An experimental investigation of ammonia combustion at subzero temperatures

Ignition and early flame propagation study of NH_3 and NH_3-blends in a custom-designed combustion chamber

Master's thesis

Alberto Francesco Aversa

Delft University of Technology

TU Delft

UMASS LOWELL

An experimental investigation of ammonia combustion at subzero temperatures

Ignition and early flame propagation study of NH_3 and NH_3-blends in a custom-designed combustion chamber

by

Alberto Francesco Aversa

to obtain the degree of Master of Science
at the Delft University of Technology,
in collaboration with the University of Massachusetts Lowell,
to be defended publicly on February 13 2026.

TU Delft Student number:	5334918	
UML Student number:	10299894	
Project duration:	March 17, 2025 – Febrary 13, 2026	
Thesis committee:	Dr. F. De Domenico,	TU Delft, Supervisor
	Dr. ir. M. Pini,	TU Delft, Chair
	Prof. dr. V. Grewe,	TU Delft
	Prof. J. H. Mack,	UML, External Supervisor

Cover: Combustion flame [1]

TU Delft

UMASS
LOWELL

Preface

For my thesis, I initially had no idea what topic to pursue and was extremely undecided until only a few months before starting. The only thing I knew was that I was really curious to see what the research environment in the USA looked like. I knew that combustion was one of the topics at the top of the list, so when the opportunity to work in the Energy and Combustion Research Laboratory at the University of Massachusetts Lowell arose, I did not think about it too much. The project seemed extremely intriguing, and the idea of achieving measurements at subzero temperatures fascinated me.

The project turned out to be very interesting (albeit with quite more work than I believed), but it also tested me as a person and as a researcher: for the first time, I had to become a true problem-solver, with every day bringing new challenges and small wins.

Overall, I'm very grateful for the experience and all that I have learnt, as I've grown so much as a person and developed such a varied set of skills that would have been difficult to achieve otherwise.

I'm very glad to have achieved results that, as the title suggests, help explain ammonia's behaviour at subzero temperatures and provide a starting point for validating kinetic models. More importantly, I'm proud to have designed the entire combustion chamber on my own, which will remain in the laboratory and will be used in the years to come, so knowing that I was the one to design and build it thrills me.

Alberto Francesco Aversa
Delft, January 2026

Acknowledgements

First of all, I want to sincerely thank my supervisors. Thank you, Prof. Mack, for giving me the opportunity to embark on this incredible overseas experience and for giving me full ownership of this project. I'm extremely grateful for this; it has taught me a lot. Of equal importance and sentiment are my thanks to Prof. De Domenico for accepting to supervise me remotely and for providing continuous, invaluable support and feedback, especially in the most challenging parts of the work.

I also wanted to thank Prof. Pini and Prof. Grewe for taking the time to read this thesis and being on my defence committee.

Volevo ringraziare la mia famiglia per credere sempre in me, per sostenermi in questa scelta e per essere sempre presenti.
Thank you to all my friends, from the new ones I made in the US to the ones I found again there and to the ones back in Europe. Your presence, even at a distance, has helped a lot and given me joy in the moments when I needed it most.

Sobre todo, quería agradecerte a ti, Cris, por haber compartido esta experiencia conmigo y por haber sido mi roca todo el rato. Gracias por tu paciencia y por darme calma y fuerzas cuando no las tenía. Esto no hubiera sido posible sin ti.

Thank you, everyone. It has been the wildest experience of my life.

Summary

The sustainability issue has sparked growing interest in investigating low-carbon alternatives to conventional hydrocarbon fuels, particularly for transport and power applications where combustion is still expected to play a role. Within this context, ammonia has emerged as a promising carbon-free energy carrier because it can be produced from renewable electricity and offers practical storage and distribution advantages. However, ammonia also poses a fundamental combustion challenge: ignition is difficult, and early flame propagation is slow, both of which can become critical during cold ambient operation and at high pressure. Experimental data in this combined regime are largely absent, making it difficult to assess feasibility and define mitigation strategies. This thesis was motivated by this gap and aimed to provide a validated experimental basis for understanding ignition and early flame propagation of ammonia and ammonia-hydrogen blends under subzero conditions.

To address this objective, a custom constant-volume chamber was designed, commissioned, and validated to enable controlled subzero operation while maintaining measurement fidelity during ignition transients. The methodology relies on high-quality pressure traces to infer laminar burning velocity during early spherical flame growth, and the setup was verified against established trends before being applied to the subzero ammonia campaign. The main result is that reducing the initial temperature from room temperature down to -50° C imposes a strong penalty on neat ammonia early performance. Laminar burning velocity decreases systematically across the investigated equivalence ratios, pushing ammonia into a very low-velocity regime where the pressure rise becomes weaker and more sensitive to losses and stretch, especially for lean mixtures. This implies a narrower operability window and a higher likelihood of intermittent ignition or early quenching at subzero temperatures, both of which are directly relevant to cold-start reliability.

Hydrogen blending was assessed at the coldest condition and produced a large improvement in combustion performance even at modest fractions. Adding 20% hydrogen increased the burning velocity markedly and brought the mixture close to the neat-ammonia performance at room temperature at stoichiometric conditions, while 30% hydrogen provided a further strong increase, moving the flame into a distinctly faster propagation regime. This shows that the low-temperature penalty is not an unavoidable limitation of ammonia-rich mixtures and can be offset by relatively small additions of a more reactive fuel, with hydrogen being especially attractive because it is carbon-free and counteracts ammonia's main weaknesses in the subzero regime.

Minimum ignition energy was investigated as a second key metric, but a reliable quantitative determination was not achievable with the present setup due to inconsistent discharge waveforms and sensitivity to noise. A qualitative trend was still evident from practical ignition threshold settings, showing that ignition becomes substantially more demanding as temperature decreases and that the coldest conditions required changes in ignition configuration to maintain reliable ignition. Overall, the work delivers a validated subzero combustion platform which shows that subzero temperatures penalise neat ammonia ignition and early flame propagation more strongly than predicted by kinetic models, while hydrogen blending enables performance recovery and improves robustness in the conditions that matter most for cold-start operation.

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations

Abbreviation	Definition
CAD	Computer-aided design
CCS	Carbon capture and storage
CDI	Capacitive discharge ignition
CPM	Constant-pressure method
CVCC	Constant-volume combustion chamber
CVM	Constant-volume method
DAQ	Data acquisition
ECRL	Energy and Combustion Research Laboratory
FEA	Finite element analysis
GHG	Greenhouse gas
ID	Inner diameter
IE	Ignition energy
LBV	Laminar burning velocity
LHV	Lower heating value
LN2	Liquid nitrogen
MAWP	Maximum allowable working pressure
MDMT	Minimum design metal temperature
MIE	Minimum ignition energy
MHSIT	Minimum hot surface ignition temperature
NIST	National Institute of Standards and Technology
OD	Outer diameter
RTD	Resistance temperature detector
STP	Standard temperature and pressure
TU Delft	Delft University of Technology
UML	University of Massachusetts Lowell

Symbols

Symbol	Definition	Unit
A	Pre-exponential factor (Arrhenius)	[–]
A_S	Bolt stress area	[mm^2]
$c_{p,u}$	Specific heat at constant pressure (unburnt mixture)	[J/(kg·K)]
cl_{washer}	Washer clearance	[mm]
cl_{wrench}	Wrench clearance	[mm]
d	Critical kernel/characteristic diameter	[m]
D	Mass diffusivity	[-]
D_b	Bolt diameter	[mm]

Symbol	Definition	Unit
D_{bc}	Bolt centreline diameter	[mm]
D_{fl}	Flange outer diameter	[mm]
$D_{o,g}$	Gasket outer diameter	[mm]
E_a	Activation energy (Arrhenius)	[J/mol]
E_{bd}	Breakdown electric field	[V/m]
F_i	Tensile force in each bolt	[N]
k	Reaction rate constant	[–]
K	Stretch coefficient	[-]
L_M	Markstein length	[m]
N	Number of moles	[mol]
N_{bolts}	Number of bolts	[–]
p	Pressure	[Pa]
p_0	Initial pressure	[bar]
$p(t)$	Chamber pressure trace	[Pa]
R	Universal gas constant	[J/(mol·K)]
$R_f(t)$	Flame radius history	[m]
S_u^0	Unstretched laminar burning velocity	[m/s]
S_u^{SD}	Stretched laminar burning velocity	[m/s]
t	Time	[s]
T	Temperature	[K]
T_0	Initial temperature	[K]
T_b	Burned gas temperature	[K]
T_u	Unburned gas temperature	[K]
V	Velocity	[m/s]
V_{bd}	Breakdown voltage	[V]
x_b	Burned mass fraction	[–]
Y	Mole fraction	[–]
α	Thermal diffusivity	[m^2/s]
δ_L	Diffusive laminar flame thickness	[m]
$\delta_{B,L}$	Blint laminar flame thickness	[m]
γ	Ratio of specific heats	[–]
κ	Thermal conductivity	[W/(m·K)]
μ	Dynamic viscosity	[Pa·s]
π	Pi constant	[–]
ρ	Density	[kg/m^3]
ρ_u	Unburned mixture density	[kg/m^3]
σ_Y	Material yield strength	[MPa]
$\sigma_{Y,B}$	Bolt yield strength	[MPa]
ϕ	Equivalence ratio	[–]

1

Introduction

1.1. Context

Since the Industrial Revolution, global energy use has steadily increased, with an exponential rise in the last century, driven primarily by two factors: population growth and technological advancements. These two factors are the primary drivers of higher energy demand for production, transportation, and domestic life. To put this into perspective, as of 2024, global energy use amounts to approximately 180 000 TWh, which is more than six times the 28 000 TWh needed in 1950, as shown in Figure 1.1. Furthermore, although the growth rate has decreased slightly, demand continues to increase at approximately 2% per year [2], indicating that without adequate control measures, energy demand will continue to grow, eventually exceeding available supply. To determine where meaningful interventions can be made, it is therefore essential to identify which sectors contribute most strongly to this increase.

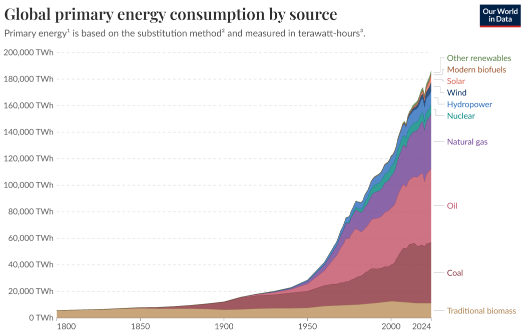

Figure 1.1: Global energy consumption over the years [2]

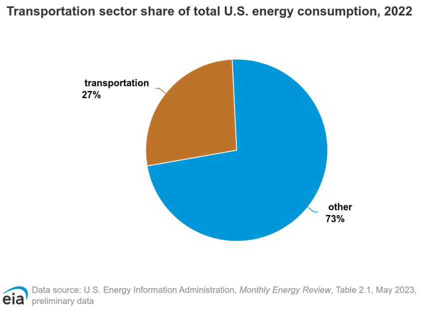

Figure 1.2: Energy consumption in the transport industry [2]

Among these, the transport sector is particularly significant, accounting for approximately 27% of total energy use globally, the largest share. In particular, road transport accounts for the bulk of this energy demand, at roughly 74%, followed by the aviation sector at 10% [3]. Maritime shipping, rail, and other transport modes constitute the remainder, but together they still represent a significant fraction of global consumption. The reason for the preeminence of this sector in energy use can be primarily explained by its needs, that is, extended ranges, high payloads, and operational flexibility, which require large amounts of energy. Despite rapidly increasing interest in electrification and alternative propulsion systems, most of these needs are still satisfied by the direct combustion of liquid fuels, due to their wide availability, high energy density, and ease of combustion. For this reason, this energy production system currently provides more than 90% of the global energy demand [4].

However, around 79% of the employed fuels are of fossil origin, including, among others, natural gas and oils, such as gasoline and kerosene, which produce harmful emissions caused by the presence of carbons in the molecular structure, which leads to the creation of CO and CO_2, among others. The latter is particularly significant, as it is one of the main contributors to global warming and climate change. Hence, fossil fuel combustion has caused a corresponding rapid growth in CO_2 emissions and accounts for more than 80% of the global artificial greenhouse gas (GHG) emissions [5]. For these reasons, in recent years, the transport sector has seen a significant increase in interest in sustainability and the need to transition away from carbon-based fuels. Recent studies indicate growing concern about climate change, whose consequences span a wide range of fields, calling for a broader investigation of strategies to reduce these emissions.

To achieve meaningful decarbonisation, enhancing combustion efficiency alone is insufficient, as it does not eliminate carbon dioxide emissions; therefore, the current most sought-after strategy is the use of alternative, low-carbon, or carbon-free fuels. Among these candidates, ammonia has attracted increasing attention as a potential carbon-free fuel due to a combination of advantages that make it highly compelling. Indeed, it is readily and widely available globally, with a considerable rise in recent years, as shown in Figure 1.3, and is used on a large scale as a fertiliser and cleaning agent, which ensures a well-established global supply chain. In addition, ammonia can be stored and transported under comparatively mild conditions, either as a refrigerated liquid at atmospheric pressure (-33° C) or as a pressurised liquid at room temperature (roughly 0.8 bar to 1 bar), which makes logistics less demanding than for cryogenic hydrogen. Finally, ammonia can be synthesised using renewable electricity via water electrolysis and nitrogen separation, enabling a near-zero life-cycle pathway when coupled to low-carbon hydrogen production, and its high volumetric hydrogen density makes it attractive both as a fuel and as a hydrogen carrier.

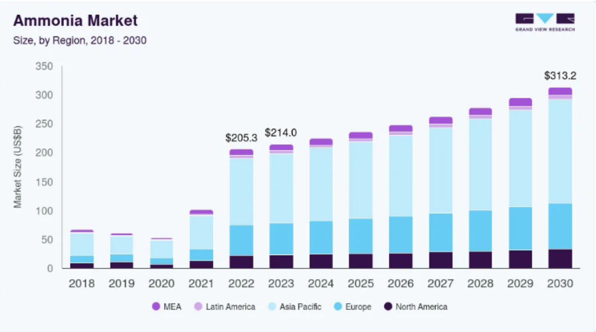

Figure 1.3: Ammonia market growth [6]

From a combustion perspective, as mentioned, ammonia produces no carbon emissions during oxidation, and its high octane rating of approximately 110 RON can be advantageous for high-compression engine operation. However, these benefits are counterbalanced by significant combustion challenges that complicate its combustion, including a very low flame speed relative to conventional hydrocarbons, high ignition energy, and narrow flammability limits. In parallel, the presence of fuel-bound nitrogen contributes to NO_x production because nitrogen is contributed not only by the oxidiser but also by the fuel itself. These limitations do not necessarily preclude ammonia from practical use, but they do imply that system-level measures are often required to ensure robust ignition and stable flame propagation. A common mitigation strategy is to blend ammonia with hydrogen, either through direct co-fuelling or by partially cracking ammonia to generate hydrogen in situ, leveraging ammonia's high volumetric hydrogen density. Consequently, while ammonia is not a drop-in replacement for hydrocarbons, its carbon-free character, scalable production routes, and compatibility with existing liquid-fuel infrastructure position it as a strong candidate for applications where high energy density and transportability are critical.

1.2. Research interest

Although research into ammonia combustion began in the 1950s [7], the recent surge in sustainability interest has sparked renewed interest over the last couple of decades. In general, studies have already characterised its ignition characteristics and pollutant formation (mainly NO_x and unburned NH_3). However, most work has been limited to atmospheric or slightly higher pressures, using constant-volume chambers, shock tubes, or rapid compression machines. As a result, existing datasets often emphasise either low-pressure flame behaviour or high-temperature chemical-kinetics regimes, such as those relevant to gas turbines, rather than high-pressure or low-temperature conditions found in some practical systems. This creates a persistent gap in experimental coverage under these conditions.

High-pressure experiments are particularly important because they determine whether ammonia can deliver acceptable ignition and propagation behaviour under engine-relevant conditions. To date, most experiments have focused on near-atmospheric pressures or high-temperature regimes relevant to gas turbines and fundamental kinetics. Only a limited number of works report ignition or flame propagation at pressures above 2 bar to 10 bar [8, 9, 10], and almost none explore the high-pressure conditions that are directly relevant to modern internal combustion engines. This scarcity limits the ability to tune and validate kinetic mechanisms and turbulent-combustion submodels under realistic constraints. It also complicates assessments of stability limits and required ignition-system capability as pressure increases. Therefore, expanding the high-pressure experimental database is a prerequisite for improving predictive confidence and accelerating technology readiness for ammonia-based combustion systems.

Low-temperature behaviour is equally important because it governs performance during non-ideal operating phases such as cold starts, where slower chemistry and altered transport properties can substantially shift ignition margins and engine requirements. In the specific case of ammonia, low-temperature operation is not just a remote case but can arise naturally from storage and delivery constraints. Ammonia can be stored as a liquid at atmospheric pressure only below -33° C, and even when stored pressurised at ambient temperature, the fuel can undergo cooling during expansion and along unheated feed lines. Consequently, unless the injection and supply path are actively heated and insulated, ammonia may enter the combustion system at subzero temperatures. This makes subzero characterisation essential to anticipate changes in minimum ignition energy, flame speed, and overall stability limits, and to inform design decisions such as injector thermal management and ignition-system sizing. In combination with high pressure, these effects define a challenging but practically relevant regime that remains underexplored.

Addressing this gap is crucial, as these conditions strongly influence ignition delay, flame stability, and pollutant pathways and will ultimately determine the feasibility of ammonia as a fuel in alternative to carbon fuels [11, 12, 13]. Therefore, the objective of this research is to characterise and quantify the initial combustion performance, specifically focusing on two of the most defining parameters for this phase, minimum ignition energy and flame propagation speed, under these two operating conditions. To achieve this, it is important to answer the following research questions:

1. How should a combustion chamber be designed to enable reliable ignition measurements at high pressure and subzero initial temperatures, while ensuring measurement fidelity, thermal control, and safe integration of instrumentation?
2. How do subzero temperatures affect the minimum ignition energy and laminar burning velocity of ammonia?
3. How much does blending ammonia with hydrogen at subzero temperatures affect

combustion performance?

Filling the gap in understanding ammonia combustion under high-pressure and low-temperature conditions can have fundamental implications. In fact, from a fundamental perspective, it could provide the data needed to validate and refine chemical kinetic models, which currently extrapolate beyond their validation range and thus provide only crude estimates of performance under these conditions. Thus, improving the models enhances predictive capability for ignition description, flame stability, and pollutant formation. Moreover, from a design perspective, the results would be critical for understanding the operation of internal combustion engines and gas turbines fuelled by ammonia. In fact, knowledge of this regime would help develop ignition systems and combustion strategies that minimise misfire and instability while also reducing NO_x emissions. Ultimately, eliminating this gap could accelerate the adoption of ammonia as a carbon-free fuel in sectors that are difficult to electrify, such as heavy-duty transportation and shipping.

1.3. Thesis structure

Before beginning the study, a thorough analysis of the current state of research on the topic is conducted to precisely pinpoint the research gap and provide data that attempts to address it. This information is summarised in Chapter 2. Following this, the research objective and related questions are reformulated and presented in Chapter 3, along with a research plan to address them and the proposed methodology to achieve them. This entails first creating an experimental setup that meets the strict temperature and pressure requirements. This includes both the design of an ad hoc combustion chamber and the setup for pressurisation and temperature conditioning, explained in Chapter 4. Once this is achieved and the set is created, the necessary experiments are conducted, starting with the verification and validation of the chamber design and progressing toward achieving the desired operating conditions. The experimental methodology, including the methods used and the data extrapolation, is described in Chapter 5, where the methods and processes employed are detailed, allowing the research to be replicated. Then, the results are presented in Chapter 6, where they are analysed against simulated values to compare trends and gain an understanding of ammonia combustion behaviour in the mentioned gap. Finally, Chapter 7 summarises the research results and answers the research questions, while Chapter 8 highlights key points of interest and potential areas for improvement in future studies.

2

Literature review

This chapter provides the theoretical background and research context for the present work by summarising the current state of combustion research and identifying the specific knowledge gap addressed in this thesis. It begins by outlining the continued significance of combustion as an energy-conversion pathway, both in current systems and in the context of the ongoing transition toward sustainable energy services. It then introduces two flame characteristics, particularly significant for the work carried out in this thesis, which will lead to a description of the main fuels available and considered. The key parameters used in this thesis to quantify early-stage combustion performance explain why these metrics are useful when comparing fuels and assessing practical operability. The discussion emphasises that these parameters are not intrinsic constants but depend strongly on environmental variables, such as initial pressure and temperature, which can shift ignition susceptibility and flame development, thereby constraining combustor design.

Building on this foundation, the review narrows from general combustion behaviour to the fuel-specific motivations for ammonia, summarising its advantages as an energy carrier and the principal combustion challenges that currently limit deployment. It then positions ammonia within the broader landscape of alternative fuels by focusing on the aspects that most directly affect ignition and laminar flame propagation. Particular attention is given to how the target operating envelope in this work, defined by elevated pressure and subzero initial temperature, influences the interpretation of the selected combustion metrics. Finally, the chapter consolidates the literature evidence to show that data remain limited for ammonia and ammonia–hydrogen blends under these combined conditions, which motivates the experimental focus of this thesis. This gap defines the context for the methodology and results presented in the subsequent chapters.

2.1. Introduction to combustion processes

Combustion is a rapid, exothermic process that converts the chemical potential energy in molecular bonds into usable thermal energy by oxidising a highly energetic fuel. This can then be translated into mechanical work, direct thrust, or electricity generation. It is one of the most widely used methods for energy conversion due to its high energy density, versatility, and straightforward implementation. For combustion to happen, three elements are necessary:

- Fuel, which is the element that donates electrons during oxidation. It exists in different phases, from solid to liquid to gaseous phases, but this thesis focuses

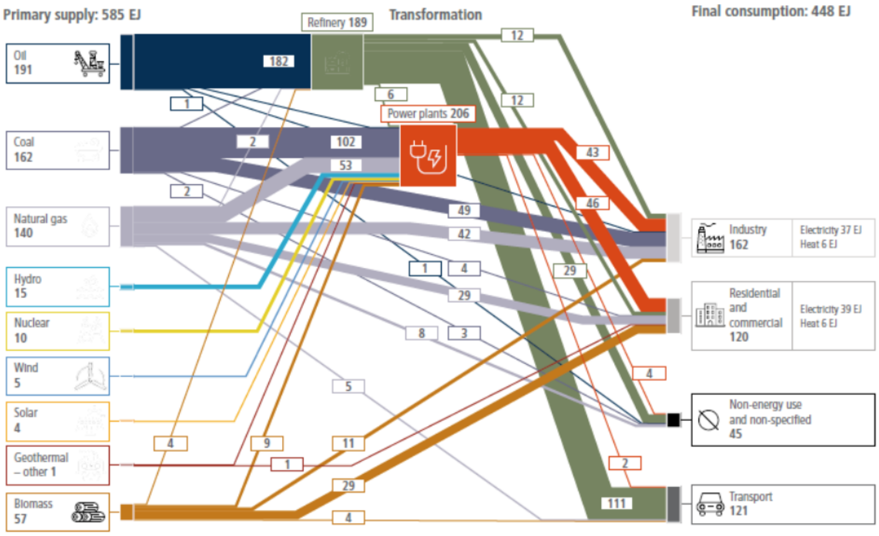

Figure 2.1: Global energy flow [14]

only on gaseous-injected fuels, as they simplify analysis and are commonly used in high-pressure combustion studies.

- Oxidiser, which is the element that accepts the electrons in the reaction and is typically air or oxygen.
- Heat or energy source, which is necessary to overcome the activation energy necessary to break the initial bonds. If this is enough, the reaction becomes self-sustaining due to the heat released by the combustion process.

Combustion is a fundamental process in modern energy systems, serving as a primary means of converting the chemical energy in fuels into various usable forms, such as heat, mechanical power, and electricity. A significant portion of the global energy supply depends on combustion, in which fuels are oxidised in engines, turbines, boilers, and industrial furnaces. The relationship between combustion processes and energy services is clearly illustrated in the global energy flow structure, shown in Figure 2.1. In this framework, primary energy sources are converted into energy carriers, including electricity and refined fuels, which are then distributed to end-use sectors such as industry, buildings, and transportation. Consequently, combustion plays a vital role in both centralised electricity and heat production and decentralised applications.

From a climate and environmental perspective, combustion is of paramount importance, as the global energy system is the most significant contributor to CO_2 emissions. Energy-related greenhouse gas emissions largely stem from the combustion of fossil fuels, primarily for heat supply, electricity generation, and transportation. This reliance on fossil fuels highlights the urgent need to decarbonise energy conversion methods while ensuring reliable energy services continue to meet our demands. As shown in Figure 2.2, global emissions have consistently increased over time, underscoring a pressing problem that cannot be ignored. Addressing these emissions is critical to achieving climate goals and promoting a

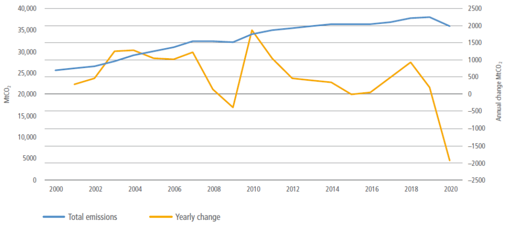

Figure 2.2: Global CO_2 emissions [14]

sustainable future.

In the context of the ongoing energy transition, future energy systems must adapt to meet stringent climate targets. This will necessitate substantial changes in how energy is produced, converted, and utilised across different sectors. Such system-level adjustments emphasise the continued relevance of combustion research, even as electrification becomes more widespread. Many hard-to-electrify applications, such as high-temperature process heat, long-duration storage, and specific segments of transportation, might still depend on low-carbon fuels and advanced conversion technologies. These technologies, which include combustion processes that minimise or manage emissions, will be crucial in bridging the gap toward a more sustainable energy landscape. In such systems, operability is often set by the earliest stage of combustion, namely, whether ignition is reliable and whether the initial flame can develop without quenching.

2.2. Fundamental flame behaviour

These combustion processes are divided into different phases. One of the fundamental ones is the ignition and early propagation phase of premixed combustion, governed not only by chemical kinetics, but also by transport and geometry-driven effects that can strongly bias measured flame speeds. This is particularly true in constant-volume experiments, where the flame develops from a small kernel and evolves through a strongly curved, stretched regime before approaching a more weakly stretched, quasi-planar behaviour. Since this thesis relies on early flame development metrics such as minimum ignition energy and laminar burning velocity, it is essential to establish the fundamental concepts that govern their interpretation in practice. In particular, curvature and differential diffusion can either suppress or enhance local burning, depending on the fuel, making early-time measurements not directly comparable unless stretch effects are accounted for.

2.2.1. Flame stretch

In the early stages of flame propagation, particularly for spherically expanding flames, the flame front is highly curved and subject to significant stretch. This parameter refers to the rate of change of the flame surface area, defined as:

$$K = \frac{1}{A}\frac{dA}{dt} \tag{2.1}$$

For a spherical flame, $A = 4\pi r^2$, and the time derivative is $\frac{dA}{dt} = 8\pi r \frac{dr}{dt}$, so the stretch scales down with the radius. Therefore, when the flame radius exceeds roughly ten times the flame thickness, the curvature becomes negligible, and the flame front behaves as locally planar.

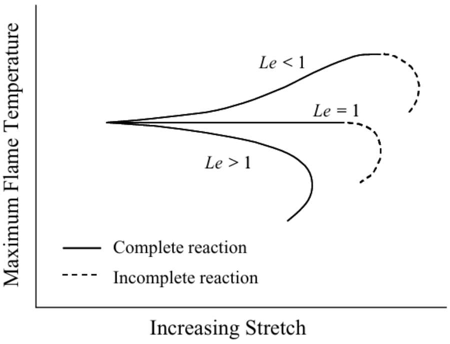

Figure 2.3: Relation between flame stretch and Lewis number [15]

Flame stretch is the sum of two main contributions:

- Strain is the deflection of the flame due to velocity gradients in the flow field
    - Normal strain is due to the compression or stretch of the flame perpendicular to itself.
    - Shear is the shearing effect of the flame, tangential to the profile.

- Curvature is the deviation from a planar profile
  - Positive when the curvature is towards the unburnt mixture, hence decreasing the flame speed
  - Negative when the curvature is towards the burnt mixture, hence increasing the flame speed

The sensitivity of the flame speed to stretch is captured by the Markstein relation:

$$S_L = S_L^0 (1 - \mathcal{L}_M K) \tag{2.2}$$

Here, $\mathcal{L}_M$ is the Markstein length, which describes how the stretch affects the flame. A positive length indicates that the stretch suppresses flame propagation, whereas a negative value suggests that it enhances it.

An important parameter affecting the stretch, at least on a microscopic level, is the Lewis number:

$$\mathrm{Le} = \frac{\alpha}{D} = \frac{\lambda/(\rho c_p)}{D} \tag{2.3}$$

This parameter compares thermal diffusivity α and mass diffusivity D, which determines the flame's response to stretch:

- If Le > 1, thermal diffusion dominates, and stretch reduces the flame temperature and speed due to heat loss.
- If Le < 1, stretch increases radical concentration, boosting reaction rates and flame speed.

The Lewis number primarily affects developing flames, but it also affects the fully developed flame; thus, it affects flame thickness and speed, though the microscopic effect has a lower impact. This is because differential diffusion can only change the equivalence ratio locally, but it needs a scalar gradient source (stretch) to maintain that imbalance. Inside a planar, steady flame, the gradients quickly settle to self-similar profiles that result from the balance between convection and diffusion. If there is no disturbance, diffusion acts to restore equilibrium and smooth any existing gradients. Without the gradient, the temperature and species profiles quickly relax to the one-dimensional steady solution, which is why the *Le* has no noticeable effect at the macroscopic level [15].

2.2.2. Spherical flames

To connect measurements to physically meaningful combustion quantities, the literature often adopts canonical flame configurations with well-defined geometry and boundary conditions. A spherically expanding premixed flame provides such a configuration because it offers a simple, analytically tractable geometry while still capturing the key physics governing early flame development. In practice, spherical expansion is realised when combustion starts from a localised region within an initially quiescent, premixed mixture, so that the flame front grows outward with near-isotropic curvature over an initial interval. During this phase, the flame area increases rapidly, and stretch effects are inherent, making the configuration particularly useful for relating early propagation to transport and preferential diffusion. This canonical geometry therefore provides a clear bridge between ignition-scale behaviour and the larger-scale flame-propagation metrics used for comparisons across fuels and conditions.

Apart from its occurrence in engines, this configuration is widely used in experimental and numerical studies to characterise fundamental flame properties under controlled conditions. In optical studies, the spherical geometry can be tracked directly to obtain stretched and unstretched flame speeds, whereas in non-optical studies, the same phase can be interpreted through the associated pressure rise and the inferred rate of heat release. Consequently, spherical flames provide a common reference point between optical and non-optical methodologies and form the conceptual bridge to the constant-volume chamber approach used in this work.

From a research perspective, these flames offer several advantages:

- They provide well-defined boundary conditions, simplifying the modelling of heat conduction, mass diffusion, and chemical kinetics.
- They closely resemble real-world ignition events, which typically begin from a point-like source.
- Their evolution enables accurate determination of parameters such as the laminar burning velocity (or flame speed) and Markstein length.

In spherical geometry, the flame expands outward due to both thermal expansion and chemical reaction, and the point-like ignition introduces significant flame stretch during the early phase. This stretch influences local flame speed through curvature effects, which becomes particularly relevant when interpreting measurements during the kernel-growth period rather than in a fully developed flame. As a result, the measured propagation in the earliest stage is not only controlled by chemistry, but also by transport and by the stability of the curved flame front. These effects become more pronounced when reactants are at subzero initial temperatures and at elevated pressures, because ignition behaviour and flame stability are more sensitive to heat losses, diffusivity, and radical-production rates.

2.3. Fuels

The choice of fuel in combustion systems is crucial for both performance and environmental impact. Fossil fuels, known for their high energy density and established infrastructure, have been the backbone of energy production. However, their significant emissions underscore the importance of exploring alternative fuels and adopting technologies such as carbon capture. Different fuels exhibit different combustion characteristics, which influence system design and operation. Factors such as ignitability, flame speed, and pollutant formation dictate how combustion systems are constructed and optimised. Understanding these properties is key to developing cleaner, more efficient combustion technologies.

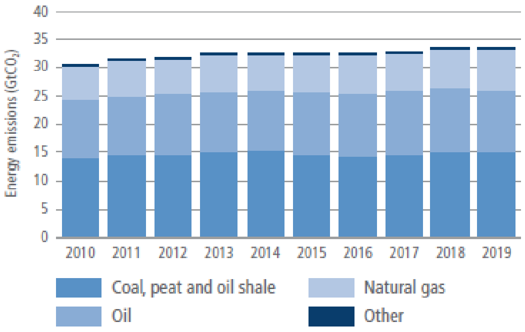

Figure 2.4: Combustion CO_2 emissions per fuel [14]

The energy sector's emissions profile varies significantly depending on the fuel source used. Coal, oil, and natural gas all contribute to greenhouse gas emissions in varying degrees, with coal being the most detrimental per unit of energy, as shown in Figure 2.4. This highlights the need for cleaner alternatives. As such, there is growing interest in alternative fuels that aim to reduce in-use impacts. However, it's essential to assess the life-cycle emissions of these alternative fuels to fully understand their overall environmental implications. Within this landscape, ammonia has gained attention as a potential carbon-free energy carrier and hydrogen vector, with prospective roles in shipping, power generation, and industrial energy supply. Its relevance to combustion research follows from a practical systems question: if ammonia is to be used in combustion devices, either pure or blended, then its ignition behaviour, flame propagation characteristics, and operability constraints must be quantified under conditions representative of real systems [14].

In this section, an overview of the fuels considered in this thesis will be provided. Methane is included as a well-characterised hydrocarbon reference fuel, providing a reliable baseline for verification and for benchmarking the facility against established literature trends. Hydrogen is considered because it is a carbon-free fuel with extensive validation data and combustion behaviour that contrasts strongly with ammonia, making it a useful comparison case for assessing method robustness across very different transport and kinetic regimes. It is also the most relevant blending partner for ammonia within a low-carbon framework, as it can substantially improve ignition and early flame propagation while maintaining a sustainable overall fuel strategy. Finally, the characteristics of ammonia are explained to understand the properties that govern its combustion behaviour.

2.3.1. Methane

Carbon-based fuels include a broad category of components that are the most common means of energy production in energy systems and exist in the three forms of matter. Among them, methane, due to its high hydrogen-to-carbon ratio, is considered the "cleanest" carbon-based fuel. Methane is widely used as a practical fuel in gas turbines and gas engines, and, for that reason, it has become a standard reference fuel in combustion research and model development. Its combustion chemistry is comparatively well understood, making it useful as a baseline for validating experimental setups or comparing fuels under controlled conditions. However, since the objective of the thesis is to investigate sustainable alternatives, after validation, the focus shifts towards cleaner fuels.

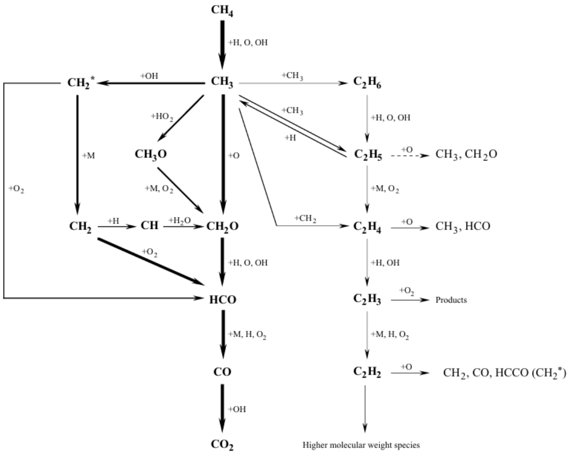

Figure 2.5: Methane reaction pathways [15]

Following the reaction pathway shown in Figure 2.5, methane oxidation in kinetic modelling is typically represented through elementary-reaction mechanisms that describe the breakdown of CH_4 into smaller radicals and intermediates, followed by rapid oxidation to the stable products CO_2 and H_2O. Compared with heavier hydrocarbons, methane generally exhibits lower soot tendency and reduced CO_2 per unit energy released, while maintaining robust premixed flame propagation in practical devices. However, despite being a well-characterised fuel, its ignition and flame behaviour can still vary substantially with operating conditions, underscoring the need to consider pressure and temperature effects when using methane data or models. Methane behaviour is also strongly dependent on the equivalence ratio, which shifts both reactivity and pollutant formation pathways [16].

2.3.2. Hydrogen

Hydrogen is an alternative energy carrier that could replace fossil fuels and is increasingly used in transport systems, both in combustion and in fuel cell engines. From a chemical standpoint, hydrogen oxidation does not form CO_2, and the primary stable product of complete combustion is water, as shown in Figure 2.6, which is one of the main reasons it is often framed as a clean fuel in the context of decarbonisation. However, this benefit is only applicable when it is produced sustainably, which is currently only achieved on a small scale due to high costs and energy requirements. The most commonly used techniques for producing hydrogen are:

- Steam methane reforming: entails the reaction of methane with steam to produce hydrogen and carbon dioxide. Since it is the cheapest, it accounts for 95% of hydrogen production.
- Electrolysis: it splits water into hydrogen and oxygen using electricity. When electricity is generated from sustainable sources, it is called green hydrogen. This process accounts for only 4% of global hydrogen production due to high electricity costs.
- Biomass gasification: in this case, organic materials are converted into hydrogen-rich syngas through high-temperature processes. This process represents another renewable alternative, but it is less developed commercially due to the considerable costs.

A defining property of hydrogen is its very high gravimetric energy density, which makes it attractive whenever mass is a dominant constraint and supports its use as an energy carrier across multiple applications. At the same time, hydrogen's low density at ambient conditions imposes a significant practical penalty on volumetric storage, meaning that equivalent energy content typically requires much larger storage volumes compared with conventional fuels. This constraint is frequently addressed through high-pressure storage or cryogenic liquid storage, but both options introduce additional complexity, cost, and design constraints for real systems. For this reason, hydrogen adoption tends to be driven not only by combustion performance but also by how realistically the fuel can be stored and delivered in the target application, particularly when space and integration constraints dominate.

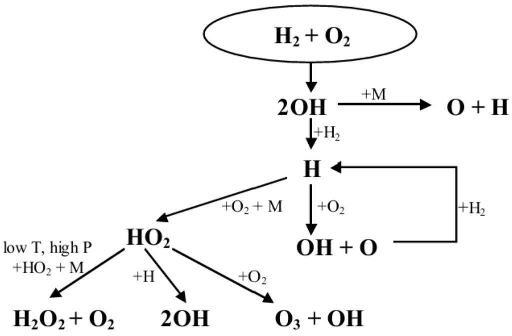

Figure 2.6: H_2 reaction pathway [17]

From a combustion perspective, hydrogen has already been explored in several established conversion platforms, including internal combustion engines and gas turbines, where it is typically used as either a neat fuel or a blending component. Blending strategies are especially relevant in transitional pathways because hydrogen can be used to improve ignition and flame development of slower-burning fuels, thereby extending operability, with ammonia commonly discussed as a carbon-free partner fuel in this context. However, hydrogen combustion is not automatically "emissions-free", since high-temperature oxidation can still lead to NO_x formation, making its mitigation a central constraint in combustor design and operation even when CO_2 is eliminated. The practical picture is therefore a trade-off: hydrogen offers strong decarbonisation potential and favourable mass-based energy content, but it also introduces storage and integration challenges and requires careful control of combustion conditions to avoid shifting the emissions burden from carbon species to nitrogen oxides.

To summarise, hydrogen presents some considerable advantages but also some drawbacks that make its large-scale use challenging:

- Advantages
  - Environmental benefits: it only emits water vapour, thus does not contribute to harmful CO_2 emissions.
  - High energy density: approximately 120 MJ/kg, which is considerably higher than any carbon fuel's.
  - Wide flammability range: hydrogen can combust with a percentage in air anywhere in between 4 and 75%.
  - High diffusivity: due to its small mass and size, it can diffuse much faster than any other molecule, which positively contributes to mixing.
  - Low MIE: the two above factors contribute to hydrogen having a very low MIE, of about 0.02 mJ.
- Disadvantages
  - Production emissions: as explained, most hydrogen is produced by emitting carbon dioxide.
  - NO_x emissions: due to its higher combustion temperatures, thermal NO_x is more concerning than in carbon fuels.
  - Storage and transport challenges: due to its very low freezing point, it must be stored and transported at cryogenic temperatures.
  - Safety concerns: its diffusivity, MIE, and flammability range are safety concerns that need to be addressed, especially in the infrastructure [18, 19].

2.3.3. Ammonia

NH_3 is another sustainable energy carrier under consideration. In fact, it is a very widespread and interesting hydrogen carrier, with a hydrogen density of 17.8% per unit weight. The most common NH_3 production process, known as the Haber-Bosch process, converts catalytically N_2 (from air) and H_2 (from industrial processes) to ammonia using a metal catalyst under high pressure (150 atm to 250 atm) and high temperature. The most common way to achieve this is via steam methane reforming, which, however, produces CO_2 emissions. However, with the rising issue of sustainability, the synthesis from sustainable energy sources is becoming of more interest, such as

- Wind power
- Nuclear energy
- Fossil fuels via CO_2 capture and storage (CCS)

This means that NH_3 can be fully carbon-free, both in production and in use.

Ammonia is widely used in the agricultural industry as a fertiliser, accounting for about 80% of its global use. Moreover, it is commonly used as a refrigerant and in the production of other industrial chemicals, including rocket fuel hydrazine. Once produced, ammonia can be stored in the liquid state at 298 K at a pressure of only 10.2 atm, which is much lower than that of natural gas, which requires 245 atm. This makes the transport of ammonia much easier, also compared to hydrogen, which needs cryogenic temperatures to remain liquid. Thus, ammonia offers several advantages, but to be used as an energy carrier, certain important drawbacks must be overcome.

- Advantages
    - High volumetric energy density: ammonia's is of about 11.2 MJ/L, which is comparable to methanol (15.6 MJ/L) and 33% higher than that of hydrogen.
    - High heating value at stoichiometric conditions: in these conditions, the heat released per unit of fuel is comparable to both methane and hydrogen.
    - High octane number: it is 110 RON, which is higher than that of natural gas and gasoline, meaning that it can be operated at high pressure without risking detonation.
    - Ease of handling and storage: due to low vapour pressure and high boiling point.
    - Carbon-free production and combustion: as explained before, ammonia does not produce any carbon emissions.
    - Low cost: the cost per unit energy is much lower than gasoline (9.9 USD/GJ vs 21.8 USD/GJ)
- Disadvantages
    - High latent heat of vaporisation: causing a significant temperature drop if injected in the liquid state.
    - Poor combustion properties: ammonia is characterised by low flame speed, high MIE, high autoignition temperature, and narrow flammability limits
    - High NO_x emissions: in addition to the thermal creation, the fuel-NO_x production must also be considered.
    - High corrosivity: ammonia is a base, with a relatively high pH, which can make the environment caustic [11, 20].

Due to its poor combustion properties, which make it difficult to use as a stand-alone fuel, ammonia is mostly used in blends. When added to carbon fuels, the fuel-NO formation process becomes dominant, overcoming the thermal one. However, as the percentage of added NH_3 increases, more of the fuel is converted to N_2 rather than to NO. The addition also leads to a longer ignition delay and lower available power due to a high ignition temperature and a low burning velocity. The pressures and temperatures achieved during the combustion of this blend are also lower, which increases CO and HC emissions while reducing CO_2 and soot. For these reasons, it is more convenient to use H_2 as a combustion promoter since it can be produced on board by dissociating NH_3 itself. As shown in Figure 2.7, the main radical formation processes for ammonia are:

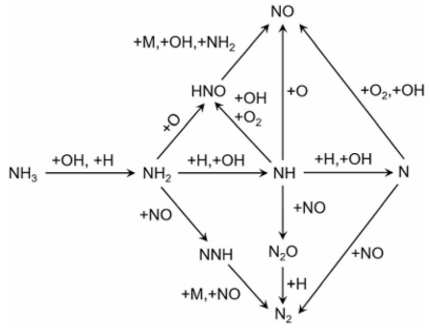

Figure 2.7: NH_3 reaction pathways [12]

- Oxidation leading to NO formation
- Reaction with NO leading to N_2 formation

The identity of the NH_i radicals responsible for the formation and destruction of NO depends on the equivalence ratio: for lean mixtures, NH_2 is the critical radical, while as the mixture becomes richer, this becomes NH first and then N [12]. The production of these radicals and products explains the buoyant behaviour of ammonia flames, which tend to rise up. This is because the density of the burnt gases, mainly made of light substances, is much lower than the unburnt ones.

This is also why, in practice, a large fraction of the ammonia-combustion literature considers NH_3 in combination with hydrogen, since H_2 is an effective combustion promoter that directly compensates for ammonia's low reactivity. From a system perspective, this strategy is particularly attractive because hydrogen can be generated on board by partially cracking ammonia, thereby improving ignition and flame propagation without introducing carbon into the fuel stream. This approach keeps storage and transport closer to conventional fuels than for pure hydrogen [20]. From a combustion standpoint, adding hydrogen expands operability and stabilises the early flame kernel, which is often the limiting stage for ammonia in spark-ignited and premixed configurations. This is also consistent with recent state-of-the-art assessments of spark-ignition engines, where hydrogen enrichment has been repeatedly identified as a practical enabler to avoid misfire and reduce cycle-to-cycle variability when ammonia is used as a primary energy carrier [21].

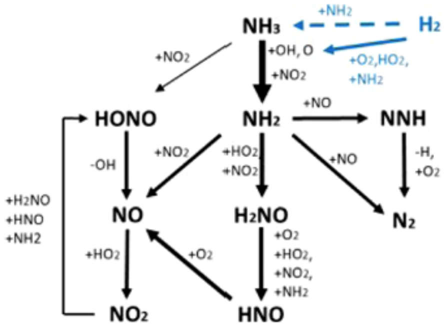

Figure 2.8: NH_3-H_2 mechanism [22]

The impact of hydrogen addition is especially evident when ignition behaviour is analysed under engine-relevant pressures and intermediate temperatures, where pure ammonia remains difficult to ignite and exhibits long ignition delays. He et al. [22] measured ignition delay times of NH_3/O_2 and $NH_3/H_2/O_2$ mixtures at 20–60 bar and 950–1150 K, and demonstrated that increasing the hydrogen mole fraction in the fuel mixture increases reactivity and shortens ignition delay times. Importantly, they also showed that the role of equivalence ratio is not universal but can change with the hydrogen fraction, meaning that trends observed for pure ammonia do not necessarily extrapolate to ammonia-hydrogen blends. A further consequence is that kinetic models that reproduce pure ammonia ignition can still struggle to capture blend behaviour across the same conditions, indicating that hydrogen enrichment introduces additional sensitivities in the radical pool and nitrogen sub-mechanism that are not fully constrained by existing datasets. This thesis reinforces the need to characterise ignition metrics within the relevant pressure-temperature envelope rather than infer them from atmospheric-pressure trends.

2.3.4. Summary

Table 2.1 summarises the differences among the three different fuels just presented, evidencing the main chemical and combustion properties. These are important for understanding the combustion behaviour of these fuels under different conditions.

Table 2.1: Comparison of combustion and transport properties of methane, hydrogen, and ammonia. All values, unless specified are at T = 298 K and p = 1 bar [18, 12, 20, 15, 23]

Property	**Methane (CH_4)**	**Hydrogen (H_2)**	**Ammonia (NH_3)**
Boiling Point [°C @ 1 atm]	–161.5	–252.9	–33.3
Storage Pressure [bar @ 25°C]	~200	350–700	~10
Lower Heating Value [MJ/kg]	50.0	120.0	18.6
Laminar Flame Speed [cm/s]	~38	~265	~7
Autoignition Temperature [°C]	537	585	651
Minimum Ignition Energy [mJ]	~0.28	~0.02	~20
Flammability Limits [vol% in air]	5–15	4–75	15–28
Adiabatic Flame Temperature [K]	~2220	~2400	~1800
Stoichiometric Air-to-Fuel Ratio (mass)	17.2	34.3	6.1
Density at STP [kg/m^3]	0.656	0.0899	0.73
Molecular weight [g/mol]	16.04	2.02	17.03
Diffusivity in Air [cm^2/s]	~0.22	~0.61	~0.20
Dynamic Viscosity [µPa·s]	11.0	8.9	9.8
Thermal Conductivity [W/m·K]	0.034	0.180	0.025
Toxicity	Low	Non-toxic	Toxic and pungent
In-use CO_2 Emissions [g/MJ]	~50	0	0

2.4. Combustion performance parameters

To correctly quantify fuel performance, it is important to analyse the initial phase of combustion, including ignition and early laminar flame propagation. This is because it is the stage that determines whether combustion starts reliably, how rapidly it transitions into a self-sustaining process, and what operating margin a device has before it misfires, blows off, or becomes unstable. In practical combustors, the system is often designed around the most demanding conditions for ignition (cold start, high dilution, high pressure, high flow, transient operation), and failures at this stage can dominate real-world operability even if the fuel burns well once fully established. Therefore, improving reliability and efficiency at ignition affects how effectively fuels can deliver energy.

Two parameters are particularly important to describe the behaviour of a flame at the initial combustion stage analysed in this work: the minimum ignition energy and the laminar burning velocity. The minimum ignition energy (MIE) reflects an ignition source's ability to create a kernel that survives losses and can grow. This growth is governed by heat release rates, radical chemistry, and the balance between chemical timescales and diffusive/convective losses. Early laminar flame propagation sets the growth rate of that kernel into a flame, and the laminar burning velocity (LBV) is a standard metric for comparing fuels and validating chemical kinetic models because it links directly to reaction rates under well-defined conditions. This is especially relevant when extrapolating to conditions where data are scarce, because a model that cannot reproduce ignition thresholds and LBV under controlled conditions is unlikely to predict real-device behaviour reliably.

Moreover, a study of these early-phase metrics directly aligns with design requirements for ignition systems and combustor stability. In fact, if MIE is high or pressure- and temperature-sensitive, the ignition system may require higher delivered energy, smaller electrode gaps, different discharge characteristics, or improved thermal management to achieve reliable light-off. On the other hand, if LBV is low, flame stabilisation becomes harder, residence-time requirements increase, and the combustor becomes more susceptible to quenching, buoyancy effects in quiescent vessels, or heat-loss-dominated extinction near walls. These effects are amplified for low-reactivity fuels such as ammonia and for ammonia–hydrogen blends in cold/high-pressure regimes, where small changes in conditions can shift the mixture from “ignitable” to “non-ignitable”.

Finally, early combustion phases are tightly coupled to safety and emissions, which are often decisive constraints in adopting new fuels. Ignition behaviour influences flammability hazards, the likelihood of unintended ignition, and the energy required to initiate combustion in controlled settings, while laminar flame development affects how long the mixture spends in low-temperature, partially reacted states that can promote undesirable intermediates. For ammonia-based fuels, early-phase behaviour also conditions the downstream formation pathways of NO_x and unburned species, because kernel growth and early flame temperature history influence residence time and local equivalence ratio evolution near the ignition region. The broader energy-systems context reinforces why these details matter: deploying alternative fuels at scale requires not only low-carbon supply pathways but also conversion processes that are reliable, safe, and compatible with practical devices.

2.4.1. Laminar burning velocity (LBV)

The laminar burning velocity is defined as the speed at which a planar, one-dimensional laminar flame propagates through a quiescent unburnt mixture. In this context, "quiescent" refers to an environment in which the air within the combustion chamber remains relatively still and non-turbulent. This stillness enables a controlled distribution of fuel, leading to more predictable combustion characteristics. The focus of this thesis will be on this stable environment, intentionally excluding the complexities introduced by turbulence. Ideally, the goal is to achieve a perfectly planar, stationary, and adiabatic flame with a consistent flame speed.

However, achieving such an ideal flame is complicated by non-uniformities, such as temperature and density gradients, which lead to local stretching effects. Due to these variations, perfectly spherical flames are rarely observed in practice. Therefore, it is crucial to distinguish between the instantaneous flame speed, denoted as S_u, which reflects local conditions, and the general flame speed, S_u^0, representing the average velocity of the entire flame. The instantaneous flame speed can vary based on local stretching and, as a result, may be either higher or lower than the average speed. This distinction helps in better understanding the behaviour of flames under different conditions.

One challenge in understanding flame speed is accurately determining the flame front, which marks the boundary between burnt gases and the unburnt mixture. This defines the flame region and introduces additional complexity regarding flame thickness. Correctly identifying this flame front is essential for analysing flame behaviour and dynamics. Understanding these characteristics provides insight into the overall combustion process, aiding in the design and optimisation of combustion systems. Thus, a comprehensive exploration of laminar burning velocity, flame dynamics, and the influencing factors is pivotal for advancing knowledge in combustion science.

2.4.1.1. Fuel dependency of LBV

The flame speed is highly dependent on the fuel’s chemical structure and behaviour, as these determine its reactivity, diffusivity, and exothermicity in the mixture. These factors directly affect how energy is released and the fuel’s kinetic behaviour. These differences in chemical composition and structure result in significant differences in flame propagation speed. In fact, for ammonia, the maximum LBV is about 6 cm/s to 7 cm/s at an equivalence ratio of 1.1. As a comparison, methane can reach 35 cm/s to 37 cm/s, while hydrogen can reach 290 cm/s. These three behaviours are shown as a function of equivalence ratio in Figure 2.9.

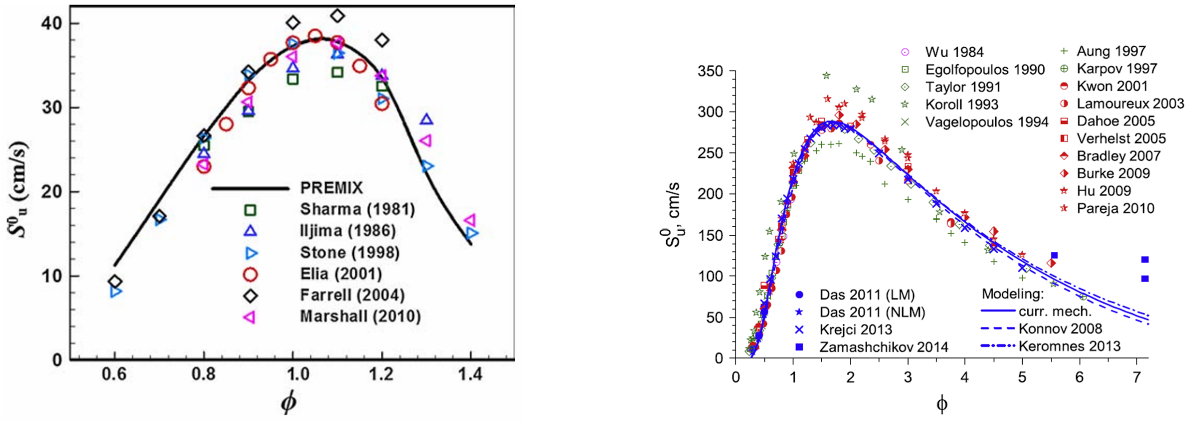

(a) CH_4 [24] **(b)** H_2 [25]

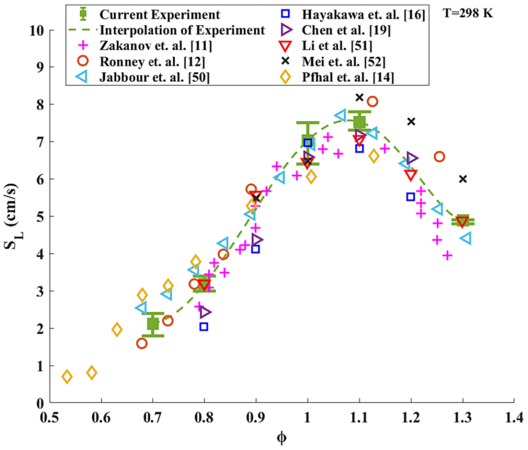

(c) NH_3 [10]

Figure 2.9: LBV of methane, hydrogen and ammonia at STP ($T_0 = 21°$ C, $p_0 = 1$ bar)

Starting with hydrogen, the reasons for its incredibly high velocity lie in the unique behaviour of its molecule. In fact, due to the very small, lightweight nature of its molecules, it diffuses much faster into the unburnt mixture than other fuels, promoting rapid transport of both heat and reactive radicals toward the region immediately ahead of the flame. A consequence of this is the thermo-diffusive instabilities, as hydrogen diffuses faster than the flame does, causing the flame front to become cellular or wrinkled, which increases the total surface area and further boosts the burning rate. In addition, hydrogen's high thermal conductivity improves reactant preheating, allowing the mixture to reach ignition-relevant temperatures more readily and reinforcing hydrogen's high flame reactivity. Finally, hydrogen oxidation mainly follows the pathway $H + O_2 \leftrightarrow O + OH$, producing large amounts of atomic hydrogen, an extremely reactive radical that drives the flame forward [11].

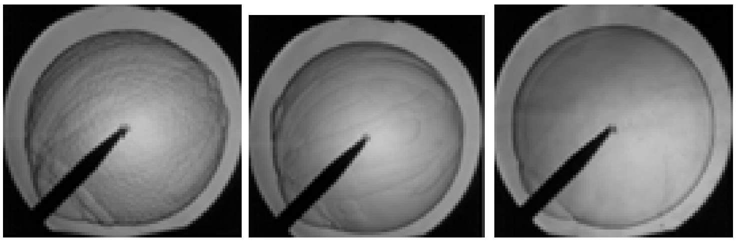

(a) ϕ = 0.8 **(b)** ϕ = 1.0 **(c)** ϕ = 1.5

Figure 2.10: H_2 Schlieren flame propagation at STP (T_0 = 21° C, p_0 = 1 bar) [26]

In the case of methane, its considerably lower speed is due to its complex chemical pathways characteristic of carbon fuels. In fact, methane combustion involves a complex multi-step reaction path. It requires the breaking of four stable C-H bonds and follows a chain-carrying process that is much slower than the rapid, radical-driven chain reactions of hydrogen. Moreover, methane flames produce fewer reactive H, O, and OH radicals in the reaction zone. In contrast, hydrogen enrichment specifically boosts burning velocity by increasing the concentration of these active species. Thus, during the reaction, some endothermic decomposition occurs, lowering the global flame speed [27, 28].

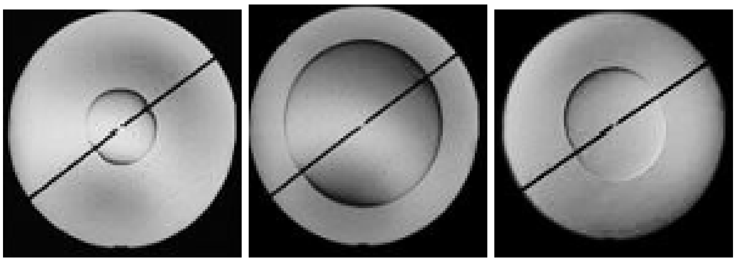

(a) ϕ = 0.8 **(b)** ϕ = 1.0 **(c)** ϕ = 1.3

Figure 2.11: CH_4 Schlieren flame propagation at STP (T_0 = 21° C, p_0 = 1 bar) [29]

The even lower flame speed of ammonia can be attributed to a greater prevalence of endothermic decomposition reactions, as ammonia combustion produces a significantly lower concentration of key reactive radicals, such as H, O, and OH. Moreover, because its chemical structure does not contain any highly energetic bonds, such as CO or CH, the heat release is much lower, as shown by its lower Lower Heating Value (LHV). In fact, ammonia's energy density is less than half that of methane. With less energy released per unit of fuel, there is less heat available to drive the flame front forward into the unburnt mixture. Finally, ammonia shows a difference between local and global flame speeds, as its lower reactivity, due to its lower molecular weight compared to air, also leads to buoyancy, causing the flame to rise. Therefore, a point in the flame expanding upwards will have a higher velocity with respect to one expanding downwards [30, 20].

Its speed can be increased by changing the mixture composition in three main ways, which can increase its value up to 5 times:

- Mixture enrichment with oxygen: this provokes an increase in the reaction rates of OH, H, O, and NH_2 radicals, which causes many more chain-branching reactions.
- Hydrogen addition: it leads to an increase in the peak pressure, as with more H_2, more H, O, and OH radical mole fractions will be released during the combustion

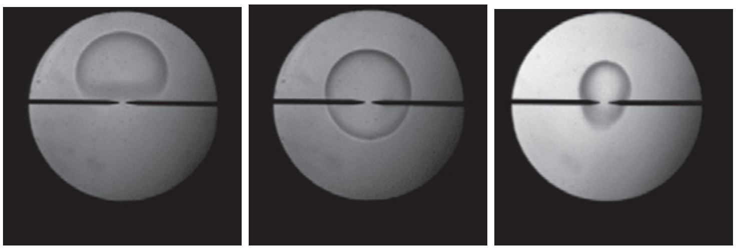

(a) $\phi = 0.8$ **(b)** $\phi = 1.0$ **(c)** $\phi = 1.2$

Figure 2.12: NH_3 Schlieren flame propagation at STP (T_0 = 21° C, p_0 = 1 bar) [8]

process, enhancing the reaction. This results in a higher velocity, with this parameter showing an exponential increase as the radical concentration increases, due to the exponential increase in chain-branching reactions.

- Carbon fuel addition: it leads to the production of more chain-branching reactions, but the increase is linear due to the more complex behaviour of these fuels [31, 32].

The results of these three changes in mixture composition are shown in Figure 2.14. Hydrogen addition will be considered primarily in this thesis as a means to increase velocity under challenging conditions, such as low temperatures and high pressures.

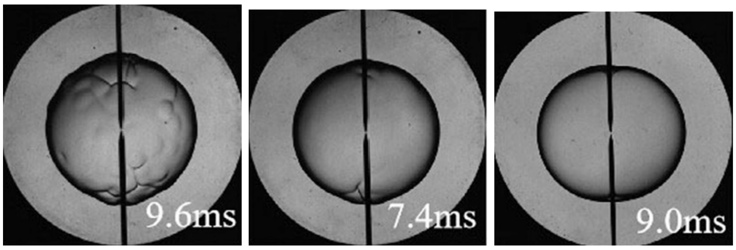

(a) $\phi = 0.8$ **(b)** $\phi = 1.0$ **(c)** $\phi = 1.3$

Figure 2.13: 50%NH_3-50%H_2 Schlieren flame propagation at STP (T_0 = 21° C, p_0 = 1 bar) [31]

2.4.1.2. Parameter dependency

The magnitude of the LBV depends on many parameters related to the combustion process; Law [15] gives a complete overview of this parameter dependence.

Initial Pressure

Increasing the initial pressure affects flame propagation through two coupled mechanisms: changes in chemical kinetics and transport properties. From a chemical standpoint, higher pressure increases the frequency of molecular collisions, which can, in general, accelerate elementary reaction rates. However, the same increase in density also enhances the importance of pressure-dependent three-body reactions, which tend to remove reactive radicals and therefore reduce the net chain-branching character of the chemistry. As a result, the radical pool that sustains the flame is depleted more efficiently at high pressure, thereby reducing the overall reaction rate within the flame zone and lowering the laminar burning velocity. This competition between radical-producing chain-branching steps and radical-removing termination/recombination steps is a generic reason why increasing pressure often leads to slower flame propagation [15].

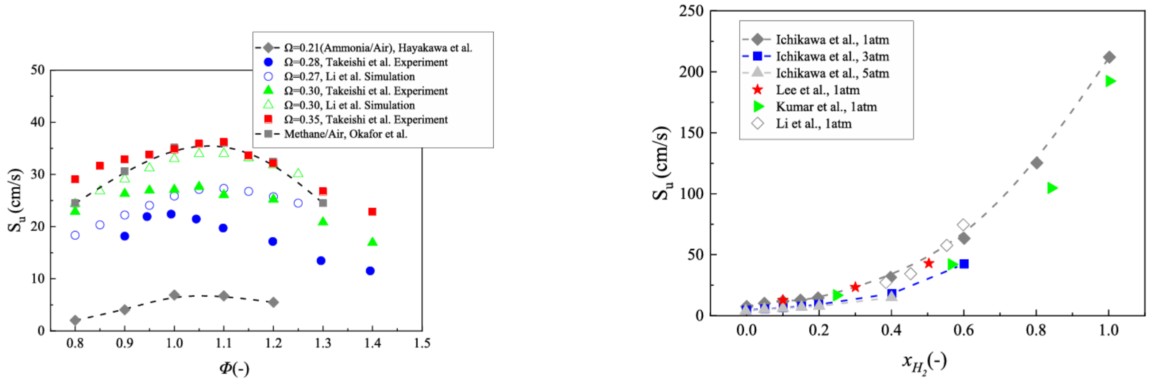

(a) Variation with increase in oxygen percentage Ω

(b) Variation with increase in hydrogen fraction

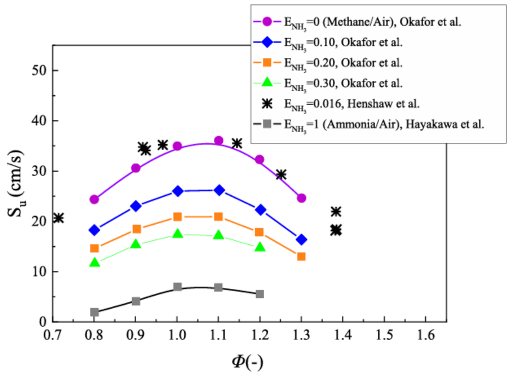

(c) Variation with increase in methane fraction

Figure 2.14: NH_3 flame speed variation with changes in mixture composition [31]

This behaviour can be illustrated using hydrogen chemistry, also present within ammonia reactions, where the balance between a chain-branching step and a three-body termination step becomes increasingly pressure-sensitive:

$$\text{chain-branching:} \quad H + O_2 \rightarrow OH + O, \tag{2.4}$$

$$\text{chain-terminating:} \quad H + OH + M \rightarrow H_2O + M. \tag{2.5}$$

The first reaction promotes a reactive radical pool by producing highly active radicals that sustain combustion, while the second reaction reduces that pool and becomes more frequent as the third-body concentration M increases with pressure. Although the specific dominant pathways differ between fuels, an analogous shift occurs for methane and ammonia because their flame propagation also depends on maintaining sufficient concentrations of key radicals (notably H and OH) in the preheat and reaction zones. Consequently, higher pressure typically shifts the chemistry toward more effective radical removal, thereby slowing the net flame chemistry and reducing the burning velocity [33].

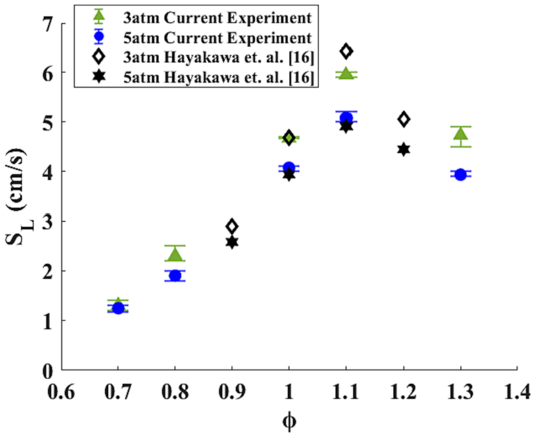

Figure 2.15: LBV variation with pressure for NH_3 [10]

In addition to kinetics, pressure modifies the transport processes that control the preheat zone ahead of the flame. Thermal diffusivity decreases as density increases, and since density is approximately proportional to pressure (for fixed initial temperature), heat diffuses less efficiently into the unburnt mixture at higher pressure. This reduces the rate at which the unburnt gases are heated to ignition-relevant temperatures, thereby lowering the flame propagation rate. Taken together, enhanced radical termination and reduced thermal diffusion provide a consistent physical explanation for why increasing initial pressure generally decreases laminar burning velocity for methane, hydrogen, and ammonia, even though the magnitude of the effect can vary by fuel and equivalence ratio [15]. The general behaviour is shown in Figure 2.15.

Initial temperature

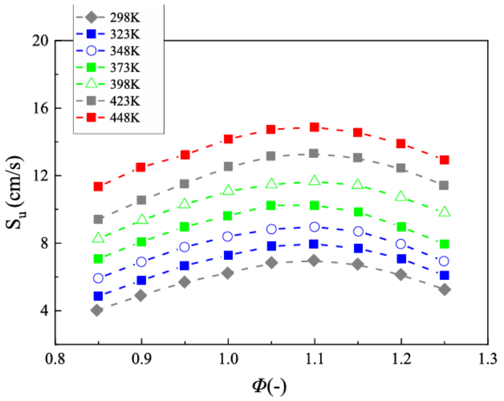

Figure 2.16: LBV variation with initial temperature for NH_3 [31]

Increasing the initial temperature generally increases the laminar burning velocity by accelerating both chemical kinetics and the transport processes that control flame propagation. From a chemical standpoint, most elementary reaction rates increase strongly with temperature through Arrhenius behaviour, so higher T_0 promotes faster radical production and faster overall heat release in the reaction zone [34]. As a result, the flame requires less time to build and sustain the radical pool that drives propagation, thereby increasing the net burning rate and, consequently, the velocity. This trend is reflected in the empirical correlations commonly used to represent laminar burning velocity, where S_L is often fitted with a positive temperature exponent of the form $S_L \propto (T_0/T_{ref})^\alpha$ [35]. Here, T_{ref} is the fixed reference temperature (ambient conditions) and T_0 is the unburnt initial temperature, showing that the speed tends to increase with increasing temperature.

Transport effects reinforce the same trend because a higher initial temperature tends to thicken the preheat zone and facilitate the transfer of heat and reactive species toward the unburnt mixture. While the exact temperature dependence of properties such as k and c_p varies with mixture composition, the key point is that flame propagation is controlled by how efficiently the unburnt gases are heated to reaction-relevant temperatures and how rapidly radicals can be sustained near the flame front. [34] Furthermore, at fixed pressure, an increase in T_0 reduces the unburnt-gas density, which increases diffusivities and generally promotes faster mixing and heat diffusion ahead of the flame. This is consistent with the same class of correlations reported in the literature, in which temperature and pressure appear as separate control variables with opposite effects on S_L (increasing with T_0 and decreasing with p_0). [35]

2.4.1.3. Calculation methods

Now that the basic physics of the laminar burning velocity and its dependence on both fuels and thermodynamic conditions are understood, the available methods for calculating it are presented. For this, a distinction needs to be made between experimental and numerical ones. The first are needed because, in non-optical contexts, velocity measurement is non-trivial; therefore, an easier way to determine it is to infer it from fundamental quantities such as pressure and temperature. On the other hand, as explained, this parameter has an intrinsic dependence on the kinetic behaviour of fuels, which can therefore be grouped into models that can be simulated at the desired thermodynamic conditions to extrapolate the velocity.

Experimental methods

In literature, many methods are found to calculate the velocity, including the Bunsen flame, counter flow or stagnation flame, burner-stabilised flat flame, hot-wire and outwardly propagating spherical flame methods [24]. However, most of these methods carry the inherent difficulty of calculating the velocity at high pressure. Therefore, since in this work the use of a non-optical chamber is considered, to ideally achieve high-pressure measurements, two main methods can be used to measure the laminar burning velocity experimentally:

- Spherically expanding flame method

  The method consists of reconstructing the entire flame history from either pressure-time or flame radius, computing the instantaneous burnt mass, and then computing the LBV. For this, the difference between the early and later stages of flame development must be understood. The former is characterised by high curvature and stretch but low pressure rise, while as the flame expands, the pressure increases and the stretch becomes negligible. Oppositely, the final stage of flame development is influenced by heat loss to the cold chamber walls. Therefore, the actual considered rise only contains the central part of the pressure curve.

  Depending on the chamber design, the constant-volume (CVM) and constant-pressure (CPM) methods can be used. The latter requires optical access because the pressure rise is very low, and the radius must be measured visually, while the former can be used with or without optical access. For this reason, the CVM method is used throughout this work. The method is applicable only when the flame radius is sufficiently large; otherwise, no appreciable pressure rise is observed. In fact, the LBV is determined by correlations based on pressure, pressure rise rate, and burnt mass fraction.

  The advantage of this method is that it can be used for high pressures, as no optical access is required. The disadvantage is that the potential for flame instabilities is not identified, which may reduce measurement accuracy. Metghalchi and Keck [36] develops a formula to calculate the LBV from the pressure measurements, first outputting the mass burning rate, based on the use of a correction factor for the equivalence ratio and the flame front area. Other authors, such as Stone, Clarke, and Beckwith [37] and Elia, Ulinski, and Metghalchi [28], use two- or multizone models that also include flame stretch. The zones in these models represent the areas where the flow is under different conditions, with the flame assumed to be infinitely thin:

  - Two-zone: the mixture is separated into an unburnt and a burnt zone, separated by an infinitesimally small flame front. While this method makes the computation much faster and easier, it misses the temperature gradients across different points within the burnt zone, as the properties are assumed to be constant within each point. In fact, the part closest to the flame front is moving faster than the part that has already been burnt and is farther away. Therefore, this method trades accuracy for computational speed. One of its uses could be to provide an initial estimate of the speed to determine whether the data point is sensible or an outlier.
  - Multizone: the mixture is separated into multiple burnt zones, represented by concentric circles, either constant in mass or radius. This is

much more accurate, as it can better predict phenomena such as heat loss to the walls and instabilities, as shown in the temperature gradients [24].

- Hot-wire method

  In 1972, Bradley and Hundy [38] developed the hot-wire method. For this, a thin hot wire is placed inside the combustion chamber at a known position. It is then heated with a constant current, and the resistance through it is constantly monitored. As the flame propagates, it pushes unburnt gases outward, creating a gas velocity field ahead of the flame. The wire measures the unburnt gas velocity before the flame reaches it via hot-wire anemometry, which relates the velocity to the wire's cooling. This cooling effect is measured by monitoring changes in electrical resistance. Once the unburnt gas speed is known, the flame speed can be measured optically or by using ionisation probes. The LBV is simply the difference between the two.

In this thesis, since the use of a hot wire is impractical, especially at the desired temperatures, and because of the added complexity of the chamber design and the electrical architecture, which introduces electrical noise, the first method is used. This also guarantees a larger body of literature results calculated with this method, which eases comparisons and removes bias when comparing methods with different assumptions.

Numerical methods

Laminar burning velocity, S_L, is most commonly obtained numerically by solving canonical premixed-flame problems in which the flame structure is resolved, and the propagation speed follows from the conservation equations coupled with chemical kinetics and transport. These models are then validated by comparing numerical results with high-fidelity experimental data collected over a range of temperatures, pressures, and equivalence ratios. For the purposes of this thesis, the most relevant numerical approaches are those that compute S_L under well-defined initial pressure and temperature and under premixed, laminar assumptions, since these conditions are directly comparable to the early-flame regime of interest. In that sense, one-dimensional premixed-flame solvers constitute the standard and most transferable methodology, while multidimensional approaches are mainly used when curvature, instability, or flow coupling must be explicitly studied [34].

- 1D freely propagating premixed flame (unstrained, planar)

  This is the most common and direct numerical route to S_L, because the flame propagates into a uniform premixed mixture without imposed strain, and the resulting S_L is obtained as part of the steady solution. Standard implementations include software such as CHEMKIN/PREMIX and open-source solvers such as Cantera, which explicitly compute the laminar flame speed for specified T_0, p_0, ϕ, and mixture composition.

- 1D burner-stabilised premixed flame (flat flame)

  In this configuration, the flame is stabilised on a burner, and the characteristic flame speed is inferred from the imposed mass flux and the computed flame structure. This approach is widely used both experimentally and numerically, and it is implemented within the same family of 1D premixed-flame solvers used for freely propagating flames.

- 1D strained premixed flame (counterflow) with extrapolation to zero strain

When stretch effects are of interest, premixed flames can be computed at controlled strain rates and then related to an unstrained burning velocity, often alongside stretch-sensitivity metrics. This method is useful to interpret conditions where curvature or strain is non-negligible, and it provides a bridge between ideal planar propagation and more practical flame configurations.

- Other approaches

  Multi-dimensional DNS can resolve laminar premixed flame propagation with high fidelity and is often used to study flame instabilities or curvature effects, but it is computationally intensive and not typically used to map S_L over broad parametric ranges [39]. In turbulent-combustion simulations, S_L is more often introduced through reduced or tabulated representations (e.g., flamelet-generated manifolds) or front-tracking methods (e.g., level-set/G-equation), and thickened-flame approaches are used in LES to make premixed flames resolvable on practical meshes; in all cases, these methods generally rely on S_L obtained from canonical 1D calculations or correlations [40, 41].

For the present work, the freely propagating 1D premixed flame calculation is the most appropriate numerical reference because it provides a well-defined definition of S_L under specified T_0 and p_0 without additional geometric or flow-induced effects. This makes it particularly suitable for comparing trends across fuels and for discussing how environmental variables, such as pressure and temperature, can shift early flame propagation in a controlled manner. The burner-stabilised and strained-flame variants remain useful as supporting frameworks, especially when interpreting departures from ideal behaviour due to stretch or heat-loss sensitivity. At the same time, it is important to distinguish these canonical S_L computations from turbulent-combustion closures, where S_L is typically an input or a tabulated quantity rather than an independently solved output. [34, 40] Consequently, when numerical results are used in this thesis, they are interpreted primarily as reference laminar quantities aligned with early-stage flame development, rather than as full representations of practical combustor operation.

2.4.2. Minimum ignition energy (MIE)

Any combustion process begins with the ignition of a specific mixture of fuel and oxidiser. Before this, the components of this mixture are generally stable, with their atoms held together by strong bonds. This binding is strong because each bond has a specific potential energy, known as the bond energy, that holds the atoms together. Thus, to initiate combustion, these bonds must be broken, which causes the release of this energy, which, when high enough, causes the breaking of new bonds and the formation of radicals, which maintain the reaction. In fact, throughout this transformation, the composition of the mixture changes, leading to the formation of products at a lower energy level and ultimately releasing energy.

Hence, initially, it takes some external energy to break the stable fuel bonds. This can come from many different source types, which are discussed in more detail in Section A.2, but the process that leads to the first rupture of bonds is equivalent and is shown in Figure 2.17. Here, the example is given for the conversion of methyl isonitrile (H_3CNC) to acetonitrile (H_3CCN), but the same principles apply to any other molecule, albeit simpler for one such as hydrogen. At first, the molecule is in a stable state and resists deformation, but once enough energy is provided, it begins to bend. Then, if this energy is sufficient, the excessive bending causes bond rupture, and this energy is called the activation energy, which allows the formation of complex forms, that is, hybrid states at a higher energy level. These unstable configurations lead to the formation of lower-energy bonds, as energy is released when bonds break. Finally, the new stable product is formed, carrying a lower potential energy than the initial component [42].

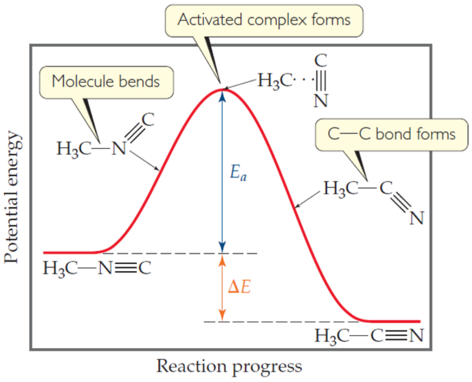

Figure 2.17: Ignition energy [42]

At this point, if the energy released is sufficient to activate other molecules and initiate a chain reaction, the external energy is called the ignition energy (IE). However, this can be much higher than the amount actually required to ignite a mixture, and it is therefore interesting to study the minimum energy required to ignite a mixture. This is both for safety reasons, as it is important to understand whether an accidental electrostatic discharge may ignite a mixture at rest, and from an efficiency standpoint, as it is beneficial to only provide the minimum necessary energy instead of wasting usable work. This parameter, called the minimum ignition energy (MIE), is a key focus of this study and is widely used to design ignition systems and assess explosion and fire hazards. Formally, it is normally defined as the minimum amount of energy leading to ignition in at least 50% of cases [43]. As with the laminar burning velocity, this parameter also shows important dependencies that must be analysed to understand its behaviour under the considered experimental conditions.

2.4.2.1. Fuel dependency

As with the LBV, the minimum ignition energy (MIE) is strongly fuel-dependent because it reflects how readily a mixture can generate and sustain a self-propagating radical pool after a finite energy deposition. Fuels that produce chain-carrying radicals rapidly and release heat efficiently during the earliest stages of kernel growth tend to exhibit low MIE, whereas fuels with slower branching chemistry and larger relative heat losses require substantially higher ignition energies. This behaviour is closely linked to intrinsic reactivity and transport properties, since diffusion and heat transfer govern how effectively the deposited energy is retained in the kernel and converted into chemical activity. Consequently, the same underlying physical factors that promote fast flame propagation also generally facilitate easier ignition, which explains why trends in MIE and LBV often remain consistent across different fuels.

This is reflected in the typical magnitude of reported values. Most light hydrocarbons (e.g., methane, propane, and butane) exhibit MIE in the order of 0.1 mJ, while hydrogen is significantly lower, around 0.019 mJ [44]. In contrast, ammonia exhibits MIE values that are orders of magnitude higher, typically of the order of 20 mJ and in some conditions approaching 200 mJ [45]. These differences highlight why hydrogen can ignite very easily and why ammonia ignition is comparatively challenging. In practice, this often implies that ammonia requires either substantially higher ignition energies, reduced ignition gaps, or the use of more reactive additives to reliably achieve ignition under comparable conditions.

2.4.2.2. Parameter dependency

Naturally, this parameter also depends on other variables related to the initial state and the combustion environment. This dependency will be explained here.

Equivalence ratio

Another parameter that affects the energy level needed to ignite a mixture is the equivalence ratio. Indeed, as shown in Figure 2.18, research shows that the relation between the MIE and this factor follows a U-shaped curve, with an optimal value at which the required energy is minimised. The sharp increase in energy at low and high equivalence ratios is due to the approach to the flammability limits. In fact, when the flame is too rich or too lean, there is not enough fuel or oxidiser to ionise enough of the mixture to start a chain reaction. In these conditions, a considerable amount of energy is required to increase the collision rate and sustain the reaction.

As ϕ is increased from very lean conditions, the needed energy decreases, as the mixture proportion is more favourable towards the production of radicals, entailing a higher energy release. This means that less input is needed to sustain the reactions, and the chain reaction occurs much more easily. A minimum MIE point is then achieved, at which the reaction is at the optimal ratio to produce the highest amount of heat. This corresponds to the point of maximum burning velocity and adiabatic temperature in a fully developed flame. Normally, this point occurs at a slightly rich equivalence ratio, around 1.1-1.2, at which the chemical kinetics are optimised.

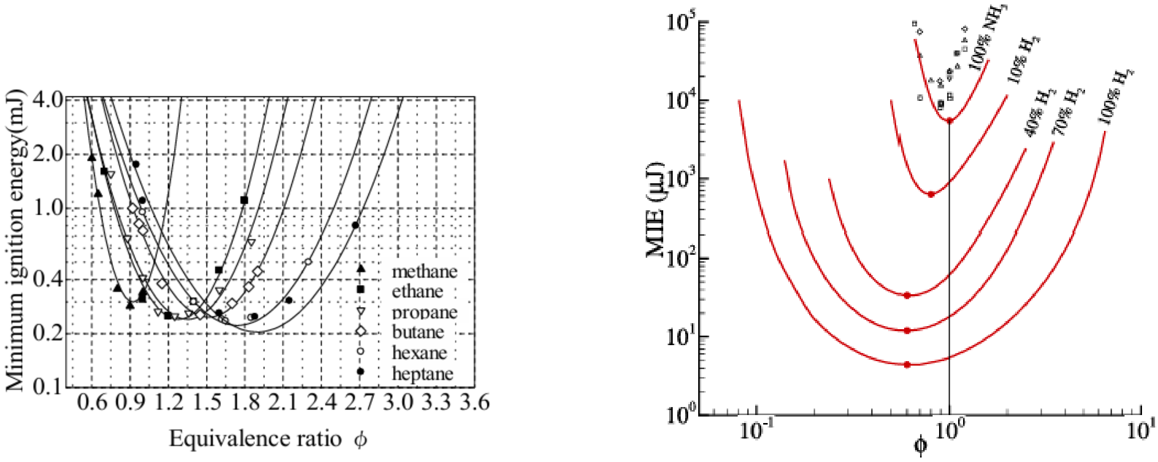

(a) CH_4 and other hydrocarbons [46]

(b) NH_3-H_2 mixtures [47]

Figure 2.18: Variation of MIE with ϕ [46]

However, from Figure 2.18, something else seems to happen, as the equivalence ratios at which the minimum energy is achieved are different from the point of maximum velocity. This is because there is a difference between the fully developed flame and the flame at the ignition point, which has not yet fully formed. In fact, at this position, the flame is a purely spherical kernel, as the energy is released at a single point almost instantaneously, resulting in a sudden increase in temperature and pressure. Then, after the brief energy surge, the flow expands equally in all directions due to the isotropic nature of the surrounding unburnt gas. Naturally, to attain perfect isotropy, complete mixing of the fuel and oxidiser must be achieved, which is not normally the case.

In reality, a sphere is still created, with some strain due to local variations in equivalence ratio. Due to the strong curvature, the flame speed is not the same as the laminar burning velocity. This, in addition to the lower temperature at this stage due to the low amount of mixture burnt, implies that the speed will be lower and will be referred to as the initial burning velocity [48]. This is evident in Figure 2.19. To understand why the MIE is minimum at a specific equivalence ratio, the Lewis number must be considered.

In a mixture, the deficient reactant controls the rate at which the mixture is supplied to the reaction zone. For instance, in a rich mixture, there is enough fuel for combustion, so the point at which it begins and the rate at which it occurs will be solely controlled by the rate at which the oxidiser diffuses to interact with the fuel. The Lewis number as a function of the equivalence ratio for different fuels is reported in the literature, with some differences due to the methods used to calculate it. This is shown in Table 2.2. These trends provide a consistent framework to interpret why the equivalence ratio associated with minimum ignition energy does not necessarily coincide with the conditions that maximise burning velocity in a fully developed flame.

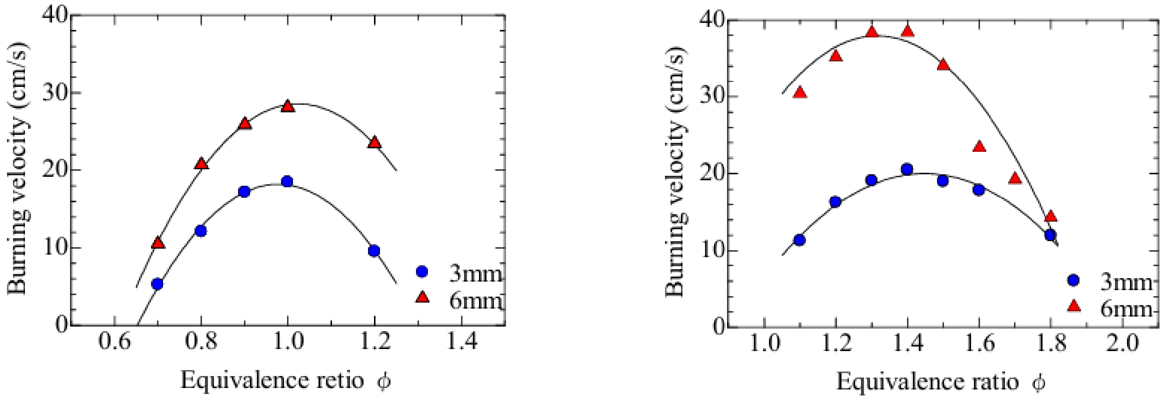

(a) Methane **(b)** Butane

Figure 2.19: Velocity variation from initial to fully expanded flame [46]

Table 2.2: Lewis number for methane and butane [46]

Methane	ϕ	0.8	0.9	1.1	1.2		
	Le	0.883	0.886	1.003	1.003		
Butane	ϕ	0.9	1.1	1.3	1.5	1.7	1.9
	Le	1.745	0.929	0.918	0.907	0.897	0.887

For methane, the mixture Lewis number is lower than 1 in lean conditions. This means that, in this setting, methane mass diffusion is higher than thermal diffusion, encouraging combustion and thus increasing the flame speed at its surface. However, for butane, the opposite behaviour occurs, with the favourable conditions found in rich conditions. This is because the higher molecular weight, due to the presence of more carbons, slows the molecule, reducing its diffusion. This is confirmed by examining the trend in the Lewis number for propane, which is higher than for methane but lower than for butane. For these fuels, the maximum initial burning velocity is therefore reached at different equivalence ratios, which explains the shift in the points of minimum ignition energy observed in Figure 2.20. Murase et al. [46] defines that the minimum radius at which the flame can be considered self-sustaining, and thus identified as a kernel, is of 3 mm. This is therefore the size at which the initial burning velocity is defined, while a radius of 6 mm corresponds to a more fully developed flame, where stretch is reduced, and the laminar burning velocity is taken. [46]

From this discussion, it is possible to extrapolate the behaviour of fuels such as hydrogen and ammonia by interpreting how their relative mass and thermal diffusion shift the Lewis number with equivalence ratio. Hydrogen has been shown to have the lowest MIE at an equivalence ratio of about 0.65 [44], which is consistent with its very high diffusivity, as described in Subsection 2.3.2. In lean mixtures, hydrogen diffuses into the reaction zone faster

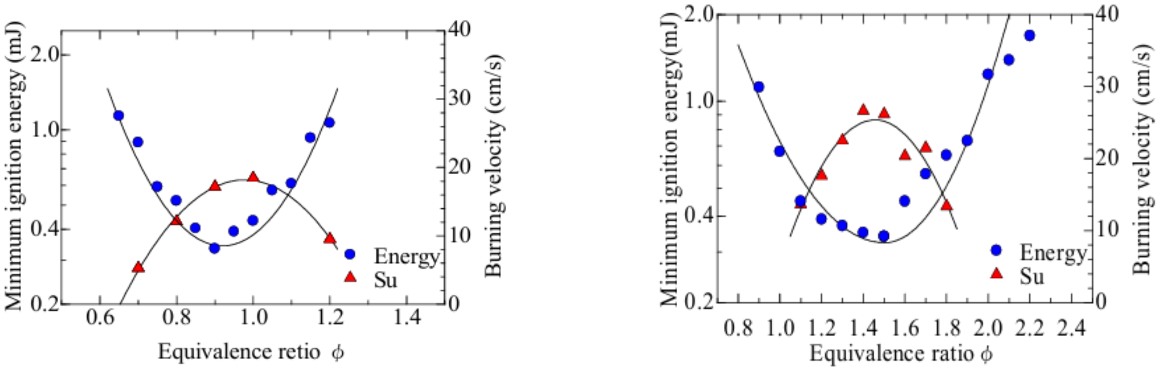

(a) Methane

(b) Butane

Figure 2.20: Relationship between LBV and MIE [46]

than the oxidiser can counterbalance it, leading to a Lewis number below 1 and, therefore, higher initial burning velocities under lean conditions. However, because hydrogen diffusion is much stronger than for carbon fuels, the equivalence ratio at which this balance becomes most favourable shifts further into the lean region, which explains why the minimum MIE occurs at a lower ϕ than for methane or propane. In contrast, ammonia is reported to have diffusion behaviour closer to methane [20], meaning that its Lewis number trends with ϕ are less extreme than hydrogen and do not push the optimum as far into lean conditions. As a result, ammonia exhibits an optimal equivalence ratio closer to unity, with values around $\phi \approx 0.9$ reported in the literature [12], which makes the extrapolation consistent with the same Lewis-number argument used above.

Initial pressure

Higher pressure favours combustion because it increases the number of collisions and reduces dissociation, according to Le Châtelier's principle. This precept is based on the fact that combustion occurs in dynamic equilibrium rather than a static state, since the reactions that produce the product are reversible. Thus, the point at which the mechanism is at equilibrium is the one at which the rates at which forward and backward reactions occur are equal. The principle states that when this dynamic equilibrium is disturbed by changing the conditions, the equilibrium position changes, consequently, to re-establish equilibrium. This means that if the pressure is increased, the reaction will try to move towards a condition of lower pressure. Therefore, mechanisms that produce fewer molecules will be prioritised, as they occupy less space and thus tend to lower pressure. This means that complete combustion will be favoured, as it produces CO_2 and H_2O, which will not dissociate into CO and O_2, or H_2 and O_2, respectively. This means that, in general, higher initial pressure results in a lower MIE. However, at very high pressures, the thermal conductivity and the number of collisions increase to a level where quenching can occur. This normally happens at pressures higher than engine-like ones, so well above 50 bar [42, 49].

Initial temperature

Similarly to pressure, a higher temperature also favours combustion. This is because it allows for faster chemical reactions, lowering the required energy. In fact, higher temperature increases the kinetic energy of the reactant molecules, leading to a higher rate and frequency of collisions. This also results in a reduction of the ignition delay time and an acceleration of the reaction rate [42].

2.4.2.3. Calculation methods

As for the LBV, the literature describes numerous methods, both experimental and numerical, for determining the MIE. In this thesis, the experimental approach relies on measuring ignition parameters (voltage, current, and discharge characteristics) and using them to quantify the energy delivered to the mixture. Numerical methods can then be used to provide a reference for comparison and to assess the consistency of the experimental results under comparable conditions. For this reason, the main available approaches are briefly reviewed below, with emphasis on those most relevant to the present setup. Based on this overview, one method is identified as the most appropriate for this work's objectives.

Experimental methods

Various experiments have been run over the years to determine the MIE necessary to ignite specific mixtures. All of these methods work similarly to the Golden-section search algorithm to find the minimum. In this section, a summary of three of the most accurate and useful will be presented. These are applied to spark ignition but can be easily modified to work with any of the previously described ignition methods.

- Method from Lesmana et al. [43]:

  In this method, the ignition energy is calculated by integrating the product of voltage and current over time, as in Equation 2.6.

  $$E = \int_0^t u(t)i(t)dt \tag{2.6}$$

  A notable feature of this method is the introduction of an inductor into the spark generator circuit to vary the duration of the spark discharge. Moreover, the possibility to vary the gap is introduced, allowing identification of MIE variation with these parameters, not only the operating and mixture conditions.

  1. Take a starting energy level and perform 5 sequential ignition attempts.
  2. If none of the attempts leads to the formation of a kernel, the energy level is raised by either varying the voltage or the capacitance.
  3. If at least one of the 5 attempts leads to the formation of a kernel, then the upper limit for the energy is set.
  4. The energy level is decreased in steps of 2 mJ until no ignition is found in 5 consecutive tries, which defines the lower energy limit.
  5. The energy level is increased again in steps of 0.1 mJ until one ignition is found, which is defined as the MIE.
  6. One more attempt is run at an energy level lower by 0.1 mJ to ensure that no ignition is found.

  One of the issues with this method is its instrumental sensitivity of 0.1 mJ to perform the steps, while guaranteeing a high repeatability of ±3-8%, with the highest scatter observed for ammonia. The disturbances observed during MIE measurements of hydrogen are attributed to electrode condensation. Moreover, it has quite a wide ignition/no-ignition band.

- Method from Proust [50]:

In this method, the author separates the breakdown and glow phases of the spark using two circuits: a high-voltage circuit and a low-voltage one. The first one is aimed at producing the plasma stream, which requires a high voltage. Once this is created, a high voltage is no longer needed; a lower one is enough to sustain the current. This is done so as to reduce the losses caused by the high current achieved in the second phase. The ignition energy is then found as the sum of the energies required by the two different circuits. The method works similarly to the one proposed by Lesmana et al. [43], with the difference that, in this case, 10 tries are performed at each energy level. However, here ignition is required in all cases to consider the energy level sufficient for ignition. If at least one trial results in no ignition, this level is considered below the MIE. This is then taken as the lowest deposited spark energy leading to ignition in all trials. For cases of overlap between ignition and non-ignition, particularly relevant for hydrogen, the band centre is taken as the MIE. The author explains that the reason for the disturbances in the case of hydrogen is due to humidity, which is, however, disproven by Ono et al. Ono et al. Therefore, a possible explanation is the condensation of water on the electrodes. The repeatability of this method is about ±5% for propane and ±15% for hydrogen. In another of his works [52], the author shows that there is a difference between the cases of no ignition and ignition, and vice versa: in the former, the MIE is found to be much higher. A possible explanation is that the combustion process chemically activates the electrode tips and removes any residual.

- Method from Ono et al. [51]:

  In this method, only one trial is used for each energy level. An interesting point is that the author accounts for energy loss in the resistor and the residual charge in the capacitor, both of which must also be considered part of the total MIE. The method results in a scatter of ±3%.

Numerical methods

To confirm the results found with experiments, some analytical methods have been developed. Here, three are summarised:

- Method from Cirrone et al. [53]:

  Since most methods in the literature rely on experimental data on either the quenching distance or the laminar burning velocity to calculate the MIE, they ultimately amount to empirical correlations. Therefore, they cannot be used in any case where experimental data are unavailable, nor can they be extrapolated, as this would result in considerable errors. Thus, since the paper's focus is the cryogenic temperature range, for which experimental data are considerably limited, the author develops a method that accurately predicts the MIE for any hydrogen composition and temperature, which can then be applied to other fuels. The procedure is as follows:

  - The unstretched laminar burning velocity S_u^0 is calculated with the "Freely propagating adiabatic 1D flame" solver in Chemkin.
  - The parameter X_{SD} is set based on the fuel concentration to account for the effect of flame stretch and preferential diffusion.
  - The stretched laminar burning velocity S_u^{SD} is calculated as the product of the two above terms.

  - All the thermophysical properties of both the unburnt and burnt mixtures are determined with the Equilibrium solver of the software Cantera.
  - The diffusive laminar flame thickness δ is calculated.
  - The flame thickness is related to the heat release via Blint's definition, resulting in δ_L^B.
  - With this, the critical diameter for kernel propagation is determined, which is assumed to be 2.5 times the flame thickness.
  - the MIE is found as

    $$MIE = \frac{1}{6}\pi d^3 \rho_u c_{p,u}(T_b - T_u) \tag{2.7}$$

    where

MIE	Minimum ignition energy
d	Critical kernel diameter
ρ_u	Unburnt mixture density
$c_{p,u}$	Unburnt mixture specific heat at constant pressure
T_b	Burnt mixture temperature
T_u	Unburnt mixture temperature

  This method introduces a novel flame stretch parameter, extending beyond the capabilities of both Cantera and Chemkin.

- Method from Nawaz and Mack [45]:

  This method is very similar to the one just presented, with the only difference being how the correction factor to account for stretching is determined. Indeed, in this case, the calculation starts from the initial experimental MIE values. In fact, initially, the method just explained is reversed to find $X_{exp,fit}$. From this, a fourth-degree interpolation is used to obtain the remaining values of this parameter, which are then used to compute the MIE. This introduces a dependency on the experimental values, which depend on conditions such as the spark gap.

- Method from Wu et al. [54]:

  This method stems from the use of an in-house code called INSFLA, which describes the kernel as spherically expanding and ignited by an external energy source, with energy density

  $$q = q_{max} e^{-(\frac{r}{r_s})^8} \tag{2.8}$$

  Where r_s is the ignition radius and q_{max} is the maximum energy density. By integrating this over the flame radius, the MIE can be found.

2.5. Non-optical diagnostics in combustion analysis

This section summarises the most commonly used experimental approaches for characterising ignition and early flame development, with emphasis on techniques compatible with the objectives and constraints of this work. Since the present thesis focuses on minimum ignition energy (MIE) and laminar burning velocity (LBV) under controlled initial conditions, the discussion is narrowed to methods that can generate repeatable, quantitative metrics rather than qualitative visual descriptions of the flame. In the literature, this typically leads to two families of approaches: optical methods, which directly measure flame geometry and propagation, and non-optical methods, which infer combustion behaviour from global chamber signals such as pressure and temperature.

Non-optical techniques are widely used in combustion experiments whenever optical access is limited, when the aim is to minimise complexity, or when robust operation is required over broad pressure and temperature ranges. Instead of tracking the flame front directly, these methods rely on signals that respond to the integrated heat release and the chamber's global thermodynamic evolution, most notably pressure traces and temperature measurements. For early-stage combustion analysis, this is particularly useful because ignition success, ignition delay behaviour, and early heat release can often be identified directly from the onset and slope of the pressure rise. In addition, non-optical measurements are typically easier to integrate in high-pressure and subzero test conditions, where windows, illumination, and image-based calibration can become limiting factors. For these reasons, the experimental framework in this thesis is built around pressure-based diagnostics, supported by temperature sensing and system monitoring, and is later used to extract both MIE and flame-propagation metrics.

Naturally, especially for non-optical combustion chambers, the quality of MIE and LBV calculations highly depends on how accurately the chamber state and the ignition event are measured. For this, accurate combustion diagnostics techniques are needed to measure and analyse the flame occurrence, along with the evolution of flow parameters such as temperature and pressure. These are particularly important in a system without optical access, as they need to map how the combustion process occurs without actually being able to see it. The main limitation of this type of sensor is its intrusiveness, as it must be placed inside the chamber. Therefore, they need to be designed and placed to minimise interference and losses affecting flame propagation. The main types of sensors applicable to non-optical-chamber measurements for LBV and MIE calculations are described below, grouped by the measured quantity.

2.5.1. Temperature

Temperature sensors are very delicate, as they need to measure the flow temperature without directly affecting it. In fact, since this type of sensor typically requires in-chamber protrusion to ensure sufficient heat transfer, measurements can be affected by errors arising from this protrusion, including conduction losses through the probe, radiative exchange with the surroundings, and finite response time. In fact, temperature sensors tend to have relatively high response times, which makes them suitable for slow combustors but less useful for capturing instantaneous ignition transients; they can instead be used to define the initial temperature. Finally, the sensor can also disturb the flow itself, as the gas must wrap around the probe, altering local velocity and, consequently, convective heat transfer at the measurement location.

Temperature sensors can be divided into two macro-categories, mechanical and electrical, depending on the way the measurement is taken. Mechanical sensors, however, are not commonly used in combustion applications and are therefore not considered here. Electrical sensors are instead the standard approach, since they convert temperature changes into measurable variations of an electrical quantity that can be recorded with appropriate instrumentation. In practice, this is typically achieved by changing resistance, voltage, or thermoelectric potential within the sensing element, which is then read as an electrical signal. This combination of relatively straightforward implementation and high measurement sensitivity explains why electrical temperature sensors remain the most common choice in combustion systems. The main types are described here.

2.5.1.1. Thermocouple

Thermocouples are formed by two different conducting materials joined together at the hot junction, while the other ends are maintained at a different temperature, as shown in Figure 2.21a. This configuration produces the Seebeck effect, namely the generation of a voltage in the circuit due to the dissimilar properties of the two metals when a temperature gradient is applied. The physical origin is that, as the junction is heated, charge carriers gain energy, establishing a net electromotive force along the conductors and leading to a measurable potential difference. Since the two materials respond differently to the same temperature gradient, the resulting voltage is directly related to the temperature difference between the hot and cold junctions. For this reason, thermocouples inherently measure a relative temperature, and the voltage must be interpreted with respect to a known reference temperature. This, however, can be accounted for by proper calibration.

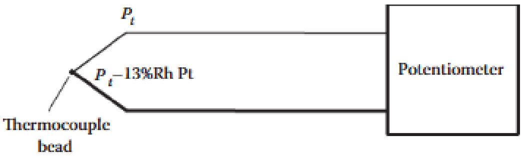

(a) Simple thermocouple

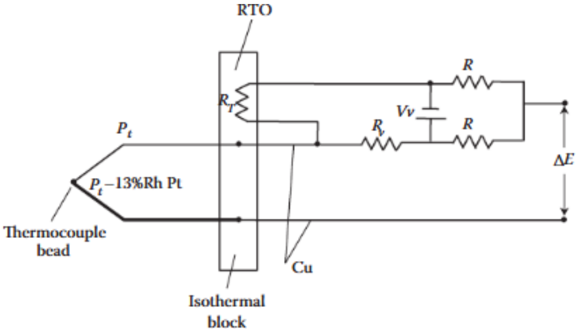

(b) Thermocouple with reference temperature

Figure 2.21: Schematics of a thermocouple [55]

The main limitations of thermocouples arise from the way measurements are produced and from the fact that the sensor itself is a thermal body. Response time is finite, as the junction must first exchange heat with the surrounding flow before the measured voltage stabilises, and this can limit the ability to capture very fast transients. In addition, radiative heat exchange at the junction can bias the measured temperature downward, particularly when the junction is exposed and the surroundings are colder. Another practical point is that the measurement circuit often includes additional metals, which introduce extra junctions and therefore additional temperature differences; this is addressed by enforcing a cold-junction reference at a known temperature or by applying cold-junction compensation, as shown in Figure 2.21b.

Despite these limitations, thermocouples remain flexible because many types exist depending on the metal pair used, such as K-type (Ni-Cr) and B/R/S types (Pt-Rh alloys). Each type is optimised for a specific temperature range, with some operating well in subzero conditions, such as E-type (NiCr-Constantan). In these environments, absolute accuracy can decrease due to lower signal levels, but thermocouples remain attractive for their robustness, relatively fast response compared with many contact sensors, and low cost. It should be emphasised that, because the thermoelectric voltage is small, careful attention to wiring, shielding, and cold-junction compensation is required to limit noise and drift, especially when long cable runs or electrically noisy ignition systems are present.

2.5.1.2. RTD

This sensor type measures temperature variation by monitoring changes in the metal's resistance. In fact, like any material, metals have an inherent resistance to conduction due to the interaction

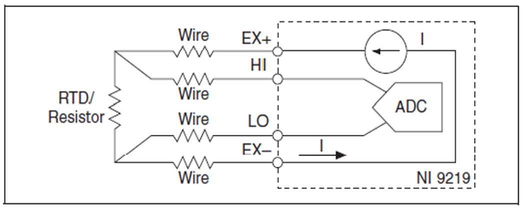

Figure 2.22: RTD working principle [56]

between the moving electrons and the atomic structure. The change in the resistance for metals can be described as:

$$R(T) = R_0 + (T - T0)\frac{dR}{dT} \tag{2.9}$$

The formula shows that the resistance behaves almost linearly, at least up to high temperatures. This type of sensor offers high stability and precision when measuring absolute temperature rather than relative temperature. To measure the voltage drop caused by a change in resistance, 3 or 4 cables are needed: two to supply current and the other one or two to measure the voltage drop across the sensors. The reason for using more than two cables is that the metal itself has low resistance, so the cable resistance is not negligible and must be eliminated using a so-called Kelvin connection. This implies that the response time is quite high, and the sensor may contribute to temperature measurements through some self-heating due to its resistance. Due to the setup, they are more expensive, while at the same time ensuring a higher precision, even at temperatures as low as -200° C.

2.5.1.3. Thermistor

This sensor is similar to an RTD, with the main difference being that the sensing element is a semiconductor rather than a metal. As a result, the resistance exhibits a strongly non-linear temperature dependence, which is commonly expressed as:

$$R(T) = R_0 e^{(B/T - B/T_0)}. \tag{2.10}$$

Since semiconductor resistance is considerably higher than that of metallic elements, the lead-wire resistance is typically negligible; hence, a two-wire measurement is usually sufficient to determine the voltage drop across the sensor. Thermistors can be very accurate, including at subzero temperatures, but they are also fragile, and their non-linearity tends to become more pronounced over time, increasing the need for calibration and drift control. For this reason, they are generally not well-suited to aggressive mixtures or chemically challenging environments, such as those containing ammonia. In addition, self-heating can be significant due to the sensing element's higher resistance, potentially biasing the measured temperature and limiting the usable measurement range if the excitation current is not carefully controlled [57].

2.5.2. Pressure

Pressure sensors can be divided into three main categories based on measurement method, measurement type, and pressure reference. The first division, as with temperature sensors, is between mechanical and electrical devices. Again, the mechanical ones are not discussed here, as they are not of interest for combustion analysis and do not readily integrate with high-speed data acquisition. The second macro-division is between static and dynamic sensors: the former measures the quasi-steady force per unit area exerted by the mixture (or the chamber) and is mainly used to track slow pressure evolution, whereas the latter is intended to resolve time-varying pressure fluctuations linked to flow unsteadiness and fast combustion phenomena, meaning that frequency response and signal conditioning become central. Finally, the pressure reference indicates whether the sensor is referenced to absolute zero pressure, to atmospheric pressure (gauge), or to a differential between two points, which is particularly relevant when pressure drops across components or when pressure drops between locations are of interest.

2.5.2.1. Piezoelectric

Piezoelectric sensors contain a crystal (commonly quartz or piezoceramics) that generates an electrical charge when mechanically stressed by pressure, so the measured signal is directly tied to the applied load. In practice, the output is read through a high-impedance/charge amplifier because the sensor behaves like a charge source in parallel with a capacitance, and any leakage paths will otherwise cause the signal to decay over time. This is why these transducers are excellent for dynamic measurements (fast pressure rise, oscillations, shocks), while truly steady, long-duration pressures are not their natural operating regime unless the full measurement chain is designed with a very long time constant. One of their main advantages in combustion work is the combination of high stiffness and high natural frequency, which supports high bandwidth, and a generally linear response with negligible mechanical hysteresis under normal loading. The main practical limitation in flames is not only bulk temperature drift but also thermal shock at the diaphragm (cyclic heating during combustion), so prolonged or direct flame exposure typically requires thermal protection strategies such as recessed mounting, protective coatings, or water-cooling/thermal barriers, depending on the required timescales and accuracy.

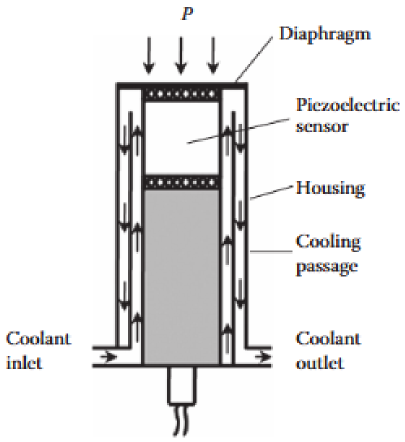

Figure 2.23: Piezoelectric transducer [55]

2.5.2.2. Piezoresistive

Piezoresistive sensors rely on the change in electrical resistance of a semiconductor (typically silicon) when it is strained by a pressure-loaded diaphragm. The resistors are usually arranged in a Wheatstone bridge to convert strain into a voltage signal. Compared with piezoelectric devices, they are well suited to static and low-frequency measurements and can be built as absolute, gauge, or differential sensors depending on how the reference side is packaged and vented. Their main drawback is that both the bridge offset and sensitivity are strongly temperature-dependent, which is why practical devices often require careful temperature compensation (either on-chip or post-processing) to maintain accuracy across operating conditions.

From a combustion perspective, the limiting factor is often the temperature range: conventional diffused silicon piezoresistive structures are commonly constrained to roughly 120–150° C because leakage mechanisms increase with temperature, which quickly becomes problematic near hot gases unless the sensing element is protected. In setups where they are still desirable (e.g., for absolute pressure level), mitigation is typically done by thermal shielding/insulation, active cooling, or remote mounting via an indicator passage line[57].

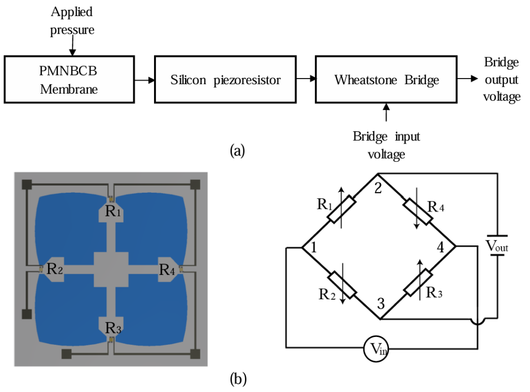

Figure 2.24: (a) Piezoresistive sensor; (b) Wheatstone bridge [58]

2.5.2.3. Capacitive

Capacitive pressure sensors measure pressure by tracking changes in capacitance between two electrodes, with one electrode a diaphragm that deflects under pressure, changing the gap (or overlap) between the plates. Because capacitance can be measured with very low power and high resolution, this approach is particularly attractive for low-pressure sensing and for applications where long-term stability and low drift are priorities.

At the same time, the measurement is sensitive to parasitic capacitances and to temperature-dependent changes in diaphragm stiffness, residual stress, and dielectric properties, so thermal compensation and robust readout electronics are often central to achieving reliable performance. The "speed" of these sensors is not inherently slow: the ultimate bandwidth is set by the mechanical resonance of the diaphragm and the readout circuit, but many practical packaged devices are optimised for stability/noise rather than for high-frequency combustion transients. As a result, they are often less convenient than piezoelectric transducers for fast in-chamber pressure dynamics, and, when used in harsh thermal environments, they typically require both thermal management and a carefully designed electronic front end to avoid temperature-driven offsets and gain changes.

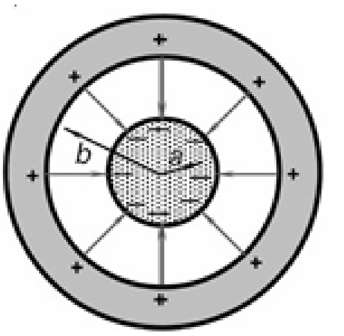

Figure 2.25: Capacitive transducer [59]

2.6. High-pressure combustion

High-pressure combustion is of broad interest because it is the regime in which most practical devices operate, as it generates the highest power release and thermal efficiency. Compared to atmospheric pressure, the chemical effect is substantial, altering chemical timescales, transport properties, flame thickness, and the balance between chain-branching and radical-termination pathways. In fact, as pressure increases, collision frequencies rise, and pressure-dependent three-body reactions become more dominant. This heightened complexity not only affects ignition dynamics but also influences pollutant formation and overall combustor performance. For these reasons, the behaviour at ignition is considerably different, and it is therefore important to study it experimentally to inform engine design, as extrapolating results from atmospheric conditions might yield inaccurate results [34].

For more widespread fuels such as methane and hydrogen, the high-pressure combustion landscape is comparatively mature, with extensive experimental databases and repeatedly validated kinetic mechanisms available across wide ranges of pressure and temperature. Methane has long served as a reference fuel in combustion science, and its pressure dependence in laminar burning velocity has been surveyed and correlated in dedicated studies, with multiple independent datasets available to support both empirical correlations and mechanism assessment [60, 34]. Similarly, hydrogen oxidation has been treated as a canonical system for combustion kinetics, and a substantial body of work exists specifically targeting high-pressure hydrogen ignition and flame behaviour, including mechanism updates and validation against shock-tube ignition delays and laminar flame speeds [61, 62, 63]. In addition, standard methane and hydrogen mechanisms are publicly documented and widely used as baselines in numerical studies. Consequently, while methane and hydrogen still require careful mechanism selection and validation for specific regimes, their high-pressure combustion properties are supported by a deeper, more consistent literature base than ammonia, for which the high-pressure database is still developing.

The high-pressure envelope is particularly critical for ammonia because its baseline reactivity is low and its combustion is strongly governed by radical-pool availability, which is precisely the part of the chemistry most sensitive to pressure through third-body and recombination pathways. At the same time, higher density reduces diffusivities and thermal transport, so kernel growth and flame development become more sensitive to heat losses and local inhomogeneities. For this reason, high-pressure studies are needed not only to quantify burning velocities but also to determine whether ignition can be achieved reliably under realistic constraints. In practice, this means the state of the art must be evaluated through three coupled metrics: ignition delay (autoignition propensity), minimum ignition energy (kernel survivability under forced ignition), and laminar flame propagation.

For premixed ammonia–air flames, the experimental literature has established a consistent qualitative outcome: laminar burning velocity decreases with increasing initial pressure, and ammonia remains markedly slower than typical hydrocarbon or hydrogen flames, as shown in Figure 2.15. A foundational dataset is provided by Hayakawa et al. [8], who measured unstretched laminar burning velocity and Markstein length for ammonia/air up to 0.5 MPa, showing a clear pressure-suppression trend. Importantly, that work also highlighted that several detailed mechanisms available at the time did not reproduce the measured pressure-dependent flame speeds with acceptable fidelity, underscoring the need to validate mechanisms against ammonia data (rather than nitrogen chemistry extrapolated from other fuels). More recently, Alvarez et al. [10] extended the known pressure range to 10 atm at 298 K, emphasising not only the nonlinear decrease of burning velocity and flame thickness with pressure but also the increasing impact of buoyancy on flame morphology as propagation becomes slow, especially near flammability limits. Together, these studies anchor the present consensus that high-pressure ammonia flames are slow, sensitive to transport/instability effects, and still challenging for mechanisms to reproduce across the full parameter space.

In terms of ignition, recent progress has been driven by rapid compression machine and shock-tube datasets at tens of bar, which have substantially improved the quantitative picture of ammonia autoignition in intermediate-to-high temperature regimes. These studies show that ignition delays are highly sensitive to mixture composition and, crucially, to hydrogen addition, which can reorganise the controlling chemistry and reduce ignition delay nonlinearly even at small fractions. This is a major step forward because it constrains the pressure-sensitive branching and termination balance in kinetic models more directly than atmospheric-pressure flames alone. However, these datasets primarily address autoignition propensity and, by themselves, do not resolve the practical question of how much external energy is required to form a self-sustaining kernel when ignition is forced rather than spontaneous. Consequently, while the literature on ignition delay at high pressure is expanding rapidly, it does not fully substitute for MIE evidence in ignition-system sizing and operability assessment [22, 64, 65].

This is where minimum ignition energy becomes a necessary complement to the high-pressure story, especially for ammonia-containing mixtures where ignition is often the limiting constraint. MIE captures the kernel-scale competition between energy deposition and quenching, which becomes more severe at elevated pressure because transport ahead of the kernel is slower and pressure-dependent termination pathways become more competitive. In ammonia, where the early radical pool is already weaker, these effects can translate into substantially higher ignition energies and a narrower practical ignition window than in more reactive fuels. Recent work has begun to quantify MIE for NH_3 and NH_3–H_2 mixtures and to show how strongly the threshold depends on both mixture composition and ambient pressure, findings directly relevant when the available ignition energy is limited. At the same time, MIE data remain far less abundant than burning-velocity and ignition-delay datasets, and are often reported under narrower condition sets or with ignition definitions that differ between studies, making cross-comparison more difficult [65, 66].

Hydrogen enrichment remains the most widely explored pathway to make high-pressure ammonia combustion feasible, because it simultaneously accelerates flame propagation and reduces the ignition-energy penalty. High-pressure flame-speed studies of NH_3–H_2 blends show consistent increases in burning velocity with hydrogen fraction while maintaining the typical decrease with increasing pressure and provide a practical mapping of the hydrogen fraction needed to recover robust early flame growth. At the ignition level, high-pressure autoignition datasets show that hydrogen shifts control toward the H_2/O_2 sub-chemistry, which can strongly shorten ignition delays, while MIE-oriented studies show that the same enrichment improves kernel survivability and reduces the required external energy. The combined implication is that blend design must be interpreted as a coupled ignition-propagation problem, rather than as a flame-speed problem alone. This is particularly relevant if hydrogen is generated by cracking, since the resulting mixture may contain substantial N_2 dilution, which affects both reactivity and transport [21, 22].

From a modelling perspective, the literature increasingly indicates that high-pressure predictability for ammonia is still not fully consolidated. Several studies report that mechanisms can match subsets of ignition-delay or flame-speed data but diverge as conditions move across pressure, temperature, equivalence ratio, and blend fraction, implying that the controlling pathways are not uniformly constrained. Mechanism comparisons at high pressure have therefore become a key research activity in their own right, and they show that ammonia validation cannot be treated as a minor extension of hydrocarbon validation. Importantly, this limitation affects not only flame-speed predictions but also kernel-scale behaviour, since ignition thresholds are inherently sensitive to the earliest radical-production and loss channels. As a result, the present state of research supports qualitative confidence in the trends (pressure slows propagation; hydrogen promotes ignition and burning), but it still leaves quantitative uncertainty that matters for design-level extrapolation [64].

Overall, the high-pressure ammonia literature is strongest in two areas: premixed laminar flame propagation up to the sub-MPa range in closed vessels and autoignition delay times at tens of bar in RCM and shock tubes. The evidence becomes thinner when the goal is to combine high pressure with additional realism drivers, such as subzero initial temperature, strong dilution, or hardware-limited ignition energies, all of which tighten the margin between kernel growth and quenching. In particular, pressure-resolved MIE datasets for NH_3 and NH_3–H_2 remain comparatively scarce and are not yet as standardised as ignition-delay reporting, even though they directly map to ignition-system requirements. This is one of the reasons why high-pressure ammonia combustion cannot yet be considered "closed" in the same way as methane or hydrogen, despite recent rapid progress. Consequently, the research frontier remains the coupled space of pressure, temperature, and ignition robustness, especially for ammonia and ammonia–hydrogen blends under conditions where both kinetics and transport penalise early flame development [66].

2.7. Low-temperature combustion

Subzero combustion already occurs in several engineering contexts, such as cold-start operation in transportation, especially in the marine sector [67, 68], and safety scenarios involving cryogenic storage and accidental releases. For this reason, it is of interest to study the chemical and kinetic behaviours that characterise low-temperature ignition and flame development, both to predict them and to improve device operability and safety margins. At reduced initial temperature, reaction rates generally slow and transport properties change, making ignition more sensitive to heat losses and mixing quality, while flame propagation typically weakens. Recent work has also highlighted that fundamental datasets at temperatures below ambient are still scarce, even for reference fuels, which limits the reliability of extrapolations of correlations or kinetic models beyond their validation range [69, 70]. This section, therefore, summarises the low-temperature picture in sufficient context, while maintaining focus on the carbon-free fuels central to this thesis.

2.7.1. Carbon fuels

For carbon fuels, low-temperature behaviour is often discussed in terms of low-temperature oxidation chemistry and the transition between weak and hot ignition. In practice, the detailed multi-regime description depends strongly on fuel class, mixture preparation, and ignition energy, and it is most relevant for engine-like conditions where two-stage ignition and cool-flame chemistry can emerge. In the present thesis, the most important point is that carbon fuels can exhibit intermediate-temperature pathways that are absent in hydrogen- or ammonia-based systems, because their oxidation proceeds through larger networks involving carbon-chain fragments and oxygenated intermediates. These pathways can lead to partial heat release prior to hot ignition and can create regions where increasing temperature does not monotonically accelerate overall reactivity (the so-called NTC behaviour) [71, 72]. However, since the objective is not to model carbon-fuel low-temperature autoignition in detail, the following discussion is intentionally kept at a higher level.

A practical limitation in subzero combustion of carbon fuels is that combustion occurs in the gas phase, while many carbon fuels are in the liquid phase at low temperatures. This means that, regardless of the detailed kinetic regime, the system must first ensure sufficient vaporisation and mixing to form a flammable gas-phase mixture. In low-temperature environments, this often requires additional thermal management, pressure reduction, or external energy coupling to support phase change and mixture preparation. Even when ignition energy is increased, this does not fully remove the vaporisation constraint, because a strong spark or plasma can only ignite the portion of the mixture that is already gaseous and within flammability limits. This is why low-temperature ignition studies of carbon fuels often need to treat mixture preparation as a first-order factor alongside chemical factors. In the context of this thesis, carbon fuels serve mainly as a reference point for the general effect of temperature on reactivity and flame propagation, rather than as the focus of the discussion of low-temperature combustion.

2.7.2. Carbonless fuels

For carbonless fuels such as hydrogen and ammonia, the carbon-fuel low-temperature framework is not directly applicable, as these fuels lack complex carbon chains and therefore do not exhibit the same cool-flame and two-stage ignition behaviour observed in low-temperature hydrocarbon oxidation. Instead, low-temperature limitations arise primarily from kinetically inhibited radical production and from transport-driven penalties that act most strongly during ignition and early flame growth. In particular, reducing the initial temperature slows the chemistry that sustains chain branching, while higher gas density and reduced diffusivities modify how heat and radicals are transported ahead of the flame. For this reason, even when the mixture remains flammable, the ignition delay increases and the burning velocity decreases as T_0 is reduced, which can push the system closer to quenching and instability thresholds. The recent subzero flame-speed literature makes this general trend explicit even for reference fuels, showing that the laminar burning velocity decreases substantially as temperatures drop from ambient to near-cryogenic levels [69, 70].

For hydrogen, low-temperature combustion is characterised by strong kinetic inhibition of the pathways that normally drive fast ignition and rapid premixed propagation. As T_0 decreases, the usual chain-branching reactions slightly mark, and the chemistry shifts toward alternative routes involving HO_2 and H_2O_2, which become comparatively more relevant below mixture temperatures of about 1000 K [61]. This produces a significant increase in ignition delay and generally raises the external ignition requirement, even though hydrogen remains highly reactive once a self-sustaining radical pool is established. Experiments at near-cryogenic temperatures show a strong reduction in the laminar burning velocity compared with ambient

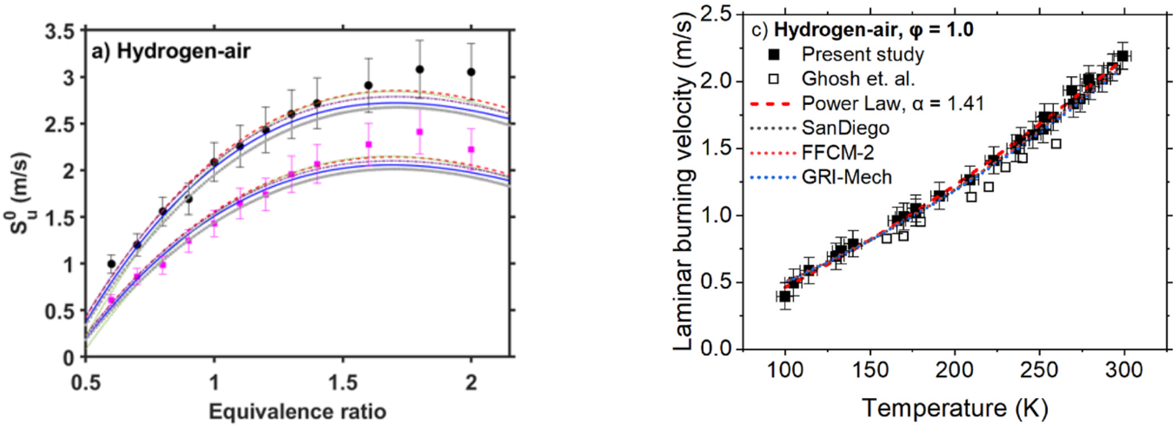

(a) Variation with ϕ at $T_0 = 240K$

(b) Variation with T_0 at $\phi = 1$

Figure 2.26: LBV of H_2 at subzero temperatures

conditions, and recent dedicated datasets extend this trend into the subzero range, reinforcing that temperature extrapolation requires care even for well-studied fuels [69, 70]. At the same time, the low-temperature increase in density and the associated hydrodynamic effects can amplify instabilities, meaning that cold hydrogen flames can display strong wrinkling and rapid acceleration once formed, even if the underlying chemistry is slower. This coupling between inhibited chemistry and enhanced instability is one of the reasons cryogenic hydrogen safety remains an active research area, especially when pressure is elevated as well [73, 74, 75].

For ammonia, the low-temperature literature is comparatively limited, and this lack of data is not only a matter of novelty but also of experimental practicality. First, ammonia already exhibits low reactivity and high ignition thresholds at ambient conditions, so lowering T_0 further penalises radical production and makes ignition more susceptible to losses, which pushes experiments toward demanding ignition hardware and narrow stability windows. Second, ammonia handling at low temperature introduces additional constraints related to condensation, material compatibility, and safety, which can discourage systematic campaigns compared with hydrogen or methane. Third, because ammonia is often proposed as an energy carrier for blends (especially with hydrogen), low-temperature behaviour is intertwined with composition effects, the possible presence of dissociation products, and dilution, so the parameter space expands rapidly. For these reasons, while low-temperature burning-velocity and ignition studies have recently advanced for hydrogen and methane down to cryogenic conditions, analogous fundamental datasets for ammonia are still sparse, and mechanism validation remains less mature in the subzero regime. [69, 20] Existing modelling approaches, such as extrapolations of MIE or ignition metrics from room-temperature data, remain promising but require further validation before they can be treated as reliable under subzero conditions. [45]

2.8. Research gap

Since ammonia is being studied as a potential carbon-free energy carrier for combustion-based power, a realistic assessment of its feasibility must include operation in cold environments and under elevated pressures. In these regimes, ignition and the first stage of flame development are often the limiting processes, because they control whether a stable kernel can form and whether early propagation can overcome heat losses and quenching. For ammonia, this is particularly important because its baseline reactivity is already low at ambient conditions, so any additional kinetic or transport penalty can push the system closer to non-ignition. Despite the growing modelling activity, the low-temperature behaviour of ammonia below approximately 21° C is still largely inferred from chemical-kinetic mechanisms rather than constrained by dedicated experiments. As a result, predictions in the subzero range are commonly obtained by extrapolating models that were primarily tuned and validated at higher temperatures, leaving uncertainty precisely in the part of the operating envelope where ignition robustness becomes most critical.

It is of particular interest to obtain experimental datasets extending to approximately -50° C, because this range represents a severe but still operationally realistic boundary for cold-region service without requiring cryogenic infrastructure in practical devices. In engineering environmental qualification practice, -51° C is widely used as a representative "severe cold" reference for real-world cold regions [67]. Similarly, maritime winterisation and ice-class requirements commonly consider low-temperature measures down to about -50° C, reinforcing that this temperature range is extremely relevant for marine and offshore systems rather than being an abstract laboratory choice [68]. Existing subzero flame-speed studies for hydrogen show that decreasing T_0 into this range produces a clear and measurable kinetic and transport penalty in early combustion metrics, motivating the expectation that ammonia will exhibit an equally pronounced reduction in burning velocity and ignition robustness Figure 2.26b. A dataset reaching this level is therefore valuable because it pushes the system far enough from ambient conditions that the trends are unambiguous and cannot be dismissed as marginal effects.

In parallel, the high-pressure regime remains essential because it is the regime where practical combustors operate and where ammonia's limiting chemistry becomes most restrictive. While methane and hydrogen benefit from mature high-pressure databases and repeatedly validated mechanisms, the high-pressure ammonia literature is thinner and more fragmented, particularly when additional realism drivers are introduced, such as limited ignition energy, dilution, or non-ambient initial temperatures. At elevated pressure, the competition between radical production and pressure-enhanced termination pathways becomes more severe, while reduced diffusivities increase quenching sensitivity during kernel growth, which directly impacts minimum ignition energy and early flame propagation. For this reason, predictions at pressures above roughly 5 bar to 10 bar remain uncertain for ammonia in many configurations, even when the qualitative pressure-suppression trend is understood. Therefore, it is important to extend the experimental database into the tens of bars, with a particular interest shown up to 50 bar, as mentioned in the approximate pressure in engine operations [76, 77].

Finally, the literature consistently shows that neat ammonia is rarely the most favourable option for practical combustion, which makes blending a central part of the feasibility discussion. Hydrogen enrichment is the most widely explored pathway because it can reduce ignition difficulty and increase burning velocity while remaining consistent with a low-carbon fuel strategy. However, the extent to which blending compensates for subzero and high-pressure penalties is not well quantified experimentally, and it is unclear whether trends observed at higher temperatures can be directly applied to cold conditions, where loss mechanisms and kernel survivability become dominant. This creates a coupled gap that is both fundamental and practical: the lack of validated low-temperature and high-pressure evidence for NH_3 and NH_3–H_2 mixtures limits confidence in kinetic mechanisms and prevents robust ignition-system and mixture-design decisions. Addressing this gap requires controlled experiments that quantify early combustion metrics under subzero and high-pressure conditions and that map how hydrogen addition shifts those limits.

3

Research objective and plan

3.1. Research objective

From the discussion just presented, a clear research gap appears in ammonia literature. In fact, the behaviour of this fuel alternative and of its blending with other fuels at low temperatures and high pressure remains largely unknown, with its understanding limited to non-validated numerical models. Therefore, the work of this thesis will be focused on filling this gap. As a starting point, a clear research objective needs to be stated, which states that:

The research objective is to characterise the ignition and initial flame propagation behaviour of ammonia and ammonia blends under subzero temperature conditions by developing and commissioning a combustion chamber suitable to withstand subzero temperatures and high pressures.

The considered research objective will be achieved by answering a number of research questions, which will lead to a unique answer to the proposed goal:

1. *How should a combustion chamber be designed to enable reliable ignition measurements at high pressure and subzero initial temperatures while ensuring measurement fidelity, thermal control, and safe integration of instrumentation?*
    - What geometrical features should the chamber include to withstand the maximum achievable pressure while minimising flame disturbance?
    - What chamber thermal design features are required to achieve a spatially uniform and stable subzero temperature initial condition at the ignition site?
2. *How do subzero temperatures affect the minimum ignition energy and laminar burning velocity of ammonia?*
    - What is the dependence of ammonia's ignition energy and flame speed on equivalence ratio at subzero conditions?
    - Do current chemical kinetic mechanisms accurately predict the laminar burning velocity of ammonia at subzero temperatures?
3. *How much does blending ammonia with hydrogen at subzero temperatures and high pressures affect combustion performance?*
    - What hydrogen fraction is sufficient to offset the subzero burning-velocity penalty for ammonia under the tested conditions?

- Does the behaviour of blended fuels at subzero temperatures relate to the one observed at higher temperatures?

3.2. Research plan

To achieve the objective of this work, a clear research plan must be developed, with a timeline shown in the Gantt chart in Figure 3.2.

3.2.1. Plan

After developing the necessary theoretical background through a literature review of the parameters of interest and the experimental conditions investigated in this work, the subsequent phases begin. The first step involves the full understanding and familiarisation with the current experimental setup in the Energy and Combustion Laboratory at the University of Massachusetts Lowell. In fact, to achieve the desired measurements, certain conditions must be met, stemming from the pressure and temperature requirements. From this initial assessment, it is evident that structural changes to the setup are required, primarily because the measurement fidelity of MIE and LBV depends on a controlled initial state, reliable delivery of ignition energy, and a chamber architecture that remains mechanically and thermally stable. The research plan is therefore structured around redesigning the core experimental hardware, quantifying measurement uncertainty, and running a staged experimental campaign that progressively de-risks the most challenging conditions.

Therefore, the bulk of the work involves designing a novel combustion chamber that meets the high-pressure, low-temperature requirements while remaining compatible with the existing laboratory infrastructure. This begins with selecting the chamber inserts that directly enable the measurements: pressure and temperature sensors to define the initial state and capture the pressure transient, an ignition system with adjustable delivered energy to allow MIE determination, and injection and exhaust ports that enable repeatable mixture preparation and safe purging. Then, the chamber structure can be designed based on the selected material and shape. For ease of integration in the current setup, the chamber design is designed to work in a similar fashion to the existing one. The selection of materials, sensors, and ignition systems is based on trade-offs that yield a design capable of safely withstanding the worst-case conditions.

In parallel, the surrounding experimental architecture is updated to support the new chamber, including the pressurisation approach, the cooling strategy for subzero operation, and the revised piping and valve layout. For these, the selection was straightforward, as only one option met the operating requirements while theoretically ensuring low lead times and costs. At the same time, the data-acquisition and post-processing workflow is prepared, so that pressure, temperature, and ignition electrical signals can be recorded consistently and reduced into MIE and LBV using established literature methods. This parallel development avoids a situation in which the hardware is available, but the measurement chain is not yet ready to produce usable results. This was achieved using an ad hoc LabVIEW environment, with outputs post-processed in MATLAB and the external software CVDART to calculate the MIE and the LBV, respectively.

Once the redesigned setup is assembled and commissioned, the experimental campaign is executed in progressive steps that serve both as data collection and as validation of the measurement approach. First, room-temperature tests are conducted with methane and hydrogen as reference fuels, because they provide strong benchmarks against established literature and enable early verification of ignition behaviour, pressure evolution, and parameter extraction. Ammonia is then included at room temperature to confirm that the chamber

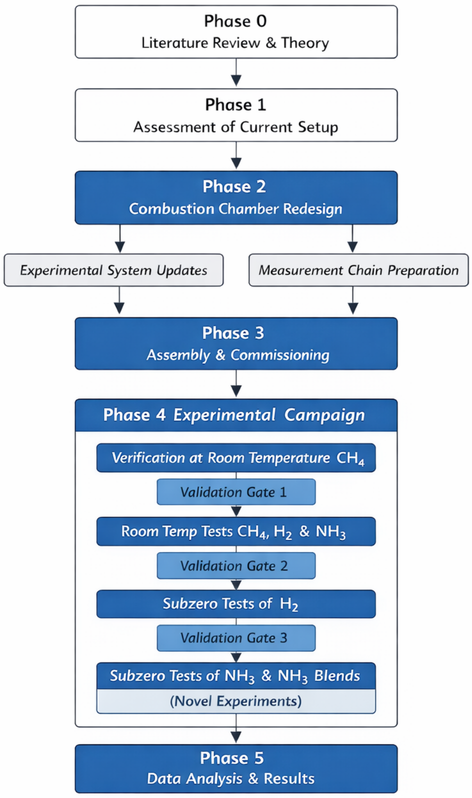

Figure 3.1: Thesis work procedure

and ignition system can support its more challenging ignition behaviour before introducing additional penalties. The second phase moves to subzero testing with hydrogen, specifically to validate the cold environment, confirm repeatable initial-state preparation, and verify that measured trends remain consistent with published low-temperature behaviour. Only after these intermediate validations are complete does the campaign move to ammonia and ammonia blends at subzero temperatures and elevated pressure, since these tests represent the combined worst-case operating point and are most likely to require iteration. This staged approach ensures that the final ammonia dataset is built on a validated measurement chain and that deviations can be interpreted with confidence.

The final phase consists of reducing the experimental traces into the parameters of interest and consolidating them into a coherent comparison across fuels, blends, and operating conditions. MIE and LBV are extracted using literature methods consistent with the chamber configuration and the pressure-based measurement approach, and the outputs are first checked against reference-fuel expectations to establish confidence in the reduction pipeline. The ammonia and ammonia-blend results are then interpreted in terms of how pressure and temperature shift ignition susceptibility and early flame development, and they are compared against the numerical models, as direct literature data are not available. Throughout the project, writing is carried out in parallel with the practical work so that methods and validation results are documented as they are refined. This structure ensures that the work remains focused on producing a defensible dataset under the target conditions while maintaining clear decision points that allow iteration without losing the overall coherence of the research plan. This plan is illustrated in Figure 3.1 as a block diagram.

3.2.2. Risk management

Risk management in this plan aims to preserve the ability to achieve the research objectives even if individual elements do not perform as expected in the first iteration. If the ignition system proves insufficient at the most penalising conditions, the mitigation is to retain a chamber interface that allows rapid replacement or adaptation of the igniter insert without redesigning the full vessel. If the cold environment does not achieve stable, repeatable subzero initial temperatures, the mitigation is to adjust insulation, sealing, and cooling implementation, and to validate performance with hydrogen before returning to ammonia. If measurement quality is limited by sensor noise, drift, or bandwidth, the mitigation is to rely on the reference-fuel phases to diagnose the issue early and adjust instrumentation or signal processing before the ammonia campaign. Finally, the optical chamber is not treated as a primary fallback for the target campaign, because it is rated only to 30 bar, and the presence of windows introduces additional uncertainty under thermal cycling due to differential thermal expansion between the window material and the metal housing. For this reason, it can be used only as a limited diagnostic tool at reduced pressure, while the main plan remains centred on achieving the required conditions in the redesigned high-pressure chamber.

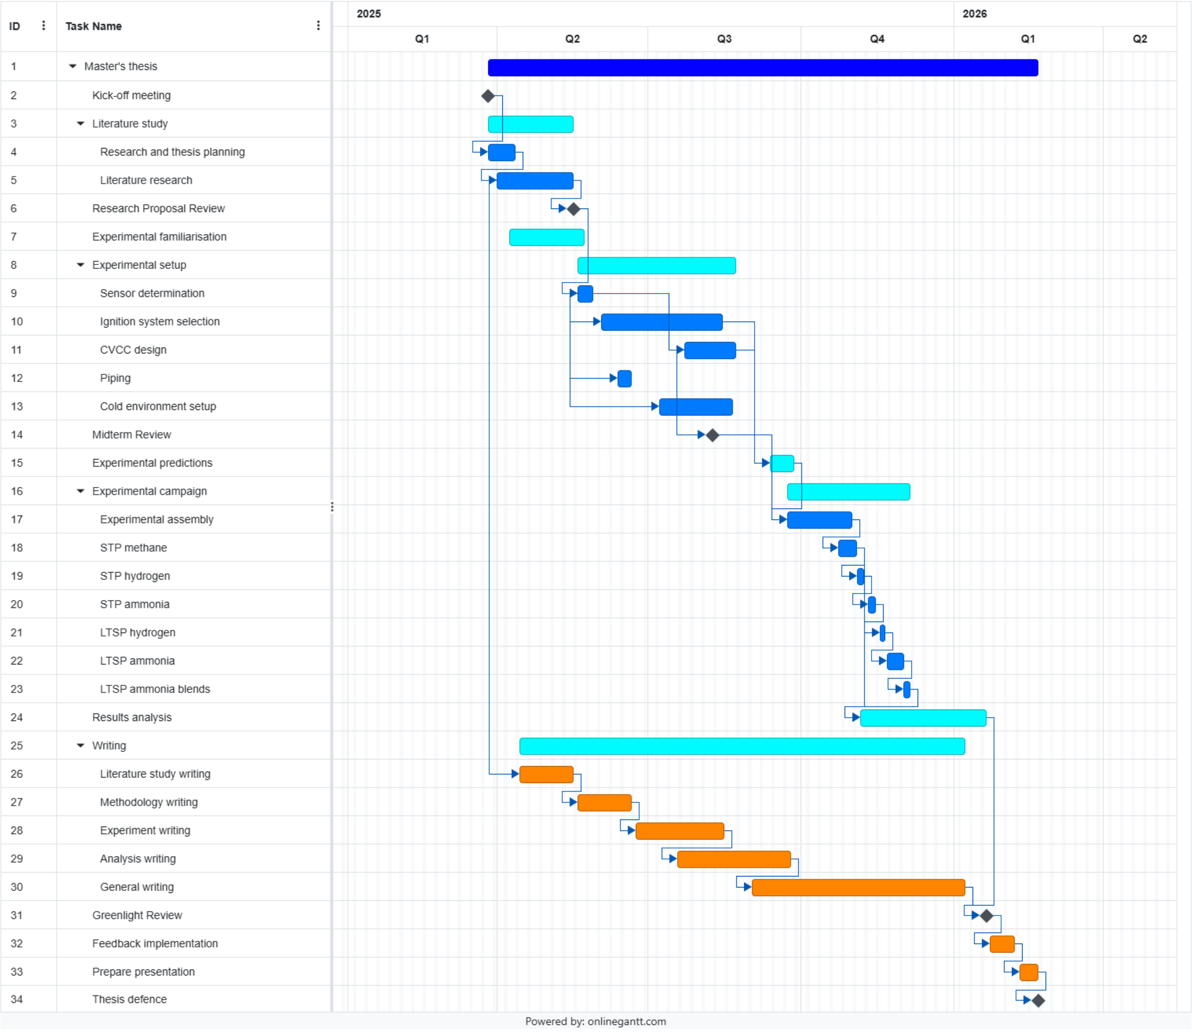

ID	Task Name
1	Master's thesis
2	Kick-off meeting
3	Literature study
4	Research and thesis planning
5	Literature research
6	Research Proposal Review
7	Experimental familiarisation
8	Experimental setup
9	Sensor determination
10	Ignition system selection
11	CVCC design
12	Piping
13	Cold environment setup
14	Midterm Review
15	Experimental predictions
16	Experimental campaign
17	Experimental assembly
18	STP methane
19	STP hydrogen
20	STP ammonia
21	LTSP hydrogen
22	LTSP ammonia
23	LTSP ammonia blends
24	Results analysis
25	Writing
26	Literature study writing
27	Methodology writing
28	Experiment writing
29	Analysis writing
30	General writing
31	Greenlight Review
32	Feedback implementation
33	Prepare presentation
34	Thesis defence

Figure 3.2: Gantt chart for this work

3.3. Proposed methodology

This section defines the methodological framework for investigating the ignition and early flame development of ammonia and ammonia-based blends. The work focuses on two complementary metrics that characterise the initial phase of combustion: minimum ignition energy (MIE), which quantifies the threshold for forming a self-sustaining flame kernel, and laminar burning velocity (LBV), which describes the subsequent kernel growth into a premixed flame. Together, these metrics provide a practical and fuel-comparable description of ignition susceptibility and early propagation, which can then be mapped across operating conditions to identify limiting regimes. This is particularly relevant for ammonia, where low baseline reactivity makes early flame development sensitive to losses and to small changes in initial state. The methodological novelty lies in applying these measurements to ammonia and ammonia-hydrogen blends at combined subzero temperature and elevated pressure, a regime for which the experimental evidence base remains limited.

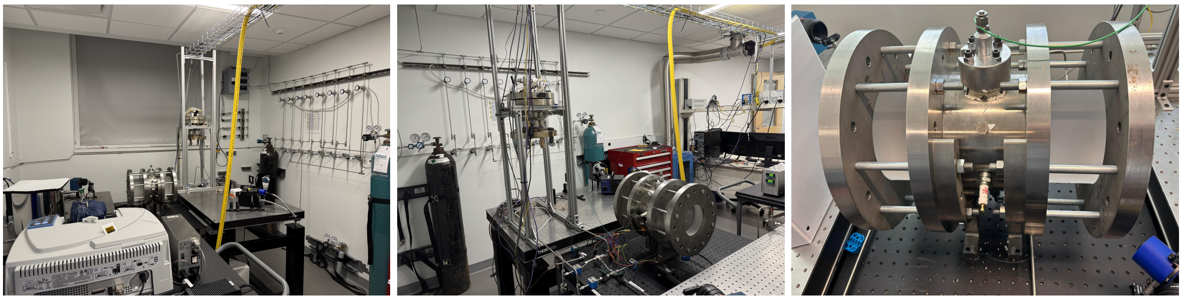

(a) Front view **(b)** Rear view **(c)** Optical chamber close-up

(d) Optical chamber window **(e)** Filling line **(f)** Injection/exhaust line

(g) FTIR (Fourier Transform Infrared Spectroscopy) machine

Figure 3.3: Initial test setup

3.3.1. Experimental concept and measurement approach

The experiments are designed around a constant-volume combustion chamber (CVCC) operated under premixed, quiescent initial conditions. A controllable-level ignition is initiated at a localised point central to the chamber, so that the early flame can be treated as an isotropic expanding spherical kernel during its initial growth. This canonical configuration is widely used because it provides well-defined boundary conditions and enables comparison between different fuels under controlled initial pressure and temperature. Because the target thermodynamic envelope requires a robust pressure boundary and repeated thermal cycling, a non-optical chamber is considered and used, meaning the combustion event is inferred from time-synchronised measured signals rather than flame imaging. Consequently, the methodology relies on pressure, temperature, and electrical ignition traces to quantify the two target metrics: MIE is determined from the delivered ignition energy and the corresponding ignition outcome, while LBV is derived from the chamber's thermodynamic response during the early and intermediate stages of combustion.

A key advantage of this approach is that it remains compatible with the operating envelope of interest, where optical access is limited by pressure rating and thermomechanical constraints. At the same time, it imposes clear requirements on instrumentation and signal quality, since the pressure trace becomes the primary observable for combustion characterisation. For this reason, it is assumed that a portion of the pressure trace always captures the behaviour of the laminar flame expansion, and that optical access is not required to determine the limits of this zone. The laminar interval is therefore treated as a trace-selection problem rather than an imaging problem, and its determination is addressed later through a consistent bounding procedure. The detailed acquisition chain, signal conditioning, and data-reduction workflow that implement this approach are presented later in Chapter 5, after the final hardware configuration has been defined.

3.3.2. Experimental requirements and test envelope

The experimental envelope is defined to quantify how ignition and early propagation vary with initial pressure, initial temperature, and fuel composition. The regime of interest extends to initial pressures up to 50 bar and initial temperatures down to approximately −50 °C, because these conditions capture a challenging, engine-relevant direction where both kinetics and transport penalise early flame development. To avoid conflating effects, pressure, temperature, equivalence ratio, and blend fraction are treated as distinct variables, and the campaign is structured so that trends can be interpreted without relying on a single-point comparison. Fuel selection follows the same logic, using well-characterised reference fuels to establish method credibility before drawing conclusions for ammonia. Methane is included primarily as a baseline for verification and for anchoring the facility against established literature trends under standard conditions. Hydrogen is included both as a subzero validation fuel with abundant comparison data and as a relevant blending partner for ammonia within a low-carbon framework.

Ammonia is then investigated as the target fuel, both neat and in ammonia–hydrogen blends, to quantify how composition can mitigate ignition difficulty and slow early propagation under penalising conditions. Blend definition is expressed as a hydrogen fraction in the fuel mixture, so it maps directly to partial-pressure mixture preparation and remains consistent across operating points. The equivalence ratio is explored independently of the thermodynamic sweeps, because mixture-strength effects can otherwise obscure the interpretation of temperature and pressure dependencies. This structure ensures that the dataset addresses both fundamental questions, such as how T_0 shifts ignition and burning velocity, and practical questions, such as how much hydrogen is required to achieve robust early combustion. The complete experimental matrix and the procedure for populating it are presented later, once the design and final hardware configuration are established.

3.3.3. Baseline setup and motivation for redesign

The University of Massachusetts Lowell has an existing experimental foundation for ammonia combustion studies through a test setup previously developed at the Energy and Combustion Research Laboratory. In Figure 3.3, the already existing experimental setup is shown, comprising its main components, namely the optical combustion chamber, the filling line, the injection and exhaust lines, and the FTIR machine for exhaust-gas analysis. Starting at the top of the experimental process, the filling line consists of several wall outlets that lead to the respective gas bottles, stored in a separate, enclosed area of the laboratory. These outlets ensure direct access to all necessary reactants, including H_2, CH_4, NH_3, O_2, and N_2, which are supplied via high-sensitivity check valves to regulate the amount of gas reaching the chamber. The limit on the delivery pressure is set by the regulators, which can release up to 200 psi to prevent the cylinders from depleting too quickly and ensure safe operation.

The filling process is controlled by piezoresistive pressure transducers from the Omega PX409 family, which are set in parallel with the line leading to the chamber. Multiple sensors with different pressure ratings are used to ensure accurate filling across the required range. To ensure a consistent baseline for mixture preparation, a vacuum pump (Welch Duoseal 1402B-01) is used to evacuate the chamber and the line prior to each experiment. As shown in Figure 3.3f, the setup features a single line for both injection and exhaust gas, meaning that the chamber is connected through a common port and the flow direction is controlled externally. The chamber closure is therefore achieved through a valve placed on the line, which is opened and closed as required during injection and exhaust. This arrangement is effective for routine testing at moderate conditions and provides a simple architecture for gas handling and evacuation.

In its original configuration, the setup provides optical access, which is beneficial for qualitative flame visualisation and for calibration under moderate conditions. However, the baseline chamber is rated for approximately 30 bar of pressure, and the windowed design introduces additional constraints during thermal cycling due to differential thermal expansion between the window material and the metal housing. Moreover, the original setup was designed for room-temperature operation and does not provide a controlled method for imposing repeatable subzero initial temperatures. For these reasons, achieving the target envelope requires redesigning the chamber and partially redesigning the surrounding architecture. In this thesis, the redesign serves both as a means to obtain experimental results and as an objective in itself, namely, presenting the necessary setup for ammonia combustion under these conditions with adequate repeatability and instrumentation compatible with non-optical data reduction. The details of the final chamber design, supporting systems, and integration are presented in Chapter 4.

3.3.4. Design methodology overview

The design approach follows a requirement-driven process that begins from the target operating envelope and the measurement needs for MIE and LBV. This includes defining the required instrumentation and interfaces (pressure and temperature sensing, ignition port, injection/exhaust ports), since these dictate the chamber layout and the measurable outputs. The chamber geometry and material are then selected to withstand repeated high-pressure operation and thermal cycling while maintaining a stable internal volume and reliable sealing. In parallel, two enabling subsystems are specified: a pressurisation solution capable of reaching the target initial pressure with sufficient control and a cooling strategy capable of achieving the target subzero temperatures with acceptable repeatability. The final architecture is then integrated with the existing laboratory infrastructure so that mixture preparation, evacuation, and data acquisition remain consistent across the full experimental matrix.

3.3.5. Verification and validation strategy

To ensure that the experimental results can be interpreted as fuel behaviour rather than facility artefacts, the overall programme is executed through a staged verification and validation (V&V) strategy before progressing to the target ammonia conditions. Verification establishes that the integrated setup operates as intended, including acceptable leak rates, repeatable evacuation and filling, correct sensor response, and reliable triggering and recording of the relevant signals. Validation then benchmarks the measurement and reduction chain against established reference behaviour, using fuels and conditions for which trends and representative values are well documented in the literature. This step is particularly important for a non-optical constant-volume approach, as it provides confidence that the selected pressure-trace intervals and the reduction workflow yield physically consistent burning velocities across different fuels. Only after this staged V&V is completed are the novel subzero ammonia and ammonia-hydrogen measurements interpreted as discovery data, since the prior stages provide the baseline needed to attribute deviations to the fuel, the operating condition, or the blend composition rather than to unresolved methodological limitations.

4

Test setup design

As discussed in Section 3.3, to achieve the required pressure and temperature conditions, an ad hoc setup is needed. In this chapter, the design process will be thoroughly explained, detailing all choices and trade-offs made and the changes to the existing setup, and demonstrating the feasibility of the design in meeting the requirements. The process was divided into three main parts. The first step involved designing the combustion chamber, as explained in Section 4.1, encompassing all geometrical aspects; then the ignition system and sensors required to achieve combustion and measure it are described in Section 4.2. Finally, using the existing setup as a starting point, the required component selection for pressurisation and temperature conditioning is described in Section 4.3.

4.1. Combustion chamber design

The design of a combustion chamber is a complex project that requires meeting multiple requirements simultaneously. Indeed, aside from the primary requirements of this work, which aim to achieve novel measurements, the chamber must also comply with the requirements characteristic of any combustion chamber. Among these, the primary objectives are to provide easy access for cleaning, ensure correct sensor placement, and have dimensions that are limited yet allow for both mobility and reduced turbulent flame behaviour. Most importantly, the chamber needs to be compatible with all the gases used, especially those such as ammonia and hydrogen, which are highly reactive with materials. Indeed, the first is highly corrosive and can lead to the complete obliteration of a surface, while the latter, over time, can cause embrittlement, resulting in micro-damages that ultimately lead to failure.

While the temperature requirement needed for the measurements is self-explanatory, with -50° C being the metal design minimum temperature (MDMT), the pressure requirement requires a bit more elaboration. In fact, it is of interest to study initial pressures up to 50 bar; in a constant-pressure chamber, this would thus be the maximum pressure for which to design the chamber. However, since a constant volume one will be used in this case, the chamber pressure will increase during combustion. Due to the scarcity of sources, including data at such high pressures, an estimate of the maximum pressure the chamber might be exposed to is achieved with the Cantera[1] software, an open-source chemical kinetics package that couples detailed reaction mechanisms with thermodynamic and transport property models to solve reacting flow problems. This calculates the adiabatic temperature by integrating the conservation equations with constitutive models provided by the selected mechanism and

---

[1] https://cantera.org/ [Accessed on 12/05/2025]

property databases, and, from it, the adiabatic pressure, ensuring a conservative estimate. In real systems, losses are always present, and thus these quantities will naturally be lower; this explains why taking this estimate as a starting point is a good approach. To ensure an adequate level of safety, however, an additional 10% safety factor is applied to the maximum adiabatic pressure.

During this work, three fuels will be used: methane, hydrogen, and ammonia. To estimate the worst-case maximum achievable pressure in the chamber, the most constraining fuel must be considered. In this case, H_2 is chosen because it has the highest reactivity and the fastest reaction rate, resulting in the largest pressure rise. Then, assuming the worst-case condition of a pure hydrogen mixture at an equivalence ratio of 1, with an initial pressure of 50 bar and an initial temperature of -50° C, the estimated adiabatic pressure with the GRI-Mech 3.0 model[2] is 707 bar. The method is known to overestimate hydrogen pressure and temperature, as it is designed for carbon fuels and overlooks some fundamental reactions that occur only in hydrogen combustion. Using, for instance, the "H_2O_2" mechanism, this maximum pressure is simulated at only 500 bar. However, since the former is the common baseline used across all studies and provides a more conservative estimate, it is chosen as the starting point, thereby allowing greater flexibility in the initial pressure as well. Considering the 10% safety factor, thus, the chamber will be designed for a maximum allowable working pressure (MAWP) of 777 bar.

All of these requirements can be grouped into the following list of requirements:

Table 4.1: List of requirements (T = Test, I = Inspection, A = Analysis, D = Demonstration)

Identifier	Requirement	V&V
REQ-01	The chamber shall withstand a pressure of at least 777 bar (MAWP)	T, A
REQ-02	The chamber shall withstand a temperature of at least -50° C (MDMT)	T
REQ-03	The chamber shall withstand a minimum of 100000 cycles	I, A
REQ-04	The chamber shall guarantee the possibility for internal access	A, D
REQ-05	The chamber should be compatible with NH_3 and H_2	D
REQ-06	The chamber shall ensure enough spacing for at least one pressure transducer with rating $\geq$ MAWP	A
REQ-07	The chamber shall ensure enough spacing for at least one temperature sensor with rating $\leq$ MDMT	A
REQ-08	The chamber shall ensure enough spacing for two igniters for symmetrical ignition	A
REQ-09	The chamber should avoid sharp transitions and corners	A
REQ-10	The chamber should be able to pull a vacuum of at least 200 mTorr	D
REQ-11	The chamber should be small enough to limit the onset of flame turbulence	A, T

4.1.1. Material selection

The first step in defining the combustion chamber design is selecting the material, which must meet all the requirements listed above. The material selection is thus driven by a series of criteria that will serve as the deciding factors in a material trade-off. However, since not all the criteria have the same importance, meaning, for instance, that a material with a shorter lead time will be prioritised over a cheaper one, weights are assigned to the criteria, which are set at the top of Table 4.3. The list of criteria considered, along with their importance, is given here:

---

[2] http://combustion.berkeley.edu/gri-mech/version30/text30.html [Accessed on 12/05/2025]

- **Strength and stiffness at design conditions (high):** adequate yield/tensile strength at 50 bar and down to −50 °C, with sufficient modulus to limit deformation and maintain sealing at minimal wall thickness. This enables a compact chamber; low-strength options drive thickness, mass, cost, and complicate port/seal integration.
- **Low-temperature toughness and ductility (high):** high impact toughness and low DBTT to prevent brittle fracture during subzero operation and thermal cycling. This is essential to the subzero objective; insulation/coatings may slow cooldown but cannot eliminate brittle-failure risk at stress concentrators, and internal coatings are not robust under ignition events.
- **Compatibility with ammonia environments (high):** resistance to NH_3 and combustion products, including SCC and ammonia-assisted pitting/crevice corrosion. Corrosion can thin walls and introduce flaws, reducing fatigue life and fracture margin, especially at threads and sealing interfaces.
- **Resistance to H_2 damage (high):** low susceptibility to hydrogen embrittlement and hydrogen-assisted cracking. Since hydrogen effects can be time- and microstructure-dependent, long-term integrity under repeated exposure and cycling requires demonstrated resistance.
- **Availability and procurement (high):** readily available in required stock forms (bar/forging) with lead times compatible with the schedule. Availability is treated as a feasibility constraint to avoid blocking iterations and testing completion.
- **Fatigue and thermal fatigue (medium):** adequate endurance and thermal cycling resistance under repeated pressurisation, ignition shocks, and gradients, with notch tolerance at ports/threads. Cyclic loading remains secondary to static strength and toughness but drives local design and surface finish.
- **Manufacturability and joinability (medium):** machinability for ports and sealing surfaces, compatibility with the sealing approach, and practical fabrication routes for the geometry. Used as a discriminator between comparable candidates due to its impact on the schedule and iteration speed.
- **Cost-effectiveness (medium):** acceptable total component cost (material, machining, tooling, post-processing) relative to achieved margin. Not primary, but relevant when performance is comparable.
- **High-temperature wall capability (low):** sufficient property retention and oxidation resistance at transient inner-wall temperatures. Combustion is millisecond-scale, so heating is local and brief; the goal is to avoid obvious incompatibilities and cumulative surface degradation over repeated runs.

Among the numerous materials available, the choice is greatly narrowed by the presence of ammonia, which is the most constraining criterion, as only a very limited number of materials can withstand it. In fact, for example, any type of aluminium cannot be considered for this reason. Therefore, the choice is reduced to the materials provided in Table 4.2.

These materials are then compared and selected based on the criteria just explained. This is shown in Table 4.3, where the scores for the weight range from 1 to 3, with 3 indicating the most important and 1 the least. On the other hand, the materials are scored on a scale of 1 to 5. To calculate the final score and select the best material, therefore, a weighted sum is performed. A more detailed explanation of the scores given in the table is presented in Appendix A.

Table 4.2: Possible material choices

Material	σ_y	Characteristics
AISI 316L SS	~170 MPa	Excellent subzero toughness; good NH_3 corrosion resistance; mitigates H_2 damage.
Duplex SS 2205	~450 MPa	Higher strength; good weldability; strong against ammonia and H embrittlement.
Super-Duplex SS 2507	~620 MPa	Even higher yield; excellent corrosion resistance in harsh environments.
Inconel 625	~380 MPa	Ni-Cr-Mo alloy; outstanding NH_3 and H_2 embrittlement resistance; fatigue resistant.
Inconel 718	~1030 MPa	Ultra-high strength; good corrosion resistance; requires careful heat treatment.
Hastelloy C-276	~275 MPa	Excellent chemical resistance; moderate strength; good low-temp toughness.
Monel 400	~250 MPa	Ni–Cu alloy; good for H_2 service; moderate strength; easy to fabricate.

Table 4.3: Material trade-off (scores: 1 = worst, 5 = best)

Material	**Strength & stiffness**	**Cryo capability**	**High-T capability**	**NH_3 compat.**	**Mfg. ease**	**H_2 compat.**	**Fatigue resistance**	**Cost**	**Avail.**	**Tot.**
Weights	3	3	1	3	2	3	2	2	3	
316L SS	1	5	2	4	3	4	3	3	4	74
SS 2205	4	3	1	4	4	3	4	4	4	79
SS 2507	4	3	1	4	3	3	4	3	3	72
Inconel 625	3	5	5	5	1	5	5	1	1	76
Inconel 718	5	5	4	4	1	4	5	1	1	75
Hast. C-276	2	4	3	5	3	5	3	1	2	71
Monel 400	2	5	2	2	4	4	3	2	3	68

Therefore, although all materials perform quite well, the chosen one is the Duplex SS2205, for which 4 companies are identified as possible sellers, as they have it in stock.

4.1.2. Geometry definition

For the geometry definition, the requirements outlined in Table 4.1 must be closely considered. This phase consists of three main parts: first, the chamber's geometric shape is selected; then, the constraints on its dimensions are set; and finally, given the selected material, the thickness is defined based on its properties.

4.1.2.1. Chamber shape

Looking at the literature, many possible shapes are widely used, among which:

- Toroidal
- Trapezoidal
- Spherical
- Cylindrical

The differences in shape are due to the design's dependence on the working environment; when turbulence is considered, the geometry plays a significant role in how the flame moves, namely in wrinkling and turbulent velocity. However, for a quiescent environment, that is, a situation where there is no or minimal external flow disturbance affecting the combustion process, the geometrical shape plays a negligible role in determining the MIE and the LBV.

Since turbulence is out of the scope of this work, spherical and cylindrical designs are more interesting, as they hinder the expansion the least and therefore leave a larger area for laminar flame propagation. Moreover, their design and stability at higher pressures are far superior to those of the more complex toroidal and trapezoidal designs with more edges. Among these two, a third concept exists: the spherical approximation by superposition of various cylinders, which can unite the benefits of both geometries. Therefore, a trade-off will be performed among these three shapes, the benefits of which will be summarised hereafter.

Spherical chambers:

+ Minimise the surface-to-volume ratio, reducing heat transfer to walls by up to 6 times as compared to a trapezoidal
+ Uniform pressure distribution
+ Able to easily handle high pressures due to its uniform distribution
+ Theoretically, the most efficient
+ Natural recirculation zones due to radial symmetry, promoting homogeneous mixing
- Manufacturing of hemispherical components is very difficult and often leads to imprecision, causing cost and complexity to go up
- More difficult to seal perfectly due to the manufacturing challenges
- Complex injector configurations to achieve mixing, potentially introducing instabilities
- Difficult to mount components at precise angles and ensure sealing
- Higher stress concentration, with greater difficulty of reinforcement due to the shape
- Especially at high pressures, they lose most of their advantage with the addition of holes

Cylindrical chambers:

+ Simpler to seal and manufacture
+ They promote predictable mixing over a distance of 5 times the chamber diameter, which makes the injector design much easier. This is because axial flow is dominant, and there are turbulent boundary layers at the walls
+ Easy to insert the injector, electrodes and probes due to the almost flat surfaces
+ Most available components are designed specifically for cylindrical chambers
+ Can be designed modularly, improving accessibility for inspection
- Higher surface-to-volume ratio, increasing the heat loss

Spherical approximation chambers: are produced by arranging multiple cylindrical tubes to form a polyhedral structure resembling a sphere.

- \+ Ease of manufacturing of cylinders while approximating a sphere
- \- Heat loss is higher compared to the sphere due to the increase in surface area
- \- Joints create localised stress points, requiring reinforced welding or gasket systems to handle high pressures
- \- Abrupt transitions between segments may disrupt spherical wave propagation
- \- Joints can reflect or scatter pressure waves, causing instability

From this discussion, it is clear that, given all requirements, the cylindrical geometry offers the greatest advantages. Within it, a further classification, depending on the position of the cylindrical axis with respect to the injection one:

1. Vertical chambers: Injection axis same as the cylinder axis
    - \+ Geometry closer to an actual engine, closer to real conditions
    - \+ Stratified filling, meaning that air and fuel can layer axially, promoting ignition
    - \+ Laminar filling, indicating a lower chance of turbulent disturbance
    - \- Less area for optical access
    - \- Poor radial mixing in the absence of swirlers
    - \- Longer filling time due to the flow having to take a long route
    - \- Axial stratification could lead to uneven combustion or ignition delays
2. Horizontal chambers: Injection axis perpendicular to cylinder axis
    - \+ Easy to design, construct, and maintain
    - \+ May allow for higher compression ratios
    - \+ Better radial mixing due to the creation of shear layers and vortices, leading to a better homogenization of gases
    - \+ Faster chamber pressurization
    - \+ Better control of local equivalence ratio
    - \- Wall impingement when the jet hits the opposite wall too directly, causing thermal or flow issues (wall erosion, thermal fatigue due to the creation of pockets)

Given the significant advantages, the horizontal cylindrical chamber is chosen.

4.1.2.2. Chamber dimensions

Diameter and length estimation

Regarding the determination of chamber size, most volumes reported in the literature are approximately 2 L [36, 28, 78]. The reason for this similarity and number comes from the superposition of two main criteria, that is:

- The chamber shall be large enough so that the flame remains in its spherical laminar flame regime

  The free propagation length should be large enough that the flame remains essentially spherical up to its maximum flame radius. If the flame radius is less than 30% of the chamber wall radius, confinement-induced flow alters the measured burning velocity by less than 3%. If, instead, 50% is taken, then uncorrected flames show errors up to ≈15%. Thus, with a value of $R_{f,max}$ = 30 mm for a hydrogen flame, that would mean that, to have minimal error, the minimum radius would need to be set to at least R = 90 mm. However, that would entail a very large chamber, which would contradict the second criterion. Therefore, the 0.5 limit is applied, accounting for the need for post-processing corrections (only for hydrogen), a widely used practice, which sets the lower limit at D = 120 mm.

- The chamber shall be small enough so that the heat loss does not considerably affect the flame

  The chamber diameter sets the timescale over which wall heat loss distorts the pressure rise. The pressure rise from ignition happens in ≈5–40 ms; after that, heat loss begins to soften the post-peak trace. Thermal-diffusion time scales roughly with D^2/α, where α is the diffusion coefficient. To maintain the "adiabatic window" at a duration of 50 ms or less, diameters significantly exceeding ≈150 mm begin to extend that period. Therefore, the upper limit is set to D = 150 mm [79].

Hence,

$$120 \leq D \leq 150\text{mm}$$

A value of 130 mm is chosen to allow for some margin on both ends, with a closer proximity to the lower end, as corrections are available. With the diameter set, the cylinder's height must be defined. Ideally, the best configuration is obtained with a 1:1 aspect ratio, for three main reasons:

- Minimises flame-boundary interactions: as a spherical flame propagates in a confined cylinder, the walls induce flow that distorts the flame and biases the measured propagation speed. Burke, Ju, and Dryer [80] showed that this distortion and the resulting error in inferred laminar burning velocity increase with chamber aspect ratio (length/diameter). A 1:1 aspect ratio maintains nearly isotropic wall effects and extends the valid range of flame-radius measurements before quenching or distortion occurs.
- Extends the undisturbed spherical flame regime: in a spherical chamber, the flame remains spherical until very close to the wall, where it is quenched. A cylinder with an aspect ratio of 1 approximates this behaviour better than a long or squat cylinder. This means that the unburned gas velocity remains near zero over a larger radial span, thereby reducing the need for complex flow correction factors.
- Symmetry for sensor integration: equal length and diameter provide a highly symmetric volume, simplifying the positioning of transducers so that each is exposed to the same acoustic and flow environment. This helps ensure that every component experiences a symmetrical pressure wave and that the bulk measurement is acceptable [79, 80].

The aspect ratio, however, must accommodate all other requirements, so a 1:1 ratio can be slightly adjusted to allow for better sensor integration if there is not enough space. Thus, initially, a height of 130 mm is selected, but this is iterated upon as the design progresses.

Cylinder thickness

Thus, once the diameter and height are known and the chosen material properties are specified, the cylinder's thickness can be determined. A distinction needs to be made between the cylindrical part and the ends. For the former, the process is straightforward, as the stress, and

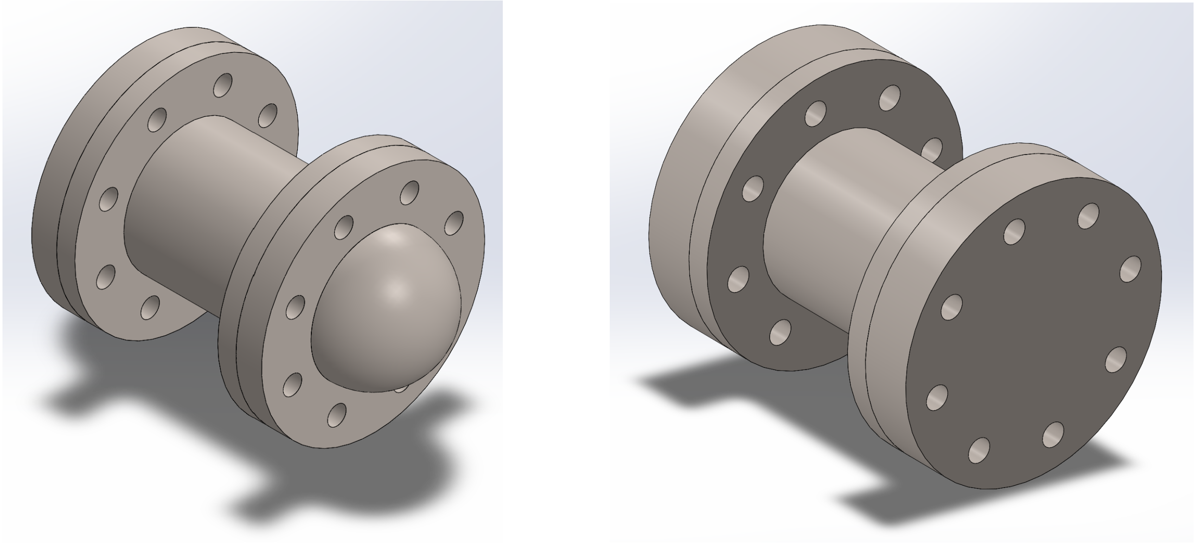

(a) Hemispherical-ended **(b)** Flat-ended

Figure 4.1: CAD sketches of considered chamber flange configurations

thus the thickness, over the radial direction needs to be considered, being the most constraining one; for the latter, the thickness depends on their shape. The ideal shape would be elliptical or semi-spherical, allowing the chamber to simulate the flame's spherical expansion as closely as possible. However, since physical access to the chamber is necessary to clean and position the igniters accurately, at least one side of the chamber must be removable. Therefore, such a design becomes very difficult to achieve, especially from a manufacturing perspective, as it would require a thick elliptical shape with a large flange at the bottom.

Thus, initially, it was considered to have one elliptical side and one removable flat side; however, due to symmetry and ease of manufacturing, this idea was quickly discarded in favour of a double flat-ended chamber with removable bolted flanges. A CAD sketch of these two configurations is shown in Figure 4.1 as a reference. As mentioned, these flanges will require a different thickness than the cylindrical part due to their flat nature, which handles pressure less effectively than a hoop structure. Since they need to be removable, these flanges will consist of two parts: one welded to the cylindrical chamber itself and the other removable. Thus, the bolt placement and number will also be important for effectively handling the pressure. The dimension calculation for the full chamber is explained hereafter.

Starting from the inner cylinder diameter set before, the outer one can be found by calculating the hoop thickness:

$$t_{cyl} = \frac{\Delta p D_i}{\sigma_y} \tag{4.1}$$

$$D_o = D_i + 2t_{cyl} \tag{4.2}$$

where

t_{cyl}	Chamber cylinder thickness
Δp	Maximum pressure differential [Pa]
σ_y	Material yield strength [MPa]
D_i	Inner chamber diameter
D_o	Outer chamber diameter

It is worth noting that throughout the entire design, a safety factor of 2 is applied to the yield strength; this means that it is assumed the material is only half as strong, reducing the strength of the SS 2205 to 225 MPa, which ensures enough safety, especially close to the insert holes.

Flange sealing

To ensure proper sealing between the two flanges, a ring or gasket must be placed, which, when compressed by the bolts, deforms to fill any voids between the flanges and thereby creates a perfect seal for any temperature and pressure within its rating. In this case, given the pressure rating and the temperature, the classical elastomer O-rings, also used in the optical chamber, cannot be used, as they will become brittle under low temperatures and cannot withstand high pressure. Therefore, the only feasible options are three types of metal gaskets: the R-ring, RX-ring and BX-ring.

- Style R (R-ring, oval or octagonal)
  - *Geometry:* Solid metal ring with either an oval or octagonal cross-section.
  - *Groove compatibility:* Matches R-type grooves; the oval style is the one used for older round-bottom grooves.
  - *Sealing mechanism:* Bolt-load (initial seating) driven.
  - *Typical use:* General service where pressure is not above 5000 psi.
- Style RX (RX-ring)
  - *Geometry:* Modified octagonal cross-section derived from Style R, with a higher profile.
  - *Groove compatibility:* Designed to fit the same groove design as Style R.
  - *Sealing mechanism:* Pressure-energised (self-energising); internal pressure increases sealing efficiency (improves sealing as pressure rises).
  - *Typical use:* More demanding services than basic Style R, but still ideally below 5000 psi.
- Style BX (BX-ring)
  - *Geometry:* Pressure-energised "BX" profile, designed to be fully confined in the groove and to allow virtual face-to-face flange contact when correctly fitted.
  - *Groove compatibility:* Only for BX flanges/grooves; not interchangeable with R/RX grooves.
  - *Sealing mechanism:* Pressure-energised and includes a pressure balance (relief) hole to equalise trapped pressure in the groove region.

  - *Pressure envelope:* Commonly specified for high-pressure systems up to 20,000 psi.
  - *Typical use:* Ultra-high-pressure wellhead and subsea-style connections, where pressure energisation and confinement are desired. [81]

Therefore, it is clear that the BX ring type is the only feasible option. Within this family, a variety of sizes are possible. The requirement for gasket placement is that its inner diameter should be larger (but as close as possible) than the chamber's inner diameter, and its outer diameter should be smaller than the bolt-circle diameter. The reason is that if the bolt diameter is smaller than the chamber's diameter, the gasket's purpose is lost, as part of the bolt compression is not applied to it, and the exposed part may introduce leaks. At the same time, it should be as close as possible to the inner chamber diameter, as otherwise the area over which the pressure is applied becomes larger, thereby increasing the required strength of the bolt and gasket. On the other hand, the bolt-circle diameter is the diameter of the imaginary line connecting the centre of the bolts, which defines the lever arm for the pressure-bending moment; naturally, if this is larger, then the flange will be thicker. Therefore, among the many options, the BX-169 is the most suitable, with an inner diameter of 131 mm and a pressure rating of 15 000 psi.

Flange thickness

Once this gasket has been selected, then the blow-out area, that is, the area over which the pressure is applied, and thus the blow-out force can be found:

$$A_b = \pi \left(\frac{D_{i,g}}{2} \right)^2 \tag{4.3}$$

$$F_b = \text{MAWP} \cdot A_b \tag{4.4}$$

where $D_{i,g}$ is the inner diameter of the gasket.

To determine the flange diameter and thickness, the only remaining step is to determine the number and type of bolts. For these, a trade-off is needed between bolt diameter and class. In fact, taking only metric bolts, they have classes ranging from 4.6 to 12.9, where these names correspond to their maximum tensile strength (4.6 $\approx$ 460 MPa, 12.9 $\approx$ 1200 MPa). At the same time, the bolt diameter also has an impact, shown in Equation 4.5.

$$A_s = \frac{\pi}{4} \cdot (d_b - 0.938194 \cdot p_b)^2 \tag{4.5}$$

$$F_i = A_s \cdot 0.2\% \ \sigma_{y,b} \tag{4.6}$$

where

A_s Stress area of the bolts [mm^2]
d_b Bolt diameter [mm]
p_b Bolt pitch [mm][3]
F_i Tensile force in each bolt [N]
$\sigma_{y,b}$ Yield Strength of the bolts [MPa]

This formula shows that a larger bolt diameter increases the stress area per bolt, as a larger surface area withstands the stress, and consequently, a larger force is applied to each bolt.

---

[3] https://help.solidworks.com/2020/english/SWConnected/cworks/r_Tensile_Stress_Area_Bolt.htm Accessed on 15/05/2025

The 0.2% yield strength is used in these calculations, as the industry standard for bolts is to design them for the proof strength. This naturally leads to the conclusion that a larger diameter and/or a higher class reduce the number of bolts required to clamp the flanges.

$$N_{bolts} = \frac{F_b}{F_i} \tag{4.7}$$

The minimum number of bolts is rounded up, and, for redundancy and safety, the required number is doubled, guaranteeing a safety factor of 2 on the flanges as well. Finally, the diameter of the bolt centre can be calculated:

$$D_{bc} = D_{o,g} + 2 \cdot (D_b + p_b + cl_{wrench} + cl_{washer}) \tag{4.8}$$

where

D_{bc}	Diameter of the bolt center line [mm]
$D_{o,g}$	OD of the gasket [mm]
D_b	Bolt diameter [mm]
p_b	Bolt pitch [mm]
cl_{wrench}	Wrench clearance [mm]
cl_{washer}	Washer clearance [mm]

With this dimension and the bolt diameter, the diameter of the flanges, which is the same for both, can be found with Equation 4.9, which also accounts for extra clearances due to the presence of the bolt head and the presence of a stress area around the holes.

$$D_{fl} = 1.5 \cdot D_b \cdot 2 + D_{bc} \tag{4.9}$$

For the thickness, a distinction needs to be made between the welded flange and the cover. Indeed, since the first one has the inner edge welded (i.e., clamped), whereas the other is free, it bends under axial load. Therefore, to calculate the thickness, Roark's plate bending formula can be used [82]. On the other hand, the cover flange is clamped by the bolt preload, meaning that the ASME VIII-1 UG-34 formula for a flat head under internal pressure can be used [83]. These two are presented hereafter:

$$t_{f,w} = a\sqrt{\frac{Kp}{\sigma_y}} + C + E_g \tag{4.10}$$

$$t_{f,c} = \sqrt{\frac{PD_{bc}}{4\sigma_y E} + C^2} + E_g \tag{4.11}$$

where

$t_{f,w}$	Thickness of the welded flange [mm]
$t_{f,c}$	Thickness of the cover flange [mm]
a	Radial span of the plate ($\approx \frac{D_{bc}}{2}$) [mm]
K	Bending coefficient (for a clamped-free plate $\sim$ 0.27 [-])
p	MAWP [Pa]
S	Allowable bending stress [MPa]
E	Joint efficiency ($\sim$ 0.7 for clamped plates) [-]
E_g	Gasket seating allowance (0.5-1) [mm]

Bolt selection

As mentioned, the procedure and final values depend on the bolt diameter and class. Therefore, a choice must be made to ensure sufficient safety and redundancy without affecting the structural integrity of the flanges. In fact, if there are too few bolts, then not only is there insufficient redundancy in case of failure, but also the clamping over the gasket is not uniform, and a much stiffer flange would be required; on the other hand, too many bolts require more holes, which reduces structural integrity, as there is too little metal surface between two adjacent holes. This trade-off is particularly relevant here because the high internal pressure imposes a large separating load on the joint, while the BX-169 gasket is a metal-to-metal ring joint seal whose performance relies on maintaining compression and proper groove contact. Designs commonly feature a pressure-relief feature to avoid pressure trapping. In practical terms, the bolt pattern and size must be sufficient to maintain a positive residual clamp load under peak pressure, with no flange separation, while keeping ligament widths and local stresses around the bolt holes within acceptable limits so that the SS 2205 flanges remain in the elastic regime.

Starting with the bolt class, to generate a high, repeatable preload, the first step was to compare property classes based on their proof stress, since this bounds the usable tightening load in an elastic design. ISO 898-1 gives nominal proof stresses of about 600 MPa for class 8.8, 830 MPa for class 10.9, and 970 MPa for class 12.9. Class 8.8 was discarded because, at a fixed diameter, the lower proof stress either reduces the preload margin (higher risk of unloading the BX gasket at the MAWP) or forces an increase in bolt count to recover clamp capacity, which directly penalises flange ligament efficiency. Class 12.9 was not selected because the joint is not governed only by bolt strength: the practical preload is limited by local bearing and crushing in the clamped parts (especially the gasket groove). Driving the design toward the higher 12.9 preload levels increases the risk of local damage in the flange and groove rather than improving joint reliability. Moreover, a higher target preload implies a higher tightening torque. Since torque control exhibits significant scatter and is sensitive to friction and galling, increasing torque levels upward increases uncertainty in the achieved bolt tension and is undesirable for a seal-critical, high-pressure joint.

For the diameter, the same approach was applied by comparing the tensile stress area and the resulting preload per bolt: for coarse threads, the stress surface increases from $245\,\text{mm}^2$ (M20) to $303\,\text{mm}^2$ (M22), $353\,\text{mm}^2$ (M24), and $459\,\text{mm}^2$ (M27). M20 and M22 bolts were therefore disfavoured because meeting the required total clamp load would generally require more fasteners, increasing hole density and reducing the ligament width between adjacent holes, thereby weakening the flange and increasing the severity of the stiffness requirements. Moving up to M27 increases the preload capacity per bolt, but it also requires larger holes and more flange real estate per fastener, and it can reduce the number of bolts that can be reasonably placed on the available pitch circle, which works against uniform gasket compression. M24 class 10.9 is the compromise that closes the clamp-load requirement with a margin while keeping bolt count, hole size, and remaining ligament widths in a range that preserves flange integrity and supports uniform BX gasket seating without relying on extreme tightening targets [84, 85].

Therefore, with this decision taken, the remaining measures can be determined, which are summarised hereafter in Table 4.4. The mass is an estimation based on the CAD version of the full combustion chamber.

Table 4.4: Initial chamber dimensions

D_i [mm]	H [mm]	t_{cyl} [mm]	N_b [-]	D_f [mm]	$t_{f,w}$ [mm]	$t_{f,c}$ [mm]	W_{ch} [kg]
130	130	22.49	8	311.51	29.35	60.99	~ 107

4.1.3. Final design

Based on the decisions taken in the previous sections, the final design can be drawn, as shown in Figure 4.2, with the sensors and a model of the igniters already inserted. Here, the chamber is shown on its supports; the design procedure for these supports is presented in Section B.2, as it is outside the scope of this work. While building the CAD model, it was noticed that a height of 130 mm would not work, as there would not be sufficient space for the nuts on the back or for the sensors. Therefore, this had to be increased initially to 140 mm. The full technical drawings for each part and assembly are given in Appendix C.

Once the CAD model is complete and fully constrained, a finite element analysis (FEA) can be performed directly in SolidWorks to calculate stresses at each point in the chamber. This, while showing that the bolts were well under their yield point, confirmed that the weak points were the insert holes. Therefore, two strategies are employed, which prove successful as shown in Figure 4.3. The first step is to redesign the holes so that all the inserts sit flush with the chamber, thereby providing stress relief through torque, as explained before. The second step is to gradually increase the chamber height to 160 mm, ensuring that all stresses remain below the yield margin.

This is because when a cylindrical pressure vessel is clamped at both ends and subjected to internal pressure, the end constraints prevent the cylinder's natural expansion, creating localised bending moments and stress concentrations near the boundaries. These boundary effects decay exponentially as you move toward the cylinder's centreline, following Saint-Venant's principle, with a characteristic decay length roughly proportional to the cylinder's radius. In a short cylinder, the disturbed stress fields from both clamped ends overlap significantly in the central region, meaning that holes located at the centreline experience a complex combination of membrane stresses, bending stresses, and boundary-induced perturbations. As the cylinder height increases, the

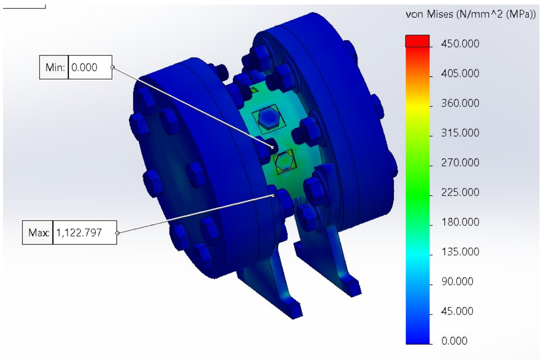

Figure 4.3: FEA stress analysis

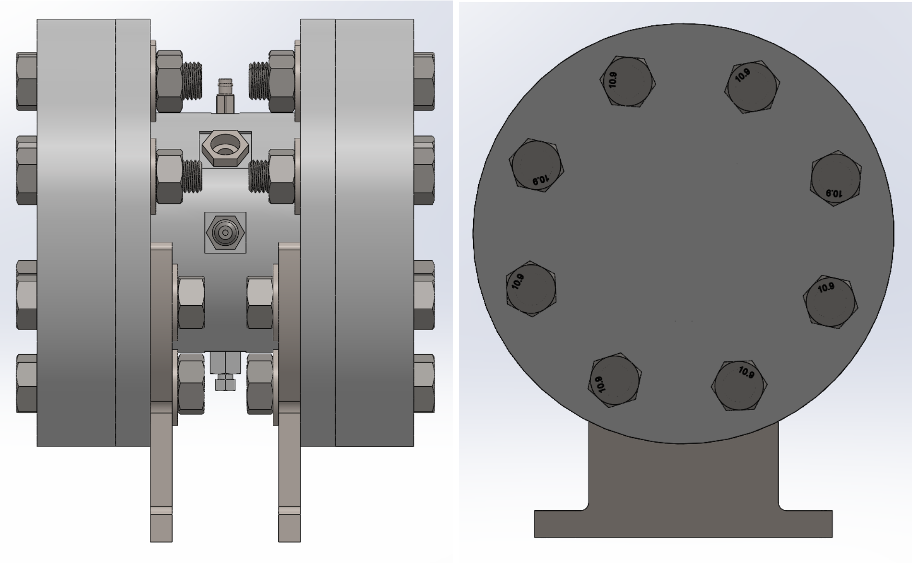

(a) Lateral view **(b)** Front view

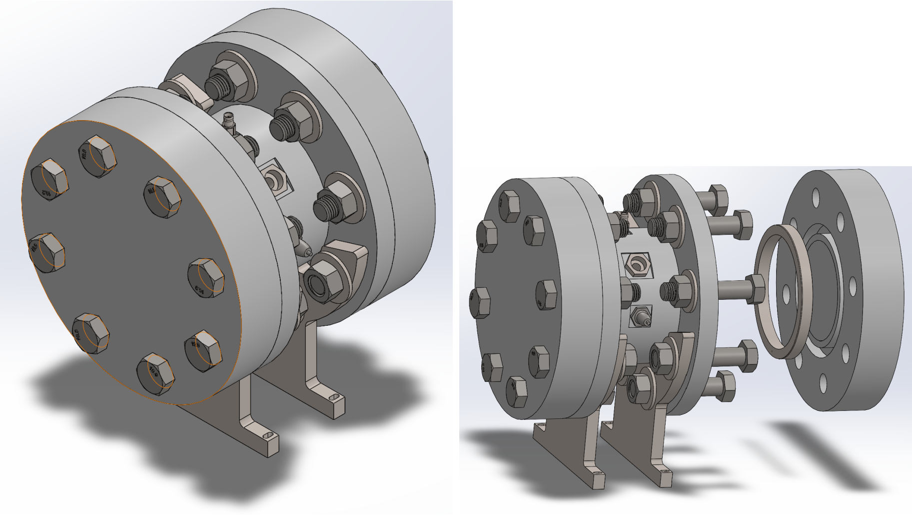

(c) Isometric view **(d)** Disassembled view

Figure 4.2: Combustion chamber CAD design

centreline moves farther from both boundaries, allowing the edge effects to fully dissipate before reaching the middle section. This results in holes at the centreline experiencing a stress state much closer to the classical thin-walled pressure vessel solution—dominated by predictable hoop and longitudinal stresses without the complicating influence of constraint-induced bending. Consequently, the stress concentrations around these holes are reduced because they're interrupting a simpler, more uniform stress field rather than one amplified by boundary effects, leading to lower peak stresses and improved structural performance.

The final flange thicknesses are slightly smaller than the designed thicknesses due to manufacturing constraints within the company's facilities, which are more than offset by the 2x safety factor used throughout the design. However, a problem arose during manufacturing: the groove for the gasket was not created according to specifications and was instead made into a simple square groove on one side of the flanges only, rendering the gasket seal ineffective. This is, in part, addressed by the second manufacturing step of the chamber, which is performed at a local company, where the groove is remade in accordance with the specifications on the other side of the flange. Thus, the external flanges lost some material, which is well below the safety factor, but this is still worth considering, especially for long-term use. The final chamber is shown in Figure 4.4, with final dimensions summarised in Table 4.5.

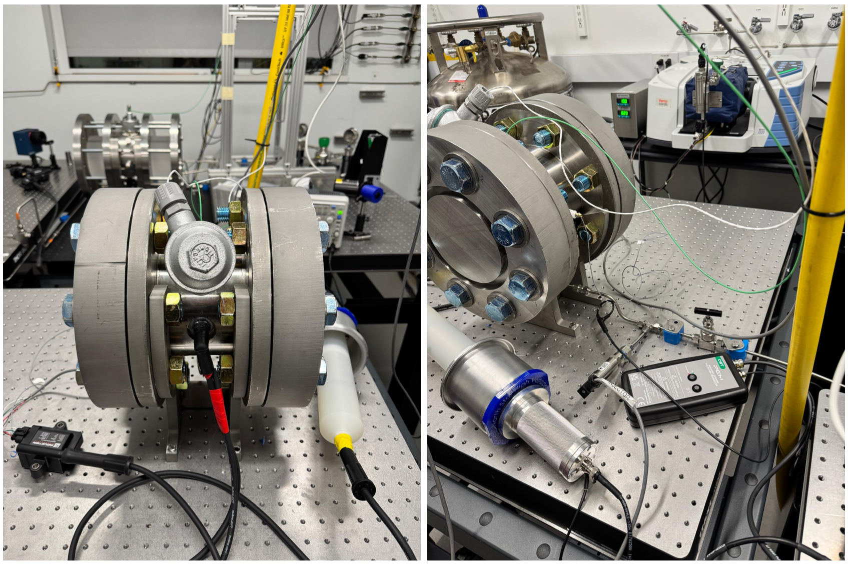

(a) Lateral view **(b)** Front view

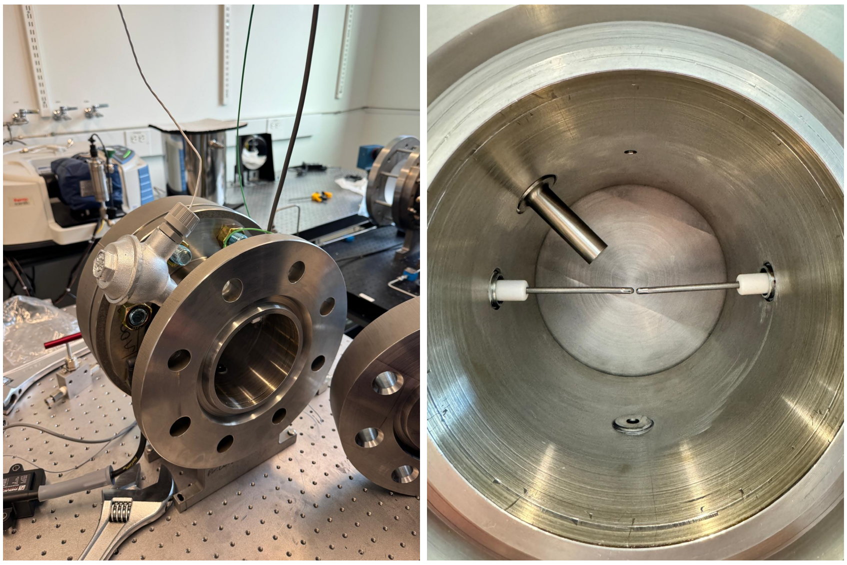

(c) Isometric view **(d)** Internal view

Figure 4.4: Non-optical combustion chamber

Table 4.5: Final chamber dimensions

D_i [mm]	H [mm]	t_{cyl} [mm]	N_b	D_f [mm]	$t_{f,w}$ [mm]	$t_{f,c}$ [mm]	W {kg}
130.05	159.77	22.35	8	311.15	25.40	57.15	≈ 104

4.2. Sensor and ignition selection

In parallel to the mechanical design of the chamber, its inserts and sensors also need to be selected. Among them, some require coupling with the ignition system to achieve chamber ignition. In this section, the ignition system design and the insert selection are described.

4.2.1. Ignition system design

For the ignition system, due to its ease of implementation and a significantly larger body of literature supporting it compared to other types of ignition, spark ignition was chosen from the beginning. However, a trade-off is still performed here to confirm the sanity of the choice, shown in Table 4.6. This clearly demonstrates that spark ignition is the superior option, particularly considering its suitability under the conditions of interest in this work. Therefore, now the best mode of spark production needs to be selected, which will be achieved in the trade-off shown in Table 4.7.

Table 4.6: Qualitative trade-off for ignition concepts (scores: 1 = worst, 5 = best).

Ignition mode	Energy control	Flame disturbance	Complexity	Cost	Suitability LTHP	Maintenance	Total
Spark	5	4	4	3	5	3	24
Plasma	3	3	2	2	4	2	16
Hot-surface	2	1	4	4	1	3	15
Pilot flame	2	2	3	4	2	4	17
Laser	5	5	1	1	4	2	18
Microwave	3	3	1	1	3	2	13

Table 4.7: Trade-off for spark ignition system architectures (scores: 1 = worst, 5 = best).

Architecture	Energy adjustability & repeatability	Energy characterization (MIE)	Breakdown headroom at MAWP	Cryogenic robustness	EMI impact on diagnostics	Integration complexity	Cost & availability	Total
CDI	5	5	5	4	5	2	2	31
TCI	4	3	4	4	3	4	4	30
Magneto	2	1	3	3	3	3	3	22
Battery	3	2	2	3	3	3	5	24

Therefore, as shown, the CDI appears to be the best option for the given requirements, especially for MIE calculations. Therefore, a full CDI circuit is initially designed and simulated in LTspice, as shown in Figure 4.5. However, implementing this circuit and procuring all the necessary components proved much more expensive and difficult than anticipated, as most components were highly specialised and had limited market availability. Therefore, after a long search, due to the time constraints of this work, it was decided to go for the second-best option to avoid wasting more time. This means a TCI system was preferred in the form of an ignition coil, the most common ignition mode in current automotive systems. An explanation of the CDI ignition system is provided for reference in Section B.3 for future work.

The requirements for this, apart from the ability to provide the required voltage identified in Subsubsection 4.2.2.3, are to also provide the necessary energy under the worst conditions, namely, for ammonia at room pressure and -50° C. For this reason, since no measurements exist, an assumption must be made. This will be based on the work of Nawaz and Mack [45],

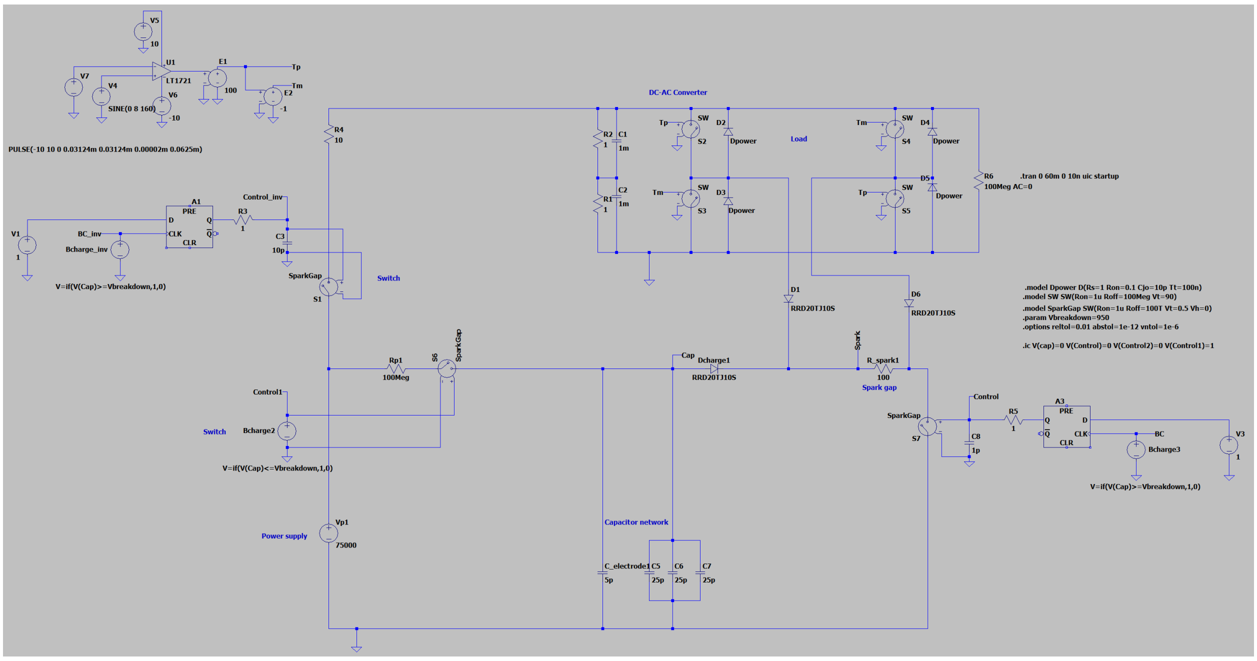

```
.tran 0 60m 0 10n uic startup
```

```
.model Dpower D(Rs=1 Ron=0.1 Cjo=10p Tt=100n)
.model SW SW(Ron=1u Roff=100Meg Vt=90)
.model SparkGap SW(Ron=1u Roff=100T Vt=0.5 Vh=0)
.param Vbreakdown=950
.options reltol=0.01 abstol=1e-12 vntol=1e-6

.ic V(cap)=0 V(Control)=0 V(Control2)=0 V(Control1)=1
```

Figure 4.5: CDI ignition system

who extrapolated that the MIE at this temperature is about 190 mJ. However, this estimation is based on chemical models that are not validated at low temperatures, especially below the melting point (~ 240 K), and may be inaccurate. However, it is considered a reasonable starting point that will need to be confirmed experimentally.

Among the various options, the IGN-1A coil from Induction Performance, shown in Figure 4.6, is the only one that meets every requirement, with a maximum output voltage of 81 kV and a relative energy release of 250 mJ. This is a smart coil, meaning the dwell time can be freely controlled between its maximum and minimum values, above and below which the coil can be damaged. The dwell time refers to the period during which the primary circuit is closed, allowing the current and, consequently, the energy to accumulate in the system. Once this ends, the primary circuit closes, creating a magnetic field that activates the secondary circuit, which provides the voltage step-up. Naturally, the energy remains the same in both (conservation of energy with minimal losses), so the current through the secondary circuit will be much lower.

Figure 4.6: IGN1-A ignition coil[4]

4.2.2. Sensor and inserts selection

Since the overall chamber geometry has now been defined, the next step is to select the inserts required to instrument the vessel and enable the combustion experiments. In practice, this includes the injection and exhaust ports, the ignition electrodes, and the sensors needed to capture the relevant diagnostics. As discussed previously, the two key quantities addressed in this work are the laminar burning velocity and the minimum ignition energy, which impose specific requirements on the measurement hardware. For this reason, the choice of probes

---

[4] https://www.inductionperformance.com/en-es/products/ign1a-smart-coils [Accessed on 20/08/2025]

is not only driven by pressure and temperature ratings but also by their interaction with the internal flow and the development of the flame. Two general constraints are therefore applied to all inserts: the threaded engagement should be as flush as possible with the inner wall, and the protrusion into the chamber volume should be minimised.

The flush-thread requirement is primarily motivated by structural considerations. Any penetration through a pressure vessel wall introduces a stress concentration that must be explicitly accounted for, since it can become the governing failure location. This effect is illustrated in Figure 4.7, where the local stresses increase markedly in the vicinity of the hole relative to the nominal membrane stress in the cylindrical wall. If these local peaks are not mitigated, failure may initiate at the penetrations at pressures significantly below the MAWP, even if the global design appears conservative. Two established strategies exist to reduce these concentrations, namely local reinforcement around the opening and the use of threaded engagement and preload to provide stress relief through compressive contact at the interface.

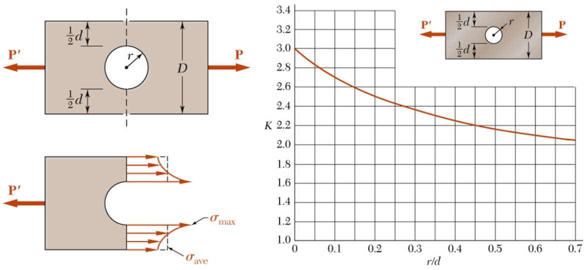

Figure 4.7: Stress concentration at hole locations [86]

While the reinforcement requirement can only be confirmed through a full stress analysis once the final hole pattern is defined, the second strategy can be incorporated directly at the insert-selection stage. In particular, inserts should be chosen with the smallest practical thread diameter to minimise geometric disturbance and with a thread engagement length as close as possible to the local wall thickness, ensuring proper load transfer without leaving an unthreaded cavity. A perfect match is rarely achievable; therefore, small external pockets can be machined to ensure the insert shoulder seats correctly while keeping the internal surface as flush as possible. This solution also enhances manufacturability by providing a flat reference for drilling and facilitating tool access during insert tightening. Additionally, a controlled seating surface helps promote repeatable tightening, which is crucial when preload is relied upon to stabilise the local stress state around penetrations. However, this cannot be overdone, as the wall would lose too much thickness and therefore structural integrity.

The second key parameter in insert selection is the intrusion into the chamber volume, since protruding features can alter the flame expansion and locally distort the pressure field. A convenient way to quantify this impact is through the blockage ratio, which measures the fraction of the chamber's cross-sectional area obstructed by the insert's projected frontal area. For a characteristic intrusion diameter d within a chamber of diameter D, the blockage ratio can be expressed as shown in Equation 4.12. This metric provides a straightforward screening criterion when comparing alternative sensor geometries, particularly for probes that extend into the flow. In the present context, it is used as a design constraint to ensure that instrumentation does not become a dominant source of systematic bias in the flame dynamics.

$$\Phi = \frac{\pi(d/2)^2}{\pi(D/2)^2} = \left(d/D\right)^2. \tag{4.12}$$

The threshold adopted for acceptable blockage is based on Li et al. [87], who report that flame-front distortion remains minimal, meaning below the experimental scatter, for $\Phi \leq 0.1$, and becomes increasingly noticeable above this value. This result is particularly relevant here because both LBV and MIE are sensitive to early flame development and to the uniformity of the pressure rise. Consequently, the blockage ratio must be checked for each candidate insert, and it becomes an explicit constraint when selecting sensor bodies, probe diameters, and any protective housings. When the blockage ratio cannot be reduced further due to sensor availability or strength requirements, the intrusion location and orientation should be chosen to minimise its interaction with the primary flame growth direction. For these reasons, Φ is treated as a practical design metric that directly informs the final selection and placement of the instrumentation.

4.2.2.1. Inserts

Starting from the general inserts, the chamber will need several components:

Injector/exhaust

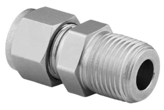

Figure 4.8: Swagelok SS-400-1-8[5]

As mentioned in Section 3.3, the optical chamber features a single line for both gas injection and exhaust, indicating that a needle valve is used right next to it, therefore extending the chamber volume slightly. To quantify this error, it is to be noted that 1/4" OD tubes are used, with a wall thickness of 0.035", meaning that the circular area will be of $\approx 23.4\,\mathrm{mm}^2$. Thus, even assuming a tube length of 300 mm between the chamber and the needle valve, the total added volume would be 0.007 L, which, for a 2 L chamber, represents less than a 0.4% error, thus quite negligible. Thus, a similar setup is also adopted for the non-optical chamber, eliminating the need for a complete redesign of the injection and exhaust lines. The error due to the possible exhaust gas remainders in the line leading from the chamber to the purge valve is also considerably small, considering the evacuation to a pressure of about 200 mTorr, but will still be considered in the calculation of the LBV uncertainty.

Therefore, a simple tube-to-chamber connector is needed for this purpose. Since this does not need to be changed very often, tapered threads are the better option given the high-pressure conditions. This is because the seal is not happening over a face, as it happens with straight ones, which rely fully on a gasket or ring. Therefore, the taper provides a stronger lock as torque increases, ensuring a better seal on the threads under high-pressure conditions. Among the different options, the SS-400-1-8, 1/4" Tube - 1/2" MNPT connector, from Swagelok, shown in Figure 4.8 is found to have the thread length closest to the wall thickness.

Igniter

Since the objective is to produce a flame as spherical and symmetric as possible, the ignition strategy must minimise any bias introduced by the ignition location and geometry. Achieving ignition in a manner similar to practical engines, that is, with a spark plug mounted at the wall, becomes problematic under this requirement. In fact, a spark plug consists of a wall insert with two electrodes in parallel, very close to each other, which reliably guarantees a consistent spark. However, the drawback is that the spark is created very close to the wall, so the initial kernel

---

[5] https://products.swagelok.com/en/c/straights/p/SS-400-1-8 [Accessed on 29/06/2025]

grows in a strongly non-uniform environment influenced by the wall and local heat losses. This tends to promote uneven early flame development and an asymmetric pressure rise, which disrupts the accurate measurement of the pressure trace. Additionally, the assumption of bulk measurement no longer holds, and the pressure sensor's positioning becomes crucial.

For this reason, a configuration similar to that shown in Figure 3.3d is selected for the optical chamber. This consists of two rod igniters placed one in front of the other, with the advantage of freely adjusting the gap distance by cutting the electrodes, which helps in conditions where the spark does not occur due to insufficient energy or excessive heat loss. Ignition is achieved by creating a voltage difference between the two igniters: one is grounded, and the other is set to a high voltage. To increase the probability of ignition, the electrodes should be tapered towards the end, thereby increasing the strength of the electric field and orienting it towards a smaller area.

An additional benefit of this geometry is that the most likely breakdown path is the spark gap itself, as it corresponds to the smallest electrode separation and therefore requires the lowest breakdown voltage. The main competing failure mode is the formation of a discharge between an electrode and the chamber wall, which can occur if the effective breakdown conditions shift, for instance, due to changes in gas composition, pressure, or temperature that affect the Paschen curve. This behaviour must be prevented because a wall discharge would reintroduce the same near-boundary ignition problems discussed above while also increasing the risk of local surface damage. The risk is reduced by selecting igniters with robust ceramic insulation on the chamber side, thereby constraining the electric field to the intended gap region. Among the available options, US Ignition R6045-type igniters, shown in Figure 4.9, are selected because they offer strong ceramic insulation and a 1/2" NPT threaded interface, which is compatible with the insert philosophy described previously and supports reliable, repeatable installation. Moreover, the electrode is made of high-strength alumina, capable of withstanding voltages up to 100 kV.

Figure 4.9: US Ignition R6045 igniter [6]

Static pressure sensor

Since the partial pressure method will be used for filling, the static pressure inside the chamber must be measured accurately. Therefore, a pressure sensor capable of measuring this is needed, with the most accurate category being the piezoresistive sensors. For this, the laboratory had a spare Omega PX419-1.5KAV, shown in Figure 4.10, capable of measuring pressures up to 1500 psi. This is, therefore, a great option to meet the required specification. However, since it is desirable to minimise the number of inserts into the chamber, it is worth noting that the static pressure remains constant in any closed system. This means that the

Figure 4.10: Omega PX419-1.5KAV transducer[7]

---

[6] https://horizonpfm.com/horizon-r6045-spark-ignitor-or-flame-rod-1-2-length-5/r6045/ [Accessed on 28/06/2025]

sensor can be freely placed on the portion of the tube between the needle valve and the chamber, while performing the same function as if it were inserted directly into the chamber. This sensor can also serve as a sanity check to verify that the dynamic sensor functions as intended, as it can measure transients, albeit with a slower response, as seen in the specifications presented in Table B.2.[7]

4.2.2.2. LBV sensors

For the LBV calculation, given the requirements, the most accurate method is the spherically expanding flame, which is the only one available for high-pressure measurements. For this reason, it will be the one used in this study. This method requires two main measurements: the combustion pressure trace and the initial temperature. Therefore, inside the chamber, two sensors are needed: a dynamic pressure sensor to measure the fast combustion pressure transient and a static temperature sensor to accurately measure the initial pressure, particularly important when the chamber is brought away from ambient temperature.

Pressure sensor

To capture the combustion event in a closed vessel, the pressure measurement must resolve fast transients at sub-millisecond timescales while operating from an initial condition as low as -50° C and an initial pressure of 50 bar. This combination immediately favours piezoelectric technology because piezoelectric sensors offer very short rise times and high natural frequencies, which are essential to avoid filtering or phase lag in the recorded pressure trace. In addition, the relevant quantity for combustion analysis is the time-varying component of pressure, rather than a perfectly stable DC reading. For this reason, the selection was primarily driven by dynamic performance (rise time, resonance, and thermal shock behaviour) and survivability across the required temperature envelope, rather than by the ability to measure static pressure indefinitely.

Within piezoelectric sensors, the key architectural decision is between PE (charge output) and IEPE (integrated electronics, voltage output). In an IEPE sensor, the charge-to-voltage conversion electronics are built into the sensor housing, which simplifies cabling and interfacing but constrains the sensor to moderate operating temperatures and a fixed measuring range. Kistler explicitly distinguishes these options and notes that IEPE variants are intended for dynamic profiles over short durations and "moderate temperatures," whereas PE sensors, with an external charge amplifier, can operate at extreme temperatures and allow more flexibility in the measurement chain [8]. Since the experiment must be achieved at temperatures as low as -50° C, and can experience large thermal gradients during ignition and combustion, the PE approach is more robust: the sensor itself remains a passive element at the chamber boundary, while the charge amplifier can be kept in a controlled ambient environment. This configuration also reduces the risk that embedded electronics become the limiting factor during cold starts or transient heating.

Based on these criteria, the initial choice was the Kistler 601CAB. The Type 601C family is designed for highly dynamic pressure measurements, featuring an optimised diaphragm that withstands thermal shock and a rise time of less than a microsecond, which is suitable for resolving the steep pressure rise associated with ignition and early flame development. Importantly for the subzero requirement, the PE versions of the 601C series specify an operating temperature range extending well below -50° C (down to -196° C), which provides a considerable

---

[7] https://assets.dwyeromega.com/pdf/test-and-measurement-equipment/pressure/pressure-transducers/PX409_Series.pdf [Accessed on 20/06/2025]

[8] https://kistler.cdn.celum.cloud/SAPCommerce_Download_original/003-199e.pdf [Accessed on 23/05/2025]

margin and reduces the likelihood of offset shifts during the initial stabilisation period. The "B" variant (601CAB) corresponds to the short-housing configuration, which is advantageous for integration into a compact vessel wall, as it minimises the structural impact on the wall and reduces the exposed sensor length. In practice, the 601CAB therefore matched both the measurement physics (fast transient) and the environmental envelope (cold start with potential thermal shock).

Later, due to procurement constraints, the sensor was switched to the Kistler 603CAA, shown in Figure 4.11. The 603C series maintains the same fundamental benefits of piezoelectric sensing for combustion transients, while offering an even higher dynamic capability (shorter rise time and higher natural frequency) and a substantially higher maximum pressure range, which increases robustness against unexpected peak pressures during energetic tests.[9] Although the 603C series generally has lower sensitivity than the 601C family (a typical trade-off when increasing full-scale range), the signal-to-noise ratio remains adequate for combustion work when paired with an appropriate charge amplifier and careful grounding/shielding. As with the 601C, the PE variants of the 603C series are specified for cryogenic operation, ensuring the initial condition is not a limiting factor. The switch to 603CAA therefore preserved the core experimental capability while maintaining pressure headroom and compatibility with the low-temperature requirement.

Figure 4.11: Kistler 603CAA transducer [9]

Combustion-dedicated cylinder pressure sensors, such as the Kistler 7061 series, used in the optical chamber, were considered but were not suitable for the full test envelope. The 7061C is a water-cooled reference sensor optimised for engine research, with excellent stability under hot combustion conditions, but its specified operating temperature range only extends down to -20° C[10]. In addition, its cooling concept and coolant specification (including mixtures suited only down to approximately -9° C as listed in the datasheet) make subzero cold-start operation operationally complex and not aligned with the requirements. In contrast, the miniature 601C/603C PE sensors are specifically designed for dynamic transients and wide-temperature operation, making them more suitable for a subzero initial state while still meeting combustion bandwidth requirements. Their main specifications are summarised in Table B.3.

Temperature sensor

For temperature, the requirement is fundamentally different: the experiment needs a reliable measurement of the static (or slowly varying) chamber temperature rather than millisecond-scale gas-temperature fluctuations during combustion. Indeed, no temperature sensor can measure a transient as quickly as in the order of hundreds of microseconds. For this reason, a platinum RTD probe, specifically the Omega PR-12E-3-100-1/4-2, was selected, as RTDs provide stable, accurate, and repeatable absolute temperature measurements. The PR-12 series is offered with Class "A" accuracy to IEC 60751 and a temperature range that covers the cold-start requirement and provides ample margin for post-test heat soak; Omega lists an "extreme" range extending to -50° C for the PR-12E configuration[11]. The 3-wire configuration

---

[9] https://kistler.cdn.celum.cloud/SAPCommerce_Download_original/003-288e.pdf [Accessed on 24/05/2025]

[10] https://kistler.cdn.celum.cloud/SAPCommerce_Download_original/003-042e.pdf [Accessed on 23/05/2025]

[11] https://www.omega.co.uk/pptst/PR-12.html [Accessed on 25/06/2025]

further reduces lead-resistance error compared to 2-wire wiring, particularly when the readout electronics are located away from the chamber in a colder laboratory environment. Finally, the 316 stainless-steel sheath is compatible with pressurised test hardware, mitigating concerns about corrosion or embrittlement compared to more reactive alloys.

However, the sensor alone cannot withstand the high pressures generated in the chamber. For this reason, the RTD was installed inside a thermowell (Omega 3/4-260A-U15/8-316SS) to decouple the electrical sensor assembly from the pressure boundary and to improve robustness in a combustion environment. A thermowell provides a sealed metallic barrier between the process and the sensing element, simplifying pressure sealing and protecting the RTD from direct exposure to combustion products, soot, and mechanical damage during assembly. The 260A series is specifically designed for limited-space installations and is dimensioned for 1/4" diameter elements with a 0.260" bore, ensuring a close match to the selected PR-12 probe diameter without excessive clearance[12].

(a) Omega PR-12-E-3-100-1/4-2 sensor [11] **(b)** Omega 3/4-260A-U15/8-316SS thermowell [12]

Figure 4.12: Temperature acquisition components

Using 316 stainless steel ensures material compatibility with typical pressure-vessel hardware and provides a strong pressure–temperature capability. The datasheet presents pressure-temperature ratings in the thousands of psi range for stainless-steel wells, placing 50 bar (about 725 psi) comfortably within typical limits. The main drawback is the added thermal mass, which increases response time; however, this is acceptable because the objective is to maintain the initial/static temperature and slow transients, rather than resolving the combustion event itself (which is handled by the piezoelectric pressure sensor). In any case, some thermal paste is applied between the sensor and the thermowell to improve measurement speed, since the sensor's response time is already quite slow.

Since the thermowell has a significant intrusion into the chamber, the blockage factor must be calculated to determine whether it might affect the measurement. With D_{th}= 12.7 mm and D = 130 mm:

$$\Phi = \left(\frac{D_{th}}{D}\right)^2 \approx 0.95\%$$

---

[12] https://assets.dwyeromega.com/pdf/hardware/thermowells/SERIES_385A_260A.pdf [Accessed on 25/06/2025]

As seen, this is one order of magnitude below the 10% threshold, so blockage effects can be considered negligible. Therefore, the specifications for the used sensor are summarised in Table B.4.

4.2.2.3. MIE sensors

The determination of the ignition energy depends on the voltage and current at ignition, as well as the duration of the streamer in the gap. Therefore, voltage and current probes are necessary and will be placed on the igniters to precisely measure the voltage across the gap and the resulting current. However, to select these probes, the worst-case conditions for voltage and current production must be estimated, so that an instrument can be chosen that not only operates under those conditions but also remains undamaged by them. Therefore, for the voltage probe, this will be the maximum breakdown voltage of the mixture at the most constraining pressure and temperature, while for the current probe, it will be the highest possible created current. Both sensors will be connected to the oscilloscope BK2190E already present in the laboratory.

The breakdown field at 50 bar was estimated by scaling from the well-known atmospheric breakdown strength of air, which is approximately 30 kV/cm (i.e., 3 kV/mm) under dry, near-STP conditions[88]. For a fixed electrode gap operated away from the Paschen minimum, Paschen's law indicates that the breakdown voltage depends primarily on the product pd, and at sufficiently large pd the breakdown voltage (and therefore the breakdown field $E_{bd} = V_{bd}/d$ scales approximately linearly with pressure[89]. Under the assumption that ammonia behaves similarly to air in terms of the reduced field needed to sustain an ionisation avalanche (i.e., comparable collision length scales and ionisation thresholds), it is reasonable to treat ($\mathrm{E_{bd}}/\mathrm{p}$) as being of the same order of magnitude between the two gases. This is supported at a first-order level by the fact that ammonia has a molecular "size" in the same range as common diatomics (e.g., kinetic diameter ~ 0.26 nm for NH_3, compared with ~ 0.35 nm to 0.36 nm for O_2/N_2[90], and an ionisation energy on the order of 10 eV, consistent with a comparable threshold for electron-impact ionisation). Applying the pressure scaling from 1 bar to 50 bar then gives

$$E^*_{\mathrm{bd}}(50\mathrm{bar}) \approx 50 \times 3\mathrm{kV/mm} \approx 150\mathrm{kV/mm}$$

, which is the value adopted for conservative insulation and electrode-gap design, subject to the usual caveat that real breakdown depends on field non-uniformity, surface finish, and local temperature and composition.

The -50° C requirement is then addressed by recognising two competing effects: density scaling pushes E_{bd} upward, while mixture thermodynamics maintains the gap behaviour close to that of air. In similarity terms, gas breakdown is often governed by the reduced field E/N, where N is the number density; since $\mathrm{N} \sim \mathrm{p/T}$, holding the critical E/N approximately constant implies $E_{bd} \propto p/T$, so cooling from ~ 293 K to 223 K would increase the breakdown field at fixed pressure by a factor ~ 293/223 ≈ 1.31[91]. However, in an ammonia–air mixture at -50° C, the ammonia partial pressure in the gas phase is limited by its saturation vapour pressure (about 40.9kPa ≈ 0.41bar) at this temperature, meaning that most of the mixture in the electrode gap remains air rather than NH_3 vapour. Under those conditions, using the air-based scaling to obtain ~ 150kV/mm remains a reasonable baseline, and it is not undermined by ammonia's intrinsic ionization threshold, which is of the same order as air constituents (NH_3 10.07 eV), O_2 12.07 eV), N_2 15.581 eV))[92]. Therefore, as a conservative estimate, 160 kV/mm can be taken as the assumed value for breakdown at these conditions. Naturally, the real value will also be dominated by electrode geometry, surface condition, and local field enhancement.

Voltage probe

Given the above discussion, the voltage probe should operate with a breakdown voltage of about 160 kV/mm. Therefore, assuming a gap value of $\sim 0.5mm$, the voltage probe will need to be able to read voltages up to 90 kV to 100 kV. Moreover, the second requirement is that it should have a bandwidth large enough to read events in the order of microseconds. Given the very limited options on the market, the PVM-5 from North Star High Voltage is selected, with specifications presented in Table B.6.[13] To obtain a measurement, this probe must have its tip in contact with a high-voltage surface. Since contact with the electrodes is impractical and not very sensible, the contact needs to occur on the pin of the igniter, through an insulated extension cable, to avoid accidental secondary discharges, as shown in Figure 4.13b.

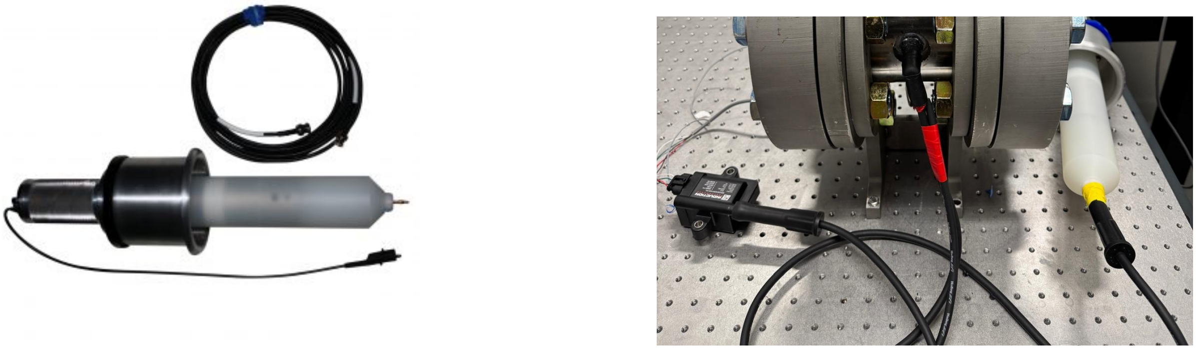

(a) North Star High Voltage PVM-5[13]

(b) Extension cable for the voltage probe connection

Figure 4.13: High-voltage reading setup

Current probe

For the current probe, the bandwidth requirement is the same as for the voltage probe, but the current requirement is much less restrictive, as only a few amperes are typically delivered across the spark gap. Therefore, the CWT Mini50HF 015/B/1/100/2 from PEM is chosen, whose characteristics are summarised in Table B.7.[14]

Figure 4.14: CWT Mini50HF probe[14]

---

[13] https://www.highvoltageprobes.com/bfd_download/voltage-probe-data-2021/ [Accessed on 17/08/2025]

[14] https://www.pemuk.com/uploads/files/instruction-sheets/CWTMini50HF-instruction-manual.pdf?v=1747394125 [Accessed on 20/08/2025]

4.3. Pipeline design

Once the chamber is designed, the remaining experimental system around it needs to be designed. This involves both redesigning the pipeline leading to the chamber and selecting the components necessary to achieve high pressures and low temperatures. The full process will be explained hereafter.

4.3.1. Booster design

A pressurisation system is required to pressurise the feed line and enable initial line pressures above the maximum achievable with the existing tank regulators. The first step is to define a clear maximum pressure requirement, because this value drives both component selection and safety margins. In the present work, the sizing criterion is based on a conservative injection-pressure rule. Specifically, the injection pressure is assumed to be a multiple of the reactant partial pressure to prevent backflow and to ensure robust injection under transients. Therefore, the system must be designed around the most demanding partial-pressure case within the intended operating envelope.

As described in Chapter 5, the injection pressure is set to six times the partial pressure of each reactant, and this factor is intentionally conservative. Under this assumption, the maximum pressure the pressurisation system must provide is 6 times the highest expected reactant partial pressure during operation. To identify that value, only equivalence ratios in the range $0.7 \leq \phi \leq 1.3$ are considered. Within this range, the highest relevant partial pressure occurs for the H_2 combustion case at a chamber pressure of 50 bar. For $\phi = 0.7$, the nitrogen partial pressure is approximately 30.5 bar, implying a maximum injection pressure of $6 \times 30.5\,\text{bar} \approx 183\,\text{bar}$.

This requirement highlights a substantial mismatch with the capabilities of the existing tank regulators. With the current nitrogen regulator, the maximum delivery pressure is 125 psi, which corresponds to approximately 8.6 bar. Comparing this value with the required injection pressure indicates that the pressurisation system must bridge a ratio of about $183/8.6 \approx 21$, i.e., on the order of a $\times 22$ pressure step. In practice, the system should be capable of exceeding this minimum ratio by a slight margin to account for uncertainty, losses, and future extensions of the operating envelope. For these reasons, the pressurisation concept is selected not only to meet the present requirement, but also to preserve flexibility for different fuels and a wider range of equivalence ratios.

Among the candidate solutions, a pneumatic gas booster offers a straightforward, robust approach to achieving the required pressure increase. A gas booster uses compressed air to drive a large-area piston, and this motion is mechanically coupled to a smaller-area piston acting on the process gas. By force balance, the pressure multiplication is approximately equal to the ratio of piston areas, so a modest drive-air pressure can generate a much higher outlet pressure on the process side. The pressurised process gas is delivered through a high-pressure outlet port, while the drive air is exhausted to the atmosphere. As a result, the booster provides high outlet pressures without requiring an electrically driven compressor to directly compress the process gas.

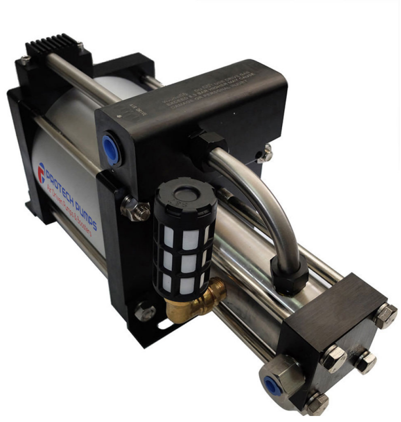

Figure 4.15: PB30 gas booster[15]

Among the commercially available options evaluated, the

Protech Pumps PB30 is selected because it offers low cost, a short lead time, and compatibility with the set of reactants considered in this work. The PB30 unit, shown in Figure 4.15, has a maximum outlet pressure of 265 bar. This rating exceeds the maximum required injection pressure of approximately 183 bar, providing margin for uncertainties and for possible future operating points. The availability of margin is particularly valuable because the present ×6 criterion is conservative, and future campaigns may explore different mixtures or chamber pressures. Consequently, the PB30 enables the required high-pressure capability while keeping the pressurisation architecture mechanically simple and modular.[15]

A practical integration issue arises because the booster outlet pressure is controlled entirely by the drive-air inlet pressure. In the present configuration, the booster operates with a drive-air pressure in the approximate range of 2–8 bar, and the drive air is supplied by a compressor. However, the compressor's minimum outlet pressure is 60 bar, which is significantly higher than the maximum pressure achievable using the existing tank regulators alone. This creates a discontinuity between the "low-pressure" regime provided by the regulators and the "high-pressure" regime produced by the booster when driven by the compressor. Without additional hardware, this gap would prevent operation at intermediate line pressures, which are expected during commissioning and during tests requiring moderate injection pressures. Therefore, an additional pressure-control element is required downstream of the booster to ensure continuous coverage of the full pressure range.

To resolve this gap, a high-pressure regulator is installed downstream of the booster, rated for an inlet pressure of at least 200 bar. This regulator allows the downstream line pressure to be set accurately to any desired value, including pressures below the raw booster outlet pressure. In operation, the system can run in a low-pressure mode by routing flow from the tanks directly to the chamber through the standard regulator path. Alternatively, it can run in a high-pressure mode by routing the gas through the booster and then using the downstream regulator to set the final line pressure. In high-pressure mode, the regulator may throttle to reduce pressure or remain fully open if the booster outlet pressure matches the required setpoint. This architecture ensures that any target pressure within the operating range can be achieved while maintaining controlled, repeatable pressurisation.

Finally, the maximum usable line pressure in the current implementation is limited by the pressure ratings of downstream components. In particular, the three-way valves are rated for a maximum pressure of 172 bar, which constrains the maximum permissible filling pressure of the line. This limit does not necessarily preclude operation at higher nominal injection pressures in all cases, because the ×6 scaling used for sizing is intentionally conservative. For specific operating points, smaller multipliers may still prevent backflow while remaining within the component ratings, provided that the backflow risk is reassessed and validated. Nonetheless, the 172 bar component rating defines the safe upper limit for routine operation and should be treated as the primary constraint for the current hardware configuration. Future upgrades to higher-rated valves would remove this constraint and allow full exploitation of the booster pressure capability.

4.3.2. Cold-environment system selection

To meet the requirements, it is necessary to maintain a temperature of -50 °C throughout the entire chamber, including the walls, not just the mixture. This is because the system is designed to replicate a real engine, which is therefore exposed to subzero external temperatures. However, reaching a chamber wall temperature as low as −50 °C is not, in itself, sufficient: the cooling

---

[15] https://www.protechpumps.com/shop/gas_boosters/pb/pb30/ [Accessed on 10/08/2025]

method must also provide repeatable, well-characterised thermal boundary conditions. This is particularly important when measuring minimum ignition energy (MIE) and laminar burning velocity (LBV), as small variations in gas density, heat losses, and near-electrode conditions can significantly alter the measurement. The selected solution must therefore cool the chamber in a reasonably short time while maintaining a uniform and stable thermal field, without introducing artefacts such as mixture stratification, local condensation, or electrical issues around the instrumentation. Additionally, the implementation should be safe, operationally simple, and compatible with repeated test campaigns.

The following criteria were used to compare candidate cold-environment concepts:

- Temperature capability, stability, and repeatability at −50 °C, including the ability to perform controlled ramps and soak periods.
- Cooling time and operational throughput, i.e., how quickly a stable thermal condition can be reached between consecutive tests.
- Thermal uniformity at the chamber and around ignition hardware, to avoid local hot/cold spots that bias MIE and LBV.
- Compatibility with sensors, cabling, and feedthroughs, avoiding frosting, water ingress, or thermal shock of components.
- Non-intrusiveness with respect to the mixture (no dilution, no phase change induced by the cooling method, and minimal risk of stratification).
- Practicality and safety (cryogen handling, oxygen-deficiency risks, and overall complexity of integration).

LN_2 bath with full submersion of the combustion chamber

Submerging the pressure vessel in a tank filled with liquid nitrogen (LN_2) offers a large temperature margin and very high instantaneous cooling capacity, which can appear attractive when targeting temperatures of −50 °C. However, heat transfer is governed by boiling at the external wall, meaning the effective heat transfer coefficient can vary substantially during cool-down and even during the soak, reducing repeatability. The temperature field is also likely to be non-uniform, as ports, flanges, and local thermal masses cool at different rates, resulting in a transient thermal gradient that is difficult to quantify without extensive instrumentation. From an experimental-integration standpoint, full submersion complicates the routing and sealing of electrical connectors, sensor cables, and high-voltage lines, and increases the likelihood of ice formation around penetrations as ambient moisture freezes on cold surfaces. Finally, this option introduces significant cryogenic handling and safety constraints (including the risk of oxygen deficiency in poorly ventilated areas) and increases the risk of thermal shock and material brittleness in seals or non-metallic parts if cool-down is not carefully controlled.

LN_2 introduced inside the combustion chamber

Using LN_2 directly in the test volume (either as a pre-cooling step or as a transient internal cooling mechanism) provides the least controlled and most intrusive boundary condition among the options considered. Any residual LN_2 or cold nitrogen gas left over during mixture preparation would significantly affect the thermodynamic state and composition, effectively invalidating the initial conditions required for MIE and LBV measurements. The method also risks local cold spots and strong buoyancy-driven flows, which promote stratification and non-uniform equivalence ratios near the igniters. In practice, condensation or freezing of trace contaminants and moisture becomes highly probable, and the resulting particles or films can alter surface properties and electrical breakdown behaviour. Moreover, rapid internal cooling can cause large and poorly controlled pressure and temperature transients, which complicate both safety and data interpretation. For these reasons, LN_2 inside the chamber was discarded early as being incompatible with controlled-ignition limit experiments.

Cooling the reactants via a heat exchanger

Pre-cooling the reactant streams through an external heat exchanger can efficiently reduce inlet-gas temperature and may appear to be a compact solution if only the mixture temperature were of interest. In this work, however, the objective is to control the chamber and near-wall environments, as wall heat losses and local temperatures at the ignition site can significantly influence ignition behaviour at high pressure. Cooling only the incoming gases does not guarantee that the vessel wall and ignition hardware reach a stable −50 °C condition, particularly if the chamber has significant thermal inertia. This option also creates additional sensitivities: any temperature-dependent adsorption/condensation of species in the supply line or at valves can shift the delivered composition and thus bias MIE/LBV. The approach further increases plumbing complexity and introduces thermal transients during filling, which can couple into pressure measurements and repeatability. Consequently, reactant-stream pre-cooling was judged useful as an auxiliary capability, but insufficient and risky as the primary cold-environment system.

Cold jacket surrounding the chamber

A dedicated cold jacket (e.g., a double-wall sleeve or clamshell) allows cooling power to be applied closer to the vessel than an environmental room, potentially reducing cool-down time. Nevertheless, achieving a stable, uniform −50 °C wall temperature with a jacket is non-trivial because coolant distribution, contact resistances, and local heat leaks at ports and fasteners create azimuthal and axial gradients. If LN_2 is used as the coolant, two-phase behaviour and flashing can make temperature control difficult unless a carefully engineered conditioning loop is implemented. A jacket also competes for volume around the chamber, potentially interfering with sensor mounting, assembly/cleaning access, and high-voltage routing, and it can worsen frosting around penetrations due to colder external surfaces. Finally, a bespoke jacket increases manufacturing effort and introduces additional leak paths and failure modes. For these reasons, a cold jacket was considered feasible but comparatively high-risk in terms of repeatability, integration, and lifecycle effort.

Environmental chamber (temperature-controlled test chamber)

Placing the entire combustion-chamber assembly within an environmental chamber provides the most controlled and repeatable cold boundary condition because the air temperature can be regulated with a closed-loop controller and maintained for long soak periods. In contrast to cryogenic baths or local jackets, the environmental chamber uniformly cools all external surfaces and interfaces, thereby reducing uncontrolled gradients caused by the selective cooling of only part of the hardware. This approach also preserves the integrity of the mixture: the chamber can be filled and mixed under steady thermal conditions, without introducing additional cold gas streams or cryogenic boil-off into the test volume. Instrumentation integration is significantly simplified since sensor cables and high-voltage lines can be routed through standard feedthroughs, thereby avoiding the risk of direct cryogenic contact with connectors. Importantly, environmental chambers can be operated with dry air or dry nitrogen purging to suppress condensation and frost on cold surfaces, improving the electrical reliability and repeatability of ignition hardware and sensors.

Figure 4.16: C900 environmental chamber[16]

For the above reasons, and because it best satisfies temperature repeatability, sensor compatibility, and non-intrusiveness to the mixture, the environmental chamber was selected as the baseline cold-environment system for testing at −50 °C. The main requirements for its selection, in addition to the ability to reach the desired temperature, are that it must provide sufficient space around the chamber to allow airflow and cooling, and withstand the weight of the combustion chamber. Additionally, it must have lateral entries for cabling and wiring, as this is necessary for combustion to occur. A feasible option is the TotalTemp Technologies C900 temperature environmental chamber, shown in Figure 4.16, which has an internal volume twice that of the chamber, thereby ensuring sufficient space for counterflows and for making adjustments to the combustion chamber as needed. This works by injecting liquid nitrogen until the temperature set in the control system is reached. Naturally, the ingested nitrogen also needs to be released, which is purged towards the vent line. For these reasons, this is the final system selected. [16]

4.4. Full experimental setup

The full experimental setup is shown in Figure 4.17. In Figure 4.17a and Figure 4.17b, the setup including only the new non-optical chamber on the new optical table is shown. This is used at the beginning to verify, as any leak or issue may significantly damage the environmental chamber. In Figure 4.17c, the power supplies used with the ignition coil and to excite the chamber piezoresistive sensor are shown, along with the National Instruments modules used to read the sensor outputs and send digital signals to the coil and oscilloscope.

The injection pressure sensors had to be updated to include two additional sensors capable of handling the higher pressures, as shown in Figure 4.17d. However, as shown by its absence in the pictures, the gas booster did not reach the laboratory in time to be used, so the actual pressure in the system remains limited, limiting the combustion analysis. Figure 4.17e shows the addition of a downstream pressure regulator, which, in the high-pressure experiments, reduces the downstream pressure to a maximum of 2 bar, which is easily handled by the vent system. Finally, the last three pictures show the environmental chamber with the LN2 Dewar attached.

---

[16] https://www.totaltemptech.com/model-c900-cryogenically-cooled-temperature-chamber/ [Accessed on 01/08/2025]

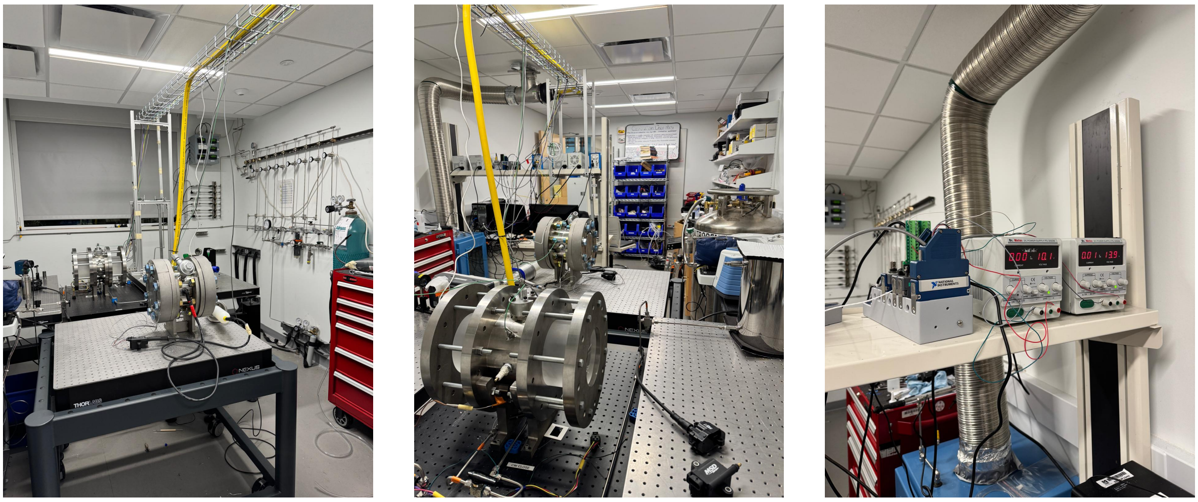

(a) Lab front view with non-optical chamber

(b) Lab rear view with non-optical chamber

(c) Power supplies and NI modules

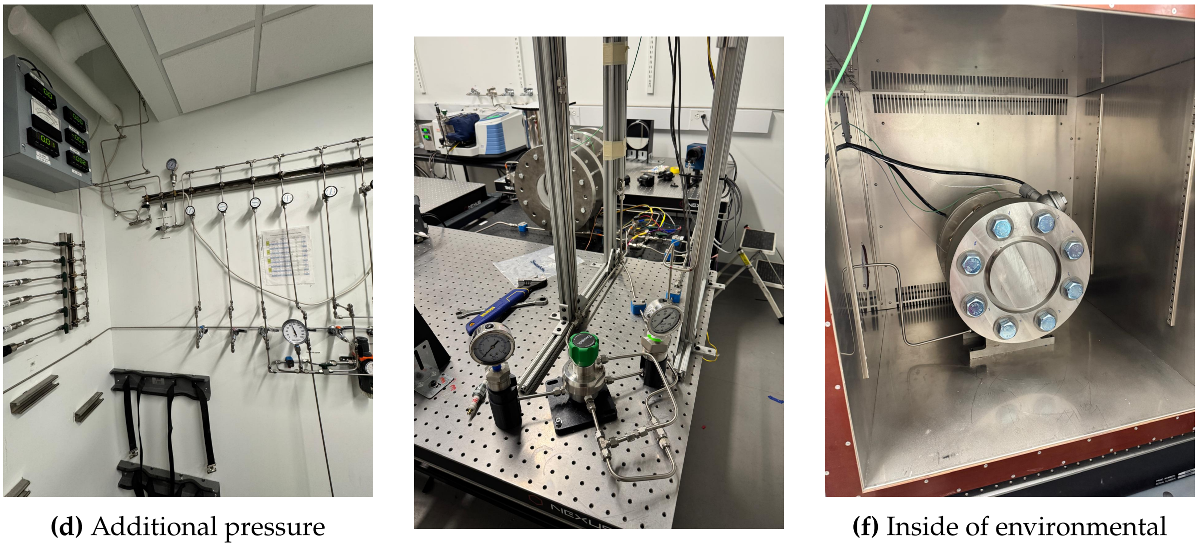

(d) Additional pressure transducers and gauges

(e) Exhaust line

(f) Inside of environmental chamber

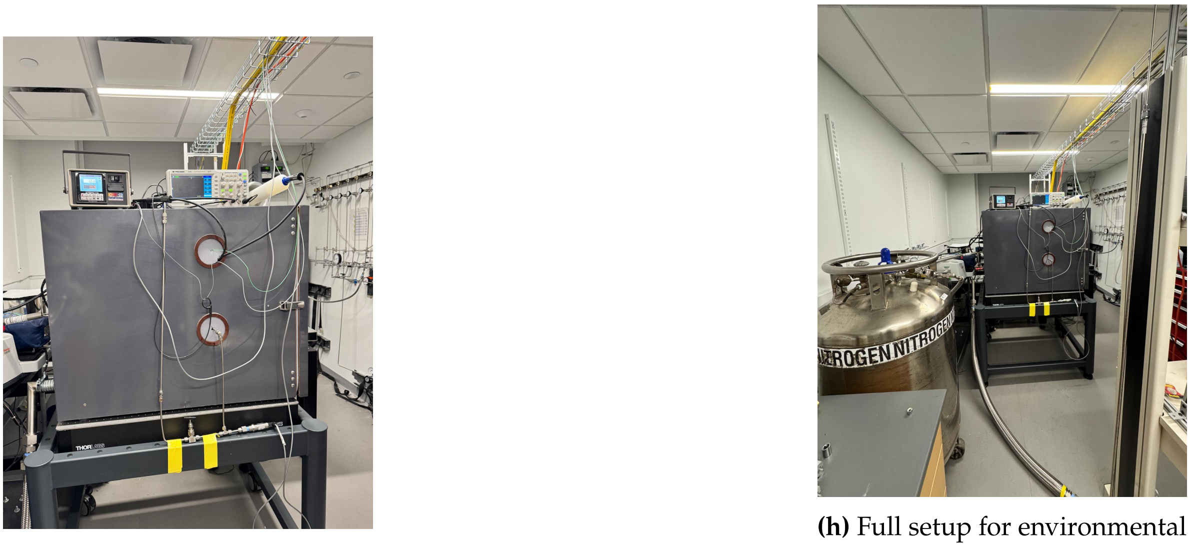

(g) Environmental chamber

(h) Full setup for environmental chamber

Figure 4.17: New setup

4.5. Requirement compliance

With the completed design, it is necessary to return to the requirements listed in Table 4.1 to verify their compliance and ensure a safe, redundant and usable design.

Table 4.8: Compliance matrix for combustion chamber requirements (from Table 4.1).

ID	Compliance	Justification
REQ-01	A: compliant; T: pending	Structural sizing of flanges/bolting/gasket using ASME/Roark-based formulas; minimum thickness confirmed by FEA. Testing is left for the future due to the absence of a pressurisation system. Refer to Section 4.1.
REQ-02	Compliant; sealing at low-T may limit vacuum	Material selection includes low-temperature toughness (DBTT) as a high-importance criterion; the experimental campaign includes operation at −50 °C. Refer to Chapter 6.
REQ-03	A: compliant, I: pending	Material choice compliant with fatigue and safety factors included. Inspection should be carried out every 1000 cycles to confirm the analysis. Refer to Subsection 4.1.1
REQ-04	Compliant	Mechanical architecture includes a bolted cover flange that enables disassembly for cleaning, insert replacement, and internal inspection. Refer to Subsubsection 4.1.2.2.
REQ-05	Compliant (by design rationale)	The material trade-off explicitly accounts for ammonia corrosion resistance and hydrogen embrittlement resistance; SS 2205 is selected for these constraints. Chamber filling with either substance shows no damage. Refer to Chapter 6.
REQ-06	Compliant	Instrumentation choice includes a high-range in-chamber pressure sensor (0 −−1000 bar) that exceeds MAWP and a packaging/insert strategy that enables installation. Refer to Subsection 4.2.2 and Subsection 4.1.3.
REQ-07	Compliant	The temperature measurement architecture uses an RTD and thermowell; the sensor temperature range meets the cold-start requirement (down to at least −50 °C), and the thermowell integrates with the chamber geometry. Refer to Subsection 4.2.2 and Subsection 4.1.3.
REQ-08	Compliant	Ignition concept selection targets symmetric ignition; the chamber includes two igniter locations/ports to enable symmetric placement. Refer to Subsection 4.2.1 and Subsection 4.1.3.
REQ-09	Compliant	Stress-concentration mitigation: fillets/chamfers and insert/port integration choices to reduce local stress risers and protrusions. Refer to Subsection 4.1.3.
REQ-10	Ambient: compliant; Low-T: partially limited	Vacuum at 100 mTorr is achieved at ambient temperature; at low temperatures, thermal contraction introduces microleaks, limiting the achievable vacuum level. Refer to Chapter 6.
REQ-11	A: compliant; T: pending	Chamber diameter selected using turbulence-onset reasoning. Full validation at the highest-pressure envelope is only possible once the system reaches the intended operating range. Refer to Subsubsection 4.1.2.2.

5

Experimental methodology

Once the design phase was completed, the work transitioned from concept and specification to the execution of the experimental programme. The first stage consisted of constructing and commissioning the updated setup, starting with the mechanical integration of the combustion chamber and its interfaces. This was followed by implementing the required modifications to the filling and exhaust lines, including any additional branches, fittings, and provisions necessary to integrate the additional instrumentation and components. In parallel, the data acquisition chain was assembled and tested to ensure that all sensors could be read reliably at the intended resolution and sampling characteristics. Only after these integration steps were completed could the facility be considered ready for commissioning activities that establish basic operational readiness.

In practice, the construction and commissioning phase was significantly impacted by procurement constraints and supplier lead times. Delays in receiving critical components affected the order in which subsystems could be integrated and reduced the extent to which assembly tasks could be parallelised. As a result, several integration steps had to be re-sequenced, with temporary configurations used where feasible to continue progress without compromising safety or measurement quality. These disruptions extended the timeline required to reach a stable, fully operational configuration and, consequently, delayed the start of the planned experimental campaign, reducing the total number of achievable experiments. Moreover, a fundamental component, the gas booster, was particularly affected by these delays and did not reach the laboratory in time to be used for any experiment. Therefore, the tests will be limited to varying initial temperatures, which is still of high interest for the study of ammonia behaviour. Thus, once the setup was fully assembled and verified, the experiments defined in the test plan could be executed systematically to generate the dataset used in the subsequent chapters.

5.1. Experimental procedure

The experimental procedure is explained in detail here, including operational steps, to ensure clarity and reproducibility. For consistency, the procedure is structured into three main phases: filling and mixture preparation, execution and control of the experiment, and exhaust release and evacuation. This breakdown reflects the actual sequence followed during testing and provides a clear procedure for future repeatability of this work. This is explained using the fluid diagram shown in Figure 5.1, which is read from right to left, starting with the gas bottles. Because the facility handles hazardous gases, both flammability and toxicological exposure

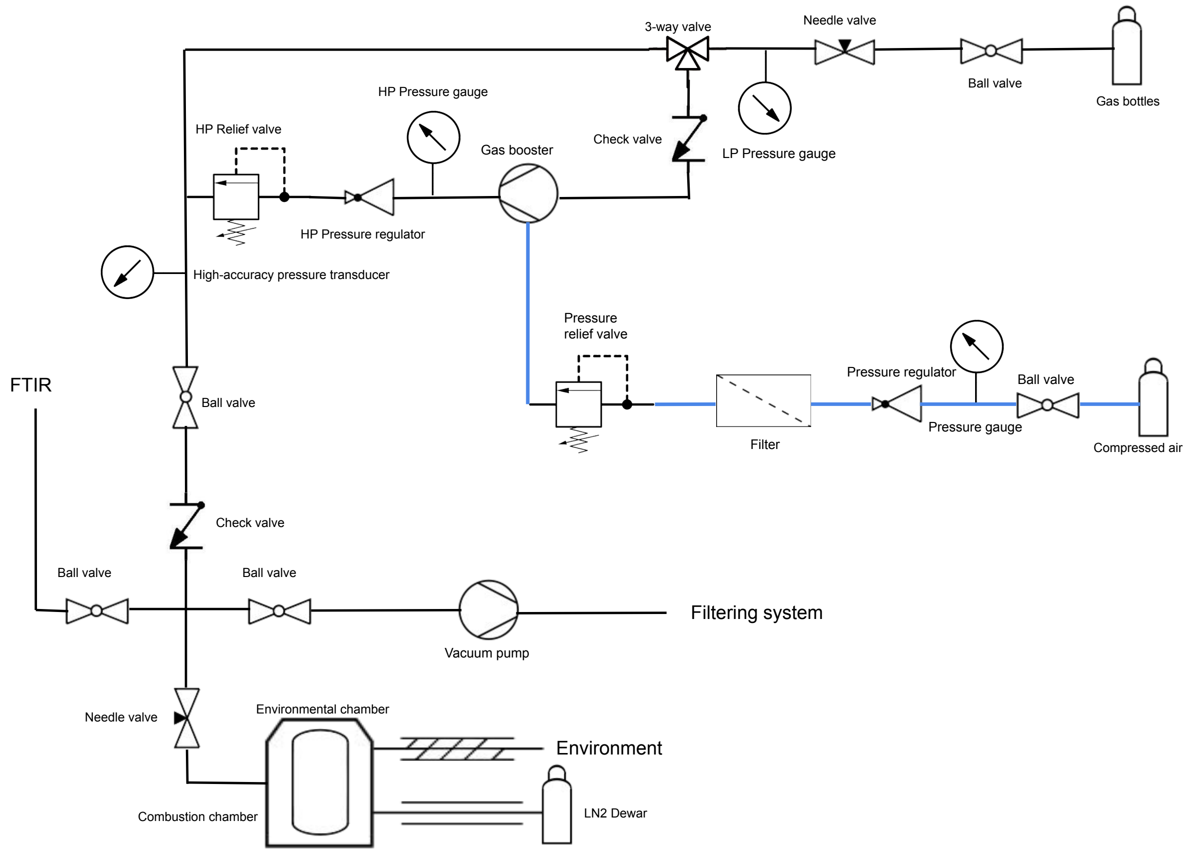

Figure 5.1: Fluid diagram (the blue line represents the compressed air line, the black the working fluid line)

risks must be considered throughout the procedure, and certain safety requirements must be in place. Thus, the safety measures and safeguards implemented for normal operation and upset conditions are also presented here, alongside the operational steps, to complete the workflow and make it auditable.

5.1.1. Procedure

Filling and mixture preparation

During this phase, the objective is to fill the chamber with the mixture at the required equivalence ratio and initial pressure. Mixture preparation is performed using the partial-pressure method, in which each reactant is introduced into the chamber to its target partial pressure relative to the desired initial total pressure:

$$Y_i = \frac{n_i}{n_T} \tag{5.1}$$

$$p_i = Y_i p_0 \tag{5.2}$$

where

Y_i	Molar fraction of a reactant [-]
n_i	Number of moles of a reactant[mol]
n_T	Total number of moles of reactants [mol]
p_0	Initial pressure
p_i	Partial pressure of a reactant

This method enables direct control of the equivalence ratio and blend fraction through measured pressure increments, provided that the initial evacuation is sufficiently repeatable. Because the mixture is prepared in a fixed-volume vessel, the method also ensures that the final mixture state is defined by the measured initial pressure and temperature, which are later used directly in data reduction.

For this method to work, the initial pressure in the chamber before filling should be as close as possible to 0 Pa. Therefore, the chamber and lines are evacuated prior to filling, with the evacuation setpoint chosen as a compromise between pump capability and minimising residual gas effects on mixture composition, set at 200 mTorr. For the first gas, since the chamber and line are both evacuated, the filling can be done with the chamber open. However, for the rest, the chamber must start from the closed state. In fact, during filling, the supply-side line is pressurised above the target chamber partial pressure, to approximately 6 times the partial pressure, to maintain a positive pressure differential and prevent backflow through the common injection/exhaust line. This is because, in the ECRL laboratory, it has been experimentally proven that this pressure level prevents backflow and, after opening the chamber and flow expansion, ensures that the pressure remains below the required level, that is, the sum of the existing reactants' pressures and the amount that can be topped up. This ensures that backflow probabilities are minimised and that the filling procedure is faster and more precise.

Thus, the gases are filled, in order of increasing pressure, from their respective wall outlets attached to the gas bottles. The filling line is common to all gases, meaning that all outlets converge into a single line, and the release is controlled by two valves: the outlet valve and a high-precision needle valve that accurately controls the outlet pressure. This is achieved via the 5 PX409 transducers, each one calibrated for a specific pressure range, going from the lowest at 0 bar to 1 bar to the highest at 0 bar to 172 bar. Naturally, the ones with a larger calibrated range do not have the same precision at a specific point as those with a more limited range, as they need to maintain the same accuracy over a larger range.

Once all the gases are filled and the chamber is at its required initial pressure, the chamber is closed, and the line is evacuated to prevent any reaction if anything goes wrong. The filling line is then closed to prevent reaction products from entering it once the chamber is opened after the reaction, and the purge line is opened to ensure that, if the chamber valve were to fail, the gases would be purged directly. To ensure that no high-pressure gases are released in the purge line and that no ambient gases enter in the reverse direction, a downstream pressure regulator followed by a check valve is placed.

Experiment execution and control

As soon as the chamber is completely filled and closed, the LabVIEW experiment control environment runs, starting a 3-minute timer before any ignition can occur. This is because sufficient time must be allowed for complete mixing; otherwise, stratified combustion may occur, invalidating the results. Ignition is achieved via the IGN1-A coil, which receives voltage from a PS-305D power supply, providing up to 32 V. Before ignition, the oscilloscope must be set to trigger mode with the trigger set to voltage increase, and the Kistler sensor's measuring mode must be activated on the charge amplifier. If the environmental chamber is used, the desired temperature is set in the controller, and the additional time needed to bring the combustion chamber to temperature must also be allowed. Thus, once mixing and the desired temperature are achieved, the ignition signal can be sent, which closes the coil circuit and generates a spark if the provided energy is sufficient.

If combustion occurs, as evinced by the rise in the pressure seen in LabVIEW, then the experiment can be considered completed, and the results are saved to the output file; otherwise, the energy is increased by varying the input voltage and coil dwell time until ignition occurs or until the power supply cannot provide more. If the maximum energy is reached and no ignition still occurs, the system is evacuated, and the electrode gap is increased to increase the electric field strength. If this does not work either, it is reduced to decrease heat loss at the electrodes, which is then the cause of the faulty ignition. Once ignition has occurred, the power supply is turned off to avoid accidental sparks, especially when operating on the chamber structure.

Exhaust release and evacuation

Once the mixture needs to be released from the chamber, it first passes through the purge line, which is directly connected to the building roof. An industrial filtering system then cleans the gases to prevent excessive emissions of soot, NOx, and other pollutants. Once the pressure in the purge line reaches atmospheric pressure, the gases can no longer flow through it, as the pressure differential is 0. Once this occurs, the exhaust three-way valve is turned from the purge direction to the vacuum direction, where the gases are pulled by the vacuum pump and pushed into the purge line. It is imperative that the gases pass through the pump only when at atmospheric pressure, as the pump is not rated for any pressure above this. Any accidental error may cause the pump to rupture, releasing gases into the room. The evacuation can be considered complete at a pressure of 200 mTorr, which is normally achieved after approximately 20-30 minutes. This is a trade-off between accuracy and experimental time. Naturally, if the evacuation is left on indefinitely, the pressure would drop until the micro-leaks in the system become dominant, but this would considerably reduce the number of experiments that can be performed.

5.1.2. Safety control

As mentioned, specific safety procedures must be in place to ensure that any system failure or human error has minimal impact and that operator safety is maintained throughout the testing process. The overall safety philosophy adopted here is based on isolation and redundancy to prevent unintended flow paths, controlled venting and evacuation to avoid the accumulation of hazardous mixtures, and rapid and unambiguous abort capability to prevent ignition when abnormal behaviour is detected. Accordingly, the facility incorporates multiple physical barriers, together with procedural steps, to reduce the likelihood that an incorrect valve configuration or an unexpected component failure can escalate into a hazardous event. Given that the system handles gases that pose both flammability and inhalation hazards, these measures are considered part of the normal operating workflow rather than exceptional provisions. The key safeguards implemented are detailed below in the same sequence as they become relevant during operation.

A first layer of protection is provided by system segmentation with ball valves, which are kept closed when the corresponding subsystem is not in use. This ensures that any leak or malfunction remains localised and that an incorrect valve actuation does not immediately propagate to other parts of the facility. A representative example is the secondary safety valve on the vacuum line, which prevents an erroneous configuration from directing a higher-than-atmospheric-pressure discharge into the vacuum pump when the exhaust valve is inadvertently set to vacuum, thereby preventing pump damage. In addition, keeping the purge line open during the experiment provides a deliberate low-resistance relief path, so that a failure in the chamber-closing system cannot result in pressure accumulation in the line. Together, these measures mitigate the consequences of incorrect valve states and provide passive protection, eliminating the need for active control logic.

A second layer is the ability to abort the test rapidly and prevent ignition if a fault is detected or suspected. The experiment can be interrupted via an Emergency Stop button integrated in the LabVIEW interface, which prevents the spark command from being issued and therefore stops ignition at the source. For additional physical redundancy, ignition can be inhibited by turning off the power supply, rendering the ignition coil inert and eliminating the possibility of unintended sparking caused by software or signal errors. An ignition delay is also implemented to ensure that operators are at a safe distance during the ignition attempt. This approach ensures that abort capability is both immediate (at the software level) and enforceable (at the hardware level), which is critical for a system where ignition is the primary escalation mechanism. Operationally, this also provides a clear decision point: any anomaly detected during filling, stabilisation, or pre-ignition checks can be handled by inhibiting ignition first, and only then proceeding to venting and evacuation in a controlled manner.

Because ammonia inhalation poses a serious health hazard, detection and containment measures are implemented specifically to prevent toxic exposure. The ammonia cylinder is stored in a hermetic container to minimise the risk of a minor leak resulting in immediate operator exposure, and two ammonia detectors are installed, one near the tank and one closer to the operator, and both are positioned as high as practicable because ammonia is lighter than air and therefore tends to accumulate in elevated regions. The detectors are configured to alarm at concentrations as low as 10 ppm, since even relatively small concentrations can be harmful over short exposure periods, and early warning is therefore essential. In addition to detection, the ammonia container is directly connected to the purge line, so any tank leak is automatically vented from the system. The room itself has roof vents that allow air to circulate and help quickly disperse leaks.

Finally, electrical and backflow precautions are implemented to reduce ignition sources and prevent unintended migration of gases or products. Cables and wires are routed away from potential leak paths to minimise the possibility that a leak produces an ignitable mixture in proximity to electrical connections, which is particularly important when hydrogen is used. All valves leading to wall outlets remain closed during experiments to provide redundancy and to prevent combustion products or reactive mixtures from reaching the gas supply infrastructure. All these systems together contribute to ensuring the safety of the system under most events, although for future experiments in high-pressure conditions, some additional safety measures will need to be implemented, such as relief valves, which were already designed and shown in Figure 5.1, but not implemented due to the absence of the gas booster.

5.2. Experimental matrix

Now that the experimental procedure has been explained, the actual experimental matrix, which provides a comprehensive overview of all the experiments to be run, is presented in Table D.1. This is divided into four main parts, depending on the objective that the specific experiment aims to achieve. Since the work is based on a new design, the first step is naturally verification and validation (V&V) of the setup, as without it there can be no confidence in the research data. Since the setup is quite complex, especially when considering the use of the environmental chamber, the validation is separated into two parts: one at standard pressure and conditions, intended to validate the chamber and lines, and one at subzero temperatures, intended to validate the subzero behaviour and setup. It is worth noting that validation is performed at two equivalence ratios: stoichiometric and lean, as these are the most common conditions encountered in operation and the literature. This is for time's sake, but validating these two data points can corroborate a complete validation, as they are the most critical and useful.

The verification process consists of checking that the setup functions as intended, ensuring there are no leaks, that all sensors are reading correctly, that the chamber withstands the pressure, and that ignition is achievable. For this step, as for the initial validation, methane is used because it is well-established and has many literature values to compare against. Moreover, since the other fuel of interest, apart from ammonia, that is, hydrogen, has some unique combustion characteristics, such as preferential diffusion, where its molecules diffuse much faster than those of other gases, methane exhibits a more neutral and predictable nature. This enables a more robust combustion behaviour, simplifying the validation process. Once the validation at STP is achieved, methane is no longer used, as it is not of interest in this work, which aims to investigate novel green alternatives.

Then, validation is also performed with H_2, as these are the fuels of interest and will be used for the remaining experiments. Once the one at STP is confirmed, the pressure is increased to the possible limits of the setup without the gas booster to investigate whether this behaviour persists at higher pressures. Afterwards, the aim is to move to low temperatures. For this, steps of about 15-20° C are taken until reaching the point of minimum temperature of interest, that is, -50° C, which allows for the creation of curves in the plots. Once this is confirmed, sufficient confidence is present to ensure the accuracy of the research data.

For the research dataset, an additional equivalence ratio was introduced for each combustion condition so that each condition is represented by three ϕ points rather than two. This provides one additional anchor point per condition, enabling interpolation of the laminar burning velocity trend with greater confidence and reduced sensitivity to single-point variability. The same experimental matrix logic was applied to the low-temperature validation campaign for both neat ammonia and ammonia blends, ensuring that differences observed between fuels are not driven by inconsistent coverage in ϕ. For the blended-fuel cases, two hydrogen enrichment levels were selected: 20% and 30% H_2 in the fuel mixture. These additions were chosen specifically to probe the regime of relatively low hydrogen enrichment, where even modest H_2 fractions are expected to modify reactivity while still remaining representative of practical "ammonia-rich" blends. The selection of the 20% and 30% points was further motivated by the availability of literature data at comparable compositions, enabling direct benchmarking, and by the fact that the two levels are sufficiently close to allow the incremental impact of a ~10% change in H_2 addition to be assessed in a controlled manner.

It is to be noted that each experimental condition is run three times to guarantee that repeatability errors are considered and evidence the presence of outliers. In fact, it is possible that, with a single measurement, this might be an outlier, and therefore, it would be difficult to understand why it is not yielding the expected result. By having three measurements per condition, this probability is lowered, and the validity of the method and setup is increased by providing the repeatability errors. For an accurate comparison, the three experiments should be run consecutively to avoid differences in conditions. The full experimental matrix is presented in Appendix D for reference, while a summary of the considered experimental parameters is given in Table 5.1.

Table 5.1: Range of experimental parameters

Parameter	Range
Fuels	CH_4, H_2, NH_3, NH_3-H_2
H_2 blend fraction	20%, 30%
ϕ	0.8, 1, 1.2
p_0	1, 2, 5
T_0 [° C]	21, 0, -15, -30, -50

5.3. Data acquisition

The data produced by the sensors in the chamber must be read by specific components that can digitise the signal and translate it into usable data. The acquisition chain is implemented in a LabVIEW environment that records sensor outputs and ignition signals, and the resulting time series are exported for post-processing. Post-processing is performed in MATLAB to apply calibration curves, filter noise where required, and align signals in time relative to the ignition trigger. For the given sensors, four modules are required, each with a specific purpose and handling different types of inputs and outputs. They are connected to a common CDAQ module, which is the interface between the software and the physical signals. The modules are explained here:

- NI-9215: Kistler 603CAA

  The Kistler 603CAA is a piezoelectric (charge-output) pressure sensor used together with a charge amplifier, the Type 5018A, that provides an analogue V-level signal proportional to dynamic pressure. The NI-9215 is required to acquire this high-bandwidth analogue voltage signal with sufficient sampling capability to resolve the fast in-chamber pressure transient during ignition and flame propagation. In practice, it provides the analogue input interface between the charge amplifier output and the DAQ system. To make sure that the combustion event is read with enough accuracy, a sampling rate of 10 000 Hz is chosen, as this would mean that a sample is taken every 0.1 ms, which is considered enough for a combustion event running for a few milliseconds. A higher sampling rate would considerably increase computation and the output file size while, at the same time, picking up more noise as more high-frequency content is captured, thereby increasing the visibility of electrical and EMI pickup noise, among others.

- NI-9217: OMEGA PR-12-E-3-100-1/4-2

  The PR-12-E probe is a Pt100 RTD, which must be excited and measured through a dedicated RTD measurement circuit to accurately infer resistance (and therefore temperature). The NI-9217 is required because it provides built-in RTD excitation and signal conditioning, with appropriate wiring support, and direct conversion to temperature, enabling stable, accurate temperature acquisition. Due to the sensor's very slow response time, a sampling rate of 10 Hz is used.

- NI-9401: ignition coil and oscilloscope

The ignition system requires a deterministic digital command to enable/disable the spark and to implement timing delays and interlocks within LabVIEW. The NI-9401 is required because it provides high-speed digital I/O for generating the ignition trigger (or gate) signal and for handling related TTL-level timing/trigger signals. It is also used to exchange synchronisation signals with the oscilloscope (e.g., trigger out/in) so that ignition events and high-speed measurements can be time-aligned.

- NI-9209: OMEGA PX419-1.5KAV

  This pressure transducer also needs an external excitation to provide a conditioned analogue voltage output in the range 0 V to 10 V, proportional to absolute pressure. The NI-9209 is required to excite the sensor and acquire the quasi-static analogue voltage signal with adequate resolution and low noise, enabling the accurate determination of the initial pressure and slowly varying pressure levels within the chamber.

The user interface of the LabVIEW structure built to control and run the experiment is shown in Figure 5.2. This displays the traces of all quantities visually, providing an immediate indication of whether ignition occurs. The output consists of two files: one containing the dynamic pressure, absolute pressure, and temperature traces, and another containing the voltage and current data. Due to differences in sampling rates, the dynamic trace is used as the constraint, as it has the harshest sampling requirement, to ensure that all curves are matched to the same time. The temperature and absolute pressure values are interpolated for the missing time steps.

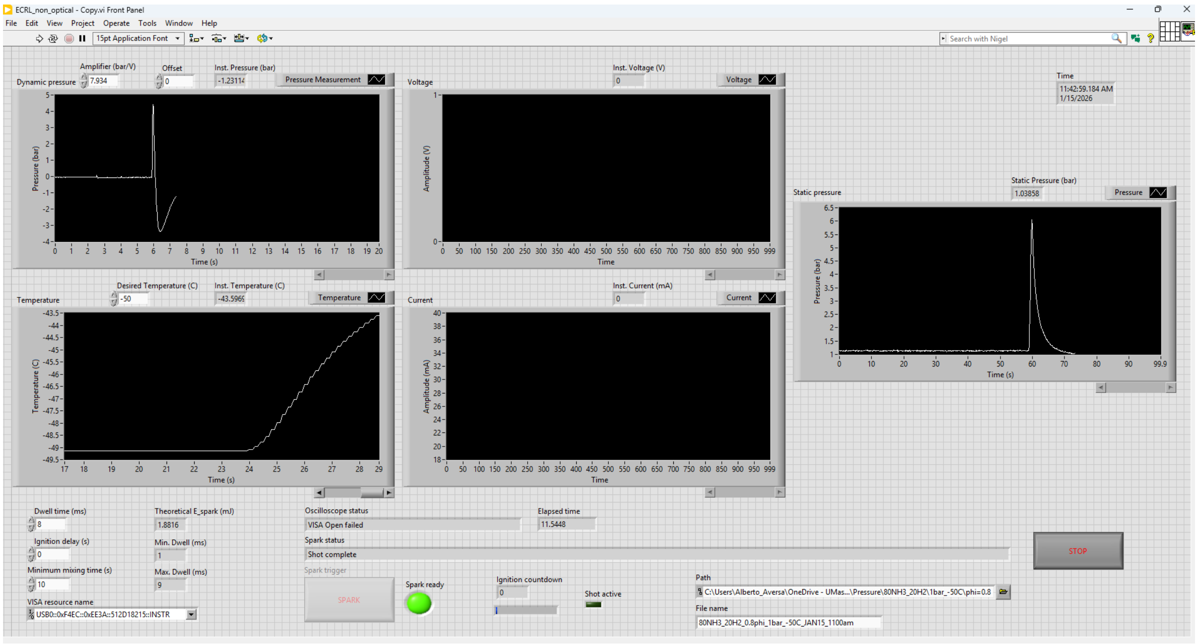

Amplifier (bar/V) 7.934
Offset 0
Inst. Pressure (bar) -1.23114
Inst. Voltage (V) 0
Time 11:42:59.184 AM 1/15/2026
Static Pressure (bar) 1.03858
Desired Temperature (C) -50
Inst. Temperature (C) -43.5969
Inst. Current (mA) 0
Dwell time (ms) 8
Ignition delay (s) 0
Minimum mixing time (s) 10
VISA resource name USB0::0xF4EC::0xEE3A::512D18215::INSTR
Theoretical E_spark (mJ) 1.8816
Min. Dwell (ms) 1
Max. Dwell (ms) 9
Oscilloscope status VISA Open failed
Spark status Shot complete
Spark trigger
Spark ready
Ignition countdown 0
Shot active
Elapsed time 11.5448
Path C:\Users\Alberto_Aversa\OneDrive - UMas...\Pressure\80NH3_20H2\1bar_-50C\phi=0.8
File name 80NH3_20H2_0.8phi_1bar_-50C_JAN15_1100am

Figure 5.2: LabVIEW user interface

5.4. Data processing

Once the sensor data has been extracted and translated to the physical measurements, the LBV and MIE can be extracted from the saved traces. The methods used to calculate them will be described here.

5.4.1. LBV method

A propagating spherical flame experiment involves centrally igniting a quiescent, premixed, homogeneous mixture inside a vessel, resulting in an approximately spherical flame front that expands outward. In its most general form, the spherically expanding method relies on tracking either the flame radius history, $R_f(t)$ (optical route), or the pressure history, $p(t)$ (non-optical route) during the event, from which the velocity can be extrapolated. This is defined as the propagation speed of an adiabatic, unstretched, planar premixed flame relative to the unburned mixture, and therefore, it is not measured directly in a closed vessel configuration.

CVM propagating spherical flame method

In the constant-volume variation of this method (CVM), the analysis is intentionally performed at large flame radii, where the pressure rise becomes significant, and the pressure trace $p(t)$ becomes the primary measured quantity used for data reduction. A central practical implication is that CVM does not intrinsically require optical access, making it compatible with thick-walled high-pressure vessels and reducing geometric constraints on chamber design. However, the main issue with the lack of optical access is the modelling assumptions used to convert $p(t)$ into the burnt mass fraction and the flame-front history, which directly affect the inferred burning velocity. For this reason, Faghih and Chen[24] emphasises that careful selection of the usable pressure interval is essential, because the earliest part of the trace is more sensitive to ignition disturbances and flame stretch, whereas later stages can be corrupted by confinement and instability effects as the flame approaches the wall.

This method includes some basic assumptions, for which various corrections exist and can be applied after the value is calculated:

- Spherical, smooth flame front
- Uniform vessel pressure
- Chemically frozen unburned mixture compressed isentropically
- Negligible buoyancy and radiation in the baseline formulation

It is essential to consider these assumptions, as their neglect can lead to significant discrepancies in the velocity calculation, particularly for ammonia, which, due to its weak flame nature, is highly susceptible to radiation and heat loss. Additionally, at lean equivalence ratios, buoyancy becomes quite significant and must be accounted for, as it significantly reduces flame speed.

In CVM, the fundamental idea is that the measured pressure rise maps the mixture's thermodynamic evolution and, through conservation laws, the progress of combustion. The analysis introduces the burnt mass fraction $x_b = m_b/m_0$, which partitions the total mass m_0 into burnt and unburnt contributions, enabling the flame radius to be linked to pressure via the unburnt-gas isentropic compression assumption. Under the standard CVM assumptions, the flame radius can be expressed in terms of x_b and p/p_0, and the burning velocity can be written in a form proportional to $\mathrm{d}x_b/\mathrm{d}t$ multiplied by a pressure- and geometry-dependent factor. Faghih and Chen[24] present a widely used formulation in which S_u is obtained once $x_b(t)$ and $x_b(p)$ are known, making the determination of x_b the central step in the CVM reduction chain. Therefore, with a measured pressure trace, a model can reconstruct the combustion progress and flame kinematics that are not directly observable in a pressure-only configuration. This modelling step is also where the dominant methodological differences between "two-zone" and "multi-zone" approaches arise.

The two-zone model idealises the vessel contents as two uniform regions (burnt and unburnt) separated by an infinitely thin flame front, yielding a compact set of conservation equations but neglecting temperature gradients in the burnt gas. In contrast, a multi-zone (shell) model divides the burnt region into multiple concentric shells, each assumed to be uniform internally but allowed to differ from the others, thereby capturing the fact that burnt gas formed earlier can be compressed for a longer period and thus reach different temperatures than burnt gas formed later. Faghih and Chen[24] highlight that, while temperature gradients in the burnt region may be modest at very early times, they can become substantial as the flame propagates outward, and this can introduce non-negligible error in the burnt mass fraction if a two-zone approximation is used throughout. Because x_b is a key intermediate variable for converting $p(t)$ into S_u, reducing bias in x_b is directly beneficial to the fidelity of the inferred burning velocity. Multi-zone methods also offer a clearer pathway to incorporate additional physics, such as wall heat transfer layers or radiation models, when required by slow-burning mixtures or strongly radiating products.

Constant Volume Data Reduction Tool (CVDART)

Initially, an ad hoc code was developed to simulate a two-zone model, but it yielded inconsistent results. For this reason, it is decided to switch to the Constant Volume Data Reduction Tool (CVDART), developed at NIST[93], which is a dedicated post-processing framework for spherical constant-volume flame experiments, where the primary measured input is the pressure trace $p(t)$. This is made possible thanks to the kind assistance of the creator, Michael J. Hegetschweiler. In CVDART, the user supplies the initial mixture state and composition, imports the pressure trace, selects a smoothing and cropping strategy, and then runs a data-reduction model (two- or multi-zone) to obtain the time histories needed to compute S_u. The tool explicitly reports, among other quantities, the burnt mass fraction $x_b(t)$, its time derivative $\mathrm{d}x_b/\mathrm{d}t$, the unburned temperature $T_u(t)$, the unburnt specific-heat ratio $\gamma_u(t)$, a reconstructed flame radius $R_f(t)$, and the resulting laminar burning velocity $S_u(t)$, which can subsequently be represented as $S_u(p, T_u)$ over the experimental trajectory.

In the two-zone formulation described by NIST, the model solves coupled mass and energy conservation equations at discrete pressure levels for the unknown burnt-gas temperature and burnt mass fraction, using chemical equilibrium calculations to estimate burnt-gas composition. The laminar burning velocity is then computed from R_f, γ_u, x_b, and the experimentally measured pressure, via a closed-form expression that uses $\mathrm{d}x_b/\mathrm{d}t$ and the isentropic compression factor. In the multi-zone MECT model, the vessel is discretised into radial shells, and three shell types are defined (unburnt shells, the currently burning shell that contains the flame front, and previously burnt shells), with no mass or energy exchange between shells; burnt shells can therefore develop distinct thermodynamic states due to their different compression histories. The algorithm advances through a pressure loop, compresses the unburned shells isentropically, updates the burned shells (including optional radiation and dissociation modelling), solves the conservation equations on the burning shell, and reconstructs the flame radius from the shell-wise burnt mass fraction distribution. The HTDR model follows the same shell treatment but differs in how the progression is discretised while maintaining the same overarching objective of reconstructing the flame history and $S_u(p, T_u)$ from pressure-only data. Here, one shell is burnt per loop, conserving mass but not energy, so it is only accurate at high shell numbers.

Since the multi-zone model can provide a significantly higher degree of accuracy by accounting for radiation and dissociation, which already addresses some of the model's criticalities, this model is chosen. In the CVDART tool, the method does not significantly increase computational time, so there are no significant drawbacks to its use. Between the MECT and HTDR methods, the former is chosen, as it is more grid-insensitive and can converge with fewer shells while guaranteeing mass and energy conservation. If either is not complied with, then the method stops with an error. Within this method, the option to select the grid for shell division is provided, allowing either equal-radius or equal-volume shells. Here, the main difference lies in the wall resolution and its effects on the flame, which is particularly relevant for ammonia. The equal-volume method converges more rapidly at larger radii and can therefore model radiation more accurately.

Determination of bounds

Even with a robust multi-zone reduction model, the inferred burning velocity is only meaningful over the portion of the trace that satisfies the CVM assumptions. Bariki et al.[94] explicitly frame the pressure-rise method as being limited at low pressures by stretch/ignition effects and at high pressures by the onset of combustion instabilities, and they propose a practical, pressure-trace-based criterion to define the usable upper bound even in spherical chambers without optical access. Specifically, they recommend defining a lower pressure limit, P_1, to exclude early-time artefacts, and an upper limit, P_2, to exclude data beyond the onset of instabilities from fitting or extrapolation.

In their recommended workflow, the pressure should first be smoothed with a moving window average, after which, P_1 can be established by checking overlap with an independent reference (when available), and, if overlap is not available, a conservative criterion of using approximately twice the initial pressure as the lower bound is suggested in line with prior literature practice. For the upper bound, propose plotting a relative pressure-change metric, $K_p = \frac{1}{p}\frac{dp}{dt}$, against p/p_0 and selecting P_2 at the inflexion point of that curve, which is interpreted as signalling the onset of instability effects in the pressure-rise-derived flame speed. This K_p-inflection approach is presented as particularly useful when optical flame imaging is not available to directly detect cellularity or other manifestations of instability.

This method can be implemented in CVDART after some initial pre-conditioning of the pressure trace in MATLAB, including the proposed smoothing. For both point determinations, the NIST software provides a help graph that shows the variation in the calculated velocity as a function of the two bounds. These need to be adjusted, if necessary, to ensure they remain in the "plateau" zone for the change, meaning that slight variations in these bounds do not cause a significant change in the velocity calculation, as shown in Figure 5.3.

Otherwise, it would mean that they are being taken in zones where the flame is unstable. Once the bounds are selected and the velocity is calculated, the code provides the variation of velocity with respect to pressure and temperature within the bounds. To map the velocity back to a reference condition, the variation needs to be correlated with the reference condition using a power-law extrapolation, yielding the final value. The accuracy of this extrapolation depends entirely on the residual value, which must be within a very small range to ensure validity[94]. An important note is that the reported velocity values correspond to unstretched estimates since the stretch sensitivity decreases as the flame expands and stretch

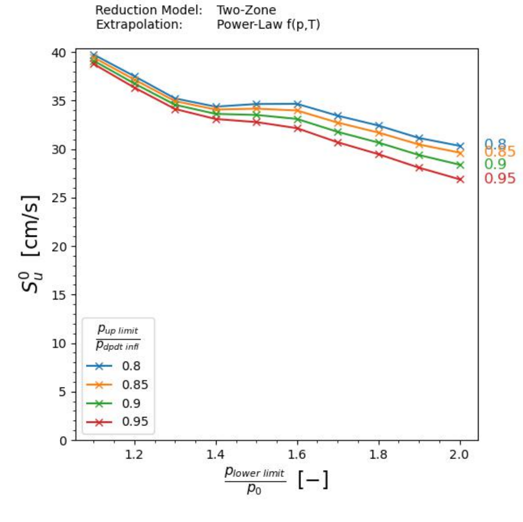

Figure 5.3: Dependency of LBV on the selected bounds for a case of CH_4 at T_0 = 21° C and p_0 = 1 bar

approaches zero over the usable interval.

5.4.2. MIE method

MIE is determined by varying the ignition input and identifying the minimum delivered energy required for successful ignition under fixed initial conditions. The delivered energy is computed from the measured voltage and current traces during the discharge, using the same definition consistently across the campaign. However, calculating the minimum ignition energy is a lengthy, time-consuming process because it requires multiple consecutive experiments to define it precisely. In fact, the canonical method used in the literature [53] involves defining a known upper bound at which ignition always occurs and subsequently lowering the energy in constant steps until a point is reached at which ignition no longer occurs. Once this occurs, the energy rises again in smaller steps until a new ignition is found. Here, many experiments are run and, if ignition occurs in at least 50% of the cases, the energy at this point is defined as the MIE. Due to time constraints, this is not achievable, and the number of repetitions per condition is set at 3, which is possible within the timeline.

Therefore, an alternative method is defined, which involves the reverse process, starting with a guessed lower bound that is not believed to cause ignition. This is because, in the case where ignition is not achieved, the energy level can be increased, allowing for a new attempt to be made without releasing the mixture and refilling the chamber, a time-consuming process. This process can continue by continuously updating the lower bound until ignition is found, which sets the upper limit. Ignition success is defined by the chamber pressure response: when the pressure trace shows a notable rise exceeding 1.5 times the initial pressure, the mixture is considered to have achieved combustion. Then, in the following two experiments, ignition is attempted again from the lower bound in smaller steps, thereby increasing the precision of the MIE definition.

Naturally, this method is not perfect, as it is possible that ignition is achieved on the first try, thereby defining an upper bound but missing the known lower bound, and requiring only two experiments to find it. Moreover, the accuracy depends entirely on the definition of the upper and lower bounds, which can be significantly wider than those in the literature due to the limited number of experiments. However, it is a good starting point, and in future iterations, it could achieve higher accuracy with only a few additional experiments, thereby improving on canonical methods.

5.5. Numerical methods

To complement comparisons with experimental literature data and to account for their absence in the research cases, a numerical component was introduced to provide a consistent baseline against which the experimental results could be interpreted, particularly when comparing trends across pressure, temperature, and equivalence ratio. All simulations were performed using Cantera. Since these simulations are based on certain modelling assumptions and partially validated by experimental results, they will be primarily discussed in terms of relative trends and closeness to experimental results, rather than as absolute values. The principal limitations relevant here arise from:

- Uncertainties and bias in the chemical kinetic mechanism itself
- Uncertainty in transport and thermodynamic data
- Model-form assumptions required to make the problem tractable (planar, laminar, ideal-gas, and typically adiabatic boundaries for freely propagating flames)

These assumptions intentionally neglect effects that may be present in the experiments, such as heat losses to walls, ignition transients, flame stretch and curvature, turbulence, and any non-uniformity in mixture preparation. As a result, agreement (or disagreement) between simulation and experiment must be interpreted through the lens of these simplifications, with particular caution when extrapolating outside the validation envelope of the underlying mechanism.

5.5.1. LBV simulation

The laminar burning velocity was computed in Cantera using premixed-flame calculations for each fuel or fuel blend at initial conditions corresponding to the experimental test matrix. For each operating point, the unburned mixture state was defined by the initial pressure, initial temperature, and equivalence ratio, after which the solver was used to obtain the steady, freely propagating flame solution and the associated unstretched LBV. This approach is appropriate for establishing a consistent baseline because it isolates chemistry-driven sensitivities to the initial state without introducing facility-specific boundary conditions. However, the resulting LBV remains a property of an idealised, planar, laminar flame, and therefore it should not be expected to capture experimental influences such as ignition energy deposition history, flame kernel development, and wall heat transfer.

Different chemical kinetic mechanisms were used depending on the fuel, reflecting the fact that no single mechanism provides uniformly optimal accuracy across all fuels under the full range of conditions explored. For methane, GRI-Mech 3.0 was adopted as the reference mechanism, given its widespread use for natural-gas combustion applications and its mature implementation in Cantera. For hydrogen, two mechanisms were employed: the canonical H2O2 mechanism and the KAUST mechanism, which was selected because it has been reported to provide a strong overall balance in predicting key hydrogen combustion targets, such as the laminar burning velocity and ignition delay [63]. Using both hydrogen mechanisms also enabled a practical assessment of mechanism-to-mechanism variability, a non-negligible contributor to uncertainty when simulations serve as a baseline for experimental comparisons.

For ammonia, three mechanisms were selected: Okafor, Otomo, and Zhang, selected based on comparative accuracy discussions in recent mechanism evaluation literature performed by Xu et al.[95]. In addition to pure NH_3, blended mixtures were simulated using the Okafor and Otomo models; this choice was driven by observed numerical robustness at elevated adiabatic temperatures as hydrogen content increases. In particular, when H_2 enrichment is sufficiently high, the Zhang mechanism fails to compute the adiabatic flame temperature, leading to a solver breakdown and a mixture reaching very high burnt-gas temperatures, for which the mechanism/property set becomes numerically problematic. The behaviour motivated excluding Zhang from blends to maintain the continuity of the parametric sweeps and avoid conflating numerical non-convergence with physical trends. Importantly, for NH_3 (especially at subzero initial temperatures), mechanism validation is comparatively sparse, so the selection reflects a best-available approach rather than a definitive claim of universal accuracy.

To map trends more clearly, the Cantera simulations were run not only at the exact experimental setpoints but also across an expanded parameter envelope surrounding the experimental range. For temperature, the experimental bounds were approximately 21° C down to -50° C, whereas the simulation sweep was extended to 30° C down to -70° C to better visualise trends and identify nonlinearities near the edge of the test domain. Similarly, the experimental pressure range of 1-2 bar and the equivalence ratio range of 0.8-1.2 were slightly expanded in the simulations to provide context for the experimental points and to reduce ambiguity when interpreting local slope changes. This strategy is useful for identifying consistent trends, such as a monotonicity with temperature, and for spotting regions where different mechanisms begin to diverge appreciably.

5.5.2. MIE simulation

Minimum ignition energy was estimated using the semi-empirical methodology proposed by Nawaz, which combines analytical relations for quenching and heat transport with kinetic and transport properties obtained from Cantera. The method, based on the approach from Cirrone et al.[53], relates MIE to a characteristic quenching distance and to the flame's ability to sustain itself against conductive losses, with the laminar burning velocity entering directly as a controlling parameter. In the present work, the LBV values used in the MIE model were those computed from the Cantera premixed-flame solutions at the relevant initial conditions, ensuring internal consistency between the chemical mechanism, transport properties, and the MIE estimate. This coupling is advantageous because it enables a systematic evaluation of how changes in temperature, pressure, and equivalence ratio affect the predicted ignition-energy requirement. However, the method also inherits limitations from both components: MIE is known to be sensitive to ignition-system details (spark gap, duration, electrode geometry), while LBV is computed for an idealised flame that does not account for ignition transients or facility heat losses.

Across all fuels, the numerical results should be interpreted as a mechanism-dependent baseline rather than a definitive prediction, particularly for NH_3 at low temperatures. For CH_4 and H_2, the employed mechanisms have extensive historical validation, so the simulations are expected to reproduce canonical trends with temperature and equivalence ratio reasonably well, even if absolute values may differ from experiment due to variations in experimental configuration. For ammonia and ammonia blends, the lack of abundant low-temperature experimental datasets means that the simulations inevitably rely on extrapolation and on the fidelity of the mechanisms outside their most-tested regime. ([MDPI][3]) In the MIE workflow, these uncertainties are significant because LBV directly influences the resulting MIE estimate, and errors in LBV can lead to systematic shifts in the inferred ignition energy. For this reason, this work emphasises trend comparisons, mechanism-to-mechanism spread, and consistency with known qualitative behaviour, while treating low-temperature NH_3 values as provisional until further experimental validation becomes available.

5.6. Uncertainty and repeatability

Uncertainty is a fundamental parameter in the analysis because both MIE and LBV depend on measured signals and on reduction procedures that introduce method-dependent variability. Three main sources of uncertainty are considered: setup-related, method-related, and repeatability. Setup-related uncertainty includes sensor calibration limits, acquisition-chain bandwidth and resolution, and systematic effects such as residual pressure after evacuation or thermal gradients during conditioning. Method-related uncertainty arises from the assumptions inherent to the selected reduction procedures, including the choice of the usable pressure bounds and the modelling assumptions used to reconstruct the flame history from pressure-only data. Repeatability uncertainty captures run-to-run variability, even under nominally identical conditions, and is quantified through repeated tests.

For quantities derived from measured inputs, such as pressure, uncertainty is propagated using first-order (linear) uncertainty propagation. If a reported quantity f depends on a set of independent inputs $\{x_i\}$, each with standard uncertainty u_{x_i}, that is, the uncertainty of a measurement result expressed as a standard deviation, then the propagated standard uncertainty is computed as

$$u_f = \sqrt{\sum_i \left(\frac{\partial f}{\partial x_i} u_{x_i} \right)^2}. \tag{5.3}$$

then, the total uncertainty u_T on a parameter can be found by also accounting for the repeatability uncertainty u_r

$$u_T = \sqrt{u_f^2 + u_r^2} \tag{5.4}$$

This approach is adopted because it is conservative for small perturbations and consistent with the objective of reporting uncertainty bands that reflect the measurement chain.

For LBV, the dominant setup-related contributions arise from dynamic pressure measurement (calibration, noise, and timing alignment), the initial-state definition (static pressure and initial temperature), and mixture preparation (partial-pressure tolerances and the finite evacuation setpoint). While uncertainties in the pressure measurement affect both the reconstructed thermodynamic trajectory and the calculated burning velocity, the remaining uncertainties affect the initial state of the mixture, thereby impacting the equivalence ratio, initial pressure, or temperature. This means that LBV plots will show both vertical error bars for the velocity and horizontal error bars for the considered parameter uncertainty. Repeatability is quantified by repeating each test condition and computing the standard deviation of the resulting burning-velocity estimates, which captures variability due to filling, thermal conditioning, and small differences in ignition kernel development. The total LBV uncertainty is then obtained by combining setup-related and repeatability uncertainties via RSS, so that the reported uncertainty band reflects both instrumentation limitations and run-to-run scatter. This treatment also makes it explicit when uncertainty is dominated by instrumentation rather than by physical variability.

For MIE, uncertainty has two distinct components. The first is the uncertainty in the delivered ignition energy, which is derived from the voltage and current traces and therefore inherits uncertainties in probe calibration and acquisition. The second is intrinsic to the threshold method: since MIE is identified by bracketing the transition between no-ignition and ignition outcomes, the true threshold lies within the interval defined by the last non-ignition case and the first ignition case. In this work, this bracketing uncertainty is reported explicitly and combined with the delivered-energy uncertainty, so that the final MIE value reflects both measurement limitations and the finite resolution of the threshold search. The bracketing uncertainty is taken as $u_{\mathrm{br}} = (E_{\mathrm{ign}} - E_{\mathrm{no}})/2$, and it is combined with the delivered-energy uncertainty using RSS to obtain the reported MIE uncertainty.

6

Results and analysis

In this chapter, the results obtained with the custom-designed experimental setup are presented and discussed. Before addressing the main objective of producing ammonia combustion data at subzero temperatures, the setup must be verified and validated to ensure the measurements are reliable and physically consistent. The verification step, described in Section 6.1, summarises the checks performed to confirm that the system behaves as expected and that the acquired signals are sensible. The validation is then divided into two parts, Section 6.2 and Section 6.3, to assess the ambient-condition and subzero configurations separately. This separation helps isolate potential sources of bias introduced by the environmental chamber and cooling approach, while also ensuring that the resulting datasets are reproducible and comparable with values reported in the literature. Finally, the novel experimental work on ammonia is presented in Section 6.4, where the results required to answer the research questions are reported and interpreted in the context of the experimental constraints and observed trends.

6.1. Verification

Verification is achieved using methane, as it is the most predictable and widely benchmarked fuel in premixed combustion experiments. This verification phase serves two purposes: it checks that the instrumentation captures the expected pressure transient shape, and it verifies that the full reduction pipeline can reproduce literature-consistent trends under controlled conditions. The process also enables practical debugging of the setup before moving to less reactive mixtures and colder conditions. In practice, it solved several issues related to signal conditioning, timing alignment, and repeatability, which are described in more detail in Appendix E. The end state of this phase is that the new chamber can be considered verified, as it produces results comparable to those obtained with the optical chamber, which is treated as the laboratory reference system.

Because the burning-velocity calculation relies primarily on the pressure trace, the behaviour of the pressure signal is the most critical element for obtaining reliable velocity results. A representative case at the end of the verification process is shown in Figure 6.1. Here, it can be seen that the sensor response is slightly erratic, as the trace does not simply rise and decay monotonically but rather exhibits an oscillatory component superimposed on the combustion-driven pressure rise. This behaviour is an inherent limitation of the 603CAA sensor under rapid thermal loading, as it undergoes thermal shock, producing transient artefacts that cannot be fully removed but are already minimised in this verification process. When the trace is cropped to the physically meaningful interval and compared against the curve measured

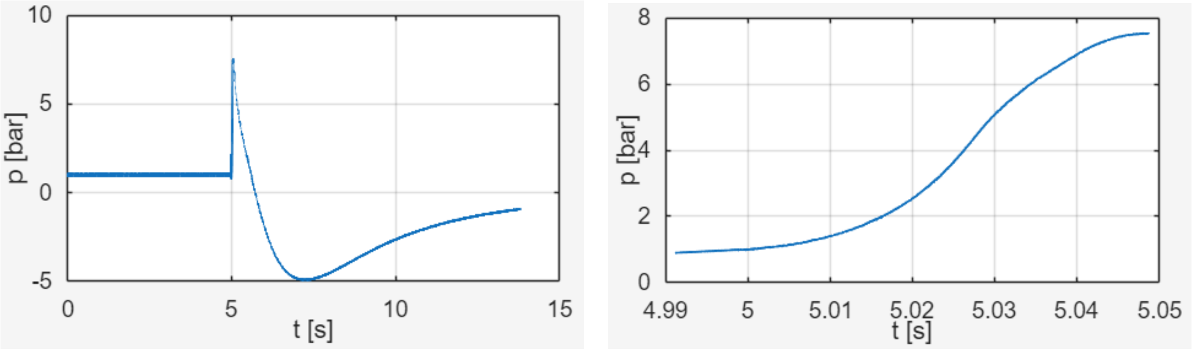

(a) Full pressure trace

(b) Combustion rise

Figure 6.1: Pressure traces after verification

by the optical chamber sensor for the same nominal conditions, the two traces are extremely similar in both shape and magnitude, supporting the validity of the measurement approach. However, to ensure consistency and still account for thermal shock and uncertainty of 10%, as specified by Lee, Choi, and Bae[96].

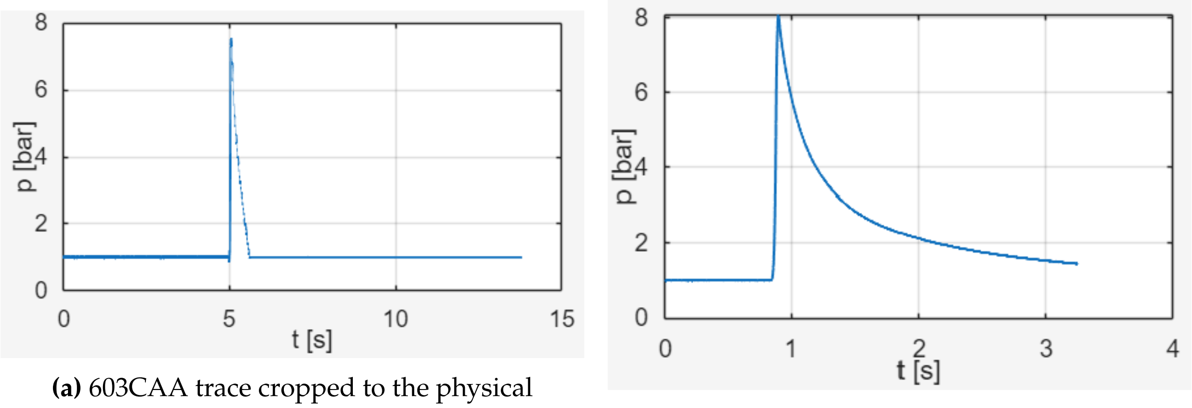

(a) 603CAA trace cropped to the physical behaviour

(b) 7061C pressure rise

Figure 6.2: Comparison pressure curves between the non-optical (603CAA) and optical (7061C) chamber

6.2. Validation at ambient temperature

The first validation step is performed at ambient conditions, as this is the most robust experimental condition and for which a large dataset of literature values is available. This is intended to demonstrate that the built setup and experimental procedure, from filling to ignition and data collection, can closely reproduce the LBV behaviour commonly observed in the literature. Therefore, this process is achieved at $T_0 = 21°$ C and $p_0 = 1$ bar. This condition is run for three fuels: the two of interest for this work, ammonia and hydrogen, plus methane, due to its predictable behaviour and large-scale validation. Therefore, the first experiments for each of these fuels are run and presented as a function of equivalence ratio in Figure 6.3.

These figures compare the obtained results with both experimental values from the literature (points) and simulations (continuous lines), showing good agreement with both data types. The larger error bars for some data points, such as ammonia at $\phi = 1$, are due to a larger gap between values calculated across different tries, which increases the repeatability uncertainty. For many conditions, outliers are found, such as for stoichiometric methane, for which one of the calculated values was 22 cm/s, considerably lower than the validated range of 36 cm/s to 40 cm/s. When such situations occur, the most likely culprit is believed to be the filling and mixing phase, as a value this low is more characteristic of leaner and richer mixtures. Therefore, an error in chamber filling may result in this type of outlier.

A further contributor to run-to-run variability is the ignition event itself, since the spark discharge is not perfectly repeatable from shot to shot. Variations in arc position and local mixture composition near the electrodes can alter the initial kernel development, introducing stochasticity that can propagate into the early pressure rise used for the velocity extraction. In these situations, the outlier is excluded from the analysis because it is clearly isolated, would disproportionately skew the fitted trend, and is more representative of a run-specific error than of meaningful physical behaviour. This filtering is applied conservatively, and for each condition, at least 2 valid repetitions are retained to ensure that repeatability uncertainty is captured in the reported error bars. Presenting the data with uncertainty bounds, therefore, supports the stated confidence interval and clarifies when scatter is dominated by experimental variability rather than a systematic shift in the underlying trend. An important thing to mention, which can be seen for $\phi = 0.8$ in Figure 6.3c, is that repeatability errors are much more considerable in magnitude with respect to setup and method ones. In fact, when the repeated measurements are very close together, the error bars are very small.

The agreement across all three fuels is particularly important, as both the pressure trace and the velocity calculation method remain robust despite very different chemical mechanisms. In fact, hydrogen is a much faster-reacting fuel, characterised by an incredibly high diffusivity, which causes the flame to expand considerably faster than with other fuels, resulting in a much steeper pressure rise. Ammonia, on the other hand, exhibits the opposite behaviour, as its diffusion is dominated by buoyancy and heat loss, leading to significantly slower expansion. Demonstrating consistent behaviour for these two extremes, as well as for the intermediate fuel, provides confidence that the overall methodology is transferable across markedly different combustion regimes. At the same time, part of the analysis remains fuel-dependent, since the admissible pressure bounds used for velocity extraction must be tailored to each fuel to remain within the model's assumptions and avoid regimes where secondary effects dominate. For this reason, stretch sensitivity cannot be ignored when comparing fuels or conditions, and the Markstein number is introduced as a compact way to interpret how flame stretch and preferential diffusion influence the burning velocity in a closed chamber.

The Markstein number is a dimensionless quantity that measures how sensitive a premixed flame's local burning velocity is to flame stretch, i.e., to curvature and flow-induced strain of the flame front. In other words, this parameter links small-scale flame structure (thickness, transport, chemistry) to the stability and dynamics of curved or strained laminar flames[97]. For a stretched laminar premixed flame, the local burning velocity S_L is often written in a linearised form as

$$S_L = S_L^0 - L_M \cdot \kappa \tag{6.1}$$

where

S_L^0 unstretched laminar burning velocity
κ flame stretch rate (from curvature and/or strain)
L_M Markstein length

Then, the Markstein number can be defined as the ratio between the Markstein length and the characteristic flame thickness, with the latter being the intermediate zone that separates the

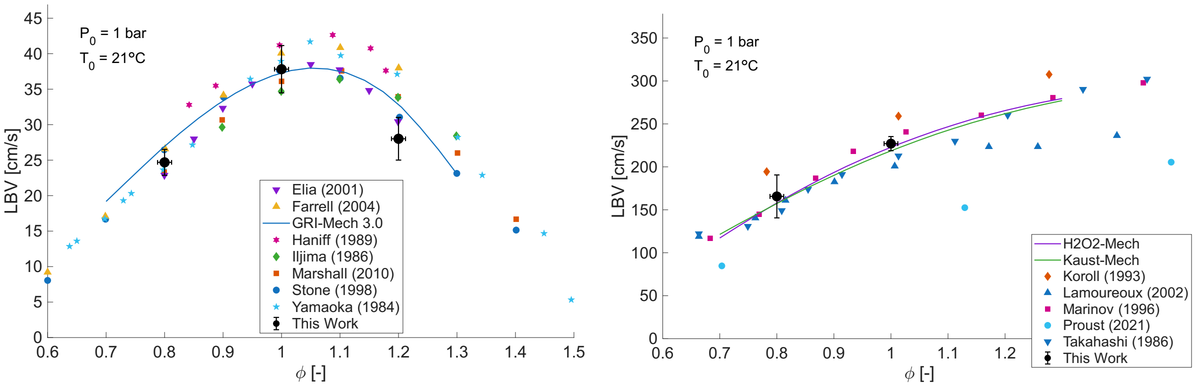

(a) CH_4 **(b)** H_2

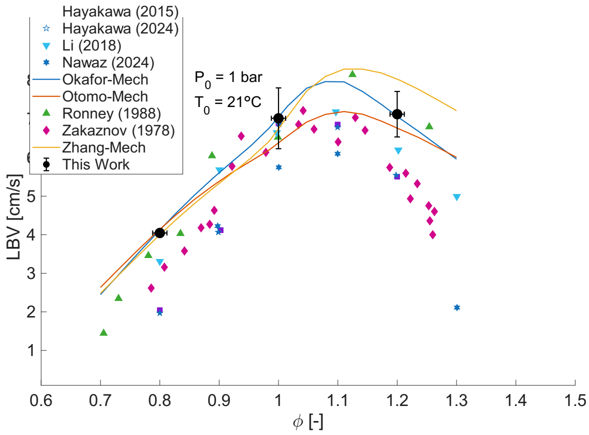

(c) NH_3

Figure 6.3: LBV of the considered fuels at $T_0 = 21°$ C and $p_0 = 1$ bar. Experimental results are denoted with points, simulations with continuous lines. Present results are described in black dots as 'This Work'

burnt from the unburnt mixture:

$$Ma = \frac{L_M}{\delta_L} \tag{6.2}$$

This quantity can describe how the curvature of the flame surface and aerodynamic strain modify the local burning velocity relative to the planar, unstretched value. The sign and magnitude of the Markstein number encode the role of diffusional-thermal effects and hydrodynamic instability. In fact, a positive number indicates that flame speed decreases with positive stretch, which tends to stabilise the flame against diffusional-thermal instabilities (cellular structures are suppressed). On the other hand, a negative Markstein indicates that flame speed increases with positive stretch, promoting cellularity and making the flame more unstable to perturbations and acceleration. The value of Ma depends on gas expansion, the effective Lewis number, and activation energy; therefore, changing the equivalence ratio or fuel/oxidiser type can alter both its magnitude and sign [97].

Methane, like most other existing fuels, has a positive Markstein number, indicating that its ignition is dominated by stretch and curvature; however, this tendency decreases as the flame expands until heat loss becomes dominant. On the other hand, hydrogen, in the lean-to-slightly-rich range, has a negative number, indicating that the stretch becomes more dominant as the flame radius increases. In fact, by the time $p/p_0 \geq 2.5$, the spherical kernel is often well into the cellular/unstable regime. Reyes et al. mention that in a constant-volume bomb used for CH_4/air and H_2/air combustion, for the latter at stoichiometry and raised pressures, it is almost impossible to maintain a purely laminar flame as pressure grows, and that the measured burning velocities are essentially in the cellular regime [98].

In a similar manner, the review from Faghih and Chen, and the many H_2 spherical-flame studies they cite, stresses that hydrogen flames go cellular at surprisingly small radii and modest pressures and that laminar S_L must be extracted by either using optically measured radii, applying stretch corrections using Markstein length, or discarding data beyond a critical radius for cellular onset. In this work, the latter option is chosen, as it can be directly implemented in CVDART using conservative values at the determined inflexion point. This means that the upper bound is not set directly at this point, but rather at 70-80% of it, depending on the condition.

Conversely, for ammonia, the created flames are very slow and thick, strongly affected by buoyancy, meaning that the flame rises in the combustion chamber due to the lower density, radiative losses, and again thermo-diffusive instabilities [8]. For this reason, it is particularly important to use a multi-zone method that includes radiation, as it significantly impacts its behaviour. Moreover, only mid-range pressure values should be used when the flame is quasi-spherical, and convection is not dominant. This ensures sufficient rise to escape ignition noise but is still far from the strong buoyancy distortion observed at larger radii.

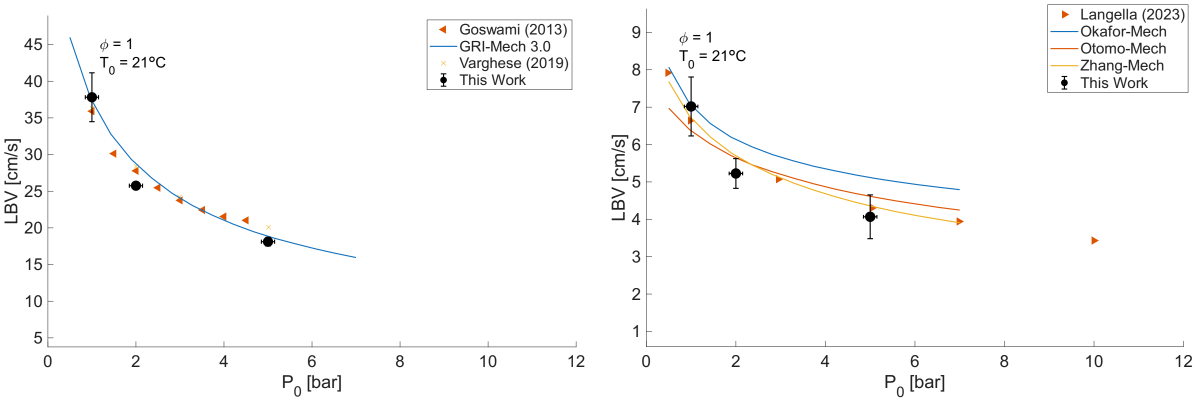

(a) CH_4 **(b)** NH_3

Figure 6.4: LBV of CH_4 and NH_3 for ϕ = 1 at varying pressure

The pressure-dependence validation at ϕ = 1 further checks whether the system reproduces the canonical reduction in laminar burning velocity with increasing initial pressure. Figure 6.4 shows that both methane and ammonia exhibit decreasing LBV as p_0 rises, which is consistent with the established expectation that higher density and enhanced third-body effects tend to suppress laminar propagation, even though chemical collision rates also increase. Methane provides a particularly clean reference in this context because its burning velocity is sufficiently high to generate a clear pressure signal while remaining relatively insensitive to secondary effects over the analysed window. Ammonia follows the same overall trend as the compared literature, but its much lower burning velocity makes it more sensitive to factors such as wall heat loss and buoyancy effects as the flame grows and the effective radius increases. This explains why, at higher pressure, the trend diverges more from the one shown in the simulations, which do not represent these factors as well, but remains quite close to the experimental results from Langella et al. [30]. Once the expected behaviour of methane is confirmed, the pressure-sweep effort is therefore focused on ammonia rather than extended to hydrogen, to prioritise the fuel that is both the most challenging for the facility and most relevant to the scope of this work.

Another notable observation during validation is that ammonia, particularly at lean equivalence ratios, produces a noticeably noisier pressure trace than methane or hydrogen. This behaviour is primarily linked to the much lower burning velocity, which increases the flame's residence time in the chamber and makes the pressure more susceptible to disturbances. Over this longer timescale, wall heat loss and buoyancy have more opportunity to influence the signal, and their impact becomes visible because the underlying pressure rise is comparatively weak. This is consistent with the work of other laboratories, which also demonstrate the occurrence of such perturbations visually [99]. As a consequence, extracting a stable burning-velocity curve for ammonia generally requires stronger filtering or smoothing of the pressure data, and this unavoidably increases the uncertainty associated with the processed result.

6.3. Validation at subzero temperatures

With the previous discussion, the design and setup can be considered validated for ambient conditions. However, before any discussion of the accuracy of ammonia results at subzero temperatures can proceed, the setup's ability to produce correct results in this range must be confirmed. The expected behaviour is that of a decrease in burning velocity, since lower temperatures have many detrimental effects on flame expansion, as explained in Chapter 2. In fact, they reduce reaction rates while increasing gas density, both of which lead to lower flame speed. Moreover, another important factor to consider is that cold walls induce even greater thermal losses relative to ambient conditions, further reducing the part of the pressure curve representing the laminar flame expansion. For this validation, hydrogen is chosen for two main reasons: first, its chemical oxidation processes, which more closely resemble those of ammonia due to the absence of carbon structures, and second, its very high flammability, which ensures robust ignition even at very low temperatures.

Both in Figure 6.5 and in Figure 6.6, the experimental results are compared to simulation data and to the dedicated low-temperature experiments reported by, most notably, Michaux et al.[69], which were performed in a facility cooled using liquid nitrogen. The expected decrease in laminar burning velocity with decreasing temperature is clearly reproduced, and the overall agreement across equivalence ratios is consistent with the literature trends. A small deviation from strict monotonicity is observed, most noticeably for $\phi = 1$ between 0° C and -15° C, where two neighbouring points do not follow the ideal decreasing sequence. Given the measurement scatter and the relatively small temperature spacing between these conditions, this local inversion is treated as within the experimental variability of the present campaign rather than as evidence of a distinct physical trend. On that basis, the dataset is considered adequate for validation, since the dominant behaviour is captured and the results fall within the same range as the reference experiments and simulations.

In general, the measured values tend to lie slightly below the simulated values, with this offset consistent with the analysis choices and the non-adiabatic nature of the experiment. First, the pressure bounds used for hydrogen are selected conservatively to avoid entering a highly stretched regime, which biases the extracted burning velocity downward relative to an extrapolated zero-stretch value. Second, numerical mechanisms typically model adiabatic flame propagation, whereas the present facility necessarily experiences heat losses that become more significant as temperature decreases and the flame speed slows. These losses reduce the effective energy available for flame growth and can therefore depress the inferred burning velocity compared with idealised predictions.

Another aspect worth noting is that, at these temperatures, the vacuum level could no longer be reduced to the target 200 mTorr and instead stabilised at a minimum of about 1000 mTorr. This behaviour is attributed to differential thermal contraction between the main chamber body and its inserts, which alters the sealing conditions as the assembly cools. As a result, micro-leaks occur at chamber-insert interfaces during subzero operation. Even so, the initial vacuum level remains low enough that this effect can be considered negligible for the present validation work, and no clear impact on the measured trends was observed within the campaign's uncertainty. The observation is nonetheless useful because it highlights that sealing performance is temperature-dependent and is a partial limitation of the designed chamber, especially when extending the temperature range to conditions closer to cryogenic.

This outcome also supports the decision not to begin the experimental campaign directly in the optical chamber to reduce delays while the setup was still being finalised and component availability was limited. In fact, because metals and quartz have very different thermal expansion coefficients, the risk of stress concentration and sealing loss during thermal cycling is higher than in the metal-only configuration used for validation. Observing temperature-driven leakage behaviour even in the current assembly indicates that starting with the optical chamber would have increased the likelihood of damage or permanent seal degradation. Moreover, any damage to the quart may have led to a catastrophic failure due to its inherent brittleness.

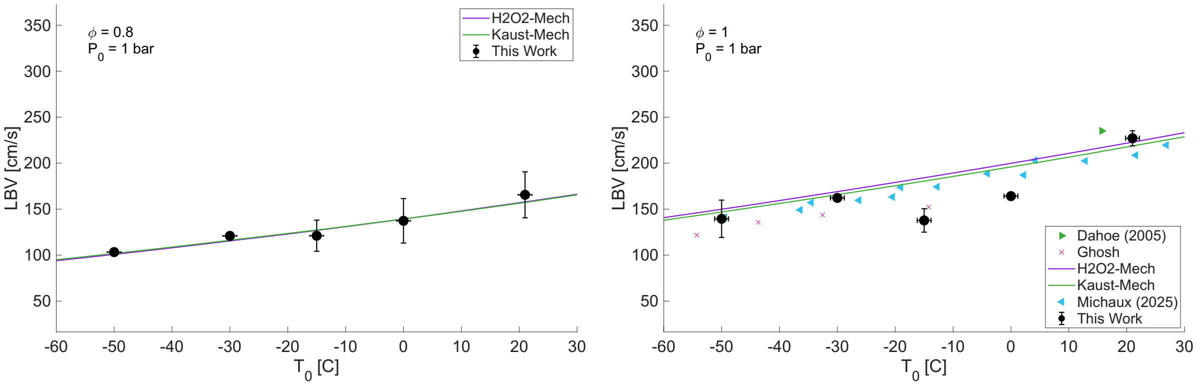

(a) $\phi = 0.8$ **(b)** $\phi = 1$

Figure 6.5: LBV of H_2 at $p_0 = 1$ bar with decreasing temperature

6.4. Novel experimental results

Once the setup is validated, the primary experiments can be carried out: combustion of ammonia and ammonia blends at subzero temperatures. Moving into this regime required operating at the practical limits of the ignition system, most notably the maximum spark energy the power supply could deliver, which was up to 33 V. To maintain reliable ignition under these colder conditions, the electrode gap was adjusted so that breakdown could occur at the available voltage, and the deposited energy was sufficient to initiate a flame kernel. This adjustment was required at 0° C and at -50° C, where stable ignition could only be achieved with a reduced spark gap of approximately 0.5 mm. At the intermediate temperatures investigated, the initial gap of 2 mm remained adequate to guarantee ignition without further changes to the ignition settings.

The need for a smaller gap at the coldest conditions is consistent with the reduced chemical reactivity and slower kernel development expected for ammonia-based mixtures as temperature decreases. Under these conditions, a weaker or more spatially dispersed discharge is more likely to quench before a self-sustaining flame is established, especially in the presence of higher heat

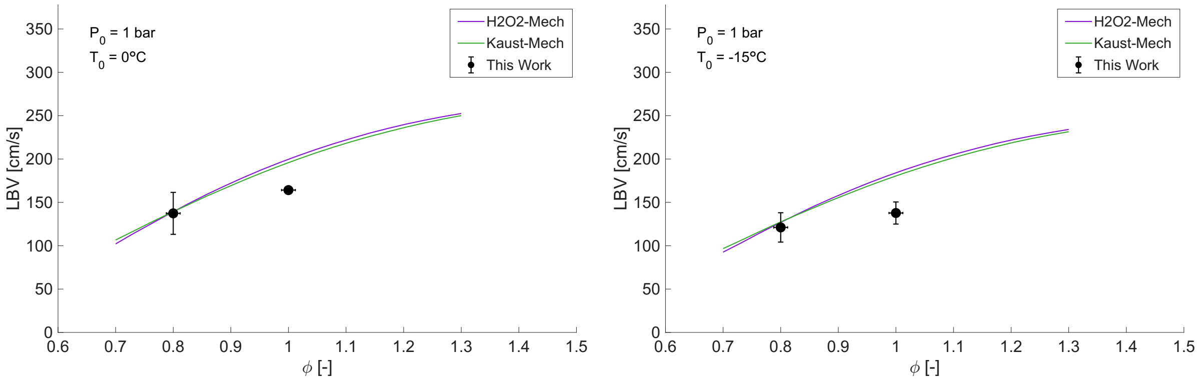

(a) T_0 = 0° C

(b) T_0 = -15° C

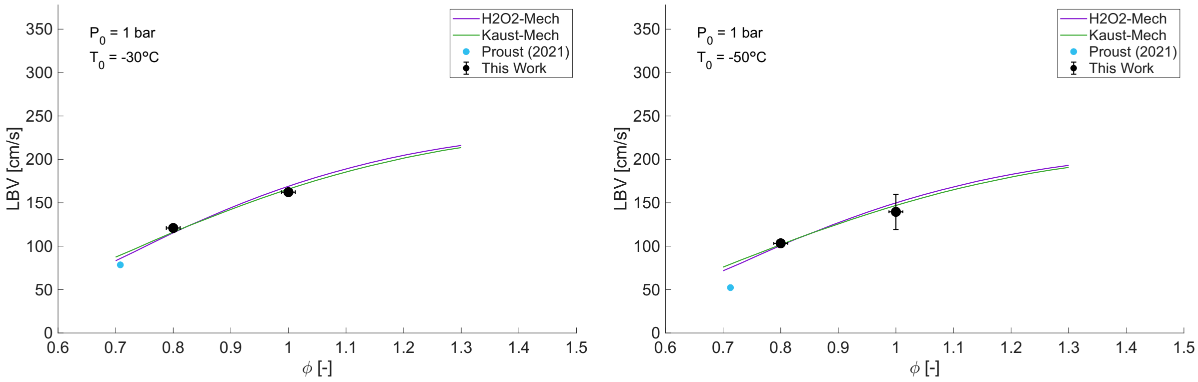

(c) T_0 = -30° C

(d) T_0 = -50° C

Figure 6.6: LBV of H_2 as a function of ϕ at p_0 = 1 bar at subzero temperatures

losses and reduced flame speeds. On the other hand, the need for a reduced gap at 0° C, despite successful ignition at -15° C and -30° C with 2 mm, is attributed to less favourable near-electrode boundary conditions around the freezing point, where condensables and thermal-control cycling can increase local quenching and variability, while colder setpoints provide a more stable and cleaner environment for kernel growth. Reducing the gap concentrates the discharge and improves repeatability by lowering the breakdown threshold and promoting a more consistent energy deposition in the near-electrode region. This modification is treated as a controlled change to the setup rather than a change in the physics being studied, and it is explicitly reported to allow comparisons across temperatures to be interpreted with the ignition configuration in mind. The remaining experimental procedures are kept identical to the validated approach to ensure that observed differences are attributable to temperature and mixture effects rather than to changes in the measurement methodology.

At these temperatures, exhaust water also began to accumulate in the pipeline leading into and out of the chamber, forcing the line to be opened several times because the vacuum pump could not clear the condensate on its own. This behaviour is attributed to condensation and freezing in the cold sections of the exhaust path, where water vapour from combustion products can condense, gradually reducing the line's conductance. While this required periodic intervention, it did not indicate ammonia condensation within the chamber under the tested conditions. In fact, the ammonia partial pressure during combustion remained sufficiently low that saturation was not reached at the wall temperature, so liquid ammonia formation was not expected. This distinction is important because it separates an operational limitation of the exhaust system from a change in the in-chamber thermodynamic state of the reactants or products.

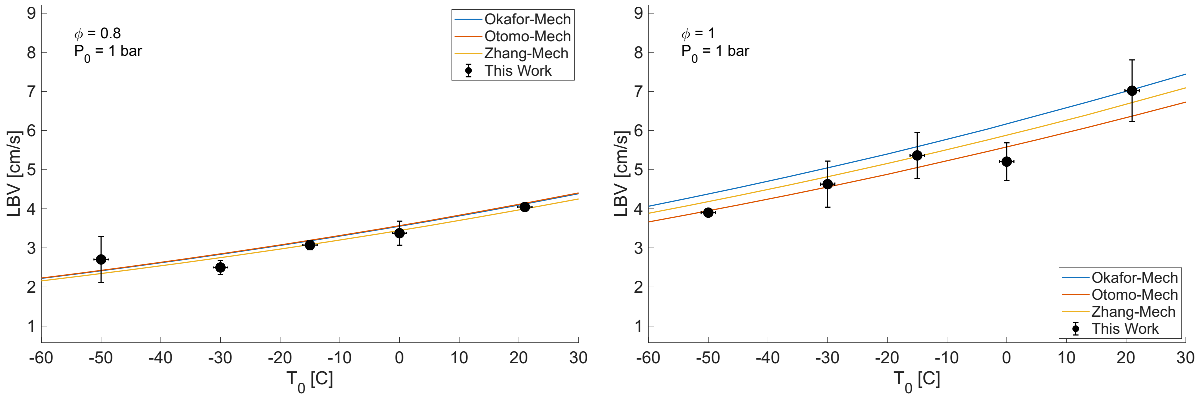

(a) $\phi = 0.8$ **(b)** $\phi = 1$

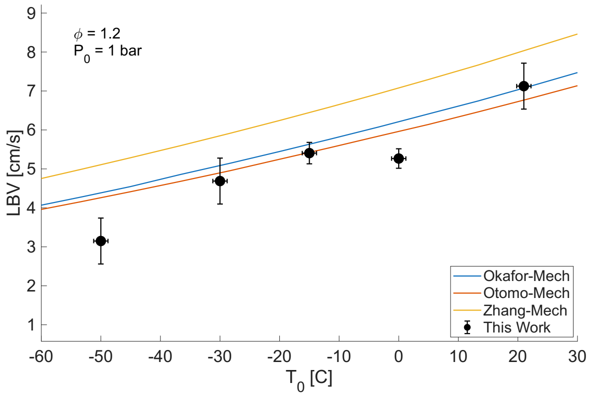

(c) $\phi = 1$

Figure 6.7: LBV of NH_3 at $p_0 = 1$ bar with decreasing temperature

As for hydrogen, ammonia also shows the expected decrease in burning velocity with decreasing initial temperature, and the experimental points follow the same overall behaviour as the simulated curves in Figure 6.7 and Figure 6.8. However, it must be noted that, for ammonia, the numerical values are not formally validated over the full temperature range because no experimental data are available in the literature. For this reason, the trends are more informative than the exact offsets, although measurements that remain close to the predicted values suggest that the simulations capture the dominant temperature dependence reasonably well. One thing to note is that the experimental slope with temperature appears steeper than in the simulations, indicating that, as temperature decreases, the velocity decreases more than expected numerically. A plausible physical explanation is that non-adiabatic effects become increasingly important as temperature decreases and the flame slows down, so heat losses, stretch sensitivity, and buoyancy-related distortions suppress propagation more strongly than an adiabatic, idealised calculation would. This interpretation is also consistent with the visibly larger scatter and uncertainty at the coldest and leanest points, where the pressure rise is weakest, and the extraction becomes more sensitive to filtering and window selection.

The two most evident deviations from monotonicity occur at $T_0 = 0°$ C and $T_0 = -50°$ C. This corresponds to the two conditions under which the electrode gap was changed, so part of the difference is attributable to this. In general, it is expected that some results do not match the prediction exactly due to experimental uncertainties, but their occurrence under these conditions is a hint that the electrode gap may affect the results. In fact, a smaller gap can bias the inferred velocity towards a lower value for two coupled reasons that do not necessarily

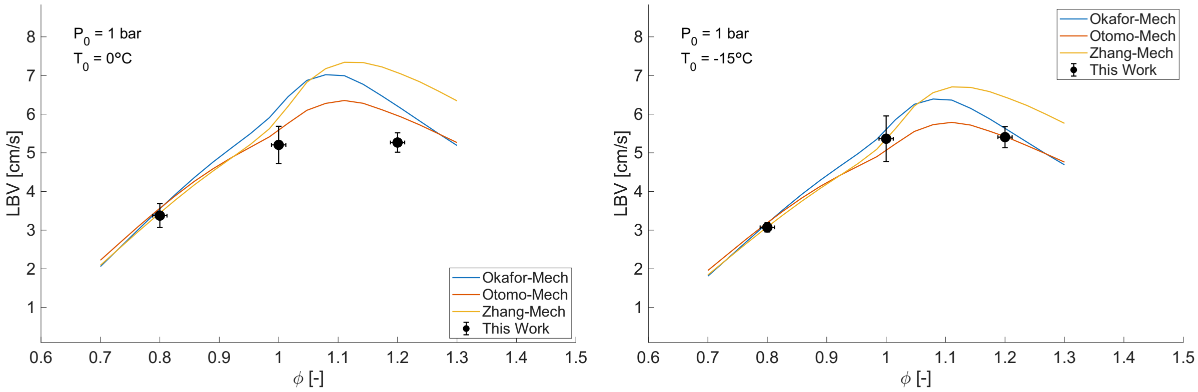

(a) T_0 = 0° C

(b) T_0 = -15° C

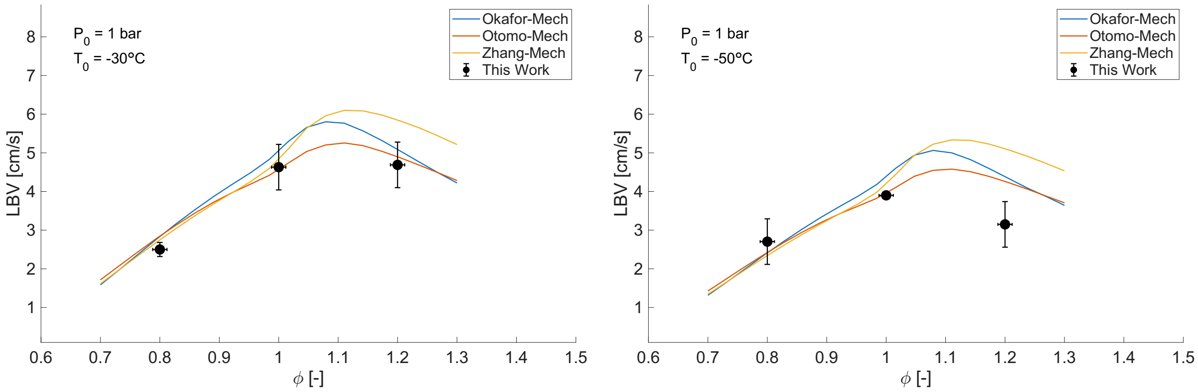

(c) T_0 = -30° C

(d) T_0 = -50° C

Figure 6.8: LBV of NH_3 as a function of ϕ at p_0 = 1 bar at subzero temperatures

reflect a true change in laminar flame behaviour. First, a tighter gap concentrates the discharge closer to the electrodes and nearby metal, increasing conductive heat losses from the nascent kernel and making the early flame more prone to partial quenching or delayed growth, even if ignition is ultimately successful. Second, the ignition kernel created by a short, intense discharge can be smaller and more highly curved in the first instants, thereby increasing flame stretch and reducing the apparent burning velocity for mixtures with positive Markstein numbers, such as ammonia.

Considering $\phi = 1$, the difference between the burning velocity at 21° C and at -50° C is approximately 3 cm/s, which is a substantial reduction given the already low values for ammonia. At the leanest conditions, the measured velocities approach values below 3 cm/s, where early quenching becomes increasingly likely and small perturbations have a larger relative impact on the extracted result. In this regime, the combination of reduced chemical reactivity, stronger heat losses, and longer residence time in the chamber can readily lead to weaker pressure rises and increased scatter. The trends in Figure 6.8 remain physically consistent, with a peak close to slightly rich mixtures and a progressive reduction of the curve as temperature decreases. The spread between kinetic mechanisms also becomes more consequential at low temperature because a similar absolute discrepancy corresponds to a larger fraction of the measured burning velocity. Overall, the dataset supports the expected temperature dependence while highlighting that ammonia at subzero conditions is close to the operational limits of ignition robustness and signal quality.

A summary of this behaviour is shown in Figure 6.9, where the horizontal error bars are

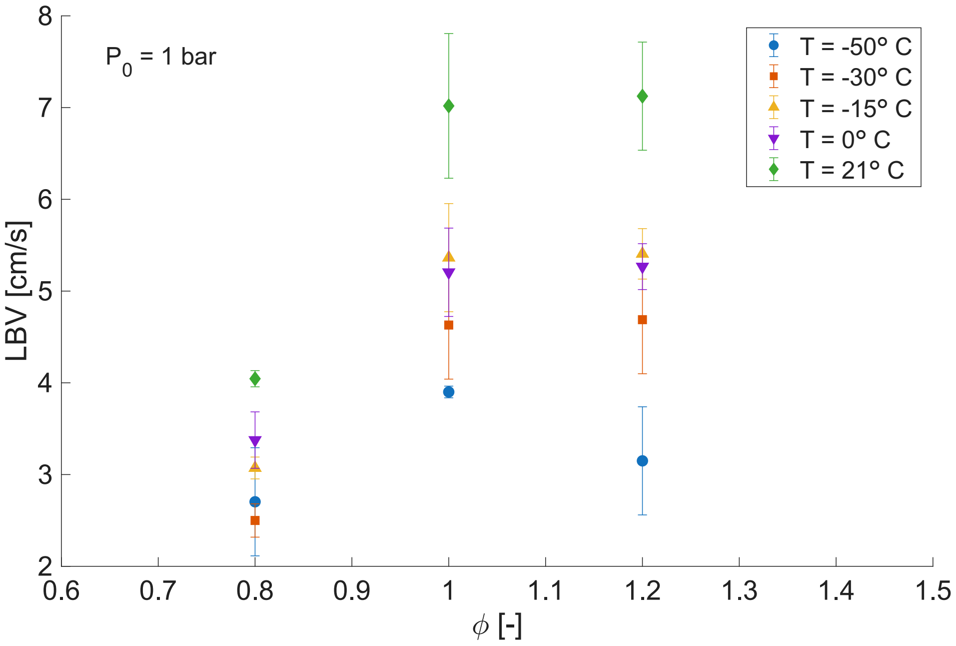

Figure 6.9: Variation of LBV with temperature as a function of ϕ

removed for clarity. Here, a compact view of the temperature sensitivity across equivalence ratios at $p_0 = 1$ bar is provided, highlighting that the reduction in burning velocity with decreasing T_0 is systematic across the three investigated ratios. The highest velocities are consistently observed near $\phi \approx 1.2$, indicating that the maximum location does not shift markedly over the explored temperature range within the resolution of the present dataset. The separation between the temperature levels is largest around stoichiometric conditions, where the absolute burning velocities are highest and the drop from 21° C to -50° C becomes most evident. On the lean and rich sides, the absolute values are smaller and the differences between neighbouring temperature points become comparable to the experimental uncertainty, which explains why the points can appear more clustered even though the trend remains decreasing.

Moreover, the largest decrease in velocity is achieved in the jump from 21 to 0° C, with the other changes in temperature causing a much lower difference in velocity. This suggests that the dominant controlling processes may change around the freezing point, so that below 0° C the burning velocity becomes less sensitive to temperature than expected from a simple extrapolation of the trend at higher temperatures. A possible explanation is that, once the flame becomes sufficiently slow, non-adiabatic and transport-related effects begin to dominate the apparent temperature dependence, effectively masking part of the kinetic acceleration that would otherwise be predicted for an adiabatic flame. In this regime, increased wall heat loss, longer residence times, and greater sensitivity to stretch can reduce the measured burning velocity in a manner that does not scale linearly with T_0, leading the curve to appear to "flatten" with further cooling.

In that sense, the observed behaviour does not necessarily imply a change in chemical mechanism, but rather a shift in the balance between kinetics and loss mechanisms in the closed vessel at low temperature. A similar behaviour is also visible for hydrogen in Figure 6.5, which also shows a more prominent jump between 21° C and 0° C. This supports the interpretation that the effect is linked to the low-temperature regime rather than being unique to ammonia chemistry. Overall, the figure supports the conclusion that lowering T_0 reduces ammonia burning velocity across the full range of ϕ considered, with a more marked difference at the beginning of the subzero domain.

This significant reduction in burning velocity at subzero temperatures motivates an investigation of hydrogen blending under the same conditions. The overall effect of hydrogen addition is shown in Figure 6.10 for the three equivalence ratios considered at $T_0 = -50°$ C and $p_0 = 1$ bar. As expected, blending ammonia with hydrogen increases the burning velocity, and even the change from 20% to 30% H_2 produces a clearly measurable rise for all three equivalence ratios. This rapid increase is consistent with hydrogen's role in increasing the radical pool and improving transport in the preheat zone, thereby promoting faster flame development and reducing the likelihood of early quenching in such cold conditions. At low hydrogen fractions, the measured points remain close to the ammonia baseline, but the upward curvature of the simulated trends indicates that larger additions should yield increasingly strong gains, rather than a linear improvement. The close match between the measurements and both mechanisms at 20% and 30% also suggests that, at least for modest blending, the dominant enhancement mechanism is captured reliably by the kinetics and transport models.

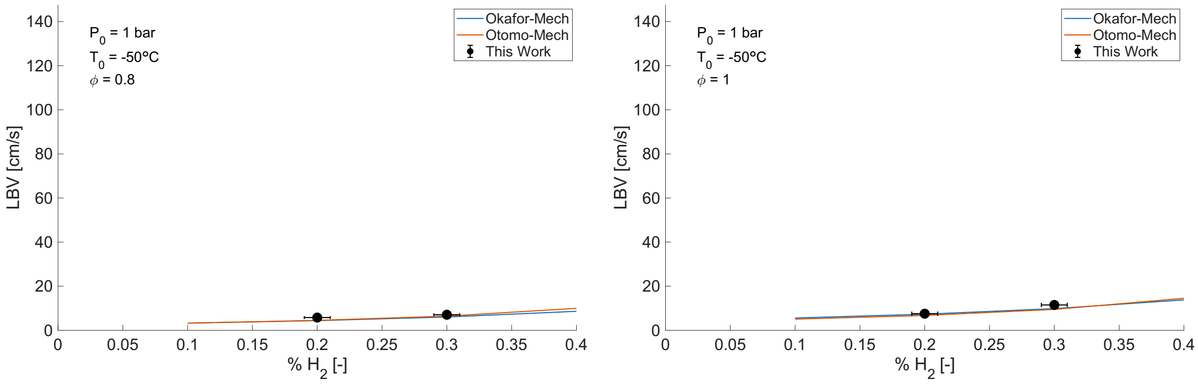

(a) $\phi = 0.8$ **(b)** $\phi = 1$

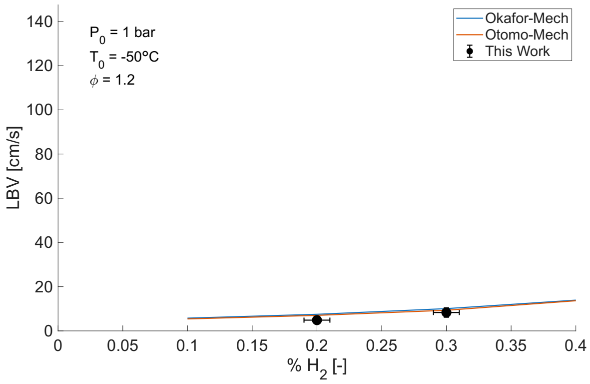

(c) $\phi = 1.2$

Figure 6.10: LBV of ammonia blends as a function of hydrogen addition at p_0 = 1 bar and T_0 = -50° C

The same behaviour is visible when the blends are viewed as a function of equivalence ratio in Figure 6.11. For both 80% NH_3 - 20% H_2 and 70% NH_3 - 30% H_2, the curve shifts upward compared with neat ammonia at the same temperature, and the peak remains close to slightly rich conditions. The 30% H_2 blend exhibits a substantially higher maximum and a clearer separation from the 20% blend, which indicates that hydrogen content becomes a primary control parameter at $T_0 = -50°$ C. The increased velocities also improve signal strength and repeatability, which is reflected in the smaller relative scatter compared with the neat-ammonia results at the same temperature. Overall, these results confirm that hydrogen blending is an effective lever for recovering flame propagation rates under subzero conditions, and they provide a quantitative baseline for assessing the amount of hydrogen required to achieve a desired level of ignition robustness and burning velocity in the present facility.

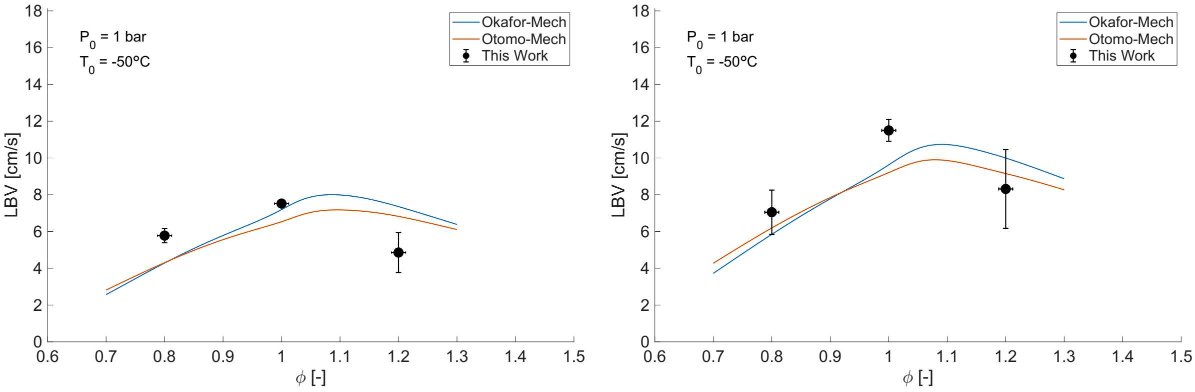

(a) 80% NH_3 - 20% H_2 **(b)** 70% NH_3 - 30% H_2

Figure 6.11: LBV of ammonia blends as a function of ϕ at T_0 = -50° C

A useful point of comparison is the hydrogen-blending effect at room temperature, where the same qualitative behaviour is expected even though no ambient-temperature blend measurements were performed in this campaign. Literature data, shown in Figure 2.14b at approximately 1 bar show a strongly non-linear increase in laminar burning velocity with hydrogen mole fraction, with relatively modest gains at low x_{H_2} and a much sharper rise once hydrogen becomes a dominant fraction of the fuel blend. This curvature is consistent across multiple datasets and indicates that hydrogen acts as more than a simple additive, reshaping the radical pool and the transport properties that control flame propagation. At room temperature, the baseline burning velocity of neat ammonia is already higher than at subzero conditions, so the absolute velocities reported for blends are also substantially larger. However, the main takeaway is that the sensitivity to hydrogen fraction remains strong, and the literature trend indicates that even moderate additions can move the mixture into a distinctly different propagation regime. In that sense, the absence of room-temperature blend points in this work does not prevent a meaningful comparison because the reference datasets already establish the typical scaling of S_u with x_{H_2} under ambient conditions.

When these room-temperature trends are compared with the present subzero results at $T_0 = -50°$ C, two differences become apparent that are consistent with the physics of slow flames in a closed vessel. First, the relative impact of modest hydrogen additions is amplified at low temperature because the ammonia baseline is very low, so the same absolute increase corresponds to a much larger fractional change and produces a clearer improvement in ignition robustness and signal strength. This is visible in Figure 6.10, where the increase from 20% to 30% H_2 yields a noticeable rise in burning velocity for all three equivalence ratios, even though the blend remains ammonia-rich. Second, subzero operation is more susceptible to non-adiabatic losses and early kernel quenching, so hydrogen helps not only by accelerating chemistry but also by making the early flame less marginal against heat loss and stretch effects. As a result, blending at -50° C acts as a practical stabilisation lever, whereas at room temperature it primarily acts as a performance lever, raising an already viable burning velocity. Taken together, the literature and the present measurements indicate that hydrogen blending increases burning velocity at both temperatures, but the operational value of modest additions becomes particularly pronounced as temperature decreases and the flame approaches the facility's stability limits.

An important summary of the magnitude of the improvement relative to neat NH_3 is shown in Figure 6.12. The figure highlights that even a modest hydrogen addition produces a substantial increase in burning velocity across the investigated equivalence ratios. At $\phi = 1$, adding 20% H_2 increases the burning velocity by roughly 4 cm/s, which is significant given that neat ammonia at $T_0 = -50°$ C is only around 4 cm/s. This brings the blend close to the room-temperature value of neat ammonia, which in this work was approximately 7 cm/s, while the corresponding 20% blend reaches about 7.5 cm/s. Increasing the hydrogen fraction to 30% yields an additional step change, raising the stoichiometric burning velocity to roughly 11.5 cm/s. The same ordering is observed at $\phi = 0.8$ and $\phi = 1.2$, albeit with a smaller magnitude increase, showing that hydrogen content acts as a strong control parameter even when the mixture remains ammonia-rich.

Overall, the comparison makes clear that relatively small hydrogen additions can compensate for a large fraction of the low-temperature penalty in burning velocity. In particular, the 20% H_2 cases already recover burning velocities comparable to the neat-ammonia values measured at 21° C, at least around stoichiometric conditions, thereby directly improving ignition robustness and reducing the likelihood of early quenching in the chamber. The 30% H_2 cases go further, moving the mixture into a distinctly faster-propagating regime, where the pressure rise is stronger, and the extracted values become less sensitive to filtering choices and small perturbations. At the same time, the larger uncertainty bars for some rich-blend points indicate that repeatability can still be limited by shot-to-shot variability at low temperature, particularly when ignition settings differ, or the kernel growth remains sensitive to local conditions. These results therefore show that hydrogen blending is an effective mitigation strategy for subzero ammonia combustion and provide a quantitative basis for selecting blend fractions, as they can not only recover room-temperature performance levels but also exceed them. Something important to mention is that, ideally, a larger number of experiments would be preferred to increase the measurement accuracy and reduce the result uncertainty, which is left for future work due to time constraints.

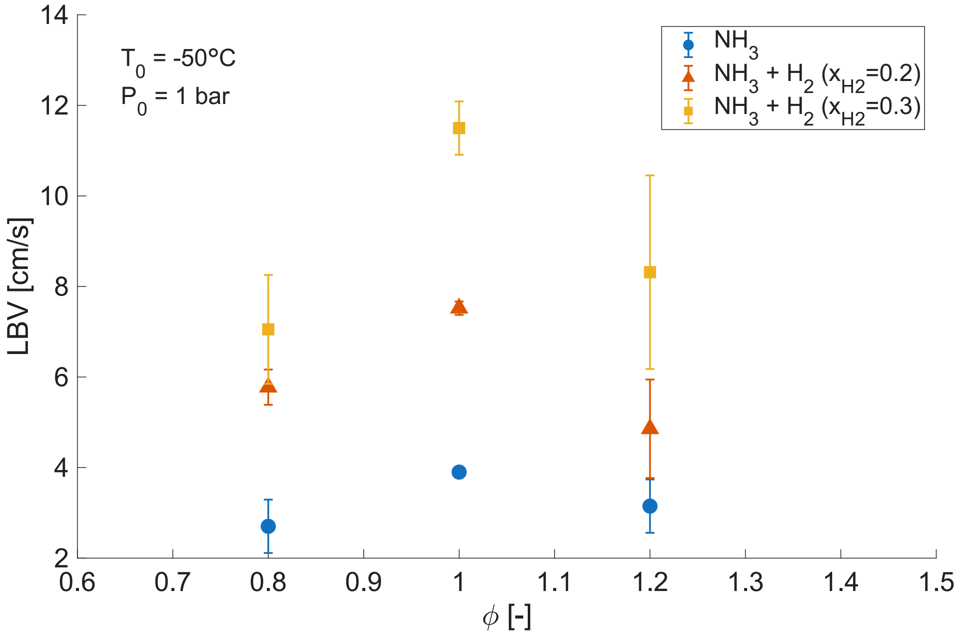

Figure 6.12: Comparison of LBV between NH_3 and NH_3 blends

6.5. MIE calculation

The calculation of the minimum ignition energy was based on the voltage and current signals acquired with the probes described in Subsection 4.2.2. Unlike the pressure measurement, this diagnostic proved highly sensitive to electrical noise and to run-to-run variability in the recorded waveforms, leading to inconsistent voltage and current traces. This behaviour is illustrated in Figure 6.13, where two consecutive runs for ammonia produce signals with markedly different shapes despite nominally identical conditions. In the first case, the trace captures a rapid voltage collapse across the gap accompanied by a current spike, followed by a return to a quasi-steady level, which is broadly consistent with the expected discharge sequence. In the second case, the signal shape is not physically consistent with a clean breakdown and appears dominated by measurement artefacts, indicating that the acquired waveform cannot be interpreted reliably as deposited spark energy. Similar inconsistencies occurred frequently and without a clear operational trigger, suggesting that the limitation lies primarily in the measurement chain rather than in the combustion conditions themselves.

As a consequence, the derived MIE values are not sufficiently robust for quantitative interpretation, and the results in Figure 6.14 should be treated as indicative rather than definitive. Under some conditions, the measured values fall within the range suggested by simulations and literature data, whereas at others they deviate strongly and reach unrealistic levels, consistent with sporadic corruption of the underlying voltage or current signals. Given this variability, a high-accuracy determination of MIE was not feasible within the present campaign, and the current methodology does not provide enough confidence to claim that the reported values represent the true minimum ignition energy. Future work should therefore prioritise improving the electrical measurement chain, for example by adopting a dedicated high-bandwidth differential probe arrangement, improving grounding and shielding, and ensuring that the acquisition bandwidth and triggering are appropriate for the discharge timescale. In addition, repeating each condition more thoroughly would improve statistical confidence by reducing sensitivity to isolated corrupted shots and enabling a more reliable identification of representative waveforms. With these changes, the MIE analysis could be revisited in a manner consistent

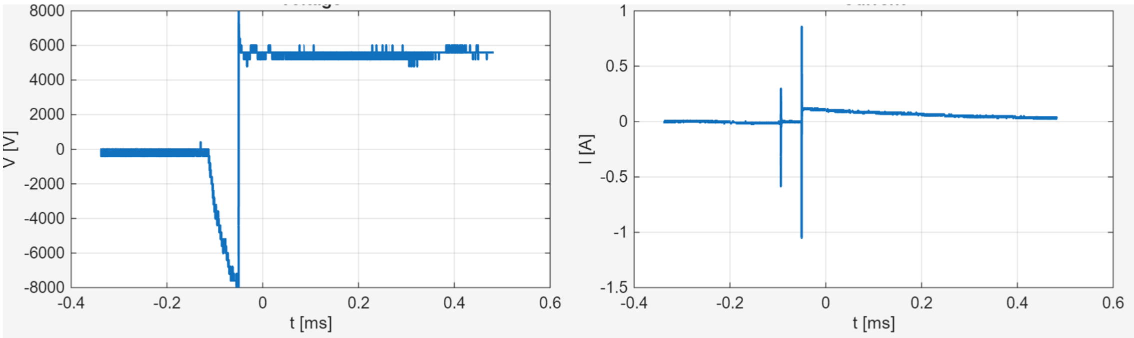

(a) First run

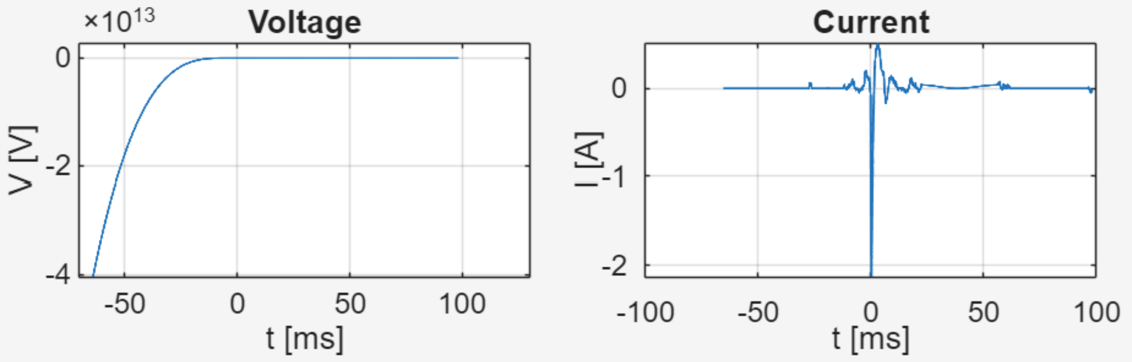

(b) Second run

Figure 6.13: Voltage and current curves for NH_3 at T_0 = 21° C and p_0 = 1 bar

with the level of confidence achieved for the pressure-based burning-velocity measurements.

While a quantitative determination of MIE is therefore not feasible with the present voltage and current diagnostics, a qualitative assessment can still be drawn from the commanded primary-coil settings, namely the supply voltage and the dwell time. These inputs do not map linearly to the energy actually deposited at the spark gap, because the coil dynamics, saturation behaviour, and discharge efficiency depend on the circuit configuration and the breakdown process. Nevertheless, the minimum primary settings required to achieve ignition provide a consistent and practical indicator of how ignition difficulty varies across operating conditions. Establishing a direct correlation with secondary-gap energy remains challenging, as the energy delivered into the gap is influenced by additional factors, such as breakdown voltage, discharge duration, arc attachment, and losses in the wiring and coil, which are not reliably captured by the present probes. For this reason, the following values are interpreted as comparative thresholds rather than as absolute energies.

For stoichiometric ammonia at ambient conditions, ignition was achieved at minimum primary settings of 14 V and 1 ms. When lowering the temperature to -30° C, the minimum required settings increased to 33 V and 8 ms, corresponding to the operational limits of the power supply voltage and the ignition-coil dwell time. This increase is consistent with reduced reactivity and slower kernel development at low temperatures, which require higher ignition intensity to prevent early quenching. At 0° C and -50° C, ignition at the original electrode spacing could not be achieved within the available supply limits, and the electrode gap was therefore reduced as described previously. With the reduced spacing, ignition at -50° C was achieved at minimum settings of 22 V and 5 ms, indicating that the geometric change effectively reduced the breakdown threshold and improved the coupling of the discharge to the near-electrode

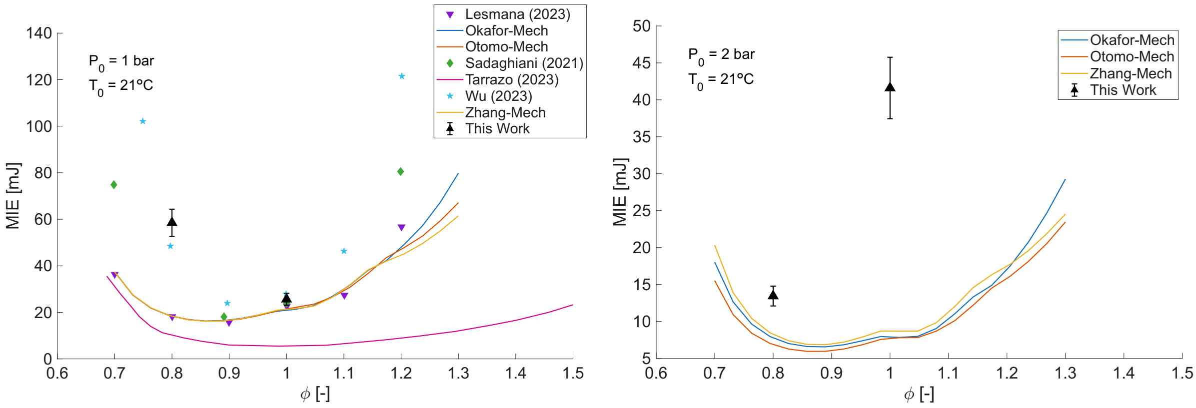

(a) MIE vs ϕ at $T_0 = 21°$ C and $p_0 = 1$ bar **(b)** MIE vs ϕ at $T_0 = 21°$ C and $p_0 = 1$ bar

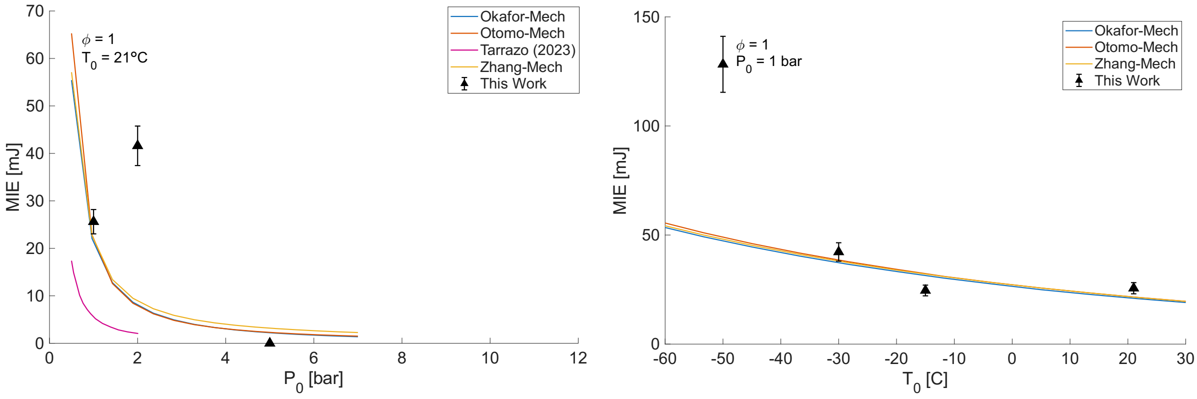

(c) MIE vs p at $T_0 = 21°$ C and $\phi = 1$ **(d)** MIE vs T at $p_0 = 1$ bar and $\phi = 1$

Figure 6.14: MIE of NH_3 at various conditions

gas. Overall, these thresholds provide a coherent qualitative picture of ignition sensitivity with temperature, even though the present instrumentation does not support a defensible quantitative MIE calculation.

7

Conclusions

The objective of this work was the characterisation of ignition behaviour and early flame propagation for ammonia and ammonia blends under subzero temperature and elevated-pressure conditions. This focus addresses a clear gap in the available experimental literature, with conditions more representative of practical combustion devices. The motivation lies in the increasing need to evaluate low-carbon fuels that could replace conventional hydrocarbons in sectors such as transportation and power generation, where the operating envelope can include both high pressures and potentially cold-start or cold-ambient conditions. In these applications, fuel suitability is not determined only by steady-state performance but also by whether reliable ignition and stable early propagation can be achieved across the full range of expected conditions. Understanding these limiting phases is especially important for ammonia, whose low reactivity and slow flame speeds can create considerable operational constraints.

For this reason, the work targeted the initial stage of combustion, which is well captured by two key parameters: the laminar burning velocity, a measure of intrinsic flame propagation, and the minimum ignition energy, which captures ignition difficulty and robustness. Due to their near independence from geometrical parameters, these quantities, together, provide a practical basis for comparing fuel performance. To perform these experiments, a custom constant-volume combustion chamber capable of subzero and high-pressure operation was designed and manufactured, and its performance and methodology were verified and validated before application to the conditions of interest. On this basis, three research questions were formulated to structure the investigation and are answered here.

7.1. Research questions

How should a combustion chamber be designed to enable reliable ignition measurements at high pressure and subzero initial temperatures while ensuring measurement fidelity, thermal control, and safe integration of instrumentation?

Only a limited number of literature setups meet the thermodynamic conditions of interest in this work, as it requires both a pressure boundary that remains mechanically reliable and a configuration that maintains diagnostic quality during rapid transients. For this reason, the first part of this work addressed the design research question directly, focusing on how to realise a chamber that can be operated at high initial pressures and subzero initial temperatures while still enabling repeatable ignition measurements. The main requirement was to withstand the worst-case operating envelope, defined by $p_0 = 50\,\text{bar}$ and $T_0 = -50^\circ$ C, while integrating sensors in a way that minimises intrusion and does not compromise the flow and flame development during the early phase. Because the pressure requirement is the dominant constraint and optical access introduces additional sealing and stress-management issues under thermal cycling, the design was intentionally developed around a non-optical chamber. This allowed the mechanical and measurement requirements to be met with fewer penetrations and more robust interfaces, while maintaining chamber volume and geometry compatible with pressure-trace-based analysis.

Starting with material selection, Duplex SS 2205 was chosen for its strong balance of high yield strength, good low-temperature performance, and chemical compatibility with ammonia-containing environments. This selection ensured that the pressure boundary could be sized conservatively while limiting unnecessary mass and maintaining a practical external envelope for integration in the cooling system. The chamber was then designed as a thick-walled constant-volume vessel with a geometry that supports the assumption of near-spherical early flame expansion. Port placement and internal features were kept as simple as possible to avoid creating preferential flow paths or strong asymmetries that would bias the first milliseconds of propagation.

Measurement fidelity was ensured by treating instrumentation as part of the chamber design. This includes a high-speed dynamic pressure sensor, which is of great importance for calculating the LBV, as this is fully dependent on the combustion pressure rise. Moreover, static temperature and pressure sensors are considered to accurately measure the initial conditions. The ignition system was implemented with an electrode geometry and mounting that prioritise a central kernel and repeatable breakdown, and it was designed to allow adjustment of the electrode gap when ignition becomes marginal at low temperatures. Together, these choices ensured that the chamber could be operated in a way that produces pressure traces dominated by combustion physics and reduced the probability of integration-induced asymmetries or measuring artefacts.

Thermal control was implemented as a system-level requirement because subzero testing is only meaningful if the initial condition is both stable and repeatable. The chamber design, therefore, favoured an external cooling approach capable of providing long soak times and maintaining uniform wall temperature so that the ignition event does not occur in a transient thermal field. This was particularly important near the electrodes, where local heat losses can determine whether a kernel survives and where small temperature gradients can increase shot-to-shot variability. Finally, safe integration was ensured by designing every penetration, seal, and connector as part of the pressure envelope and by minimising the number of interfaces exposed to both ammonia and thermal cycling. This resulted in a chamber that can achieve the required p_0 and T_0 envelope, supports repeatable ignition and early-flame measurements through high-quality pressure diagnostics, and remains robust to the operational realities of subzero, high-pressure testing with multiple sensors and high-voltage hardware. Due to the final absence of the pressurising component, the analysis is limited only to the subzero range.

How do subzero temperatures affect the minimum ignition energy and laminar burning velocity of ammonia?

This thesis experimentally demonstrated, for the first time, a decrease in the burning velocity with decreasing temperature for ammonia. In fact, similar to the trends shown for other fuels, including hydrogen, subzero temperatures led to clear degradation of ammonia combustibility in this work, as evidenced by a strong reduction in the laminar burning velocity and a marked increase in ignition difficulty. For pure NH_3 at p_0 = 1 bar, the laminar burning velocity decreased as T_0 was reduced from 21° C down to -50° C across the explored equivalence ratios, with the largest effect shown for stoichiometric mixtures. For this equivalence ratio, the measured values at -50° C fell to about 4 cm/s, with the flame becoming slow enough that heat losses, buoyancy, and stretch sensitivity had a much larger influence on the pressure trace and on the extracted velocity. The largest step decrease in velocity is found to occur between ambient temperature and 0° C, indicating that the two-phase behaviour of some radicals and products causes a large impact on the magnitude of the flame propagation, with the effect of temperature becoming less impactful below this value.

With regard to the kinetic models used for comparison with the experimental results, the temperature dependence was also steeper in the experiments than in the adiabatic simulations, consistent with non-adiabatic suppression becoming more dominant as the flame slows at low temperatures. However, the general trends of these models do predict the behaviour of the experimental data quite well, which, given the lack of validation data, puts them in a good position for use in the prediction of velocity in subzero conditions. However, the three models used tend to diverge as the mixture becomes richer, and the experimental results consistently yield values on the lower end of these predictions, with a more pronounced difference compared to leaner conditions. Overall, the most important result is that the actual velocity decrease is steeper than that predicted by these models, most likely due to higher heat losses in real systems at these temperatures. Due to limited time, the number of considered conditions per case ranges from 2 to 4, resulting in a degree of uncertainty that should be reduced in future work.

For minimum ignition energy, subzero operation increased the required ignition intensity, but a quantitative MIE trend could not be established with confidence because the voltage and current measurements were frequently corrupted and inconsistent between nominally identical runs. Despite this limitation, the primary-coil ignition thresholds (voltage and dwell time) provided a robust qualitative indicator of the temperature effect. At stoichiometric NH_3 and ambient conditions, ignition was achieved with minimum settings of about 14 V and 1 ms, whereas at -30° C the minimum required settings rose to about 33 V and 8 ms, which corresponded to the hardware limits of the power supply and dwell time. At 0° C and -50° C, reliable ignition at the original 2 mm gap could not be achieved within those limits and required a reduced electrode spacing of roughly 0.5 mm, after which ignition occurred at much lower energy, as per the case of -50° C, where ignition was obtained at about 22 V and 5 ms. These observations show that subzero temperatures substantially increase the ignition requirement for ammonia, to the point that geometry changes were needed to maintain ignition at the coldest condition, even though the exact MIE in mJ could not be reported with high accuracy.

How much does blending ammonia with hydrogen at subzero temperatures and high pressures affect combustion performance?

Based on these novel results, blending ammonia with hydrogen at subzero temperatures produced a large and immediate improvement in combustion performance, primarily by increasing laminar burning velocity and making ignition more robust. In particular, at T_0 = -50° C and p_0 = 1 bar, adding 20% H_2 shifted the burning velocity upward by several cm/s across the tested equivalence ratios, and at $\phi = 1$ it increased LBV by about 4 cm/s, bringing the value close to the room-temperature LBV of neat ammonia. Increasing the blend to 30% H_2 produced a second-step change, with the stoichiometric LBV rising to roughly 11.5 cm/s, a distinct propagation regime compared with neat NH_3 at the same temperature. This improvement is not just a small correction, because at subzero conditions the neat-ammonia flame is slow enough that heat losses, stretch, and buoyancy strongly penalise early propagation, whereas hydrogen addition moves the system away from that quenching-prone regime and yields stronger, cleaner pressure rises. In practical terms, the blends reduce the "cold penalty" substantially and can recover, or even exceed, room-temperature-like burning velocities at the coldest condition tested.

The behaviour of blended fuels at subzero temperatures is consistent with the established room-temperature literature trends. Room-temperature data show a strongly non-linear increase in laminar burning velocity with hydrogen mole fraction, with modest increases at low x_{H_2} and a sharper acceleration as hydrogen becomes more prominent. The present subzero results follow the same qualitative pattern and align with the simulated curvature, but the operational impact is amplified at low temperature because the neat-ammonia baseline is very low. As a result, at -50° C, a given absolute increase in LBV corresponds to a much larger fractional improvement, and it directly translates into better ignition success and reduced sensitivity to losses during kernel growth. In that sense, the subzero blending behaviour relates directly to the higher-temperature behaviour in its underlying mechanism, but it becomes more valuable as a mitigation strategy because it moves the system away from a regime where non-ideal effects dominate, and ignition becomes unreliable.

7.2. General outlook and implications

From the results, it is clear that subzero temperatures impose a substantial penalty on the early combustion performance of ammonia, in both flame propagation and ignition robustness. The measured laminar burning velocities decrease significantly, making the flame more vulnerable to stretch and losses. In parallel, ignition becomes markedly more demanding, to the point that, in some cases, maintaining reliable ignition requires modifying the ignition configuration. Taken together, these outcomes imply that, under subzero initial conditions, neat ammonia is close to the practical boundary where ignition becomes intermittent and early propagation becomes highly sensitive to disturbances. In an engine context, this translates into increased misfire risk, greater cycle-to-cycle variability, and a narrower operability window, particularly during cold starts and at lean conditions, where ammonia is already inherently slow.

For this reason, the results support the conclusion that blending is a practical mitigation strategy when ammonia is used in cold environments, at least during the ignition and early propagation phases. Hydrogen blending, in particular, showed a strong benefit at subzero temperatures, with modest additions producing a large recovery in burning velocity and noticeably improving the margin against quenching. This suggests a realistic operational concept in which a small fraction of a more reactive fuel, like hydrogen, is used to ensure a robust kernel and to carry the initial flame development through the most loss-sensitive period, after which the energy release can transition toward ammonia-dominant combustion. Such a staged strategy is consistent with how ignition assist is implemented in other difficult-to-ignite systems, and it is especially relevant here because the early phase is where ammonia is most penalised by low temperature.

Hydrogen is an attractive blending partner not only because it can be produced sustainably, but also because its combustion characteristics complement ammonia, addressing the observed limitations directly. Hydrogen's high reactivity and favourable transport properties accelerate radical production and heat release during kernel growth, helping overcome the non-adiabatic losses that dominate at low temperatures and low burning velocity. At the same time, the experiments indicate that even modest hydrogen fractions can cancel a large part of the low-temperature penalty, recovering burning velocities comparable to those of neat ammonia at room temperature under the tested conditions. This is a key practical finding, because it indicates that subzero operability does not necessarily require hydrogen-dominant mixtures, but can be achieved while keeping the blend ammonia-rich. In addition, faster early propagation improves signal quality and repeatability in the constant-volume context, which implies that blending may also simplify control and diagnostics in practical devices by making ignition events less marginal. These results therefore justify further exploration of blends as a means to widen the ignition and stability envelope of ammonia at low temperature and high pressure.

8

Recommendations

Based on the conclusions presented, this chapter outlines the main recommendations and directions for future work, taking into account the limitations encountered and the opportunities identified during the campaign. The overarching goal is to increase measurement robustness so that additional parameters can be quantified with even higher confidence and to extend the operating envelope toward higher, more engine-relevant pressures and more representative combustion regimes. The highest-priority improvement is the minimum ignition energy methodology, since the current and voltage measurements were not sufficiently repeatable to support a reliable quantitative analysis. The main limitation is that the inferred energy depends strongly on the integration window and on waveforms that were often corrupted by electrical noise and non-physical artefacts. Future work should therefore treat MIE as a dedicated measurement system, improve shielding and grounding, ensure stable probe mounting, and enforce a consistent, physically motivated discharge-window definition for integration. With these changes, ignition trends could be quantified credibly across temperature, pressure, and mixture composition and compared directly with the burning-velocity results. At the same time, the accuracy of these latter results should be improved by running a larger set of experiments to reduce the error bars accompanying the measurements.

A second key recommendation is to treat the spark gap as a controlled experimental parameter, rather than an adjustment to apply only when ignition becomes problematic. In the present work, the electrode gap had to be reduced to achieve ignition for NH_3 at 0° C and at -50° C, and this change is likely to have influenced early kernel development and, in turn, the pressure-based velocity extraction. If this parameter is not controlled explicitly, it becomes difficult to separate the true temperature effect on combustion from an ignition-system effect that varies only at specific setpoints. Future campaigns should therefore fix the electrode gap across the entire NH_3 subzero campaign and accept that some conditions may not ignite within the available energy limits, yielding a clean interpretation of temperature trends and ignition limits.

Once ignition diagnostics and configuration control are improved, the next step is to expand the hydrogen blending campaign beyond the two blend fractions investigated at the coldest condition. The present results already show that hydrogen blending provides a large performance gain at -50° C, with a strong increase in burning velocity between 20% and 30% H_2, but two points cannot define the shape of the response curve. Extending the blend sweep to intermediate and higher hydrogen fractions would allow the improvement to be described in terms of curvature and sensitivity, enabling determination of whether the benefit is approximately linear, saturating, or strongly non-linear in the subzero regime. A denser sweep would also improve model comparison, since kinetic mechanisms can appear equally plausible when only two blend ratios are available but diverge significantly when curvature and threshold behaviour are resolved.

A parallel priority is to confirm design consistency and safe handling as the operating envelope is extended toward higher pressures. Although the chamber and supporting hardware were designed to withstand severe conditions, the main validation and novel datasets were acquired at lower pressures, and a dedicated high-pressure campaign is required to ensure that the filling procedure, sealing behaviour, and instrumentation remain reliable as p_0 increases. High-pressure operation should also be used to reassess the diagnostics, because faster pressure rises and higher absolute pressures can change the effective signal bandwidth requirements and can expose resonance or mounting sensitivity that is not visible at 1 bar. In addition, the practical handling of ammonia at high pressure requires specific attention to purging effectiveness, exhaust-line condensation, and operational contingencies, particularly during long subzero holds where water and other condensables can accumulate in the system. Establishing these aspects would allow the current subzero conclusions to be confidently extended to a broader range of conditions.

Finally, while this thesis appropriately focuses on ignition and early laminar propagation, fully assessing the feasibility of ammonia in real systems requires additional parameters beyond these initial metrics. Practical combustors operate with turbulence, longer burn durations, and strong interactions with confinement and heat transfer, all of which can shift the performance relative to early laminar behaviour and can determine stability and operability in practice. This extension is particularly relevant for blends, since hydrogen may affect not only laminar burning velocity but also quenching propensity, flame stability, and sensitivity to stretch and instabilities as the flame grows. In addition, emissions were not addressed in this work, but they are indispensable for any realistic assessment of ammonia as a sustainable fuel. In particular, NO_x production, especially the fuel-based portion, must be quantified and compared against the carbon-emission benefit to determine whether improved sustainability is achieved in practice, rather than only in fuel chemistry. For these reasons, to fully define ammonia's operation and understand its actual feasibility as a fuel alternative, future studies should combine early-phase metrics with later-phase combustion behaviour and emissions characterisation to provide a complete, application-oriented picture of ammonia and ammonia–hydrogen operation.

References

[1] Victor Nikolaevich Kondratiev. *Combustion | Definition, Reaction, Analysis, & Facts | Britannica*. Encyclopedia Britannica. 2025. URL: `https://www.britannica.com/science/combustion` (visited on 01/29/2026).

[2] Hannah Ritchie, Pablo Rosado, and Max Roser. "Energy Production and Consumption". In: *Our World in Data* (July 10, 2020). URL: `https://ourworldindata.org/energy-production-consumption` (visited on 09/09/2025).

[3] IEA. *World Energy Outlook 2024*. Paris: IEA, 2024. URL: `https://www.iea.org/reports/world-energy-outlook-2024`.

[4] Katharina Kohse-Höinghaus. "Combustion in the future: The importance of chemistry". In: *Proceedings of the Combustion Institute. International Symposium on Combustion* (Sept. 25, 2020). ISSN: 1540-7489. DOI: `10.1016/j.proci.2020.06.375`. URL: `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7518234/` (visited on 09/09/2025).

[5] Xiangmei Meng, Wiebren de Jong, and Tadeusz Kudra. "A state-of-the-art review of pulse combustion: Principles, modeling, applications and R&D issues". In: *Renewable and Sustainable Energy Reviews* 55 (Mar. 1, 2016), pp. 73–114. ISSN: 1364-0321. DOI: `10.1016/j.rser.2015.10.110`. URL: `https://www.sciencedirect.com/science/article/pii/S1364032115011892` (visited on 09/09/2025).

[6] Grand View Research. *Ammonia Market (2025 - 2030)*. 2025.

[7] R. C. Murray and A. R. Hall. "Flame speeds in hydrazine vapour and in mixtures of hydrazine and ammonia with oxygen". In: *Transactions of the Faraday Society* 47.0 (Jan. 1, 1951), pp. 743–751. ISSN: 0014-7672. DOI: `10.1039/TF9514700743`. URL: `https://pubs.rsc.org/en/content/articlelanding/1951/tf/tf9514700743` (visited on 01/17/2026).

[8] Akihiro Hayakawa et al. "Laminar burning velocity and Markstein length of ammonia/air premixed flames at various pressures". In: *Fuel* 159 (Nov. 1, 2015), pp. 98–106. ISSN: 0016-2361. DOI: `10.1016/j.fuel.2015.06.070`. URL: `https://www.sciencedirect.com/science/article/pii/S001623611500650X` (visited on 01/19/2026).

[9] Alka Karan et al. "Laminar flame speeds of ammonia mixtures at high pressure and temperature conditions: New experimental results and performance of different kinetic models". In: *Proceedings of the 17th International Conference on Flow Dynamics*. Proceedings of the 17th International Conference on Flow Dynamics. Sendai, Japan: Institute of Fluid Science, Tohoku University, Oct. 2020. URL: `https://hal.science/hal-03556828` (visited on 09/09/2025).

[10] Luis F. Alvarez et al. "Laminar burning velocity of Ammonia/Air mixtures at high pressures". In: *Fuel* 363 (May 1, 2024), p. 130986. ISSN: 0016-2361. DOI: `10.1016/j.fuel.2024.130986`. URL: `https://www.sciencedirect.com/science/article/pii/S0016236124001327` (visited on 01/19/2026).

[11] Jun Li et al. “A Review on Combustion Characteristics of Ammonia as a Carbon-Free Fuel”. In: *Frontiers in Energy Research* 9 (Oct. 6, 2021). ISSN: 2296-598X. DOI: `10.3389/fenrg.2021.760356`. URL: `https://www.frontiersin.orghttps://www.frontiersin.org/journals/energy-research/articles/10.3389/fenrg.2021.760356/full` (visited on 04/11/2025).

[12] Herry Lesmana et al. “NH3 as a Transport Fuel in Internal Combustion Engines: A Technical Review”. In: *Journal of Energy Resources Technology* 141.70703 (Mar. 11, 2019). ISSN: 0195-0738. DOI: `10.1115/1.4042915`. URL: `https://doi.org/10.1115/1.4042915` (visited on 04/07/2025).

[13] A. Valera-Medina et al. “Review on Ammonia as a Potential Fuel: From Synthesis to Economics”. In: *Energy & Fuels* 35.9 (May 6, 2021), pp. 6964–7029. ISSN: 0887-0624. DOI: `10.1021/acs.energyfuels.0c03685`. URL: `https://doi.org/10.1021/acs.energyfuels.0c03685` (visited on 04/25/2025).

[14] “Energy Systems”. In: *Climate Change 2022 - Mitigation of Climate Change*. Ed. by Intergovernmental Panel On Climate Change (Ipcc). 1st ed. Cambridge University Press, Aug. 17, 2023, pp. 613–746. ISBN: 978-1-009-15792-6. DOI: `10.1017/9781009157926.008`. URL: `https://www.cambridge.org/core/product/identifier/9781009157926%23c6/type/book_part` (visited on 01/18/2026).

[15] Chung K. Law. *Combustion physics*. In collab. with Internet Archive. Cambridge : Cambridge University Press, 2010. 750 pp. ISBN: 978-0-521-15421-5. URL: `http://archive.org/details/combustionphysic0000lawc` (visited on 05/26/2025).

[16] Niklas Zettervall et al. “Evaluation of Chemical Kinetic Mechanisms for Methane Combustion: A Review from a CFD Perspective”. In: *Fuels* 2.2 (May 24, 2021), pp. 210–240. ISSN: 2673-3994. DOI: `10.3390/fuels2020013`. URL: `https://www.mdpi.com/2673-3994/2/2/13` (visited on 01/18/2026).

[17] Alexander Starik, Nataliya Titova, and Boris Loukhovitski. “Control of combustion and detonation by means of resonance laser radiation. Analysis and potentialities”. In: Proceedings of SPIE - The International Society for Optical Engineering. Vol. 4948. Aug. 1, 2003. DOI: `10.1117/12.516937`.

[18] Ram B. Gupta. *Hydrogen Fuel: Production, Transport, and Storage*. Google-Books-ID: UlVhpI0XnRcC. CRC Press, July 30, 2008. 626 pp. ISBN: 978-1-4200-4577-2.

[19] Mohamed A. Habib et al. “Hydrogen combustion, production, and applications: A review”. In: *Alexandria Engineering Journal* 100 (Aug. 1, 2024), pp. 182–207. ISSN: 1110-0168. DOI: `10.1016/j.aej.2024.05.030`. URL: `https://www.sciencedirect.com/science/article/pii/S1110016824005003` (visited on 05/18/2025).

[20] Hideaki Kobayashi et al. “Science and technology of ammonia combustion”. In: *Proceedings of the Combustion Institute* 37.1 (Jan. 1, 2019), pp. 109–133. ISSN: 1540-7489. DOI: `10.1016/j.proci.2018.09.029`. URL: `https://www.sciencedirect.com/science/article/pii/S1540748918306345` (visited on 04/11/2025).

[21] Bao-zhi Jin et al. “Experimental and numerical study of the laminar burning velocity of NH3/H2/air premixed flames at elevated pressure and temperature”. In: *International Journal of Hydrogen Energy* 47.85 (Oct. 15, 2022), pp. 36046–36057. ISSN: 0360-3199. DOI: `10.1016/j.ijhydene.2022.08.198`. URL: `https://www.sciencedirect.com/science/article/pii/S0360319922038046` (visited on 01/18/2026).

[22] X. He et al. “Auto-ignition kinetics of ammonia and ammonia/hydrogen mixtures at intermediate temperatures and high pressures”. In: *Combustion and Flame* 206 (Aug. 1, 2019), pp. 189–200. ISSN: 0010-2180. DOI: `10.1016/j.combustflame.2019.04.050`. URL: `https://www.sciencedirect.com/science/article/pii/S0010218019302020` (visited on 01/18/2026).

[23] B. A. Younglove and J. F. Ely. “Thermophysical Properties of Fluids. II. Methane, Ethane, Propane, Isobutane, and Normal Butane”. In: *Journal of Physical and Chemical Reference Data* 16.4 (Oct. 1, 1987), pp. 577–798. ISSN: 0047-2689. DOI: `10.1063/1.555785`. URL: `https://doi.org/10.1063/1.555785` (visited on 05/22/2025).

[24] Mahdi Faghih and Zheng Chen. “The constant-volume propagating spherical flame method for laminar flame speed measurement”. In: *Science Bulletin* 61.16 (Aug. 1, 2016), pp. 1296–1310. ISSN: 2095-9273. DOI: `10.1007/s11434-016-1143-6`. URL: `https://www.sciencedirect.com/science/article/pii/S209592731630055X` (visited on 04/17/2025).

[25] Oras Obayes and Haroun Shahad. “A Review of Laminar Burning Velocity and Flame Speed of Gases and Liquid Fuels”. In: 7 (Feb. 1, 2017), pp. 183–197.

[26] Yu Xie et al. “Intrinsic cellular instabilities of hydrogen laminar outwardly propagating spherical flames”. In: *Fuel* 327 (Nov. 2022), p. 125149. ISSN: 00162361. DOI: `10.1016/j.fuel.2022.125149`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0016236122019901` (visited on 01/19/2026).

[27] X. J. Gu et al. “Laminar burning velocity and Markstein lengths of methane–air mixtures”. In: *Combustion and Flame* 121.1 (Apr. 1, 2000), pp. 41–58. ISSN: 0010-2180. DOI: `10.1016/S0010-2180(99)00142-X`. URL: `https://www.sciencedirect.com/science/article/pii/S001021809900142X` (visited on 12/01/2025).

[28] M. Elia, M. Ulinski, and M. Metghalchi. “Laminar Burning Velocity of Methane–Air–Diluent Mixtures”. In: *Journal of Engineering for Gas Turbines and Power* 123.1 (June 23, 2000), pp. 190–196. ISSN: 0742-4795. DOI: `10.1115/1.1339984`. URL: `https://doi.org/10.1115/1.1339984` (visited on 05/16/2025).

[29] Mohd Suardi Suhaimi et al. “FLAME PROPAGATION AND BURNING RATES OF METHANE-AIR MIXTURES USING SCHLIEREN PHOTOGRAPHY”. In: *Jurnal Teknologi (Sciences & Engineering)* 78.10 (Oct. 4, 2016). ISSN: 2180-3722. DOI: `10.11113/jt.v78.9664`. URL: `https://journals.utm.my/jurnalteknologi/article/view/9664` (visited on 01/19/2026).

[30] G. Langella et al. “Ammonia as a Fuel for Gas Turbines: Perspectives and Challenges”. In: *Journal of Physics* (2023).

[31] Huizhen Li, Huahua Xiao, and Jinhua Sun. “Laminar burning velocity, Markstein length, and cellular instability of spherically propagating NH3/H2/Air premixed flames at moderate pressures”. In: *Combustion and Flame* 241 (July 1, 2022), p. 112079. ISSN: 0010-2180. DOI: `10.1016/j.combustflame.2022.112079`. URL: `https://www.sciencedirect.com/science/article/pii/S0010218022000980` (visited on 04/07/2025).

[32] Minh Tien Nguyen et al. “Effect of Ignition Energy and Hydrogen Addition on Laminar Flame Speed, Ignition Delay Time, and Flame Rising Time of Lean Methane/Air Mixtures”. In: *Energies* 15.5 (Jan. 2022). Number: 5, p. 1940. ISSN: 1996-1073. DOI: `10.3390/en15051940`. URL: `https://www.mdpi.com/1996-1073/15/5/1940` (visited on 04/17/2025).

[33] Stig R. Sellevåg, Yuri Georgievskii, and James A. Miller. "The Temperature and Pressure Dependence of the Reactions H + O2 (+M) → HO2 (+M) and H + OH (+M) → H2O (+M)". In: *The Journal of Physical Chemistry A* 112.23 (June 1, 2008), pp. 5085–5095. ISSN: 1089-5639. DOI: `10.1021/jp711800z`. URL: `https://doi.org/10.1021/jp711800z` (visited on 01/19/2026).

[34] Alexander A. Konnov et al. "A comprehensive review of measurements and data analysis of laminar burning velocities for various fuel+air mixtures". In: *Progress in Energy and Combustion Science* 68 (Sept. 1, 2018), pp. 197–267. ISSN: 0360-1285. DOI: `10.1016/j.pecs.2018.05.003`. URL: `https://www.sciencedirect.com/science/article/pii/S0360128517301417` (visited on 01/19/2026).

[35] Khizer Saeed and C. R. Stone. "Measurements of the laminar burning velocity for mixtures of methanol and air from a constant-volume vessel using a multizone model". In: *Combustion and Flame* 139.1 (Oct. 1, 2004), pp. 152–166. ISSN: 0010-2180. DOI: `10.1016/j.combustflame.2004.08.008`. URL: `https://www.sciencedirect.com/science/article/pii/S0010218004001610` (visited on 01/19/2026).

[36] M. Metghalchi and J. C. Keck. "Laminar burning velocity of propane-air mixtures at high temperature and pressure". In: *Combustion and Flame* 38 (Jan. 1, 1980), pp. 143–154. ISSN: 0010-2180. DOI: `10.1016/0010-2180(80)90046-2`. URL: `https://www.sciencedirect.com/science/article/pii/0010218080900462` (visited on 05/16/2025).

[37] R. Stone, A. Clarke, and P. Beckwith. "Correlations for the Laminar-Burning Velocity of Methane/Diluent/Air Mixtures Obtained in Free-Fall Experiments". In: *Combustion and Flame* 114.3 (Aug. 1, 1998), pp. 546–555. ISSN: 0010-2180. DOI: `10.1016/S0010-2180(97)00329-5`. URL: `https://www.sciencedirect.com/science/article/pii/S0010218097003295` (visited on 05/16/2025).

[38] D. Bradley and G. F. Hundy. "Burning velocities of methane-air mixtures using hot-wire anemometers in closed-vessel explosions". In: *Symposium (International) on Combustion*. Thirteenth symposium (International) on Combustion 13.1 (Jan. 1, 1971), pp. 575–583. ISSN: 0082-0784. DOI: `10.1016/S0082-0784(71)80059-0`. URL: `https://www.sciencedirect.com/science/article/pii/S0082078471800590` (visited on 05/15/2025).

[39] Thierry Poinsot and Denis Veynante. *Theoretical and Numerical Combustion*. Google-Books-ID: cqFDkeVABYoC. R.T. Edwards, Inc., 2005. 544 pp. ISBN: 978-1-930217-10-2.

[40] J.A. VAN OIJEN and L.P.H. DE GOEY. "Modelling of Premixed Laminar Flames using Flamelet-Generated Manifolds". In: *Combustion Science and Technology* 161.1 (Dec. 1, 2000). _eprint: https://doi.org/10.1080/00102200008935814, pp. 113–137. ISSN: 0010-2202. DOI: `10.1080/00102200008935814`. URL: `https://doi.org/10.1080/00102200008935814` (visited on 01/19/2026).

[41] V. Moureau, B. Fiorina, and H. Pitsch. "A level set formulation for premixed combustion LES considering the turbulent flame structure". In: *Combustion and Flame* 156.4 (Apr. 1, 2009), pp. 801–812. ISSN: 0010-2180. DOI: `10.1016/j.combustflame.2009.01.019`. URL: `https://www.sciencedirect.com/science/article/pii/S0010218009000224` (visited on 01/19/2026).

[42] Theodore L. Brown. *Chemistry: the central science*. 12th ed. OCLC: 670248458. Boston: Prentice Hall, 2012. ISBN: 978-0-321-69672-4.

[43] Herry Lesmana et al. “An experimental investigation into the effect of spark gap and duration on minimum ignition energy of partially dissociated NH3 in air”. In: *Combustion and Flame* 241 (July 2022), p. 112053. ISSN: 00102180. DOI: `10.1016/j.combustflame.2022.112053`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0010218022000724` (visited on 03/27/2025).

[44] Bernard Lewis and Guenther Von Elbe. *Combustion, Flames and Explosions of Gases*. 3rd ed. New York: Academic Press, June 4, 1987. ISBN: 978-0-12-446751-4. URL: `https://shop.elsevier.com/books/combustion-flames-and-explosions-of-gases/lewis/978-0-12-446751-4` (visited on 04/25/2025).

[45] Behlol Nawaz and J. Hunter Mack. “Calculating minimum ignition energy of ammonia-air mixtures at sub-zero temperatures”. In: 14th U.S. National Combustion Meeting. Boston, Mar. 16, 2025.

[46] Eiichi Murase et al. “A Study on Equivalence-Ratio Dependence of Minimum Ignition Energy Based on Initial Burning Velocity”. In: *TRANSACTIONS OF THE JAPAN SOCIETY OF MECHANICAL ENGINEERS Series B* 79.805 (2013), pp. 1839–1847. ISSN: 0387-5016, 1884-8346. DOI: `10.1299/kikaib.79.1839`. URL: `https://www.jstage.jst.go.jp/article/kikaib/79/805/79_1839/_article/-char/ja/` (visited on 04/22/2025).

[47] E. Fernández-Tarrazo, R. Gómez-Miguel, and M. Sánchez-Sanz. “Minimum ignition energy of hydrogen–ammonia blends in air”. In: *Fuel* 337 (Apr. 2023), p. 127128. ISSN: 00162361. DOI: `10.1016/j.fuel.2022.127128`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0016236122039527` (visited on 08/29/2025).

[48] Peng Chen et al. “Initial flame propagation characteristics of the hydrogen spherical premixed flame”. In: *International Journal of Hydrogen Energy* 48.92 (Nov. 2023), pp. 36112–36121. ISSN: 03603199. DOI: `10.1016/j.ijhydene.2023.05.348`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0360319923027659` (visited on 04/22/2025).

[49] Sameer Vidyadhar Bondre. “Effects of Pressure and Equivalence Ratio on Minimum Ignition Energies of Methane-Air Mixtures and Effect of Pressure on Breakdown Threshold Energies of Methane and Air”. In: *TRACE:: Tennessee Research and Creative Exchange* (2004).

[50] Christophe Proust. “A new technique to produce well controlled electrical sparks. Application to MIE measurements”. In: 13th International symposium on hazards, prevention, and mitigation of industrial explosions (ISHPMIE). Brawnschewig, Germany, 2021.

[51] Ryo Ono et al. “Minimum ignition energy of hydrogen–air mixture: Effects of humidity and spark duration”. In: *Journal of Electrostatics* 65.2 (Feb. 2007), pp. 87–93. ISSN: 03043886. DOI: `10.1016/j.elstat.2006.07.004`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0304388606000726` (visited on 04/24/2025).

[52] Christophe Proust and Didier Jamois. “Some fundamental combustion properties of “cryogenic” premixed hydrogen air flames”. In: *ID32*. 9th International Conference on Hydrogen Safety (ICHS 2021). 2021.

[53] D. Cirrone et al. “Minimum ignition energy of hydrogen-air mixtures at ambient and cryogenic temperatures”. In: *International Journal of Hydrogen Energy* 48.43 (May 2023), pp. 16530–16544. ISSN: 03603199. DOI: `10.1016/j.ijhydene.2023.01.115`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0360319923001829` (visited on 04/03/2025).

[54] Chunwei Wu et al. “Numerical and experimental studies on minimum ignition energies in primary reference fuel/air mixtures”. In: *Proceedings of the Combustion Institute* 39.2 (Jan. 1, 2023), pp. 1987–1996. ISSN: 1540-7489. DOI: `10.1016/j.proci.2022.08.043`. URL: `https://www.sciencedirect.com/science/article/pii/S1540748922003352` (visited on 05/07/2025).

[55] D. P. Mishra. *Experimental Combustion: An Introduction*. 0th ed. CRC Press, May 12, 2014. ISBN: 978-0-429-08703-5. DOI: `10.1201/b17027`. URL: `https://www.taylorfrancis.com/books/9781466517363` (visited on 05/26/2025).

[56] C Nandhini and M Jagadeeswari. “MEASUREMENT OF TEMPERATURE USING RTD AND SOFTWARE SIGNAL CONDITIONING USING LABVIEW”. In: 20.2 (2016).

[57] *Motores de combustión interna alternativos*. In collab. with Francisco Payri González and José María Desantes Fernández. 1ª edición, 6ª reimpresión. OCLC: 1322973895. Valencia, Barcelona: Universitat Politècnica de València ; Reverté, 2017. ISBN: 978-84-8363-705-0.

[58] Anh Vang Tran, Xianmin Zhang, and Benliang Zhu. “Mechanical Structural Design of a Piezoresistive Pressure Sensor for Low-Pressure Measurement: A Computational Analysis by Increases in the Sensor Sensitivity”. In: *Sensors* 18.7 (June 24, 2018). ISSN: 1424-8220. DOI: `10.3390/s18072023`. URL: `https://www.mdpi.com/1424-8220/18/7/2023` (visited on 01/21/2026).

[59] Jacob Fraden. *Handbook of Modern Sensors*. 4th ed. Springer, 2010. ISBN: 978-1-4419-6465-6. DOI: `10.1007/978-1-4419-6466-3`.

[60] Mayuri Goswami et al. “The effect of elevated pressures on the laminar burning velocity of methane + air mixtures”. In: *Combustion and Flame* 160.9 (Sept. 1, 2013), pp. 1627–1635. ISSN: 0010-2180. DOI: `10.1016/j.combustflame.2013.03.032`. URL: `https://www.sciencedirect.com/science/article/pii/S0010218013001326` (visited on 01/21/2026).

[61] Guido Saccone and Marco Marini. “Chemical Kinetic Analysis of High-Pressure Hydrogen Ignition and Combustion toward Green Aviation”. In: *Aerospace* 11.2 (Jan. 25, 2024). ISSN: 2226-4310. DOI: `10.3390/aerospace11020112`. URL: `https://www.mdpi.com/2226-4310/11/2/112` (visited on 01/21/2026).

[62] Kazuya Shimizu et al. “Updated kinetic mechanism for high-pressure hydrogen combustion”. In: *Journal of Propulsion and Power* 27.2 (2011), pp. 383–395. ISSN: 0748-4658. DOI: `10.2514/1.48553`. URL: `https://tohoku.elsevierpure.com/en/publications/updated-kinetic-mechanism-for-high-pressure-hydrogen-combustion/` (visited on 01/21/2026).

[63] Xinlei Liu et al. “Evaluation of chemical kinetic mechanisms for hydrogen engine modeling”. In: *Chemical Engineering Journal* 524 (Nov. 15, 2025), p. 169082. ISSN: 1385-8947. DOI: `10.1016/j.cej.2025.169082`. URL: `https://www.sciencedirect.com/science/article/pii/S1385894725099243` (visited on 12/22/2025).

[64] Liming Dai et al. “Experimental and numerical analysis of the autoignition behavior of NH3 and NH3/H2 mixtures at high pressure”. In: *Combustion and Flame* 215 (May 2020), pp. 134–144. ISSN: 00102180. DOI: `10.1016/j.combustflame.2020.01.023`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0010218020300353` (visited on 09/09/2025).

[65] Florian Hurault et al. "Experimental and numerical ignition delay times comparison for ammonia mechanisms at high pressure". In: *Proceedings of the Combustion Institute* 40.1 (Jan. 1, 2024), p. 105625. ISSN: 1540-7489. DOI: `10.1016/j.proci.2024.105625`. URL: `https://www.sciencedirect.com/science/article/pii/S1540748924004334` (visited on 09/09/2025).

[66] Jundie Chen et al. "Effect of hydrogen blending on the high temperature auto-ignition of ammonia at elevated pressure". In: *Fuel* 287 (Mar. 1, 2021), p. 119563. ISSN: 0016-2361. DOI: `10.1016/j.fuel.2020.119563`. URL: `https://www.sciencedirect.com/science/article/pii/S001623612032559X` (visited on 09/09/2025).

[67] Department of Defense. *MIL-STD-810H, Method 502.7: Low Temperature*. 2019. URL: `https://cvgstrategy.com/wp-content/uploads/2019/08/MIL-STD-810H-Method-502.7-Low-Temperature.pdf`.

[68] Swedish Polar Research Secretariat. *Technical Requirements of the New Swedish Heavy Polar Icebreaking Research Vessel*. 2025. URL: `https://www.polar.se/media/a3wk0lud/technical-requirement-240613-rev-4-2025-06-24.pdf`.

[69] Sylvain L. Michaux et al. "Laminar burning velocities of hydrogen-air and methane-air flames from ambient to cryogenic temperatures at different equivalence ratios". In: *International Journal of Hydrogen Energy* 100 (Jan. 27, 2025), pp. 608–616. ISSN: 0360-3199. DOI: `10.1016/j.ijhydene.2024.12.114`. URL: `https://www.sciencedirect.com/science/article/pii/S0360319924053163` (visited on 12/06/2025).

[70] Anupam Ghosh et al. "Laminar burning velocity of hydrogen, methane, ethane, ethylene, and propane flames at near-cryogenic temperatures". In: *Applications in Energy and Combustion Science* 12 (Dec. 2022), p. 100094. ISSN: 2666352X. DOI: `10.1016/j.jaecs.2022.100094`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S2666352X22000371` (visited on 04/22/2025).

[71] L. Coniglio et al. "Combustion chemical kinetics of biodiesel and related compounds (methyl and ethyl esters): Experiments and modeling – Advances and future refinements". In: *Progress in Energy and Combustion Science* 39.4 (Aug. 1, 2013), pp. 340–382. ISSN: 0360-1285. DOI: `10.1016/j.pecs.2013.03.002`. URL: `https://www.sciencedirect.com/science/article/pii/S0360128513000166` (visited on 05/13/2025).

[72] Valentin Soloiu et al. "Constant volume combustion chamber (CVCC) investigations of aerospace F-24 and Jet-A in low-temperature heat release and negative temperature coefficient regions". In: *Energy Conversion and Management* 263 (July 1, 2022), p. 115687. ISSN: 0196-8904. DOI: `10.1016/j.enconman.2022.115687`. URL: `https://www.sciencedirect.com/science/article/pii/S0196890422004836` (visited on 04/18/2025).

[73] Hiroaki Kobayashi et al. "Experimental study on cryo-compressed hydrogen ignition and flame". In: *International Journal of Hydrogen Energy* 45.7 (Feb. 7, 2020), pp. 5098–5109. ISSN: 0360-3199. DOI: `10.1016/j.ijhydene.2019.12.091`. URL: `https://www.sciencedirect.com/science/article/pii/S0360319919346270` (visited on 05/25/2025).

[74] M. Kuznetsov et al. "Flame propagation regimes and critical conditions for flame acceleration and detonation transition for hydrogen-air mixtures at cryogenic temperatures". In: *International Journal of Hydrogen Energy* 47.71 (Aug. 19, 2022), pp. 30743–30756. ISSN: 0360-3199. DOI: `10.1016/j.ijhydene.2022.07.024`. URL: `https://www.sciencedirect.com/science/article/pii/S0360319922030385` (visited on 05/13/2025).

[75] Jac Clarke et al. "Cryogenic Hydrogen Jet and Flame for Clean Energy Applications: Progress and Challenges". In: *Energies* 16.11 (Jan. 2023). Number: 11, p. 4411. ISSN: 1996-1073. DOI: `10.3390/en16114411`. URL: `https://www.mdpi.com/1996-1073/16/11/4411` (visited on 05/25/2025).

[76] J. Würmel and J. M. Simmie. "CFD studies of a twin-piston rapid compression machine". In: *Combustion and Flame* 141.4 (June 1, 2005), pp. 417–430. ISSN: 0010-2180. DOI: `10.1016/j.combustflame.2005.01.015`. URL: `https://www.sciencedirect.com/science/article/pii/S0010218005000519` (visited on 01/29/2026).

[77] Sebastian Schuh et al. "Experimental Investigation and Benchmark Study of Oxidation of Methane–Propane–n-Heptane Mixtures at Pressures up to 100 bar". In: *Energies* 12.18 (Sept. 4, 2019), p. 3410. ISSN: 1996-1073. DOI: `10.3390/en12183410`. URL: `https://www.mdpi.com/1996-1073/12/18/3410` (visited on 01/29/2026).

[78] Herry Lesmana et al. "Experimental and kinetic modelling studies of laminar flame speed in mixtures of partially dissociated NH3 in air". In: *Fuel* 278 (Oct. 2020), p. 118428. ISSN: 00162361. DOI: `10.1016/j.fuel.2020.118428`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0016236120314241` (visited on 04/07/2025).

[79] M P Burke, Y Ju, and F L Dryer. "Effect of Cylindrical Confinement on the Evolution of Outwardly Propagating Flames". In: *Chemical & Physical Processes in Combustion* (2007).

[80] Michael P. Burke, Yiguang Ju, and Frederick L. Dryer. "Effect of flow field perturbations on laminar flame speed determination using spherical flames". In: *AIAA Aerosp. Sci. Meet. Exhib.* 46th AIAA Aerospace Sciences Meeting and Exhibit. Journal Abbreviation: AIAA Aerosp. Sci. Meet. Exhib. 2008. ISBN: 978-1-56347-937-3.

[81] Flexitallic. *Ring Type Joints*. 2022.

[82] Warren C. Young, Richard G. Budynas, and Raymond J. Roark. *Roark's formulas for stress and strain*. 7. ed., reprint. McGraw-Hill international edition, General engineering series. New York: McGraw-Hill, 2011. 852 pp. ISBN: 978-0-07-072542-3.

[83] American Society of Mechanical Engineers. *ASME Boiler and Pressure Vessel Code, Section VIII, Division 1: Rules for Construction of Pressure Vessels*. Version 1968. New York, 1968.

[84] International Organization for Standardization (ISO). *Mechanical properties of fasteners made of carbon steel and alloy steel — Part 1: Bolts, screws and studs with specified property classes — Coarse thread and fine pitch thread*. Version 4. Geneva, Switzerland, 2009. URL: `https://www.iso.org/standard/43383.html`.

[85] European Cooperation for Space Standardization (ECSS). *Space engineering: Threaded fasteners handbook*. ECSS-E-HB-32-23A Rev.1. Noordwijk, The Netherlands: ECSS, Feb. 6, 2023. URL: `https://ecss.nl/wp-content/uploads/2023/02/ECSS-E-HB-32-23A-Rev.1%286February2023%29.pdf` (visited on 01/03/2026).

[86] Walter Pilkey and Deborah Pilkey. "Stress Concentration Analysis and Design". In: *Peterson's Stress Concentration Factors*. John Wiley & Sons, Ltd, 2007, pp. 457–512. ISBN: 978-0-470-21110-6. DOI: `10.1002/9780470211106.ch6`. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470211106.ch6` (visited on 01/03/2026).

[87] Shengnan Li et al. "Effect of low blockage ratio obstacle on explosion characteristic in methane/air mixture". In: *Arabian Journal of Chemistry* 17.9 (Sept. 1, 2024), p. 105890. ISSN: 1878-5352. DOI: `10.1016/j.arabjc.2024.105890`. URL: `https://www.sciencedirect.com/science/article/pii/S1878535224002922` (visited on 01/03/2026).

[88] *Macmillan encyclopedia of physics*. In collab. with Internet Archive. New York : Simon & Schuster Macmillan, 1996. 522 pp. ISBN: 978-0-02-864586-5. URL: `http://archive.org/details/macmillanencyclo0004unse` (visited on 01/04/2026).

[89] Friedrich Paschen. "Ueber die zum Funkenübergang in Luft, Wasserstoff und Kohlensäure bei verschiedenen Drucken erforderliche Potentialdifferenz". In: *Annalen der Physik* 273.5 (1889). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.18892730505, pp. 69–96. ISSN: 1521-3889. DOI: `10.1002/andp.18892730505`. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.18892730505` (visited on 01/04/2026).

[90] D. Freude. "Size, Mass and Kinetics of Molecules". In: *Molecular Physics*. Leipzig: Universität Leipzig, 2004.

[91] Yangyang Fu et al. "Electrical breakdown from macro to micro/nano scales: a tutorial and a review of the state of the art". In: *Plasma Research Express* 2.1 (Feb. 7, 2020), p. 013001. ISSN: 2516-1067. DOI: `10.1088/2516-1067/ab6c84`. URL: `https://iopscience.iop.org/article/10.1088/2516-1067/ab6c84` (visited on 01/04/2026).

[92] R. Locht et al. "The dissociative electroionization of ammonia and ammonia-d3. I. The NH^ and NH_2^ dissociation channels". In: 123 (1988), p. 443.

[93] Michael Hegetschweiler and Gregory T. Linteris. "Data Reduction Tool for Spherical Constant Volume Flame Experiments". In: *NIST* (Apr. 2, 2021). Last Modified: 2022-11-29T09:11-05:00. URL: `https://www.nist.gov/publications/data-reduction-tool-spherical-constant-volume-flame-experiments` (visited on 01/08/2026).

[94] Chaimae Bariki et al. "A well-defined methodology to extract laminar flame speeds at engine-relevant conditions". In: *Combustion and Flame* 268 (Oct. 1, 2024), p. 113612. ISSN: 0010-2180. DOI: `10.1016/j.combustflame.2024.113612`. URL: `https://www.sciencedirect.com/science/article/pii/S0010218024003213` (visited on 12/17/2025).

[95] Shaowen Xu et al. "Investigation and Optimization of Existing Laminar Flame Velocity Reaction Mechanisms in NH3 and NH3/H2 Flames". In: *Processes* 13.2 (Feb. 2025), p. 466. ISSN: 2227-9717. DOI: `10.3390/pr13020466`. URL: `https://www.mdpi.com/2227-9717/13/2/466` (visited on 12/21/2025).

[96] Seok-Hwan Lee, Wook Choi, and Choong-Sik Bae. "Quantification of Thermal Shock in a Piezoelectric Pressure Transducer". In: *Transactions of the Korean Society of Automotive Engineers* 13.5 (2005), pp. 96–103. ISSN: 1225-6382. URL: `https://www.koreascience.kr/article/JAKO200504840637990.page` (visited on 12/13/2025).

[97] Serdar Bilgili, Vitaly Bychkov, and V'yacheslav Akkerman. "Impacts of the Lewis and Markstein numbers on premixed flame acceleration in channels due to wall friction". In: *Physics of Fluids* 34.1 (Jan. 1, 2022), p. 013604. ISSN: 1070-6631, 1089-7666. DOI: `10.1063/5.0067222`. URL: `https://pubs.aip.org/pof/article/34/1/013604/2845966/Impacts-of-the-Lewis-and-Markstein-numbers-on` (visited on 12/08/2025).

[98] Miriam Reyes et al. "Experimental characterization of burning velocities of premixed methane-air and hydrogen-air mixtures in a constant volume combustion bomb at moderate pressure and temperature". In: *Applied Thermal Engineering* 130 (Feb. 2018), pp. 684–697. ISSN: 13594311. DOI: `10.1016/j.applthermaleng.2017.10.165`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S1359431117306907` (visited on 12/08/2025).

[99] Jerric R. Delfin et al. “Thermoacoustic parametric instability of premixed ammonia flames propagating downwards in an open-closed tube”. In: *Fuel* 373 (Oct. 1, 2024), p. 132344. ISSN: 0016-2361. DOI: `10.1016/j.fuel.2024.132344`. URL: `https://www.sciencedirect.com/science/article/pii/S0016236124014923` (visited on 01/08/2026).

[100] Robert DI SALVO, Stephen Doherty, and Alton J. REICH. “Constant volume combustion chamber”. Pat. WO2012177523A2. Llc Streamline Automation. Dec. 27, 2012. URL: `https://patents.google.com/patent/WO2012177523A2/en` (visited on 04/18/2025).

[101] Bernard Robic. “Constant-volume combustion (cvc) chamber for an aircraft turbine engine including an intake/exhaust valve having a spherical plug”. U.S. pat. 20150204240A1. SNECMA SAS. July 23, 2015. URL: `https://patents.google.com/patent/US20150204240A1/en` (visited on 04/18/2025).

[102] Nguyen Tuan Nghia et al. “A Comparison in Combustion Characteristics in a CVCC with Biodiesel Blends”. In: *Energies* 15.3 (Jan. 2022). Number: 3, p. 1017. ISSN: 1996-1073. DOI: `10.3390/en15031017`. URL: `https://www.mdpi.com/1996-1073/15/3/1017` (visited on 05/16/2025).

[103] Rudolf R. Maly and Rüdiger Herweg. “Spark Ignition and Combustion in Four-Stroke Gasoline Engines”. In: *Flow and Combustion in Reciprocating Engines*. Ed. by C. Arcoumanis and Take Kamimoto. Berlin, Heidelberg: Springer, 2009, pp. 1–66. ISBN: 978-3-540-68901-0. DOI: `10.1007/978-3-540-68901-0_1`. URL: `https://doi.org/10.1007/978-3-540-68901-0_1` (visited on 04/14/2025).

[104] Thomas H. Pratt. *Electrostatic Ignitions of Fires and Explosions*. Google-Books-ID: bmI69XjsvtMC. John Wiley & Sons, Sept. 7, 2010. 198 pp. ISBN: 978-0-470-93514-9.

[105] V. A. W. Hillier. *Fundamentals of motor vehicle technology*. 3rd ed. OCLC: 1110269467. 471 p: Ill, 1998. ISBN: 978-0-09-143161-7.

[106] Jacob Millman and Samuel Seely. *Electronics*. 2nd ed. New York: McGraw-Hill Book Company, 1941. 234-297.

[107] Stefan Essmann et al. “Ignition characteristics of hydrogen-enriched ammonia/air mixtures”. In: *Applications in Energy and Combustion Science* 17 (Mar. 1, 2024), p. 100254. ISSN: 2666-352X. DOI: `10.1016/j.jaecs.2024.100254`. URL: `https://www.sciencedirect.com/science/article/pii/S2666352X24000098` (visited on 05/20/2025).

[108] J. F. Henderson. “Magneto Ignition”. In: *Proceedings of the Institution of Automobile Engineers* 9.2 (June 1, 1915), pp. 443–480. ISSN: 0369-9838. DOI: `10.1243/PIAE_PROC_1914_009_027_02`. URL: `https://doi.org/10.1243/PIAE_PROC_1914_009_027_02` (visited on 05/11/2025).

[109] Normand M. Laurendeau. “Thermal ignition of methane-air mixtures by hot surfaces: A critical examination”. In: *Combustion and Flame* 46 (Jan. 1, 1982), pp. 29–49. ISSN: 0010-2180. DOI: `10.1016/0010-2180(82)90005-0`. URL: `https://www.sciencedirect.com/science/article/pii/0010218082900050` (visited on 05/12/2025).

[110] Shalom Eliezer and Yaffa Eliezer. *The fourth state of matter: an introduction to plasma science*. 2nd ed. Bristol Philadelphia: Institute of physics publ, 2001. ISBN: 978-0-7503-0740-6.

[111] Bruce Chehroudi. “Laser Ignition For Combustion Engines”. In: (2004). DOI: `10.13140/2.1.1845.1206`. URL: `http://rgdoi.net/10.13140/2.1.1845.1206` (visited on 05/12/2025).

[112] Hirokazu Kojima, Eiichi Takahashi, and Hirohide Furutani. “Breakdown plasma and vortex flow control for laser ignition using a combination of nano- and femto-second lasers”. In: *Optics Express* 22.101 (Jan. 13, 2014), A90–A98. ISSN: 1094-4087. DOI: `10.1364/OE.22.000A90`. URL: `https://opg.optica.org/oe/abstract.cfm?uri=oe-22-S1-A90` (visited on 05/12/2025).

[113] Azer P. Yalin. “High power fiber delivery for laser ignition applications”. In: *Optics Express* 21.106 (Nov. 4, 2013), A1102–A1112. ISSN: 1094-4087. DOI: `10.1364/OE.21.0A1102`. URL: `https://opg.optica.org/oe/abstract.cfm?uri=oe-21-S6-A1102` (visited on 05/12/2025).

[114] M. F. Ivanov, A. D. Kiverin, and M. A. Liberman. *Ignition of Deflagration and Detonation Ahead of the Flame due to Radiative Preheating of Suspended Micro Particles*. July 2, 2015. DOI: `10.48550/arXiv.1406.4058`. arXiv: `1406.4058[physics]`. URL: `http://arxiv.org/abs/1406.4058` (visited on 05/12/2025).

[115] H. G. Wolfhard and D. S. Burgess. “The ignition of combustible gases by flames”. In: *Combustion and Flame* 2.1 (Mar. 1, 1958), pp. 3–12. ISSN: 0010-2180. DOI: `10.1016/0010-2180(58)90016-6`. URL: `https://www.sciencedirect.com/science/article/pii/0010218058900166` (visited on 05/12/2025).

[116] Jiannan Shen et al. “Experimental Research on Microwave Ignition and Combustion Characteristics of ADN-Based Liquid Propellant”. In: *Micromachines* 13.4 (Mar. 25, 2022), p. 510. ISSN: 2072-666X. DOI: `10.3390/mi13040510`. URL: `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9027818/` (visited on 05/07/2025).

[117] David. A Frank-Kamenetskii. *Towards temperature distributions in a reaction vessel and the stationary theory of thermal explosion*. Vol. 18. Doklady Akademii Nauk SSSR, 1938.

[118] M. Kono, S. Kumagai, and T. Sakai. “The optimum condition for ignition of gases by composite sparks”. In: *Symposium (International) on Combustion* 16.1 (Jan. 1977), pp. 757–766. ISSN: 00820784. DOI: `10.1016/S0082-0784(77)80369-X`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S008207847780369X` (visited on 04/03/2025).

[119] Xiao Zhang and Lin Chen. “The Synergy Effect of Ignition Energy and Spark Plug Gap on Methane Lean Combustion with Addressing Initial Flame Formation and Cyclic Variation”. In: *ACS Omega* 8.7 (Feb. 21, 2023), pp. 7036–7044. DOI: `10.1021/acsomega.2c07897`. URL: `https://doi.org/10.1021/acsomega.2c07897` (visited on 04/03/2025).

[120] ABDELKADER FRENDI and MERWIN SIBULKIN. “Dependence of Minimum Ignition Energy on Ignition Parameters”. In: *Combustion Science and Technology* 73.1 (Sept. 1, 1990). _eprint: https://doi.org/10.1080/00102209008951659, pp. 395–413. ISSN: 0010-2202. DOI: `10.1080/00102209008951659`. URL: `https://doi.org/10.1080/00102209008951659` (visited on 03/27/2025).

[121] Joonsik Hwang et al. “Microwave-assisted plasma ignition in a constant volume combustion chamber”. In: *Combustion and Flame* 167 (May 1, 2016), pp. 86–96. ISSN: 0010-2180. DOI: `10.1016/j.combustflame.2016.02.023`. URL: `https://www.sciencedirect.com/science/article/pii/S0010218016000791` (visited on 05/07/2025).

[122] Liming He et al. “Research progress of microwave plasma ignition and assisted combustion”. In: *Chinese Journal of Aeronautics* 36.12 (Dec. 1, 2023), pp. 53–76. ISSN: 1000-9361. DOI: `10.1016/j.cja.2023.04.029`. URL: `https://www.sciencedirect.com/science/article/pii/S1000936123001498` (visited on 05/07/2025).

[123] INTERNATIONAL ELECTROTECHNICAL COMMISSION. *Explosive atmospheres. Part 0, Equipment–general requirements*. Ed. 7.0. OCLC: 1038502822. Geneva, Switzerland: International Electrotechnical Commission, 2017. ISBN: 978-2-8322-5065-5.

[124] Ciprian Dumitrache, Carter Butte, and Azer Yalin. "Resonant dual-pulse laser ignition technique based on oxygen REMPI pre-ionization". In: *Scientific Reports* 10.1 (Nov. 16, 2020), p. 19916. ISSN: 2045-2322. DOI: `10.1038/s41598-020-76968-5`. URL: `https://www.nature.com/articles/s41598-020-76968-5` (visited on 05/06/2025).

[125] Xu He et al. "Super adiabatic flame temperature phenomenon for NH3/O2/N2 mixtures". In: *Fuel* 346 (Aug. 15, 2023), p. 128264. ISSN: 0016-2361. DOI: `10.1016/j.fuel.2023.128264`. URL: `https://www.sciencedirect.com/science/article/pii/S0016236123008773` (visited on 01/19/2026).

A

Additional background information

A.1. Constant-volume combustion chambers (CVCC)

A constant-volume combustion chamber is a device in which fuel and oxidiser are injected and then sealed, allowing combustion to occur at nearly constant volume and greatly increasing pressure and temperature. Only after combustion is complete is the chamber opened, releasing the high-pressure products, which, through expansion, produce work, usable either as direct thrust or as drive for a piston. It approximates the idealised thermodynamic Otto cycle for constant-volume combustion. This is different from constant-pressure combustion, found in gas turbines and some rockets, where heat is produced as gases continually enter and exit the chamber, keeping the pressure almost constant, with some losses. CVCCs are used in:

- Internal combustion engines (ICEs), where it is the zone between the piston and the end of the piston chamber.
- Laboratory devices to study ignition, used to study the combustion and emission characteristics of fuels, as in the case of this research as well.
- Pulse detonation engines, which mimic constant volume combustion by rapidly burning the fuel-air mixture before exhaust. However, true CVCCs use mechanical valves or pintles to seal the chamber during combustion.
- Rocket propulsion, where some advanced rockets use CVCC to improve the I_{sp} at given supply pressures [100].
- Aircraft engines, as there is active research and some patented designs for implementing CVCCs in aircraft engines to replace traditional Brayton cycle combustors with Humphrey cycle (constant volume with expansion to atmospheric pressure) combustors for improved efficiency [101].

This type of system has several benefits and drawbacks over its constant-pressure counterpart, which are described here:

- Advantages
    - High thermal efficiency: Otto cycle is more efficient than the Diesel and Brayton cycle, because the combustion process transforms internal energy into pure heat, while in the other cycles some goes into work and some into heat

  - Good for combustion process analysis: it allows for precise control of combustion parameters like the equivalence ratio, initial pressure, and temperature, which allows for flexibility and easier combustion analysis of the flame propagation
  - Simplicity of design and operation: they have no moving parts and they do not require any flow control or continuous operation since they consist of an enclosed environment [102]
- Disadvantages
  - Potential for detonation and combustion instabilities: the pressure inside the chamber rises steeply, which may lead to detonation or pressure oscillations.
  - High heat rejection to the chamber walls: since it is a fully closed chamber, more heat is dissipated to the surrounding walls.
  - Complex ignition systems: the combustion is not continuous, so it needs rapid and repeated ignition.

For these reasons, constant-volume combustion chambers (CVCC) are frequently used in fundamental combustion research because they enable well-controlled initial pressure, temperature, and mixture composition, while providing a closed environment where the global pressure response can be measured with high repeatability. Under quiescent, premixed conditions, the chamber approximates a constant-volume system during the early and intermediate phases of the event, and the pressure evolution can be related to the integrated heat release of the flame. This makes CVCC experiments particularly suitable for comparing fuels and operating conditions, since the initial state can be precisely set and the same ignition hardware can be used across cases. For the objectives of this thesis, the CVCC also provides a practical way to explore conditions rarely covered in the existing literature, such as elevated initial pressure combined with subzero initial temperature, while maintaining a consistent measurement approach. The next subsection, therefore, focuses on the instrumentation that enables these chamber conditions and responses to be measured reliably.

A.2. Ignition modes

When a premixed mixture is introduced into the combustion chamber, some conditions must be met for it to actually combust. In fact, as explained in Section 2.1, three conditions are necessary for this process to start. Namely, these are an oxidiser, a fuel, and an ignition source that provides the initial heat required to break the bonds and initiate the chain reaction. Many different possibilities exist and are found in the literature, among which the most common are described here.

A.2.1. Spark ignition

Spark ignition is the most widely used system for igniting a mixture. It consists of an electrical spark to release the necessary energy. However, this spark can also occur unintentionally due to lightning strikes and electrostatic discharge [104]. Thus, it is important to study this system and the mixture's MIE to avoid accidental ignition. Many different systems exist to give this electric spark, which will be discussed here, but they all work with the general circuit shown in Figure A.1. This shows the basic principle

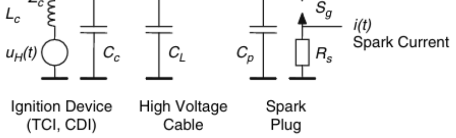

Figure A.1: Electrical circuit for spark ignition [103] of this ignition mode: creating a high voltage between two electrodes and, therefore, an electric field. The circuit is connected to a power supply, which can be a battery for transportation systems or a land supply for static systems, providing the voltage. This voltage passes through the circuit and is delivered to the mixture via a spark plug, a device that protrudes into the combustion chamber and ends in two electrodes, between which the voltage energises the mixture. Once this is high enough, normally between 12000 and 45 000 V, it is able to break the dielectric barrier separating them and allowing for the ionisation of the mixture. In fact, air acts as an insulator, and without a force strong enough to break its barrier, electrons are not allowed to pass through. When this is broken, a plasma flow is created, with highly energetic electrons that break molecular bonds, ionise the mixture, and initiate reactions [105]. This type of ignition can be divided into five main phases:

1. Pre-breakdown phase: energy, and thus voltage, is accumulated at the electrodes; this is, however, still not enough to break the dielectric barrier.
2. Breakdown phase: a very short phase, in the order of nanoseconds, in which the resistance of air is broken and the current rises quickly, as the flow of electrons is now allowed to move between electrodes. It generates a sudden burst of energy, creating a plasma stream of electrons [103]. The voltage at which the electric field is sufficient to allow flow is called the breakdown voltage. It is not a constant for every mixture but varies with operational conditions. This voltage is related to the mean free path of the mixture molecules, that is, the average distance between collisions with other molecules. This quantity is inversely proportional to the pressure at constant temperature; when the pressure is higher, the molecules are closer together, allowing more collisions. For electrons, the mean free path is larger than the molecular mean free path due to their smaller size.

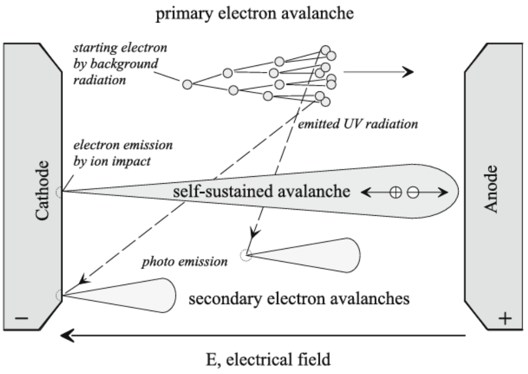

Figure A.2: Electron cascade effect [103]

When a voltage is applied between the two electrodes, some electrons enter the electric field, where they will be accelerated and energised. If this energy is high enough, then they can ionise the present mixture molecules, releasing more electrons in a cascade effect, as seen in Figure A.2.

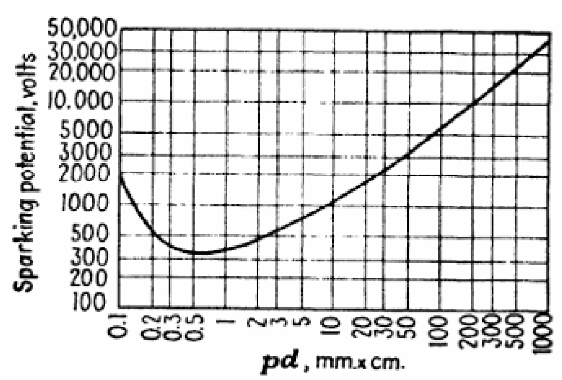

Figure A.3: Paschen curve [106]

As mentioned, if the pressure rises, more collisions will occur due to the lower mean free path. However, when such an event occurs, the electrons can be scattered in any direction, which may be different from the one of the electric field. This means that these electrons are not energised by it. Moreover, they tend to lose energy with every collision, as kinetic energy is used to break the molecular bonds. Thus, pressure is an important factor relating the cascade effect to the required external voltage. However, another parameter also needs to be considered: the gap distance. This affects both the length of the path of the electrons and the strength of the electric field, which is related to the voltage according to

$$E = \frac{V}{d} \tag{A.1}$$

Therefore, the breakdown voltage can be related to the product of the pressure and the gap distance. This relation is shown in the so-called Paschen curve, shown in Figure A.3. This indicates that an optimal value for this product exists that minimises the breakdown voltage. Indeed, if $p \cdot d$ is too high, the electrons can collide with more particles due to the higher concentration and higher length of the path, losing their energy and changing their direction. On the other hand, if this is too small, the mean free path is longer than the gap, requiring higher energy to sustain the cascade. However, this also means that, if the electrodes are not designed correctly with respect to the initial design pressure in the chamber, the breakdown voltage may occur at a position not coincident with the smallest distance between electrodes [106].

3. Arc phase: once the breakdown voltage is achieved and electrons are allowed to flow freely between the two electrodes, the area in between becomes conductive plasma. This is highly conductive, so the voltage drops; there is no need for a large voltage, and the current can flow easily. This phase runs for about 1 µs, during which the current stabilises around a constant value. In this phase, a few millijoules of energy are transferred to the surrounding gas, thereby increasing its temperature. This release of energy is visible with the arc created between the electrodes.

4. Glow phase: This phase lasts several milliseconds, during which most energy is released. As this is released, the current decreases, and the voltage increases again to maintain a constant flow of energy against dissipation to the environment. The spark raises the mixture temperature to 6000 K, above the auto-ignition temperature, initiating combustion via thermal dissociation.

5. Flame kernel formation: If the energy the spark delivers is sufficient, then a flame kernel is created, in which the energy produced within its reactions contributes to its propagation.

Within the category of spark ignition, many subcategories can be identified depending on how energy is delivered to the electrodes. The main ones will be briefly discussed here.

- Transistor Coil Ignition (TCI): This is the most common type of spark-ignition system, operating on the principle of magnetic induction. Spark ignition requires a high voltage to break the dielectric barrier of the mixture. In this system, however, the voltage is stepped up rather than delivered directly from a battery or magneto, which would require a large power supply. The coil acts as a transformer, with two separate coil windings: one with many turns and the other with fewer. The latter is connected to the power supply, thus creating a low-voltage circuit. Once this reaches its maximum allowed energy, the circuit opens, allowing energy to transfer to the secondary coil. Since this has more turns, the voltage increases while the total energy remains constant, thereby reducing the current. Using a transistor, an electronic switch, allows for better control and timing of when circuits need to be opened and closed, reducing energy losses. Since it has no moving parts, it can switch much faster. This also ensures that the magnetic field collapses much more quickly, increasing the current and, hence, the voltage. This ignition system is very efficient, as the coils release energy slowly over about 2 ms, ensuring a strong, consistent spark. This can help guarantee a more reliable ignition, but may be slower under cold conditions, where the system is also more sensitive to coil saturation [105].
- Capacitive Discharge Ignition (CDI): This type of ignition is very similar to the one explained above, with the difference that, in this case, the energy is stored in a capacitor that, when the circuit to the spark plug closes, releases it quickly, generating a strong, short spark. This could be useful in cold conditions, especially for the case of ammonia, as a fast and strong discharge is better for igniting cold mixtures, as the kernel can be quenched easily, and a fast discharge reduces losses [107].
- Magneto Ignition: This system also features a coil, but the difference from the above is how the voltage is supplied to it. Indeed, in this case, a battery is not necessary, as the voltage is provided by a magnet. This spins around the primary coil, generating a strong magnetic force. Faraday's Law states that a changing magnetic field induces a current in a wire, generating a voltage difference. When the contact breakers are opened, the secondary coil steps up this voltage. This ignition mode depends on the magnet's rotation rate; at low speeds, the magnetic force is weaker, and consequently, the spark is weaker. It generally produces a lower spark energy than CDI, and the cold cranking, that is, the ability to deliver sufficient current in cold conditions, is very weak, as it takes more energy to rotate the magnet [108].
- Battery Ignition: This system is very similar to the TCI, with the only difference that, instead of a transistor, mechanical breakers physically open and close the connection. This is much slower, resulting in a lower-intensity spark. Moreover, breakers are subject to much higher wear due to the continuous movement. The low spark energy provided, along with the voltage drop and the slow current rise in cold conditions, makes it unsuitable for ignition in these cases [105].

A.2.2. Plasma ignition

Although plasma streams are already produced in spark ignition, once the breakdown voltage is reached and electrons are allowed to move freely between electrodes, they can also be used directly to ignite a mixture. A distinction has to be made between equilibrium and non-equilibrium plasma. The streams occurring in spark plugs are made of so-called equilibrium plasma, that is, where the temperature of the introduced electrons and that of the mixture is the same. In fact, in this case, the thermal energy is shared among all particles, and therefore, the electrical energy input between the two electrodes heats the mixture. On the other hand, in a non-equilibrium plasma, the electrons have a considerably higher temperature than the surrounding mixture. This is because the electrical energy input into the stream increases its kinetic energy, and thus its thermal energy. However, due to their reduced size, they cannot dissipate energy by conduction. Instead, they excite the surrounding molecules by impact, hence breaking the bonds and creating highly reactive radicals, such as O and OH. In fact, combustion can be initiated by thermal heating as well as by chemically active pathways. Thus, these metastable species can activate additional species and initiate chain reactions, even at temperatures below the auto-ignition temperature. The most relevant parameter defining which reactions are dominant in the mixture is the reduced electric field E/n, which indicates the ratio of the electric field strength to the number density of the mixture gas molecules. This latter parameter indicates the number of particles per cubic metre. E/n expresses the average energy gain of electrons per collision. In fact, electrons lose energy after every collision, which is used to break molecular bonds. However, the electric field energises them; thus, the parameter indicates their ability to remain energetic depending on the number of collisions. When this number is high, electrons can dissociate many molecules, even after collisions, as they still have energy to collide with additional particles and break their bonds. Conversely, when this is low, the cross section, that is, the probability of a process during collisions, for vibrational excitation, is much higher. This means that electrons will increase the vibrational energy of the surrounding molecules, especially N_2, which can store this energy and help provoke the dissociation of other molecules, as the potential energy is much higher. This ignition mode has a much shorter ignition delay and works well with cold combustion and lean mixtures, as the mixture does not need to reach the auto-ignition temperature.

A.2.3. Hot surface ignition

This ignition process works with the same principle as spark ignition, meaning that ignition occurs when the temperature reaches the self-ignition one. In this case, an external resistive element, normally made of silicon carbide or nitride, is introduced in the combustion chamber. This is heated via an electrical current, which increases its temperature via resistive dissipation, according to

$$P = I^2R \tag{A.2}$$

The surface temperature can reach 1500 K or higher, thereby increasing the temperature of the adjacent mixture, primarily through conduction and convection. This means that the ignition is purely thermal and does not involve the creation of radical species. In this case, the ignition delay and the energy are related to the surface temperature, as this entirely controls when combustion will occur [109]. Naturally, in cold conditions, especially for a high-MIE fuel such as ammonia, more energy must be supplied to the surface to reach the autoignition temperature, making it less useful in these conditions.

A.2.4. Laser ignition

LASER is a word that stands for Light Amplification by Stimulated Emission of Radiation. As the name suggests, its function is to force the emission of photons. In fact, in a molecule, electrons can either be in their ground state, which is stable as it carries the lowest potential energy, or in an excited state, in which the electrons absorb sufficient external energy to rise to a higher-energy level. However, in the second configuration, they tend to go back to the lower state by releasing this energy. This can occur spontaneously or in response to stimulation. In the former case, the particle simply releases the energy it acquired by emitting a photon in a random direction. In the latter case, a photon interacts with the excited electron, causing it to emit a photon with the same frequency, direction, and phase as the incident photon.

The difference between these two cases can be seen in Figure A.4. For a laser to work, stimulated emission must be predominant over spontaneous one, as the wanted result is a focused beam in which all photons move in the same direction. In fact, once coherent photons are emitted, they can excite additional electrons, leading to a cascade that produces more photons. Once the energy level is high enough, they can leave the partially reflective mirror at one end of the laser chamber, resulting in a highly focused, energetic beam. This is able to initiate combustion in three main ways, namely:

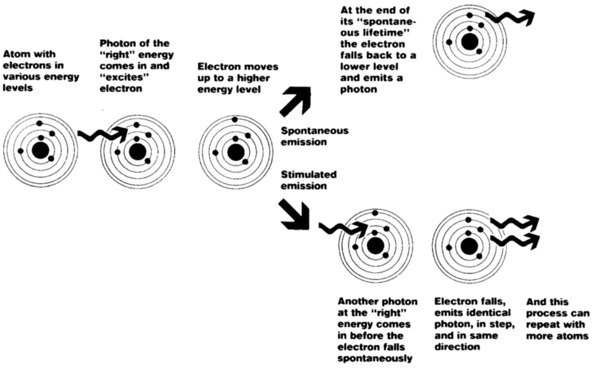

Figure A.4: Electron emission [110]

- Laser-Induced Breakdown (LIB): The laser beam is introduced into the combustion chamber, where the absorption of photons by the molecules leads to their ionisation. The electrons that are liberated gain energy through collisions with photons, thus allowing them to ionise molecules as well. This produces a hot plasma that ignites the mixture, similarly to spark ignition.
- Thermal ignition: The beam directly heats a specific region, with the molecules absorbing the photon energy and converting it into heat mostly via molecular vibration. In this case, plasma is not formed, and ignition occurs as soon as the mixture achieves its self-ignition temperature.
- Photochemical ignition: the specific wavelengths of the laser beam break the molecular bonds, creating highly-reactive radicals, which initiate the chain reactions leading to ignition. In this case, no plasma forms because the provided wavelength carries insufficient energy to induce ionisation. Thus, in this case, only dissociation occurs. The difference between the two processes is that the first entails removing electrons from a molecule, creating positive ions and electrons, while the latter forms plasma. On the other hand, dissociation entails breaking bonds, creating smaller neutral radicals that are chemically active but not charged [111].

The advantage of using a laser is the elimination of electrodes from the combustion chamber, which reduces heat loss and thus increases the effectiveness of this method at high pressures, as stronger pressure waves enhance this phenomenon [112]. One of the main problems is the need for optical components to deliver the laser beam, which would normally require windows and thus an optical chamber. However, recent advancements have demonstrated the feasibility of delivering it via fibre optics, which is useful for the scope of this research [113]. The first and third ignition modes may work for cold ignition, but would require a large amount of laser energy.

A.2.5. Flame-induced ignition

This mode, as the name suggests, requires ignition to occur as a result of interaction with another flame. For this to occur, the pre-existing flame needs to be ignited in a separate channel, often an external pre-chamber, which increases the design complexity, as not only does this additional chamber need to be created, but two separate mixtures also need to be created. In this case, ignition can occur mainly through thermal conduction, driven by the Zeldovich gradient mechanism, which explains how temperature gradients ahead of a flame can lead to spontaneous ignition. This is because the self-ignition temperature can be achieved, igniting the mixture [114]. Further, this temperature can also be achieved through pressure waves generated by the pre-chamber flame, which compress the flow and raise its temperature. Finally, ignition can also be caused by radical diffusion from the flame into the unburnt mixture, lowering its ignition energy and temperature. Normally, ignition occurs as a result of the sum of the three specified mechanisms [115]. This mechanism can be useful at low temperatures, especially if the initial flame is doped with a high concentration of hydrogen, which is easier to ignite. To ignite it, one of the other systems described here must be used, but creating a smaller flame with a higher hydrogen content is easier.

A.2.6. Microwave ignition

In this mode, microwaves of frequencies normally around 2.45 GHz are introduced into the combustion chamber by means of antennas or waveguides. These waves are substantial oscillating electric fields that interact with the electrons of the mixture molecules, imparting oscillatory motion and transferring kinetic energy, allowing them to collide with other molecules and creating plasma. However, since the combustion chamber acts as a resonant cavity, waves can rebound and interact, leading to different intensities at different points in the chamber. This means that multiple-point ignition can occur with microwave introduction, leading to higher pressure and temperature peaks and faster combustion, as the flame does not need to expand as much. Moreover, since it does not rely on thermal ignition mechanisms, it is suited for cold combustion [116].

A.2.7. Autoignition

This ignition mode is the only one that does not require an external energy source. In fact, this occurs when the mixture reaches the critical autoignition temperature, at which the rate of heat generation from exothermic reactions exceeds heat losses, leading to ignition. This normally occurs when the pressure is high enough to raise the temperature above this limit. To understand when this occurs, Frank-Kamenetskii developed a criterion to establish whether a system will undergo thermal self-ignition, based on the balance between heat generation and heat loss. This introduces a dimensionless parameter, which varies depending on the combustion chamber shape and is dependent on the Damköhler number, which is the ratio between the flow and the chemical timescale:

$$Da = \frac{t_f}{t_{chem}} \tag{A.3}$$

This means that if this is much greater than 1, the reaction occurs much faster than the flow can carry it away, so there is not enough time for conduction to remove the heat, and ignition can occur. On the other hand, if the parameter is low, heat loss dominates, and the system enters a steady state, requiring an external energy source to ignite. Thus, the Frank-Kamenetskii number will be the critical Damköhler number at which autoignition begins. This needs to be taken into account in the design of the combustion chamber, as ignition may not occur as expected and may instead start before the ignition source is applied [117].

A.3. MIE dependency on ignition mode

The MIE is also dependent on the ignition method, as different ignition modes operate through different physical principles and therefore exhibit different sensitivities to system-specific variables. In this work, the most relevant ignition approach is spark ignition, as it is widely used and allows for a relatively direct determination of MIE. However, it is important to note that other ignition modes exist and may yield different apparent ignition thresholds, either because they deposit energy in different spatial or temporal patterns or because they promote ignition via non-thermal pathways. For completeness, the main ignition modes are briefly summarised hereafter, with emphasis on the parameters that most strongly influence the ignition threshold.

- Spark ignition:

  This is the most common mode, and its MIE depends primarily on the spark-gap size, the discharge duration, and the electrode configuration. An optimum spark gap generally exists, as very small gaps produce small kernels that are more susceptible to heat losses and electrode quenching, while very large gaps distribute the energy over a larger volume and require more energy to sustain kernel growth [118, 119, 43]. Similarly, an optimum discharge duration is typically observed, since too short a discharge may not allow sufficient radical formation, while overly long discharges increase conductive losses and may induce strong pressure disturbances [43, 120]. Finally, electrode geometry affects heat losses and field distribution, and smaller electrodes tend to minimise quenching provided the arc remains stable [118].

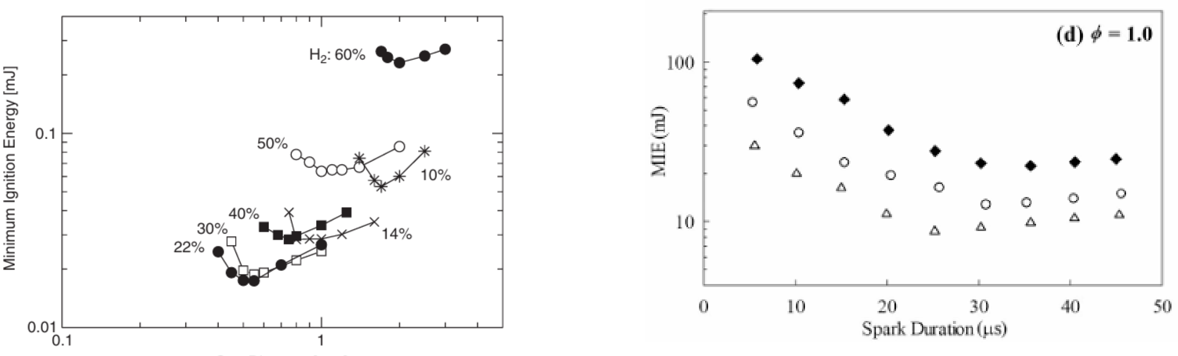

(a) Spark gap with increasing percentage of H_2 in air [51]

(b) Spark duration (squares are pure NH_3, circles are 5% addition of H_2 and triangles are 10% addition) [43]

Figure A.5: Relation of the MIE with spark ignition parameters

- Plasma and microwave-assisted ignition:

  These approaches can lower the effective ignition threshold by enhancing radical production and energising the mixture through non-equilibrium effects. Their performance generally depends on discharge conditions (e.g., temperature and density for plasmas, or frequency and pulse duration for microwaves), as well as on electrode or cavity configuration, which controls field intensity and energy coupling. [121, 122]

- Hot-surface ignition and autoignition

  In these cases, ignition is more naturally characterised by a temperature threshold rather than a delivered energy, since combustion begins once the local temperature is high enough for reactions to proceed. For hot-surface ignition, the minimum hot-surface ignition temperature depends mainly on surface material properties and the effective heat-transfer area. [123] For autoignition, the outcome is governed by the initial thermodynamic state and mixture composition, which determine whether the mixture can reach ignition-relevant temperatures and reaction rates without external excitation.

- Laser ignition and flame-induced ignition

  Laser ignition deposits energy locally and can be tuned by pulse duration and focusing, but it often requires higher nominal energy than spark ignition due to optical losses and the highly localised nature of the kernel. [124] Flame-induced ignition depends on the properties of an adjacent flame and the resulting thermal and radical fluxes that can initiate combustion in the unburnt mixture.

A.4. LBV dependency on adiabatic temperature and Lewis number

The adiabatic flame temperature is directly dependent on the heat released during combustion. In fact, a reaction liberating more heat will logically result in a higher flame temperature, meaning that fuels with a larger heat of combustion will propagate faster. The relation between the speed and the temperature is given by the Arrhenius kinetics mechanism:

$$k = Ae^{-\frac{E_a}{RT}} \tag{A.4}$$

where

k	Reaction rate constant
A	Pre-exponential factor representing frequency of useful collisions
E_a	Activation energy
R	Universal gas constant
T	Temperature

This equation essentially states that increasing the adiabatic temperature also increases the reaction rate and, hence, the flame speed. This direct relationship is evident in Figure A.6a, with both speed and temperature curves peaking between 1.05 and 1.1. However, the one for hydrogen is quite different, with the flame speed peaking at an equivalence ratio of 1.75 while the temperature at 1.07. The difference lies in the extremely high diffusivity of hydrogen and the scaling of velocity with the square of the Lewis number:

$$S_u^0 \sim \sqrt{Le} \tag{A.5}$$

where

S_u^0	Unstretched burning velocity
Le	Lewis number

The shift towards richer mixtures is due to the fact that, with a very small Lewis number, mass diffusion is too high for reactions to occur; thus, the burning velocity will be lower. The peak will be achieved when Le is closer to 1, which occurs in rich hydrogen conditions. On the other hand, the Le of methane is almost 1 over the equivalence ratios of 1.05-1.1. A similar behaviour to that of methane is shown by ammonia, since NH_3 shows a Lewis number varying between 0.94 for lean and 1.09 for rich mixtures, as shown in Figure A.6c, which shows a curve for the adiabatic flame temperature very similar to the methane one [8].

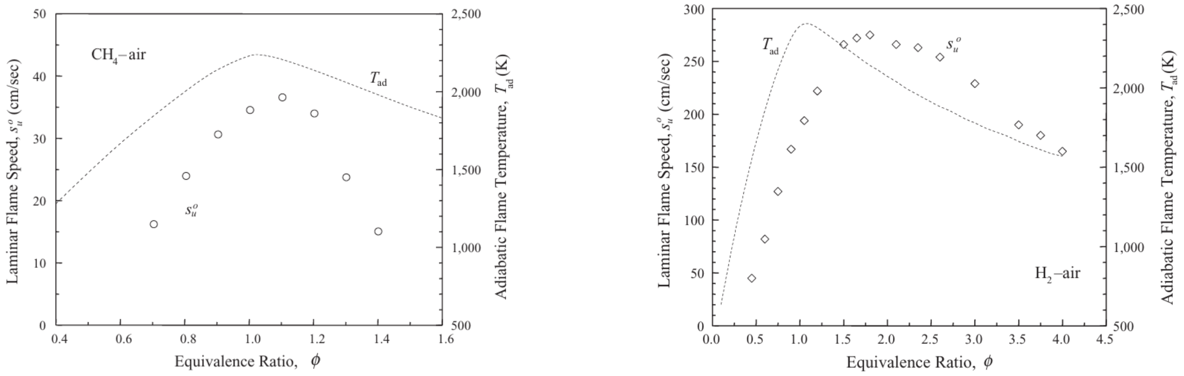

(a) Methane [15] **(b)** Hydrogen [15]

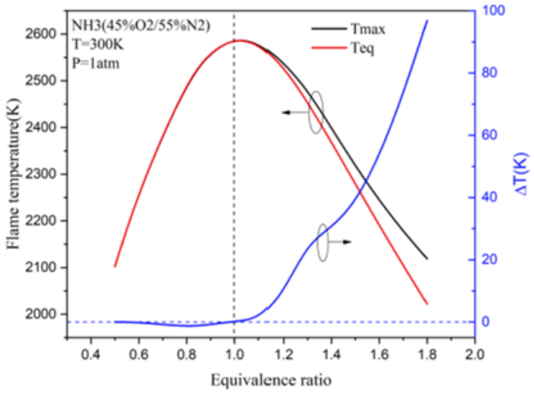

(c) Ammonia [125]

Figure A.6: LBV variation with adiabatic flame temperature

B

Additional design details

B.1. Material choice

Hereafter, a more detailed explanation of the scores given in Table 4.3:

- 316L SS: Austenitic stainless steel with excellent low-temperature toughness and stable ductile behaviour, which supports the subzero-capability score. Its main drawback is the comparatively low yield strength, which drives thicker walls for a given pressure rating and therefore penalises strength/stiffness. It is widely used for ammonia service in industrial handling, which supports a good NH_3-compatibility score, and it remains attractive from a manufacturability and availability standpoint due to standard stock forms and well-understood machining.
- SS2205 (duplex): The high yield strength motivates a high strength/stiffness score and typically correlates with good fatigue strength for pressure-cycling applications. The main limitation is microstructural sensitivity to sustained exposure at elevated-temperature precipitation ranges, which can degrade toughness and corrosion performance, thereby lowering the high-temperature-capability score. In addition, welding and heat input require control to avoid phase imbalance, which increases manufacturing risk and constrains repair/iteration even if the raw machinability is acceptable.
- SS2507 (super duplex): Similar to 2205 but with higher strength and generally improved pitting resistance, supporting high scores in strength and NH_3-environment resistance. However, the same precipitation/embrittlement concerns apply, and the processing window is typically narrower, so the high-temperature score remains limited, and manufacturability is penalised relative to standard austenitic grades. The tighter welding control requirement also increases schedule risk in prototyping contexts.
- Inconel 625: Nickel-based alloy with broad service capability from cryogenic to high temperature, supporting high scores in subzero capability and high-temperature capability. It also provides strong general corrosion resistance in aggressive environments, which supports the NH_3-compatibility score and makes it robust in the presence of condensates. Its drawbacks are cost and procurement lead time, and it can be more demanding to machine than stainless steels, which lowers manufacturability despite excellent overall performance.
- Inconel 718: High-strength precipitation-hardened nickel alloy with excellent mechanical performance, particularly in yield and fatigue, which supports strong scores in strength-/stiffness and fatigue resistance. Its usable temperature range is narrower than 625 at the

upper end, and its processing route (solution/age heat treatments) is more prescriptive, increasing manufacturing complexity and schedule sensitivity. Machinability is also more challenging than that of austenitic stainless steels, so manufacturability is scored lower, even though the material capability is high.

- Hastelloy C-276: Typically selected when corrosion dominates the design driver, which supports the very high NH_3-compatibility score and good long-term chemical robustness. Mechanical properties are strong but not necessarily superior to the highest-strength nickel alloys, so strength/stiffness is not maximised. Cost and availability are commonly limiting factors for small-batch procurement, and manufacturability can be penalised by machining effort and tool wear.
- Monel 400: Nickel–copper alloy with good general corrosion resistance and good low-temperature toughness, which supports a strong subzero score and reasonable hydrogen compatibility. However, ammonia compatibility can be condition-dependent and sensitive to impurities or specific exposure modes, so the NH_3 score is conservatively reduced. Its strength is moderate compared with duplex steels and precipitation-hardened nickel alloys, which lowers the strength/stiffness score, while manufacturability and availability are typically acceptable but not as favourable as common stainless grades.

B.2. Chamber support

Once the chamber's mass and overall dimensions are fixed, the support legs must be sized to ensure the chamber stands stably and robustly on any surface during handling and operation. Therefore, the requirements for this are that it allow easy mounting and dismounting of the chamber, withstand the full weight of the chamber while resisting accidental pushes or disturbances, and be clamped to the laboratory optical table or any other supporting surface. As mentioned, the complete chamber (cylinder, flanges, and bolts) has a mass of approximately $m = 107$ kg, giving a weight of $W = mg \simeq 1.05$ kN. Initially, an intuitive approach would be to size the legs based solely on the vertical load; however, this is rarely the driving factor for short, metallic supports. The vertical load is distributed between the two legs, so each leg carries roughly $W/2 \simeq 0.525$ kN in compression, which is small compared with the compressive capacity and buckling resistance of typical steel plates used in laboratory fixtures. For this reason, the design is instead governed by lateral effects, because horizontal disturbances generate bending moments and can lead to tipping, both of which are more constraining than pure compression.

The key difference is that vertical loads primarily produce axial compression, while horizontal disturbances produce bending and overturning. In axial compression, the stresses scale with W/A (where A is the cross-sectional area), so even modest plate widths and thicknesses yield low stress levels. Conversely, a horizontal load applied at the centre of gravity (CG) produces a moment that scales with the lever arm, i.e. the CG height above the ground. Even if the horizontal force is significantly smaller than the weight, the moment $M \sim Fh_{\mathrm{CG}}$ can become the limiting factor, because h_{CG} is typically on the order of a few decimetres. Practically, horizontal disturbances are also the loads most likely to occur in a laboratory environment, for instance, due to accidental pushes, contact with tooling, or minor impacts during installation and maintenance. Therefore, the leg sizing was approached by separating the problem into two requirements: stability (to avoid tipping) and stiffness/strength (to limit bending and provide a robust attachment).

A conservative stability check can be performed using a rigid-body tipping model. A lateral disturbance is represented through an equivalent horizontal acceleration a_{lat} acting at the CG, generating an overturning moment about the outer edge of the footprint. Denoting with h_{CG} the vertical height of the CG above the ground and with B the total footprint width (distance between the outermost ground contact edges), the restoring moment is approximately $W(B/2)$. Imposing a safety factor N against tipping yields

$$W\frac{B}{2} \geq N\,(m a_{\mathrm{lat}})\,h_{\mathrm{CG}}, \qquad \Rightarrow \qquad B \geq 2N\left(\frac{a_{\mathrm{lat}}}{g}\right)h_{\mathrm{CG}}. \tag{B.1}$$

For equipment operated around people, a practical design envelope is $a_{\mathrm{lat}} = 0.1g$ to $0.3g$, where the upper bound captures handling-related disturbances and accidental pushes. By taking $N = 1.5$ and a representative range $h_{\mathrm{CG}} = 0.25$–0.35 m, the resulting requirement is $B \gtrsim 0.23$–0.32 m. This shows that the stand should have a footprint width of 250–320 mm to provide a clear stability margin.

After achieving stability, the leg cross-section was sized to prevent excessive bending and ensure adequate stiffness. Under a lateral acceleration, the total horizontal force is $F = m a_{\mathrm{lat}}$ and, assuming symmetric load sharing, each leg carries approximately $F/2$. If the leg behaves as a cantilever of height H (from the floor to its attachment on the cylindrical part), the bending moment at the attachment is $M \approx (F/2)H$. For a representative case $a_{\mathrm{lat}} = 0.3g$ and $H \approx 0.20$ m, this yields $F \approx 338$ N and $M \approx 34$ N m per leg. This moment is not large in terms of yielding, which confirms that the design is not primarily strength-limited; nevertheless, it is relevant for stiffness, as thin plates can deflect noticeably, making the assembly feel unstable. Additionally, a thicker leg enhances the quality and robustness of the attachment (whether welded or bolted), increases the available weld length, and reduces sensitivity to local stress concentrations.

Finally, because the legs are placed directly on the cylindrical part, the interface must avoid local point loading of the pressure boundary. Even if the global loads are modest, a small contact area can introduce local indentation or high bending stresses in the cylinder wall. For this reason, the leg design should include either a saddle surface that matches the cylinder's outer diameter or a reinforcement pad between the cylinder and the leg, ensuring the load is distributed over a larger area. This requirement also favours a thicker leg section, as it provides a stiffer load path and reduces local compliance at the contact point. Based on these considerations, a leg thickness of $t = 16$ mm was selected as a conservative solution: it provides high stiffness under lateral disturbances, offers a robust attachment region, and supports the use of a saddle or pad without introducing manufacturing complexity. Therefore, the final choice of 16 mm is justified not because the vertical load demands it, but because it ensures stable and repeatable support behaviour under the more critical horizontal disturbance and interface-loading conditions.

Regarding the geometry, inspiration was drawn from the existing supports for the optical chamber. This design features two small holes at the bottom to facilitate bolting the chamber to the optical table, ensuring it remains securely in place. The above discussion on the minimum leg thickness is particularly important when this is not possible, as the chamber might tilt due to vibrations or accidental pushes. Further changes to the optical supports concerned the hole position and size, as the bolts need to pass through the supports and be tightened to the nuts on the backside, and the radius of the curved part supporting the chamber. The final support design is illustrated in Figure B.1.

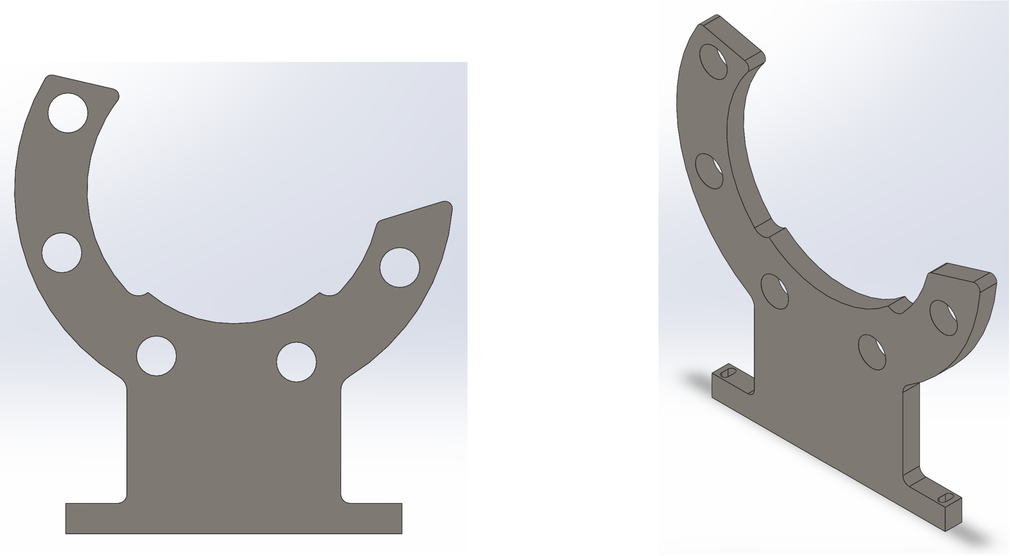

(a) Front view **(b)** Isometric view

Figure B.1: Chamber support leg

B.3. CDI ignition system

The ignition system shown in Figure 4.5 is explained here. It is important to highlight that some components are included for the simulation but are not real components that would be present in a real circuit. Starting from the high-voltage power supply on the left, it charges the array of capacitors via the high-resistance resistor R_p. This ensures both a slow RC charge, beneficial for the capacitors, and that this discharge does not cause ignition at the electrodes. In fact, an ignition caused directly by the high-voltage supply would be quite slow and might not lead to the development of a self-sustaining kernel, as the energy is not concentrated enough.

The system capacitance is distributed across an array of parallel capacitors, allowing it to be dynamically changed by simply activating or deactivating some. Once the capacitor reaches its maximum voltage, the NI9401 module generates an electrical pulse that activates both switches: the first closes the current path from the supply to the capacitor, and the second opens the one from the capacitor to the electrodes. Energy is then released from the capacitor through the small resistor to avoid surges in the coil (and to ensure RC discharge). At the same time, when this occurs, the other side of the circuit path is opened, which features a DC-AC converter. This allows the arc to be sustained and the spark duration to be controlled. In fact, the discharge from the capacitor might be enough to cause dielectric breakdown and initiate the first plasma stream, but for the actual arc, a longer discharge duration might be needed, which is provided by the converter.

B.4. Component specifications

In this section, all the important specifications of the sensor and components used in the setup are provided as a reference.

B.4.1. Ignition coil

Table B.1: Main specifications of the IGN-1A smart coil.

Specification	**Value**
Coil type	Inductive smart coil with integrated igniter (IGBT)
Output voltage (max.)	Up to 81 kV ± 10%
Output energy (max.)	Up to 250 mJ ± 7% (8 ms dwell @ 14 V)
Peak secondary current	100 mA ± 7%
Arc duration	3.2 ms ± 10% (3.0 ms dwell target)
Max continuous dwell	9 ms (not to exceed 40% duty cycle)
Max intermittent dwell	88% duty cycle for 9 s max

B.4.2. Static pressure transducer

Table B.2: Main specifications of the Omega PX419-1.5KAV transducer.

Specification	**Value**
Pressure type / range	Absolute, 0–1500 psi
Output signal	10 mV/V (ratiometric), “100 mV @ 10 Vdc”
Excitation / supply (mV/V models)	5–10 Vdc ratiometric; 10 Vdc nominal
Accuracy (combined)	±0.08% BSL (linearity + hysteresis + repeatability)
Temperature-compensated range	−29 to 85°C
Operating temperature range	−45 to 121°C
Proof pressure	300% minimum
Response time	< 1 ms
Electrical connection / IP rating	Mini-DIN 43650 C; IP65

B.4.3. Dynamic pressure transducer

Table B.3: Main specifications of the Kistler piezoelectric pressure sensors used in this work.

Specification	**Kistler 601CAB**	**Kistler 603CAA**
Sensor type / output	Piezoelectric (charge)	Piezoelectric (charge)
Pressure range	0–250 bar	0–1000 bar
Sensitivity (typ.)	−37 pC/bar	−10 pC/bar
Linearity (max.)	≤ ±0.1% FSO	≤ ±0.1% FSO
Operating temperature range	−196 to 350°C	−196 to 200°C
Rise time (max.)	≤ 1.4 μs	≤ 0.9 μs
Natural frequency (approx.)	215 kHz	430 kHz
Weight	1.9 g	7.8 g

B.4.4. Temperature sensor

Table B.4: Main specifications of the Omega PR-12-E RTD probe (configured as PR-12-E-3-100-1/4-2)

Specification	Value
Sensor type	Platinum RTD (Pt100)
Nominal resistance / TCR	100 Ω at 0°C; TCR = 0.00385 Ω/Ω/°C
Accuracy class	Class A (IEC 60751)
Temperature range (E = Extreme)	−200 to 600°C
Wiring	3-wire
Probe diameter / length (as configured)	1/4 in diameter; 2 in probe length
Process connection	1/2 in NPT
Ingress protection	IP67
Response time (T90, 1/4 in probe, PR-12E)	10.6 s

B.4.5. Oscilloscope

Table B.5: Main specifications of the B&K Precision BK2190E (Model 2190E) digital storage oscilloscope

Specification	Value
Instrument type	Digital storage oscilloscope (DSO)
Number of channels	2
Analog bandwidth	100 MHz
Real-time sampling rate	1 GSa/s (single ch.); 500 MSa/s (dual ch., fast timebase)
Record length (memory)	Up to 40,000 points (timebase-dependent)
Rise time (typ.)	< 3.5 ns
Vertical resolution	8-bit
Vertical sensitivity	2 mV/div to 10 V/div (1-2-5 sequence)
Maximum input voltage	400 V (DC+AC pk-pk, 1 MΩ input, X10), CAT I
Timebase range	2.5 ns/div to 50 s/div
Input impedance	1 MΩ±2% || 16 pF ±3 pF

B.4.6. Voltage probe

Table B.6: Main specifications of the PVM-5 probe

Specification	Value
Max DC / pulsed voltage (PVM-5)	60 / 100 kV
Standard divider ratio	1000:1
Rise time (nom., standard ratio)	3.0 ns
DC to 2 Hz accuracy	0.15%
Input resistance / capacitance	400 MΩ / 12 pF
Cable length	15 ft / 4.5 m
Probe length	19 in / 47 cm

B.4.7. Current probe

Table B.7: Main specifications of the CWT Mini50HF Rogowski current probe (model 015/B/1/100/2)

Specification	Value
Probe type	Shielded Rogowski coil + matched integrator (AC/pulsed current)
Bandwidth (-3 dB)	150 Hz to 50 MHz
Sensitivity / range (this model)	200 mV/A, 30 A
Accuracy (calibrated, reading)	±0.2% (conductor centred in coil)
Linearity (with current magnitude)	0.05% of reading
Coil geometry (Mini50HF)	100 mm coil circumference; 3.5 mm cross-section
Coil peak voltage isolation	2 kV peak
Temperature limits	Coil/cable: −40 to 125°C; electronics: 0 to 40°C
Power	4×AA alkaline batteries (integrator)

B.4.8. Gas booster

Table B.8: Main specifications of the PB30 gas booster

Specification	Value
Pump type	Air-driven gas booster (single acting)
Nominal pressure ratio	32:1
Displacement per cycle	80.42 cc (4.91 cu in)
Maximum outlet pressure	3843 psi (265 barg)
Inlet / outlet connections	1/4 in NPT / 1/4 in NPT

B.4.9. Environmental chamber

Table B.9: Main specifications of the TotalTemp Technologies C900 temperature chamber.

Specification	Value
Equipment type	Bench-top forced-convection temperature chamber
Temperature range	−100 to +200°C (−148 to 392°F)
Interior material	Stainless steel
Exterior material / finish	Cold rolled steel; grey hard powder coat
Capacity	9 ft^3 (255 L)
Interior dimensions (W×H×D)	25 × 25 × 25 in (63.5 × 63.5 × 63.5 cm)
Exterior dimensions (W×H×D)	30 × 38 × 29 in (76.2 × 96.5 × 73.6 cm)
Approximate weight	234 lb (as stated)
Heating rate (standard motor)	8 °C/min (208 VAC, 15 A)
Cooling rate (standard motor)	15 °C/min using 75 psi liquid nitrogen
Controller	Synergy Nano colour touchscreen
Coolant (standard)	High-pressure LN_2 (75–125 psi) supplied via dewars

C

Technical drawings

This appendix contains the technical drawings of the parts designed and used in this work.

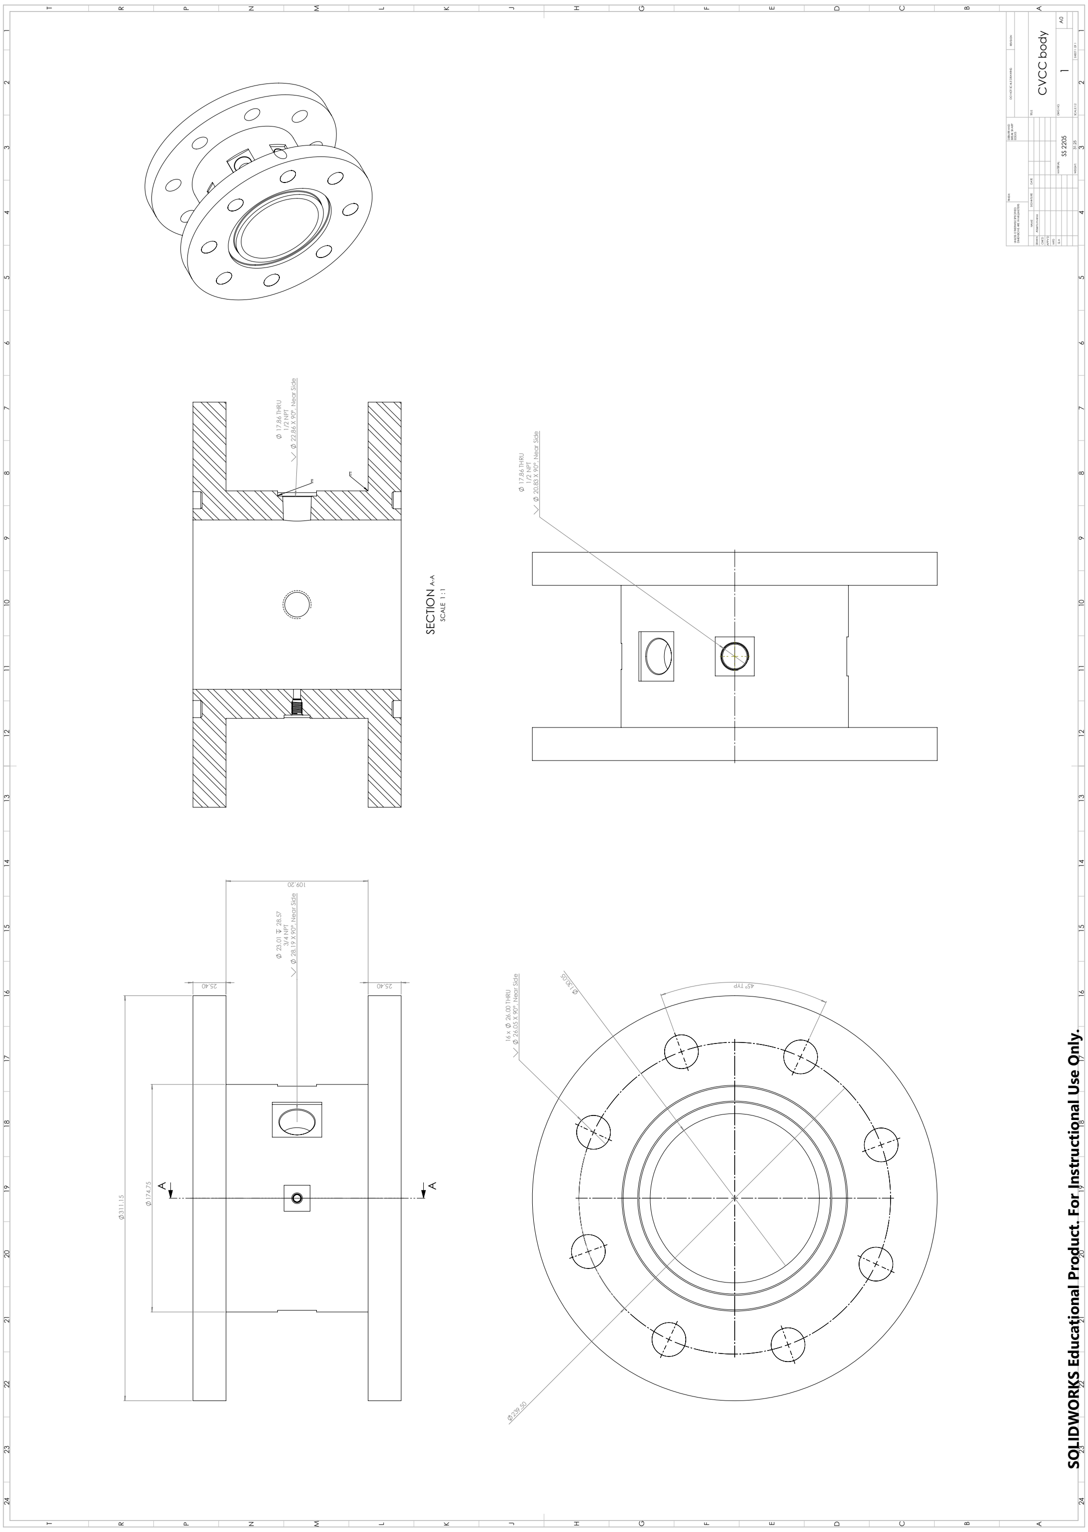

Figure C.1: Combustion chamber body

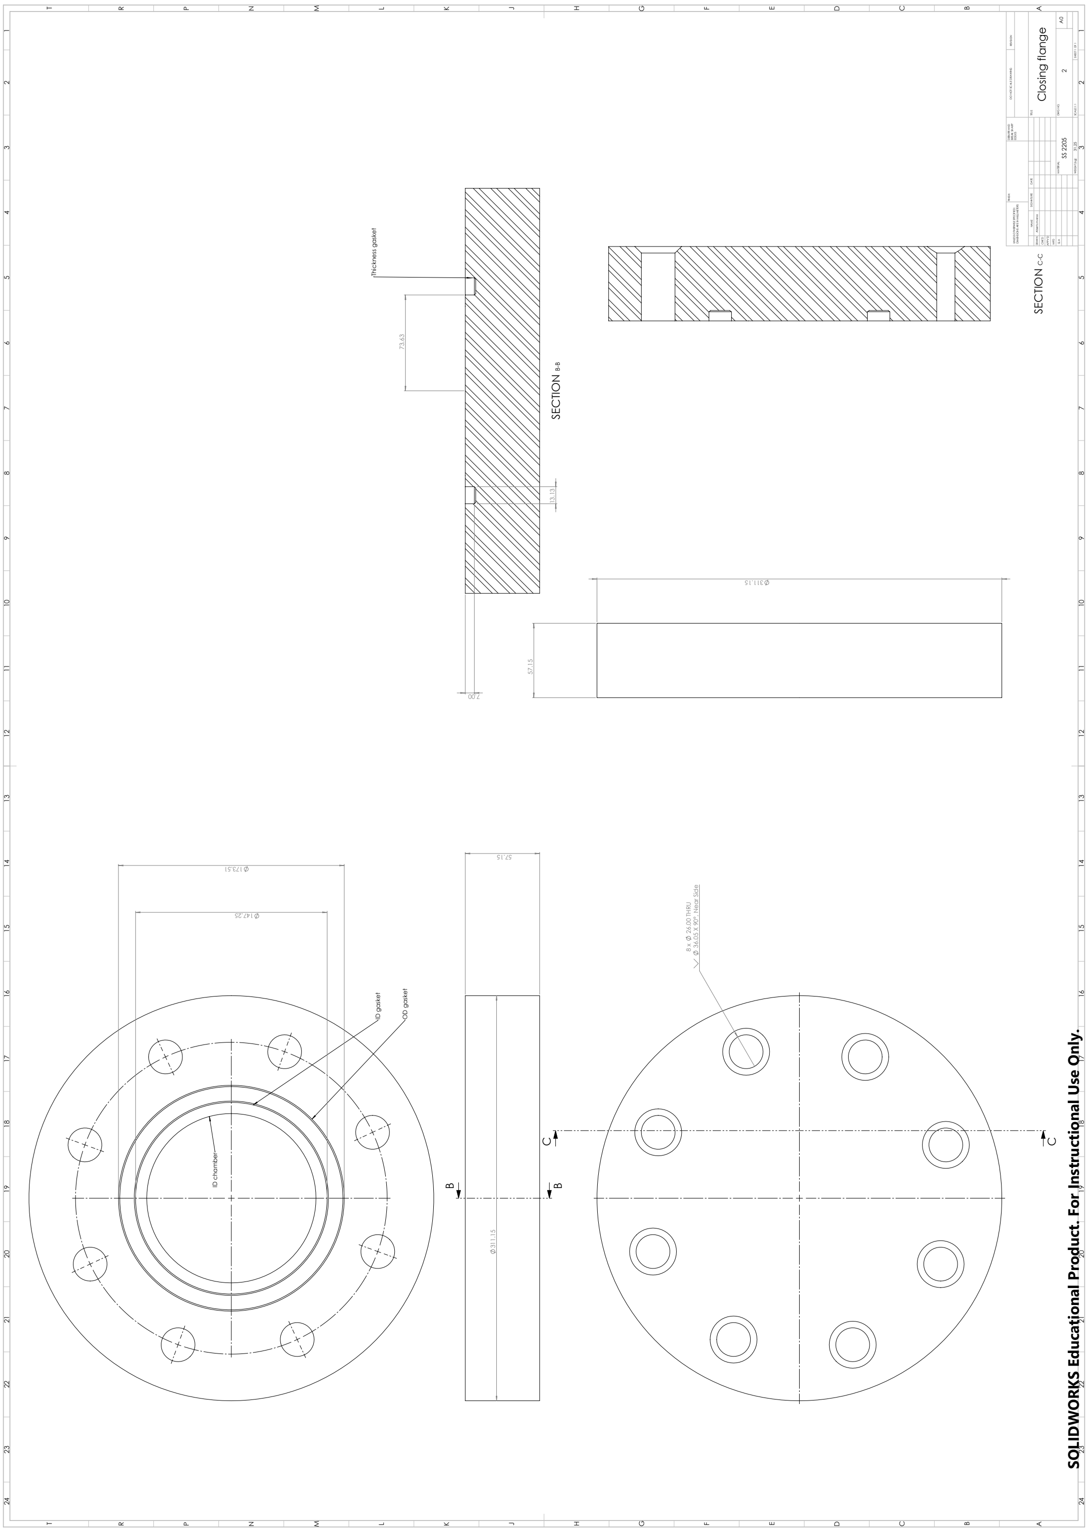

Figure C.2: Closing flange

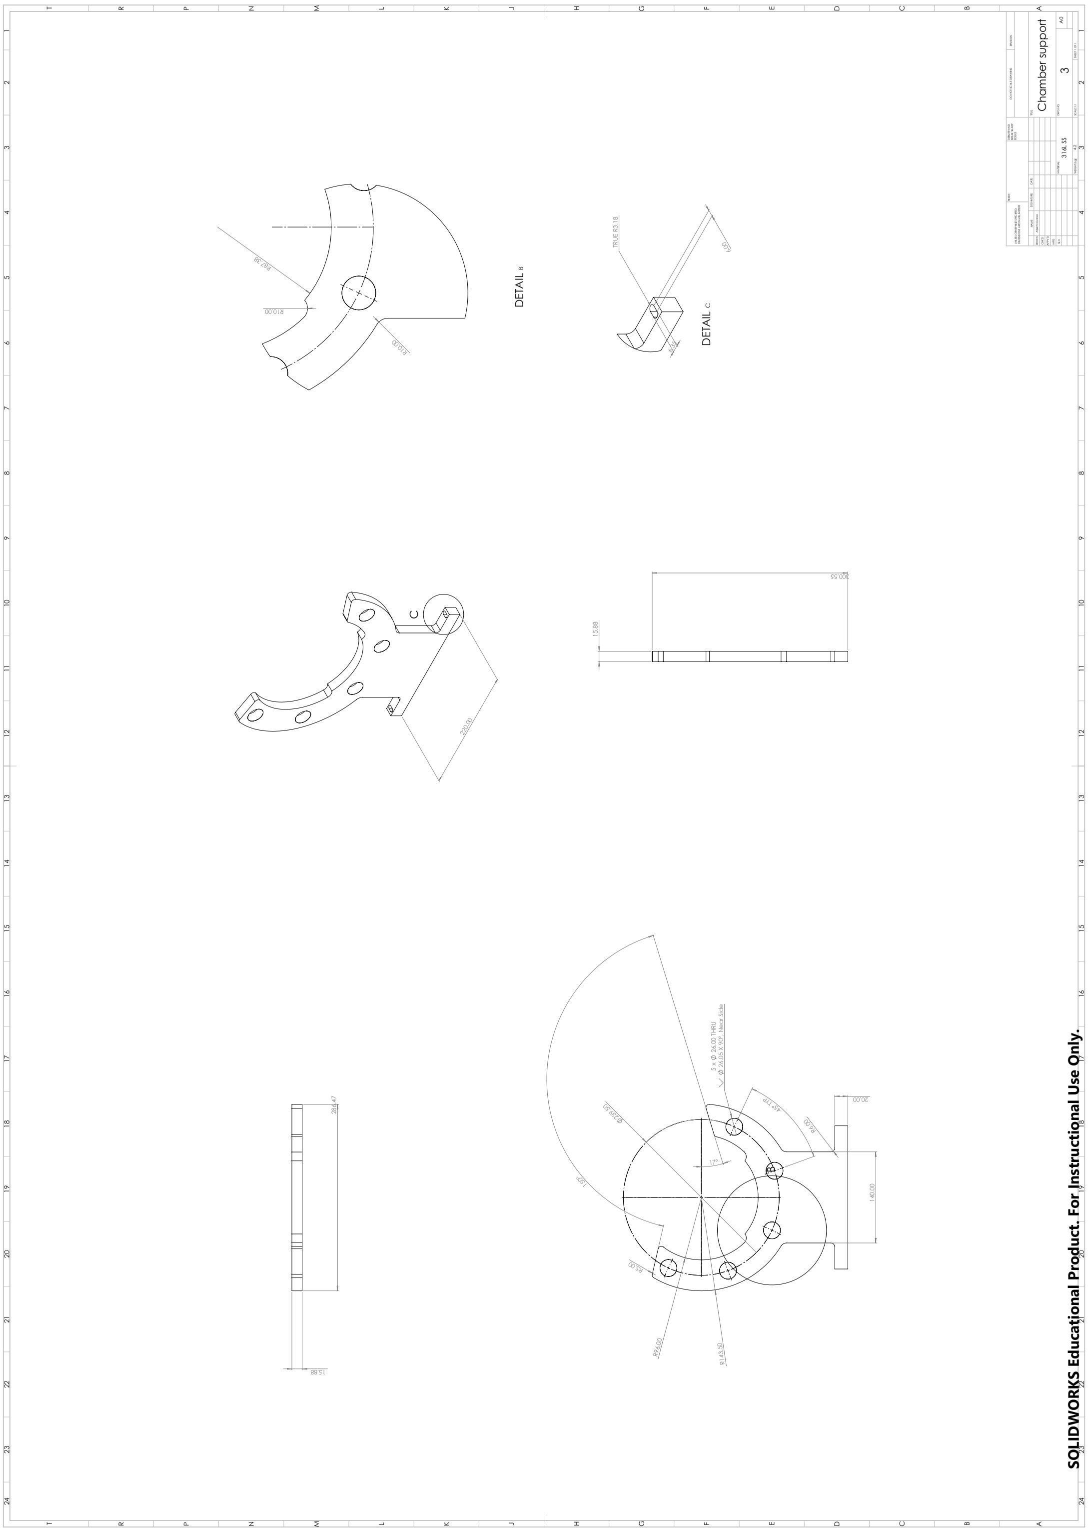

Figure C.3: Combustion chamber support

D

Experimental matrix

This appendix includes the list of all experiments performed in this work, in the form of an experimental matrix. This includes the initial parameters for each experiment, along with the run's objective.

Table D.1: Experimental matrix (A: Verification, B: Validation at STP, C: Validation of low-T setup, D: Research)

Exp. ID	Objective	Fuel	ER [-]	Set T [C]	Set p [bar]
VER1	A	CH_4	1	21	1
VAL1	B	CH_4	1	21	1
VAL2	B	CH_4	0.8	21	1
VAL3	B	CH_4	1.2	21	1
VAL4	B	H_2	1	21	1
VAL5	B	H_2	0.8	21	1
VAL6	B	NH_3	1	21	1
VAL7	B	NH_3	0.8	21	1
VAL9	B	H_2	1	21	2
VAL10	B	H_2	0.8	21	2
VAL13	B	NH_3	0.8	21	2
VAL14	B	NH_3	1	21	2
VAL17	C	H_2	0.8	0	1
VAL18	C	H_2	1	0	1
VAL19	C	H_2	0.8	-15	1
VAL20	C	H_2	1	-15	1
VAL21	C	H_2	0.8	-30	1
VAL22	C	H_2	1	-30	1
VAL23	C	H_2	0.8	-50	1
VAL24	C	H_2	1	-50	1
DIS1	D	NH_3	1	0	1
DIS2	D	NH_3	0.8	0	1

Table D.1 continued from previous page

Exp. ID	Objective	Fuel	ER [-]	Set T [C]	Set p [bar]
DIS3	D	NH_3	1.2	0	1
DIS4	D	NH_3	1	-15	1
DIS5	D	NH_3	0.8	-15	1
DIS6	D	NH_3	1.2	-15	1
DIS7	D	NH_3	1	-30	1
DIS8	D	NH_3	0.8	-30	1
DIS9	D	NH_3	1.2	-30	1
DIS10	D	NH_3	1	-50	1
DIS11	D	NH_3	0.9	-50	1
DIS12	D	NH_3	0.8	-50	1
DIS13	D	NH_3	1.1	-50	1
DIS14	D	NH_3	1.2	-50	1
DIS15	D	NH_3	1	-50	2
DIS16	D	NH_3	0.8	-50	2
DIS17	D	NH_3	1.2	-50	2
DIS20	D	NH_3/H_2 ($0.2H_2$)	0.8	-50	1
DIS21	D	NH_3/H_2 ($0.2H_2$)	1	-50	1
DIS22	D	NH_3/H_2 ($0.2H_2$)	1.2	-50	1
DIS23	D	NH_3/H_2 ($0.3H_2$)	0.8	-50	1
DIS24	D	NH_3/H_2 ($0.3H_2$)	1	-50	1
DIS25	D	NH_3/H_2 ($0.3H_2$)	1.2	-50	1
DIS26	D	NH_3/H_2 ($0.2H_2$)	0.8	-50	2
DIS27	D	NH_3/H_2 ($0.2H_2$)	1	-50	2
DIS28	D	NH_3/H_2 ($0.2H_2$)	1	-50	2
DIS29	D	NH_3/H_2 ($0.3H_2$)	0.8	-50	2
DIS30	D	NH_3/H_2 ($0.3H_2$)	1	-50	2
DIS31	D	NH_3/H_2 ($0.3H_2$)	1.2	-50	2

E

Chamber verification

Before arriving at the results shown in Figure 6.1, this process took considerable time at first because the chamber had numerous leaks, primarily around the igniters. These turned out to be only insulated electrically via the ceramic protection, with no fluid insulation guaranteed. For this reason, a layer of TorrSeal had to be applied to guarantee a complete fluid seal. Aside from this, it was found that the 603CAA initially exhibited very unusual and unexpected behaviour, as shown in Figure E.1. In fact, the expectation was of a rapid pressure increase, followed by a slower decay to the initial pressure as the chamber temperature lowered again once the flame quenched at the wall. However, the curve, after presenting the initial rise, exhibits erratic behaviour, with a decrease below 0, followed by a new increase and a final decrease to the initial pressure, almost as if it overshoots after the first rise. This behaviour is later explained as due to thermal shock in the sensor. This possibility was mentioned by the company during the sensor selection process, but was deemed quite remote. Nevertheless, the sensor was chosen for its extended temperature range and high acquisition speed. It is also evident that, compared to the curve shown in Figure 6.2b, which is one obtained with the sensor 7061C in the optical chamber, the maximum pressure achieved is much lower, and considerably lower than the approximate 8 bar value of the adiabatic pressure for this mixture.

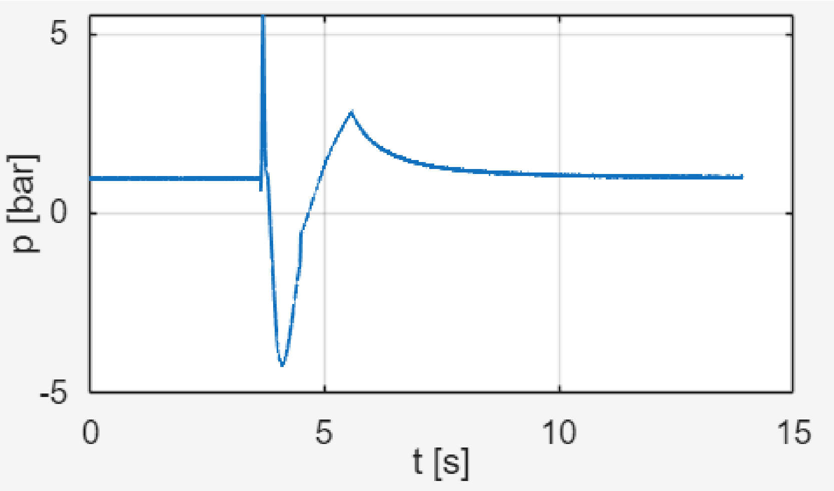

Figure E.1: Initial pressure measurements for CH_4 at $\phi = 1$, $T_0 = 21°$ C and $p_0 = 1$ bar

For this reason, the first attempt to solve this problem involved switching the sensor to the

601CAB, which arrived at the laboratory after a considerable delay and has the same dimensions, thereby avoiding modification of the hole geometry. On paper, this sensor appears to have better resistance to thermal shock. The sensor threads, however, broke the moment they were inserted, necessitating the re-boring of the entire hole to ensure the sensor could be repositioned. The 601CAB could be repositioned after smoothing the threads, but it could no longer guarantee the pressure seal. The results obtained under these conditions underline that this sensor is also susceptible to thermal shock, although its results are much closer to the optical ones, as shown in Figure E.2, but due to the use of hydrogen and ammonia, it would be a risk to continue using it if there were a leak.

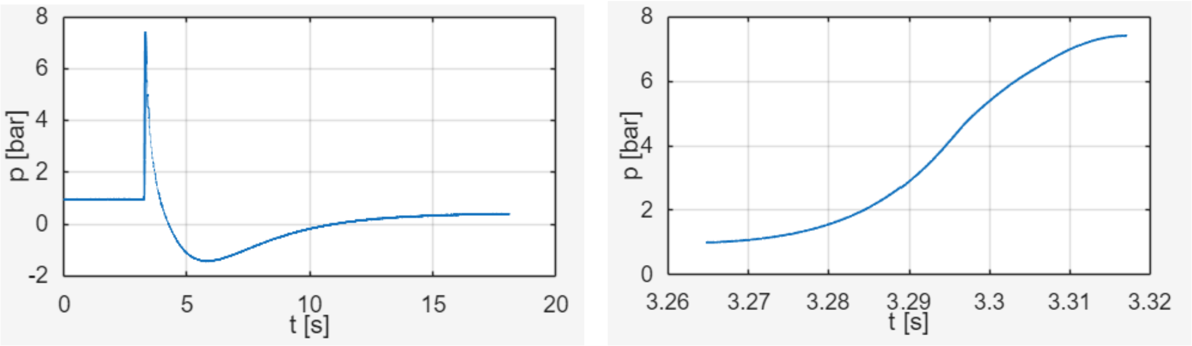

(a) Full pressure trace **(b)** Trace cropped around the combustion event

Figure E.2: 601CAB pressure trace

After re-boring the hole, the 603CAA is placed again with 1.5 mm of RTV material, which, in principle, is the best fix for thermal shock, as the RTV absorbs the bulk of the temperature rise, reducing the load on the sensor. However, despite its use, the sensor still undergoes thermal shock. Still, after the work on the hole, the sensor shows a pressure rise much more consistent with the one obtained in the optical chamber, It is therefore believed that the combination of the presence of the RTV and the fact that the hole contained some chips from manufacturing, probably causing the sensor to sit at an angle as compared to the chamber shoulder, the sensor can measure the combustion pressure rise quite accurately.