9          spectra(k).amps = flip(spectra(k).amps);
10     end
11 end
12
13 for k = 1:numel(spectra)
14     y = spectra(k).amps(:); % Copy spectrum amplitudes into a vector
15     L = length(y); % Calculate the total number of available data points
16
17     deltaF = deltaF / 1e12;
18     df = median(abs(diff(spectra(k).freqs))); % Compute the median frequency spacing of
        the sampled spectrum
19     winMax = round(deltaF / df); % Define the window size for constructing the upper
        envelope
20     winMax = 2*floor(winMax/2) + 1; % Force the window size to be an odd number
21
22     winMax = max(winMax, 3); % Ensure the window contains at least 3 points
23
24     % Ensure the window size does not exceed the total available data length
25     if winMax > L
26         winMax = L;
27         winMax = 2*floor(winMax/2) + 1;
28     end
29
30     spectra(k).df = df; % Save the median sampling frequency interval (df)
31     spectra(k).winMax = winMax; % Save the upper envelope window size (winMax)
32
33     y0 = smoothdata(y, 'movmedian', winMedian); % Apply a moving median filter over
        winMedian points

```

```

34     y0 = movmax(y0, winMax); % Construct the upper envelope using a moving maximum
        window of winMax size
35     y0 = smooth(y0, winSmooth, 'rloess'); % Smooth the envelope using a robust local
        regression (rloess) filter
36
37     D = diff(speye(L), 2); % Create a second-order difference matrix to evaluate
        baseline curvature
38     w = ones(L,1); % Initialize data point weights
39
40     for iter = 1:niter % Execute the iterative asymmetric least squares routine for
        niter loops
41         W = spdiags(w, 0, L, L); % Construct a diagonal sparse matrix containing the
            point weights
42         A = W + lambda * (D' * D); % Define the fitting matrix combining data fidelity
            and baseline smoothness
43         b = A \ (w .* y0); % Compute the baseline solution via sparse matrix division
44
45         % Dynamic weight adaptation:
46         % Points below baseline = absorption features to reject -> assign small weight
            p
47         % Points above baseline = informative envelope baseline -> assign large weight
            1-p
48         w = p * (y0 < b) + (1 - p) * (y0 >= b);
49     end
50
51     spectra(k).baseline = b; % Save the calculated baseline into the spectra structure
52     spectra(k).amps_corr = y ./ b; % Compute and save the normalized transmittance
        spectrum amplitudes
53 end

```

C.1.5. Python PYNQ Online Acquisition Script

The following section presents the script executed within the Jupyter Notebook environment to manage the wave acquisition with the board RFSoc 4x2. This script runs on the internal ARM processing system of the AMD RFSoc 4x2 board via the PYNQ framework. It handles the high-level initialization of the base hardware overlay and uses the basic data acquisition functions from the base Overlay.

```

1 import plotly.graph_objs as go
2 from plotly.subplots import make_subplots
3 import numpy as np
4 import ipywidgets as ipw
5 import matplotlib.pyplot as plt
6 from pynq.overlays.base import BaseOverlay

```

```

7
8 base = BaseOverlay('base.bit')
9
10 base.init_rf_clks()
11
12 c_r = 3 #receiver channel
13 #c_t = 0 #transmitter channel
14 probe_len = 256
15 N_samples = 30000
16 treshold = 0.1 #da tarare
17 sample_freq = 2457.6e6
18
19 #Set the transmitted frequencies
20
21 #base.radio.transmitter.channel[c_t].dac_block.MixerSettings['Freq'] = 1100 # MHz
22
23 #Set the transmitter gain and enable it
24
25 #base.radio.transmitter.channel[c_t].control.gain = 0.5
26 #base.radio.transmitter.channel[c_t].control.enable = False
27
28 #Set receiver's mixer freq. (default = 1228.8 MHz)
29 adc = base.radio.receiver.channel[c_r].adc_block
30 adc.MixerSettings['CoarseMixFreq'] = 0 # forza bypass del coarse
31 adc.MixerSettings['Freq'] = 30 # MHz
32 adc.UpdateEvent(1)
33
34 #probe
35 probe = np.array(base.radio.receiver.channel[c_r].transfer(probe_len))
36 peak = np.max(np.abs(probe))
37
38 #acquisition
39 if peak > treshold:
40     cdata = base.radio.receiver.channel[c_r].transfer(N_samples)
41
42 #Plot channel in time domain
43 t = np.arange(0, N_samples / sample_freq, 1 / sample_freq)
44 fig = go.Figure()
45 fig.add_trace(go.Scatter(x=t, y=np.real(cdata), name='I'))
46 fig.add_trace(go.Scatter(x=t, y=np.imag(cdata), name='Q'))
47 fig.update_layout(title=f'Time Domain ADC Channel {c_r}',
48                   xaxis_title='Seconds',
49                   yaxis_title='Amplitude')
50 fig

```



```
51
52 base.radio.receiver.channel[c_r].adc_block.MixerSettings
53 #nel grafico ho una freq. che esattamente il doppio di quella che imposto sul
   waveform generator
54
55 #Fourier Transform
56
57 f = np.fft.fftshift(np.fft.fftfreq(N_samples, d=1/sample_freq)) #asse frequenze
58
59 spectrum = np.fft.fftshift(np.fft.fft(np.real(cdata)))
60
61 # modulo (ampiezza)
62 amp = np.abs(spectrum)
63
64 # plot
65 fig = go.Figure()
66 fig.add_trace(go.Scatter( x=f, y=amp, mode='lines',name='|FFT|'))
67 fig.update_layout(
68     title=f'Frequency Domain  ADC Channel {c_r}',
69     xaxis_title='Frequency [Hz]',
70     yaxis_title='Amplitude')
71
72 fig.show()
```


List of Figures

1.1	Principle of molecular spectroscopy using an optical frequency comb: the source spectrum is modified by the absorption lines of the sample. Adapted from [18].	4
1.2	Time-domain representation of a pulse train. The electric field $E(t)$ and time t are expressed in arbitrary units (a.u.). Adapted from [18].	5
1.3	Overview of an optical frequency comb. The top panel shows the pulse train in the time domain, while the bottom panel illustrates the corresponding teeth in the frequency domain, shifted by the offset frequency f_{ceo} . Adapted from [18].	8
1.4	Schematic representation of the Chirped Pulse Amplification (CPA) technique, illustrating the stretching, amplification, and compression stages of the pulse. Adapted from (Schimpf, 2010).	10
1.5	Example of supercontinuum generation through numerically simulated spectral broadening in a standard telecom fiber. The simulation shows the evolution of 1 ps pulses with 50 nJ energy as they propagate through the non-linear medium. Adapted from (RP Photonics).	11
1.6	Schematic diagram of the complete dual-parameter active feedback loop for frequency comb stabilization. The error signals for both f_{rep} and f_{ceo} are processed by electronic controllers to generate real-time correction voltages. These voltages drive physical actuators within the cavity for cavity length adjustments (f_{rep}) and a pump power modulation system for carrier-envelope offset control (f_{ceo}).	14
1.7	Schematic block diagram of the $f - 2f$ interferometer setup using a dichroic mirror to split the octave-spanning spectrum into longer-wavelength and shorter-wavelength arms.	15

1.8	Spectral representation of the $f - 2f$ self-referencing technique for the two modes where the condition $2m = n$ is satisfied. The fundamental comb (red) contains modes ν_1 and ν_2 . The mode ν_1 is frequency-doubled via SHG (SH-OFC, blue). The heterodyne beating between $2\nu_1$ and the native fundamental mode ν_2 isolates the carrier-envelope offset frequency f_{ceo} . Adapted from Keller, U. (2021). Optical Frequency Comb from Modelocked Lasers. In: Ultrafast Lasers. Graduate Texts in Physics. Springer.	16
1.9	Scheme of a FTIR Spectrometer.	17
1.10	Scheme of a VIPA Spectrometer.	18
2.1	Dual Comb Spectroscopy general scheme.	19
2.2	Interferograms (bursts) observed at the photodiode, resulting from the interference between pulses from the two lasers.	21
2.3	Spectrum of the two optical frequency combs (three modes are shown for simplicity). The corresponding beatings spectrum is reported below. Beatings arising from the interference between modes of the same order are highlighted in red.	22
2.4	Visual representation of timing jitter in a dual-comb interferogram train. The lower trace illustrates the ideal, unperturbed scenario where bursts occur at perfectly equidistant intervals with a constant repetition period $T_r = 1/\Delta f_{rep}$. The upper trace demonstrates the effect of timing jitter, where fluctuations in Δf_{rep} lead to random variations in the arrival time of consecutive bursts, causing a dynamic distortion of the temporal mapping axis.	25
2.5	Visual representation of phase jitter in dual-comb spectroscopy bursts, assuming an underlying harmonic RF comb configuration ($f_0 = 0$). (A) In the presence of phase jitter, the carrier oscillations randomly shift beneath the pink envelope. (B) In the phase-stabilized scenario, the phase relation is rigidly locked, ensuring that the fringes remain perfectly stationary and identical from burst to burst.	26
2.6	Example of an optical feed-forward stabilization scheme using an Acousto-Optic Modulator (AOM). In this configuration, the relative fluctuations of f_{rep} and f_{CEO} are measured via auxiliary beat-notes, typically requiring gratings to isolate spectral portions and high-speed detectors. These error signals are processed by a control unit that drives an external AOM to shift the frequency of the slave comb in real-time, ensuring mutual coherence without affecting the laser cavity parameters [19].	28

2.7	Impulse response $\hat{x}_D[n]$ of the discrete Hilbert transformer.	32
2.8	Modern FPGA architecture showing different basic components (Adapted from [2])	34
3.1	Scheme of the dual-comb spectroscopy setup. The pulses from two independent 1-GHz Erbium solid-state lasers are temporally stretched by polarization-maintaining dispersion-compensating fibers (PM-DCF) and amplified via Erbium-Doped Fiber Amplifiers (EDFA). After re-compression through diffraction gratings, supercontinuum (SC) generation is performed. The broadened output of the first comb interrogates the reference gas cell, whereas the second comb acts as a local oscillator. To provide optical frequency calibration, a CW laser is combined alongside the combs via a beam splitter, subsequently the signal goes through a narrow-band filter and is detected by a photodiode and saved via an oscilloscope. The multi-heterodyne beating signal goes instead through a $4f$ filter and is then captured by a fast photodiode and recorded using a digital oscilloscope.	38
3.2	Portion of the real setup dedicated to chirped pulse amplification.	39
3.3	Commercial supercontinuum generation module (Cosmo by Octave Photonics) used in the experimental setup.	39
3.4	Experimental optical spectrum of the generated dual-comb supercontinuum for different EDFA drive currents. The emission spans over a full octave, covering a broad spectral window from approximately 1000 nm up to 2400 nm (corresponding to a wavenumber range from roughly 4160 cm^{-1} to 10000 cm^{-1}). This broadening completely covers the target molecular absorption bands, including the acetylene (C_2H_2) transitions centered around $1.55 \mu\text{m}$ ($\approx 6450 \text{ cm}^{-1}$).	40
3.5	Photograph of the initial section of the experimental setup, showing the two laser sources and the beam entering the molecular gas absorption cell.	41
3.6	View of the final section of the dual-comb spectrometer layout, capturing the path from the exit of the gas cell down to the high-speed photodetector	42
3.7	The Teledyne LeCroy HDO6054 high-definition digital oscilloscope used as the benchmark data acquisition system for offline multi-heterodyne signal processing.	44
3.8	The AMD Xilinx Zynq UltraScale+ RFSoc 4x2 development board configured with a Gen 3 ZU48DR device.	45

4.1	Experimental profiles of the tracked instabilities across the 20 ms acquisition window. (Top panel) Relative timing jitter.(Bottom panel) Cumulative carrier-envelope phase jitter $\Phi(t)$	50
4.2	Continuous carrier-envelope phase correction vector $\Phi(t)$ obtained after linear interpolation across the entire 20 ms window.	51
4.3	Comparison of the multi-heterodyne interferogram train over a localized temporal window. (Top panel) The original uncorrected waveform. (Bottom panel) The fully corrected interferogram waveform highlights the homogeneity and uniform periodicity achieved after the computational time and phase corrections.	51
4.4	Zoomed view of a single burts before and after the timing and phase jitters corrections. It is evident how the carrier envelope phase was canceled. . . .	52
4.5	Overlay of individual multi-heterodyne interferogram bursts. Due to the action of the correction all sequential bursts are successfully forced into complete mutual coherence, showing excellent constructive alignment. . . .	52
4.6	Comparison demonstrating the impact of coherent time-domain averaging on the interferogram tails. (Top panel) Overlay of the un-averaged individual segmented traces; far from the main burst center, the molecular signature is completely masked by the high noise floor. (Bottom panel) Resulting coherently averaged single burst profile; the random noise contributions are actively suppressed, revealing the molecular absorption lines of the sample.	53
4.7	Multi-heterodyne frequency comb structure before (blue) and after (red) digital computational correction. In the uncorrected spectrum (blue), the individual longitudinal modes are severely broadened. Following the self-correction processing (red), the teeth collapse back into sharp discrete spectral components.	53
4.8	Direct and Inverted RF spectrum	55
4.9	Possible positions of the CW laser with respect to the different combs	56
4.10	Overlap of the calibrated optical spectra after the cross-correlation alignment processing.	58
4.11	Final averaged and corrected optical spectrum of the sample (cell H). By linearly averaging the previously aligned datasets, the result is a drastically improved signal-to-noise ratio (SNR).	59
4.12	Acetylene spectrum (blue) with the calculated non-flat laser baseline (red). . . .	60
4.13	Final normalized transmittance spectrum of the acetylene sample, showing a unity transmittance.	61

4.14	Normalized experimental spectrum of acetylene overlaid with the HITRAN database simulation.	61
5.1	Uncompensated raw signal.	66
5.2	Corrected output stream.	66
5.3	Comparison of the complete dual-comb interferogram time-domain sequences before and after the algorithmic execution, plotted against the number of samples. The analysis was conducted using the reference acetylene dataset provided in the authors' official GitHub repository [17]	66
5.4	Laboratory experimental setup with the external Hewlett Packard 33120A waveform generator connected via SMA cable to the RFSoc 4x2 receiver channel input.	67
5.5	Acquisition with low sample size ($N_{\text{samples}} = 256$, $f_s = 2457.6$ MHz) for a 100 MHz single-tone input. The short observation time window captures fewer harmonic cycles in the time domain and results in broad, less defined frequency peaks in the FFT.	68
5.6	Acquisition with increased sample size ($N_{\text{samples}} = 30000$, $f_s = 2457.6$ MHz) for a 100 MHz single-tone input. The longer time axis allows the system to capture a dense train of oscillations, drastically narrowing the symmetric FFT peaks into ultra-sharp lines at ± 100 MHz.	68
5.7	Diagram of the three major functional components inside the Zynq RFSoc device architecture.	70
5.8	Full block design of the manufacturer base overlay in AMD Vivado.	71
5.9	Detail of the Zynq UltraScale+ MPSoC processing system block.	72
5.10	Detailed view of the radio subsystem block containing the RF Data Converter configuration.	73
B.1	Spectra of the Real and the Analytic Signals.	85

List of Tables

4.1 Summary of sign assignments for the CW beatings corresponding to the
scenarios illustrated in Fig.4.9. 55

Acknowledgements

I would like to thank my advisor, Prof. Alessio Gambetta, and my co-advisors for the opportunity to work on this project and for providing the laboratory resources.

I would also like to extend a special thank you to my parents, Pierluigia and Graziano, as none of this would have been possible without their help. You have always actively supported me throughout my years of study, you have literally listened to me revising before tests and exams for 10 years, so this achievement of mine is yours too.

I'd also like to say a huge thank you to my boyfriend, Edoardo, who has always been there for me whenever I've had to study and who has looked after me through all my difficult times in his own unique way.

This thesis marks the end of a chapter, and I can't wait to see what the next one has in store for me, knowing that you will all continue to shower me with love in the future, just as you always have.

