



Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG



DEVELOPMENT OF HIGH-AVERAGE-POWER
LASER PLASMA ACCELERATORS DRIVEN
BY INDUSTRIAL YB LASERS

*Dissertation zur Erlangung des Doktorgrades
an der Fakultät für Mathematik, Informatik und Naturwissenschaften
Fachbereich Physik der Universität Hamburg*

vorgelegt von:
BONAVENTURA FARACE

HAMBURG
2024

Gutachter der Dissertation:

Prof. Dr. Wim Leemans
Dr. Kristjan Pöder

Zusammensetzung der Prüfungskommission:

Prof. Dr. Wim Leemans
Dr. Kristjan Pöder
Dr. Christoph Heyl
Prof. Dr. Sven-Olaf Moch
Prof. Dr. Arwen Pearson

Vorsitzende der Prüfungskommission:

Prof. Dr. Sven-Olaf Moch

Vorsitzender Fach-Promotionsausschuss Physik:

Prof. Dr. Markus Drescher

Leiter des Fachbereichs Physik:

Prof. Dr. Wolfgang J. Parak

Dekan der Fakultät MIN:

Prof. Dr. Ing. Norbert Ritter

Datum der Disputation:

2. September 2024

Declaration of oath

I hereby declare and affirm that this doctoral dissertation is my own work and that I have not used any aids and sources other than those indicated. If electronic resources based on generative artificial intelligence (gAI) were used in the course of writing this dissertation, I confirm that my own work was the main and value-adding contribution and that complete documentation of all resources used is available in accordance with good scientific practice. I am responsible for any erroneous or distorted content, incorrect references, violations of data protection and copyright law or plagiarism that may have been generated by the gAI.

Hamburg, 29/07/2024



Bonaventura Farace

A Cecilia, grazie per esserci sempre stata. Senza di te sarei l'ombra di quel che sono ora.

A Maria Luisa, dove il mio viaggio finisce inizia il tuo. Anche se non te lo dico mai, sono fiero di te.

Contents

1	Introduction	11
2	Principles of laser plasma acceleration	15
2.1	Basics of plasmas	15
2.1.1	Debye shielding and plasma oscillations	16
2.2	Laser-driven plasma waves	17
2.2.1	Single particle motion	18
2.2.2	The ponderomotive force	18
2.2.3	Laser wakefields	20
2.3	Trapping and injection	24
2.3.1	Electron trajectories	24
2.3.2	Injection in the wakefield	26
2.4	Nonlinear plasma optics	30
2.4.1	Ionisation-induced defocusing	30
2.4.2	Relativistic self-focusing	33
2.4.3	Photon acceleration	34
2.5	High average power LPA	38
2.5.1	Scaling of the blowout regime	39
2.5.2	State-of-the-art of high-average-power LPA	40
3	VERA: an industrial laser for LPA	43
3.1	Post-compression to few-cycle duration	44
3.1.1	Spectral broadening	44
3.1.2	Dispersion compensation	45
3.2	The VERA laser	47
3.2.1	A two-stage cascaded MPC setup	47
4	Plasma sources for high-average-power LPA	53
4.1	Flow models	53
4.2	Gas nozzles	54
4.2.1	Subsonic nozzles	57
4.2.2	Supersonic nozzles	58
4.3	A novel tunable microfluidic source: the VERA source	63
4.3.1	The VERA source: first designs and tests	68

5	Modelling and optimisation of high-average-power LPA	71
5.1	Modelling LPAs	71
5.1.1	The particle-in-cell code FBPIC	71
5.1.2	Principles of Bayesian optimisation	72
5.2	Bayesian optimisation of high-average-power LPA	73
5.3	CFD-aided Bayesian optimisation of advanced plasma sources	76
5.3.1	Source geometry optimisation with ANSYS-FLUENT	76
5.3.2	Bayesian optimisation of the VERA source	78
6	Wakefields driven by industrial lasers	85
6.1	The experimental setup	85
6.2	On-target pulse characterisation	89
6.3	Plasma source characterisation	92
6.3.1	Interferometry	92
6.3.2	Density characterisation and simulations benchmarking	98
6.4	Ionisation-induced effects	103
6.4.1	Density clamping	103
6.4.2	Peak-intensity-reduction effects on the excited wakefields	105
6.5	Towards the first LPA driven by industrial Yb lasers	111
7	Conclusions and future perspectives	115
	Bibliography	117
A	Wavebreaking injection	131
B	Wigner distributions	133
C	Chromatic dispersion	137
D	Ansys-Fluent	139
D.1	Guidelines for fluid modeling	139
D.2	Turbulence models	141
E	Spectrally-resolved pulse characterisation	145
F	Interferometry	151
F.1	Taylor expansion of the plasma refractive index	151
F.2	Peak density and density scans	152
G	Mode decomposition of real focal spot images	155

Abstract

Over the past two decades, laser plasma accelerators (LPAs) have emerged as a groundbreaking technology with immense potential for electron acceleration. Their ability to sustain exceptionally high accelerating gradients, on the order of 100 GV/m, and provide electron bunches with only few femtoseconds duration, promises a compact, cost-effective solution for numerous industrial, commercial, and medical applications. Yet, to transition from experimental setups to practical applications, it is crucial to enhance their robustness, reliability, and repetition rate. In this context, industrial-quality Ytterbium:Yttrium-aluminium-garnet (Yb:YAG) lasers present an ideal, economically-efficient option, thanks to their inherently small quantum defect, high slope efficiency and high average power.

This thesis explores the feasibility of using industrial Yb:YAG lasers as drivers for plasma accelerators. Typically, these lasers deliver pulses of relatively long duration, extending to the picosecond-level. However, to excite a plasma wave, femtosecond durations are usually required. Hence, the temporal compression of the laser output to a few optical cycles is the first critical aspect under analysis. To address this challenge, an efficient double multi-pass-cell (MPC) post-compression scheme is employed, achieving the highest-ever compression factor to date for a 10 mJ-level pulse.

In addition to the driver laser, the plasma source plays a key role in the laser-plasma interaction, shaping the plasma density profile. Therefore, an extensive analysis of the most common sources for high-average-power LPA is presented. To overcome the limitations usually faced, a novel microfluidic source is proposed with unique capabilities for precise tailoring of the plasma profile along the laser axis, at the μm level. Its exceptional fine-tuning ability is demonstrated through a pioneering Bayesian optimisation tool, combining fluid dynamics and particle-in-cell simulations. This novel optimisation approach holds the promise to significantly boost the performances of LPA, particularly in application-oriented scenarios.

Finally, the post-compressed laser output is used to demonstrate, for the first time, a plasma wakefield driven by an industrial Yb:YAG laser. The laser-plasma interaction is thoroughly analysed and a clear path towards the first industrial-laser-driven electron acceleration is presented.

Zusammenfassung

In den letzten Jahrzehnten haben sich Laser-Plasma-Beschleuniger als revolutionäre Technologie für die Elektronenbeschleunigung etabliert.

Sie sind in der Lage, extrem hohe Beschleunigungsgradienten von bis zu 100 GV/m aufrechtzuerhalten, und ermöglichen die Erzeugung von Elektronenstrahlen mit Pulsdauern von wenigen Femtosekunden. Dies macht sie zu einer kompakten und kosteneffizienten Lösung für zahlreiche industrielle, kommerzielle und medizinische Anwendungen. Um jedoch den Übergang von experimentellen Aufbauten zu praktischen Anwendungen zu ermöglichen, ist es entscheidend, ihre Stabilität, Zuverlässigkeit und Wiederholrate zu verbessern. In diesem Zusammenhang bieten industrielle Yb-Laser dank ihres geringen quantenmechanischen Defekts, herausragender Effizienz und hoher durchschnittlicher Leistung eine vielversprechende, wirtschaftlich effiziente Alternative.

Diese Arbeit untersucht die Eignung industrieller Yb:YAG-Laser als Treiber für Plasma-Beschleuniger. Ein zentraler Aspekt ist die zeitliche Kompression des Laser-Ausgangspulses von seiner typischen Pikosekunden-Dauer auf nur wenige optische Zyklen. Zur Bewältigung dieser Herausforderung wird ein effizientes Kompressionschema eingeführt, das aus zwei Mehrfachdurchgangszellen besteht und den bisher höchsten Kompressionsfaktor für einen Laserpuls im Bereich von 10 mJ erreicht. Neben dem Laserpuls spielt die verwendete Plasmaquelle eine entscheidende Rolle bei der Laser-Plasma-Wechselwirkung, da sie das Plasmadichteprofil vorgibt. Es wird eine umfangreiche Analyse der gebräuchlichsten Quellen präsentiert, um anschließend eine neuartige mikrofluidische Plasmaquelle mit außergewöhnlichen Feinabstimmungsmöglichkeiten zu entwickeln. In diesem Kontext wird ein innovatives Werkzeug zur Optimierung von Plasmaquellen konzipiert, das Strömungsmechanik und Particle-in-Cell-Simulationen in einem Bayes'schen Optimierungsalgorithmus vereint und das Potenzial hat, die Leistungsfähigkeit von Plasma-Beschleunigern signifikant zu steigern, besonders in anwendungsorientierten Szenarien. Schließlich wird der komprimierte Laserpuls verwendet, um erstmals Plasma-wellen zu demonstrieren, die von einem industriellen Yb:YAG-Laser angetrieben werden. Die Laser-Plasma-Wechselwirkung sowie ihre charakteristischen Merkmale und Grenzen werden ausführlich untersucht, und es wird ein klarer Weg aufgezeigt, um die erste Elektronenbeschleunigung im Bereich von wenigen MeV mithilfe industrieller Yb-Laser zu erreichen.

CHAPTER 1

Introduction

Rarely has a tool in the modern era revolutionised society as profoundly as particle accelerators have.

Engines of scientific discovery, accelerators have been playing a pivotal role in advancing knowledge, contributing (to date) to a total of 46 Nobel prizes: 33 in physics, 10 in chemistry and 3 in medicine [1]. Aside from their contribution to fundamental research, it is perhaps less widely recognized that over 97% of the approximately 40 000 particle accelerators worldwide are used for industrial and commercial purposes [2]. Their impact, hence, extends far beyond pure scientific investigation, driving technological innovation and transforming a myriad of applications.

In material science, particle accelerators enable precise manipulation of matter at atomic and molecular scales, promoting advancements in electronics, semiconductor manufacturing, nanotechnology, and materials engineering. In healthcare, they have completely transformed medical imaging, radioisotope production and cancer treatment. They play a vital role in preserving cultural heritage, enabling non-destructive analysis of artworks and artifacts and they are used in environmental applications, for food sterilization and wastewater purification. Moreover, their use as neutron sources and in transmutation processes offers promising solutions for managing nuclear waste and advancing progress in both nuclear fission and fusion reactors. All these examples only scratch the surface, as the applications of particle accelerators continue to expand and evolve.

While large research accelerators keep growing in size and complexity, the ever-increasing amount of medical and industrial applications calls for more compact and cost-effective solutions. However, the accelerating field of traditional radiofrequency (RF) accelerators, the most widely used today, is ultimately limited by the electrical breakdown of their accelerating cavities, which restricts their field gradients well below 1 GV/m. This constraint not only limits the accelerator compactness, but also affects their overall costs. A promising alternative to overcome this limitation and reduce both the size and price of the machine involves the use of plasmas as accelerating medium, chiefly motivated by their ability to sustain fields 1000 times higher than RF cavities.

As initially proposed by Tajima and Dawson in 1979 [4], strong fields in plasmas, on the order of 100 GV/m, can be excited by an intense laser pulse, a concept which, nowadays,



Figure 1.1.: Left: Gordon Isaacs, the first pediatric patient treated with radiotherapy, 1957 (picture from the National Cancer Institute, via Stanford University). Right: the “Ritratto Trivulzio” by Antonello da Messina during a particle induced X-ray emission (PIXE) analysis at the INFN’s Laboratory for Cultural Heritage and Environment, Florence, Italy [3].

is known as “laser plasma acceleration” (LPA). Driven by these field gradients, a significant portion of the LPA community focuses on accelerating electrons to very high energies [5]. In the early 2000s, plasma accelerators were already able to produce high-quality electron beams in the 100 MeV¹ range [6, 7, 8]. Shortly thereafter, the GeV scale was reached [9], setting the stage for the multi-GeV energy regime [10, 11, 12, 13]. However, given that many of the aforementioned applications make use of low-energy electrons beams (ranging from few MeV to 10-20 MeV), laser plasma accelerators can unveil exciting avenues also in the low-energy regime. Typically, the laser limits the operation of LPAs at Hz repetition rates. Nonetheless, when targeting low energies, these machines can function with significantly reduced laser pulse energies. Consequently, their repetition rate can be increased by several order of magnitudes, possibly extending to the multi-kHz regime. These accelerators usually exploit high-average-power lasers and, therefore, are commonly referred to as “high-average-power LPAs”.

In recent years, advancements in this particular field have largely been driven by Titanium:Sapphire (Ti:Sa) lasers [14, 15, 16]. However, while these systems provide the highest pulse peak power off the shelf, they are particularly energy inefficient and their susceptibility to heat poses severe limitations on their day-to-day stability. Therefore, to facilitate the advent of LPA technology in industrial and commercial applications, a more cost-effective and scalable solution is necessary, with the potential to reach even higher repetition rates. Here, industrial Ytterbium:Yttrium-aluminium-garnet (Yb:YAG) lasers present a compelling alternative, offering a balanced peak power and a significantly higher efficiency and relia-

¹The electronvolt (eV) is the preferred unit of measure of the energy when treating the atomic scale. By definition, 1 eV is the energy gained by an electron accelerated by a potential difference of 1 V. As an example an electron accelerated by a field gradient of 100 GV/m would reach 100 MeV in only 0.001 m, so in 1 mm.

bility [17].

This thesis, specifically, explores the potential use of industrial-quality, Yb lasers as drivers for laser plasma acceleration. These lasers have never been employed for such purposes, yet their implementation could represent a groundbreaking advancement, providing a substantial boost to the practical use of LPAs in industrial and commercial settings. In medicine, for instance, their typically high peak electron flux holds great promise for radiotherapy [18], while in material science their ability to accelerate electron bunches of a few femtoseconds, inherently synchronized with the driving laser, unveils exciting opportunities in ultra-high temporal resolution measurements at the atomic scale [19].

In detail, chapter 2 covers the basic principles of the physics underlying laser plasma acceleration, with special emphasis on nonlinear plasma optics, particularly relevant when industrial lasers are considered. Chapter 3 introduces the industrial, post-compressed Yb:YAG “VERA” laser, the primary tool for all the experimental work presented in the thesis. Chapter 4 analyses, from a fluid dynamics perspective, the most common plasma sources used for high-average-power laser plasma acceleration. Notably, it introduces a novel microfluidic source, with unique fine-tuning capabilities and a transformative potential to improve the performance of such accelerators. Chapter 5 focuses on the main computational tools for modeling the laser-plasma interaction and reassesses the described sources in terms of the electron beams they can generate. In particular, a pioneering combination of computational fluid dynamics (CFD) and particle-in-cell (PIC) simulations in a Bayesian optimisation loop exploits machine learning algorithms in the optimisation of plasma source designs. Chapter 6 is dedicated to the experimental results and addresses the main challenges that arise employing industrial lasers in LPA. It provides, for the first time, evidence of a plasma wakefield excited by an industrial Yb:YAG laser and propose a clear path towards the first electron acceleration. Finally, chapter 7 summarises the main findings of this work, highlighting the importance of developing industrial-lasers-based LPA setups as reliable, compact and cost-effective few-MeV electron sources.



CHAPTER 2

Principles of laser plasma acceleration

Imagine compressing the energy generated by a hundred thousand (100,000) lightning bolts into a space no larger than the tip of a ballpoint pen. That is the peak intensity delivered by a single high-power laser pulse when focused down to a few micron spot size, about 10^{18} W/cm². In a laser plasma accelerator (LPA), these ultra-intense pulses interact with a tenuous plasma and can excite exceptionally strong electric fields, orders of magnitude higher than the maximum field achievable in radiofrequency particle accelerators. Consequently, unlike their more traditional counterpart, these compact and versatile devices could fit on a tabletop, opening up a realm of possibilities for scientific discovery, medical advancements, and industrial applications.

To begin with, the upcoming chapter is devoted to the theoretical aspects underpinning the laser-driven acceleration process. The laser-plasma interaction will be treated extensively, from the excitation of plasma waves up to the injection and acceleration of plasma particles, with particular emphasis to high-average-power lasers.

SECTION 2.1

Basics of plasmas

From stars and interstellar medium to the ionosphere and lightning phenomena, plasmas pervade the observable universe. They encompass an extraordinarily extensive spectrum of particle density and temperature, spanning approximately twenty-five orders of magnitude for the former and seven for the latter¹. Plasmas employed in accelerators are only a tiny fraction of this vast sea. Therefore, it becomes extremely important to introduce and delineate the main parameters describing the “fourth state of matter”, essential to gain a deeper insight on the underlying physics governing laser-plasma interaction.

A plasma can be described as a quasi-neutral gas of charged particles showing collective behaviour. The quasi-neutrality refers to the presence of a local balance between positive and negative charges, whereas the collective behaviour arises because of the long range

¹For comparison, consider that from the density of air to the one of neutron stars there are approximately eighteen orders of magnitude involved.

nature of the Coulomb potential. Generally, then, macroscopic fields dominate over microscopic fluctuations, so that as soon as there's a net charge imbalance $\rho = e(Z n_i - n_e)$, an electrostatic field $\nabla \cdot \vec{E} = \rho / \epsilon_0$ will appear. Likewise a current density $J = e(Z n_i v_i - n_e v_e)$ will give rise to $\nabla \times \vec{B} = \mu_0 \vec{J}$ ². These internally driven fields largely determine the plasma dynamics, including its response to an externally applied field as, for instance, a laser pulse.

2.1.1. Debye shielding and plasma oscillations

The characteristic length describing how far the electrostatic potential of a charge carrier is felt by other particles, the Debye length (λ_D), plays a fundamental role in defining a plasma. It can be expressed as:

$$\lambda_D = \left(\frac{\epsilon_0 k_B T_e}{e^2 n_e} \right)^{1/2} \simeq 743 \left(\frac{T_e[\text{eV}]}{n_e[\text{cm}^{-3}]} \right)^{1/2} \text{ cm}, \quad (2.1)$$

where ϵ_0 is the vacuum permittivity, k_B is the Boltzmann constant, n_e is the electron number density and T_e is the electron temperature³. Equation 2.1 allows the plasma parameter g , inversely proportional to the number of plasma particles in a Debye sphere, to be defined:

$$g = \frac{1}{n_e \lambda_D^3}. \quad (2.2)$$

When a Debye sphere is densely populated ($g \ll 1$) collective effects dominate over collisions between particles and the medium can actually be addressed as “plasma”.

The electron density n_e , moreover, is extremely important to define how fast plasmas can respond to an externally applied perturbation. An intuitive picture can be given considering a quasi-neutral plasma slab where an electron layer has been displaced by a distance δ , as depicted in figure 2.1.




Figure 2.1.: Scheme representing the initial position of an oscillating electron layer in a quasi-neutral plasma slab.

The displacement creates two slabs with a surface charge $\sigma = \pm e n_e \delta$, negative for the displaced electrons and positive for the charge left behind. The electron layer is hence accelerated back towards the positive one by a force:

²Here e is the elemental charge, n_e , v_e , n_i and v_i are respectively the electrons and the ions number densities and velocities, Z is the ions charge, \vec{E} and \vec{B} are the electric and magnetic fields and ϵ_0 and μ_0 are the permittivity and permeability in vacuum.

³Typical LPA plasma temperature is in the few eV range (10^4 Kelvin)

$$F \equiv m_e \frac{dv}{dt} = -m_e \frac{d^2\delta}{dt^2} = -eE = e^2 n_e \delta / \epsilon_0. \quad (2.3)$$

Equation 2.3 can be rewritten in the form of a harmonic oscillator as:

$$\frac{d^2\delta}{dt^2} + \frac{e^2 n_e}{\epsilon_0 m_e} \delta = 0, \quad (2.4)$$

where the (electron) plasma frequency is defined as:

$$\sqrt{\frac{e^2 n_e}{\epsilon_0 m_e}} = \omega_p. \quad (2.5)$$

The characteristic time⁴ over which the plasma evolves, then, can be expressed as the inverse of the characteristic frequency:

$$t_D = 1/\omega_p. \quad (2.6)$$

This response time (and therefore the plasma frequency) is of fundamental importance in the interaction with externally applied fields. An external electromagnetic radiation, in fact, will be “shielded out” if its period is longer than the plasma response time, or equivalently if $\omega_p/\omega_L > 1$. Hence $\omega_p/\omega_L = 1$ defines the density over which radiation with frequency ω_L cannot propagate in a plasma with characteristic frequency ω_p , the “critical density”.

$$\omega_p^2/\omega_L^2 = 1 = \frac{e^2 n_e}{\epsilon_0 m_e} \frac{\lambda^2}{4\pi^2 c^2}, \quad (2.7)$$

$$n_c = \frac{e_0 m_e 4\pi^2 c^2}{e^2 \lambda^2} \simeq 10^{21} \lambda_{[\mu\text{m}]}^{-2} \text{cm}^{-3}. \quad (2.8)$$

A laser of a given wavelength λ , hence, can penetrate a plasma with density $n < n_c$ and, if enough intense, can give rise to fields particularly suitable for electron acceleration.

SECTION 2.2

Laser-driven plasma waves

Propagating in a tenuous plasma, an intense laser pulse is able to generate plasma oscillations commonly known as “(laser-driven) plasma waves”. These waves can sustain electric field as high as 100 GeV/m, thereby showing groundbreaking potential for electron acceleration over extremely short distances. In the following the physics underlying the wave excitation is outlined, starting from a 1D, single particle approximation.

⁴In the following equation t_D refers to “Debye”, as the response time can also be expressed as the ratio of the Debye length divided by the thermal velocity $t_D = \lambda_D/v_{te}$.

2.2.1. Single particle motion

In the interaction with a laser pulse, the plasma response is usually dominated by the electrons motion. Thus, here, the dynamics of a single electron in presence of an electromagnetic field will be briefly described, as it underpins most of the laser-plasma acceleration physics.

The fundamental interaction between electromagnetic fields and charged particles is described by the Lorentz equation:

$$\frac{d\vec{p}}{dt} = -e(\vec{E} + \vec{v} \times \vec{B}). \quad (2.9)$$

Specifically, the electric field \vec{E} will give rise to an in-phase oscillation, while the magnetic component $\vec{v} \times \vec{B}$ (relevant only for relativistic electron velocities) will result in a particle drift along the longitudinal direction. Often the electromagnetic field can be more effectively described in terms of its potentials \vec{A} and ϕ , defined according to $\vec{E} = -\partial\vec{A}/\partial t - \nabla\phi$ and $\vec{B} = \nabla \times \vec{A}$. For instance, assuming a linearly polarised plane wave travelling along z , the vector potential \vec{A} can be expressed as:

$$\vec{A}(\omega, k) = (A_x, A_y, A_z) = (\pm A_0 \cos \varphi, 0, 0), \quad (2.10)$$

where the wave phase φ is equal to $\omega t - kz$. From the conservation laws of canonical momentum and energy (assuming an electron initially at rest, before the interaction with the laser field), the electron momentum $p = (p_x, p_y, p_z)$ and its trajectories can be derived as [20, 21]:

$$x = a_0 \sin \varphi, \quad (2.11)$$

$$z = \frac{a_0^2}{4} \left(\varphi + \frac{\sin 2\varphi}{2} \right). \quad (2.12)$$

Figure 2.2 represents the electron dynamics in the laboratory frame. It is characterised by an oscillatory motion in the transverse direction, commonly referred to as “quiver” motion, and a drift in the longitudinal one, with relative importance depending on the parameter a_0 . This parameter is crucial, as it expresses the (maximum) normalised momentum of the electron oscillation $a_0 = p_q/c m_e$ or, equivalently, the normalised amplitude of the vector potential $a_0 = \frac{eA_0}{m_e c^2}$. The value $a_0 = 1$, for instance, marks the limit over which the electron motion becomes relativistic, profoundly impacting the laser-plasma interaction⁵.

2.2.2. The ponderomotive force

A linearly polarised plane wave is a valid simplification to portray the oscillatory quiver motion of an electron, however when dealing with real laser fields, temporal and spatial gradients have to be considered. Generally, the field of a laser pulse can be described by an

⁵In this context a Lorentz factor related to the quiver electron motion can be defined as $\gamma_{\perp} = (1 + a^2)^{\frac{1}{2}}$, which will be of use in the next sections.

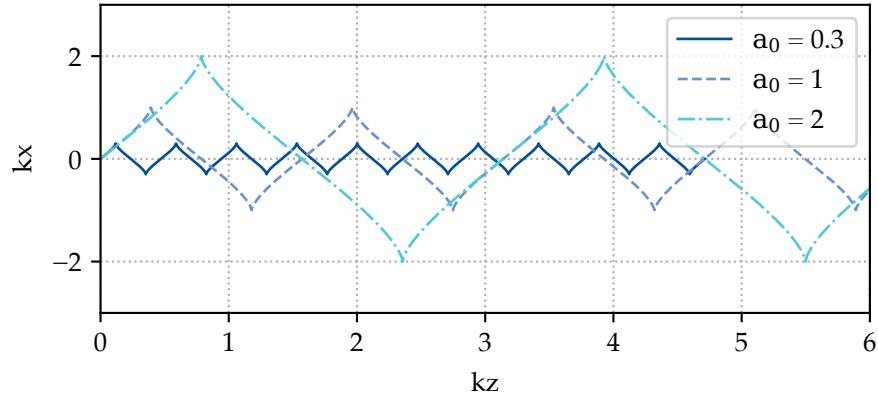


Figure 2.2.: The (quiver) motion of a test electron under the action of a monochromatic plane wave in the laboratory frame, in vacuum, for 3 different values of a_0 .

envelope and an oscillating function [22]:

$$\vec{E}(\vec{r}, t) = \frac{1}{2} \vec{E}_{\text{env}}(\vec{r}, t) e^{-i\omega t} + c.c., \quad (2.13)$$

where \vec{r} and t are the space-time variables, ω is the central laser frequency and *c.c.* stands for “complex conjugate”. The envelope function varies on a time scale much longer than the oscillation period, therefore $\langle \vec{E}(\vec{r}, t) \rangle_T \simeq 0$, while $\langle \vec{E}_{\text{env}}(\vec{r}, t) \rangle_T \neq 0$ ⁶. A slowly varying term (“*s*”) related to the envelope and a fast oscillating one (“*o*”) can then be identified, so that the electron motion can be expressed as:

$$r(t) = r_s(t) + r_o(t), \quad (2.14)$$

and as $\langle r_o(t) \rangle_T = 0$,

$$\langle r(t) \rangle_T = \langle r_s(t) \rangle_T = r_s(t). \quad (2.15)$$

Over a period, hence, the averaged electron motion is due to the “slow” envelope term and can be described by a slowly varying force: the ponderomotive force.

Assuming that:

- ◇ the electron motion is non-relativistic and
- ◇ the spatial variation of the envelope field over a light cycle is negligible⁷,

the ponderomotive force can be defined as [22]:

$$F_p = -\frac{e^2}{2m_e \omega^2} \nabla \langle \vec{E}^2(\vec{r}_s, t) \rangle_T = -\nabla \Phi_p, \quad (2.16)$$

⁶Here the notation $\langle x \rangle_T$ stands for the mean value of the variable x over the oscillation period $T = 2\pi/\omega_L$.

⁷this assumption is not strictly valid for ultrashort pulses, composed by only few light cycles [23].

where e and m_e are the electron charge and mass⁸. As equation 2.16 is similar to the generic relation between a force and a potential, Φ_p is often addressed as ponderomotive potential and express the oscillation energy of the quiver motion of the electron averaged over an oscillation period.

A charged particle in a laser field, hence, is expelled from the high intensity regions and pushed towards regions of lower potential. This is the driving mechanism for the excitation of plasma waves during a laser-plasma interaction.

2.2.3. Laser wakefields

The collective behavior of charged particles in a plasma can give rise to a wide range of oscillations. These can be electrostatic or electromagnetic and their dynamics can be very different depending on the plasma temperature and on the characteristic of the incident radiation [24]. However, in laser-plasma acceleration temperature effects can often be neglected and frequently only linearly polarised pulses are employed. Therefore, in the following, the plasma response to these specific conditions will be analysed, focusing on its ability to sustain accelerating structures.

1D Linear regime When the quiver motion of the plasma electrons is subrelativistic ($a_0 < 1$), the plasma response can be assumed linear and can be described by a simple 1D model. Starting from the plasma fluid equations one can express the density perturbation (n_1), the related electric field (\vec{E}) and the correspondent pseudo scalar potential (ϕ) as a function of the laser field strength [25, 26]:

$$\frac{\partial^2 n_1}{\partial \xi^2} + k_p^2 n_1 = \frac{1}{2} n_0 \frac{\partial^2 a^2}{\partial \xi^2}, \quad (2.17)$$

$$\frac{\partial^2 E}{\partial \xi^2} + k_p^2 E = -\frac{1}{2} k_p^2 \left(\frac{m_e c}{e} \right)^2 \frac{\partial a^2}{\partial \xi}, \quad (2.18)$$

$$\frac{\partial^2 \phi}{\partial \xi^2} + k_p^2 \phi = -\frac{1}{2} k_p^2 \left(\frac{m_e c}{e} \right)^2 a^2, \quad (2.19)$$

where, for simplicity, the description is referred to a frame co-moving with the laser field, at the speed of light, with coordinates $\xi = c t - z$ and $\tau = t$.

Equation 2.17, in particular, describes the oscillation of the density perturbation n_1 with frequency ω_p , excited by the ponderomotive force of the laser pulse, where $a = a_0 \exp(-\frac{\xi^2}{2L^2})$. As shown in figure 2.3, this oscillation follows the pulse propagation, with a phase velocity equal to the laser group velocity and results in an oscillating longitudinal electric field E . This structure is commonly referred to as “laser wakefield” and can be exploited to accelerate electrons during the laser-plasma interaction.

Stronger pulses are able to excite stronger wakefields, exceeding 100 GV/m, with ground-breaking potential for particle acceleration. However, the accelerating field cannot increase

⁸Here we are referring directly to the electrons. In reality each charged particle in a laser field will feel a ponderomotive force, but the $1/m$ scaling makes it negligible (most of the times) for heavier particles.




Figure 2.3.: Linear wakefield generated by a gaussian pulse with $a_0 = 0.3$ (red) and pulse length $L = \lambda_p/4$ propagating in an underdense plasma. The electron density perturbation $\delta n = n_1 - n_0$ (blue) and the resulting longitudinal electric field E (green) are normalised, respectively, to the initial plasma density n_0 and to the cold wavebreaking limit $E_0 = \frac{m_e c \omega_p}{e}$.

indefinitely. When the wake amplitude becomes excessively large, the electrons can lose their coherence, causing a local breakdown of the fluid description, a phenomenon known as “wavebreaking”. For a cold plasma, in the linear regime, the maximum achievable field can be intuitively estimated assuming that all the plasma electrons are oscillating with a frequency ω_p , as [27, 28]:

$$E_0 = \frac{m_e c \omega_p}{e}. \quad (2.20)$$

1D Nonlinear regime While in the linear regime the plasma wave is a simple sinusoidal oscillation with frequency ω_p , increasing the a_0 the quiver motion becomes relativistic and the plasma response highly nonlinear. For an intense enough driver with the focal spot size on the order of $\lambda_p = \frac{2\pi c}{\omega_p}$, the ponderomotive force acts both longitudinally and transversely. It expels all plasma electrons from the laser propagation region, leaving only the much more massive ions in place. The restoring attractive force of the ions then pulls the electrons back, creating a high density electron spike at the back of the driver and resulting in electric fields particularly suitable for electron acceleration [28].

To fully understand the field structure of nonlinear plasma waves and their importance in laser-plasma acceleration, a multidimensional approach is needed. However, assuming that the driver is not evolving (i.e. the quasi-static approximation can be applied), both the pulse and the plasma fluid quantities can be expressed as a function only of ξ and the nonlinear dynamics can be treated analytically in 1D. Such description, even if limited, provides a good overview about nonlinear plasma waves, which are the core of most laser plasma acceleration processes.

In 1D, the evolution of the normalised pseudo scalar potential in response to the laser field can be written as [26, 28]:

$$\frac{1}{k_p^2} \frac{\partial^2 \phi}{\partial \xi^2} = \frac{1}{2} \left[\frac{1+a^2}{(1+\phi)^2} - 1 \right]. \quad (2.21)$$

From equation 2.21 the evolution of both the density perturbation and of the longitudinal field can be derived as:

$$\frac{n_1}{n_0} = \frac{1}{k_p^2} \frac{\partial^2 \phi}{\partial \xi^2} + 1, \quad (2.22)$$

$$E = \frac{1}{k_p} \frac{\partial \phi}{\partial \xi}. \quad (2.23)$$

As shown in figure 2.4, the plasma wave loses its sinusoidal nature and evolves into a more complex structure, with peaks and periodicity depending on the driver a_0 . A more intense laser results in a stronger ponderomotive force and, thus, a longer λ_p , higher density spikes and, therefore, stronger wakefields.




Figure 2.4.: Nonlinear 1D wakefield generated by laser pulses propagating in an underdense plasma. The top panel shows the electron density perturbation $\delta n = n_1 - n_0$ (blue) normalised to the initial plasma density n_0 and the resulting longitudinal electric field (green), normalised to the cold wavebreaking limit $E_0 = m_e c \omega_p / e$, for a pulse with $a_0 = 2$ (red). The bottom is relative to a pulse with $a_0 = 2.5$.




Figure 2.5.: Example of a plasma wakefield in the nonlinear regime ($a_0 = 2$) simulated with the particle-in-cell code FBPIC [31]. The top panel depicts the general wave structure in the blowout regime. The bottom shows the longitudinal and transverse fields within the first cavity of the plasma structure, outlined by the dotted ellipse.

3D Nonlinear plasma waves and the blowout regime A 1D treatment provides a good qualitative picture, yet usually the transverse evolution of the laser driver is comparable in size to the plasma wavelength. Therefore the electron dynamics cannot be treated only in the longitudinal dimension and a 3D description is needed. A fully analytical model is given in [29, 30], where the electron trajectories are described considering all the forces acting during the interaction: the ponderomotive force, the electrostatic field due to the background plasma ions and the force due to the effect of radial plasma currents. In the following, though, the description is limited to a more phenomenological model, highlighting the main characteristics of nonlinear plasma waves.

In the nonlinear regime, the laser drives a wakefield which resembles a spherical cavity⁹

⁹The trajectory of the innermost electrons layer in the cavity can be described by $r_b \frac{d^2 r_b}{dz^2} + 2 \left(\frac{dr_b}{dz} \right)^2 + 1 = 0$ [29, 30]. Aside from the 2 factor, this is the equation of a sphere. The blowout regime, indeed, is often addressed as “bubble regime”, from the shape of the plasma cavity.

of radius r_b , with a strong density peak at the back of it, as shown in figure 2.5. To excite these fields, it is particularly favourable for the laser to be resonant (or “matched”) with the characteristic distances of the plasma dynamics, both longitudinally and transversely [29, 30]. Along the laser axis, matched propagation is achieved when the pulse length is on the order of the radius of the plasma cavity:

$$c \tau \sim \frac{2}{3} r_b, \quad (2.24)$$

where c is the speed of light in vacuum and τ is the pulse duration. On the other hand a pulse is matched transversely when the transverse ponderomotive force of the laser is balanced by the restoring force for the ions. This simply requires the laser spot size w_0 to be similar to r_b , which translates in a condition on the laser field strength a_0 :

$$k_p w_0 = k_p r_b = 2\sqrt{a_0}, \quad (2.25)$$

where $k_p = 2\pi/\lambda_p$ is the plasma wavenumber.

Under matched conditions, the laser propagation becomes particularly stable and the plasma creates a field structure very favourable for electron acceleration. With reference to figure 2.5, the high electron density at the back of the cavity generates a strong negative longitudinal electric field (E_z), which expands for almost half of the bubble. An electron placed in this region would hence feel a positive force and be accelerated forward. On the other hand, in the transverse direction, the linear gradient in E_y , in combination with the azimuthal magnetic field B_θ , acts as a focusing force, trying to keep the negatively charged particles on the propagation axis. These two fields together can accelerate and guide the electrons during the laser plasma interaction and constitute the main advantage of the blowout regime.

SECTION 2.3

Trapping and injection in a plasma wave

While a plasma wave has the capacity to sustain strong accelerating fields, their excitation alone does not inherently entail an electron acceleration process. In fact, in order to be trapped within the wakefield and acquire energy from it, an electron needs to be in the correct phase with respect to the driver and with sufficient initial energy. Hence, in the upcoming section the main techniques employed for electron “injection” into a plasma wave will be outlined, along with the necessary conditions for their successful “trapping” and acceleration.

2.3.1. Electron trajectories

The dynamics of a single electron in a plasma wave can be described by its Hamiltonian [32, 33]. In one-dimension:

$$\mathcal{H}(u, \psi) = \sqrt{\gamma_\perp^2 + u_z^2} - \beta_p u_z - \phi(\psi), \quad (2.26)$$

where u_z is the longitudinal momentum normalised to $m_e c$, $\beta_p \simeq 1 - \omega_p^2/2\omega_L^2$ is the normalised phase velocity of the plasma wave, $\phi(\psi) = e\Phi(\psi)/m_e c^2$ is the normalised scalar potential, $\gamma_\perp = (1 + u_\perp^2)^{\frac{1}{2}}$ is the transverse Lorentz factor and u_\perp is the normalised transverse momentum. A constant hamiltonian $\mathcal{H} = H_0$ describes the motion of an electron on a distinct orbit in the plasma wave.

Solving the equation for the longitudinal momentum, the electron trajectories in the phase space (u_z, ξ) can be derived as [33]:

$$u_z(\xi) = \beta_p \gamma_p^2 (H_0 + \phi) \pm \gamma_p \sqrt{\gamma_p^2 (H_0 + \phi)^2 - \gamma_\perp^2}, \quad (2.27)$$

where $\gamma_p = (1 - \beta_p^2)^{-\frac{1}{2}}$ is the Lorentz factor associated with the plasma wave phase velocity and u_z represents an electron orbit of constant total energy for a given set of $a(\xi)$ ¹⁰, $\phi(\xi)$ and H_0 . Figure 2.6, for instance, is showing different trajectories in the linear regime, for a laser with $a_0 = 0.3$. Electrons initially at rest, pushed by the ponderomotive force of the laser without gaining any substantial momentum, will follow the orbit defined by $H_0 = 1$, often addressed as “fluid orbit”, and will constitute the plasma wave. Likewise, electrons with too low initial momentum will not be confined in the (u_z, ξ) phase space and will follow an “open” trajectory. On the other hand, a particle moving with the plasma wave will remain within the same ξ region and describe a “closed” orbit. Therefore it can be “trapped” in the wakefield and accelerated. The first closed trajectory is commonly referred to as “separatrix”, as it effectly separates open and closed electron orbits.

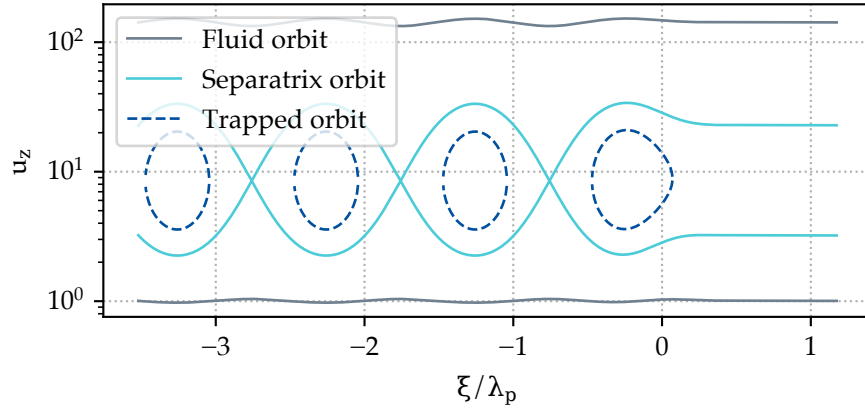


Figure 2.6.: Electron trajectories in the linear regime for a matched pulse with $a_0 = 0.3$ propagating from left to right.

An electron in front of the laser pulse ($a_0 = \phi = u_\perp = 0$) will be trapped if its momentum is higher than the longitudinal momentum of the electrons describing the separatrix $u_{z_{sep}}(+\infty) = \beta_p \gamma_p^2 H_{sep} - \gamma_p \sqrt{\gamma_p^2 H_{sep}^2 - 1}$. Energy-wise, this translates in a trapping thresh-

¹⁰The conservation of the canonical momentum for an electron in a laser field implies $u_\perp = a$ and, therefore $\gamma_\perp = (1 + a^2)^{\frac{1}{2}}$ [33].

old [34]:

$$E_{\text{trap}} = m_e c^2 \left(\sqrt{1 + (u_{z\text{sep}}(+\infty))^2} - 1 \right), \quad (2.28)$$

which is easier to achieve for higher amplitude plasma waves and for slower wake phase velocities (higher plasma densities)¹¹, as shown in figure 2.7.




Figure 2.7.: Dependence of the trapping threshold on the laser field strength a_0 and on the plasma density n_0 . Matched (longitudinal) nonlinear propagation is assumed.

2.3.2. Injection in the wakefield

While trapped electrons can gain energy from the plasma wave and be accelerated, the position and instant at which these electrons are initially “injected” in the wave are extremely important, as they largely define the properties of the resulting electron beam. Localising the injection point in both time and space, for instance, ensures that the particles start the acceleration process simultaneously, experiencing nearly identical fields. This, in turn, would result in the production of a monoenergetic electron beam. Nowadays a considerable amount of injection methods exist [35, 36, 37, 38, 39]. Among these “ionisation injection” [33] and “density-downramp injection” [40] have emerged, thanks to their ability to finely tune the injection position and their simple experimental implementation.

Ionisation injection The use of gases with high atomic number (Z) can be advantageous when injecting electrons in a plasma wave. In a high- Z gas (e.g. nitrogen) the electrons of the outer shells have a relatively low ionisation potential. During the laser propagation these electrons will be ionised at lower intensities ($\sim 10^{16}$ W/cm²), they will follow the fluid trajectories and form the plasma wave. The inner shell electrons, on the other hand, can only be ionised by high laser intensities ($\sim 10^{18}$ W/cm²). Therefore, these will be born inside the plasma wave, in regions of strong wakefields, with phase depending on the ionisation instant. Under certain conditions they can follow a trapped orbit and, therefore, be accelerated.

¹¹Note that for high a_0 the trapping threshold tends to 0: this describes the onset of wavebreaking.

In a 1D Hamiltonian derivation, the change in the (normalised) transverse momentum of an electron ionised inside the laser pulse with respect to the wave phase will follow the one of the vector potential [33]:

$$\frac{\partial}{\partial \xi} u_{\perp} = \frac{\partial}{\partial \xi} a_{\perp}. \quad (2.29)$$

Assuming that electrons are born at rest at the position $\xi = \xi_i$ (and so $u_{\perp}(\xi_i) = 0$), their transverse momentum will vary as:

$$u_{\perp}(\xi) = a_{\perp}(\xi) - a_{\perp}(\xi_i), \quad (2.30)$$

where $a_{\perp}(\xi_i)$ is the vector potential at the ionisation instant. After the passage of the laser, $a_{\perp}(\xi) = 0$ and the electrons will keep a residual transverse momentum which depends on their ionisation instant $u_{\perp}(\xi) = -a_{\perp}(\xi_i)$ ¹². Following this path, the trapping condition for the ionised electrons can be written as [33]:

$$1 + \phi_{\min} - \phi_i < \frac{(1 + a_{\perp}^2(\xi_i))^{1/2}}{\gamma_p}, \quad (2.31)$$

where ϕ_{\min} is the minimum electrostatic potential of the plasma wave (which is directly related to the laser a_0 , cf. section 2.2.3) and ϕ_i is the potential at the ionisation position. This trapping condition, hence, relates the phase at which the electron is ionised (ξ_i) with the laser a_0 .




Figure 2.8.: Threshold a_0 for ionisation injection (blue) with respect to ξ for a pulse centered on $\xi = 0$. Pulses with a_0 ranging from 1 to 4 are shown as examples. Electrons ionised in the shaded areas will fulfil the trapping condition and be injected in the plasma wave.

Figure 2.8 depicts (in blue) the a_0 needed to fulfil the ionisation injection condition as a function of ξ . Superimposing such condition on pulses of different intensity (red), one

¹²Note that for an electron ionised at the peak of the pulse $a_{\perp}(\xi_i) = 0$. The further away from the peak an electron is ionised, the bigger will be its residual momentum once the laser is passed.

could retrieve the region of a given pulse in which, when an ionisation event happens, the corresponding electron will be trapped in a closed orbit and injected in the plasma wave¹³. As the the laser peak corresponds to the maximum ionisation rate, the amount of electrons born and trapped thereafter, with $\xi_i < 0$, will be minimal. Evidently, figure 2.8 shows how for more intense pulses the interval over which the injection is possible is larger, leading to a higher trapped charge. On the other hand, though, the bigger is the region over which electrons can be injected, the more will be their relative phase difference. Therefore different sections of the same injected bunch will feel different fields, resulting in a higher energy spread and, generally, a beam of suboptimal quality [33]. An intuitive way around this is to use in the plasma source a mixture of heavy and light gases, localising as much as possible the one with high atomic number. This way the control of the injection region relies mainly on the plasma source itself rather than on the pulse intensity and a better quality beam can be obtained without excessively sacrificing the charge [33, 41].

Density downramp injection Another effective method to induce injection in a localised plasma region is to make use of density transitions [40, 42, 43]. During a downward transition (or “downramp”), for instance, the plasma wave elongates and its phase velocity decreases. Consequently electrons shift their position forward in phase with respect to the wakefield and the trapping threshold decreases. As a result, a fraction of them can have enough momentum to follow a closed orbit and be injected and accelerated in the wakefield.

In a density transition the wakefield potential cannot be written as only a function of ξ but it must be treated in two variables, $\phi(z, \xi)$. In order to understand the mechanism, hence, it is easier to follow a fluid model, where the phase of the plasma wave can be expressed as a function of the density gradient [34].

The phase velocity of the plasma wave can be computed from the wake phase¹⁴ [34, 44]:

$$\frac{v_p}{c} = -\frac{\frac{\partial \psi}{\partial c t}}{\frac{\partial \psi}{\partial z}} = \frac{1}{1 + \frac{\xi}{k_p} \frac{dk_p}{dz}}, \quad (2.32)$$

where $\psi = k_p \xi$ and $k_p = \omega_p/c$. The variation of the plasma wavenumber along z can be written as:

$$\frac{dk_p}{dz} = \frac{dk_p}{dn} \frac{dn}{dz} = \frac{k_p}{2n} \frac{dn}{dz}, \quad (2.33)$$

hence equation 2.32 becomes:

$$\frac{v_p}{c} = \frac{1}{1 + \frac{\xi}{2n} \frac{dn}{dz}}. \quad (2.34)$$

Behind the laser pulse $\xi < 0$. Therefore in a density downramp ($dn/dz < 0$) v_p decreases

¹³It is interesting to notice that (within the limits of this 1D model) a minimum a_0 of ~ 1.5 is needed for ionisation injection, meaning that this method cannot be applied in the linear regime (cf. section 2.8).

¹⁴The phase velocity can be retrieved comparing the phase change with respect to $c\tau$ to the phase change in z (the minus sign is because of the wake frame ξ).

and the trapping threshold is lowered¹⁵. Additionally, this decrease in phase velocity depends on ξ and, as a result, injection gets easier further behind the laser pulse. This generally means that injection in a downramp is less challenging to achieve, especially for low laser intensities, when the trapping condition can be difficult to fulfill. Moreover, as the electrons get injected along the density transition, a shorter ramp will result in a more localised injection.

It is worth mentioning that the same process can be seen from the injected electron's point of view. Within a downramp electrons get pushed forward in phase: those which are on (or close to) the separatrix can then end up being on a trapped trajectory. Considering a density downramp starting from an initial position i and extending up to a final position f , the change in the electron phase due to the ramp is

$$\begin{aligned}\Delta\psi &= \psi_i - \psi_f = k_{p_i}\xi - k_{p_f}\xi = \\ &= \psi_i \left[1 - \left(\frac{n_f}{n_i} \right)^2 \right] \simeq \psi_i \frac{\Delta n}{2n_i},\end{aligned}\tag{2.35}$$

assuming $\Delta n = n_i - n_f \ll n_i$. This forward push is shown in figure 2.9, where an electron bunch (in pink) is injected and trapped in the first plasma cavity.




Figure 2.9.: Example of a particle in cell (FBPIC) simulation illustrating downramp injection. The top panel shows the normalised longitudinal density profile, while the bottom zooms in the density ramp region. Electrons in pink are pushed forward in phase and get injected in the first cavity of the plasma wave.

¹⁵Density downramp injection fits in a group of “wavebreaking injections”, where electrons are injected from the plasma wave itself, once the wavebreaking threshold is overcome. Such threshold is generally a function of γ_{\perp}/γ_p (cf. appendix A), therefore the more v_p decreases the more electrons can be trapped.

Ionisation-aided downramp injection An injection method does not preclude the others. A combination of techniques can often be a viable solution to reach the desired beam parameters, as in [45], where the presence of a localised region of nitrogen is used both to build a density downramp and to induce ionisation injection.

SECTION 2.4

Nonlinear plasma optics

Thus far, the influence of laser propagation on the plasma dynamics has been analysed, along with its capacity to excite plasma waves, and its potential to trigger the injection and acceleration of plasma electrons. Conversely, the plasma exerts a substantial impact on the evolution of the laser pulse, which can lead to significant changes in its spectral composition, duration, and intensity.

The change of the refractive index of the plasma during the interaction determines the evolution of the laser pulse. In an unmagnetised plasma, the index of refraction (considering linearly polarised light wave) can be expressed as:

$$\eta = \left(1 - \frac{\omega_p^2}{\gamma \omega_L^2} \right)^{1/2}, \quad (2.36)$$

where ω_L is the laser central frequency and the Lorentz factor associated to the quiver motion of single electrons $\gamma_{\perp} = (1 + a_0^2/2)^{1/2}$ has been written simply as γ to simplify the notation. From equation 2.36 it is evident that the index of refraction of the plasma can change due to:

- ◇ the plasma density
- ◇ the laser amplitude
- ◇ the laser frequency

Such changes in η can give rise to a great variety of linear and nonlinear phenomena and highly affect the laser propagation both in the longitudinal and transverse direction.

2.4.1. Ionisation-induced defocusing

When the driver laser is responsible for the generation of plasma, the gas ionisation rate can vary substantially in the transverse direction and it can highly affect the pulse propagation. An ideal gaussian pulse, for instance, ionises the species on its propagation axis before the ones on its wings. This could ultimately result in a reduction of the refractive index on axis and, therefore, a defocusing gradient. Such effect, commonly known as “ionisation-induced defocusing”, when relevant, can limit the peak intensity reached by the laser pulse and be detrimental for the laser plasma interaction.

Generally, a change in the spot size of a laser beam can always be traced back to its wavefront curvature. In order for a wavefront to curve in a medium (and hence for a laser to get focused or defocused), the pulse phase velocity should not be constant along the beam transverse direction. The acceleration of the spot size w due to transverse variations of the phase velocity can be expressed as [46]:

$$\frac{\partial^2 w}{\partial \tau^2} = -c \left(\frac{v_{\phi_2} - v_{\phi_1}}{w} \right), \quad (2.37)$$

where v_{ϕ_1} is the phase velocity on the optical axis, v_{ϕ_2} is the corresponding velocity far from the optical axis and $v_{\phi_2} \neq v_{\phi_1}$. If $v_{\phi_2} > v_{\phi_1}$ the beam will be focusing, conversely if $v_{\phi_2} < v_{\phi_1}$ defocusing will take place. As $v_{\phi} \equiv c/\eta$, a refractive index change in the transverse direction is likely to cause one of the two effects.

During the propagation the central (and more intense) part of the laser starts ionising the gas. The refractive index of the resulting plasma, assuming $\gamma = 1$ and an electron density $n_e \ll n_c$ where n_c is the critical density, can be expanded as:

$$\eta_1 = \left(1 - \frac{\omega_p^2}{\gamma \omega_L^2} \right)^{1/2} = \left(1 - \frac{n_e}{n_c} \right)^{1/2} \simeq 1 - \frac{n_e}{2n_c}. \quad (2.38)$$

At the same time, the pulse wings are still propagating in a neutral gas ($\eta_2 \simeq 1$). As a result $v_{\phi_1} > v_{\phi_2}$ and the plasma acts as a defocusing lens.

A simple analytical estimate can be instrumental to predict when this effect can be relevant during the pulse propagation. The phase changes over a distance L for the two considered points can be expressed as:

$$\Delta\varphi = \varphi_2 - \varphi_1 = \frac{\omega_L L}{c} (\eta_2 - \eta_1) = \frac{\omega_L L}{2c n_c} n_e. \quad (2.39)$$

A phase change of $\pi/2$ corresponds to doubling the diffraction limited beam divergence. The corresponding length, hence, can be defined as “defocusing length” [47, 48, 49]:

$$L_D = \frac{\pi c n_c}{\omega_L n_e} = \frac{\lambda_L n_c}{2n_e}, \quad (2.40)$$

where λ_L is the central laser wavelength. Therefore, ionisation defocusing will play a non-negligible role when L_D is shorter than the characteristic length of the focusing apparatus, the Rayleigh length Z_R ¹⁶:

$$n_e > \frac{\lambda_L n_c}{2Z_R}. \quad (2.41)$$

As this effect depends on the refractive index gradient, higher atomic number gases (where multiple levels can be ionised) can have a bigger impact. Figure 2.10, for instance, analyses

¹⁶Note that this is just an estimate for a single ionisation level. In reality this is not a threshold effect and its impact of the pulse peak intensity is difficult to evaluate analytically, especially when other effects (as relativistic self-focusing, cf. section 2.4.2) become relevant.




Figure 2.10.: An ideal gaussian pulse (0.5 mJ, 15 fs) focusing into a constant atomic nitrogen density ($2 \times 10^{19} \text{ cm}^{-3}$), simulated with the particle-in-cell code FBPIC. The top panels show the generated electron density and the pulse profile at the entrance of the gas. Two pulse sections (front: dash-dot line - centre: solid line) are highlighted in red. The radially-dependent refractive index η experienced in the two sections is plotted in the bottom left panel. As a result, the front of the pulse undergoes a strong ionisation-induced defocusing, as shown in the last panel.

the propagation of a 0.5 mJ, 15 fs gaussian pulse focusing in constant nitrogen atomic density of $n_g = 2 \times 10^{19} \text{ cm}^{-3}$. The (vacuum) spot size position is set at $z = 0$ ¹⁷. The top panels show the electron density generated at the entrance of the gas and the correspondent pulse shape. The ionisation level varies across different pulse transverse sections, giving rise to different radially-dependent refractive indexes. Two particular sections (highlighted in red in the top panels) are compared in the bottom left plot. The front of the pulse (dash-dot line) experiences a refractive gradient all along its transverse size and, therefore, a strong defocusing effect. On the other hand in the centre of the pulse (solid line) ionisation has already took place. As a result the central and most intense part of the pulse propagates in a flat η and keeps focusing during the propagation. The last panel, indeed, depicts the same pulse after 20 μm . Ionisation-induced defocusing is visible at the pulse front: it effectively reduces the energy contained in the focal spot and, consequently, the peak a_0 achieved during the propagation.

2.4.2. Relativistic self-focusing

Once the plasma has been formed, the index of refraction is still susceptible to modification. When only small modulations and weakly relativistic drivers are considered, equation 2.36 can be expanded and the influence of the different terms can be analysed [46]:

$$\eta = 1 - \frac{\omega_p^2}{2\omega_L^2} \left(1 + \frac{\delta n}{n} - \frac{a^2}{2} - 2 \frac{\delta \omega_L}{\omega_L} \right), \quad (2.42)$$

where δn is the density variation and $\delta \omega_L$ is the laser bandwidth. Specifically, in the transverse direction, the relativistic term in equation 2.42 ($a^2/2$) can give rise to “relativistic self-focusing” and counteract (or even dominate over) the ionisation defocusing effect.

Equation 2.37 can be rewritten in terms of the refractive index, considering only the relativistic term in its expansion. The resulting equation expresses how the spotsize of the laser w evolves due to the change of the refractive index in the transverse direction, after the plasma is formed and when the change is mainly due to the laser intensity (a_0):

$$\frac{\partial^2 w}{\partial \tau^2} = -\frac{1}{8} \frac{\omega_p^2 c^2 a_0^2}{\omega_L^2 w}. \quad (2.43)$$

During the laser propagation, though, diffraction takes place as well and the laser spot size evolves as

$$w = w_0 \sqrt{1 - \frac{\tau^2}{t_R^2}}, \quad (2.44)$$

where $t_R = \frac{k_0}{2c} w_0^2$ is the Rayleigh time and w_0 is the focal spotsize. Close to the focus nearly planar wavefronts can be assumed and the above equation can be differentiated twice to

¹⁷The parameters have been chosen in order to exclude other transverse nonlinear effects while still have sufficient peak intensity to ionise five nitrogen levels.

get [46]:

$$\frac{\partial^2 w}{\partial \tau^2} \simeq \frac{4c^2}{k_0^2 w_0^3}. \quad (2.45)$$

Adding this contribution to equation 2.43, the evolution of the laser spotsize can be described as:

$$\frac{\partial^2 w}{\partial \tau^2} = \frac{4c^2}{k_0^2 w_0^3} \left(1 - \frac{1}{32} \frac{\omega_p^2}{c^2} a_0^2 w_0^2 \right). \quad (2.46)$$

Self-focusing occurs if the term in the brackets is negative, i.e. if the change in the spotsize due to the transverse change of the v_ϕ overcomes the diffraction term, so when:

$$\frac{1}{32} \frac{\omega_p^2}{c^2} a_0^2 w_0^2 > 1. \quad (2.47)$$

The term $a_0^2 w_0^2$, moreover, is proportional to the laser power. The same condition can then be expressed as a power threshold, leading to the well-known definition of “critical power” for self-focusing [50]:

$$P_{cr} = \frac{8\pi\epsilon_0 m_e^2 c^5 \omega_L^2}{e^2 \omega_p^2} \simeq 17 \frac{\omega_L^2}{\omega_p^2} [GW] = 17 \frac{n_c}{n_e} [GW]. \quad (2.48)$$

2.4.3. Photon acceleration

On the other hand, when looking at the change of the refractive index along the laser propagation direction, phenomena that alter the pulse spectrum are particularly interesting, as they can provide insights into the energy transfer in the interaction, the laser pulse compression and more.

A change in the longitudinal refractive index of the medium in a given region will result in a different local phase velocity, an effect often referred to as “photon acceleration”¹⁸. Following the argument of [46], the rate of change of the local frequency of the pulse and its dependence on the variation of the refractive index in the “wave frame” (ξ) can be written as:

$$\frac{1}{\omega_L} \frac{\partial \omega_L}{\partial \tau} = - \frac{1}{\eta^2} \frac{\partial \eta}{\partial \xi}, \quad (2.49)$$

where $\xi = ct - z$ and $\tau = t$.

Therefore, each of the terms in equation 2.42 has the potential to modify the local frequency of the pulse¹⁹.

¹⁸Note: in the following we are assuming the η to be stationary along the driver during the propagation.

¹⁹The term related to the laser bandwidth ($\delta\omega_L$) in equation 2.42 often gives rise to a much smaller contribution with respect of the others and can be neglected. However, this is not entirely true when few-cycle pulses are considered, as in high average power LPAs [51].

Density modulations Density modulations are always present. Equation 2.42 shows that the change of the refractive index in the wake frame is opposite with respect to the change in the density modulation:

$$\frac{\partial \eta}{\partial \xi} \sim -\frac{\partial(\frac{\delta n}{n})}{\partial \xi}. \quad (2.50)$$

The pulse local frequency will then change as:

$$\frac{\partial \omega_L}{\partial \tau} \sim \frac{\partial(\frac{\delta n}{n})}{\partial \xi}, \quad (2.51)$$

therefore, it will increase (“blue-shift”) for the portion of the laser pulse placed in a region where $\frac{\partial(\frac{\delta n}{n})}{\partial \xi} > 0$, while decrease (“red-shift”) where $\frac{\partial(\frac{\delta n}{n})}{\partial \xi} < 0$.

Laser field strength The term of the refractive index which is related to the laser field strength ($a_0/2$) acquires increasing importance the more the interaction becomes relativistic. It is negligible in the linear regime, while it becomes dominant the ultra-relativistic one (or “complete blowout”). From equation 2.42:

$$\frac{\partial \eta}{\partial \xi} \sim \frac{\partial a^2}{\partial \xi}, \quad (2.52)$$

and hence

$$\frac{\partial \omega_L}{\partial \tau} \sim -\frac{\partial a^2}{\partial \xi}. \quad (2.53)$$

This, physically, is due to the effect that the relativistic mass increase of the electrons has on the refractive index. The higher is a , the faster are the electrons in their quiver motion and the more relevant this term is. Therefore, a more truthful dependence here is on the electron’s γ -factor:

$$\frac{\partial \eta}{\partial \xi} \sim \frac{\partial \gamma}{\partial \xi} \rightarrow \frac{\partial \omega_L}{\partial \tau} \sim -\frac{\partial \gamma}{\partial \xi}. \quad (2.54)$$

Hence the portion of the laser pulse placed in a region where $\frac{\partial \gamma}{\partial \xi} > 0$ will red-shift, while the section where $\frac{\partial \gamma}{\partial \xi} < 0$ will blue-shift.

Figure 2.11 adopts a 1D model (cf. 2.2.3) to show how these two contributions add up in the case of a matched pulse propagating in an underdense plasma with $a_0 = 2$ ²⁰.

The top panel shows the different contributions coming from the $\frac{\partial \gamma}{\partial \xi}$ and the $\frac{\partial(\frac{\delta n}{n})}{\partial \xi}$ terms. The bottom depicts their combined influence in $\frac{\partial \eta}{\partial \xi}$. The regions where $\frac{\partial \gamma}{\partial \xi} > 0$ will result in red shift, while the regions where $\frac{\partial \gamma}{\partial \xi} < 0$ will give rise to a blue-shift. The opposite happens with the plasma density oscillation : when $\frac{\partial(\frac{\delta n}{n})}{\partial \xi} > 0$ a blue-shift contribution rises, whereas $\frac{\partial(\frac{\delta n}{n})}{\partial \xi} < 0$ corresponds to a redshift.

²⁰The model used to describe photon acceleration relies on a linear expansion of the refractive index η and, therefore, is not accurate to describe nonlinear propagation ($a_0 > 1$). Nonetheless, it gives a very good qualitative understanding of the physical picture.




Figure 2.11.: Overview of the contributions to the photon acceleration (in 1D) considering a matched pulse with $a_0 = 2$. The top panel shows the different contributions coming from the $\partial\gamma/\partial\xi$ (orange) and the $\partial(\frac{\delta n}{n})/\partial\xi$ (blue) terms along the laser envelope (red). The bottom illustrates their superposition in $\partial\eta/\partial\xi$ (green). Shaded areas represent the local frequency shift.

- ◇ At the very front of the pulse ($\xi \gtrsim 2.4$) the change in electron density would produce a blueshift while the change in γ causes a redshift. As a result the two contributions cancel out and $\frac{\partial\eta}{\partial\xi} \sim 0$.
- ◇ Going towards the center of the pulse, for $a_0 > 1$, the growth rate of the γ -factor typically prevails, leading to $\frac{\partial\eta}{\partial\xi} > 0$ and, thereby, a local redshift.
- ◇ Around the pulse center both contributions result in a redshift ($\frac{\partial\gamma}{\partial\xi} > 0$ and $\frac{\partial(\frac{\delta n}{n})}{\partial\xi} < 0$). Consequently the center of the laser pulse, where most of the intensity resides, shifts towards lower frequencies. This usually brings the average shift of the whole pulse towards the red side of the spectrum after the interaction.
- ◇ At the back of the pulse the two terms are again opposite. In case of matched longitudinal propagation with $a_0 > 0$, the γ -factor prevails and the pulse is shifted towards the blue.

It is important to point out here that while the γ variation is due to the pulse itself and follows




Figure 2.12.: PIC simulation showing the evolution of an ideal gaussian pulse with $a_0 = 2$ propagating in a matched hydrogen density. For each panel the spectral content (vertical axis) of the pulse and its temporal envelope (horizontal axis) are depicted through Wigner plots, where W is the pulse chronocyclic intensity (cf. appendix B). The first panel represents the beginning of the simulation, while the second one shows the pulse after $\sim 130 \mu\text{m}$.

it, the plasma wavelength λ_p (and so the characteristic length of the density oscillations) is not directly related to the pulse and depends mostly on the plasma density. This means that, for a non matched propagation, the contribution to the photon acceleration related to the plasma density variation can be very different. A longer pulse propagating in the same plasma density, for instance, would experience at its back the electron density peak of the back of the bubble regime. That steep up-density gradient would lead to a much stronger blue-shift at the back of the pulse [52].

To further validate the model, figure 2.12 represents the result of a PIC simulation (FBPIC) where an ideal gaussian pulse with $a_0 = 2$ propagates in a hydrogen plasma with a matched plasma density. The spectral and temporal evolution of the pulse after $\sim 130 \mu\text{m}$ of propagation is shown with Wigner plots²¹. After the propagation the very front of the pulse (negative times) does not experience any photon acceleration. Going towards the pulse center a

²¹Wigner plots are an extremely useful tool to visualise the pulse evolution [53] (cf. appendix B)

prominent red-shift is visible, pulling the overall mean pulse frequency towards lower values. The opposite shift, on the other hand, is evident on the trailing edge (positive times).

It is interesting to note that, as a result of the nonlinear interaction, the temporal structure of the pulse evolves as well. In fact, while the pulse spectrum gets broadened, with the pulse front redshifting and its trailing edge blueshifting, plasmas generally have an anomalous dispersion relation, where $\frac{\partial v_g}{\partial \lambda} < 0$ or equivalently $\frac{d\eta}{d\omega} < 0$. Therefore, the blue frequencies of the pulse travel faster than the red components. Depending on the initial pulse duration and spectral phase, the dispersion can compensate the positive chirp formed by the nonlinear interaction and compress the pulse during the propagation, even beyond its initial transform limited duration [51, 54, 55]²².

Ionisation-induced blueshift

Describing the different aspects of photon acceleration in this section, we have implicitly assumed that the plasma was already formed. However, when gas ionisation is taken into account, also the presence of an ionisation front can affect the pulse and modify its spectrum [56].

Recalling equation 2.51, the change in the local frequency of the pulse is proportional to the change in the local plasma density as:

$$\frac{\partial \omega_L}{\partial \tau} \sim -\frac{\partial \eta}{\partial \xi} \sim \frac{\partial n_e}{\partial \xi}, \quad (2.55)$$

where n_e is the (electron) plasma density. Along the propagation axis the pulse front, ionising the gas, experiences an increase in plasma density. Consequently, its local frequency content gets shifted towards higher energies, a phenomenon commonly known as “ionisation-induced blueshift”. The relative importance of this frequency shift on the laser propagation generally depends on the pulse peak intensity. In many cases such intensity is orders of magnitude higher than the ionisation threshold of the chosen gas, therefore the up-shift is only experienced by the very front edge of the laser and its effects can often be neglected. However, in the case of high-average-power LPA the pulse energy is more contained and ionisation-induced blueshift can be relevant (cf. section 6.4.2).

SECTION 2.5

High-average-power laser plasma accelerators

Having laid the groundwork for the general understanding of laser plasma accelerators, the emphasis is now shifted on the most notable state-of-the-art advancements in this ever evolving field, particularly focusing on high average power.

High quality electron beams in the 100 MeV range were already produced in plasma accelerators nearly two decades ago [6, 7, 8]. The GeV scale was reached soon after [9], paving the way towards the multi-GeV regime [10, 11, 12, 13]. Nowadays laser-plasma acceleration

²²The interplay of spectral broadening and pulse dispersion leading to temporal compression is more extensively treated in chapter 3.

has branched out in multiple directions, addressing all the challenges that could potentially limit their practical use [57, 58, 59, 60, 61]. Among these, increasing the repetition rate of the machine is essential, in order to drive applications where many pulses per second and high average electron current are needed [14].

Most LPA experiments currently rely on lasers that produce pulses with peak power in the tens of terawatts. Nevertheless, these systems are restricted to few-hertz operation, which can be a significant impediment when envisioning the broader application of this technology beyond the realm of research. To facilitate such expansion, approaching the kilohertz operational frequency is highly desirable:

- ◇ at these repetition rates lasers are usually more stable, often resulting in an enhanced electron beam stability,
- ◇ a higher average electron current can be generated, which is a key parameter for many applications and
- ◇ a higher repetition rate gives the chance to have active feedback control loops, which can be particularly relevant when aiming at a less operator-dependent machine.

Given the present laser technology, though, increasing the repetition rate inevitably results in a reduced pulse energy and systems operating at kilohertz typically produce pulses in the millijoule range. This fundamentally limits the maximum energy gain of the accelerated electron beam up to ~ 10 -20 MeV. However in many applications high energy is not strictly required and one can generally benefit from lower energies combined with improved stability and signal-to-noise ratio. Yet employing few mJ pulses as drivers is not trivial and introduces its own set of complexities.

2.5.1. Scaling of the blowout regime

The majority of plasma-based accelerators are operated in the blowout regime, where the wakefields are particularly suitable for electron acceleration (cf. 2.2.3). When the maximum pulse energy is limited to only few mJ this regime is still accessible, provided that all the other parameters are correctly scaled [30]. As an example, table 2.1 outlines a comparison between the characteristics of an LPA driven by 10 millijoule pulses, operating at a kilohertz, and a more conventional Ti:Sa laser²³ that delivers pulses at the joule level but with a lower frequency. Here are the main differences:

- ◇ The pulse duration (τ). In order to drive a nonlinear wave, high peak-power is a key parameter, calling for pulse durations of only few femtoseconds, in the few-cycle regime.
- ◇ The plasma densities (n_e). Matched propagation (cf. 2.2.3) with such short pulses needs very high plasma densities, on the order of 10^{20} electrons/cm³.

²³Titanium-sapphire laser, the most used to date. The name refers to the lasing medium: a crystal of sapphire doped with Ti³⁺ ions.

- ◇ The energy gain (ΔE). The tight focusing needed in order to achieve sufficient intensity leads to an even stronger defocusing. All the distances are reduced, meaning that the range over which electrons can actually be accelerated is much less, as well as their final energy.

Rep. rate	E [J]	τ [fs]	n_e [cm^{-3}]	ΔE [MeV]
Hz	1	30	5×10^{18}	500
kHz	0.01	10	1×10^{20}	10

Table 2.1.: Parameters needed in order to achieve the blowout regime with a Joule-level, Hz, laser compared to a 10 mJ-level, kHz system. In the order: pulse energy (E), pulse temporal duration (τ), plasma density (n_e) and energy gain (ΔE).

Nonetheless, despite these challenges, high-average-power LPA, operating at kHz, has emerged in recent years as a promising avenue for ultrafast dynamics research, cancer treatment, medical imaging, space radiation studies and several other applications. In the following, hence, a concise overview of the state-of-art of the technique is presented, highlighting the key achievements and paving the way for the upcoming chapters.

2.5.2. State-of-the-art of high-average-power LPA

Early experiments with high repetition rates predominantly emphasized the experimental demonstration, rather than the fine-tuning of the electron beam characteristics. In this context, density downramp injection (cf. 2.3.2) emerged as a practical solution, notably lowering the threshold for electron injection. Low charge (10-20 fC/shot) electron acceleration was consistently demonstrated, first at 500Hz [62] and then at 1 kHz [63], from sub-TW Ti:Sa laser injecting at the end of a subsonic gas jet. This approach, though, leaves only a limited amount of plasma available for the actual acceleration process, resulting in low energies (~ 100 KeV) and a thermal distribution.

The alternative path to inject electrons with low peak power pulses (< 1 TW) is to exploit self-focusing in plasma (cf. 2.4.2). Near critical plasma densities, for instance, were used to demonstrate for the first time acceleration to energies higher than 1 MeV (with a bunch charge of ~ 1 pC) from a 10 mJ, 30 fs Ti:Sa laser running at 1kHz [64]. The mismatch between the pulse length and the plasma wavelength, though, lead to operation in the self-modulated regime [65] and in a fairly divergent electron beam.

Shortly thereafter, the advent of laser post-compression allowed to compress the pulse temporal duration down to few fs, while preserving the laser energy. Pulses approached the TW regime, leading to a more reliable electron acceleration. In Paris, the Salle Noire Ti:Sa laser system, delivering 2-3 mJ pulses compressed down to 3.5 fs [66], allowed to reach matched propagation in the blowout regime with plasma densities on the order of 1×10^{20} cm^{-3} . Therefore, ionisation injection could be properly exploited to achieve electron energies as high as 9 MeV from a subsonic gas jet, with a charge of ~ 0.5 pC/shot and a relatively

small divergence of ~ 45 mrad fwhm [67].

Yet self-focusing in plasma plays a major role in the interaction, hence the more confined the plasma source is the better²⁴. Supersonic gas jets, in this sense, are a promising alternative, providing sharper density profiles and lower gas loads in the vacuum chamber. Employing them, a striking increase in bunch charge was obtained (>20 pC/shot) while still accelerating in the few MeV range, with a divergence < 60 - 70 mrad fwhm [68].

Following these proof of principle experiments, recent years have seen a transition towards more methodical studies, analysing the dependence of the electron beam properties on various interaction parameters, such as the laser pulse duration [69], the focal spot position with respect to the plasma density profile and the laser polarisation [16]. Moreover density downramp has been proven as one of the most controllable injection methods, providing long term stable injection in the few MeV regime with a shocked gas jet [69]. To date, charges as high as 24 pC/shot, energies extending to 15 MeV and beam divergence < 10 mrad have been separately achieved.

These parameters appear as an incredibly promising avenue for different industrial, commercial and medical applications where low energies but high electron currents are needed, especially when delivered in bunches as short as few-fs (imaging, radiotherapy, pump-probe studies, space radiation studies...). Despite the substantial progress that has been made, though, the realisation of a practical industrial application is contingent upon the availability of a highly stable and reliable laser driver, which could work over extended periods of times. Currently, despite the proven hour-long stability, the accelerated electron beams continue to exhibit noticeable day-to-day fluctuations, primarily attributable to the laser system's performances [69]. It is worth noting that all the successful experiments mentioned here made use of Ti:Sa laser drivers. These systems provide the highest pulse peak power off the shelf, however they come with a substantial price tag and the challenge related to heat management, driven by their significant quantum defect, poses severe limitations on their average power, repetition rate and day-to-day stability. Thus, it would be highly desirable to transition toward a more cost-effective solution, capable of operating in a stable and reliable way at multi-kHz repetition rates.

Building upon this idea, the next chapter will introduce a novel, industrial-quality driver for high-average-power LPA, holding the promise for a more practical table-top accelerator.

²⁴The more confined is the plasma, the less the laser pulse is sensible to ionisation induced defocusing

CHAPTER 3

VERA: an industrial laser for LPA

The previous chapter revealed the exceptional potential of LPAs, chiefly driven by their ability to sustain accelerating fields order of magnitudes higher than radiofrequency accelerators. However, looking towards industrial and medical applications, more efficient, reliable and scalable solutions are needed.

In this context, bulk, industrial Ytterbium:Yttrium-aluminium-garnet (Yb:YAG) lasers emerge as a compelling option: their quasi-three level scheme ensures a minimal quantum defect and the availability of high power pumping diodes in their absorption band (~ 940 nm) allows for very high slope efficiency¹. In recent years, moreover, novel gain-media designs have enabled high power scalability [17], up to the level where their output could now be exploited for multi-kHz, low energy electron acceleration.

The upcoming chapter, hence, is devoted to the introduction of the “VERA” Yb:YAG laser², the primary tool for all the experimental work presented in this thesis. Table 3.1 summarizes the laser parameters. It has the capability to operate at repetition rates as high as 20 kHz, delivering an average power of almost 190W. However, to manage the possible thermal effects on the optics, it is currently being operated at 1 kHz.

λ_0 [nm]	E [mJ]	τ [ps]	f [kHz]	W_{\max} [W]	M^2	σ_E [%]
1030	> 10	> 1.1	20	215	1.05×1	< 1
1030	9.45	1.2	1	190	1.22×1.39	0.2

Table 3.1.: Factory-acceptance-test specifications (first row) and parameters used in this work (second row) for the VERA laser, namely the central wavelength (λ_0), the pulse energy (E), the FWHM pulse duration (τ), the repetition rate (f), the maximum average power (W_{\max}), the beam quality (M^2) and the root mean square energy stability (σ_E).

¹Slope efficiency is a measure of the laser conversion efficiency, obtained plotting the laser output power against the input pump power.

²Amphos, model 3502.

Nonetheless, such high average power comes at the cost of a narrow spectral bandwidth. The system emits relatively long pulses and its output peak power is far from being suitable for driving an LPA³. Therefore, compressing the pulses to durations on the order of few femtoseconds, preserving their energy, is crucial in order to reach the required intensity. Starting from 1.2 picoseconds, this would mean obtaining a compression factor higher than 100, an ambitious and challenging task which can be achieved through recently pioneered post-compression techniques.

SECTION 3.1

Post-compression to few-cycle duration

In pulse compression the time-bandwidth product is pivotal: an ultrashort pulse has to be supported by a large bandwidth⁴. Therefore, a typical post-compression setup consists of two main steps [70,71]:

- ◇ a spectral broadening stage, where nonlinear interaction with a specific medium is employed to achieve the required bandwidth, and
- ◇ a dispersion compensation line, where the phase distortions introduced in the first step are compensated, so that the pulse would have the shortest possible temporal duration for the new, larger bandwidth⁵.

The alternative: OPCPA

In the early 90s, a different way to increase the pulse peak power was developed based on direct amplification of ultra-broadband pulses: the optical parametric chirped pulse amplification (OPCPA). It fuses the chirped pulse amplification (CPA) technique together with optical parametric amplification (OPA): a signal pulse is temporally stretched to match the pump duration, amplified in a suitable non-linear crystal and then recompressed. However, while these system reach great performances in terms of peak-power, their low energy conversion efficiency needs expensive and complex pump architectures and the output beam can often have poor spatial characteristics [72, 73, 74].

3.1.1. Spectral broadening

During the interaction of a high intensity pulse with a nonlinear medium, the pulse spectrum can be broadened by different physical phenomena [70,71]. In this context, the medium polarisation (P) is central:

$$P(t) = \epsilon_0 \chi E(t) = \epsilon_0 (\chi^{(1)} E(t) + \chi^{(2)} E^2(t) + \chi^{(3)} E^3(t) + \dots), \quad (3.1)$$

³At the laser output the peak power is of only ~ 8 GW. As a rule of thumb, pulses on the order of a TW are needed to efficiently drive a nonlinear plasma wave. A 10 mJ pulse, for instance, would need to be as short as 10 - 20 fs.

⁴An ideal gaussian pulse of 5 fs, for instance, requires a bandwidth of ~ 300 nm.

⁵cf. appendix C.1.

where $E(t)$ is the laser field, ϵ_0 is the vacuum permittivity, $\chi^{(n)}$ is the n th order of the medium susceptibility and a power series expansion has been applied. The third order $\chi^{(3)}$, in particular, can give rise to a nonlinear contribution oscillating at the same frequency of the applied field, ultimately leading to a nonlinear (and intensity-dependent) modification of the medium refractive index [75]:

$$\eta(I) = \eta_0 + \eta_2 I(r, t), \quad (3.2)$$

where η_0 is the usual, linear refractive index, I is the pulse intensity and η_2 is a nonlinear coefficient, which can be assumed to a good approximation as linearly dependent on the density of the medium⁶. This phenomenon, commonly known as optical Kerr effect, can have a profound impact on the laser propagation.

Namely, the temporal dependence of its intensity profile $I(t)$ gives rise to a phase shift⁷:

$$\varphi = \omega_0 t - kz = \omega_0 t - \frac{\omega_0}{c} [\eta_0 + \eta_2 I(t)] z, \quad (3.3)$$

where ω_0 is the laser central frequency, k is its wavenumber and z is the direction of propagation. The shift, in turns, modifies the pulse frequency spectrum as:

$$\omega(t) = \frac{\partial \varphi}{\partial t} = \omega_0 - \frac{\omega_0 \eta_2 z}{c} \frac{\partial I}{\partial t}, \quad (3.4)$$

where $\omega(t)$ is the time-dependent carrier frequency (or the “instantaneous frequency”). The $\frac{\partial I}{\partial t}$ factor is responsible for the “self-phase modulation” (SPM) of the pulse, for which new frequencies are generated wherever intensity gradients are present [71]. The leading edge of the pulse shifts towards lower frequency values (redshift), while the trailing towards higher (blueshift): the spectrum gets, then, broadened⁸. Bell-shaped intensity profiles, as the gaussian depicted in figure 3.1, results in a linear time-dependence of the instantaneous frequency in the most intense part of the pulse. This is namely a positive linear chirp, which can then be compensated through a strategic use of chromatic dispersion.

3.1.2. Dispersion compensation

As described in the previous section, during the broadening stage the interaction produces spectral phase distortions and, as a result, the pulse instantaneous frequency becomes time-dependent. Consequently, even with a significantly wider spectrum, the pulse is still not temporally compressed. To achieve the actual compression, a dispersion compensation system is needed, where a chirp opposite to the one induced before is applied. The spectral phase would then be flattened again and, finally, the transform-limited duration corresponding to the new bandwidth could be obtained (cf. appendix C).

⁶ η_2 is directly related to the third order susceptibility as $\chi^{(3)} = 4\epsilon_0 c \eta_2 \eta_0^2 / 3$ [75].

⁷Note that the intensity profile has a spatial dependence as well ($I(r)$) which can lead to self-focusing effects [76].

⁸A remarkably similar effect on the laser pulse is described in section 2.4.3, discussing the influence of the laser field strength on photon acceleration in plasma. For this reason, that specific contribution to photon acceleration is often addressed as “relativistic SPM”.




Figure 3.1.: SPM-induced broadening example: a 9.45 mJ, 1.2 ps pulse (cf. 3.2.1) propagating in 34 m of Argon at 1.2 bar, assuming a constant spot size of ~ 2.1 mm radius. The change in the instantaneous frequency ($\omega - \omega_0$) is plotted in red along the normalised pulse intensity profile (in grey).

Dispersion compensation is usually achieved by altering the optical path of each wavelength to add a required wavelength-dependent delay. Nowadays, almost all pulses from ultrabroadband systems are compressed with chirped mirrors [77], reflective optics composed of alternating thin layers of two different materials, as schematically shown in figure 3.2. Depending on the optical path length of the layers, each boundary causes the (partial) reflection of a specific wavelength. Therefore, increasing the layer thickness with the penetration depth, a frequency-dependent group delay can be introduced by reflecting different spectral components of the incident radiation at different positions inside the structure. Being able to sustain high intensities, such mirrors provide ultra-broadband dispersion control with minimal losses [78].




Figure 3.2.: Sketch of a chirped mirror. Longer wavelengths (red) follow a longer optical path, inducing a negative group delay dispersion on the pulse.

Within the VERA project, the combined use of SPM and dispersion compensation allowed to reach a record compression factor of ~ 120 , reducing the laser pulse duration from 1.2 ps

to below 10 fs [79].

SECTION 3.2

The VERA laser

In a post-compression setup the spectral broadening stage plays a fundamental role. It is essential for it to happen gradually in time and as uniformly as possible along the beam, so that the accumulated phase can be more easily compensated. Self-phase-modulation (SPM) in gaseous media arises as the optimal choice: gases, in fact, allow for a continuous broadening along the beam path and provide an easy way to tune the strength of the nonlinear interaction⁹.

Nowadays, the great majority of post-compression schemes aiming at high peak power in the few-femtoseconds range make use of Ti:Sa lasers and exploit SPM in gas-filled hollow-core-fibers (HCF) [80]. HCF can sustain relatively large broadening factors and produce an excellent beam quality [81] but, on the other hand, they are extremely prone to damage and often result in very high energy losses. Furthermore, the complexity of the fiber design makes it challenging to produce several-meters-long straight fibers, limiting the length over which the nonlinear interaction can be sustained and, thereby, the achievable compression factor and power scalability. Thus, when dealing with low initial laser energy and long pulse duration, as it is often the case with industrial Yb-based lasers,¹⁰ a less fragile setup which could be folded would be highly advantageous. To this end, multipass-cells (MPCs) are the ideal tool [82].

In the following, the working principle of a MPC will be briefly presented, particularly focusing on the scheme adopted in the VERA project.

3.2.1. A two-stage cascaded MPC setup

Similar to an optical cavity, a Herriot-type multi-pass cell (MPC) consists of two concave mirrors, between which a laser beam is reflected multiple times [83]. Each pass imprints a small nonlinear phase modulation on the pulse allowing, overall, for large spectral broadening. If mode-matched with the transverse eigenmodes of the cell, the pulse preserves its beam quality and, moreover, its repetitive refocusing through the cell's mirrors redistributes the spatial beam content, enabling high spatio-spectral homogeneity [82]. With MPC technology, in the working frame of VERA, the output of an industrial Yb:YAG laser ($\lambda_0 = 1030$ nm) delivering 9.45 mJ, 1.2 ps pulses at multi-kHz (cf. table 3.1) is compressed down to sub-10 fs duration, with an overall efficiency above 70% [79].

When such large compression factors are needed, a multi-stage, cascaded post compression configuration is to be preferred, as it alleviates the need for a prolonged nonlinear interaction, which would ultimately compromise the temporal quality of the compressed

⁹the nonlinear refractive index of a medium (η_2) depends on the medium density and hence, in the case of gases, on their pressure.

¹⁰as the VERA laser, used in this work.

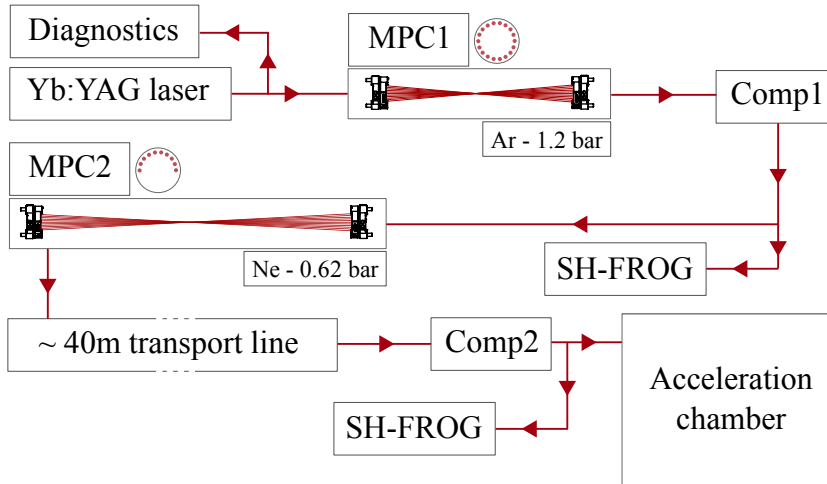


Figure 3.3.: Scheme of the experimental post-compression setup. Two gas-filled MPC stages are operated in series. Each MPC is followed by a compressor (Comp1 and Comp2), after which the pulse is characterised with a SH-FROG. The beam pattern obtained at the MPCs mirrors is indicated schematically for both cells.

pulse in a more common single-stage approach [84]. Fig. 3.3 shows an outline of the experimental, two-stage VERA laser setup. In the first stage the laser beam goes through a 2 m-long MPC filled with Argon at 1.2 bar. After seventeen round trips, its spectrum gets broadened to a transform limited FWHM duration of ~ 50 fs, with a peak energy throughput of $95.2\% \pm 0.2\%$. Then, to compensate for the SPM-induced chirp, the beam is sent through a single-pass transmission grating compressor (Comp1 in figure 3.3)¹¹.

A 3.45 m-long cell constitutes the second broadening stage. Here, ten round trips in 0.62 bar of Neon yield an output spectrum corresponding to a transform-limited FWHM duration of only 7 fs, with an exceptional throughput of $87\% \pm 0.3\%$. The pulse, then, is sent through a pair of fused silica wedges for fine dispersion control and along a ~ 40 m transport line, up to the radiation laboratory. There, it is compressed by four pairs of chirped mirrors, and sent to the experimental chamber for electron acceleration. Before the transport line, the beam can be diverted, compressed by a secondary dispersion compensation line (which resembles the main one) and measured by a second harmonic frequency-resolved optical gating (SH-FROG) device [86]. The measured pulse spectra and the retrieved pulse durations for each stage are shown in figure 3.4. The spatial quality of the beam is conserved and the peak power of the pulse has been increased by a factor 40 [79].

It is interesting to note here how the retrieved pulse after the second compression stage presents a distinct amount of pre and post pulses. In fact, the nonlinear interaction during the spectral broadening stage gives rise to high-order phase components that are extremely

¹¹Similarly to a chirped mirror (cf. 3.1.2), a grating compressor alters the optical path of different wavelengths to introduce negative dispersion [85]. A transmissive grating is a compact solution which can be used when the spectral bandwidth is of limited extent at the cost of a lower throughput (with respect to chirped mirrors). The compressor used in this first stage has a throughput of $\sim 96\%$.

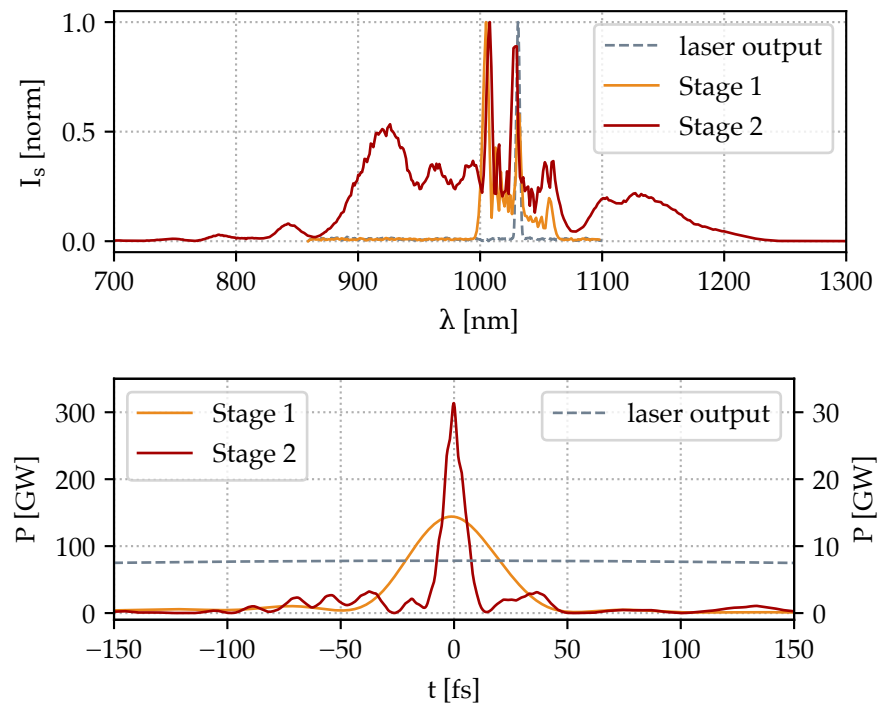


Figure 3.4.: Typical measured spectra (top) and FROG-retrieved pulse temporal profiles (bottom). The outputs of the first and second compression stages (measured before the transport line) are shown respectively in orange and red, compared to the initial laser parameters (grey).




Figure 3.5.: The temporal structure (left) and the spectral phase (right) of the measured pulse (red - M) are compared to their transform-limited counterpart.

challenging to perfectly compensate experimentally. These residual components can then have a non negligible impact on the resulting pulse, modifying its temporal structure. During the optimisation of the compression, this usually leads to a tradeoff between the shortest possible pulse duration and the highest possible intensity contrast [84]¹². To visualise their effect, figure 3.5 compares the measured pulse at the end of the VERA post-compression setup (red) with its corresponding ideal transform limited pulse (blue).

The post-compression results achieved are exceptionally noteworthy. The parameters obtained in each stage are summarised in table 3.2, while figure 3.6 locates the VERA laser within the vast post-compression landscape. It stands out as a strong contender to drive few-MeV, multi-kHz electron acceleration.

However, while the laser is pivotal in the acceleration process, the heart of a LPA lies in the plasma source, where particles are injected and accelerating fields are excited. Hence, the upcoming chapter will center its discussion on the most common sources for high-average-power LPA, highlighting the critical aspects to take into account when sub-TW pulses are considered as drivers.

	E_{tot} [mJ]	τ [fs]	P_{peak} [GW]	M^2
Laser output	9.45 ± 0.02	1200 ± 14	7.8	$1.22 \pm 0.12 \times 1.39 \pm 0.14$
Stage 1	9.00 ± 0.01	44.00 ± 0.26	144	$1.34 \pm 0.13 \times 1.27 \pm 0.13$
Stage 2	6.70 ± 0.02	9.60 ± 0.07	313	$1.11 \pm 0.11 \times 1.31 \pm 0.13$

Table 3.2.: Overview of the main laser parameters achieved in the cascaded MPC post-compression setup: the total pulse energy (E_{tot}), the FWHM pulse duration (τ), the peak power computed taking into account the FROG-retrieved pulse structure (P_{peak}) and the beam spatial quality (M^2).

¹²figure C.1 in appendix C, for instance, shows the effect of a small residual third order component on a 10 fs transform limited pulse




Figure 3.6.: Overview of the major results obtained with the most common post-compression methods: solid MPC (blue), gaseous MPC (orange) and HCF (green), adapted from [71] and [82]. The compressed energy (E_{out}) is plotted against the compressed FWHM pulse duration (τ_{out}). The output parameters of the compressed pulse obtained in two VERA stages are highlighted in red.

CHAPTER 4

Plasma sources for high-average-power LPA

In a laser plasma accelerator, the acceleration mechanism heavily relies on the region where the plasma resides, commonly referred to as “plasma source”. Its geometry and its characteristics play a key role, shaping several parameters of the resulting electron beam: the energy, the energy spread, the divergence and more. In the following, a detailed overview of plasma sources for high-average-power LPA will be given, from a fluid dynamics perspective. The prevalent choices in the field will be assessed and compared using the computational fluid dynamics (CFD) code ANSYS-FLUENT and a novel source will be introduced, with the potential to significantly enhance the tunability of such accelerators.

SECTION 4.1

Flow models

In high-average-power LPA, the plasma formation is predominately assigned to the laser driver. Consequently gas flow dynamics is pivotal in the source design, as the plasma density obtained during the interaction is intrinsically connected to the gas density available before the ionisation process. Therefore, in the following, the most important quantities that govern the flow regime in which a source needs to be operated are presented.

Generally, fluid motion can be dominated either by viscous (laminar flow) or inertial forces (turbulent flow). Their ratio, the Reynolds number, has a crucial role in predicting the flow dynamics [87]. It is defined as:

$$Re = \frac{F_i}{F_v} = \frac{\rho u D}{\nu_d}, \quad (4.1)$$

where ρ is the fluid density, u its velocity, D the characteristic dimension of the flow and ν_d is the fluid dynamic viscosity¹.

¹Generally, a flow can be considered laminar when $Re < 2000$ and turbulent when $Re > 4000$. However, for CFD purposes, the choice of the flow model to be used really depends of the physics to be analysed. Often

In designing a plasma source, the Reynolds number can be particularly useful to anticipate scalings and transitions between different flow regimes. In turbulent flow, for instance, a transition to laminar regime always happens in the vicinity of a wall, in the so called “boundary” layer, where the velocity rapidly drops from free stream to zero. Boundary layers generally determine the gradients of the major flow quantities as, for instance, the gas density. Therefore, foreseeing their thickness can be extremely beneficial in modelling and simulating a plasma source both in terms of computational costs and convergence validation.

Another dimensionless fluid quantity, fundamental when flow dynamics evolves on the micrometer scale, is the Knudsen number, defined as :

$$Kn = \frac{\lambda_{\text{free}}}{D}, \quad (4.2)$$

where λ_{free} is the molecular mean free path and D is a characteristic length representative of the physical scale of the problem. The value $Kn = 0.01$ marks the limit over which the continuum flow approximation is not valid anymore and transitional (or molecular) flow has to be taken into account. However, when sources for high average power LPA are considered, the high pressure needed to reach the desired density often results in a high Re and a $Kn \ll 0.01$. Therefore, in the following parametric study, continuum flow approximation can be assumed and turbulences have to be considered.

SECTION 4.2

Gas nozzles

As discussed in chapter 2, the blowout regime and its advantageous fields are accessible with mJ-level, high repetition rate lasers, given the correct scaling of all the relevant parameters. In this scenario, the need for relativistic intensities calls for a tight focusing geometry, leading to spotsizes of only few microns and an exceptionally short Rayleigh length (Z_R). Moreover, matched laser propagation relies on remarkably high plasma densities resulting in an abrupt reduction of the dephasing length (L_{dep})².

P_{peak} [TW]	n_e [cm^{-3}]	w_0 [μm]	z_R [μm]	L_{dep} [μm]
0.5	3.7×10^{19}	2.5	19	47

Table 4.1.: Characteristic interaction lengths for an ideal gaussian pulse with $a_0 = 2$ and 0.5 TW peak power (P_{peak}). Here n_e is the desired electron density for transverse matched propagation, w_0 is the laser spot size, z_R and L_{dep} are the resulting Rayleigh and dephasing length.

a laminar description of a turbulent problem can effectively describe mean evolution of the major fluid quantities (cf. appendix D.2)

² L_{dep} is the length over which an injected electron is in the correct phase with respect to the accelerating field of the plasma wave and can hence be accelerated.

In table 4.1 a 0.5 TW (5.3 mJ, 10 fs) ideal Gaussian pulse with $a_0 = 2$ serves as an example: the necessity for a precise and extremely compact source becomes evident.

With high densities and short Rayleigh lengths, laser refraction can significantly influence the propagation. Therefore, to ensure optimum coupling between the laser and the plasma it is essential for the source to provide steep density gradients. Additionally, the demand for high plasma density, combined with the kHz operation, results in a substantial gas load in the vacuum chamber, frequently difficult to manage. Effective gas confinement, hence, is of critical importance.

A commonly used method to fulfill these requirements and tackle the challenges just outlined are gas nozzles [14, 16, 64]. In its most basic configuration, a nozzle is a straight, narrow channel connected to a high pressure reservoir. Despite the high backing pressure usually needed, such geometry keeps the gas flow velocity below the sound speed and, hence, it is frequently referred to as “subsonic”. While subsonic nozzles have proven to be a viable option, their ability to confine the gas flow is heavily limited by the sudden transverse expansion happening at their outlet. As a result, to attain the necessary gas (and plasma) density, an exceedingly large amount of gas is needed, consequently severely degrading the experiment vacuum levels.




Figure 4.1.: Axial evolution of Mach number (green), gas molecular density (blue) and temperature (red) within the typical de Laval geometry sketched on the top, with throat diameter (t), length (L) and outlet diameter (D).

To enhance the confinement, one could employ supersonic jets [88, 89]. A sketch of a standard “de Laval” supersonic configuration is presented in figure 4.1, where a converging

section is followed by a conical diverging one. In the first region the gas undergoes isentropic expansion, leading to a continuous decrease in density and temperature. Entering the diverging section, additional expansion causes the gas to accelerate and cool down, gradually increasing its Mach number³ and ultimately reaching supersonic velocities. At the nozzle exit (set at $z = 0$ in the figure) low gas pressure and low temperature lead to a much more contained vacuum expansion with respect to a subsonic geometry, resulting in a more confined flow.

The upcoming parametric study, performed with ANSYS-FLUENT and assuming nitrogen (N_2) as gas, aims at assessing the most important scaling laws and highlights the criticalities of both types of sources. Therefore, it is instrumental, here, to introduce some definitions to quantify the relative influence of the considered parameters on the gas distribution at the nozzle outlet, where the laser plasma interaction takes place. Figure 4.2 considers, as an example, the typical gas plume at the exit of a de Laval nozzle (left panel). In particular, the confinement angle α is defined as complementary to the angle containing 95% of the mass flow at the jet output. Therefore, a larger α indicates a better gas confinement⁴. On the other hand, in the right panel, the radial density profile 50 μm above the nozzle tip serves as an example to define the width of the density gradient Δl at a given height, i.e. the length over which the gas density changes from 90% to 5% of its peak value.

With these quantities now in place, an analysis of their variations with respect to different input parameters will be conducted to characterise the density profiles achievable using both subsonic and supersonic nozzles.




Figure 4.2.: Typical 2D axisymmetric normalised gas density plume at the exit of a supersonic de Laval nozzle with $D = 100 \mu\text{m}$ (left). The normalised radial density trend 50 μm above the nozzle is plotted in blue in the right panel. The parameters α and Δl (as defined in the main text) are highlighted in red.

³The Mach number (M) is a dimensionless quantity expressing the ratio between the fluid velocity and the local speed of sound. $M > 1$ indicates a supersonic flow.

⁴The angle is derived from a straight line fitting, considering the gas density distribution at the jet exit up to an height of 50 μm above its tip.

Mach number and supersonic nozzles

A gas expansion in vacuum always results in supersonic speed. This is due to the decrease of the speed of sound due to the temperature drop ($v_s \propto \sqrt{T}$), rather than an exponential increase of the absolute gas velocity. The different terminology between “subsonic” and “supersonic”, hence, originates from the fact that in a de Laval geometry a first expansion happens already inside the nozzle. This leads to a Mach number > 1 before the jet exit, as shown in figure 4.1. The same does not happen in a straight-channel geometry.

4.2.1. Subsonic nozzles

In the case of a straight subsonic nozzles the density scaling with respect to different input parameters is rather intuitive. Figure 4.3 condenses the main dependences.

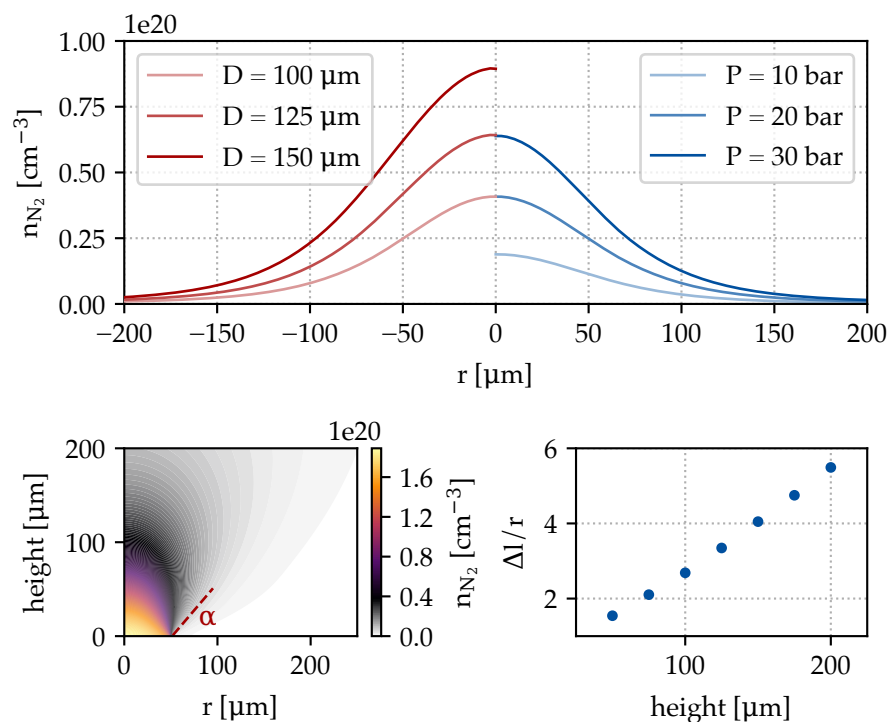


Figure 4.3.: Overview of the major parametric dependencies for subsonic nozzles. A nozzle with $D = 100 \mu\text{m}$ backed by $P = 20 \text{ bar}$ is taken as a reference. The top panel shows the dependence of the density profile $100 \mu\text{m}$ above the nozzle tip on source diameter (red) and backing pressure (blue). The bottom-left panel depicts the typical gas plume. The bottom-right shows how the density gradient (Δl , normalised to the outlet radius r) changes with respect to the height above the nozzle.

As shown in the top panel, where density trends are traced $100 \mu\text{m}$ above the nozzle tip,

the nature of the gas expansion give rise to a wide gaussian-like profile, which scales linearly with the backing pressure P and quadratically with the output diameter D . The width of the density gradient (Δl) rapidly grows with the height above the nozzle, soon reaching multiples of the outlet radius r and denoting a sub-optimal gas confinement. As a further evidence, the bottom-left panel represents a typical gas plume at the exit of a 100 μm diameter nozzle, backed with 20 bar of nitrogen. The resulting confinement angle (α) measures only around 50° and it is determined mainly by the expansion in vacuum, rather than by the backing pressure or the source geometry.

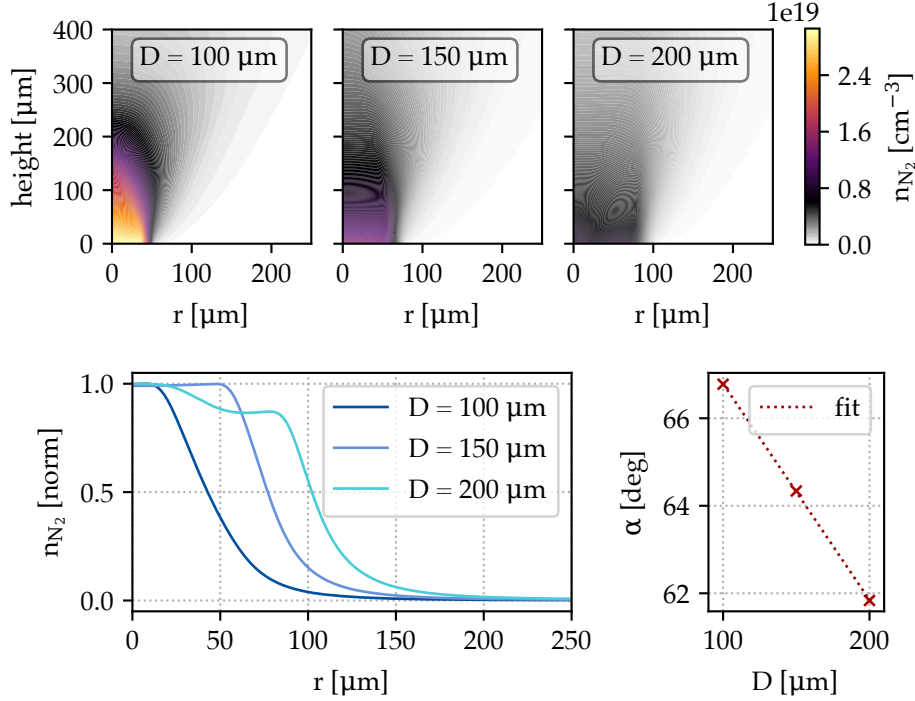


Figure 4.4.: Gas density plume above the nozzle for different outlet diameters (top), with the correspondent radial density profiles 100 μm above the nozzle (bottom left) and confinement angle α (bottom right). A linear fit (dotted red) has been added to highlight the α trend. The remaining parameters are fixed, with $t = 40 \mu\text{m}$, $L = 200 \mu\text{m}$ and backing pressure $P = 20 \text{ bar}$.

4.2.2. Supersonic nozzles

Supersonic nozzles, on the other hand, can provide significantly better confinement. However, the expansion plume at their output is highly affected by the source geometry, which can consequently have a major impact on shaping the resulting density profile.

With reference to the geometry introduced in figure 4.1, the nozzle output diameter D is the main factor affecting the total width of the density profile. The top panel of figure 4.4

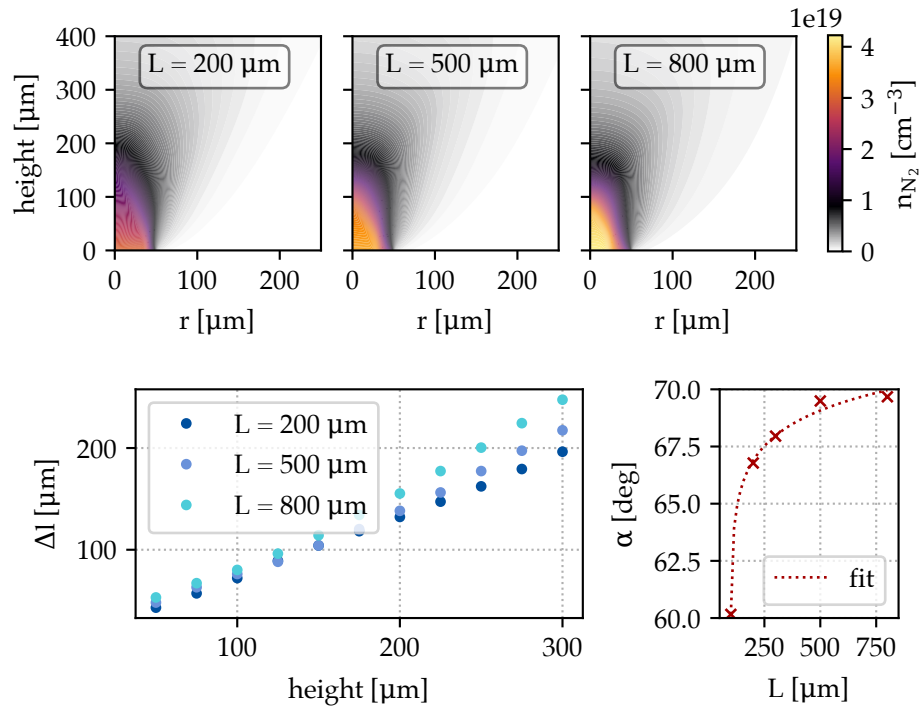


Figure 4.5.: Gas density plume above the nozzle for different nozzle lengths L (top). The corresponding ramp lengths Δl are shown in the bottom left panel. The bottom right plot depicts the evolution of the confinement angle α , where the results from two additional simulations ($L = 100 \mu\text{m}$ and $L = 300 \mu\text{m}$) have been added for a better logarithmic fit (red dotted line). The remaining parameters are fixed, with $t = 40 \mu\text{m}$, $D = 100 \mu\text{m}$ and $P = 20 \text{ bar}$.

shows the influence of different outlets on the gas plume at the exit of a de Laval nozzle with $t = 40 \mu\text{m}$, $L = 200 \mu\text{m}$ and backing pressure $P = 20 \text{ bar}$. The mass-flow rate through the geometry is fixed by the inlet conditions, therefore a smaller diameter results in a higher peak density and a more confined flow. The bottom panels depict the normalised density profiles traced $100 \mu\text{m}$ above the nozzle tip (left) and the angle of confinement α for the different geometries (right). Increasing the diameter D the confinement is reduced, while the total width of the profile follows the opposite trend. Hence, given the characteristic lengths outlined in table 4.1, opting for a smaller outlet represents a better alternative.

Interestingly, the density profile corresponding to $D = 200 \mu\text{m}$ shows a small rise close to the nozzle axis. This sort of modulations are known as “mach waves” and arise from the region soon after the nozzle throat where the flow, expanding at supersonic velocity, needs to adapt to the wall of the diverging section. The larger the diverging angle⁵, the more abrupt is the flow expansion and the more these waves become relevant in the flow [89].

⁵This refers to the expansion angle inside the nozzle structure, hence to $\beta = \arctan((D - t)/2L)$. In the current scenario a larger β simply means a bigger nozzle outlet D .

This behaviour, even being fully turbulent, can be exploited to have density rises at specific locations [90].

The gas expansion and evolution, on the other hand, is dominated by the length of the diverging section L , which defines the flow conditions at the nozzle tip. Figure 4.5 provides an overview of how the parameter affects the density at the nozzle outlet. 2D density contours are plotted in the top panel for lengths L of 200, 500 and 800 μm , keeping the remaining geometrical parameters fixed ($t = 40 \mu\text{m}$, $D = 100 \mu\text{m}$) and $P = 20 \text{ bar}$. The respective ramp length Δl and confinement angle α are shown in the bottom. A longer nozzle generally increases the flow directionality and enhance the confinement, but it tends to elongate the ramps, which could be detrimental when coupling the laser energy into the plasma.

Such behaviour can be explained looking at the relative influence of the boundary layers on the flow. As the free stream velocity keeps increasing in the nozzle diverging part, the layers tend to grow in size with its length, affecting the initial width Δl of the ramps at the nozzle exit. Figure 4.6 shows the normalised gas density contour lines within the diverging section of the nozzles under investigation, where the nozzle throat is placed at $L = 0$ and the nozzle outlet is at the vertical end. The growth of the layers is particularly accentuated in the third case, where the density ramp is already $\sim 20 \mu\text{m}$ long even before exiting the nozzle⁶. However, as the gain in the confinement angle α rapidly saturates with increasing L (see bottom-right panel of figure 4.5), a nozzle length of $\sim 300 \mu\text{m}$ appears a valuable trade off.




Figure 4.6.: Normalised density contour lines in the diverging section of the nozzle for different nozzle lengths L . The boundary layer grows with the development of the flow along the section. The other geometrical parameters are fixed ($t = 40 \mu\text{m}$ and $D = 100 \mu\text{m}$). The backing pressure is 20 bar.

Last but not least the nozzle throat t , being the smallest section, has a major impact on the mass flow rate through the jet and on the dynamics of the first inner expansion, which in turns highly affects the flow directionality at the outlet. The top panels of figure 4.7 show

⁶Note that a shock front is again evident where Mach waves can develop, i.e. for the shortest L .

how its dimension affects the density plume in a nozzle with $L = 200 \mu\text{m}$, $D = 100 \mu\text{m}$ and backing pressure $P = 20 \text{ bar}$. On the other hand, the bottom panels depict how both the confinement angle α and the ramp length Δl get significantly better with smaller throats.

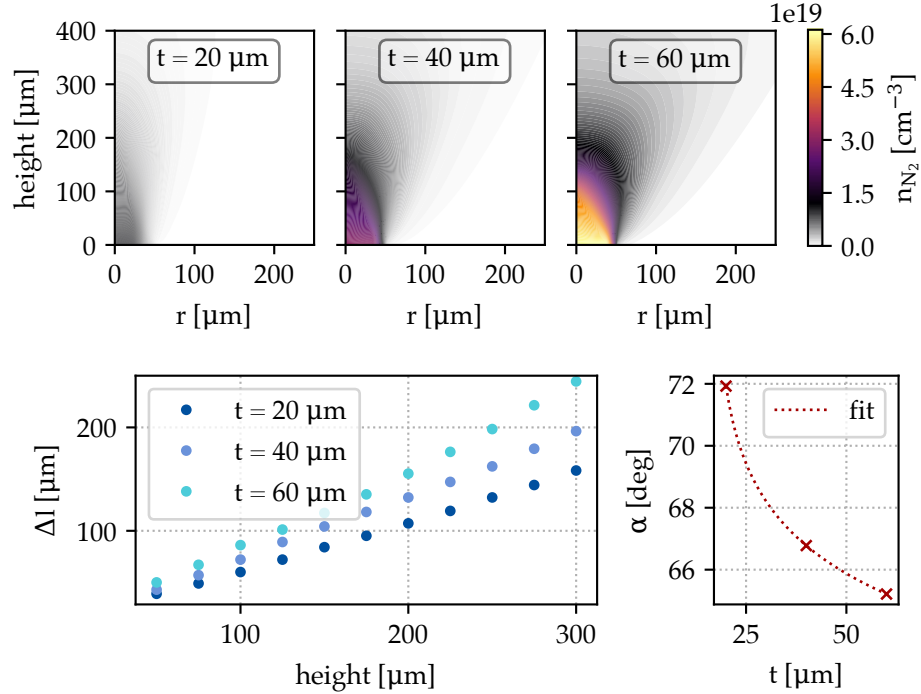


Figure 4.7.: Gas density plume above the nozzle for different throat diameters t (top). The respective ramp length Δl and confinement angle α are plotted in the bottom panels, for $L = 200 \mu\text{m}$, $D = 100 \mu\text{m}$ and $P = 20 \text{ bar}$. The red dotted line shows a logarithmic fit of the data.

This behaviour is linked to the angle of the inner diverging section of the nozzle: a larger angle (and hence a smaller throat) leads to a more sudden inner expansion and to a higher temperature drop along the nozzle length. As a result, once the flow reaches the nozzle outlet, the transverse expansion in vacuum is contained and the flow is more collimated. Therefore, thinner throats are the most suitable option, provided that the backing pressure can be scaled accordingly. In this context, figure 4.8 compares two radial density profiles $100 \mu\text{m}$ above the nozzle tip for two different throat diameters, $20 \mu\text{m}$ and $40 \mu\text{m}$, where the backing pressure has been scaled in order to have a similar mass flow rate, thereby a similar peak density (left panel). The temperature trend for the two geometries in the diverging section of the nozzle is plotted on the right, where the throat is placed at $z = -200 \mu\text{m}$ and the nozzle outlet at $z = 0$. A thinner throat results in a lower temperature at the nozzle exit and, hence, a more confined gas plume.

However, when passing from the design to the realisation of these sources, an excessively small, round feature can be difficult to manufacture and more sensitive to the absolute sur-




Figure 4.8.: The radial density profiles 100 μm above the tip for different combinations of t and P , resulting in a similar peak density (left panel). The corresponding temperature drop along the nozzle diverging section is plotted in the right panel, where the nozzle outlet is set at $z = 0$ and the throat at $z = -200 \mu\text{m}$.




Figure 4.9.: An “optimal” supersonic geometry with $t = 40 \mu\text{m}$, $L = 300 \mu\text{m}$ and $D = 100 \mu\text{m}$ (blue) backed with $P = 20 \text{ bar}$ is compared to the respective one resulting from a subsonic channel with the same outlet diameter and backing pressure. Their resulting normalised gas density profile (nitrogen) is plotted at 100 μm (left) and 150 μm (right) above the nozzle tip.

face roughness, which can ultimately lead to flow instabilities. The thinnest possible throat, then, heavily depends on the manufacturing capabilities. A good compromise is to opt for $t \sim 40 - 50 \mu\text{m}$, a dimension within the reach of the most common production techniques, selective laser etching (SLE) and micro electrical discharge machining (EDM).

As a result from the analysis above, figure 4.9 compares the normalised molecular density profile for a standard subsonic channel with $D = 100 \mu\text{m}$ and $P = 20 \text{ bar}$ to an optimised de Laval geometry with equal outlet diameter and backing pressure, $t = 40 \mu\text{m}$ and $L = 300 \mu\text{m}$. As expected, the supersonic geometry leads to a substantially improved gas confinement, with steeper density gradients. Moreover it ensures a reduced mass flow in the vacuum chamber ($\sim 1/4$) while reaching a comparable absolute peak density.

Nonetheless, despite the numerous advantages just highlighted, supersonic nozzles suffer from a remarkable lack of tunability. In fact, while the gas confinement is highly dependent on the jet geometry, changing the available parameters yields minimal effects on the overall functional form of the density profile. As a consequence, within the laser-plasma interaction, electron injection is poorly controllable and the accelerating length is minimal. Therefore a more tunable source, able to enhance both the energy and the quality of the accelerated electron beam, would be highly desirable.

SECTION 4.3

A novel tunable microfluidic source: the VERA source

Generally, decoupling the injection from the acceleration process can significantly enhance the quality of the accelerated beam [43, 45]. When gas jets are employed, shaping the density usually involves shock fronts, either embedded in the geometry [90], or formed with a thin blade above the nozzle tip [91]. However, both options offer limited tuning possibilities and, moreover, resorting to shock fronts may give rise to flow instabilities which will then affect the electron beam properties. Attempting to tackle these issues, this section proposes a novel plasma source able to supply a widely tunable density profile, while employing subsonic flows and therefore limiting the flow-related instabilities.

Sketched in figure 4.10, this novel source consists of a series (two in this specific case) of hollow, thin-walled, micro-capillary tubes, where gas is flowing at a controlled pressure and confinement is ensured by a closed system for capture and recirculation. In each capillary, two transverse pinholes (of diameter d) enable the passage of the driver laser. Intuitively, fine tuning of the gas (and therefore the plasma) density can be obtained by stacking capillaries of different geometries in series.

As an example, a typical density profile along the laser propagation axis (z , indicated by the red dashed line in the figure) is shown in the right panel, assuming two capillaries with inner diameters $D_1 = D_2 = 50 \mu\text{m}$, separated by a distance $L = 20 \mu\text{m}$. The μm -scale diameter and the reduced size of the capillary transverse pinholes allows for a subsonic, (almost) laminar flow in the tubes, generally suggesting great flow stability.

Having introduced the source (which from now on will be addressed as “VERA source”), a parametric study is now outlined to assess its main dependences and tuning capabilities.

The source is evidently not axi-symmetric, hence the analysis requires full 3D simula-




Figure 4.10.: Vertical cross-section of the novel microfluidic source proposed, as described in the main text (left). A typical (normalised) gas density profile along the laser propagation axis z is shown on the right panel.

tions. Meshing 3D geometries is nontrivial, as the computation time rapidly scales with both the volume and the mesh refinement. Therefore, to save computational resources, in the following only a single capillary is considered, which is the fundamental building block of the source. Additionally, the simulations presented are limited only to a reduced section, centered around the transverse pinholes and of 1 mm length⁷.

Figure 4.11 illustrates the main geometrical dependences, considering nitrogen as gas. The influence of the inner capillary diameter on the resulting density profile is depicted in the top panel, assuming a symmetric pinholes with diameters $d_1 = d_2 = 40 \mu\text{m}$ and 0.5 bar of backing pressure. Conversely, the middle pane focuses on the dependence on the pinholes size, while assuming an inner diameter $D = 50 \mu\text{m}$. Specifically the left pinhole is fixed at $d_1 = 40 \mu\text{m}$, while the right (d_2) varies between 30 and 50 μm .

As shown in the bottom left panel, the overall FWHM of the density profile is clearly proportional to the capillary diameter, which, on the other hand, has a minimal impact on the density ramps. Smaller inner diameters usually result in slightly longer ramps, as the relative influence of the boundary layers on the flow increases. Nevertheless, narrower capillaries offer, in the overall picture, a finer tuning parameter and therefore are usually to be preferred. On the other hand, as the last panel illustrates, the pinhole size significantly affects the length of the corresponding density ramp Δl (defined as in section 4.2), while it has a minimal effect on the width of the density profile. It is worth mentioning here that the thickness of the capillary walls can also impact the resulting density. Thinner walls intuitively give rise to steeper gradients and, therefore, are generally a better option. Specifically, in all the simulations presented in this section a wall thickness of 10 μm is considered,

⁷A convergence test with increasing mesh sizes (starting from 0.5 μm) showed that a relatively coarse mesh, with a refinement size of 3 μm is enough to capture the (minor) turbulent behaviour of the flow while still converging in an affordable time.

enough to withstand the gas pressures needed, assuming silica glass as material.

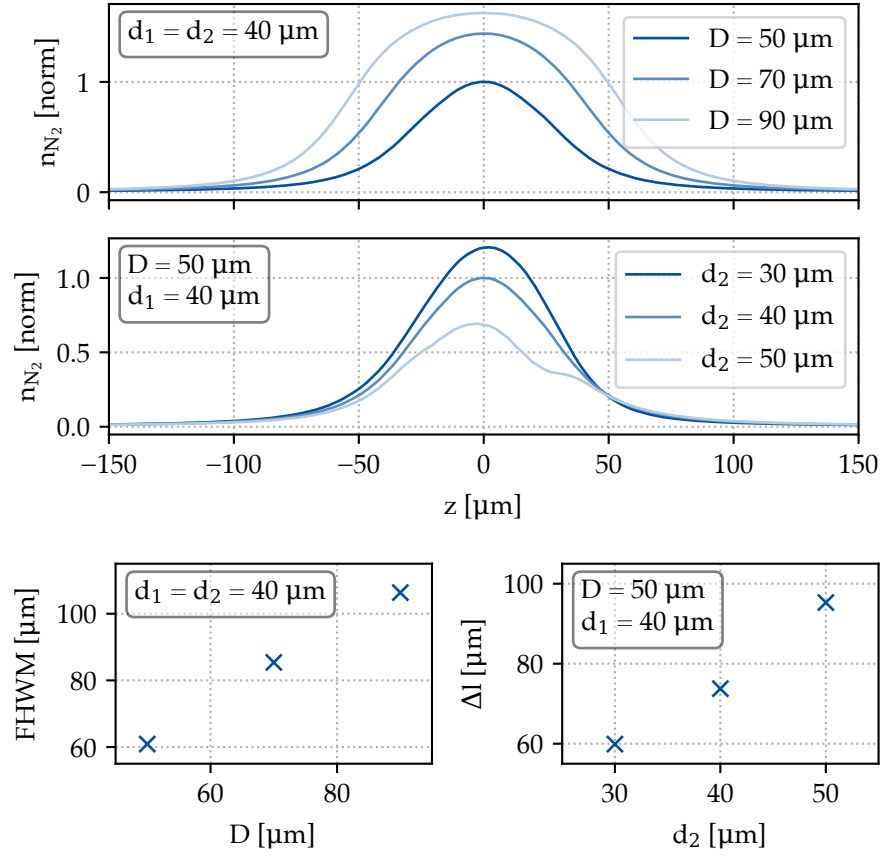


Figure 4.11.: Simulated gas density lineouts for a single building block of the VERA source along the pinholes axis, normalised to a reference geometry with $D = 50 \mu\text{m}$ and $d_1 = d_2 = 40 \mu\text{m}$. Top panel: inner diameter dependence (D) with constant pinholes diameter ($d_{1,2}$). Central panel: dependence on the right pinhole diameter (d_2) with constant d_1 and D . The corresponding FWHM of the profile and the length of the density ramps Δl are represented in the bottom row. The backing pressure is fixed at $P = 0.5$ bar for all the simulations.

Extending the analysis, the influence of different input pressures on the flow is of particular interest. The top panel of figure 4.12 depicts how different backing pressures affect the density lineout for a capillary with inner diameter $D = 50 \mu\text{m}$ and symmetric pinholes with $d_1 = d_2 = 40 \mu\text{m}$. For lower pressures, the peak density is directly proportional to the inlet conditions. However, for $P > 0.5$ bar the proportionality decreases and a shock front appears in the density lineout on both sides of the geometry. As the bottom panel shows, such behaviour is a consequence of the partial vacuum expansion happening around the pinholes: the fluid impinges on their top wall and initiates a shock, similar to those generated by the insertion of blades on top of gas nozzles. Once steady state is reached, the shocks on

both sides of the capillary build up and give rise to a single transverse front, as shown in the CFD simulated contours. Such shock is a possible source of flow instabilities which, ideally, should be prevented. Moreover, from the driver laser perspective, it can generate a sharp transverse gradient with respect to the laser propagation axis, which could be detrimental for the acceleration process. Particle-in-cell (PIC) codes could be employed to estimate its effective impact on the LPA process (cf. section 5.1.1). However, while computationally efficient PIC codes usually work on a 2D cylindrically symmetric grid, this sort of density profile clearly breaks the symmetry around the optical axis. Therefore, this scenario appears particularly difficult to simulate and would require a full 3D description. Depending on the laser parameters and on electron injection method used, the shock may have a negligible effect on the overall acceleration process. However, at this stage, it will be considered detrimental for the source stability and, thus, avoided. In light of this discussion, the presence of a marked shock front acts, in practice, as an upper limitation for the achievable density. Consequently, depending on the gas and the geometry adopted, the gas density is usually constrained to $n_{\text{gas}} < 1 \times 10^{19} \text{ cm}^{-3}$. This elucidates why nitrogen has been considered as gas up to this point: high- z gases, given their multiple ionisation levels, can provide higher plasma densities from equal backing pressure and, hence, help in mitigating the issue.

The various trends highlighted thus far depict the outstanding tunability of the VERA source. An additional degree of freedom, when stacking multiple capillaries in series, is their separation L (cf. figure 4.10), which governs the gas mixing regions. In this context, figure 4.13 serves as an example, considering a double-capillary configuration, where $D_1 = D_2 = 50 \text{ }\mu\text{m}$, the backing pressures are respectively $P_1 = 0.5 \text{ bar}$ and $P_2 = 0.15 \text{ bar}$ and the pinholes diameters (numbered from left to right) are $d_1 = d_4 = 50 \text{ }\mu\text{m}$ and $d_2 = d_3 = 40 \text{ }\mu\text{m}$ ⁸. Lastly, while so far circular capillary cross-sections have been assumed, this is not a stringent requirement and different shapes can provide even more versatility. A rectangular cross-section, for instance, offers the same longitudinal tunability while providing greater flexibility in transverse pinhole sizing⁹. Moreover, avoiding curved surfaces, it facilitates optical access for the plasma density diagnostics during the operation.

In conclusion, this novel microfluidic plasma source, while generally being more complex than the commonly used nozzles, has the potential to enable an exceptionally precise control of the plasma density profile. Critically, its great tunability can have a tremendous impact on the characteristics and the stability of the accelerated electron beam, as will be discussed in chapter 5.

⁸This sort of pinholes schemes mimic a situation where the driver laser is focused in the middle, between the capillaries. The two central pinholes $d_{2,3}$, hence, can be smaller than the outer ones $d_{1,4}$.

⁹The pinholes, in fact, have to provide access to the driver laser. Their diameter is highly correlated to the driver spot size, its Rayleigh length and the overall source longitudinal dimension.

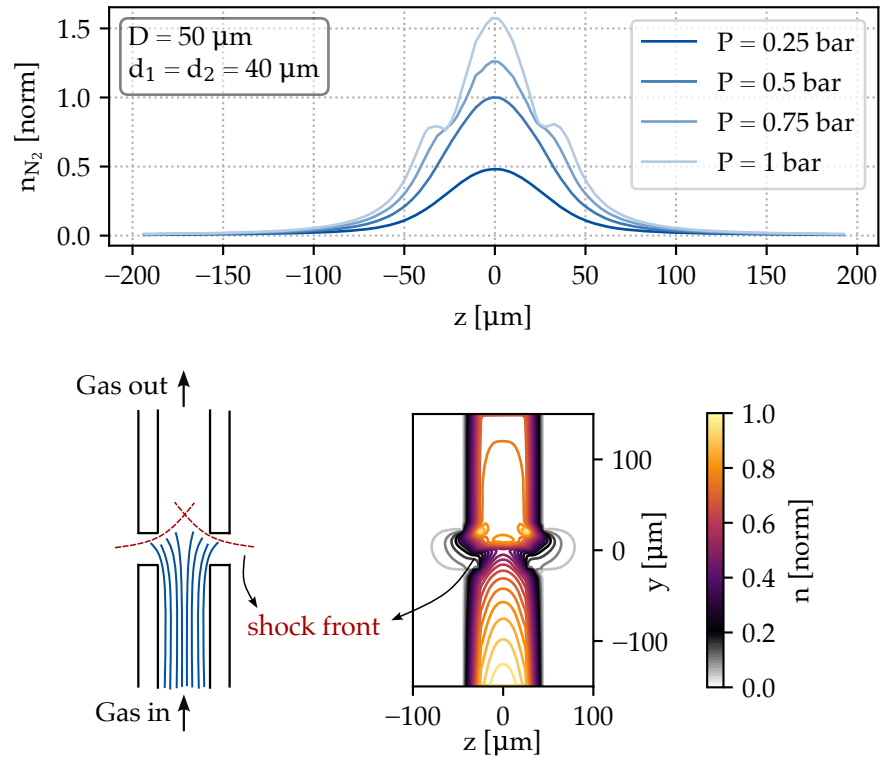


Figure 4.12.: Normalised gas density lineouts for a single capillary along the pinholes axis, as a function of the backing pressure (P), with $d_{1,2} = 40 \mu\text{m}$ and $D = 50 \mu\text{m}$ (top panel). The bottom pane illustrates the formation of the shock front in the highest pressure case ($P = 1$ bar), as described in the main text.




Figure 4.13.: Normalised gas density lineouts along the pinholes axis for a double capillary configuration, as a function of the distance between the capillaries L . The backing pressures are $P_1 = 0.5$ bar and $P_2 = 0.15$ bar and the remaining geometrical parameters are fixed, with $D_{1,2} = 50 \mu\text{m}$, $d_{1,4} = 50 \mu\text{m}$ and $d_{2,3} = 40 \mu\text{m}$.

4.3.1. The VERA source: first designs and tests

Besides the extensive simulation study, first prototypes of the VERA source have been designed and tested.

Thin-walled, quartz circular capillaries with inner diameter of $\sim 100 \mu\text{m}$ are routinely used in material science (X-ray crystallography, X-ray diffraction...) and, therefore, can be bought off the shelf. Figure 4.14 shows, for instance, a microscope image of the first capillary tested, where the transverse pinholes have been drilled directly by the VERA laser (after meticulous alignment). The capillary shown, specifically, has an inner diameter of $110 \mu\text{m}$ and an outer diameter (at the position of the pinholes) of $180 \mu\text{m}$. The diameter of the transverse pinholes resulted to be roughly $60 \mu\text{m}$ on both sides after ~ 24.5 minutes of continuous operation at 1 kHz.




Figure 4.14.: Microscope image showing the input and output pinholes at the end of the first test for a single building block of the VERA source, as described in the main text. A quartz capillary was used, it withstood 1.5×10^6 laser pulses (24.5 minutes at 1 kHz).

Aside from simple capillaries, more advanced designs can be manufactured with extreme precision through SLE (selective laser etching). To illustrate the capabilities of the etching technique, figure 4.15 shows the most recent prototype, designed in order to experimentally test single-capillaries of different geometries. Three channels of different size, with rectangular cross-section, are embedded in a cubic structure. A single inlet can supply gas to all the channels, in parallel. Each channel has three different shooting positions along its height, marked by inward facing cones to ease the alignment and guide the transverse driver laser in and out of the capillaries. This prototype is already available in the laboratory and it will be soon subjected to extensive testing. In figure 4.15, the rightmost panel shows a close-up for one of the shooting position. A section of the design (top) is compared to a real microscope image of the source (bottom).

The novel source discussed in this section is the subject of a patent application submitted to the European Patent Office and currently under review for approval [92].




Figure 4.15.: 3D rendering of the most recent single-building-block prototype, as described in the main text. The bottom right panel show a real microscope image of one of the “shooting positions” of the source, manufactured through SLE.

Slip-flow in microchannels

The analysis of the proposed new microfluidic source often referred to extremely narrow channels, with inner diameters extending only over few tens of μm . Such small ducts, under given conditions, can give rise to a particularly high Knudsen number (introduced in section 4.1 as $Kn = \lambda_{\text{free}}/D$) and, ultimately, modify the flow regime of the source. For gases, the mean free path λ_{free} (at constant temperature) is inversely proportional to the gas pressure P , therefore the Knudsen number could be expressed as:

$$Kn \propto \frac{1}{PD}. \quad (4.3)$$

Generally in fluid dynamics, when the continuous approximation is valid, the “no-slip” boundary condition assumes that the fluid velocity at a solid boundary is equal to the velocity of the boundary itself, typically zero. Yet, when low gas pressures and microchannels are considered, equation 4.3 may result in $Kn \gtrsim 0.01$. The corresponding flow, hence, would be located in a intermediate regime between a continuous and molecular flow, commonly known as “slip-flow regime”. Here rarefaction effects start to influence the flow, as the collision frequency is not enough to ensure thermodynamic equilibrium. As a consequence, a discontinuity arises in the fluid velocity (and temperature) at the wall-boundary, causing the gas to effectively glide over the solid surface [93]. The novel source proposed in the previous section hence, in addition to the numerous features discussed, could theoretically enable operation in the slip-flow regime. From a plasma source perspective, a fluid velocity at the wall-boundary different from zero could significantly reduce the extension of the boundary layers, resulting in an improved gas confinement and, generally, steeper density gradients.

CHAPTER 5

Modelling and optimisation of high average power LPA

The previous chapter introduced the most common plasma sources for high average power acceleration from a fluid dynamics perspective. However, when considering the accelerated electron beam, it is crucial to investigate how a given density profile would affect the accelerated electrons properties.

In this context a systematic comparison of different plasma sources can be particularly complex, due to the intricate, highly nonlinear nature of the laser-plasma interaction. Nonetheless, the combined use of simulation tools and optimisation algorithms can significantly streamline the process and assist in the assessment and evaluation of the accelerated bunch properties.

SECTION 5.1

Modelling laser plasma accelerators

5.1.1. The particle-in-cell code FBPIC

To fully characterize the complex dynamics of laser-plasma interaction, computational tools are essential. Here, particle-in-cell (PIC) codes emerge as ideal solutions, effectively capturing the interplay between charged particles and electromagnetic fields. To handle the computation load, these codes typically use “macroparticles”, which aggregate many physical particles based on their phase space (i.e. position and momentum). The electromagnetic fields, on the other hand, are usually discretised onto a grid. The system, then, loops through discrete time steps, where the fields push the particles forward¹ and the macroparticles deposit their charge and current to neighboring grid points, advancing the fields.

This thesis, specifically, employs the quasi-3D PIC code FBPIC [31]. While in a generic PIC a 3D cartesian grid is used, FBPIC takes advantage of the (close to) cylindrical symme-

¹The field felt by a macroparticle is a weighted sum of the field at neighboring grid points

try of most of the LPA-related scenarios and describes the fields in a computationally more efficient cylindrical grid. Moreover, rather than exclusively modelling radially symmetric fields, the code incorporates their azimuthal Fourier decomposition, enabling the simulation of fields with azimuthal dependencies². In most of the cases, such decomposition can be truncated to the first few modes (typically 2 or 3), resulting in a simulation which capture all the relevant 3D effects at the cost of few 2D grids.

The simulations presented in the following sections, if not differently stated, are performed with a grid size $\Delta z \times \Delta r = 0.02 \times 0.2 \mu\text{m}^2$ and three azimuthal modes³. Moreover, as ionisation-induced defocusing could significantly affect the laser propagation (cf. section 2.4.1), such simulations are always initialised with neutral atomic density (32 macroparticles per cell) and laser ionisation is modelled according to the Ammosov-Delone-Krainov (ADK) model [94]⁴. For reference, table 5.1 outlines the laser intensity needed to ionise a given species of helium or nitrogen. These intensities can be computed setting the peak of the laser field equal to the “critical” field, which would suppress the Coulomb potential of a bound electron in an atom, a theoretical model commonly known as “barrier-suppression ionisation” [95].

Species	$I [\frac{\text{W}}{\text{cm}^2}]$	Species	$I [\frac{\text{W}}{\text{cm}^2}]$
He ¹⁺	1.46×10^{15}	N ³⁺	2.25×10^{15}
He ²⁺	8.77×10^{15}	N ⁴⁺	9.01×10^{15}
N ¹⁺	1.78×10^{14}	N ⁵⁺	1.47×10^{16}
N ²⁺	7.68×10^{14}	N ⁶⁺	1.03×10^{19}

Table 5.1.: Barrier suppression ionisation levels of Helium and Nitrogen (only the first six) in terms of the laser intensity.

5.1.2. Principles of Bayesian optimisation

When trying to assess the capabilities of different plasma sources, relying only on PIC codes can be restrictive, particularly if aiming at a quantitative comparison. The acceleration process, in fact, is extremely sensitive to a wide range of parameters and even minor changes can yield markedly different properties of the resulting electron beam. Consequently, comparing sources demands an extensive parametric exploration and, therefore, substantial computational resources and time. In this scenario, the application of machine learning algorithms, especially when rooted in bayesian optimisation, emerges as a highly valuable solution that can simplify the parameters’ exploration and offer great efficiency gain.

²as, for instance, the one of the driver laser, usually linearly polarised.

³Assuming $\lambda_L = 1.03 \mu\text{m}$ and a plasma density $n_e = 1 \times 10^{20} \text{cm}^{-3}$, this grid corresponds to 52 points per laser wavelength on z .

⁴The model is implemented in FBPIC by default.

Bayesian optimisation (BO) [96] is an iterative strategy to find the optimal set of parameters for a given objective function f and can be particularly effective in the evaluation of black-box functions, where no functional form is assumed a priori. Its approach is based on a “surrogate model”, which generally serves as a probabilistic approximation of the true objective in the given parameter space. The iterative process involves evaluating f for a specific parameter set, updating the surrogate model accordingly and, hence, refining the space exploration. As the approximation of f is iteratively improved, a global maximum (or minimum) can be identified, pointing at the optimal solution for the problem. The selection of parameters in each iteration is guided by an “acquisition function”, which assesses the potential utility of the next measurement in locating the global optimum. This function typically aims at balancing the sampling of regions where few data points are present (exploration) and the search for an optimum where high objective values are predicted (exploitation)⁵.

With a surrogate model and an acquisition function in place, an iteration of the optimisation routine consists of:

- ◇ an initialisation step, where few evaluations of the function f with random parameters are performed to seed the surrogate model,
- ◇ an optimisation step, where the acquisition function is maximised to locate the most promising set of parameters for the next evaluation,
- ◇ an evaluation step, where the function f is determined in the location suggested before and
- ◇ an update step, where the surrogate model is updated according to the evaluation.

Therefore, the algorithm methodically navigates the parameter space, ultimately pinpointing the global optimum of the objective function.

With the main tools for modeling laser-plasma interaction now introduced, the subsequent sections will explore the integrated use of PIC simulations and Bayesian optimization algorithms to thoroughly assess the capabilities of the plasma source described in the previous chapter from the accelerated beam’s perspective.

SECTION 5.2

Bayesian optimisation of high-average-power LPA

Chapter 4 showed how subsonic and supersonic nozzles, the most common sources in high-average-power LPA, exhibit substantial differences from a fluid dynamics perspective. On the other hand, evaluating and comparing these sources in terms of their electron acceleration capabilities is challenging, as the nonlinear nature of laser-plasma interaction, especially with sub-TW drivers, is highly sensitive even to minor parameter changes. In this context Bayesian optimisation can be particularly effective. The algorithm fine-tunes the interaction in order to have the best possible outcome according to a pre-defined fitness

⁵A common and relatively simple acquisition function is “expected improvement” (EI) [96].

function. Therefore, it excels in mitigating biases and enables a fair comparison between the sources.

In the following, specifically, the results of two different optimisations are considered⁶, where the electron beam obtained using a 100 μm -diameter subsonic nozzle is analysed alongside the “optimal” supersonic nozzle geometry presented in figure 4.9. In both cases nitrogen is used and the longitudinal density profile assumed in the PIC simulations is computed at a height of 100 μm above the source tip⁷. The parameters to be optimised are selected to address experimentally accessible variables, namely the peak gas density (directly proportional to the nozzle input pressure) and the laser vacuum focal plane, relative to the density profile.

Optimas, a Python package that allows for scalable Bayesian optimisation on high-performance computing systems [97], is used in combination with FBPIC. The optimiser has been set to maximise the spectral density of the accelerated electron beam [98], i.e. to minimise the objective function:

$$f \propto -\sqrt{Q} \frac{E}{\Delta E}, \quad (5.1)$$

where Q denotes the total bunch charge above 1 MeV, E is the median energy of the electron beam and ΔE represents its median absolute deviation⁸. To reduce the complexity relative to the laser driver, an ideal gaussian pulse is assumed in the simulations with $E_L = 6.7$ mJ, $\tau_{\text{FWHM}} = 9.6$ fs, resembling the parameters of the VERA laser after the second compression stage. The laser (vacuum) spot size is assumed to be $w_0 = 3.5$ μm ($1/e^2$ intensity radius), according to the experimental focusing conditions (cf. section 6.2), leading to a peak normalised potential in vacuum of $a_0 \simeq 1.6$.

Due to the limited peak power of kHz drivers, relativistic self-focusing often plays a key role in the acceleration process. It allows to overcome the detrimental effects of ionisation-induced defocusing and inject electrons (either by self-injection or ionisation injection) around the peak of the density profile. Conversely, though, the nonlinear effects due to the high plasma density heavily affect the pulse temporal structure and the strong focusing generates a rapid defocusing after the injection, typically resulting in electron energies extending only to few MeV. Figure 5.1, outlines the results of the best PIC simulation for both optimisation runs. Here the evolution of the laser’s a_0 is shown in red, the injected charge q in blue and the plasma density profile is represented by the grey-shaded area. In both cases the near-critical plasma density (2.8×10^{20} and $1.95 \times 10^{20} \text{ cm}^{-3}$) promotes strong self-focusing before the density peak. The sharper density gradients produced by the supersonic jet result in a reduced ionisation-induced defocusing. Consequently, self-focusing can be triggered later in the propagation and the energy coupling between the laser and the plasma is enhanced. In particular, the peak of the vector potential a_0 (and therefore the

⁶Each optimisation run is composed of 80 different simulations.

⁷The most common approach with nozzles is to set the laser as close as possible to the nozzle tip, so to interact with steeper density gradients. Here 100 μm is assumed to be a typical “safe” height to avoid damaging the nozzle tip, based on previously published experimental results.

⁸The presence of the square root in f assigns a reduced relative importance to the bunch charge, prioritising beams with smaller energy spread.

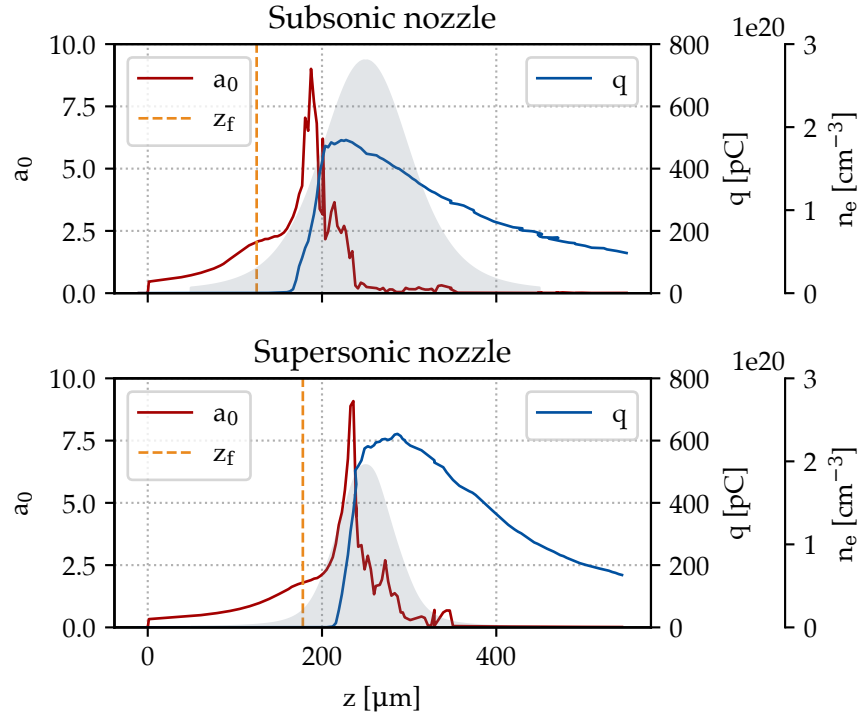


Figure 5.1.: Best simulations for a subsonic (top) and a supersonic nozzle (bottom), with geometries as defined in the main text. The evolution of the a_0 and of the injected charge q are shown in red and blue respectively. The longitudinal plasma density (n_e) profile is represented by the grey shaded area, while the orange dashed line indicates the vacuum focus position z_f .

peak intensity) is achieved at a longitudinal position where the plasma density is higher, generally resulting in a significantly higher injected charge. Following this path, figure 5.2 illustrates the electron spectra obtained for the two geometries, averaged over the 5 best simulations of the corresponding optimisation. Aside from the larger total injected charge Q , the spectra are remarkably similar, with a low energy peak around 2 MeV and a long tail extending to higher energies.

The results from these two optimisations illustrate a key point already mentioned before: despite their ease of production and use, nozzles often suffer from an inherent lack of tunability. This makes it challenging to optimise the accelerated electron beam for specific targets as, for instance, high spectral density, which is a critical parameter for various applications.

In this context, more flexible and advanced sources (as the novel source introduced in section 4.3) can offer substantial benefits. Their improved flexibility, though, frequently comes with a particular sensitivity to specific geometric parameters, where even minor adjustments can significantly influence the resulting density profile. Thus, employing Bayesian




Figure 5.2.: Electron spectra at the end of the propagation for the subsonic (top) and supersonic (bottom) geometry, averaged over the 5 best simulations. The inset shows the averaged spectral parameters: the bunch charge (Q), the median energy (E) and the median absolute deviation (ΔE).

optimisation in the design phase of these sources can be transformative: iteratively adjusting the source geometry during an optimisation enables precise tailoring of the density profile, that can target specific characteristics of the accelerated beam.

SECTION 5.3

CFD-aided Bayesian optimisation of advanced plasma sources

When more structured sources are considered, an optimization routine that relies exclusively on PIC simulations is severely limited, as it cannot incorporate the modifications in the gas density profile that only CFD simulations can accurately capture. Following this path, the next section focuses on the integration of ANSYS-FLUENT simulations into a bayesian optimisation loop with FBPIC for the design of advanced plasma sources. This pioneering, holistic method allows to optimise the source's geometry in order to achieve specific desired characteristics of the accelerated electron beam, holding the potential to significantly boost the performances of LPA, particularly in application-oriented scenarios.

5.3.1. Source geometry optimisation with ANSYS-FLUENT

Generally, integrating CFD simulations in an optimisation routine including PIC simulations is nontrivial, as it requires cross talk between geometrical modeling, meshing and fluid

dynamics to generate a density profile which could then be imported in FBPIC. Nonetheless, ANSYS Workbench allows all these separate building blocks to be included in a single workbench-project, which greatly simplify the process. In the following, a brief outline of the workflow will be presented, before introducing the VERA plasma source as a demonstration of the excellent optimisation ability of the method.

As FLUENT is mainly devoted to flow modeling, it requires the source geometry (and the corresponding mesh) to be imported separately. Here specifically, as with all other fluid simulations discussed thus far, ANSYS Design Modeler is used to build the source geometry, and ANSYS Meshing is employed to construct its mesh. To streamline the optimisation and minimize the eventual errors in geometrical modeling, a pre-built fluid simulation (whose convergence has been properly checked) is used as main “template”. This includes:

- ◇ The geometry and the mesh files, which describe the base source geometry. In the geometry a “parametrisation” can be applied to any source dimension or feature. This will generate geometrical parameters which can then be separately modified through a simple text file and, therefore, easily included in the optimiser. The use of a geometry-template generally allows to save computation time and, crucially, it enables the use of a pre-built mesh, which can be separately tested.
- ◇ The FLUENT setup file, that defines all the relevant flow properties and the appropriate solver settings. Even in this case, the use of a setup-template allows to test the CFD solver separately and provides a simple way to modify fluid parameters during the optimisation process.

Once a base template is prepared, the optimisation loop will go through the following steps:

- ◇ The optimiser selects specific values for the parameters to be optimised, now extending to the source geometry or to specific conditions of the CFD simulation (i.e. gas type, inlet conditions, flow model...)
- ◇ The geometry template is copied in a new ANSYS project and modified according to the specified parameters. A new mesh, with the same structure of the mesh template, is generated and the geometry is saved.
- ◇ The new geometry is imported into the FLUENT setup file, where the flow properties and the solver settings are defined. The optimisation parameters related to the flow conditions, if present, are now added.
- ◇ A FLUENT simulation is run, until a certain convergence criterion is met or for a given number of iterations.
- ◇ The resulting gas density profile (along a pre-defined axis) is extracted and saved in a separate file, while the rest of the newly generated ANSYS project is deleted.
- ◇ The simulated gas density profile is used to initialise the atomic density of the subsequent FBPIC simulation.

- ◇ The outcome of the FBPIC simulation is then analysed. The optimiser updates the surrogate model and a new set of parameters is chosen.

Having outlined the general procedure, this novel tool will guide the optimisation of the VERA source, extracting its full potential in terms of electron acceleration. The upcoming section will demonstrate the wide range of tunable parameters and the importance of the inclusion of CFD-aided BO in the design phase of a new source, particularly when aiming at specific characteristics of the accelerated electron beam.

5.3.2. Bayesian optimisation of the VERA source

The VERA source, shown again in figure 5.3 for convenience, is relatively complex from the geometrical perspective.




Figure 5.3.: Cross-section of the novel microfluidic source proposed, as previously described in section 4.3 (left). A typical (normalised) gas density profile on the laser propagation axis is shown on the right panel.

Given the micrometer scale of its double-capillary structure, all parameters of the source are closely interconnected. Modifying just a single parameter, as for instance the inner diameter of the first capillary D_1 , can directly affect the whole density profile along the laser propagation axis:

- ◇ it will alter the mass-flow-rate of the first capillary, generally impacting the peak gas density reached,
- ◇ it will influence the flow mixing between the capillaries, thereby modifying the central density ramps and potentially leading to a different peak density in the second capillary and
- ◇ assuming that the laser focal position remains fixed and the transverse pinholes are

drilled directly by the driver laser, it will modify the dimensions of the first two pinholes (d_1 and d_2), consequently affecting the first two density gradients.

Considering these small yet critical changes in the resulting density profile without the aid of proper CFD modeling is virtually impossible, particularly when multiple parameters are varied.

Theoretically, every single parameter of the source could be optimised. However such optimisation routine would require an increasing number of simulations to account for an increasing number of parameters. Moreover, as for this source the CFD simulations have to be performed on a 3D grid, this will ultimately result in an unaffordable increase of the computation time. Therefore, in order to demonstrate the capability of the method, the following optimisation will only focus on the following parameters:

- ◇ the inner diameter of the first capillary D_1 ,
- ◇ the inlet pressure in the first capillary P_1 ,
- ◇ the inner diameter of the second capillary D_2 ,
- ◇ the inlet pressure in the second capillary P_2 ,
- ◇ the laser vacuum focus position z_f .

Parameter	D_1 [μm]	P_1 [Pa]	D_2 [μm]	P_2 [Pa]	z_f [μm]
min.	30	40000	40	10000	-20
max.	50	60000	70	40000	20

Table 5.2.: Search space for the VERA source optimisation, where the parameters are defined in the main text. The centre of the geometry is placed at $z = 0$.

Table 5.2 outlines the search space for each parameter, where the centre of the geometry is placed at $z = 0$. Importantly, the transverse pinholes of the capillaries (d_{1-4}) are assumed to be drilled directly by the laser. Therefore, during an optimisation loop, each pinhole's dimension is computed separately, taking into account its distance from the laser focal plane, the real Rayleigh length of the VERA laser in the experimental setup ($z_R = 37.4 \mu\text{m}$) and assuming an intensity of $1 \times 10^{14} \text{ W/cm}^2$ as the lower limit for which a femtosecond pulse could drill into glass [99]. As cylindrical capillaries are assumed, the total length of the geometry is always chosen in order to never have a pinhole diameter larger than its corresponding capillary diameter⁹. Moreover, the optimiser is forced to keep the maximum inlet pressure in both capillaries below a given pressure limit (which depends on the diameter and the length of the capillaries), to avoid a marked shock front, as identified in section 4.3 (cf. figure 4.12).

⁹This condition ($d > D$) will ultimately cause the capillary to fail in practical applications. However, as discussed in section 4.3, capillaries cross-sections of different shapes can allow for more flexibility on the pinholes' size.




Figure 5.4.: Best simulation from the optimisation. The evolution of the laser a_0 and the injected charge q are plotted respectively in red and blue. The corresponding longitudinal plasma density profile n_e is represented by the grey shaded area.

The remaining geometrical parameters are kept fixed, with the distance between the two capillaries being $L = 20 \mu\text{m}$ and the capillaries' wall thickness assumed to be $10 \mu\text{m}$. Finally, as in the previous optimisations, nitrogen is assumed as gas and, to reduce the complexity derived from the driver, in the PIC simulations an ideal gaussian pulse is assumed, with $E_L = 6.7 \text{ mJ}$, $\tau_{\text{FWHM}} = 9.6 \text{ fs}$ and a vacuum spotsize of $3.5 \mu\text{m}$ ($1/e^2$ intensity radius), leading to a peak (vacuum) $a_0 \simeq 1.6$.

The original design of this two-stage source naturally suggests exploiting downramp injection around the center of the geometry. This injection method, while much more controllable than self-injection or ionisation injection, generally results in electrons being injected in various buckets along the downramp, which, consequently, will experience different accelerating and focusing fields in the second stage of the source (i.e. in the second capillary). Therefore, choosing an appropriate optimisation function is crucial: simply optimising for the maximum spectral density may not fully exploit the source's potential and could result in suboptimal performances due to the different conditions experienced by the injected electrons.

A practical approach is to divide the energy spectrum into distinct intervals and assign each interval a function to compute its spectral density. A weight factor can then be assigned to each one, so to prioritise the optimisation of the spectral density within a desired energy range while suppressing the others. In the following, the energy spectrum is divided in three different intervals, namely $\varepsilon_1 = [0.5, 2) \text{ MeV}$, $\varepsilon_2 = [2, 5) \text{ MeV}$ and $\varepsilon_3 > 5 \text{ MeV}$. For each interval, the corresponding spectral density is computed as $s_i \propto -\sqrt{Q_i} \frac{E_i}{\Delta E_i}$, where Q_i is the total bunch charge, E_i is the median energy and ΔE_i the median absolute deviation relative to the interval ε_i . The objective function is then formulated as:

$$f = \sum_i w_i s_i, \quad (5.2)$$

where w_i is the weight factor and the three weights are respectively 0.01, 1 and 0.01, so to maximise spectral density between 2 and 5 MeV. Dividing the spectrum in separate inter-

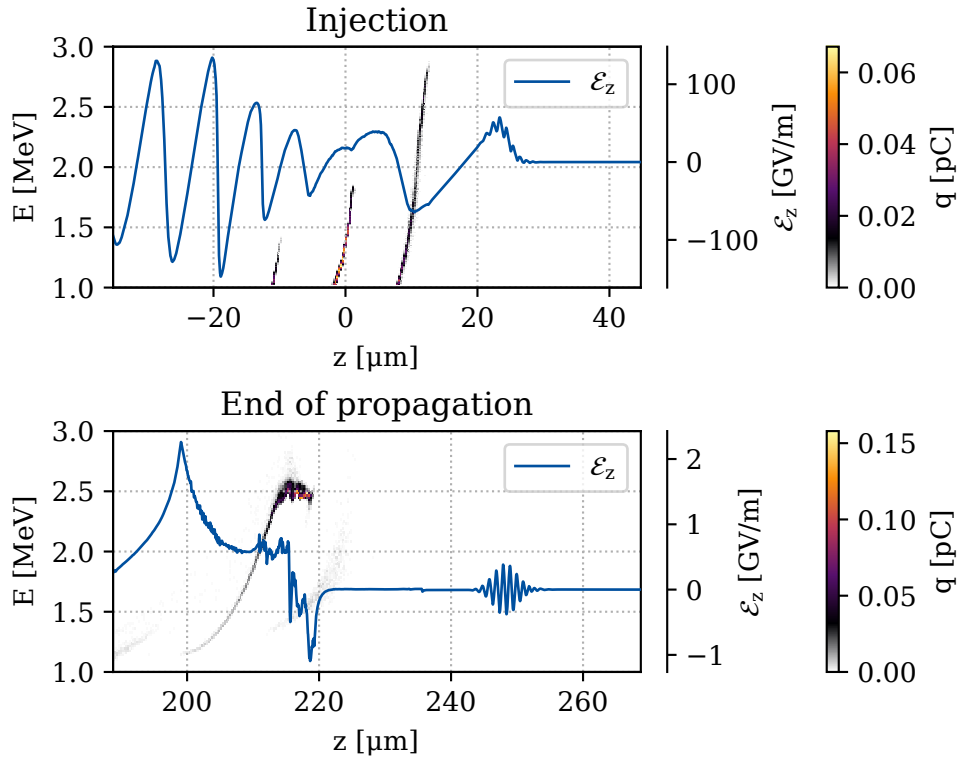


Figure 5.5.: Longitudinal phase space plot of the injected electrons with energy above 1 MeV at the injection instant (top) and at the end of the propagation (bottom). The corresponding longitudinal electric field \mathcal{E}_z is plotted in blue.

vals may initially appear redundant as, in theory, an equal result could be achieved considering only the desired energy range (e.g. simply defining $f = s_2$). However this approach would exclude from the analysis all the charge outside the specified range, possibly seeding the surrogate model with an excessive number of points where $f = 0$. This condition can easily mislead the algorithm and direct it away from the true optimal condition. Thus, incorporating a set of different energy intervals with small weight factors is essential to build a surrogate model capable of effectively guiding the optimiser toward the global maximum¹⁰.

Figure 5.4 depicts the evolution of the laser a_0 (red) and the injected charge q (blue) in the best simulation, after the optimisation procedure. The corresponding optimised density profile is represented by the grey-shaded area. Unlike the optimisations discussed in section 5.2, the limited gas density prevent self-focusing and electrons are injected through downramp injection after the first capillary. The longitudinal phase space of such electrons

¹⁰Having various optimisation parameters and a complex objective function to optimise for, a wide exploration of the search space is particularly important. A straightforward way to widen the exploration in a bayesian optimisation routine is to increase the number of “initial guesses” over which the surrogate model is built. A multiple of the number of the optimisation parameters (n_p) is usually preferred in order to have a suitable starting point for the optimiser. Generally $n = 5 * n_p$ is a safe choice, where n is the number of initial seeds.




Figure 5.6.: Electron spectrum at the end of the propagation. The characteristics of the energy peak and the total injected charge are highlighted in the text box.

at the injection instant (top panel) and at the end of the propagation (bottom panel) is depicted in figure 5.5 alongside the longitudinal electric field (\mathcal{E}_z). As shown in the first panel, injection happens in three consecutive bunches behind the laser pulse. However, during the propagation through the second stage (i.e. the second capillary), a single electron bunch is selected: the energies outside the desired range are suppressed, while the electrons in the target spectral density undergo efficient phase space rotation, as illustrated in the bottom panel. Such behavior demonstrates the excellent ability of the optimiser to navigate the search space and identify the best possible set of parameters to maximise the desired spectral density. The resulting monoenergetic electron spectrum at the end of the propagation is depicted in figure 5.6, showing a distinct peak at $E_{\text{peak}} = 2.48$ MeV, with a corresponding energy spread ΔE_{fwhm} of only 0.11 MeV and a divergence $(\theta_x, \theta_y) = (5.7^\circ, 3.4^\circ)$.




Figure 5.7.: Electron spectra resulting from a subsonic (left) and a supersonic (right) nozzle, applying the objective function of equation 5.2.

For comparison, figure 5.7 illustrates the outcome of applying the same objective function to the density profiles at the outlet of a subsonic and a supersonic nozzle, over the same

parameter space of section 5.2). While a well-defined objective function is crucial, it alone is not sufficient to achieve a monoenergetic spectrum. Nozzle geometries neither support controlled injection nor precise phase space manipulation, consequently, the “optimised” spectra still exhibit an exponentially decreasing trend, similar to the results shown in figure 5.2.

To further demonstrate the tunability of the VERA source, a second optimisation is conducted using the same initial parameters and search space (cf. table 5.2). The key difference in this optimization is the objective function f , designed to maximize the spectral density in the energy range $\varepsilon = [4, 5]$ MeV and, therefore, effectively aiming at doubling the energy of the spectral peak. The electron spectrum resulting from the best simulation of this second optimisation is presented in figure 5.8. It illustrates a promising electron beam which, aside from a low energy background, presents a monoenergetic peak at $E_{\text{peak}} = 4.49$ MeV with a remarkably similar relative energy spread with respect to the the previous optimisation, i.e. $\Delta E_{\text{fwhm}} = 0.21$ MeV¹¹.




Figure 5.8.: Electron spectrum at the end of the propagation for the higher energy optimisation, as described in the main text

However, the peak charge density in this second case is significantly lower. This is an expected outcome, as achieving similar percentage energy spread while doubling the electron energy is generally non-trivial, particularly when maintaining fixed driving laser parameters.

In terms of plasma source design, table 5.3 is of particular interest, as it outlines the optimal parameters for the two optimisations discussed in this section. In both cases, as expected, the optimizer generally favors a mismatch between the diameters of the two capillaries: a thinner, high density first stage facilitates downramp injection, while a wider second stage is responsible for acceleration and phase space rotation. Notably, the optimal diameter of the second capillary D_2 is at upper limit of the available range, suggesting that

¹¹The proposed objective function can efficiently select the beam energy, yet different alternatives could be tested in order to have a more effective background subtraction. For instance, the use of negative weights for the undesired energy intervals could lead to promising results in terms of background subtraction, while probably reducing the charge density peak.

an even longer second stage could be beneficial. On the other hand, it can be counterintuitive how the first capillary is not set at its minimum diameter. This is to be attributed to fluid dynamics reasons: an excessively narrow diameter would result in large boundary layers in the first capillary, which would negatively impact the density downramp and, eventually, the injection process. This highlights how a CFD-aided Bayesian optimisation approach can be extremely valuable in designing complex sources, ensuring that all aspects of the fluid dynamics and of the laser plasma interaction are optimally balanced. Furthermore, it is noteworthy that while the two optimisations target different electron energies, the optimal source geometry (represented by D_1 and D_2) remains remarkably similar. Conversely, the inlet pressures P_1 and P_2 differ significantly. This proves again the exceptional tuning capabilities of the novel microfluidic source proposed: by maintaining the same geometry and simply varying the input conditions (and the laser focal plane position), the same source can yield electron beams with distinctly different energies, yet both exhibit an energy spread of approximately 2 %.

Optimisation	D_1 [μm]	P_1 [Pa]	D_2 [μm]	P_2 [Pa]	z_f [μm]
$f \rightarrow 2-5$ MeV	43.5	44737	70	10293	9
$f \rightarrow 4-5$ MeV	44.8	54804	70	18209	2

Table 5.3.: Optimal parameters for the two optimisation cases presented in this section, where the centre of the geometry is placed at $z = 0$.

In conclusion, this section highlighted even more the crucial importance of accurately and flexibly modelling fluid dynamics in laser plasma acceleration. Bayesian optimisation (and machine learning algorithms in general) has already been proven to be exceptional for efficiently navigating the multitude of parameters that influence the laser-plasma interaction [60, 98]. However, the importance of the inclusion of CFD simulations in these algorithms cannot be overstated, given the major impact of gas dynamics on the acceleration process. Therefore, combining CFD and PIC simulations in a Bayesian optimisation framework, this section introduced a novel method for optimising the design plasma sources, aiming at specific characteristics of the accelerated electron bunch. This innovative approach, exemplified by the VERA source, holds the potential to significantly boost the performances of LPA, particularly in application-oriented scenarios where tailored sources can be adapted to meet different target properties.

Having discussed the capabilities of various plasma sources through simulations, the upcoming chapter will focus on the VERA experimental apparatus for electron acceleration. It will introduce the setup and evaluate the feasibility of using industrial, post-compressed Yb lasers for high-average-power electron acceleration. By addressing the most significant challenges and presenting new, encouraging results, it will lay the groundwork for using Yb-driven LPA in industrial, commercial and medical applications.

CHAPTER 6

Wakefields driven by industrial lasers

The previous chapters have provided an extensive overview of high-average-power LPA, particularly focusing on the possibility to use industrial Yb lasers as drivers. However, the limited pulse energy inherent in these systems can trigger significant ionization-induced effects, subsequently influencing the pulse propagation and the acceleration process. To experimentally investigate the origins of these effects and assess their impact, the upcoming chapter will be centered on the VERA experimental apparatus. The laser pulse and the plasma density profile will be characterised at the interaction point, while both CFD and PIC simulations will guide the discussion on the influence of ionisation-induced effects on the excited wakefields.

SECTION 6.1

The experimental setup

Being part of an electron acceleration experiment, the VERA target chamber has been built in a laboratory where radiation is allowed to be generated. The whole setup is hence divided in two different areas which, for simplicity, will be addressed as “laser laboratory” and “radiation laboratory”. The laser and most of its post-compression apparatus (presented in section 3.2.1) resides in the laser laboratory. The pulse is then transported through 35 meters of vacuum into the radiation laboratory, up to the interaction point. The transport line and the experimental setup in radiation laboratory has been designed and built from scratch during the first year of the project, including the experimental chambers, the diagnostics, the gas delivery system and most of the electronics needed during the operation.

A 2D schematic of the radiation laboratory setup is shown in figure 6.1. It is composed of two $50 \times 70 \text{ cm}^2$ vacuum chambers: a “compression” chamber, hosting the dispersion compensation line responsible for the final temporal compression of the laser pulse¹, and

¹As described in section 3.2.1 the VERA laser is composed of two cascaded post-compression stages. An usual procedure in high peak power compression schemes is to apply the last dispersion compensation step (and hence the final compression) only once the pulse is close to the target position. This approach minimises the pulse susceptibility to eventually detrimental nonlinear effects when it gets transported over long distances.

an “acceleration” chamber, where the interaction point is located.

At the entrance of the first chamber, an ILAO Star 50 deformable mirror (Imagine Optics, with 19 actuators and dynamic range above 20 μm) works in closed loop with a HASO4 First wavefront sensor and it is able to flatten the pulse wavefront to RMS values below 20 nm. After the adaptive optics, the laser is sent to a compressor composed of 4 pairs of dielectric chirped mirrors (DCMs), designed to compensate the group delay dispersion (GDD) accumulated during the second spectral broadening stage². Before the last DCM a magnifying telescope expands the beam from the initial size of $7.3 \times 5.1 \text{ mm}^2$ ($1/e^2$ radius) of a factor 1.875, allowing for a smaller focal spot at the interaction point. Such pulse is finally sent to the acceleration chamber.

Here, a 50 mm, 28.1 degrees, silver coated off-axis parabola (OAP) with $\lambda/8$ surface accuracy and 80 mm effective focal length (Shanghai Optics), typically focuses the beam to a $1/e^2$ spot size radius of $\sim 3.5 \mu\text{m}$ (cf. section 6.2). The focusing optic is mounted on a H-811 miniature hexapod (PI), allowing for 6-axes fine-tuning and high repeatability³. At the interaction point, a 100 μm -scale plasma source (described in section 6.3.2) can be controlled by 3-axis picomotor stages⁴. The source position, moreover, can be monitored and adjusted thanks to the imaging setup of the interferometric probe, used for the plasma density characterisation and described in section 6.3.1.

In the compression chamber a reflective attenuator can be driven in the laser path to enable the focal spot size optimisation and measurement. The focal spot imaging system (1.5 μm resolution) is composed of a 20x Mitutoyo objective, a tube lens and a series of folding mirrors directing the laser light to an Allied Vision Prosilica GT 1290 focal camera. Three Owis stepper motors are responsible for the movement of both the focal camera and of the out-coupling parabola, in a configuration which allows for a simple, fast interchange between the two, whose 3D rendering is shown in figure 6.2.

The electron diagnostics line closely follows. A Bergoz integrating current transformer (turboICT) with a detection limit of 50 fC is used for the electron beam charge measurement. The beam energy, on the other hand, can be monitored by an electron spectrometer composed of a 200mT permanent dipole magnet and a LANEX scintillating screen, optimised for 5 MeV⁵. A second screen can be added to measure the beam profile and divergence.

Besides the in-vacuum setup, a number of additional diagnostics have been built from scratch and are used to monitor and measure both the laser pulse and the laser-plasma interaction. At the entrance of the compression chamber the pulse is stabilised by a simple two-mirror nearfield-farfield beam stabilisation apparatus. Entering the acceleration chamber, a 7 mm-diameter pick-off mirror is slightly inserted in beam path sideways. Its reflection is used as probe for a standard Mach-Zehnder interferometer, serving both as plasma diagnostics and as plasma source positioning monitor (described in section 6.3.1).

²Each DCM reflection applies a negative GDD equivalent to 3 mm of fused silica. The fine-tuning of dispersion, on the other hand, is addressed in the laser laboratory, where a couple of silica wedges can be driven in the pulse to introduce the desired amount of in-glass propagation.

³The minimum incremental motion on x, y and z is respectively 0.2 μm , 0.2 μm and 0.08 μm , the minimum for their corresponding angles is 2.5 μrad , 2.5 μrad and 5 μrad .

⁴While the repeatability of such motors is limited, they offer an extremely compact solution.

⁵The detectable energy range is 1-10 MeV.

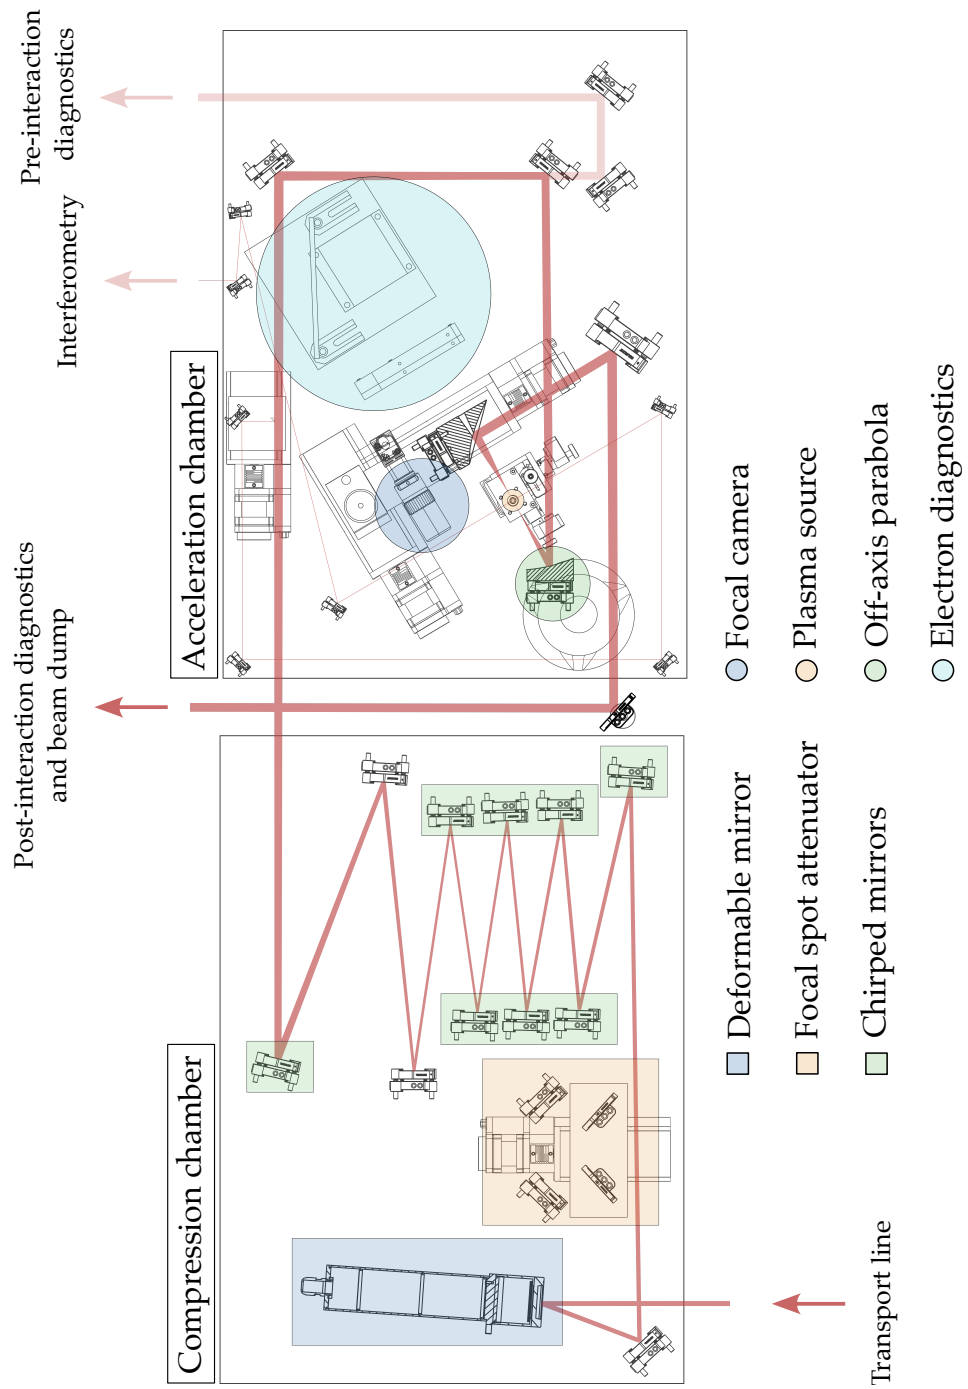


Figure 6.1.: Sketch of the compression and the acceleration chambers in the radiation laboratory.




Figure 6.2.: 3D rendering of the compact focal camera and out-coupling parabola configuration. Respectively isometric, side and front view.

Going towards the interaction point, a further stabilisation system sets the references and ensures stable positioning and pointing on the OAP. Finally, after the interaction, most of the laser energy gets dumped, while a wedge reflection guides the remaining part to the laser post-interaction diagnostics. Here an in-house built second-harmonic FROG is responsible for the pulse temporal reconstruction and an ultra-broad band spectrometer (790 - 1620 nm) monitors the laser post-interaction spectrum on shot.

Correctly measuring such broad spectra, extending to the infrared, is not straightforward. Silicon-based detectors, the most common in the field, can only detect wavelengths whose corresponding photon energy is larger than their bandgap energy (i.e. $\lambda_d < 1100$ nm). Therefore, usually, these are employed in combination with a shortwave infrared (SWIR) detector, typically Indium Gallium Arsenide (InGaAs)⁶. However a post-interaction spectrometer, devoted to the detection of spectral shifts induced by the plasma, would greatly benefit from streamlining the process, using a single detector with a uniform quantum efficiency in the desired wavelength range. To this end, an Andor Shamrock spectrograph has been combined with the Allied Vision Goldeye G-130 TEC1 camera. The camera sensor (Sony IMX990) is a typical InGaAs, yet it is built with an exceptionally thin protective layer, which allows visible wavelengths to reach the material and get detected. The sensor detection range, hence, extends from 400 to 1650nm, with an efficiency which varies only $\sim 10\%$ in the wavelength range of interest. On the other hand, resolving the small spectral variations due to the laser-plasma interaction (cf. section 6.4.2) can be challenging, as the sensor area is limited to only 6.4×5.1 mm² and this poses remarkable constraints on the spectrometer apparatus. To give more visibility to the spectral wings, the range around the central wavelength (1030 nm, with higher intensity) would need to be blocked. Notch filters, commonly used for these purposes, have limited bandwidth with a fixed central wavelength. Therefore, in order to block exactly the desired wavelength range, the spectrometer grating has been chosen in order to disperse the pulse over a length of 4 cm and a scattering screen with a mechanical filter is placed at the spectrometer exit, as shown in figure 6.3. The scattering plate is then imaged on the camera sensor with a spectral resolution of 0.7 nm,

⁶The spectra presented in section 3.2.1, for instance, are measured this way.

similar to what would have been obtained having the whole wavelength range dispersed on the length of the camera chip. As an example, the right panel of figure 6.3 shows a typical spectral image registered during the operation and its corresponding spectrum, once the proper calibration has been applied.




Figure 6.3.: Sketch of the post-interaction spectrometer built to monitor plasma-induced spectral changes. A typical camera image is shown on the right, together with the correspondent normalised spectrum.

Having introduced the experimental apparatus, the upcoming section is centered on the characterisation of the laser pulse at the focal position, crucial for the acceleration process.

SECTION 6.2

On-target pulse characterisation

The ultra-broadband, ultrafast laser system presented in this thesis is heavily relying on nonlinear interactions and aims at record compression factors. Therefore it can be particularly sensitive to even small spatio-temporal couplings, which can be detrimental during the acceleration process. Hence, it is crucial to characterise the pulse as close as possible to the interaction point.

The energy throughput from the output of the second MPC in the laser laboratory to the interaction point in the radiation laboratory was measured to be $69.2\% \pm 1.6\%$, theoretically enabling a total pulse energy of 4.6 mJ on target. However, extended high-power operation led to carbonization of the first MPC's mirror surface, significantly reducing the cell's throughput. Consequently, the effective pulse energy on target during the experimental campaign detailed in this chapter was limited to 3.7 mJ. Close to the interaction point, the laser can be diverted and measured in a SH-FROG. The corresponding typical temporal reconstruction is presented in figure 6.4.

The pulse is then focused to a spot size $w_0 \sim 3.5 \mu\text{m}$ ($1/e^2$ radius), as shown in figure 6.5. Its peak intensity at focus approaches $5 \times 10^{17} \text{ W/cm}^2$, computed integrating the light impinging on the focal camera and considering only the main peak of the temporal pulse reconstruction shown in figure 6.4, having a FWHM temporal duration of 14.7 fs and an




Figure 6.4.: FROG-retrieved pulse temporal structure at the interaction point. The main peak has a FWHM pulse duration of 14.7 fs a peak power of 125 GW.

energy of 2.22 mJ ⁷.




Figure 6.5.: Typical laser focus at the interaction point. The gaussian fit of the spot is highlighted in red, with $1/e^2$ radial dimensions of $3.5 \times 3.6 \mu\text{m}^2$. The pointing stability (computed as the standard deviation of the focus center of mass over 50 laser pulses) is $1.2\% \times 2.4\%$ of the spot size, respectively on x and y.

Around the focus, a high-dynamic-range focal scan allows to image the beam along the propagation axis for different exposure times, i.e. integrating over a different number of pulses. Figure 6.6 depicts some example measurements at positions separated by $54.5 \mu\text{m}$, where the camera exposure time has been scaled in order to have a similar contrast with

⁷The temporal structure of the pulse at the interaction point is significantly different from the structure measured in the laser laboratory and presented in section 3.2.1. While the specific cause of this change is still to be investigated, the issue is partially addressed in a preliminary study in appendix E.

respect to the background.




Figure 6.6.: High-dynamic-range focal scan imaging the pulse going towards the focal position.

The possibility to simultaneously scan both the camera position and exposure has two major advantages. First, as shown in figure 6.7, it allows to measure the size of the lower intensity contours, key to avoid severe damages to the plasma source. In the simple case of a nozzle made of glass, for instance, the radius of the 10^{14} W/cm² contour can be assumed as the absolute minimum distance that the laser axis can have from the jet [99]. In the figure, specifically, two different exposures are adopted (400 ms and 1 ms) to show such contour (left) compared to the the 1×10^{16} W/cm² intensity contour depicted on the right.




Figure 6.7.: Focal spot images taken with different exposures. The left image integrates 400 laser pulses and shows the 1×10^{14} W/cm² intensity contour. The right panel integrates a single pulse and outlines the 1×10^{16} W/cm² intensity contour.

On the other hand, these scans allow to estimate the spatial quality of the focusing beam, or its M^2 . In this context, figure 6.8 depicts the (gaussian fitted) size of the pulse while getting focused on both axis and compares them to an ideal pulse with $M^2 = 1$ ⁸. While standard beam quality measurements would require a much finer scan, this approach can

⁸Interestingly, both axis are getting focused on the same point ($z=0$), demonstrating an effective OAP optimisation routine.

provide a good measurement of the beam profile at the entrance on the plasma (hundreds of μm before the focal position) and a qualitative estimation of the effective Rayleigh length. Such measurements can be pivotal in high-average-power LPA, where the interplay of optical focusing, ionisation-induced defocusing (cf. 2.4.1) and relativistic self-focusing (cf. 2.4.2) can highly affect the beam propagation.




Figure 6.8.: Beam quality measurement along the focal scan. The focus position is set at $z = 0$. For each position 10 camera images were acquired. The corresponding measurement error is within the plot markers.

SECTION 6.3

Plasma source characterisation

Besides the laser pulse, knowing the (electron) plasma density during the interaction is key in the acceleration process, as it affects both the driver propagation and the accelerated beam properties. In this context, when the plasma source provides optical access, laser interferometry is often the method of choice. The following section, hence, will briefly introduce the fundamentals of the technique, emphasising the specific needs of the VERA setup. The typical source employed in the experimental campaign will be presented and characterised. Its interferometric measurements will be used to benchmark CFD simulations and a range of attainable densities and density profiles will be outlined.

6.3.1. Interferometry

In an interferometry setup a laser probe is guided through the plasma. Here, it acquires a phase shift which depends on the plasma refractive index η and, therefore, analysing the shift, the plasma density can be retrieved.

Generally, when two plane waves interfere at a shallow angle θ , they give rise to an interferometry pattern whose fringes are separated by $\lambda_f = \lambda_L/\theta$, where λ_L is the central laser

wavelength. The presence of plasma in one of the two paths will distort these fringes introducing a phase shift $\Delta\phi$ [100]:

$$\Delta\phi(x, y) = \frac{2\pi}{\lambda_L} \int [\eta(x, y, z) - 1] dz, \quad (6.1)$$

where z is the wave propagation direction. It is intuitive how, measuring the shift, it is possible to retrieve the plasma density⁹.

Figure 6.9 outlines a schematic of the Mach-Zehnder transverse interferometer set up in the radiation laboratory, where the probe is split after the passage through the plasma. In this configuration, a magnifying telescope images the plasma to the camera location, where a “clean” section of the probe, which did not interact with the plasma, can be used as reference, greatly simplifying the timing between the two interferometer arms.

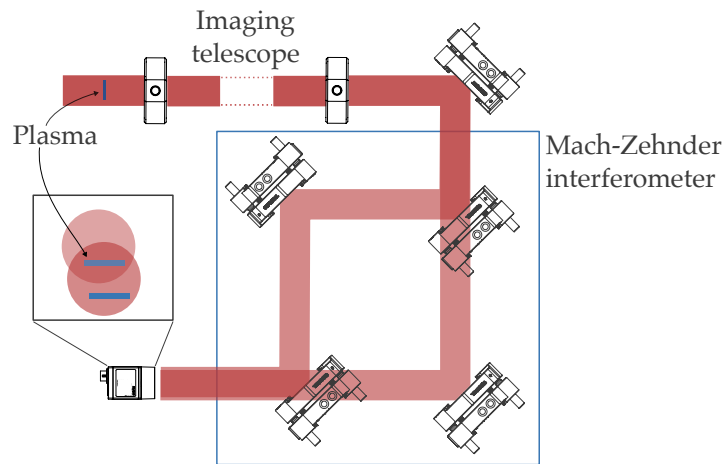


Figure 6.9.: Schematic of the transverse interferometry configuration. The inset on the left serves as an example of the typical camera image, with the two beams displaced on the chip.

A raw interferometric image is essentially a sinusoidal signal, hence its analysis typically takes place in the Fourier space. While different algorithms exist, each one with its own advantages and limitations¹⁰, this thesis follows the Fourier-based phase retrieval described in [101]. To start with, the left panel of figure 6.10 shows a typical raw image. The hard edges of the picture can easily give rise to unwanted complexity in the Fourier space therefore, most commonly, these images are filtered with a low amplitude modulation (“Hann-window”), which preserves the phase profile but smooths out the boundary (right panel).

⁹The density retrieval procedure described in this section assumes that the plasma refractive index can be expressed as a Taylor expansion of the plasma density, cut to the first order as: $\eta = 1 - n_e/2n_c$. While this assumption is often considered valid only for $n_e \ll n_c$, appendix E.1 shows how the estimation error is minimal even for densities of the order of $1 \times 10^{20} \text{ cm}^{-3}$.

¹⁰A selection of these algorithms can be found in [100]




Figure 6.10.: Typical interferometric raw image (left) shown alongside with the outcome of a Hann-filter (right).

The presence of the fringes provides a way to separate, in the reciprocal space, the information related to the phase shift $\Delta\phi$ from the general fluctuations in background intensity inherent to experimental data. Specifically, the data concerning the background is located around $(k_x, k_y) = (0, 0)$, whereas (for vertical fringes) the details regarding the plasma manifest in a pair of sidebands at $(k_x, k_y) = (\pm k \sin \theta, 0)$, where k_x and k_y are the axes of the Fourier space¹¹. In order to filter out the background, a mask is applied to the transformed image to isolate one of the sidebands, as shown in figure 6.11. The size of this Fourier-mask is crucial: ideally a larger mask enhances spatial resolution, however dealing with noisy data can pose challenges and a smaller option can help with noise reduction, particularly when high plasma densities are considered. The selected sideband is then re-centered in the reciprocal space and back-transformed. The argument of the resulting complex image is then considered to retrieve the phase¹².

While the calculated phase of the complex values is bounded between $[-\pi, \pi)$, a physical phase shift can easily exceed 2π . Consequently, an unwrapping algorithm is applied to remove the sharp phase transitions. However, as shown in the first panel of figure 6.12, the image obtained this way will still present a relatively strong phase background. This is generally due to imperfections in the imaging system which, for instance, can cause fringe bending or stretching and has to be subtracted. In many cases it is enough to register a reference interferometric image without the plasma and use its phase as a reference. Some experimental setups though, depending on the plasma size and the density to measure, can be particularly sensitive even to small pointing fluctuations and will end up having a constant phase background which can be difficult to quantify and can strongly vary from shot to shot. Therefore, figure 6.12 presents a more reliable background subtraction method, which allows for a faster analysis. Starting from the initial retrieved phase image, the region where the plasma is present is masked out. Then, a 2D third order polynomial fit is used to reconstruct the background in the masked region and the resulting image is used for the background subtraction. Finally, the $\Delta\phi$ due to the presence of plasma is obtained (last

¹¹Note that a different fringe angle would result in a contribution even on the k_y axis.

¹²Re-centering the masked sideband is one of the ways to deal with background subtraction. It is equivalent to consider a set of reference fringes separated by λ_f , without any fringe shift, and use it as a background image [100].




Figure 6.11.: Fourier peaks in the (k_x, k_y) space. A typical mask for the sideband selection in highlighted in red.

panel).




Figure 6.12.: Steps for the background subtraction method, as described in the main text.

It is of particular interest here to show how different Fourier-masks in the reciprocal space, will affect the density retrieved at the end of the process. Figure 6.13 shows the results of the analysis of the same raw interferometric image for masks of decreasing widths on the k_x axis. The first row shows the output of an excessively big window. In this case the mask is so big that it includes part of the main Fourier peak and, while the plasma is still visible in the corresponding phase image, the retrieval is completely misleading. The bottom row depicts the opposite scenario: a very narrow window effectively catches the plasma phase shift, but the noise reduction is so strong that filters out the actual plasma shape. Moreover, it leads to the appearance of a marked positive phase shift along the plasma axis, which




Figure 6.13.: Influence of the Fourier mask size on the retrieved phase. The same raw image is analysed with different Fourier windows, marked in red on the left column. The corresponding phase retrieval is shown on the right.




Figure 6.14.: Density retrieved from the inverse Abel transform of the phase image in 6.12. The top row shows the same retrieval with (right panel) and without (left panel) a gaussian filter with $\sigma = 2$ pixels. The corresponding density lineouts along the symmetry axis are depicted in the bottom panel, where the camera pixel calibration has been added.

would then correspond to unphysical negative electron densities after the retrieval process. On the other hand, the two central rows show the results of a more efficient masking, one emphasising a better spatial resolution, the other enhancing the noise reduction. Generally, it is difficult to find a “perfect” Fourier window and its size needs to be adapted to the specific raw image to analyse.

Yet, the result obtained is line-integrated, i.e. a 2D projection of a 3D object. However, assuming the plasma to be cylindrically symmetric, a symmetry axis can be found and inverse Abel transforms [102] can be exploited to reconstruct the actual plasma shape¹³. While often noise in the raw data can be managed in the Fourier space, for high plasma densities ($n_e \gtrsim 5 \times 10^{19} \text{ cm}^{-3}$) it can be difficult to find a good trade-off between resolution and noise filtration. A narrow gaussian filter (i.e. a convolution with a 2D gaussian function with standard deviation of few pixels) within the Abel process, hence, can help in retrieving more meaningful high-density data, while having no major impact for the lower and more easily retrievable densities, as shown in figure 6.14¹⁴. An alternative path for high density measurements would be to estimate the functional form of the density profile through CFD simulations and fit the retrieved phase shift or Abel-inverted image. However, this requires

¹³A number of different algorithms exist to perform inverse Abel transforms. This thesis employs the (symmetric) recursive Hansen-Law method [102], which has previously been proven to be both robust and computationally efficient.

¹⁴In this thesis, when employed, the gaussian filter is a 15×15 pixels kernel with a standard deviation of 2 pixels.

precise knowledge of both the source geometry and the laser parameters and an accurate benchmarking.

The resolution of the measurement, on the other hand, depends on the specific setup. The phase resolution, or the minimum retrievable density, correlates directly with the smallest observable phase shift. Intuitively, such shift would be of the size of a single camera pixel, therefore the phase resolution is essentially a function of the number of pixels across a single fringe, namely $\text{res}_\phi = 2\pi/N$, where N is the pixel count. Conversely, the fringes dimension usually corresponds to the spatial resolution¹⁵. Generally, hence, setting up an interferometer eventually results in a trade-off between phase (large fringes) and spatial resolution (narrow fringes) and it is heavily dependent on the specific measurement to perform. In the following, the VERA interferometer has been configured in order to have, in most cases, a phase resolution of ~ 60 mrad, resulting in a minimum retrievable plasma density as small as 10^{17} cm^{-3} . Such configuration allows relatively low densities to be measured with a good signal-to-noise ratio and makes the plasma phase-shift visible on the camera even at low pressure (key in order to monitor the laser position during the operation). On the other hand, the spatial resolution is reduced to $\sim 50 \mu\text{m}$. However, in nozzle-like density profiles, this will impact mainly on the ability to resolve the density ramps, whose functional form is known and can be inferred through CFD simulations.

6.3.2. Density characterisation and simulations benchmarking

Having introduced the general concept of plasma interferometry, it is crucial now to benchmark the fluid simulations with the retrieved density data, so that simulated density profiles could, later, guide the discussion on the complex, nonlinear laser-plasma interaction.

Given their historical success and simple implementation, gas nozzles have been initially selected as sources for the experimental campaign. The plasma source employed in all the following measurements, in particular, is shown in 6.15. It is a simple silica glass subsonic nozzle, modelled so as to minimise the pressure losses before the outlet while still providing a gas density suitable for high-average-power LPA experiment. A 2 mm-wide, 3 cm-long channel smoothly shrinks down to an outlet diameter on the order of $100 \mu\text{m}$ ¹⁶. Further refinement is then usually applied at the jet outlet in order to reduce the wall size of the jet to few multiples of the output hole, making it less prone to laser damage. As discussed previously, subsonic jets are sub-optimal for the application (cf. chapter 4). Yet they are easier to model, manufacture and test and they can provide the required density with a lower backing pressure¹⁷. Furthermore, in order to deal with the gas load and maintain the

¹⁵To be more precise the spatial resolution is usually a convolution of the fringe dimension and the optical resolution of the imaging setup of the interferometry. However an usual procedure is just to assume the largest as the actual resolution which, in many cases, coincides with the fringes resolution.

¹⁶A range of different outlets is available, spacing from 75 to 200 μm .

¹⁷More complex solutions, like supersonic nozzles, typically require a higher backing pressure to achieve the plasma density needed (10^{20} cm^{-3}) at the interaction point. The current valves of the gas delivering system (pulsed solenoid valves, Parker Series 9) are limited to a maximum backing pressure of only 51.7 bar. New valves, able to sustain pressures up to 200 bar, are already available and will be implemented in the next update of the experimental setup (cf. section 6.5).

vacuum levels in the target chamber below few 10^{-2} mbar, the gas is usually operated in burst mode, delivering pulses of few milliseconds duration.




Figure 6.15.: Schematic of the gas nozzle employed as plasma source.

The limited backing pressure plays a role also in the gas selection. While light gases are generally less prone to induce ionisation defocusing (cf. 2.4.1), high-Z gases can provide a remarkably higher plasma density for the same backing pressure. In this context, however, it is interesting to benchmark the measurements obtained in both cases in a density region in which defocusing has a minimal effect and full ionisation of the gas can be assumed. To this end, helium (doped with 1% of Argon) and nitrogen are considered in the following study.

For each density measurement a 3D scan of the laser (vacuum) focus position has been performed. In a single scanned point the interferometry analysis results in a 1D density profile. Here, then, the retrieved peak density is considered as the density value for that specific scan step. This way a vertical, a transverse and longitudinal scan can be gridded together to reconstruct a 3D map of the plasma density plume above the nozzle¹⁸.

As an example, figure 6.16 compares the density measurements with 25 bar of helium to the corresponding CFD simulation. The agreement is excellent. The top panel depicts the vertical evolution of the on-axis density with respect to the laser height above the nozzle (h), while the bottom panel represent the transverse density profile. The error in the density retrieval is computed as the standard deviation of the fringe analysis over 10 camera shots, where each shot integrates a single pulse of the driver laser. On the other hand, estimating the laser height above the nozzle and its correspondent error from the interferometric image is non-trivial, as the imaging plane of the probe is set on the plasma channel, rather than on the nozzle tip. For simplicity, the height of the point with the maximum fringe shift is considered. Its corresponding error is plotted as a horizontal bar in the vertical trend (top panel), while it results in a confidence interval in the simulations when considering the transverse profile (bottom panel).

On the other hand, figure 6.17 benchmarks the fluid simulations for molecular Nitrogen. Here a backing pressure of only 5 bar is considered and the height range has been moved further away from the nozzle in order to minimise the effect of ionisation defocusing on the driver laser. While being slightly offset, also in this case the agreement is remarkable.

¹⁸The procedure through which a density profile is retrieved is explained step by step in appendix E2.




Figure 6.16.: Comparison between simulated (blue) and retrieved (red) density data for 25 bar of Helium and the nozzle geometry showed in figure 6.15. The top panel shows the vertical peak density evolution with respect to the height above the jet. The bottom panel compares the transverse profiles at an estimated height of $147 \pm 6 \mu\text{m}$.

Addressing the small mismatch is challenging, especially when the results are compared to the ones obtained with helium. Some of the possible reasons could be found in:

- ◇ small modifications in the inner nozzle geometry. Especially if close to the nozzle output these can easily lead to few % differences and are hard to catch with simulations.
- ◇ Turbulence modeling. Simulating turbulent flows expanding in vacuum is generally a nontrivial problem and different flow models or set of parameters could be more suitable to gases with different molecular weights
- ◇ The Fourier mask in the interferogram analysis. The retrieved density is often a function of the width of the mask and different sizes can often result in a few % difference.
- ◇ Gas ionisation. In CFD, generally, a gas is simulated and the resulting density is then




Figure 6.17.: Comparison between simulated (blue) and retrieved (red) density data for 5 bar of molecular Nitrogen and the nozzle geometry showed in figure 6.15. The top panel shows the vertical, axial, density evolution with respect to the height above the jet. The bottom compares the transverse profiles at an estimated height of $215 \pm 6 \mu\text{m}$.

scaled up assuming a given ionisation level. For nitrogen 5 ionised electrons per atom are assumed. While the 5th ionisation level should comfortably be within the reach of the measured vacuum peak intensity, the interferometric measurement is line integrated and, therefore, the simulation can result in a slightly overestimated density.

- ◇ The spatial fringe resolution of the measurement. The nitrogen dataset analysed here has a fringe resolution of $\sim 50 \mu\text{m}$, lower than the helium case of figure 6.16 ($\sim 30 \mu\text{m}$). A Richardson-Lucy deconvolution algorithm [103] has been implemented and tested for this resolution and resulted in a deconvolved signal which can easily result in peak values with few % difference with respect to the observed maxima.

Nevertheless, the disparity between the retrieved and simulated data only amounts to 10-15%, an excellent outcome, particularly considering the inclusion of turbulent flow in the

simulations.

Having benchmarked the experimental data, it is instrumental now to use the outcome of the fluid simulations to estimate the (ideal) maximum available electron density with the current configuration.

With this aim, the minimal laser height above the jet to avoid severe damage to the source is assumed to be $100\ \mu\text{m}$ ¹⁹. At this height, the peak density for both gases has been extrapolated accounting for the exponential dependence shown in the measurements of the precedent section, while the width of the measured density profile has been scaled according to the simulations results. The outcome is depicted in figure 6.18, where a backing pressure of 30 bar is considered²⁰.

The range of available plasma densities appears promising, however, as the next section will discuss, when such long ramps are present ionisation induced defocusing can severely affect the interaction.




Figure 6.18.: Maximum experimentally obtainable plasma density for nitrogen (N_2) and helium (He), estimated as described in the main text. For both gases the backing pressure is assumed to be 30 bar and the height above the nozzle $100\ \mu\text{m}$. The red dotted line depicts the critical density needed for relativistic self-focusing of the VERA pulse (shown in figure 6.4).

¹⁹Estimated considering the $1 \times 10^{14}\ \text{W}/\text{cm}^2$ intensity contour at the focal spot, measured in figure 6.7 and adding a “safety” factor of 2.

²⁰Experimentally, 30 bar was found to be the maximum pressure available for a steady and repeatable operation of the high pressure solenoid valve currently installed before the plasma source. This limitation was primarily due to the rigidity of the valve mount, which, under higher backing pressures, exhibited tens-of-micrometer-scale movements that hindered source operation. A new, more robust, mounting option has already been implemented and will be tested soon (cf. section 6.5).

Ionisation-induced effects

To excite a nonlinear wakefield with the pulse presented in section 6.2, relativistic self-focusing in the plasma is crucial. The onset of this phenomenon, for the current VERA pulse, would require plasma densities above $1.4 \times 10^{20} \text{ cm}^{-3}$ which, as described in the previous section, is within the reach of the plasma source adopted, especially when pure nitrogen is adopted. However, the experimentally measured plasma density rarely exceeds $6 \times 10^{19} \text{ cm}^{-3}$ and, consequently, the excited wakefield is limited. In the following section, the results achieved experimentally with high nitrogen backing pressure will serve as an example to investigate the reason behind such limitations, while still demonstrating the first evidence of a linear plasma wakefield excited by an industrial, post-compressed Yb laser.

6.4.1. Density clamping

The limited pulse energy available may easily result in non-negligible ionisation-induced defocusing [47, 49]. Because of defocusing, the laser peak intensity gets reduced and its position moves earlier in the propagation. The lower effective intensity can ultimately limit the amount of ionised electrons, resulting in the density being clamped, independent of the gas backing pressure. Moreover the longitudinal position of the plasma density peak tends to move upstream with increasing density, following the evolution of the laser spot size [104]. Therefore, tracking the plasma density profile under different initial conditions can provide valuable insights into the laser-plasma interaction.




Figure 6.19.: Peak density (n_e , left panel) and corresponding longitudinal position of the density peak ($z_{n_{max}}$, right panel) for different heights (h) above the plasma source. Two different nozzle backing pressures are considered, 20 and 30 bar.

Experimentally, a behavior similar to the one just mentioned becomes particularly evident with high nitrogen backing pressures. As an example, the left panel of figure 6.19 illustrates how the plasma density evolves as a function of the laser height above the nozzle (h)

for two different nitrogen backing pressures, 20 and 30 bar. The right panel, on the other hand, shows the corresponding longitudinal position (along the laser propagation axis) of the plasma density peak, $z_{n_{\max}}$. The two panels show an opposite trend, where the position of the density peak moves upstream for increasing density. Furthermore, considering a single laser height, the higher backing pressure (30 bar) does not result in a proportional density increase with respect to the lower case (20 bar), as fluid simulations would suggest (cf. section 4.2).

The same correlation can be found analysing transverse density scans. Following the same scheme, the left panel of figure 6.20 depicts how the peak density value evolves along the scan for the two backing pressures considered previously. Conversely, the right pane presents the corresponding longitudinal position of the density peak²¹. Analogously, a higher density results in an earlier density peak²².




Figure 6.20.: Typical peak density (n_e , left panel) and corresponding longitudinal position of the density peak ($z_{n_{\max}}$, right panel) along a transverse scan. Two different nozzle backing pressures are considered, 20 and 30 bar. Where not visible, the error is smaller than the marker size.

These trends strongly point toward relevant ionisation-induced defocusing, which can have a significant impact on the laser-plasma interaction. In fact, a reduced peak intensity during the propagation essentially restricts the achievable a_0 and the presence of a maximum obtainable density ultimately hinders the onset of relativistic self-focusing. Consequently the excited wakefield will be limited. In the next section, hence, a more quantitative analysis of these limiting factors is presented, guided by measurements of the laser

²¹Note that both figure 6.19 and 6.20 are illustrating a relative change in $z_{n_{\max}}$, emphasising the magnitude of the shift during a given scan, rather than the absolute longitudinal position of the plasma density peak.

²²Determining the position of the plasma peak from an interferometry measurement is not always trivial as, depending on the raw image and on the amount of phase shift, oscillations from the Fourier space can affect the plasma density shape. However CFD simulations suggest that subsonic nozzles (like the one used in this experimental campaign) give rise to bell-like profiles, which tend to a gaussian at heights above the source similar to the ones discussed here. Hence, to estimate the plasma peak density position along the propagation axis in each case a gaussian fit over the retrieved density profile has been used.

post-interaction spectrum and supported by detailed FBPIC simulations.

6.4.2. Peak-intensity-reduction effects on the excited wakefields

The variations in the driver post-interaction spectrum can be extremely valuable to quantify the relative relevance of gas ionisation in the laser propagation and in the wakefield excitation.

During the interaction, as described in section 2.4.3, the laser spectrum can be subjected to a series of shifts and modulations. When low energies are considered, as in the case under analysis, such spectral changes mainly arise either because of gas ionisation (ionisation-induced blueshift) or as a result of plasma density modulations.

Figure 6.21 shows the typical spectrum registered in the experimental campaign with the ultra-broad band spectrometer described in section 6.1. As already mentioned, a filter placed along the camera imaging setup blocks the most intense part of the spectrum, around the central wavelength, between 980 and 1100 nm, in order to give more visibility to the spectral wings. Moreover, a 850 nm long-pass filter (shown in red in the figure) is placed at the spectrometer entrance to block short wavelengths, whose second order grating reflection would otherwise appear in the camera image and be indiscernible from the real spectrum²³.




Figure 6.21.: Typical spectrum registered during the experiment, with a detection range extending from 790 to 1620nm. The wavelengths from 980 to 1100 nm are filtered out and a long-pass 850 nm filter is placed at the entrance of the spectrometer (red dashed line). The red shaded area indicates the section considered to estimate the spectral redshift ($\lambda > 1200$ nm).

In the following, the spectral content in the region where $\lambda > 1200$ nm (highlighted by shaded area in figure 6.21) is considered. Specifically, the variations in its energy fraction are analysed, defined as the integral under the section divided by the integral under the

²³e.g. the second order reflection of $\lambda = 700$ nm would appear as $\lambda = 1400$ nm.

whole spectral curve:

$$\frac{E_r}{E_{\text{tot}}} = \frac{\int_{1200}^{1620} I_s d\lambda}{\int_{850}^{1620} I_s d\lambda}. \quad (6.2)$$

Such definition allows quantifying the red-shift resulting from the nonlinear interaction without resorting to a “mean wavelength”, which can be difficult to trace with the filtering applied to the spectrum and can be misleading when ionisation-induced blueshift plays a non-negligible role.

In this context, figure 6.22 represents the typical spectral energy variation E_r/E_{tot} obtained scanning the laser (vacuum) focus along the plasma source²⁴. The corresponding retrieved density is shown in grey. It is evident how the redshift (or the fractional amount of energy included in the post-interaction spectrum for $\lambda > 1200$ nm) follows a trend opposite to the measured density.

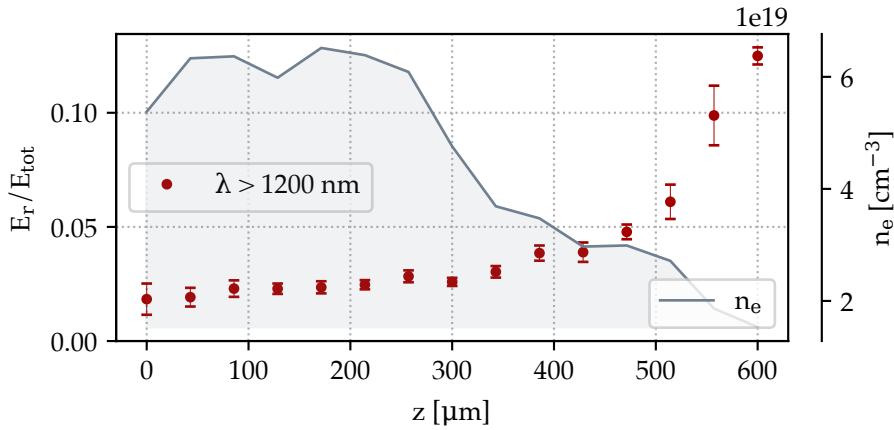


Figure 6.22.: Typical redshift trend (red) compared to the measured peak density (gray) along an exemplary scan, with backing pressure $P = 20$ bar and measured pulse energy on target $E = 3.7$ mJ.

To visualise the spectral shift, figure 6.23 compares the spectrum with the maximum measured redshift (red) to the initial laser spectrum (blue), registered without the plasma. The strong blueshift on the left side is attributed to gas ionisation, while the presence of two small additional peaks (around 1250 and 1330 nm) are a signature of a modest redshift and, therefore, an indication that a plasma wake has been excited during the interaction. Interestingly, such redshift corresponds to the last step of the longitudinal scan shown in figure 6.22, where the measured peak density is only $1.5 \times 10^{19} \text{ cm}^{-3}$. Increasing the plasma density, the fraction of spectral energy present in the red part of the spectrum rapidly decreases,

²⁴Here a longitudinal scan of the focus position is considered where the source backing pressure is set to 20 bar. However similar trends have been found for different backing pressures and scanning directions.

again suggesting that gas ionisation acts as a limiting factor: a larger amount of ionised gas shifts the overall spectrum towards lower wavelengths and restricts the achievable effective peak intensity (and consequently the peak a_0).




Figure 6.23.: The spectrum with the maximum registered redshift of the scan showed in figure 6.22 (red) is compared to the initial laser spectrum (blue)

Due to the complexity of the interaction, PIC simulations serve as a valuable tool for a deeper analysis, investigating the origin of the measured spectral modifications. The code FBPIC, introduced in section 5.1.1, allows the driver pulse to be simulated as a superposition of a radial and a longitudinal profile. In the following simulations, hence, while a perfect gaussian radial profile is considered, the longitudinal laser profile has been initialised from the measured spectrum and the FROG-retrieved spectral phase, shown in figure 6.4.

To analyse the redshift dependence on both the peak density and the pulse energy, a simple grid scan is presented. In each simulation, the density profile considered is taken from a real interferometry measurement. However, in order to account for gas ionisation, the PIC simulations are initialised with gas (rather than electron) density. The measured electron density, hence, is scaled down to its corresponding atomic gas density and the ADK ionisation model [94] (already included in the code) is exploited to take into account gas ionisation²⁵. The vacuum laser focus is set to be at the position of the gas density peak.

For each simulation the laser spectrum at the end of the plasma density profile is extracted and analysed applying the same filtering as for the experimental conditions and computing the spectral redshift according to equation 6.2. Figure 6.24 outlines the results of the grid scan for the total pulse energy used during the experimental campaign, 3.7mJ, the maximum energy available on target in the case of ideal transmission of the laser post-compression apparatus, 4.6mJ, and a lower energy, 2.8mJ, chosen for symmetry. While the

²⁵To pass from electron to atomic density (and viceversa) it is necessary to assume a given number of ionisation degree. Knowing the ionisation energy levels of different gases, this is essentially a function of the laser intensity. In the current case, 5 ionised electrons are assumed per nitrogen atom. In reality, this is not completely true, especially when the pulse is far from its focal position. However the (vacuum) peak intensity of the VERA pulse is at least an order of magnitude above the fifth ionisation level of nitrogen ($\sim 1.46 \times 10^{16}$ W/cm²) and, therefore, the assumption holds.

highest energy shows a general increase of the amount of redshift with increasing peak density, lower energies depict a trend similar to the one found in the experiment, where a redshift peak is reached around $1 - 1.5 \times 10^{19} \text{ cm}^{-3}$.




Figure 6.24.: FBPIC grid scan showing the redshift dependence on the peak density and on the total pulse energy. Grid points related to the same energy are connected as a guide for the eye.

Analysing the simulations it is evident how, in these latter cases, the decreasing trend for high densities is to be attributed mostly to gas ionisation. Given the limited pulse peak intensity, a higher gas density generally results in a larger fraction of the pulse affected by ionisation-related effects. This gives rise both to a severe spectral shift towards higher frequencies and to non-negligible ionisation induced defocusing, limiting the maximum a_0 available and, therefore, the amplitude of the excited plasma wake.

To picture the correlations between time and frequency, figure 6.25 represents the Wigner distributions at the beginning (left column) and at the end (right column) of the propagation for two simulations where the total pulse energy is 3.7 mJ and the peak densities are respectively 1.5×10^{19} (top row) and $4.8 \times 10^{19} \text{ W/cm}^2$ (bottom row)²⁶. Complex pulses, as in VERA, always give rise to intricate Wigner distributions (cf. appendix B), however focusing on the main temporal peak, where most of the Wigner intensity resides, it is possible to understand how its spectral content evolves during the propagation and to capture the underlying physics.

In the lower density case (top row) the immediate pulse front is responsible for gas ionisation, experiencing a small shift towards higher frequencies. The most intense part of the main peak, though, remains unaffected and keeps enough energy to drive a linear plasma wake. The wake modulation then results in a redshift just before the pulse peak, visible as a long tail toward lower frequencies around $t = 0$. In the second scenario, the most intense part of the pulse shifts as a whole towards higher frequencies, leaving a much reduced fraction on the red side of the spectrum. Furthermore, as shown in figure 6.26, the higher

²⁶According to the usual convention, the pulse front is at negative times, hence the pulses in the figure are propagating from right to left.




Figure 6.25.: Wigner distributions of the simulated VERA pulse at the beginning (left column) and the end (right column) of the propagation for two different peak densities: 1.5×10^{19} (top row) and $4.8 \times 10^{19} \text{ cm}^{-3}$ (bottom row). The total pulse energy is 3.7 mJ].

density results in a stronger ionisation induced defocusing, a lower peak intensity (or a_0) and, therefore, a more extended region where ionisation induced blueshift is dominant²⁷.

It is particularly interesting, now, to compare the measured maximum redshift to the corresponding simulation. Figure 6.27 compares the measured (left) and the simulated (right) data, with the same applied filtering, showing the initial and the post-interaction spectra respectively in blue and red. The agreement is very good. The small bump visible experimentally around 1250 nm appears as an evident red peak in the simulation, with a secondary satellite extending up to 1400 nm which exactly resembles the experiment. The simulation, moreover, results in a clear blueshift of the overall spectrum, which matches the experimental one for $\lambda < 980$ nm.

Relatively small differences are obviously still present. The initial laser spectrum measured experimentally with the broadband spectrometer displays a slightly different inten-

²⁷The spot size evolution in the simulations is computed through a gaussian fit performed at the longitudinal position of the maximum pulse intensity, for each simulation step. The higher density affects the pulse temporal shape after the focus, impacting on the way the spot size (and hence the a_0) is computed. For this reason the high-density curves in figure 6.26 show a less smooth behaviour for $z > 0$.




Figure 6.26.: Laser waist and a_0 evolution during the pulse propagation for the two simulations considered in figure 6.25. For reference, the corresponding evolution in vacuum is represented by the dotted grey line.




Figure 6.27.: Measured (left) and simulated (right) spectral evolution. In both cases the initial spectrum (blue) is compared to the spectrum after the interaction (red).

sity distribution compared to the initial spectrum used in the simulation, although their total bandwidth is the same (blue lineouts in figure 6.27). This can often be caused by minimal day-to-day changes in the spectral broadening stage of the post-compression apparatus, which can be accentuated by the suboptimal resolution of the spectrometer²⁸. However, the biggest deviation from the experimental conditions in the simulation is the radial profile of the laser pulse, which is assumed to be a perfect gaussian. This ideal radial shape most probably results in a more uniform blueshift in the pulse and in a slightly higher peak intensity during the interaction, leading to more evident red peaks in the spectrum. The possibility to simulate the measured radial profile is currently under implementation in FBPIC and it is expected to significantly enhance the fidelity of simulations (cf. appendix G).

Nevertheless, the good correspondence between the experimental and simulated data suggests that, while gas ionisation is limiting the interaction, a wake has been excited in the plasma. The PIC simulation shows that the excited plasma density oscillation amounts to $\delta n/n_0 = 0.21$ and the corresponding longitudinal (linear) accelerating field exceeds 30 GV/m.

As for the current scientific literature, this is the first ever evidence of a plasma wakefield being excited by an industrial, Yb-based laser system. Given the high-average-power scalability of Yb-based technology and the multi-kHz regime at which these systems can be operated, this finding represents an extremely promising step: it paves the way towards a more cost-effective, industrial solution for laser-plasma acceleration, potentially transformative to bridge the gap between laboratory experiments and practical applications.

SECTION 6.5

Towards the first LPA driven by industrial Yb lasers

In the previous section first, promising evidences of a plasma wakefield excited by industrial, post-compressed Yb:YAG lasers have been presented. In addition to this, ionisation-induced defocusing has been identified as the main limiting factor for the energy transfer from the driver pulse to the plasma, ultimately hindering electron injection. Therefore, effectively addressing this issue is of paramount importance for achieving successful electron injection. This concluding section, hence, will detail the planned upgrades to the VERA experimental apparatus, specifically designed to mitigate ionisation-induced defocusing and enhance the energy coupling between the driver and the plasma.

Increasing the peak intensity at the interaction point is crucial. As already mentioned, when high compression factors are considered, the laser pulse can be particularly sensitive to even small spatio-temporal couplings. These can develop during long propagation distances and reduce the achievable intensity at focus. Based on the preliminary wavelength-dependent measurements presented in appendix E, a comprehensive pulse characterisation on target is planned, to identify potential couplings and optimise the pulse compression. In addition to the thorough pulse characterisation, a more straightforward approach to raise the peak intensity involves reducing the focal length of the focusing optic to reach

²⁸This is a drawback due to the ultrabroad bandwidth to be measured.

a smaller spot size. A shorter Rayleigh length, in fact, can balance the relative influence of ionisation defocusing. Therefore, the VERA acceleration chamber is being modified to accommodate a new off-axis parabola with a focal length of only 5.08 cm, significantly shorter than the previous (8 cm). Such modification will lead to a proportional reduction in the (vacuum) spot size, ideally resulting in a $1/e^2$ radius of 2.2 μm . Consequently, the corresponding peak intensity is expected to raise of a factor 2.5, reaching $I_{\text{peak}} > 1 \times 10^{18} \text{ W/cm}^2$.

A second approach is to modify the density profile, as reducing the length of the density gradients also positively affects the energy coupling efficiency. In this context, supersonic nozzles can be an effective, simple to implement solution. Yet, these jets have to be operated with a significantly higher backing pressure to achieve the required density. Therefore, the gas delivering system, currently limited to a maximum backing pressure of approximately 50 bar, is being upgraded with a new solenoid valve able to withstand up to 206 bar (Peterpaul, series 20, model H22). Additionally, a more rigid source mounting structure has been implemented, to enhance the source's stability under high pressure. In particular, supersonic nozzles designed according to the optimal geometry discussed in section 4.2.2 have been produced (through selective laser etching, at LightFab) and are ready to be tested in the improved experimental setup. The design, illustrated in figure 6.28, is etched from a single silica cube with a 1 cm side. The bottom of the source allows for easy mounting at the end of a 6 mm Swagelok pipe. On the other hand, the jet tip has a wall thickness of only 10 μm , to minimise potential damages from the laser low-intensity wings during the experiment.




Figure 6.28.: Microscope image of a supersonic nozzle designed according to the optimal geometry discussed in section 4.2.2, shown alongside with a 3D rendering of the design and a close up sketch of the De Laval structure at the nozzle tip.

To illustrate the impact of these modifications, two separate 2D grid scans have been conducted in FBPIC to locate and compare the parameter space where electron injection is feasible, in the current and in the “upgraded” setup. The scanned parameters are the peak nitrogen density, ranging from 1×10^{19} to $5 \times 10^{19} \text{ cm}^{-3}$, and the focal plane position along the propagation axis z_f , scanned from -125 μm to +125 μm , with $z = 0$ being the position of the density peak. The grid step sizes are $4 \times 10^{18} \text{ cm}^{-3}$ and 25 μm respectively, resulting in a total of 121 simulations for each scan. The grid scan related to the current experimen-

tal apparatus assumes the on-target pulse characterisation outlined in section 6.2 and the subsonic nozzle used in the experimental campaign, described in section 6.3.2. Whereas, the upgraded setup assumes an ideal transmission of both multi-pass cells, the supersonic nozzle depicted in figure 6.28 as plasma source and the reduced laser vacuum spot size, estimated according to the newly implemented focusing optics. Figure 6.29 presents the results of these scans, highlighting the regions where electron injection is possible for each setup. Since the widths of the subsonic and supersonic density profiles differ significantly, the plot shows the dependence on z_f/w_D , where w_D is the $1/e^2$ radial width of the considered profile ($\sim 109 \mu\text{m}$ for the subsonic jet and $\sim 65 \mu\text{m}$ for the supersonic one²⁹). As shown in figure 6.29, the area over which injection is possible with the new setup is significantly larger, spanning a much broader range for the focal plane position. Additionally, while with the current apparatus injection is possible only for (peak) densities above $n_e = 1 \times 10^{20}$, the results of the scans suggests that with the upgrade in place injection will be generally feasible even with lower densities, owing to the significantly higher peak intensity. This outcome is extremely encouraging, indicating that the new additions can effectively tackle the ionisation-induced effects and will substantially increase the probability of electron injection.




Figure 6.29.: FBPIC grid scan identifying the regions where electron injection is possible for the current setup and the ongoing setup upgrade, as described in the main text.

Therefore, the upcoming implementation of the upgraded VERA setup promises to reach a significant milestone in high-average-power LPA, showcasing the first electron acceleration driven by an industrial Yb:YAG laser. Moreover, with the laser's repetition rate extending into the multi-kHz regime, it will enable the generation of electron currents an order of magnitude higher than current records [14]. The proposed novel microfluidic source, the VERA source, will further this progress by operating at reduced plasma densities. By exploiting downramp injection, it will alleviate the challenges posed by ionisation-defocusing and lower the demands for both self-focusing and high peak density, thereby also minimising the gas load in the vacuum chamber. With the novel tools introduced, the encouraging

²⁹A comparison of these profiles can be found in figure 4.9.

results already obtained and the future upgrades proposed, the VERA project promises to substantially enhance the stability, tunability and repetition rate of the accelerated beam. Ultimately, it will mark a major step towards bridging the gap between laboratory experiments and industrial applications, where high current, reliability, and efficiency are essential.

CHAPTER 7

Conclusions and future perspectives

In recent years, LPAs have consistently progressed. Besides holding the promise of a table-top accelerator, these machines stand out for their unique beam characteristics, which can be transformative in various application-oriented scenarios. For instance, their ability to accelerate femtosecond electron bunches, intrinsically synchronised with the driving laser, can offer ultra-high temporal resolution at the atomic scale in ultrafast electron diffraction (UED) experiments [19]. Furthermore, in medicine, the high peak electron flux of LPAs is extremely promising for FLASH radiotherapy (RT), which is showing significant benefits over conventional RT in terms of DNA damage and response, due to ultra-high dose rate irradiation [18]. To date, high-average-power LPAs driven by Titanium:Sapphire lasers have separately achieved an accelerated charge as high as 24 pC/shot, energies extending to 15 MeV and beam divergence < 10 mrad at a maximum repetition rate of 1 kHz [14, 16, 69]. However, despite the substantial progress that has been made, transitioning from laboratory experiments to practical industrial and medical applications requires substantial improvements in beam stability, reliability, repetition rate, and cost-effectiveness. For all these points, the introduction of industrial-quality lasers into the realm of laser plasma acceleration can represent a major breakthrough. These lasers set the stage for the development of more energy-efficient, stable few-MeV electron sources that can operate at repetition rates significantly higher than current records.

The results presented in this thesis have explored the feasibility of using industrial-quality, high-average-power, post-compressed Ytterbium:Yttrium-aluminium-garnet (Yb:YAG) lasers as drivers for LPA.

Achieving the required intensities for the acceleration process with industrial Yb systems necessitates a substantial temporal compression. In this context, multi-pass cells (MPC) are the optimal tool, enabling large compression factors without strictly requiring a proportional increase in the size of the setup. The Yb laser employed in this work, specifically, makes use of a double-MPC apparatus, which has demonstrated the highest temporal compression factor to date with mJ-class lasers [79].

Extensive simulations and experimental investigations have been conducted to assess the potential of using the compressed laser output to drive an LPA. The limited available pulse peak power highlighted relativistic self-focusing as the main condition to be satisfied

during the interaction. This generally requires high plasma densities, in excess of $1 \times 10^{20} \text{ cm}^{-3}$. Consequently, pure nitrogen has been identified as an appealing gas choice due to its ability to provide ten plasma electrons per molecule and the possibility to be efficiently removed by the vacuum pump, especially when compared to other, lighter gases. Nonetheless, the multiple ionisation levels of nitrogen can result in strong ionisation-induced defocusing, particularly when using sub-TW lasers as drivers. This defocusing can substantially reduce the pulse peak intensity during interaction and, therefore, must be efficiently mitigated through stronger optical focusing or advanced plasma density tailoring.

In this context, this thesis presented an extensive analysis of the most common plasma sources for high-average-power LPA and proposed a novel, microfluidic, multi-stage source, which can efficiently decouple the injection from the acceleration process. This source, potentially able to produce mono-energetic electron beams with an absolute energy spread lower than 0.2 MeV, could significantly enhance the tunability of LPAs, offering substantial benefits for arguably all the potential applications. The analysis was driven by a pioneering tool for the optimisation of plasma sources, which combines computational fluid dynamics and particle-in-cell simulations in a Bayesian optimization loop. This approach results particularly advantageous when dealing with multiple variables and is fundamental to explore the parameter space of both simple and complex sources, in order to optimise specific characteristics of the accelerated electron beam. Specifically, Bayesian optimisation has been employed to optimise the geometry of the novel source proposed, aiming at mono-energetic electron spectra of different energies. The source, when combined to the newly developed optimisation tool, demonstrated exceptional and unique tuning ability, especially in the few-MeV range.

Finally, this work marks a significant milestone by experimentally demonstrating the first ever plasma wakefield driven by an industrial Yb laser and provided a concrete path for improvements of the current setup that will soon enable the injection and acceleration of the first electron beam.

The achievements obtained thus far hold great promise, yet, looking ahead, numerous tasks and challenges remain to be addressed.

Section 3.2.1 discussed how post-compressed pulses always exhibit minor pre and post-pulses. While the highest temporal contrast is typically considered optimal, a small, correctly timed prepulse could preionise the plasma and enhance the energy coupling [105]. Investigating this effect, starting from simulations, would be highly valuable. Therefore, incorporating MPC post-compression into a Bayesian optimization algorithms could represent a major step forward, identifying the optimal temporal structure and post-compression settings to achieve specific electron beam properties. Moreover, given the significant compression factors discussed in this thesis, a comprehensive spatio-temporal pulse characterization at the interaction point is crucial. Building on the preliminary analysis in appendix E, a detailed study of the spatio-temporal couplings and their sources could greatly advance the understanding of the dynamics at the interaction point.

Furthermore, employing more advanced source structures will be instrumental in controlling the injected electrons' characteristics. The planned system upgrade (cf. section 6.5) already includes the use of supersonic gas nozzles, yet an extensive testing and refinement the novel microfluidic source proposed could result in superior electron beam properties.

In this context, implementing Bayesian optimisation in the real experiment can uncover new possibilities. It can aid both in balancing the two compression stages for specific temporal pulse characteristics and optimising plasma source operation during experimental campaigns, eventually leading to a more autonomous system.

Lastly, it is evident from the experimental work presented that future advancements in high-average-power LPA will need to address the critical issue of a limited pulse energy. In this context, the rapid development of industrial Yb:YAG technology towards compact systems capable of delivering increasingly higher pulse energies is particularly promising and motivates even more this thesis work. As an example, the VERA laser will soon be subjected to a major upgrade (Amphos, model 3257): compared to the current system, the laser will maintain similar footprint and average power (~230 W), but it will reduce the maximum repetition rate (12.5 kHz)¹ to increase the output pulse energy, nearly doubling it (17 mJ)². A higher available energy will significantly streamline the overall experimental process:

- ◇ a more modest temporal compression factor would be required to achieve the needed pulse peak power,
- ◇ the “excess” energy would allow for additional filtering during the pulse compression, generally resulting in a smoother spectrum and a more controllable pulse spectral phase [106, 107] and
- ◇ the necessary peak intensity at the interaction point could be achieved with less stringent focusing requirements. Consequently, the dependence on self-focusing would be considerably reduced, improving the control over the accelerated beam and enabling longer acceleration distances and better phase space manipulation.

In conclusion, the adoption of industrial Yb lasers in laser-plasma acceleration holds the potential to significantly enhance the machine efficiency, reduce the costs, and greatly increase the accessibility. Coupled with the exceptional energy scalability of MPC-based post-compression, this technology paves the way for the widespread use of compact few-MeV electron sources across universities, research facilities, and industries. With multi-kHz repetition rates, extremely short electron bunch durations, and high electron flux, these systems could provide substantial advancements in fields such as crystallography, medical sciences, and radiography, promising to unveil new insights into fundamental physics and delivering practical applications that benefit society.

¹The repetition rate would still be more than an order of magnitude higher than the current record with Titanium:Sapphire (Ti:Sa) laser technology.

²For comparison, table 3.1 indicates the current output parameters of the VERA laser.

Bibliography

- [1] Enzo F. Haussecker and Alexander W. Chao. The influence of accelerator science on physics research. *Physics in Perspective*, 13(2):146–160, May 2011. doi:10.1007/s00016-010-0049-y.
- [2] International atomic energy agency, accelerators knowledge portal. URL: <https://nucleus.iaea.org/sites/accelerators/Pages/default.aspx>.
- [3] Novella Grassi. Differential and scanning-mode external pixe for the analysis of the painting ritratto trivulzio by antonello da messina. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 267(5):825–831, March 2009. doi:10.1016/j.nimb.2008.12.018.
- [4] T. Tajima and J. M. Dawson. Laser electron accelerator. *Physical Review Letters*, 43(4):267–270, July 1979. doi:10.1103/PhysRevLett.43.267.
- [5] Felicie Albert, M E Couprie, Alexander Debus, Mike C Downer, Jerome Faure, Alessandro Flacco, Leonida A Gizzi, Thomas Grismayer, Axel Huebl, Chan Joshi, M Labat, Wim P Leemans, Andreas R Maier, Stuart P D Mangles, Paul Mason, Francois Mathieu, Patric Muggli, Mamiko Nishiuchi, Jens Osterhoff, P P Rajeev, Ulrich Schramm, Jorg Schreiber, Alec G R Thomas, Jean-Luc Vay, Marija Vranic, and Karl Zeil. 2020 roadmap on plasma accelerators. *New Journal of Physics*, 23(3):031101, March 2021. doi:10.1088/1367-2630/abcc62.
- [6] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick. Monoenergetic beams of relativistic electrons from intense laser–plasma interactions. *Nature*, 431(7008):535–538, sep 2004. doi:doi:10.1038/nature02939.
- [7] C. G. R. Geddes, Cs. Toth, J. van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans. High-quality electron beams from a laser wake-field accelerator using plasma-channel guiding. *Nature*, 431(7008):538–541, sep 2004. doi:10.1038/nature02900.
- [8] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J.-P. Rousseau, F. Burgy, and V. Malka. A laser–plasma accelerator producing monoenergetic electron beams. *Nature*, 431(7008):541–544, sep 2004. doi:10.1038/nature02963.

- [9] W. P. Leemans, B. Nagler, A. J. Gonsalves, Cs. Tóth, K. Nakamura, C. G. R. Geddes, E. Esarey, C. B. Schroeder, and S. M. Hooker. GeV electron beams from a centimetre-scale accelerator. *Nature Physics*, 2(10):696–699, sep 2006. doi:[10.1038/nphys418](https://doi.org/10.1038/nphys418).
- [10] Hyung Taek Kim, Ki Hong Pae, Hyuk Jin Cha, I Jong Kim, Tae Jun Yu, Jae Hee Sung, Seong Ku Lee, Tae Moon Jeong, and Jongmin Lee. Enhancement of electron energy to the multi-GeV regime by a dual-stage laser-wakefield accelerator pumped by petawatt laser pulses. *Physical Review Letters*, 111(16):165002, oct 2013. doi:[10.1103/PhysRevLett.111.165002](https://doi.org/10.1103/PhysRevLett.111.165002).
- [11] Xiaoming Wang, Rafal Zgadzaj, Neil Fazel, Zhengyan Li, S. A. Yi, Xi Zhang, Watson Henderson, Y.-Y. Chang, R. Korzekwa, H.-E. Tsai, C.-H. Pai, H. Quevedo, G. Dyer, E. Gaul, M. Martinez, A. C. Bernstein, T. Borger, M. Spinks, M. Donovan, V. Khudik, G. Shvets, T. Ditmire, and M. C. Downer. Quasi-monoenergetic laser-plasma acceleration of electrons to 2 GeV. *Nature Communications*, 4(1), jun 2013. doi:[10.1038/ncomms2988](https://doi.org/10.1038/ncomms2988).
- [12] W. P. Leemans, A. J. Gonsalves, H.-S. Mao, K. Nakamura, C. Benedetti, C. B. Schroeder, Cs. Tóth, J. Daniels, D. E. Mittelberger, S. S. Bulanov, J.-L. Vay, C. G. R. Geddes, and E. Esarey. Multi-GeV electron beams from capillary-discharge-guided subpetawatt laser pulses in the self-trapping regime. *Physical Review Letters*, 113(24):245002, dec 2014. doi:[10.1103/PhysRevLett.113.245002](https://doi.org/10.1103/PhysRevLett.113.245002).
- [13] A. J. Gonsalves, K. Nakamura, J. Daniels, C. Benedetti, C. Pieronek, T. C. H. de Raadt, S. Steinke, J. H. Bin, S. S. Bulanov, J. van Tilborg, C. G. R. Geddes, C. B. Schroeder, Cs. Tóth, E. Esarey, K. Swanson, L. Fan-Chiang, G. Bagdasarov, N. Bobrova, V. Gasilov, G. Korn, P. Satorov, and W. P. Leemans. Petawatt laser guiding and electron beam acceleration to 8 GeV in a laser-heated capillary discharge waveguide. *Physical Review Letters*, 122(8):084801, feb 2019. doi:[doi:doi:10.1103/PhysRevLett.122.084801](https://doi.org/10.1103/PhysRevLett.122.084801).
- [14] J Faure, D Gustas, D Guénot, A Vernier, F Bohle, M Ouillé, S Haessler, R Lopez-Martens, and A Lifschitz. A review of recent progress on laser-plasma acceleration at kHz repetition rate. *Plasma Physics and Controlled Fusion*, 61(1):014012, nov 2018. doi:<https://doi.org/10.1088/1361-6587/aae047>.
- [15] L. Rovige, J. Huijts, I. Andriyash, A. Vernier, V. Tomkus, V. Girdauskas, G. Raciukaitis, J. Dudutis, V. Stankevicius, P. Gecys, M. Ouille, Z. Cheng, R. Lopez-Martens, and J. Faure. Demonstration of stable long-term operation of a kilohertz laser-plasma accelerator. *Physical Review Accelerators and Beams*, 23(9):093401, sep 2020. doi:<https://doi.org/10.1103/PhysRevAccelBeams.23.093401>.
- [16] F Salehi, M. Le, L. Railing, M. Kolesik, and H. M. Milchberg. Laser-accelerated, low-divergence 15-MeV quasimonochromatic electron bunches at 1 kHz. *Physical Review X*, 11(2):021055, jun 2021. doi:<https://doi.org/10.1103/PhysRevX.11.021055>.

-
- [17] P. Russbueldt, T. Mans, G. Rotarius, J. Weitenberg, H. D. Hoffmann, and R. Poprawe. 400W Yb:YAG Innoslab fs-Amplifier. *Optics Express*, 17(15):12230, July 2009. doi:<https://doi.org/10.1364/OE.17.012230>.
- [18] Maria Grazia Andreassi, Andrea Borghini, Silvia Pulignani, Federica Baffigi, Lorenzo Fulgentini, Petra Koester, Monica Cresci, Cecilia Vecoli, Debora Lamia, Giorgio Russo, Daniele Panetta, Maria Tripodi, Leonida A. Gizzi, and Luca Labate. Radiobiological effectiveness of ultrashort laser-driven electron bunches: Micronucleus frequency, telomere shortening and cell viability. *Radiation Research*, 186(3):245–253, September 2016. doi:[10.1667/RR14266.1](https://doi.org/10.1667/RR14266.1).
- [19] Z.-H. He, B. Beaurepaire, J. A. Nees, G. Galle, S. A. Scott, J. R. Sanchez Perez, M. G. Lagally, K. Krushelnick, A. G. R. Thomas, and J. Faure. Capturing structural dynamics in crystalline silicon using chirped electrons from a laser wakefield accelerator. *Scientific Reports*, 6(1), November 2016. doi:[10.1038/srep36224](https://doi.org/10.1038/srep36224).
- [20] Paul Gibbon. *Short Pulse Laser Interactions with Matter*. Imperial College Press, sep 2005. doi:[10.1142/P116](https://doi.org/10.1142/P116).
- [21] J. N. Bardsley, B. M. Penetrante, and M. H. Mittleman. Relativistic dynamics of electrons in intense laser fields. *Physical Review A*, 40(7):3823–3835, oct 1989. doi:<https://doi.org/10.1103/PhysRevA.40.3823>.
- [22] Andrea Macchi. *A Superintense Laser-Plasma Interaction Theory Primer*. Springer Netherlands, 2013. doi:[10.1007/978-94-007-6125-4](https://doi.org/10.1007/978-94-007-6125-4).
- [23] E. N. Nerush and I. Yu. Kostyukov. Carrier-envelope phase effects in plasma-based electron acceleration with few-cycle laser pulses. *Physical Review Letters*, 103(3):035001, jul 2009. doi:[10.1103/PhysRevLett.103.035001](https://doi.org/10.1103/PhysRevLett.103.035001).
- [24] T. J. M. Boyd and J. J. Sanderson. *The Physics of Plasmas*. Cambridge University Press, jan 2003. doi:<https://doi.org/10.1017/CB09780511755750>.
- [25] P. Sprangle, E. Esarey, A. Ting, and G. Joyce. Laser wakefield acceleration and relativistic optical guiding. *Applied Physics Letters*, 53(22):2146–2148, nov 1988. doi:<https://doi.org/10.1063/1.100300>.
- [26] Jonathan Wood. *Betatron radiation from laser wakefield accelerators and its applications*. PhD thesis, 2016. doi:<https://doi.org/10.25560/58282>.
- [27] John M. Dawson. Nonlinear electron oscillations in a cold plasma. *Physical Review*, 113(2):383–387, jan 1959. doi:<https://doi.org/10.1103/PhysRev.113.383>.
- [28] E. Esarey, C. B. Schroeder, and W. P. Leemans. Physics of laser-driven plasma-based electron accelerators. *Reviews of Modern Physics*, 81(3):1229–1285, aug 2009. doi:<https://doi.org/10.1103/RevModPhys.81.1229>.

- [29] W. Lu, C. Huang, M. Zhou, W. B. Mori, and T. Katsouleas. Nonlinear theory for relativistic plasma wakefields in the blowout regime. *Physical Review Letters*, 96(16):165002, apr 2006. doi:<https://doi.org/10.1103/PhysRevLett.96.165002>.
- [30] W. Lu, M. Tzoufras, C. Joshi, F. S. Tsung, W. B. Mori, J. Vieira, R. A. Fonseca, and L. O. Silva. Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3d nonlinear regime. *Physical Review Special Topics - Accelerators and Beams*, 10(6):061301, jun 2007. doi:[10.1103/PhysRevSTAB.10.061301](https://doi.org/10.1103/PhysRevSTAB.10.061301).
- [31] Rémi Lehe, Manuel Kirchen, Igor A. Andriyash, Brendan B. Godfrey, and Jean-Luc Vay. A spectral, quasi-cylindrical and dispersion-free particle-in-cell algorithm. *Computer Physics Communications*, 203:66–82, jun 2016. doi:<https://doi.org/10.1016/j.cpc.2016.02.007>.
- [32] Eric Esarey and Mark Pilloff. Trapping and acceleration in nonlinear plasma waves. *Physics of Plasmas*, 2(5):1432–1436, may 1995. doi:<https://doi.org/10.1063/1.871358>.
- [33] M. Chen, E. Esarey, C. B. Schroeder, C. G. R. Geddes, and W. P. Leemans. Theory of ionization-induced trapping in laser-plasma accelerators. *Physics of Plasmas*, 19(3), mar 2012. doi:[10.1063/1.3689922](https://doi.org/10.1063/1.3689922).
- [34] J. Faure. Plasma injection schemes for laser plasma accelerators. *CERN Yellow Reports*, pages Vol 1 (2016): Proceedings of the 2014 CAS–CERN Accelerator School: Plasma Wake Acceleration, 2016. doi:<http://dx.doi.org/10.5170/CERN-2016-001.143>.
- [35] S. Corde, C. Thaury, A. Lifschitz, G. Lambert, K. Ta Phuoc, X. Davoine, R. Lehe, D. Douillet, A. Rousse, and V. Malka. Observation of longitudinal and transverse self-injections in laser-plasma accelerators. *Nature Communications*, 4(1), February 2013. doi:[10.1038/ncomms2528](https://doi.org/10.1038/ncomms2528).
- [36] V. Malka, J. Faure, C. Rechatin, A. Ben-Ismaïl, J. K. Lim, X. Davoine, and E. Lefebvre. Laser-driven accelerators by colliding pulses injection: A review of simulation and experimental results. *Physics of Plasmas*, 16(5), March 2009. doi:[10.1063/1.3079486](https://doi.org/10.1063/1.3079486).
- [37] Ming Zeng, Alberto Martinez de la Ossa, and Jens Osterhoff. Ponderomotively assisted ionization injection in plasma wakefield accelerators. *New Journal of Physics*, 22(12):123003, December 2020. doi:[10.1088/1367-2630/abc9ee](https://doi.org/10.1088/1367-2630/abc9ee).
- [38] L.-L. Yu, E. Esarey, C.B. Schroeder, J.L. Vay, C. Benedetti, C.G.R. Geddes, M. Chen, and W.P. Leemans. Two-color laser-ionization injection. *Physical Review Letters*, 112(12):125001, March 2014. doi:[10.1103/PhysRevLett.112.125001](https://doi.org/10.1103/PhysRevLett.112.125001).
- [39] C.B. Schroeder, C. Benedetti, E. Esarey, M. Chen, and W.P. Leemans. Two-color ionization injection using a plasma beatwave accelerator. *Nuclear Instruments and Meth-*

- ods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 909:149–152, November 2018. doi:[10.1016/j.nima.2018.01.008](https://doi.org/10.1016/j.nima.2018.01.008).
- [40] S. Bulanov, N. Naumova, F. Pegoraro, and J. Sakai. Particle injection into the wave acceleration phase due to nonlinear wake wave breaking. *Physical Review E*, 58(5):R5257–R5260, November 1998. doi:[10.1103/PhysRevE.58.R5257](https://doi.org/10.1103/PhysRevE.58.R5257).
- [41] J. S. Liu, C. Q. Xia, W. T. Wang, H. Y. Lu, Ch. Wang, A. H. Deng, W. T. Li, H. Zhang, X. Y. Liang, Y. X. Leng, X. M. Lu, C. Wang, J. Z. Wang, K. Nakajima, R. X. Li, and Z. Z. Xu. All-optical cascaded laser wakefield accelerator using ionization-induced injection. *Physical Review Letters*, 107(3):035001, jul 2011. doi:[10.1103/PhysRevLett.107.035001](https://doi.org/10.1103/PhysRevLett.107.035001).
- [42] C. G. R. Geddes, K. Nakamura, G. R. Plateau, Cs. Toth, E. Cormier-Michel, E. Esarey, C. B. Schroeder, J. R. Cary, and W. P. Leemans. Plasma-density-gradient injection of low absolute-momentum-spread electron bunches. *Physical Review Letters*, 100(21):215004, May 2008. doi:[10.1103/PhysRevLett.100.215004](https://doi.org/10.1103/PhysRevLett.100.215004).
- [43] A. J. Gonsalves, K. Nakamura, C. Lin, D. Panasenkov, S. Shiraishi, T. Sokollik, C. Benedetti, C. B. Schroeder, C. G. R. Geddes, J. van Tilborg, J. Osterhoff, E. Esarey, C. Toth, and W. P. Leemans. Tunable laser plasma accelerator based on longitudinal density tailoring. *Nature Physics*, 7(11):862–866, August 2011. doi:[10.1038/nphys2071](https://doi.org/10.1038/nphys2071).
- [44] G. Fubiani, E. Esarey, C. B. Schroeder, and W. P. Leemans. Improvement of electron beam quality in optical injection schemes using negative plasma density gradients. *Physical Review E*, 73(2):026402, feb 2006. doi:[10.1103/PhysRevE.73.026402](https://doi.org/10.1103/PhysRevE.73.026402).
- [45] Manuel Kirchen, Soren Jalas, Philipp Messner, Paul Winkler, Timo Eichner, Lars HÄCEbner, Thomas HÄCElsenbusch, Laurids Jeppe, Trupen Parikh, Matthias Schnepf, and Andreas R. Maier. Optimal beam loading in a laser-plasma accelerator. *Physical Review Letters*, 126(17):174801, apr 2021. doi:[10.1103/PhysRevLett.126.174801](https://doi.org/10.1103/PhysRevLett.126.174801).
- [46] W.B. Mori. The physics of the nonlinear optics of plasmas at relativistic intensities for short-pulse lasers. *IEEE Journal of Quantum Electronics*, 33(11):1942–1953, 1997. doi:[10.1109/3.641309](https://doi.org/10.1109/3.641309).
- [47] W. P. Leemans, C. E. Clayton, W. B. Mori, K. A. Marsh, P. K. Kaw, A. Dyson, C. Joshi, and J. M. Wallace. Experiments and simulations of tunnel-ionized plasmas. *Physical Review A*, 46(2):1091–1105, July 1992. doi:[10.1103/PhysRevA.46.1091](https://doi.org/10.1103/PhysRevA.46.1091).
- [48] S. C. Rae. Ionization-induced defocusing of intense laser pulses in high-pressure gases. *Optics Communications*, 97(1-2):25–28, March 1993. doi:[10.1016/0030-4018\(93\)90611-8](https://doi.org/10.1016/0030-4018(93)90611-8).

- [49] W. P. Leemans, C. E. Clayton, W. B. Mori, K. A. Marsh, A. Dyson, and C. Joshi. Plasma physics aspects of tunnel-ionized gases. *Physical Review Letters*, 68(3):321–324, January 1992. doi:[10.1103/PhysRevLett.68.321](https://doi.org/10.1103/PhysRevLett.68.321).
- [50] Guo-Zheng Sun, Edward Ott, Y. C. Lee, and Parvez Guzdar. Self-focusing of short intense pulses in plasmas. *The Physics of Fluids*, 30(2):526–532, feb 1987. doi:[10.1063/1.866349](https://doi.org/10.1063/1.866349).
- [51] B Beaurepaire, A Lifschitz, and J Faure. Electron acceleration in sub-relativistic wake-fields driven by few-cycle laser pulses. *New Journal of Physics*, 16(2):023023, feb 2014. doi:[doi:10.1088/1367-2630/16/2/023023](https://doi.org/10.1088/1367-2630/16/2/023023).
- [52] C. D. Murphy, R. Trines, J. Vieira, A. J. W. Reitsma, R. Bingham, J. L. Collier, E. J. Divall, P. S. Foster, C. J. Hooker, A. J. Langley, P. A. Norreys, R. A. Fonseca, F. Fiuza, L. O. Silva, J. T. Mendonça, W. B. Mori, J. G. Gallacher, R. Viskup, D. A. Jaroszynski, S. P. D. Mangles, A. G. R. Thomas, K. Krushelnick, and Z. Najmudin. Evidence of photon acceleration by laser wake fields. *Physics of Plasmas*, 13(3), mar 2006. doi:<https://doi.org/10.1063/1.2178650>.
- [53] J. Paye. The chronocyclic representation of ultrashort light pulses. *IEEE Journal of Quantum Electronics*, 28(10):2262–2273, 1992. doi:[10.1109/3.159533](https://doi.org/10.1109/3.159533).
- [54] J. Schreiber, C. Bellei, S. P. D. Mangles, C. Kamperidis, S. Kneip, S. R. Nagel, C. A. J. Palmer, P. P. Rajeev, M. J. V. Streeter, and Z. Najmudin. Complete temporal characterization of asymmetric pulse compression in a laser wakefield. *Physical Review Letters*, 105(23):235003, dec 2010. doi:[10.1103/PhysRevLett.105.235003](https://doi.org/10.1103/PhysRevLett.105.235003).
- [55] M.V. Streeter, S. Kneip, M.S. Bloom, R.A. Bendoyro, O. Chekhlov, A.E. Dangor, A. D’Ápp, C.J. Hooker, J. Holloway, J. Jiang, N.C. Lopes, H. Nakamura, P.A. Norreys, C.A.J. Palmer, P.P. Rajeev, J. Schreiber, D.R. Symes, M. Wing, S.P.D. Mangles, and Z. Najmudin. Observation of laser power amplification in a self-injecting laser wake-field accelerator. *Physical Review Letters*, 120(25):254801, jun 2018. doi:[10.1103/PhysRevLett.120.254801](https://doi.org/10.1103/PhysRevLett.120.254801).
- [56] E. Esarey, G. Joyce, and P. Sprangle. Frequency up-shifting of laser pulses by co-propagating ionization fronts. *Physical Review A*, 44(6):3908–3911, September 1991. doi:[10.1103/PhysRevA.44.3908](https://doi.org/10.1103/PhysRevA.44.3908).
- [57] Carl A. Lindstrøm. Staging of plasma-wakefield accelerators. *Physical Review Accelerators and Beams*, 24(1):014801, jan 2021. doi:[10.1103/PhysRevAccelBeams.24.014801](https://doi.org/10.1103/PhysRevAccelBeams.24.014801).
- [58] S. Diederichs, C. Benedetti, E. Esarey, J. Osterhoff, and C. B. Schroeder. High-quality positron acceleration in beam-driven plasma accelerators. *Physical Review Accelerators and Beams*, 23(12):121301, dec 2020. doi:[10.1103/PhysRevAccelBeams.23.121301](https://doi.org/10.1103/PhysRevAccelBeams.23.121301).

- [59] Markus Buscher, Anna Hutzen, Liangliang Ji, and Andreas Lehrach. Generation of polarized particle beams at relativistic laser intensities. *High Power Laser Science and Engineering*, 8, 2020. doi:[10.1017/hpl.2020.35](https://doi.org/10.1017/hpl.2020.35).
- [60] R. J. Shalloo, S. J. D. Dann, J.-N. Gruse, C. I. D. Underwood, A. F. Antoine, C. Arran, M. Backhouse, C. D. Baird, M. D. Balcazar, N. Bourgeois, J. A. Cardarelli, P. Hatfield, J. Kang, K. Krushelnick, S. P. D. Mangles, C. D. Murphy, N. Lu, J. Osterhoff, K. Pöder, P. P. Rajeev, C. P. Ridgers, S. Rozario, M. P. Selwood, A. J. Shahani, D. R. Symes, A. G. R. Thomas, C. Thornton, Z. Najmudin, and M. J. V. Streeter. Automation and control of laser wakefield accelerators using bayesian optimization. *Nature Communications*, 11(1), dec 2020. doi:<https://doi.org/10.1038/s41467-020-20245-6>.
- [61] Wentao Wang, Ke Feng, Lintong Ke, Changhai Yu, Yi Xu, Rong Qi, Yu Chen, Zhiyong Qin, Zhijun Zhang, Ming Fang, Jiaqi Liu, Kangnan Jiang, Hao Wang, Cheng Wang, Xiaojun Yang, Fenxiang Wu, Yuxin Leng, Jiansheng Liu, Ruxin Li, and Zhizhan Xu. Free-electron lasing at 27 nanometres based on a laser wakefield accelerator. *Nature*, 595(7868):516–520, jul 2021. doi:[doi:10.1038/s41586-021-03678-x](https://doi.org/10.1038/s41586-021-03678-x).
- [62] Z-H He, B Hou, J A Nees, J H Easter, J Faure, K Krushelnick, and A G R Thomas. High repetition-rate wakefield electron source generated by few-millijoule, 30 fs laser pulses on a density downramp. *New Journal of Physics*, 15(5):053016, may 2013. doi:[doi:10.1088/1367-2630/15/5/053016](https://doi.org/10.1088/1367-2630/15/5/053016).
- [63] B. Beaurepaire, A. Vernier, M. Bocoum, F. Bohle, A. Jullien, J-P. Rousseau, T. Lefrou, D. Douillet, G. Iaquaniello, R. Lopez-Martens, A. Lifschitz, and J. Faure. Effect of the laser wave front in a laser-plasma accelerator. *Physical Review X*, 5(3):031012, jul 2015. doi:[10.1103/PhysRevX.5.031012](https://doi.org/10.1103/PhysRevX.5.031012).
- [64] F. Salehi, A. J. Goers, G. A. Hine, L. Feder, D. Kuk, B. Miao, D. Woodbury, K. Y. Kim, and H. M. Milchberg. MeV electron acceleration at 1 kHz with < 10 mJ laser pulses. *Optics Letters*, 42(2):215, jan 2017. doi:<https://doi.org/10.1364/OL.42.000215>.
- [65] A. J. Goers, G. A. Hine, L. Feder, B. Miao, F. Salehi, J. K. Wahlstrand, and H. M. Milchberg. Multi-MeV electron acceleration by subterawatt laser pulses. *Physical Review Letters*, 115(19):194802, nov 2015. doi:[10.1103/PhysRevLett.115.194802](https://doi.org/10.1103/PhysRevLett.115.194802).
- [66] Frederik Bohle, Martin Kretschmar, Aurélie Jullien, Mate Kovacs, Miguel Miranda, Rosa Romero, Helder Crespo, Uwe Morgner, Peter Simon, Rodrigo Lopez-Martens, and Tamas Nagy. Compression of CEP-stable multi-mJ laser pulses down to 4 fs in long hollow fibers. *Laser Physics Letters*, 11(9):095401, jun 2014. doi:[doi:10.1088/1612-2011/11/9/095401](https://doi.org/10.1088/1612-2011/11/9/095401).
- [67] D. Guénot, D. Gustas, A. Vernier, B. Beaurepaire, F. Bohle, M. Bocoum, M. Lozano, A. Jullien, R. Lopez-Martens, A. Lifschitz, and J. Faure. Relativistic electron beams driven by kHz single-cycle light pulses. *Nature Photonics*, 11(5):293–296, apr 2017. doi:[DOI:10.1038/NPHOTON.2017.46](https://doi.org/10.1038/NPHOTON.2017.46).

- [68] D. Gustas, D. Guénot, A. Vernier, S. Dutt, F. Bohle, R. Lopez-Martens, A. Lifschitz, and J. Faure. High-charge relativistic electron bunches from a kHz laser-plasma accelerator. *Physical Review Accelerators and Beams*, 21(1):013401, jan 2018. doi:[10.1103/PhysRevAccelBeams.21.013401](https://doi.org/10.1103/PhysRevAccelBeams.21.013401).
- [69] L. Rovige, J. Huijts, I. A. Andriyash, A. Vernier, M. Ouillé, Z. Cheng, T. Asai, Y. Fukuda, V. Tomkus, V. Girdauskas, G. Raciukaitis, J. Dudutis, V. Stankevici, P. Gecys, R. Lopez-Martens, and J. Faure. Optimization and stabilization of a kilohertz laser-plasma accelerator. *Physics of Plasmas*, 28(3), mar 2021. doi:<https://doi.org/10.1063/5.0040926>.
- [70] B. Farace. Post-compression of high energy femtosecond pulses. Master's thesis, 2020. URL: <https://hdl.handle.net/10589/164722>.
- [71] Tamas Nagy, Peter Simon, and Laszlo Veisz. High-energy few-cycle pulses: post-compression techniques. *Advances in Physics: X*, 6(1), November 2020. doi:[10.1080/23746149.2020.1845795](https://doi.org/10.1080/23746149.2020.1845795).
- [72] Donna Strickland and Gerard Mourou. Compression of amplified chirped optical pulses. *Optics Communications*, 56(3):219–221, December 1985. doi:[10.1016/0030-4018\(85\)90120-8](https://doi.org/10.1016/0030-4018(85)90120-8).
- [73] A. Dubietis, G. Jonusauskas, and A. Piskarskas. Powerful femtosecond pulse generation by chirped and stretched pulse parametric amplification in BBO crystal. *Optics Communications*, 88(4–6):437–440, April 1992. doi:[https://doi.org/10.1016/0030-4018\(92\)90070-8](https://doi.org/10.1016/0030-4018(92)90070-8).
- [74] C. M. Lazzarini, G. M. Grittani, P. Valenta, I. Zymak, R. Antipenkov, U. Chaulagain, L. V. N. Goncalves, A. Grenfell, M. Lamac, S. Lorenz, M. Nevrkla, V. Sobr, A. Spacek, W. Szuba, P. Bakule, G. Korn, and S. V. Bulanov. 50 MeV electron beams accelerated by a terawatt scalable kHz laser. 2023. doi:<https://doi.org/10.48550/arXiv.2302.11415>.
- [75] R. W. Boyd. *Nonlinear optics*. Elsevier, 2020. doi:<https://doi.org/10.1016/C2015-0-05510-1>.
- [76] E. Garmire, R. Y. Chiao, and C. H. Townes. Dynamics and Characteristics of the Self-Trapping of Intense Light Beams. *Physical Review Letters*, 16(9):347–349, February 1966. doi:<https://doi.org/10.1103/PhysRevLett.16.347>.
- [77] Robert Szipocs, Christian Spielmann, Ferenc Krausz, and Karpát Ferencz. Chirped multilayer coatings for broadband dispersion control in femtosecond lasers. *Optics Letters*, 19(3):201, February 1994. doi:[10.1364/ol.19.000201](https://doi.org/10.1364/ol.19.000201).
- [78] Felix Frank. *Generation and Application of Ultrashort Laser Pulses in Attosecond Science*. PhD thesis, 2011. doi:<https://doi.org/10.25560/7025>.

- [79] Supriya Rajhans, Esmerando Escoto, Nikita Khodakovskiy, Praveen K. Velpula, Bonaventura Farace, Uwe Grosse-Wortmann, Rob J. Shaloo, Cord L. Arnold, Kristjan Poder, Jens Osterhoff, Wim P. Leemans, Ingmar Hartl, and Christoph M. Heyl. Post-compression of multi-millijoule picosecond pulses to few-cycles approaching the terawatt regime. *Optics Letters*, 48(18):4753, September 2023. doi:<https://doi.org/10.1364/OL.498042>.
- [80] M. Nisoli, S. De Silvestri, and O. Svelto. Generation of high energy 10 fs pulses by a new pulse compression technique. *Applied Physics Letters*, 68(20):2793–2795, May 1996. doi:[10.1063/1.116609](https://doi.org/10.1063/1.116609).
- [81] Marie Ouille, Aline Vernier, Frederik Bohle, Maimouna Bocoum, Aurelie Jullien, Magali Lozano, Jean-Philippe Rousseau, Zhao Cheng, Dominykas Gustas, Andreas Blumenstein, Peter Simon, Stefan Haessler, Jerome Faure, Tamas Nagy, and Rodrigo Lopez-Martens. Relativistic-intensity near-single-cycle light waveforms at kHz repetition rate. *Light: Science and Applications*, 9(1), March 2020. doi:[10.1038/s41377-020-0280-5](https://doi.org/10.1038/s41377-020-0280-5).
- [82] Anne-Lise Viotti, Marcus Seidel, Esmerando Escoto, Supriya Rajhans, Wim P. Leemans, Ingmar Hartl, and Christoph M. Heyl. Multi-pass cells for post-compression of ultrashort laser pulses. *Optica*, 9(2):197, February 2022. doi:<https://doi.org/10.1364/OPTICA.449225>.
- [83] D. Herriott, H. Kogelnik, and R. Kompfner. Off-axis paths in spherical mirror interferometers. *Applied Optics*, 3(4):523, April 1964. doi:<https://doi.org/10.1364/AO.3.000523>.
- [84] Esmerando Escoto, Anne-Lise Viotti, Skirmantas Alisauskas, Henrik Tunnermann, Ingmar Hartl, and Christoph M. Heyl. Temporal quality of post-compressed pulses at large compression factors. *Journal of the Optical Society of America B*, 39(7):1694, June 2022. doi:<https://doi.org/10.1364/JOSAB.453901>.
- [85] E.B. Treacy. Compression of picosecond light pulses. *Physics Letters A*, 28(1):34–35, October 1968. doi:[https://doi.org/10.1016/0375-9601\(68\)90584-7](https://doi.org/10.1016/0375-9601(68)90584-7).
- [86] Rick Trebino, Kenneth W. DeLong, David N. Fittinghoff, John N. Sweetser, Marco A. Krumboltz, Bruce A. Richman, and Daniel J. Kane. Measuring ultrashort laser pulses in the time-frequency domain using frequency-resolved optical gating. *Review of Scientific Instruments*, 68(9):3277–3295, September 1997. doi:[10.1063/1.1148286](https://doi.org/10.1063/1.1148286).
- [87] Jiyuan Tu, Chaoqun Liu, and Guan-Heng Yeoh. *Computational Fluid Dynamics A Practical Approach*. Elsevier, 2019. doi:[10.1016/C2015-0-06135-4](https://doi.org/10.1016/C2015-0-06135-4).
- [88] S. Semushin and V. Malka. High density gas jet nozzle design for laser target production. *Review of Scientific Instruments*, 72(7):2961–2965, July 2001. doi:[10.1063/1.1380393](https://doi.org/10.1063/1.1380393).

- [89] K. Schmid and L. Veisz. Supersonic gas jets for laser-plasma experiments. *Review of Scientific Instruments*, 83(5), May 2012. doi:10.1063/1.4719915.
- [90] L. Rovige, J. Huijts, A. Vernier, I. Andriyash, F. Sylla, V. Tomkus, V. Girdauskas, G. Raciukaitis, J. Dudutis, V. Stankevicius, P. Gecys, and J. Faure. Symmetric and asymmetric shocked gas jets for laser-plasma experiments. *Review of Scientific Instruments*, 92(8), August 2021. doi:10.1063/5.0051173.
- [91] K. Schmid, A. Buck, C. M. S. Sears, J. M. Mikhailova, R. Tautz, D. Herrmann, M. Geissler, F. Krausz, and L. Veisz. Density-transition based electron injector for laser driven wakefield accelerators. *Physical Review Special Topics - Accelerators and Beams*, 13(9):091301, September 2010. doi:10.1103/PhysRevSTAB.13.091301.
- [92] Farace B., Shaloo R.J., P oder K., and Leemans W. Plasma source device and plasma accelerator apparatus for laser-driven acceleration of electrons and creation of X-rays, 2023.
- [93] Stephane Colin. Rarefaction and compressibility effects on steady and transient gas flows in microchannels. *Microfluidics and Nanofluidics*, 1(3):268–279, August 2004. doi:10.1007/s10404-004-0002-y.
- [94] Maxim V. Ammosov, Nikolai B. Delone, and Vladimir P. Krainov. Tunnel ionization of complex atoms and atomic ions in electromagnetic field. October 1986. doi:10.1117/12.938695.
- [95] S. Augst, D. Strickland, D. D. Meyerhofer, S. L. Chin, and J. H. Eberly. Tunneling ionization of noble gases in a high-intensity laser field. *Physical Review Letters*, 63(20):2212–2215, November 1989. doi:10.1103/PhysRevLett.63.2212.
- [96] Donald R. Jones, Matthias Schonlau, and William J. Welch. *Journal of Global Optimization*, 13(4):455–492, 1998. doi:10.1023/A:1008306431147.
- [97] A. Ferran Pousa, S. Jalas, M. Kirchen, A. Martinez de la Ossa, M. Thevenet, S. Hudson, J. Larson, A. Huebl, J.-L. Vay, and R. Lehe. Bayesian optimization of laser-plasma accelerators assisted by reduced physical models. *Physical Review Accelerators and Beams*, 26(8):084601, August 2023. doi:10.1103/PhysRevAccelBeams.26.084601.
- [98] Soren Jalas, Manuel Kirchen, Philipp Messner, Paul Winkler, Lars Hubner, Julian Dirkwinkel, Matthias Schnepf, Remi Lehe, and Andreas R. Maier. Bayesian optimization of a laser-plasma accelerator. *Physical Review Letters*, 126(10):104801, March 2021. doi:10.1103/PhysRevLett.126.104801.
- [99] M. Lenzner, J. KrÄEger, S. Sartania, Z. Cheng, Ch. Spielmann, G. Mourou, W. Kautek, and F. Krausz. Femtosecond optical breakdown in dielectrics. *Physical Review Letters*, 80(18):4076–4079, May 1998. doi:10.1103/PhysRevLett.80.4076.

-
- [100] Jason Cole. *Diagnosis and application of laser wakefield accelerators*. PhD thesis, 2016. doi:10.25560/42222.
- [101] Mitsuo Takeda, Hideki Ina, and Seiji Kobayashi. Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry. *Journal of the Optical Society of America*, 72(1):156, January 1982. doi:10.1364/JOSA.72.000156.
- [102] Eric W. Hansen and Phaih-Lan Law. Recursive methods for computing the abel transform and its inverse. *Journal of the Optical Society of America A*, 2(4):510, April 1985. doi:10.1364/JOSAA.2.000510.
- [103] L. B. Lucy. An iterative technique for the rectification of observed distributions. *The Astronomical Journal*, 79:745, June 1974. doi:10.1086/111605.
- [104] T. Auguste, P. Monot, L. Lompre, G. Mainfray, and C. Manus. Defocusing effects of a picosecond terawatt laser pulse in an underdense plasma. *Optics Communications*, 89(2-4):145–148, May 1992. doi:10.1016/0030-4018(92)90148-K.
- [105] M. Tzoufras, F.S. Tsung, W. B. Mori, and A. A. Sahai. Improving the self-guiding of an ultraintense laser by tailoring its longitudinal profile. *Physical Review Letters*, 113(24):245001, December 2014. doi:10.1103/PhysRevLett.113.245001.
- [106] N G Khodakovskiy, M P Kalashnikov, V Pajer, A Blumenstein, P Simon, M M Toktamis, M Lozano, B Mercier, Z Cheng, T Nagy, and R Lopez-Martens. Generation of few-cycle laser pulses with high temporal contrast via nonlinear elliptical polarisation rotation in a hollow fibre compressor. *Laser Physics Letters*, 16(9):095001, August 2019. doi:10.1088/1612-202X/ab3852.
- [107] Viktor Pajer and Mikhail Kalashnikov. High temporal contrast ultrashort pulses generated by nonlinear ellipse rotation in multipass cells. *Laser Physics Letters*, 18(6):065401, May 2021. doi:10.1088/1612-202X/abff49.
- [108] C. B. Schroeder, E. Esarey, B. A. Shadwick, and W. P. Leemans. Trapping, dark current, and wave breaking in nonlinear plasma waves. *Physics of Plasmas*, 13(3), mar 2006. doi:https://doi.org/10.1063/1.2173960.
- [109] Daniela Dragoman. Phase-space interferences as the source of negative values of the wigner distribution function. *Journal of the Optical Society of America A*, 17(12):2481, December 2000. doi:10.1364/JOSAA.17.002481.
- [110] Tsan-Hsing Shih, William W. Liou, Aamir Shabbir, Zhigang Yang, and Jiang Zhu. A new k-e eddy viscosity model for high reynolds number turbulent flows. *Computers and Fluids*, 24(3):227–238, March 1995. doi:10.1016/0045-7930(94)00032-T.
- [111] Spencer W Jolly, Olivier Gobert, and Fabien Quere. Spatio-temporal characterization of ultrashort laser beams: a tutorial. *Journal of Optics*, 22(10):103501, September 2020. doi:10.1088/2040-8986/abad08.

- [112] Antoine Jeandet, Antonin Borot, Kei Nakamura, Spencer W Jolly, Anthony J Gonalves, Csaba Toth, Hann-Shin Mao, Wim P Leemans, and Fabien Quere. Spatio-temporal structure of a petawatt femtosecond laser beam. *Journal of Physics: Photonics*, 1(3):035001, June 2019. doi : [10.1088/2515-7647/ab250d](https://doi.org/10.1088/2515-7647/ab250d).

APPENDIX A

Wavebreaking injection

The mechanisms for trapping and injection in a plasma wakefield have already been discussed in section 2.3. However delving a little further into the physics of wavebreaking can be helpful to elucidate its relation with particles injection.

In the 1D nonlinear regime, a more general expression of evolution of the normalised scalar potential ϕ in the co-moving frame ξ , including normalised phase velocity of the plasma wave β_p , can be written as [108]:

$$\frac{\partial^2 \phi}{\partial \xi^2} = k_p^2 \gamma_p^2 \left[\beta_p \left(1 - \frac{\gamma_\perp^2}{\gamma_p (1 + \phi)^2} \right)^{-1/2} - 1 \right], \quad (\text{A.1})$$

where $\gamma_p = (1 - \beta_p^2)^{-1/2}$ and $\gamma_\perp = (1 + \frac{a^2}{2})^{1/2}$ are the two Lorentz factors associated respectively with the plasma wave phase velocity and the electron quiver velocity. Equation A.1 breaks down when a singularity occurs, i.e when $\frac{\gamma_\perp}{\gamma_p} = 1 + \phi$. This is the definition of the cold relativistic wavebreaking, which in terms of peak electric field can be expressed as [108]:

$$E_{\text{WB}} = [2\gamma_\perp(\gamma_p - 1)]^{1/2} E_0, \quad (\text{A.2})$$

where $E_0 = \frac{m_e c \omega_p}{e}$ is the linear cold wavebreaking limit.

Considering the dynamics of a single electron in a plasma wave the same condition defines the first closed orbit, i.e. the separatrix, in the (u_z, ξ) phase space, where u_z is the normalised longitudinal electron momentum (cf. section 2.3). Recalling equation 2.27:

$$u_z(\xi) = \beta_p \gamma_p^2 (H_0 + \phi) \pm \gamma_p \sqrt{\gamma_p^2 (H_0 + \phi)^2 - \gamma_\perp^2}, \quad (\text{A.3})$$

the Hamiltonian defining the separatrix can be written as [33]:

$$H_{\text{sep}} = \frac{\gamma_\perp}{\gamma_p} - \phi. \quad (\text{A.4})$$

As we have already seen, for an electron to be trapped in a wakefield its hamiltonian $H_0 \leq H_{\text{sep}}$. A general trapping condition can hence relate the electron quiver motion to

the plasma phase velocity as:

$$\frac{\gamma_{\perp}(\xi)}{\gamma_p} \geq \phi(\xi) + H_0. \quad (\text{A.5})$$

This generally indicates that, at a position ξ , the electron fluid velocity exceeds that of the plasma wave and it is, namely, the onset of wavebreaking¹. Consequently, injection mechanisms targeting this condition can be classified as “wavebreaking injections”. While this nomenclature is usually associated with self-injection, a number of different methods are part of this class. As an example, density downramp injection (particularly relevant for high average power LPA) has the purpose of inducing injection by reducing the plasma wave phase velocity and, thus, its corresponding Lorentz factor.

¹for a cold plasma electron $H_0 = 1$ and the equation exactly reduces to the definition given before.

APPENDIX B

Wigner distributions

In ultrashort pulses the spectral and temporal description are highly entangled, therefore a tool to describe their interplay can be particularly advantageous. For a long time the scientific community has relied on spectrograms. A spectrogram function, indeed, can give a good understanding of pulse dispersion (group delay dispersion, third order dispersion, etc.¹) but, being always positive, makes it difficult to extrapolate the actual pulse spectrum and temporal profile. This is where Wigner functions come into the picture.

Given a complex electric field $E(t)$, the Wigner function of the field $W(E; t, \omega)$ (or its “chronocyclic intensity”) is the Fourier transform of the “mutual temporal intensity”, which is related to the temporal intensity but written in the quantum notation as $E(t) \times E^*(t)$ (where the star indicates the complex conjugate) [53]. It represents, simultaneously, how the energy of a pulse is distributed both along its temporal envelope and among its spectral components:

- ◇ the integral of the function over the whole (ω, t) domain is equal to the total pulse energy, meaning that the time-bandwidth product of the pulse is implicitly defined and
- ◇ the marginals of the function are the spectral and temporal intensity:

$$\int W dt = I(\omega), \quad (\text{B.1})$$

$$\frac{1}{2\pi} \int W d\omega = I(t). \quad (\text{B.2})$$

It is interesting to point out that aside from gaussian functions, for which the Wigner distribution is always positive, any other pulse functional form will inherently give rise to region of negative chronocyclic intensity, whose physical meaning is often counterintuitive. Because of their mathematical definition, Wigner functions are bilinear functions of time and frequency therefore they capture the correlations between the two domains [109].

¹cf. appendix C

When the signal in one domain interfere (in the Wigner space) with its own complex conjugate in the other domain, the resulting interference pattern can give rise to negative values of intensity. These values though, instead of representing unphysical negative intensities, indicate phase relationships between frequency components at a specific time that results in a decrease in the overall signal's intensity at that particular frequency, when compared to the neighbouring frequencies.

To clarify the concept, let's briefly analyse a practical example². Consider a 1D odd field distribution:

$$E(x) = C_0 x e^{-\pi x^2}, \quad (\text{B.3})$$

where C_0 is a normalisation constant. Figure B.1 illustrate the field $E(x)$ alongside its corresponding Wigner distribution $W(E)$ in the (x, k_x) space. Odd field distributions always results in negative Wigner intensities in the centre of the phase space, therefore they are particularly instrumental in the current scenario.




Figure B.1.: Example of an odd field distribution $E(x)$ and its corresponding (normalised) Wigner distribution W , as described in the main text.

$E(x)$ has been chosen in order to be effectively approximated by two gaussians, suitably displaced and dephased:

$$E(x) = E_1(x) + E_2(x) = C_0 e^{\pi d^2} e^{-\pi(x-d)^2} - e^{-\pi(x+d)^2}, \quad (\text{B.4})$$

where $d = 1/4\pi$ is the “displacement” parameter. As illustrated in figure B.2, although each gaussian field has an always positive Wigner function, the Wigner distribution of the total field $E_1(x) + E_2(x)$ exhibits negative values due to their interference in the phase space:

$$W(E_1 + E_2) = W(E_1) + W(E_2) + W_{\text{int}}(E_1, E_2), \quad (\text{B.5})$$

where $W_{\text{int}}(E_1, E_2)$ includes the “cross-terms” of the Wigner intensity, which in this case are the ones proportional to $E_1(x) \times E_2^*(x)$ and $E_2(x) \times E_1^*(x)$.

²Adapted from [109].

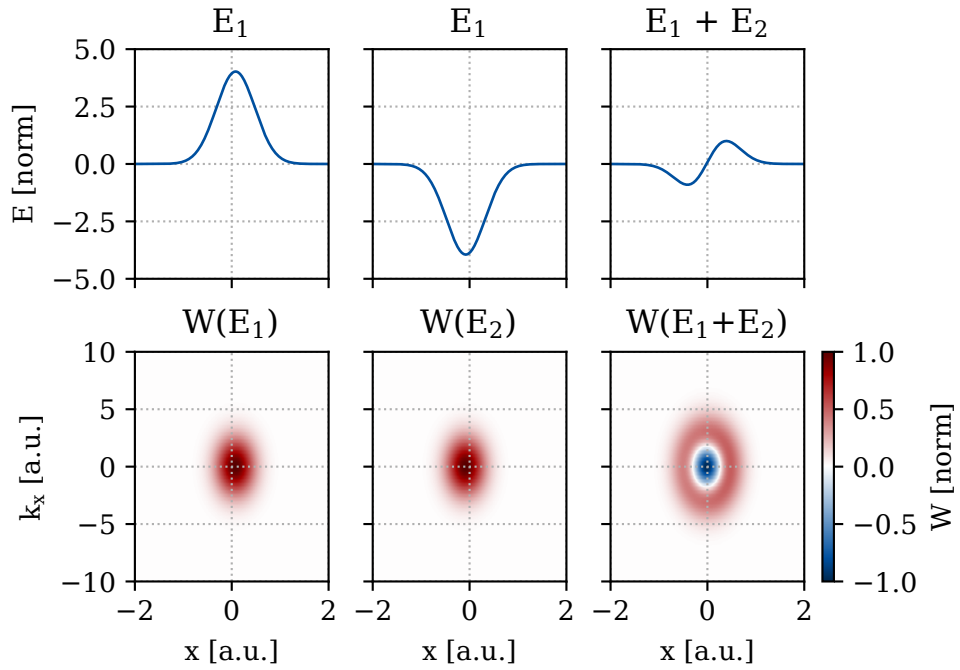


Figure B.2.: The field $E(x)$ plotted as a decomposition of gaussian fields E_1 and E_2 (top row). The bottom row shows the corresponding normalised Wigner distributions W .

Having discussed the origin of the negative values in Wigner distributions, let's consider again the example of section 2.4.3 and see how this applies to a laser pulse. During the propagation in the plasma the pulse evolves, loses its initially gaussian nature and gives rise to regions of negative Wigner intensity. The bottom panels of figure B.3 show two lineouts of the Wigner distribution depicted on the top, respectively for a given frequency (1.6×10^{15} rad/s) and a given time (7 fs). The considered frequency (blue lineout) adds up in phase with the neighbouring frequencies in all the time instants for which its Wigner intensity is positive, while it goes out of phase when it becomes negative. The lineout, therefore, represents how the energy at that given frequency is distributed along the pulse. The integral under the curve is then representative of the spectral intensity of the frequency over the whole pulse. Equivalently, when a single time instant is consider (red), the Wigner intensity lineout depicts which frequencies add up in phase at that given time (positive) and which one are out of phase (negative). Again the integral under the curve is proportional to the overall pulse intensity at that specific time.




Figure B.3.: Wigner distribution of the laser pulse considered in 2.4.3. Two exemplary line-outs for a given frequency (1.6×10^{15} rad/s) and a certain time (7 fs) are shown in the bottom panels, respectively in blue and red.

APPENDIX C

Chromatic dispersion

Generally, when a laser pulse interacts with a dielectric medium, its response depends on the frequency of the radiation ω . This dependence is usually addressed as “chromatic dispersion” and it is responsible for several optical properties of the medium.

The time-dependent carrier frequency of the pulse, the instantaneous frequency, is linked to the time derivative of the phase as:

$$\omega_{(t)} = \omega_0 - \frac{d\varphi}{dt}, \quad (\text{C.1})$$

therefore the effects of dispersion can be described expanding the spectral phase around the central pulse frequency ω_0 :

$$\varphi(\omega) = \varphi(\omega_0) + \frac{\partial \varphi}{\partial \omega}(\omega - \omega_0) + \frac{1}{2} \frac{\partial^2 \varphi}{\partial \omega^2}(\omega - \omega_0)^2 + \frac{1}{6} \frac{\partial^3 \varphi}{\partial \omega^3}(\omega - \omega_0)^3 + \dots \quad (\text{C.2})$$

- ◇ The zero order term just adds a constant phase all along the laser spectrum, which does not impact the instantaneous frequency.
- ◇ The first order term, or the “group delay”, is the inverse of the group velocity ($\text{GD} = 1/v_g = dk/d\omega$) and results in a delay in the pulse, which shifts in time.
- ◇ The second order is referred to as “group velocity dispersion” ($\text{GVD} = \frac{d}{d\omega} \text{GD}$). Differently from the ones before, this introduces a frequency-dependent delay, thus changing the pulse temporal profile. In the case of an initial transform-limited¹ gaussian pulse, it will introduce a quadratic component, resulting in a linear dependence of the group velocity of the pulse with respect to the frequency ($v_g = v_g(\omega)$). Different frequencies, hence, will propagate at different velocities inside the medium. Depending on the GVD sign, the high-frequency end of the pulse will be slower ($\text{GVD} > 0$) or faster ($\text{GVD} < 0$) than the low-frequency components and during the propagation the pulse will be temporally stretched.

¹i.e. with a constant phase

- ◇ The “third order dispersion” ($\text{TOD} = \frac{d}{d\omega}\text{GVD}$) is often negligible, as usually $\text{TOD} \ll \text{GVD}$. However, when dealing with ultrashort pulses (and therefore with very broad bandwidths) its relative importance is higher and it can heavily affect the pulse temporal profile. In the case of an initial transform limited gaussian pulse, the presence of TOD makes the pulse asymmetric and develops an oscillatory structure on one of the sides of the two sides, depending on its sign.

Fig. C.1, for instance, illustrates the effects of (pure) GVD and (pure) TOD on an initially transform-limited ultrashort pulse². These two simple examples clearly indicate that precise control of dispersion is crucial in the context of ultrashort lasers.




Figure C.1.: The effect on the pulse temporal structure of a 10 fs^2 accumulated GVD (blue) and of a 100 fs^3 accumulated TOD (green) on a 10 fs transform-limited pulse (red).

²Note that chromatic dispersion affects only the pulse temporal structure, while the spectrum remains unchanged.

APPENDIX D

Ansys-Fluent

SECTION D.1

Guidelines for fluid modeling

Ansys-Fluent, along with other CFD softwares, is designed to be user-oriented and relatively easy to use. To a new user, it can initially appear as a plug-and-play program capable of producing reliable fluid dynamics results effortlessly. However, it is not rare to create ill-posed problems where convergence is achieved only under specific initial conditions, or not at all. Ultimately, the reliability of the results depends on how well the fluid problem is structured. Therefore, a brief “getting started” guide to CFD simulations is presented below, with a particular focus on the use of Ansys-Fluent. Note that this is not intended to be a comprehensive guide, but rather an overview of the most important parameters to consider when building a simulation and checking its convergence. A more extensive and complete treatment can be found in [87].

Meshing a CFD problem

Meshing is probably the most critical aspect of CFD simulations, strongly impacting both the accuracy and the efficiency of the solution. An incorrect mesh can easily lead to simulations particularly sensitive to the initial conditions, with excessively long computation times and convergence issues.

When a new problem is set, it is generally advisable to start with a coarse mesh. This approach allows for the rapid evaluation of simulation time and computational cost, it helps in identifying challenging geometric features (such as flow restrictions or sharp corners), and it enables multiple quick test runs to optimise solver settings and convergence criteria. In fact, while mesh resolution can play a role, typically well-posed problems converge even with a coarse mesh. Yet progressive mesh refinement is necessary to enhance the solution accuracy.

To avoid ill-posed problems, ensuring the mesh quality is key. Mesh quality is mainly governed by two factors: the aspect ratio (AR) of the cells and their skewness (S). Depending on the solver used, an excessively large AR could degrade the solution accuracy and eventually induce divergence. A good guideline is to keep $0.2 < AR < 5$ in areas with free-flow

conditions, i. e. in the inner parts of the geometry. This interval gets wider in the vicinity of the walls, where the AR is typically increased along the flow direction to resolve the fluid boundary layers. On the other hand, the skewness relates to the angles between the mesh lines. It is particularly important near the walls, where it can lead to numerical instabilities. Ideally, right angles between the grid lines are preferred, however angles $45^\circ < \theta < 135^\circ$ are generally considered acceptable for accurate simulations. Due to skewness considerations, quadrilateral meshes (in 2D) and polyhedral meshes (in 3D) are typically more robust compared to their triangular and tetrahedral counterparts.

Solver settings

Typically in softwares like Fluent, a wide variety of solvers and solver settings are available, making it challenging to identify the optimal configuration given conditions. Each flow problem may require different settings and a common guideline is to experiment different solvers to determine the most effective approach for a specific case. However, considering plasma sources design in Fluent, some settings appear more robust than others. These will not guarantee convergence in every situation, yet they can serve as a reliable starting point.

In Fluent, simulations can be performed using either a pressure-based or a density-based solver. Pressure-based solvers typically use a segregated approach, solving the governing equations for each flow variable sequentially. This approach makes them generally fast and robust, however they may struggle with compressible flows, which are often critical in plasma source simulations. A density-based solver, on the other hand, employs a coupled approach, solving the continuity, momentum, and energy equations simultaneously. This makes it more suitable for modeling compressible flows and shocks, yet more complex to manage. To bridge the gap, Fluent offers coupled methods within the pressure-based solver. This approach combines the robustness of pressure-based solvers with the accuracy of coupled methods, making it the preferred choice for most of the fluid simulations discussed in this thesis.

In addition to their approach to solving the fluid equations, solvers can be differentiated based on their interpolation methods, as cell-based or node-based. Cell-based solvers compute variables on the center of each mesh cell, while using interpolation methods to reconstruct the quantities throughout the fluid volume. Conversely, node-based solvers solve equations at the mesh nodes. The main distinction between these approaches lies in computational efficiency and accuracy, driven by the simple statement that a cell has one centre, but multiple nodes. Indeed, while the node-based approach requires more computational resources and time, it provides greater accuracy, particularly in resolving flow features near the walls (in the boundary layers), where steep gradients occur.

Convergence monitoring

Monitoring convergence is pivotal, particularly to avoid wasting computational resources on unsuccessful runs. In this context, achieving a qualitative convergence can initially suffice when testing solver settings and convergence criteria, while a more quantitative convergence can be pursued in later stages for accuracy.

A practical method to monitor qualitative convergence is through residuals. In CFD, the conservation law for any transport variable ϕ in a mesh element e can be written as a

balance equation involving the contributions from neighboring elements. This can be expressed as:

$$a_e \phi_e = \sum_{\text{nb}} a_{\text{nb}} \phi_{\text{nb}} + C, \quad (\text{D.1})$$

where a_e and a_{nb} are coefficients that represent the influence of the central element and its neighboring elements, respectively, and C is a constant that may arise from source terms or boundary conditions. During the iterative solving process these coefficients and the value of the transport variable are progressively updated to minimise the imbalance. The “residual” for a given mesh element R_e represents the net imbalance:

$$R_e = a_e \phi_e - \sum_{\text{nb}} a_{\text{nb}} \phi_{\text{nb}}. \quad (\text{D.2})$$

Residuals are essential for monitoring simulation convergence as they provide a measure of how well the numerical solution satisfies the governing equations. Qualitative convergence is typically achieved when all relevant residuals have decreased by 3 or more orders of magnitude. Furthermore, to enhance the stability of the iterative solution process, “under-relaxation” factors (β_{ur}) can be used. They control the extent to which newly computed values of flow variables are applied within each iteration. Therefore, they effectively damp rapid changes in the variables that might lead to divergence or oscillations¹. The updated value of a flow variable ϕ for a mesh element e is computed as:

$$\phi_e^{\text{new}} = \phi_e^{\text{old}} + \beta_{\text{ur}}(\phi_e^{\text{computed}} - \phi_e^{\text{old}}). \quad (\text{D.3})$$

On the other hand, quantitative convergence is best monitored by observing specific variables at designated points within the geometry until their values stabilise. It is good practice to monitor different variables at various locations, particularly at critical points, such as areas with strong expansion, flow restrictions, or other complex flow features. Once a converged solution is obtained, conducting a grid-independence study is essential to ensure that the solution does not depend on the specific mesh used. The usual approach is to double the mesh refinement twice, ensuring consistency across three consecutive solutions. If the results remain consistent, the initial converged solution is considered accurate and further mesh refinements are unnecessary.

SECTION D.2

Turbulence models

A high average power LPA often relies on near-critical electron density. As discussed in chapter 4, such densities generally necessitate high gas pressures and, as a result, the plasma sources devoted to this specific application have to manage high Reynolds number within their flow regime. Consequently, to capture the turbulent behaviour inherent to the flow, the application of turbulence models is essential.

¹Typically, $0 < \beta_{\text{ur}} < 1$. Note that smaller under-relaxation factors enhance the simulation stability, however they slow down the convergence since each iteration makes a smaller adjustment to the flow variables.

Turbulence is generally associated to random fluctuation in the fluid quantities. The fluid velocity $u(t)$ in a given point, for instance, could be expressed as $u(t) = \langle u \rangle + u'(t)$, where $\langle u \rangle$ is the mean velocity and $u'(t)$ is its statistical, random fluctuation. To accurately describe turbulent flows, without any approximation, a direct solution of the transport equation would be needed. However, in most of the cases, this calls for a very refined grid (in order to resolve even the smallest fluctuation) and, hence, results in an extremely time consuming simulation.

Nonetheless, insight about the mean quantities are often enough to describe the flow evolution. In this context, the Reynolds-Averaged Navier Stokes models (RANS) aim at predicting the effects of turbulences on the mean flow properties by averaging the Navier-Stokes equation over time. RANS models typically parametrise the fluctuating components of the fluid quantities by means of additional transport equations for “turbulent quantities”, namely:

- ◇ the turbulent kinetic energy $k = \frac{1}{2} u'_i u'_i$ and
- ◇ the turbulent dissipation rate $\varepsilon = \nu_T \left\langle \frac{\partial u'_i}{\partial x_j} \frac{\partial u'_i}{\partial x_j} \right\rangle$,

where ν_T is the turbulent (“eddy”) viscosity, x is the position of the fluid element, i and j are two separate axes and $\langle \rangle$ denotes a mean value².

Among the models:

- ◇ the *Standard k-ε* model is one of the earliest, most robust and most widely used. It can adapt to a wide variety of scenarios, being only slightly dependent of the simulated geometry or on the input flow conditions. However, while being accurate on the free-flow, it cannot properly resolve the fluid boundary layers and can lead to misleading results in the near-walls regions.
- ◇ The *Standard k-ε* model includes empirical constants derived from a set of “well-known” fluid problems compared with their experimental results. These coefficients affect the behavior of the turbulence model and play a crucial role in determining the accuracy of the model predictions. However, a set of “realisability” constraints can be applied to the model in order to express the same constants with a mathematical formulation which is independent from the specific problem under analysis. Following this path, the *Realizable k-ε* model [110] ensures that both k and ε are always positive and prevent unphysical behaviour of the turbulent viscosity ν_T , providing more physical consistency.
- ◇ The *Standard k-ω* model, on the other hand, has been developed specifically to perform at best where strong gradients are present, namely near the fluid walls. Equivalently, though, it delivers less accuracy on free stream quantities and it can be extremely sensible to the input conditions.

²note that in some models the turbulent dissipation rate is expressed as $\omega \propto \varepsilon/k$.

- ◇ To overcome the limitations of these methods, the *Shear-stress transport (SST) k - ω* model combines the two standard models just introduced, adopting a blending function to interface them in the intermediate region between free stream flow and boundary layers.
- ◇ Alternatively, to handle the near-wall region in a more efficient way, *wall functions* can be implemented in k - ε models. They provide a simplified treatment of the boundary layers through predefined dependencies of specific fluid quantities with respect to the distance from the wall. Taking again the fluid velocity as an example, the “standard” wall function $U^+ = f(y^+)$ is defined by parts:

$$\begin{cases} U^+ = y^+ & \text{viscous layer } (y^+ < 5) \\ U^+ = \frac{1}{c_1} \ln(c_2 y^+) & \text{turbulent layer } (y^+ > 30) \end{cases}, \quad (\text{D.4})$$

where U^+ is the fluid velocity u normalised to the so called “friction velocity” u_T , y^+ is the distance from the wall normalised to the viscous length scale of the flow (ν/u_T) and c_1 and c_2 are wall-dependent constants (for smooth walls, for instance, $c_1 \simeq 0.4$ and $c_2 \simeq 9.8$)³.

As an example, figure D.1 compares the results from different flow models for a 2D supersonic de Laval nozzle with backing pressure $P = 20$ bar, throat $t = 40 \mu\text{m}$, outlet diameter $D = 100 \mu\text{m}$ and length $L = 200 \mu\text{m}$. The first two panels show the molecular (nitrogen) number density and Mach number along the nozzle axis, where the throat is placed at $z = -200 \mu\text{m}$ and the outlet at $z = 0$. The bottom panels, on the other hand, zoom in the region where a shock is visible for some of the models. All the analysed models accurately depict the mean evolution of the fluid quantities and lead to very similar results. In a counterintuitive way, this includes the laminar model which, for relatively simple geometries as the one considered here, can provide a good description of a supersonic, turbulent flow with a lower computational cost. Moreover, the model correctly locate the shock position which, conversely, it is completely missed by both the standard turbulence models (k - ε and k - ω). It is interesting to notice how the shock amplitude predicted in the laminar model is higher with respect to the results of the more advanced SST and Realizable models. This can generally be traced back to how the boundary layer is modeled. While in a laminar scenario viscous effects dominate, leading to a gradual transition from wall to free stream quantities, turbulent simulations include turbulent mixing and dissipation rate that mitigate the growth of the shock and generally points towards a more reliable result.

During the years, both the SST k - ω and the Realizable k - ε (with wall function) model have been proven to accurately predict turbulent flows in a wide variety of problems and, specifically, in gas jets. However, as the k - ω model does not take advantage of wall functions, its accuracy is heavily dependent on the correct mesh resolution of the boundary layers. On average, hence, it requires a more refined mesh and a longer computation time. While

³It is interesting to note here that, in the simulation domain, these functions are discretised according to the mesh, which must be fine enough to properly resolve the near-wall fluid quantity gradients. While different rules of thumb exist, in practice a progressive near-wall mesh refinement is the best to find the optimal trade-off between accuracy and computation time.

only slightly impacting 2D problems, this can be detrimental for 3D simulations. Therefore, when turbulent flow is discussed, this thesis adopts the Realizable $k-\varepsilon$ model with standard wall functions.




Figure D.1.: Number density and Mach number evolution along the axis of a de Laval nozzle with $P = 20$ bar, $t = 40 \mu\text{m}$, $D = 100 \mu\text{m}$ and $L = 200 \mu\text{m}$ for different flow models. The nozzle throat is placed at $z = -200 \mu\text{m}$, the outlet at $z = 0$. The bottom panels zoom in the shock region for a better visibility. The realizable $k-\varepsilon$ model includes the standard wall function of equation D.4,

APPENDIX E

Spectrally-resolved pulse characterisation

A careful reader may have noticed how the temporal structure of the FROG-retrieved pulse in the laser laboratory, registered by diverting the beam directly at the exit of the second multi-pass cell (MPC2), is significantly different to the profile measured at the interaction point, in the radiation laboratory. Specifically, the pulse FWHM temporal duration increases of a factor ~ 1.5 between the two laboratories, as shown in figure E.1. Addressing the causes of such change is non trivial and it is still subject of investigation. However, in the following, a preliminary study is presented, with the aim of characterising eventual wavelength-dependent properties, which could ultimately limit the pulse compressibility in the radiation laboratory, close to the interaction point.




Figure E.1.: Comparison of the retrieved pulse temporal structure in the two laboratories. The FWHM duration of the main peaks are respectively 9.6 fs and 14.7 fs.

When compression factors as high as the ones discussed in this thesis are considered, long propagation distances (as between the two laboratories) could affect the pulse temporal duration. In this context, characterising the different pulse frequency components close to the interaction point can provide valuable insights on the pulse compressibility.

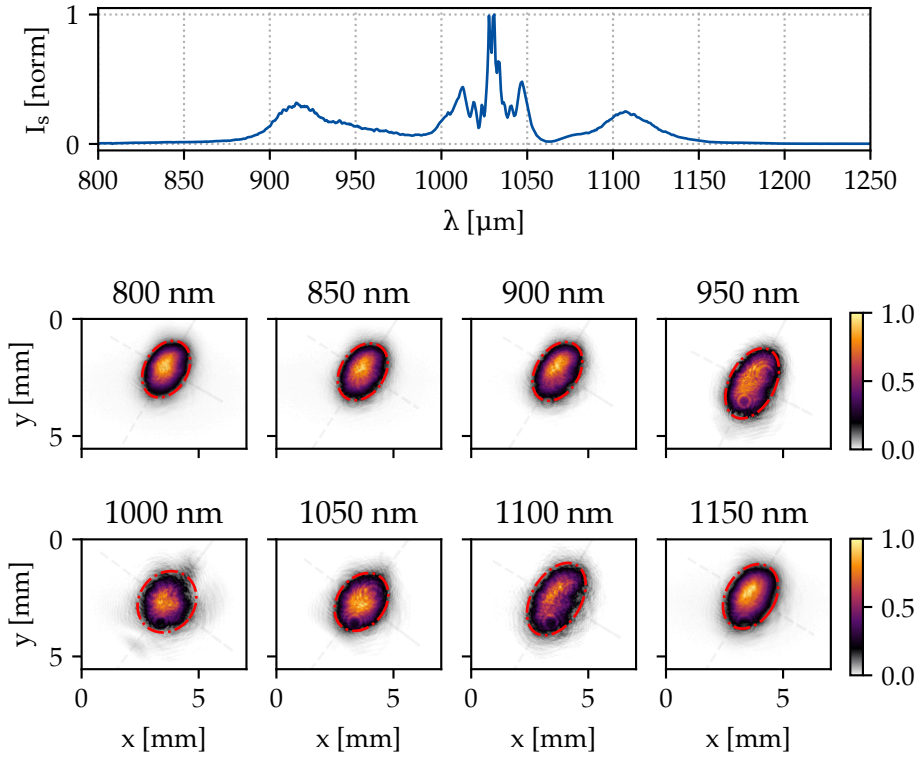


Figure E.2.: Normalised pulse near field images for different (10 nm wide) bandpass filters. Their relative intensity can be inferred according to the spectrum in the top panel. The red ellipses represent the corresponding $1/e^2$ intensity widths.

λ [nm]	800	850	900	950	1000	1050	1100	1150
x_c [mm]	3.28	3.35	3.35	3.35	3.31	3.30	3.30	3.28
y_c [mm]	1.95	2.05	2.04	2.50	2.42	2.43	2.32	2.21
ϵ	1.41	1.40	1.41	1.64	1.16	1.31	1.6	1.42

Table E.1.: Characteristic parameters for the red ellipses in figure E.2, namely the ellipses centre (x_c, y_c) and the ellipticity (ϵ) .

Figure E.2, for instance, shows the normalised, background-subtracted near field images registered for different bandpass filters (10 nm band), whose relative intensity can be inferred according to the spectrum in the top panel¹. For each wavelength, five different images integrating a single laser pulse were acquired and averaged. The red ellipse in each panel represent the corresponding $1/e^2$ intensity width. The centre of the ellipses (x_c, y_c)

¹Note that, in the images, the beam has been demagnified to entirely fit the camera sensor.

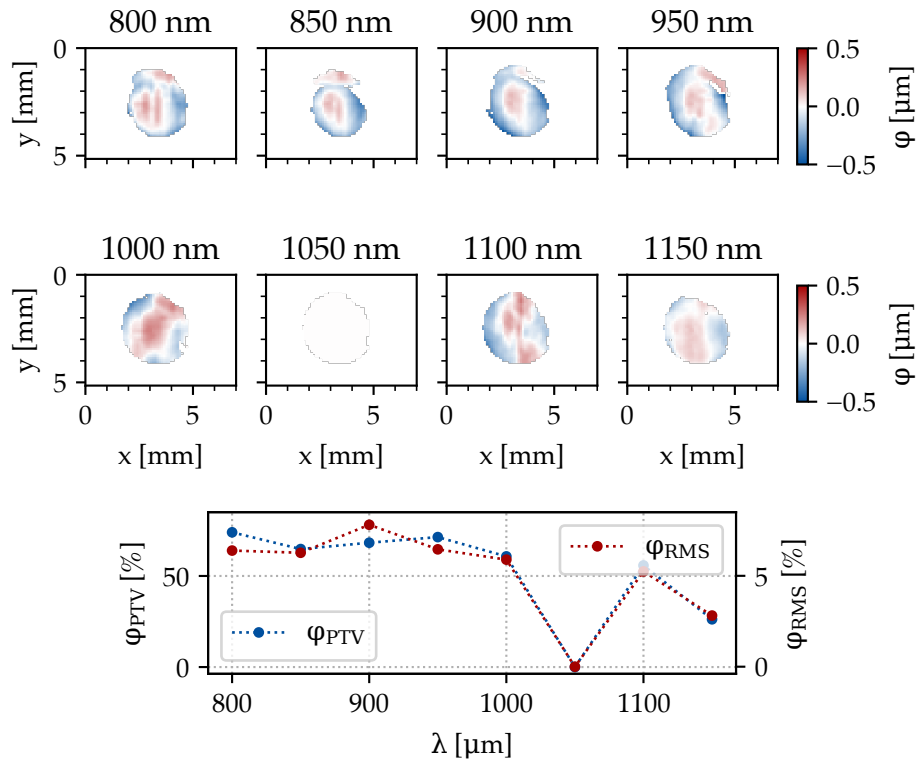


Figure E.3.: Spectrally-resolved wavefront measurement with reference to the wavefront at $\lambda = 1050$ nm. The top panel shows the relative wavefront value in μm , the bottom panel depicts the corresponding peak-to-valley (PTV, blue) and root mean square (RMS, red) values as a percentage of their respective wavelength.

and their respective ellipticity (ϵ) are reported in table E.1.

Analysing the near field images, a strong ellipticity change is evident at 950 nm and 1100 nm. Away from the laser central wavelength (i.e. at 800, 850, 900 and 1150 nm) the pulse appears shifted on the vertical axis. Moreover, for the same wavelengths, the peak of the intensity profile follows a similar trend, moving towards the top-half of the beam. This simple measurement already highlights a spectral dependence of the spatial properties of the pulse, commonly known as “chromatism”. When relevant, particularly for ultra-broad band pulses, chromatism can affect also the temporal properties of the beam, giving rise to a series of phenomena named “spatio-temporal (or spatio-spectral) couplings” (STC) [111]. To properly assess the influence of chromatism on ultrashort pulses, an in-depth investigation would be needed, involving metrology able to fully characterise the spatio-temporal structure of the pulse. Conversely, this preliminary study aims at addressing their presence, rather than quantifying their effect. Therefore following the methodology of [112], the filter scheme introduced before is used for a spectrally-resolved modal decomposition of the pulse wavefront through Zernike polynomials. For each wavelength, the pulse wavefront is

registered with a HASO4 wavefront sensor (Imagine optics) and decomposed to analyse the wavelength-dependence of the wavefront aberrations.

Comparing through Zernike polynomials pulses of different sizes and shapes is particularly challenging. The coefficients of the decomposition, in fact, significantly depend on the radius of the unit circle over which the polynomials are defined and on its centre. Therefore, for a proper comparison, it is of utmost importance for the wavefronts to be defined over the same unit circle. In the following measurement, the circle inscribed in the near field profile of the broadband pulse (without any bandpass filter) is considered as unit circle for all the analysed wavelengths. This approach allows the comparison, yet it restricts the measurement to only a fraction of the pulse, generally reducing the absolute value of the wavefront aberrations. Furthermore, the impact of a nearfield shift, evident for some of the wavelengths in figure E.2, will be underestimated, as the centre of the unit disk will eventually not coincide with the geometrical centre of the measured pulse². In this context, figure E.3 illustrates the wavelength-dependent wavefronts, relative to $\lambda = 1050$ nm, taken as a reference as it is the most intense among the measured bands. The bottom panel, specifically, outlines the corresponding peak-to-valley (PTV, blue) and root mean square (RMS, red) values as a percentage of the wavelength. The relative RMS value is always below 8%, which generally is an indication of overall contained STCs.




Figure E.4.: Wavelength-dependence of the first three Zernike coefficients, decomposed from the wavefronts in figure E.3 and normalised to have modulus π .

The Zernike decomposition of the measured wavefronts, provides a more complete picture, addressing single aberrations. Figure E.4, for instance, plots the evolution of the first three Zernike coefficients (normalised to have modulus π over the unit disk). In this representation, a significant STC would appear as a clear dependence of the coefficient with respect to the wavelength [111].

²Generally, the unit circle should be as big as the pulse in order to fully catch all the aberrations. However here we are dealing with elliptical beams of different sizes. A too big unit circle would mean to extrapolate the Zernike coefficients in the region where no (or minimal) light is present, which can easily lead to misleading results, particularly in a comparative analysis.

As mentioned before, the unit disk chosen for the wavefront measurement generally reduces the magnitude of the retrieved aberrations. While being small, though, their value is non-constant, suggesting the presence of small couplings. The vertical tilt, in particular, exhibits a (small) linear dependence on the pulse spectrum, that could eventually point towards a small angular dispersion on the vertical axis. Yet, in this appendix, the discussion remains limited to a qualitative estimation, as more points and a wider unit disk would be needed to properly quantify the possible angular dispersion.

The simple preliminary measurements discussed, hence, show that the vertical axis of the beam could be subjected to both low and high order couplings which could affect its compressibility in the radiation laboratory, after a long vacuum propagation. As mentioned before, however, the causes leading to a different compressibility in the two laboratories are still under investigation and a complete characterisation of the spatio-temporal structure of the pulse is planned before the next experimental campaign.

APPENDIX F

Interferometry

SECTION E1

Taylor expansion of the plasma refractive index

As described in section 6.3.1, interferometry typically measures the plasma refractive index, rather than directly its density. To retrieve the density value then, a linear relation is assumed, expanding the plasma refracting index as a Taylor series and considering only the first term:

$$\eta = \left(1 - \frac{n_e}{n_c}\right)^{1/2} \simeq 1 - \frac{n_e}{2n_c}, \quad (\text{E.1})$$

where n_e is the plasma density and n_c is the critical plasma density. Often this approximation is considered valid only when $n_e \ll n_c$, which may suggest considering also the quadratic term in the expansion when high plasma densities (on the order of $1 \times 10^{20} \text{ cm}^{-3}$) are considered. However, as shown in figure E.1, the error due to the linear assumption is minimal and stays below 1 % for $n_e < 2.5 \times 10^{20} \text{ cm}^{-3}$. For reference, the critical density for a laser with central wavelength $\lambda_0 = 1030 \text{ nm}$ is $n_c = 1.04 \times 10^{21} \text{ cm}^{-3}$.




Figure E.1.: The exact plasma refractive index (blue) is compared to its first order approximation (red dashed) in dependence of the plasma density n_e .

Peak density and density scans

All along the thesis, whenever a density measurement is considered, it is usually a combination of different camera shots, registered scanning the focal position along the source and with different scaling factors applied. As a result, it can be difficult for the reader to understand what measurement and what points are considered to trace a density profile. Therefore, this section briefly outlines the procedure used in the analysis.




Figure F2.: Schematic of a transverse (left) and a vertical (right) scan.




Figure E3.: Typical density profiles obtained analysing a single scan step. The first 3 steps of a vertical scan are shown as an example.

Generally a density profile is registered scanning the main laser around the source. Figure F2 serves as an example. It outlines the schematic of a typical transverse and vertical

scan. Each scan is composed of a number of steps and for each step different camera shots are acquired (usually 10), each one integrating a single pulse of the main laser. A single interferometric image, once analysed, will then result in a line integrated density profile. Averaging over all the profiles acquired in a single step, the density of that specific step can be retrieved. Figure E3, for instance, depicts the first 3 steps of a vertical scan in 25 bar of Helium. The first step is taken 147 μm above the nozzle, the subsequent ones are spaced by 10 μm . As expected, each step results in a gaussian-like density profile, whose peak density decreases moving away from the nozzle.




Figure E4.: Example showing the evolution of the (peak) density profile along a vertical scan.




Figure E5.: Interpolated 3D density profile (normalised) 150 μm above a subsonic nozzle.

When characterising a source, the density evolution along the scan is usually more important than the profile at a specific location. Therefore, in a single step the peak density (highlighted in red in figure E3) is extracted, with its correspondent standard deviation. Figure E4, then, shows the vertical evolution of the (peak) plasma density, where errors have been omitted for simplicity.

Lastly, during a single scan, given the typical dimensions of both the plasma source and the laser focal spot, the main laser can be slightly misaligned with respect to the centre of the nozzle. Hence a transverse and a longitudinal scan are performed and analysed in the same way so that the peak density registered in the first scan can be scaled according to the peak of the others. As a result trends like the ones depicted in section 6.3.2 can be obtained. In addition to this, on the basis of the 3 different scans a 3D grid can be meshed and a 3D plasma density profile can be interpolated. As an example, figure F.5, shows the normalised 3D peak density 150 μm above the tip of a subsonic nozzle.

APPENDIX G

Mode decomposition of real focal spot images

All along the thesis, the PIC simulations discussed have assumed the radial intensity distribution of the laser pulse to be an ideal gaussian. Yet high-power lasers frequently suffer from intensity inhomogeneities and wavefront aberrations, often resulting in less ideal spatial profiles. Asymmetries, particularly around the focal plane, can highly affect the laser-plasma interaction, the wakefield excitation and, eventually, the characteristics of the accelerated electron beam. While gaussian-pulses can closely match experimental results, as demonstrated in section 6.4, incorporating spatial asymmetries can significantly enhance the fidelity of PIC simulations.

To maintain the computational efficiency of cylindrical symmetry (as opposed to full 3D propagation), a real focal spot image of the laser pulse could be decomposed into various azimuthal modes and incorporated into the simulation. However, directly decomposing a raw camera image is complex and can lead to misleading results. Real images, in fact, are prone to noise and artifacts coming from the imaging system and, moreover, the limited dynamic range of a camera sensor may not capture fine intensity variations, particularly in low-intensity regions.

In this context, reconstructing raw focal images using Hermite-Gaussian (HG) modes can be particularly beneficial. By using a series of orthogonal functions, HG decomposition accurately represents the spatial profile of the pulse and filters out noise and artifacts. Additionally, it addresses the dynamic range limitations of the sensor, capturing finer details. Subsequently the reconstructed image can be further processed and decomposed in azimuthal modes for being used in PIC simulations, ensuring that the laser profile is modeled with greater accuracy. Figure G.1 serves as an example. Here, a typical laser focal spot image is reconstructed with 40 HG modes: the image noise gets evidently reduced and, as shown in the last panel, the reconstruction error is only of few %. The reconstructed image can then be efficiently decomposed in azimuthal modes. Figure G.2 illustrates the radial (normalised) intensity profile constructed using five azimuthal modes. The decomposition of the reconstructed HG spatial profile (red) is compared to the equivalent decomposition

done directly on the original image. The difference is particularly evident around the pulse peak, where the camera noise gives rise to artificial modulations which could then affect the accuracy of the subsequent PIC simulation.




Figure G.1.: HG modes reconstruction of the VERA laser focus. The first panel shows the normalised intensity profile (registered image), the second depicts its corresponding HG reconstruction, where 40 modes have been used and the last outlines their relative percentage error.




Figure G.2.: The result of the azimuthal modes decomposition of the HG reconstructed profile (red) is compared to the outcome of the modes decomposition applied directly on the registered camera image (grey). In both cases 5 modes are used.

Having decomposed the pulse into azimuthal modes, it is of particular interest to illustrate its evolution in vacuum, as it can be significantly different from an ideal gaussian pulse. Figure G.3 serves as an example comparing the two propagations, where the laser focal position is assumed to be at $z_f = 200 \mu\text{m}$. The two evolutions are markedly different. This kind of difference can be crucial in the laser plasma interaction, especially with sub-TW pulses, when ionisation induced defocusing plays a non-negligible role.




Figure G.3.: Normalised peak intensity evolution in vacuum for mode-decomposed pulse (red) compared with an ideal gaussian pulse (orange, dotted).

Acknowledgments

Surviving a PhD is hard and it is definitely a team sport. Without the incredible support I had, none of this would have come to fruition. Therefore, here is a heartfelt thank you to everyone who stood by me with their encouragement, wisdom, and (mostly terrible) humor.

Wim, it is unbelievable how inspiring and motivating a short meeting with you can be. Thank you for your guidance and trust over these years. Keep inspiring the next generation of scientists!

Inga and **Katalin**, you truly deserve special recognition for all the help you've provided. We probably wouldn't have been able to organize the defence without you.

The most sincere thank you to **Jens** and **Ilka**, for having always been available for a chat, offering their invaluable support and wise advice.

A huge shoutout is due to the entire VERA team, particularly to **Supriya**, **Andy**, **Nikita**, **Tae**, and **Christoph**. I am particularly happy with my PhD, largely because of the people I have worked with. You guys have been amazing and none of this would have been possible without your help. If I'm honest, one of you is a bit too Russian...the others though, they're class.

The whole plasma group at DESY is extraordinary. Here, you really feel like part of a big family, among friends rather than just colleagues. A special thank you to my fellow PhD mates™ **Felipe**, **Advait** and **Juditha**, because shared pain always feels less painful. Thanks to **Martin**, for being the best office mate one could ask for, to **Felix**, which is now basically fluent in Italian, and to **Paul**, for having showed me the Tarifa way. I reckon I should mention here **Cristina** as well, I don't want to deal with her complaining about not having been included...

A massive thank you to **Antonio** and **Pippo**, because you always need a little bit of Italy, even when you live abroad.

A very big hug to **Simon**, **Jon**, **Richard**, and **Lewis**. Without you guys, covid would have hit much harder. Thanks for all the good moments, the jokes, and the pints.

Last but not least, **Kris** and **Rob**. You probably know better than me how hard it is to move to a new country, adapt to a new culture, and start everything from scratch. You made me feel at home since day one, and I will never be grateful enough for that. As your first PhD student, I believe it is important to let you know that you did an amazing job in every single aspect (though there's still a looot of room for improvement, but we can discuss that later). If I'm not the worst at my job today, it is mainly because of you. Thanks, lads. I guess now we can actually have the one we came for.

