



LUND UNIVERSITY

Extreme Electron Beams and Brilliant X-rays

Generation, Manipulation and Characterization of Relativistic Electron Beams for and from Plasma-Based Accelerators

Björklund Svensson, Jonas

2020

Document Version:

Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (APA):

Björklund Svensson, J. (2020). *Extreme Electron Beams and Brilliant X-rays: Generation, Manipulation and Characterization of Relativistic Electron Beams for and from Plasma-Based Accelerators*. [Doctoral Thesis (compilation), Lund University]. Department of Physics, Lund University.

Total number of authors:

1

Creative Commons License:

Unspecified

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00



Extreme Electron Beams and Brilliant X-rays

Generation, Manipulation and Characterization of Relativistic
Electron Beams for and from Plasma-Based Accelerators

JONAS BJÖRKLUND SVENSSON | FACULTY OF ENGINEERING | LUND UNIVERSITY





LUND UNIVERSITY
Faculty of Engineering, LTH
Department of Physics
Division of Atomic Physics

ISBN: 978-91-7895-618-0 (print)

ISBN: 978-91-7895-619-7 (pdf)

ISSN 0281-2762

Lund Reports on Atomic Physics, LRAP 567 (2020)



Extreme Electron Beams and Brilliant X-rays

Generation, Manipulation and Characterization of
Relativistic Electron Beams for and from
Plasma-Based Accelerators

Jonas Björklund Svensson



LUND
UNIVERSITY

Academic thesis which, by due permission of the Faculty of Engineering at Lund University, Sweden, will be publicly defended on September 25th, 2020, at 9.15 in the Rydberg hall at the Department of Physics, Professorsgatan 1, Lund, for the degree of Doctor of Philosophy in Engineering.


Faculty opponent:

Dr. Patric Muggli

Max-Planck-Institut für Physik, Munich, Germany

Organization LUND UNIVERSITY Atomic Physics Department of Physics P.O. Box 118 SE-112 00 Lund, Sweden		Document name Doctoral Dissertation	
Author Jonas Björklund Svensson		Date of issue September 1 st , 2020	
		Sponsoring organization	
Title and subtitle <i>Extreme Electron Beams and Brilliant X-rays - Generation, Manipulation and Characterization of Relativistic Electron Beams for and from Plasma-Based Accelerators</i>			
Abstract <p>This thesis is based on work done by the author on the development of plasma-based electron accelerators driven by ultra-intense laser pulses and dense electron bunches. Plasma based accelerators have several benefits, such as accelerating fields around 1000 times stronger than in "conventional" radio-frequency accelerators, which can allow for shrinking the overall footprint of the accelerator. They can also allow for generating electron beams with unprecedented peak currents and ultra-low emittances, meaning that a large number of electrons can be packed into a very short time duration and that the quality of the bunches is high. They can also be used to generate X-ray pulses with durations only otherwise achievable at a few large accelerator facilities, using a laboratory setup the size of a large living room. These characteristics make plasma-based accelerators interesting as a technology for future particle colliders and free-electron lasers, as well as, for example, smaller and more available X-ray sources with particular source characteristics such as ultra-short pulse durations.</p> <p>This thesis describes both numerical and experimental studies on plasma-based accelerators. The experimental work has mainly been on generating electron bunches and X-ray pulses using a laser-wakefield accelerator, as well as applications of the generated X-rays. The results from this branch of the research include the identification and demonstration of a new regime for laser-driven X-ray generation, which produces pulses with significantly reduced divergence compared to the standard method, simplifying the subsequent use of such pulses in applications.</p> <p>The numerical work has been focused towards conventional radio-frequency accelerators, concerning the shaping of electron bunches from such an accelerator for use in electron beam-driven plasma-wakefield acceleration. The main point in this research has been removing or circumventing detrimental effects that occur during acceleration and transport, to create bunches which can drive stable wakes. One of the results from this research is an optimization strategy for certain bunch compressors, leading to a decrease in chromatic and geometric aberrations in the bunch. The common thread through both experimental and numerical work is plasma-based acceleration of electrons, and as such there is a larger overlap between these two parts than might initially be seen.</p>			
Key words Laser, plasma, wakefield, acceleration, accelerator, relativistic, electrons, x-rays, ultra-fast, betatron			
Classification system and/or index terms (if any)			
Supplementary bibliographical information		Language English	
ISSN and key title 0281-2762, Lund Reports on Atomic Physics, LRAP 567		ISBN 978-91-7895-618-0 (print) 978-91-7895-619-7 (pdf)	
Recipient's notes		Number of pages 255	
		Security classification	

I, the undersigned, being the copyright owner of the abstract of the above-mentioned dissertation, hereby grant to all reference sources permission to publish and disseminate the abstract of the above-mentioned dissertation.

Signature 

Date: 2020-08-20

EXTREME ELECTRON BEAMS AND BRILLIANT X-RAYS

GENERATION, MANIPULATION AND CHARACTERIZATION OF RELATIVISTIC
ELECTRON BEAMS FOR AND FROM PLASMA-BASED ACCELERATORS

Jonas Björklund Svensson

Doctoral Thesis
2020



LUND
UNIVERSITY

EXTREME ELECTRON BEAMS
AND BRILLIANT X-RAYS

Paper I: © 2019 CC BY 4.0
Paper II: © 2020 The Authors
Paper III: © 2018 IEEE
Paper IV: © 2020 The Authors
Paper V: © 2018 CC BY 3.0
Paper VI: © 2020 The Authors
Paper VII: © 2017 AIP Publishing
Paper VIII: © 2018 The Optical Society of America
Paper IX: © 2020 The Optical Society of America
Paper X: © 2020 The Authors

For all other material: © 2020 Jonas Björklund Svensson
All rights reserved
Printed in Sweden by Media-Tryck, Lund, 2020

Division of Atomic Physics
Department of Physics
Faculty of Engineering, LTH
Lund University
P.O. Box 118
SE-221 00 Lund
Sweden
www.atomic.physics.lu.se

ISSN 0281-2762
Lund Reports on Atomic Physics, LRAP 567 (2020)
ISBN 978-91-7895-618-0 (PRINT)
ISBN 978-91-7895-619-7 (PDF)

“You have to be realistic about these things.”

LOGEN NINEFINGERS

ABSTRACT

This thesis is based on work done by the author on the development of plasma-based electron accelerators driven by ultra-intense laser pulses and dense electron bunches. Plasma based accelerators have several benefits, such as accelerating fields around 1000 times stronger than in “conventional” radio-frequency accelerators, which can allow for shrinking the overall footprint of the accelerator. They can also allow for generating electron beams with unprecedented peak currents and ultra-low emittances, meaning that a large number of electrons can be packed into a very short time duration and that the quality of the bunches is high. They can also be used to generate X-ray pulses with durations only otherwise achievable at a few large accelerator facilities, using a laboratory setup the size of a large living room. These characteristics make plasma-based accelerators interesting as a technology for future particle colliders and free-electron lasers, as well as, for example, smaller and more available X-ray sources with particular source characteristics such as ultra-short pulse durations.

This thesis describes both numerical and experimental studies on plasma-based accelerators. The experimental work has mainly been on generating electron bunches and X-ray pulses using a laser-wakefield accelerator, as well as applications of the generated X-rays. The results from this branch of the research include the identification and demonstration of a new regime for laser-driven X-ray generation, which produces pulses with significantly reduced divergence compared to the standard method, simplifying the subsequent use of such pulses in applications.

The numerical work has been focused towards conventional radio-frequency accelerators, concerning the shaping of electron bunches from such an accelerator for use in electron beam-driven plasma-wakefield acceleration. The main point in this research has been removing or circumventing detrimental effects that occur during acceleration and transport, to create bunches which can drive stable wakes. One of the results from this research is an optimization strategy for certain bunch compressors, leading to a decrease in chromatic and geometric aberrations in the bunch. The common thread through both experimental and numerical work is plasma-based acceleration of electrons, and as such there is a larger overlap between these two parts than might initially be seen.

POPULAR SCIENCE SUMMARY

Relativistic electrons, meaning electrons travelling at close to the speed of light, have many interesting uses. The foremost use-case today is probably for light-sources, where energetic “bunches” of electrons are used to create pulses of light within a broad spectral range, or range of “colors”. The short wavelength of this electromagnetic radiation, in the ultra-violet or X-ray range, allows for studying small and/or fast-moving things, such as the molecular structure of pharmaceuticals, or the electrons moving around in molecules. The work-horses among the light-source facilities are the synchrotron light sources, typically circular accelerators accommodating dozens of simultaneous research groups. The next step up in the hierarchy of light sources are free-electron lasers, which can generate even brighter and shorter bursts of light than the circular machines (which allows taking snapshots of ultra-fast phenomena), but can only serve a few simultaneous research groups.

The other “big” use-case is for particle colliders using, for example, electrons and their anti-matter cousins, positrons. Particle colliders are used for fundamental physics research, to gain knowledge about what our universe is made of and what holds it together at a fundamental level. However, the radiation that is so useful in light sources is detrimental if you want to build a particle collider. For the latest while, protons have been used in circular colliders, but future machines often aim to use electrons and positrons. Because of the issues with radiation from these particles, the machine designs are often linear instead of circular.

The common problem for these two applications is that the different facilities are often very large and expensive - one of the proposed designs for a future circular collider is 100 km in circumference, which dwarfs the already gigantic 27-km-circumference Large Hadron Collider (LHC) which is already the largest ever machine built by humans. Therefore, researchers have been looking towards alternative accelerator concepts, which might help shrinking these machines to reduce cost and complexity. One of these concepts is plasma-based accelerators, which this thesis is centered around. A plasma is an ionized material, where the atomic nuclei and their orbiting electrons have been separated and turned into a kind of “soup”. As it turns out, it is possible to make waves in this soup which provide a near-perfect environment for accelerating electrons at an un-matched rate - the accelerating electric fields are around 1000 times stronger than using conventional technology - which has the potential to shrink the machines by a similar factor. These waves can be driven by ultra-intense laser pulses and electron bunches. However, the technology is not yet mature and there are fundamental issues to be solved still. The results presented in this thesis are contributions

to this end - to shrink electron accelerators and enable all the benefits that comes with it.

The first of the topics studied within this thesis is how to manipulate and shape the electron bunches from “conventional” accelerators to be suitable for use in plasma-based acceleration. Conventional accelerators can provide stable and repeatable bunches, but since plasma-based accelerators are often very sensitive to the initial parameters, a fine level of control of different aspects is required to yield stable and repeatable bunches also after acceleration in the plasma. For example, a bunch which is meant to drive a plasma wave must be straight when propagating into the plasma - more like, say, a cigar than a snake or a banana. If not, the whole bunch can begin to jiggle around, creating a turbulent environment which is not suitable for acceleration. This level of control of the electron bunches is also an advantage in free-electron lasers, which can operate with different types of bunches, but which nonetheless benefit from the increased level of control. Included in this first topic is also the generation and manipulation of electron beams from a laser-driven plasma-based accelerator, where the challenges are different from those in a conventional accelerator, for example because the generated beams are more divergent.

The second topic is the creation of “betatron” X-rays using electrons from laser-driven plasma-based accelerators. When inside the plasma, the electrons perform rapid sideways oscillations, which allows for creating ultra-short X-ray pulses with photon energies in the same range as large synchrotron facilities, but using a laboratory-scale setup. Other aspects of the betatron radiation differs much from the X-rays generated using synchrotron facilities, and some might say that the quality is not as good, but the usability of the radiation is entirely dictated by the application. In the thesis, both generation and applications of this radiation are studied. Among other things, a novel regime of laser-plasma-based X-ray generation was identified and demonstrated, yielding X-ray pulses with significantly smaller divergence than betatron X-rays, and betatron X-rays were used for one of the first applications with a truly industrial connection as a diagnostic for atomizing sprays.

Finally, the third topic is regarding how to characterize, or measure, different parameters of the bunches going into and coming out of plasma-based accelerators. To be able to draw any conclusions regarding the behavior of the bunches and their interaction with the plasma, their properties must be measured. Properly measuring the properties of bunches for and from plasma-based accelerators is not easy and incurs several difficulties which must be considered when designing the diagnostics setups. Within this thesis, mainly measurements of the energy spectrum and time structure of such bunches are considered.

POPULÄRVETENSKAPLIG SAMMANFATTNING

Relativistiska elektroner, alltså elektroner som rör sig med nära ljusets hastighet, har många intressanta användningsområden. Den främsta tillämpningen idag är troligtvis i ljuskällor, där högenergetiska "buncher" av elektroner används för att skapa ljuspulser inom ett brett spektrum. Den korta våglängd, ofta ultraviolett eller röntgen, som denna elektromagnetiska strålning har tillåter studier av små och/eller snabba objekt, som den molekylära strukturen i ett läkemedel, eller elektronerna som rör sig runt i molekyler. Arbetshästen bland den här typen av ljuskällor är synkrotronljuskällan, typiskt stora anläggningar med en cirkulär accelerator som kan försörja flera forskargrupper med ljus samtidigt. Ett steg upp i hierarkin finns de linjära frielektronlasrarna, som kan generera kortare och ljusstarkare ljuspulser än de cirkulära maskinerna, men som bara kan försörja ett fåtal forskargrupper samtidigt.

Det andra "stora" användningsområdet är partikelkolliderare (colliders), som kan använda sig av till exempel elektroner och deras antimateriakusiner, positroner. Colliders används för grundläggande fysikforskning, för att öka kunskapen om vad vårt universum är gjort av och vad som håller ihop det, på en fundamental nivå. Dock är den strålning som är så användbar i ljuskällor en käpp i hjulet när det gäller colliders. Den senaste tiden har protoner använts i cirkulära colliders, men framtida maskiner siktar ofta på att använda elektroner och positroner. På grund av problemen med strålning från dessa partiklar är designerna för framtida maskiner ofta linjära istället för cirkulära.

Det gemensamma problemet för båda tillämpningarna är att dessa anläggningar ofta är väldigt stora och dyra - ett av förslagen på en framtida cirkulär collider är 100 km i omkrets, vilket kan jämföras med den redan gigantiska Large Hadron Collider (LHC), som med sina 27 km i omkrets redan är den största maskinen som människan någonsin byggt. Därför har forskare letat efter alternativa accelerationsmetoder, som skulle kunna hjälpa till att krympa dessa maskiner och på så sätt minska kostnad och komplexitet. Ett av dessa koncept är plasmabaserade accelerators, vilket är vad den här avhandlingen handlar om. Ett plasma är ett joniserat material, där atomkärnorna och deras elektroner har separerats och bildat ett slags soppa. Det har visat sig att det är möjligt att bilda vågor inuti den här soppan som utgör nästan perfekta förhållanden för att accelerera elektroner otroligt fort - de accelererade elektriska fälten är omkring 1000 gånger starkare än med konventionell teknik - vilket har potentialen att krympa maskinerna med motsvarande faktor. Dessa vågor kan drivas med både

ultra-intensiva laserpulser och elektronbuncher. Dock är tekniken inte mogen än, och det finns fortfarande fundamentala problem att lösa. Resultaten som presenteras i den här avhandlingen är bidrag till det här ändamålet - att krympa elektronacceleratorer och tillgängliggöra alla de fördelar som kommer med detta.

Den första kategorin av ämnen som studerats inom avhandlingen handlar om hur man kan manipulera och forma elektronbuncher från "konventionella" acceleratorer för att göra dem lämpliga för användning i plasmabaserad acceleration. Konventionella acceleratorer kan producera stabila och repeterbara elektronbuncher, men eftersom plasmabaserade acceleratorer ofta är väldigt känsliga för bunchernas initiala parametrar krävs en finkänslig kontrollnivå för att buncherna ska vara stabila även efter acceleration i plasmata. Till exempel bör en bunch som ska driva en plasmavåg vara helt rak när den propagerar in i plasmata - mer som, säg, en cigarr än en orm eller banan. Om den inte är det kan hela bunchen börja svänga okontrollerat, vilket skapar en väldigt turbulent miljö som inte är lämplig för acceleration. Denna nivå av kontroll över buncherna är också en fördel för frielektronlasrar, som visserligen använder andra typer av buncher, men som oavsett drar nytta av den ökade kontrollen. I denna första kategori ingår även generation och manipulation av elektronstrålar från en laserdriven plasmabaserad accelerator, där utmaningarna skiljer sig från en konventionell accelerator bland annat på grund av att strålarna är mer divergenta.

Den andra kategorin handlar om skapandet av "betatronstrålning", en typ av röntgenstrålning, från laserdrivna plasmaacceleratorer. Inuti plasmata oscillerar dom accelererande elektronerna snabbt i sidled, vilket producerar ultrakorta röntgenpulser med fotonenergi i samma område som stora synkrotronanläggning, men från en uppställning i laboratoriestorlek. Andra aspekter av betatronstrålningen skiljer sig från strålningen som produceras vid synkrotroner, och vissa skulle nog säga att kvalitén inte är lika bra, men strålningens användbarhet bestäms helt och hållet av tillämpningen. I den här avhandlingen har både produktion och tillämpningar av laser-plasmagenererad röntgenstrålning studerats. Bland annat har en ny regim för laserbaserad röntgenproduktion upptäckts och demonstrerats, vilken kan ge röntgenstrålning med betydligt lägre divergens än betatronstrålning, och betatronstrålning har använts för en av de första tillämpningarna med en faktisk anknytning till industri, för diagnostik av atomiserande sprayer.

Slutligen handlar den tredje och sista kategorin om hur man karakteriserar, eller mäter, olika parametrar hos elektronbuncherna som åker in i och ut ur plasmabaserade acceleratorer. För att kunna dra några slutsatser kring hur buncherna uppför sig och interagerar med plasmata måste man mäta deras egenskaper. Det är inte enkelt att mäta plasmaaccelererade elektronbunchers egenskaper på ett bra sätt och det medför flera svårigheter som måste övervägas när man designar mätuppställningarna. Inom den här avhandlingen har framförallt mätningar av energispektrum och tidsstruktur hos dessa buncher gjorts.

LIST OF PUBLICATIONS

This thesis is based on ten papers which are appended at the end of the thesis, and referred to by their Roman numerals.

I Third-order double-achromat bunch compressors for broadband beams

J. Björklund Svensson, T.K. Charles, O. Lundh, and S. Thorin.
Physical Review Accelerators and Beams **22**, 104401 (2019).

II Arc-like variable bunch compressors

P. H. Williams, G. Pérez-Segurana, I. R. Bailey, S. Thorin, B. Kyle, and J. Björklund Svensson.
arXiv:2006.15051v1, Manuscript under review (2020).

III Beamline design for plasma-wakefield acceleration experiments at MAX IV

J. Björklund Svensson, H. Ekerfelt, O. Lundh, E. Mansten, J. Andersson, M. Kotur, F. Lindau, S. Thorin and T.K. Charles.
Proceedings of the Advanced Accelerator Concepts Conference (AAC), Breckenridge, CO, USA (2018).

IV Start-to-end simulations of the MAX IV linear accelerator towards plasma-wakefield acceleration experiments

J. Björklund Svensson, J. Andersson, T.K. Charles, H. Ekerfelt, J. Ferri, M. Kotur, F. Lindau, E. Mansten, S. Thorin and O. Lundh.
Manuscript in preparation (2020).

V Double-bunches for two-color soft X-ray free-electron laser at the MAX IV laboratory

J. Björklund Svensson, J. Andersson, F. Curbis, M. Kotur, F. Lindau, O. Lundh, E. Mansten, S. Thorin and S. Werin.
Proceedings of the Free-Electron Laser Conference (FEL), Santa Fe, NM, USA (2017).

VI Low-divergence, femtosecond X-ray pulses from a passive plasma lens

J. Björklund Svensson, H. Ekerfelt, J. Ferri, I. Gallardo González, D. Guénot, A. Persson, K. Svendsen, L. Veisz and O. Lundh.
Manuscript under review (2020).

VII Highly efficient angularly resolving X-ray spectrometer optimized for absorption measurements with collimated sources

M. Šmíd, I. Gallardo González, H. Ekerfelt, J. Björklund Svensson, M. Hansson, J. C. Wood, A. Persson, S. P. D. Mangles, O. Lundh and K. Falk.
Review of Scientific Instruments **88**, 063102 (2017).

VIII Optimization of soft X-ray phase-contrast tomography using a laser wakefield accelerator

K. Svendsen, I. Gallardo González, M. Hansson, J. Björklund Svensson, H. Ekerfelt, A. Persson and O. Lundh.
Optics Express **26**, 33930 (2018).

IX Simultaneous laser-driven X-ray and two-photon fluorescence imaging of atomizing sprays

D. Guénot, K. Svendsen, J. Björklund Svensson, H. Ekerfelt, A. Persson, O. Lundh and E. Berrocal.
Optica **7**, 131-134 (2020).

X Electron acceleration in merging laser wakefields

I. Gallardo González, H. Ekerfelt, J. Björklund Svensson, G. Gatti, A. Gonoskov, D. Guénot, M. Hansson, M. Marklund, J. A. Pérez-Hernández, C. Salgado, E. Wallin, G. Zeraoui and O. Lundh.
Manuscript in preparation (2020).

Other related publications by the author not included in this thesis:

Driver-witness-bunches for plasma-wakefield acceleration at the MAX IV linear accelerator

J. Björklund Svensson, J. Andersson, F. Curbis, H. Ekerfelt, M. Kotur, F. Lindau, O. Lundh, E. Mansten, S. Thorin, S. Werin.

Proc. IPAC'17, Copenhagen, Denmark, 2017.

First experimental measurements of the caustic nature of trajectories in bunch compressors

T. K. Charles, J. Björklund Svensson, A. Latina, S. Thorin.

Proc. IPAC'19, Melbourne, Australia, 2019.

Control of laser wakefield accelerated electron beams and betatron radiation using multistage gas targets

V. Tomkus, V. Girdauskas, J. Dudutis, P. Gečys, V. Stankevič, G. Račiukaitis, I. Gallardo González, D. Guénot, J. Björklund Svensson, A. Persson and O. Lundh.

Manuscript under review (2020).

CONTENTS

Introduction	1
1 Interactions between charged particles and electromagnetic fields	5
1.1 Charged particles and their propagation	5
1.1.1 Fields, forces and relativistic effects	5
1.1.2 The electromagnetic fields of a charged particle	7
1.1.3 Single-particle propagation	10
1.1.4 Particle beam propagation	12
1.2 Electromagnetic waves and their propagation	18
1.2.1 Ray optics	18
1.2.2 Maxwell's equations and potentials	20
1.2.3 Electromagnetic beam propagation	21
1.2.4 Electromagnetic pulse propagation	23
1.3 Plasmas	26
1.3.1 What is a plasma?	26
1.3.2 Electromagnetic waves in plasmas	27
1.3.3 Laser-driven plasma waves	31
1.3.4 Particle beam-driven plasma waves	38
1.4 Summary	42
2 Radio-frequency-based accelerators	43
2.1 Acceleration in radio-frequency waves	43
2.1.1 Radio-frequency cavities	43
2.1.2 Different types of radio-frequency structures	43
2.1.3 Energy gain in a time-harmonic voltage	44
2.2 Electron sources	46
2.3 Beam control	47
2.3.1 Longitudinal properties	47
2.3.2 Transverse properties	52
2.4 The MAX IV linear accelerator	55
3 Plasma-based accelerators	57
3.1 Acceleration in plasma waves	57
3.1.1 Limitations to the acceleration length	58
3.1.2 Betatron radiation	60
3.2 Electron sources	63

3.2.1	Internal injection mechanisms	64
3.2.2	Hybrid plasma-based accelerators	68
3.3	Beam control	69
3.3.1	Longitudinal properties	69
3.3.2	Transverse properties	71
3.4	The laser-wakefield accelerator at the Lund High-Power Laser Facility	75
4	Beam instrumentation	79
4.1	Laser diagnostics	79
4.1.1	Pulse energy	79
4.1.2	Longitudinal properties	80
4.1.3	Transverse properties	81
4.2	Electron diagnostics	82
4.2.1	Longitudinal properties	82
4.2.2	Transverse properties	94
4.3	X-ray diagnostics	97
5	Summary and outlook	99
	Article summaries and the author's contributions	103
	Acknowledgements	109
	References	111

Papers

I	Third-order double-achromat bunch compressors for broadband beams	129
II	Arc-like variable bunch compressors	138
III	Beamline design for plasma-wakefield acceleration experiments at MAX IV	149
IV	Start-to-end simulations of the MAX IV linear accelerator towards plasma-wakefield acceleration experiments	155
V	Double-bunches for two-color soft X-ray free-electron laser at the MAX IV laboratory	170
VI	Low-divergence, femtosecond X-ray pulses from a passive plasma lens	177
VII	Highly efficient angularly resolving X-ray spectrometer optimized for absorption measurements with collimated sources	200
VIII	Optimization of soft X-ray phase-contrast tomography using a laser wakefield accelerator	211
IX	Simultaneous laser-driven X-ray and two-photon fluorescence imaging of atomizing sprays	225
X	Electron acceleration in merging laser wakefields	231

INTRODUCTION

Particle accelerators have been useful tools for scientists for over 100 years; from the early electrostatic Cockroft-Walton accelerators [1], through Ising and Widerøe's radio-frequency linear accelerators [2] and Lawrence's cyclotron [3], to the big machines like the Large Hadron Collider (LHC) [4] and the Linac Coherent Light Source (LCLS) [5] of today, the applications, direct or indirect, of accelerated charged particles for science and engineering are vast. However, machines such as the ones mentioned above can be very large and expensive (for example, the LHC is 27 km in circumference and has cost around 6 billion Euro so far [4]), and their successors are projected to, in some cases, be significantly larger and more expensive [6]. The large sizes and high costs have prompted research into advanced accelerator concepts which could deliver higher accelerating fields, and thus decrease the accelerator size, ideally while providing higher-quality particle pulses (normally called *bunches*) for the given application at a higher efficiency. One such technique, and indeed the topic of this thesis, is plasma-based accelerators (PBAs), which utilize the properties of plasmas to sustain incredibly strong accelerating fields.

In this thesis, the broader field of study is electron accelerators. Electrons, being much lighter than protons and other ions, reach relativistic velocities at lower energies, but the low electron mass also means that they radiate much more strongly since their trajectories are more easily bent. This radiation can be both beneficial and detrimental. Many of the large-scale light source facilities such as synchrotrons [7] and free-electron lasers (FELs) [8, 9] are designed specifically for using electrons to generate X-rays, which can be subsequently used as, for example, a powerful microscope to study the molecular structures of proteins, medical drugs and viruses (such as in the current SARS-CoV-2 pandemic [10]). However, the energy that the electrons lose by radiating impose limits on the particle energies and machine sizes that can be reached with *circular* machines. These *synchrotron radiation* losses is one reason that the designs for future electron/positron (lepton) colliders, such as the Compact Linear Collider (CLIC) [11] and International Linear Collider (ILC) [12], are often based on *linear* accelerators (linacs), where much less energy is lost since the particles are not forced onto a circular path. The other alternative is to make the machine huge, such as the 100-km-circumference Future Circular Collider-ee (FCC-ee) [13]. These main collider designs (as well as other projects) considered by the European Strategy Group (ESG) are all based on conventional radio-frequency (RF) accelerating technology, where oscillating electro-magnetic (EM) fields in hollow, metallic cavities are used for acceleration. Whilst CLIC aims to provide a more compact option by using high accelerating fields in advanced structures, the final design is still ~ 50 km long

[11]. Therefore, the ESG has recently stated that further research and development on advanced technologies, including PBAs, should be done in an effort to reach the ambitious goals of future colliders [14].

The maximum electric field strength in conventional accelerating structures is limited by so-called RF breakdown, where the electric fields on the inner surface of the metallic structures becomes too strong, and rips electrons out from the surface, producing arcs which negate the accelerating fields, and can damage the expensive accelerating structures. In contrast, in a plasma (a “soup” of free electrons and ionized atomic nuclei), the breakdown has already occurred, and plasmas therefore have the ability to sustain much higher electric fields, normally between 100 and 10,000 times stronger than conventional RF accelerating structures. The potential thus exists for shrinking the accelerator by a corresponding factor, significantly lowering the cost and footprint of the facility. The plasma can also be used as a source of the electron bunches themselves, with the capability of producing very short bunches of high quality (high *brightness*), which could be excellent for operating for example particle colliders or FELs, but also more directly for e.g. medical applications.

But if everything just worked without any obstacles, this thesis would not have been written, because of course there are challenges with this technology, some of which I have worked on and will present in more detail in the following Chapters. One of the big challenges concerns bunch brightness preservation once the bunches have been generated. The brightness can suffer from e.g. aberrations in the focusing electron optics when transporting the bunches from the source to applications. Chromatic aberrations in electron bunches are often similar to those in photon optics, where for example the focal length of a given lens typically varies with the wavelength (or energy) of the light passing through it, or a prism disperses the light into a spectrum, separating the different wavelengths transversely. These aberrations are important to address, because another obstacle is that the energy spread of the PBA-generated bunches can be very large - a narrowband plasma-generated electron bunch can have some few percent of relative energy spread, meaning that the electrons’ energies on average deviate from the mean by a few percent, whereas for example FELs normally require less than a tenth of that. These issues are being addressed from multiple directions by researchers, from decreasing the intrinsic energy spread of the bunches [15, 16], to designing electron optics which can handle bunches with higher energy spread [17, 18], and even inventing undulator radiation and FEL schemes which circumvent the typical energy spread requirements [19, 20]. For this thesis, the two former points have been part of the work, together with work on how to characterize the electron bunches to measure different bunch parameters, such as the energy spread and time duration. The electron bunch attributes are also the reason behind the title: the bunches that have been studied are ultra-short, high-brightness, sometimes ultra-broadband and/or highly divergent, and never quite “ordinary” - typically “extreme” in one sense or another.

Outline of the thesis

This thesis is rather long, but there is a good (to my mind, at least) reason for it. The work described herein spans what has historically been two or three separate fields of research - “conventional” RF accelerators, beam-driven plasma accelerators and laser-driven plasma accelerators, where mainly the research of laser-driven plasma

acceleration was originally more separate as it originated from the field of lasers and photonics. The former accelerator type uses oscillating electromagnetic fields inside RF cavities for acceleration, whereas the two latter use the quasi-static fields inside plasma waves driven by intense particle bunches or laser pulses, respectively. As these research fields have evolved, the lines between them have started to become more and more blurry - ideas and concepts are frequently shared between the three, and the topic is now inherently multi-disciplinary. This means that many different topics must be introduced in this thesis, from accelerator physics to plasma and laser physics, and from theoretical to practical aspects. It also means that I have been fortunate to work with many different people, coming from different backgrounds, which has sometimes resulted in communication difficulties. Since the purpose of the thesis introduction (i.e. not this text, but Chapter 1-5) in a Swedish thesis is to introduce the reader to the research topic(s) of the presented Papers, I have made an effort to make it understandable for people coming from either background. This has in most cases meant going back to the fundamental physics, to then present the theory in what I hope is a comprehensible way to all. I hope that this also allows people who are not originally from an accelerator physics or photonics background to follow along to a greater extent, and thereby make the work available to a larger audience. To my mind, science should be more open and accessible, and this is one way that I could contribute to this end. I have also tried to produce a thesis I would have liked to have had access to when I was new to the field, and hopefully some future student can find useful information herein.

In Chapter 1, I introduce the basic physics that is needed to understand the different phenomena which dictate how the particle accelerators in question function. The forces on charged particles as well as the electromagnetic fields that emanate from them is covered first, followed by an introduction to the description of charged particle propagation. Following this is a description of some important general features of light optics and how intense laser pulses propagate in matter. The final puzzle piece before describing the accelerators is concerning plasmas and how the non-linear waves therein, which we want to use for accelerating electrons, can be excited. Chapter 2 and 3 applies these concepts to electron acceleration in RF accelerators and plasma-based accelerators, respectively. In Chapter 2 I show some basics concerning the acceleration in RF fields and bunch manipulation with bunch compressors, as well as some background on the origins of beam tilt and how to manage it. These manipulations are done using conventional accelerator components, such as magnets. Chapter 3 focuses on the generation of the electron bunches and X-rays in the PBAs, as well as the important interactions between the accelerating electron bunches and the plasma structures. Comparisons to conventional technology is made, and new types of strongly focusing electron optics are discussed. Chapter 4 is devoted to the description of some of the diagnostics used to study the electron bunches and X-ray pulses that can be produced with these accelerators. In particular, the spectrum and time structure of the bunches have been of interest. In Chapter 5, I give a short summary and outlook based on my work, before the research Papers are displayed in the Appendix.

INTERACTIONS BETWEEN CHARGED PARTICLES AND ELECTROMAGNETIC FIELDS

Before we get started on describing the functions of the actual particle accelerators in Chapter 2 and 3, we need to examine some fundamental concepts concerning the motion and behavior of relativistic charged particles when they interact with electromagnetic fields - and vice versa. This Chapter presents many fundamental concepts with the aim to give a feeling for the different phenomena that play a role in an accelerator. Included concepts are for example the propagation of electromagnetic waves, mainly in the form of laser pulses, as they are not only crucial for the operation of laser-driven plasma accelerators, but for many aspects of particle accelerators in general. The broad range of subtopics makes for a lengthy Chapter, and many of the fine details have been omitted for the sake of clarity, but the contents should provide a decent base for at least understanding the following Chapters in this thesis, and the Papers herein. Throughout this thesis, SI units are used.

1.1 Charged particles and their propagation

1.1.1 Fields, forces and relativistic effects

To alter the magnitude and/or direction of the velocity of an object, i.e. to accelerate it, we need to apply a force to it. The electromagnetic force, \mathbf{F} , on a charged particle in an external field is given by the Lorentz force [21],

$$\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \quad (1.1)$$

where q is the charge of the particle, \mathbf{v} is its velocity and \mathbf{E} and \mathbf{B} are the electric and magnetic field respectively. The magnetic field does no work on the particle [22], and so to increase the magnitude of its velocity, which is what is colloquially meant by “acceleration”, an electric field is required. The increase in particle energy, or energy

gain, $\Delta\mathcal{E}$, is [22]

$$\Delta\mathcal{E} = q \int_{\mathbf{r}_1}^{\mathbf{r}_2} \mathbf{E} d\mathbf{r} = qU, \quad (1.2)$$

where \mathbf{r}_1 and \mathbf{r}_2 are two points along the particle motion, and U is the voltage between these two points. Since the scope of this thesis is electron acceleration, we will henceforth use the electron charge, $q_e = -e$, where e is the unit charge, meaning that negative voltages are required for electron acceleration.

That the magnetic field does not do any work, in the physical sense, does not mean that it can be disregarded in particle accelerators - quite the contrary, in fact. Particularly when the particles reach high velocities, $v = |\mathbf{v}| \lesssim c$, where c is the speed of light, the force on the particle which is perpendicular to its direction of motion, \mathbf{F}_\perp , is easily dominated by the $\mathbf{v} \times \mathbf{B}$ -term in Eq. 1.1 from a likewise perpendicular magnetic field. For static fields, $B = |\mathbf{B}| = 1$ T is easily produced, which would require $E = |\mathbf{E}| = 300$ MV/m to generate the same force, as $c \approx 3 \times 10^8$ m/s. In electron accelerators, these high velocities are reached over a short distance because of the low mass, m_e , of the electrons. This means that relativistic effects must be taken into account almost immediately, and that quantities such as velocity become less meaningful. Instead, the kinetic energy, \mathcal{E} , or total energy, \mathcal{E}_{tot} , of the electron are more relevant quantities. The kinetic energy is associated with the relativistic gamma, or Lorentz, factor [23],

$$\gamma_b = \frac{1}{\sqrt{1 - \left(\frac{v}{c}\right)^2}} = \frac{1}{\sqrt{1 - \beta_b^2}} = \frac{\mathcal{E} + \mathcal{E}_e}{\mathcal{E}_e} = \frac{\mathcal{E}_{\text{tot}}}{\mathcal{E}_e}, \quad (1.3)$$

where $\beta_b = v/c$ is the normalized particle velocity and $\mathcal{E}_e = m_e c^2 = 0.511$ MeV is the electron rest energy. The subscript b is here used to disambiguate these quantities with respect to others which are introduced later, and the subscript e symbolizes fundamental, invariant quantities such as electron mass. While the velocity barely increases once the particle moves at close to c , further acceleration will increase the kinetic energy and the relativistic momentum,

$$\mathbf{p} = \gamma_b m_e \mathbf{v}, \quad (1.4)$$

of the particle, and instead of moving faster the particle will behave as though it has gained mass and become heavier. The gamma factor is virtually omnipresent when concerning electron accelerator theory.

One relativistic effect is that the particle's equations of motion must be modified according to its relativistic momentum. In Newtonian mechanics, the equation of motion is simply acquired from Newton's second law [24],

$$\mathbf{F} = m_e \mathbf{a} = \frac{d\mathbf{p}}{dt}, \quad (1.5)$$

where \mathbf{a} is the acceleration of the particle and \mathbf{F} is the force applied to it, typically the Lorentz force (Eq. 1.1) in the context of accelerators. When relativistic effects come into play, Eq. 1.5 becomes less simple, and when inserting Eq. 1.1 and 1.4 becomes

$$\frac{d\mathbf{p}}{dt} = -e(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \iff \frac{d(\gamma_b \mathbf{v})}{dt} + \frac{e}{m_e}(\mathbf{E} + \mathbf{v} \times \mathbf{B}) = 0, \quad (1.6)$$

where we remind that $\gamma_b = \gamma_b(v)$ and $\mathbf{v} = \mathbf{v}(t)$. In some cases, Eq. 1.6 can be solved analytically, normally by considering only a pure, homogeneous electric or magnetic field at a time, but in general, a numerical solution is required. In one of these simple cases, the particle moves in a plane perpendicular to a homogeneous magnetic field. Then, the motion describes a circle, see Fig. 1.1a, and the Lorentz force is balanced by the (relativistic) centripetal force [22]:

$$\frac{\gamma_b m_e v^2}{R} = evB \iff R = \frac{\gamma_b m_e v}{eB} = \frac{p}{qB}, \quad (1.7)$$

where R is the circle radius and $p = |\mathbf{p}|$. This circular motion is called cyclotron motion [22, 25], after one of the early types of accelerator [3], but has applications which are much more general. Of particular interest here is that such a homogeneous magnetic field is relatively easily achieved, most commonly in a so-called dipole magnet, which will be treated further in Chapter 2. The foremost use of dipole magnets are to bend the trajectory of the traversing particles, redirecting their propagation axis. It is also readily seen from Eq. 1.7 that particles of different energies (different gamma factors/momenta) will travel on paths of different radii, with the trajectories of higher-energy particles being “bent” less, see Fig. 1.1b. This difference in bending radius is used for example to determine the energy of the particles in spectrometers, which will be described further in Chapter 4, and to tune the temporal structure of particle bunches using magnetic bunch compressors, which are treated in Chapter 2.




Figure 1.1: Motion in homogeneous magnetic fields. **a**, Trajectory of a negatively charged particle in a homogeneous magnetic field pointing out of the page. **b**, Trajectories of particles of three different energies, as they pass through a (sector) dipole magnet. The green arrow shows the trajectory of the nominal energy, as it exits the magnet with the same angle as it enters.

1.1.2 The electromagnetic fields of a charged particle

Stationary charged particles are surrounded by electric fields, and when they move they also generate a magnetic field. For simplicity, we will first only consider the

electric field, which for a stationary electron has the form [23]

$$\mathbf{E} = -\frac{e}{4\pi\epsilon_0} \frac{\hat{\mathbf{r}}}{r^2}, \quad (1.8)$$

where ϵ_0 is the vacuum electric permittivity and $r = |\mathbf{r}|$ is the radial distance to the charge, and the electric field is pointing in the radial direction, $\hat{\mathbf{r}}$. The field distribution in Eq. 1.8 is visualized in Fig. 1.2a. Because of relativistic effects, when the charge moves at constant velocities approaching c , space and time around it deforms and the electric field distribution attains more of a disc-shape, see Fig. 1.2b. The equation for the electric field of a moving charge can be obtained through for example a Lorentz transformation, and becomes [23]

$$\mathbf{E} = -\frac{e}{4\pi\epsilon_0} \frac{1 - \beta_b^2}{(1 - \beta_b^2 \sin^2 \theta)^{3/2}} \frac{\hat{\mathbf{r}}}{r^2}, \quad (1.9)$$

where $\sin^2 \theta = \rho^2 / (z^2 + \rho^2)$, and θ is an angle between the particle trajectory and an observation point in the transverse plane. Equation 1.9 tells us that the magnitude of the electric field in Eq. 1.8 is modified in such a way as to have it decrease in the direction of propagation ($\hat{\mathbf{z}}$) and increase in the transverse direction ($\hat{\boldsymbol{\rho}} = \hat{\mathbf{x}} \cos \phi + \hat{\mathbf{y}} \sin \phi$). At the two angular extrema, $\theta = 0$ and $\theta = \pi/2$, Eq. 1.9 yields

$$\mathbf{E}_{\theta=0} = -\frac{e}{4\pi\epsilon_0} \frac{1}{\gamma_b^2} \frac{\hat{\mathbf{z}}}{z^2} \quad \text{and} \quad \mathbf{E}_{\theta=\pi/2} = -\frac{e}{4\pi\epsilon_0} \gamma_b \frac{\hat{\boldsymbol{\rho}}}{\rho^2}. \quad (1.10)$$

It is evident that where the magnitude of the field was earlier equal in all directions, there is a factor of γ_b^3 between the electric field in the transverse direction and in the direction of motion when the particle's velocity starts approaching c . Values of γ_b in the hundreds and thousands are not unusual in electron accelerators. The magnetic field which is generated by the moving charge has the form [23]

$$\mathbf{B} = \frac{1}{c^2} \mathbf{v} \times \mathbf{E} = -\frac{vE}{c^2} \hat{\boldsymbol{\phi}}, \quad (1.11)$$




Figure 1.2: The fields of a positively charged particle. **a**, Electric field of a stationary positive charge (chosen positive for visual purposes), pointing radially outwards in all dimensions in space. **b**, Electric field of the charge moving at close to c , where the fields are squeezed into more of a disc-shape transversely to the direction of motion. **c**, Magnetic field for the same case as **b**. The magnetic field is azimuthal, just as the field from a wire carrying a current, but has a large increase in magnitude close to the particle, just like the electric field. A negatively charged particle of the same charge magnitude just sees the directions of the arrows reversed, with nothing else changed.

so the generated magnetic field is azimuthal, just as for wire carrying a current, see Fig. 1.2c. An interesting and important relativistic effect is that two identical, co-propagating charged particles will at small γ_b repel each other, but as γ_b increases, the magnitude of this mutually repulsive force begins to vanish as $1/\gamma_b^2$ [8]. The decreased repulsion can be viewed as the azimuthal, restoring magnetic field canceling the repulsive electric field as $v \rightarrow c$, or that the repulsion is slowed down because of relativistic time dilation. This repulsive *space charge*-effect and its cancellation is paramount in electron sources, as will be discussed further in Chapter 2 and 3.

When the trajectories of the particles are bent, i.e. accelerated laterally, they lose a fraction of their energy in the form of electromagnetic radiation. This radiation can be viewed in a quantum electrodynamical picture as a photon “recoil”, through conservation of momentum, or in a classical electrodynamical picture as a time-dependent change in the electric and magnetic fields surrounding the particle. More specifically, the escaping energy is in the *radiative field*, while the disc-shaped field described above makes up the *velocity field*, which stays with the particle [23]. The power radiated during acceleration parallel to the direction of motion, such as during the acceleration we exploit to increase the particle momentum in the forward direction, is typically many orders of magnitude smaller than the power radiated from acceleration perpendicular to the direction of motion [21, 23], such as during propagation through a dipole magnet. Furthermore, the stronger the acceleration and the higher the particle energy, the more radiation will be emitted and the higher energy the radiated photons will have. The strong scaling of radiated power with particle energy and acceleration is readily seen in the relativistic Liénard-Larmor formula for acceleration perpendicular to the direction of motion [21, 23], where the radiated power is

$$P_{rad} = \frac{e^2}{6\pi\epsilon_0 c^3} \gamma_b^4 \left(\frac{d\mathbf{v}}{dt} \right)^2 = \frac{e^3}{6\pi m_e^2 \epsilon_0 c} \gamma_b^2 B^2, \quad (1.12)$$

where $d\mathbf{v}/dt = \mathbf{a}$ is the lateral acceleration. Equation 1.12 illuminates a large reason behind many future lepton collider designs being either linear or circular with large circumferences - at some combination of particle energy and ring circumference, the beam will lose so much energy from being forced onto a circular orbit that the energy one needs to add to keep the beam stable becomes unfeasible.

However, the emitted radiation is not all bad; as a matter of fact, one of the most common uses of highly relativistic electron beams is to produce X-rays for secondary use, for example in scientific applications at synchrotron light sources. That this radiation is so useful is an effect of special relativity [23]; when accelerated perpendicular to their direction of motion, the electrons can be viewed as little antennae that are moving very rapidly in the forward direction, see Fig. 1.3a. When the antenna is stationary, the radiated power forms a “donut” shape around it, but when it is highly relativistic, the donut is squeezed into a cone in the forward direction. This confines the emitted radiation (to $\sim 95\%$) within a cone of opening angle $\theta_0 = 1/\gamma_b$ around each electron’s direction of motion, as shown in Fig. 1.3b. It is the narrow opening angle of the radiation that makes it easier to handle and use for applications.

As the particle accelerates laterally, the radiation cone sweeps as a searchlight around its direction of motion, not unlike the headlights of a car going around a corner. The time it takes the light cone passing across the observer is related to the spectrum that the observer will measure. We will not go into detail on the radiation spectra,




Figure 1.3: Relativistic dipole radiation. **a**, Radiation pattern of a dipole “antenna” moving at three different velocities perpendicular to the acceleration, such as in e.g. a dipole magnet (see Fig. 1.1b). Note the difference in magnitude between the three cases! **b**, Angular distribution of the three cases. The angle $\pm\theta_0 = 1/\gamma_b$ corresponding to the first zero of the distribution for $\beta_b = 0.9$ has been marked with a red arrow. The angular distribution in the y -direction (out of the page) is somewhat similar, but has a slightly wider distribution as it is in the plane of the “donut”.

but there are some features worth noting. When the particle takes curved trajectories, this type of radiation is called *synchrotron radiation*. Synchrotron radiation has a characteristic broadband spectrum, which has a peak at approximately $0.3\mathcal{E}_c$, where \mathcal{E}_c is called the critical (photon) energy and is given by [21]

$$\mathcal{E}_c = \hbar\omega_c = \frac{3}{2}\gamma_b^3\frac{a}{c} = \frac{3}{2}\frac{e}{m_e}\gamma_b^2B, \quad (1.13)$$

where $\hbar = h/(2\pi)$ is the reduced Planck’s constant, ω_c is the critical angular frequency and $a = |\mathbf{a}|$. Just like the radiated power (Eq. 1.12), the critical energy increases quickly with particle energy, but has a weaker dependence on the lateral acceleration.

1.1.3 Single-particle propagation

If the particles in an accelerator would be left to propagate freely, effects such as space charge would eventually cause them to diverge and be lost by hitting the vacuum chamber. Therefore, we need to apply an external restoring force which keeps the particles close to the desired propagation axis. If the applied restoring force is linear with distance from the axis, $F \propto x$, where x is the horizontal spatial coordinate, we can get a harmonic oscillation around the axis, effectively keeping our particles confined. Since magnetic fields are useful for controlling the transverse motion of particles, we assume that this force is magnetic (we will get into that a bit more in Chapter 2 and 3), and use Eq. 1.5 and 1.6 to write

$$F = \frac{dp}{dt} = -evg_q x, \quad (1.14)$$

where g_q is the magnetic gradient. We want to solve this equation in space rather than time, so that we can follow it along its spatial longitudinal coordinate, s , and since $dt \approx ds/v$, we can write (assuming constant γ_b)

$$F = \frac{dp}{dt} = v\frac{dp}{ds} = v\gamma_b m_e \frac{dv}{ds} = v\gamma_b m_e \frac{d}{ds} \frac{dx}{dt} = v^2\gamma_b m_e \frac{d^2x}{ds^2} = -evg_q x, \quad (1.15)$$

which we can rewrite into the transverse equation of motion as

$$\frac{d^2 x}{ds^2} - \frac{eg_q}{p} x = x'' - \kappa(s) \cdot x = 0, \quad (1.16)$$

where we have introduced the *normalized focusing strength*

$$\kappa(s) = \frac{eg_q}{p}. \quad (1.17)$$

The differential equation in Eq. 1.16 is known as Hill's equation [22, 25], and differs from the simple harmonic oscillator in that the proportionality constant in the equation of motion can in principle vary with position. Here, we will only examine the cases of constant focusing and no focusing, which can easily be solved analytically. For an arbitrary $\kappa(s)$, the analytical solution is either more complex or non-existent and the equation must be solved numerically.

The case of no focusing is simple: $\kappa = 0$ gives the solution $x(s) = x_f = x_0 + x'_0 s$, where x_0 is the initial position and $x'_0 = dx_0/ds_0 = \tanh(p_x/p_z) \approx p_x/p_z$ is the initial angle with respect to the propagation axis. As there is no force on the particle, the angle is constant, $x'(s) = x'_f = x'_0$. This set of propagation equations can be expressed in matrix form,

$$\begin{pmatrix} x_f \\ x'_f \end{pmatrix} = \mathbf{x}_f = \begin{bmatrix} 1 & L \\ 0 & 1 \end{bmatrix} \begin{pmatrix} x_0 \\ x'_0 \end{pmatrix} = \mathbf{R}_d \begin{pmatrix} x_0 \\ x'_0 \end{pmatrix} = \mathbf{R}_d \mathbf{x}_0, \quad (1.18)$$

where \mathbf{R}_d is the transfer matrix of the so-called drift-section with length $L = s_f - s_0$, and \mathbf{x}_f and \mathbf{x}_0 are the final and initial coordinate vectors, respectively. When we instead have constant focusing along s , we obtain [22]

$$\mathbf{R}_F = \begin{bmatrix} \cos(\sqrt{\kappa}L) & \frac{1}{\sqrt{\kappa}} \sin(\sqrt{\kappa}L) \\ -\sqrt{\kappa} \sin(\sqrt{\kappa}L) & \cos(\sqrt{\kappa}L) \end{bmatrix} = \begin{bmatrix} \cos(\Delta\Psi) & \frac{1}{\sqrt{\kappa}} \sin(\Delta\Psi) \\ -\sqrt{\kappa} \sin(\Delta\Psi) & \cos(\Delta\Psi) \end{bmatrix}, \quad (1.19)$$

where we have introduced the *betatron phase advance*, $\Delta\Psi = \sqrt{\kappa}L$, which is the phase advance of this harmonic motion between s_0 and s_f . Since we are describing an oscillation in the spatial domain, we can also identify the wavenumber, $k_\beta = \sqrt{\kappa}$, which is the spatial equivalent of the angular frequency, $\omega_\beta = k_\beta v$, in the temporal domain. We can note two features of Eq. 1.19; firstly, if we set $\kappa = 0$, we retrieve \mathbf{R}_d from Eq. 1.18, and secondly, if we expand Eq. 1.19 for small $\Delta\Psi$, i.e. we have a short and weak focusing section that only imparts a “kick” to the propagation angle of the particle, we obtain the thin-lens approximation of our focusing section,

$$\mathbf{R}_{F,\text{thin}} = \begin{bmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{bmatrix}, \quad (1.20)$$

where $f = 1/\kappa L$. The transfer matrix notation is convenient for calculating the final coordinates after a series of elements, since we can multiply the matrices together as

$$\mathbf{x}_f = \mathbf{R}_n \dots \mathbf{R}_2 \mathbf{R}_1 \mathbf{x}_0 = \mathbf{R}_{\text{tot}} \mathbf{x}_0, \quad (1.21)$$

where \mathbf{R}_1 is the transfer matrix of the first element passed through, \mathbf{R}_2 the second, et cetera.

There are other coordinates that we might also want to propagate from one location to the other. x and x' can be used to describe the motion in one transverse plane, and we can do the same in the perpendicular transverse plane with y and y' and find the corresponding equation of motion. To fully describe the propagation of the particle, we also need the local longitudinal motion, along z . Instead of also using an angle z' , which might be a bit less intuitive to envision, the particle kinetic energy, \mathcal{E} , momentum, p , or relative energy deviation, $\delta = (\mathcal{E} - \langle \mathcal{E} \rangle) / \langle \mathcal{E} \rangle$, where $\langle \mathcal{E} \rangle$ is the average beam energy, is often used. That gives us a total coordinate vector $\mathbf{x} = (x, x', y, y', z, \delta)$, and the corresponding transfer matrix is now 6×6 instead of 2×2 , but Eq. 1.21 is still valid.

Sometimes one wants to examine the propagation of one particular coordinate when varying one of the initial coordinates. In that case, it might suffice to extract individual elements of \mathbf{R} ; for example, we might want to examine how the horizontal position depends on an deviation in energy - for that we can simply extract the corresponding matrix element, R_{16} in this case, and calculate $x_f = R_{16}\delta$. The subscripts 1 and 6 refer to the position of the coordinate in the vector \mathbf{x} , so R_{16} maps the 6th coordinate (δ) to the 1st coordinate (x), with the matrix element being located at the 1st row and 6th column in the matrix.

Worth mentioning here is also that the above treatment has only been for terms which are linear with any element in \mathbf{x} . There are, of course, more complicated cases than this, with non-linear and cross-terms that scale with e.g. δ^2 , xy and $yy'z$. The transfer matrices then need to be higher-rank tensors, and the situation inevitably becomes more complicated. However, we can still study the propagation of one of the coordinates using different input coordinates. The general form of this coordinate mapping can be written [26, 27]

$$\mathbf{x}_{i,f} = \mathbf{x}_i + \sum_j R_{ij} \mathbf{x}_j + \sum_{jk} T_{ijk} \mathbf{x}_j \mathbf{x}_k + \sum_{jkl} U_{ijkl} \mathbf{x}_j \mathbf{x}_k \mathbf{x}_l + \dots, \quad (1.22)$$

where we have omitted the subscript 0 on the initial coordinate vectors for clarity, and where R_{ij} , T_{ijk} and U_{ijkl} are the first- second- and third-order transfer matrix elements, respectively. Equation 1.22 might look daunting, but as an example, we can limit ourselves to look at the final horizontal position as function of only initial relative energy deviation, which then simplifies Eq. 1.22 to

$$x_f = x_0 + R_{16}\delta + T_{166}\delta^2 + U_{1666}\delta^3 + \mathcal{O}(\delta^4), \quad (1.23)$$

with $i = 1$ and $j = k = l = 6$.

1.1.4 Particle beam propagation

In the typical case, we are dealing with an ensemble of particles, maybe containing millions or billions of particles, and not individual ones. These ensembles can often be represented by statistical distributions that capture the many different combinations of for example x and x' that exist in such an ensemble, see Fig. 1.4. A convenient and common way of doing this is by using the *Courant-Snyder* (CS) parameters [22, 25]

$$\beta_u = \frac{\langle u^2 \rangle}{\varepsilon_u}, \quad \alpha_u = -\frac{\langle uu' \rangle}{\varepsilon_u} = -\frac{1}{2}\beta'_u, \quad \gamma_u = \frac{\langle u'^2 \rangle}{\varepsilon_u}, \quad \beta_u \gamma_u = 1 + \alpha_u^2, \quad (1.24)$$

where $u = x, y$ denotes either transverse coordinate and the primes denote their propagation angle with respect to the propagation axis as before. ε_u is the *geometric* transverse root-mean-square (rms) trace-space emittance, and $\langle \cdot \rangle$ denotes the central moment [28] of the enclosed quantity. Here, the central moments are equal to the rms values, $\langle u^2 \rangle = \sigma_u$ and $\langle u'^2 \rangle = \sigma_{u'}$. The β function relates to the spatial beam envelope, γ relates to the angular beam envelope, and α describes the correlation between the two. The particle ensembles can be visualized in their $u - u'$ *phase space* distributions - an example of such a phase space distribution is shown in Fig. 1.4. The area of the distribution in phase space is related to the geometric emittance; the smaller the area, the smaller the emittance. A very important property of the geometric emittance is that it, by virtue of Liouville's theorem, is conserved under influence of conservative, linear forces [22, 25, 29], but not under acceleration (since then p_z increases). To this end one can define a corresponding *normalized* emittance, $\varepsilon_{n,u}$, which is conserved also under acceleration. These emittances can be calculated statistically as [30, 31]

$$\begin{aligned}\varepsilon_u^2 &= \langle u^2 \rangle \langle u'^2 \rangle - \langle uu' \rangle^2 \\ \varepsilon_{n,u}^2 &= \langle u^2 \rangle \langle \beta_b^2 \gamma_b^2 u'^2 \rangle - \langle \beta_b \gamma_b uu' \rangle^2.\end{aligned}\tag{1.25}$$

For beams of small divergence and energy spread, the emittances are related by $\varepsilon_{n,u} = \varepsilon_u \beta_b \gamma_b$ [30, 31]. The emittance quantifies how well correlated the particles' positions are with their propagation angles. A high degree of correlation is important if one wants to, for example, focus a beam to a small size with a linear lens after which the outgoing angle is linearly dependent on the transverse position, see e.g. Eq. 1.20. A smaller value of the emittance, meaning a high degree of correlation, is usually deemed better for many applications. The CS parameters and emittances are normally evaluated in the two perpendicular transverse directions, x and y , for




Figure 1.4: Transverse phase space distribution. A representation of the $u - u'$ phase space distribution and the phase space ellipse enclosing the distribution. The geometric emittance of the beam is related to the area of the enclosing ellipse and the CS parameters are related to the ellipse's orientation and aspect ratio.

reasons that will become apparent in Chapter 2. It should be noted that there are two ways of characterizing the parameters of a bunch, as *projected* parameters and as *slice* parameters. The projected parameters concern the entire bunch, where all the particles have been projected onto the examined plane (e.g. $x - x'$). The slice parameters are usually obtained by “slicing” the beam into thin sections, for example in time, to obtain the variation of different parameters along the bunch. The expressions in Eq. 1.24 and 1.25 hold for both applications, the difference lies in which set of particles is examined.

The CS parameters can be manipulated using focusing optics in a corresponding way to the single-particle case in Section 1.1.3, and a corresponding set of transport matrices can be found for this purpose [22, 25]. This way of propagating a representation of the entire ensemble at once can be useful when designing electron transport systems, for example. The propagation of the CS parameters is governed by [25]

$$\frac{1}{2}\beta'' + \kappa(s)\beta - \frac{1}{\beta} \left[1 + \left(\frac{1}{2}\beta' \right)^2 \right] = 0, \quad (1.26)$$

where $\beta = \beta(s)$ and derivatives are again with respect to s , we have omitted the subscript u since these equations are valid for either axis. Equation 1.26 does not generally have analytical solutions for arbitrary $\kappa(s)$, but we can find solutions numerically if we know $\kappa(s)$.

As in Section 1.1.3 we will examine the two cases of free drift and constant focusing to make some interesting observations, but alter the former slightly to look at the behavior of the β function around a focal plane, or *beam waist*. At a beam waist, the slope of the β function is 0, see Fig. 1.5a, and so $\alpha = -1/2\beta' = 0$. Assuming no




Figure 1.5: β function propagation. **a**, Envelope β functions around a focus (or beam waist), as given by Eq. 1.27, for three different waist radii. **b**, β functions going into a constant focusing region with $\kappa = 1 \text{ m}^{-2}$ starting at $s = 0$, as given by Eq. 1.29 for $s > 0$. The matching condition is simple in this case, and the matched beam (red), exhibits no oscillations.

external focusing force ($\kappa = 0$), Eq. 1.26 then has the deceptively simple solution

$$\beta(s) = \beta^* + \frac{s^2}{\beta^*}, \quad (1.27)$$

where $\beta^* = \beta(0)$ is the beam waist β function. Examining this equation, we can see that the smaller our beam waist, i.e. the smaller our β^* and the more we “squeeze” the beam together, the quicker the beam will expand outside of the focus, see Fig. 1.5a. This has the implication that if we want to focus to a smaller focal spot, we need have a larger initial beam size before focusing. The physical rms beam radius around the waist location we get from Eq. 1.24 as

$$\sigma = \sqrt{\varepsilon\beta} = \sqrt{\varepsilon\beta^*} \sqrt{1 + \left(\frac{s}{\beta^*}\right)^2} = \sigma^* \sqrt{1 + \left(\frac{s}{\beta^*}\right)^2}. \quad (1.28)$$

While individual particles propagate in straight lines in a drift, their ensemble evidently behaves differently. The other case we will examine is again when we have a constant $\kappa(s)$. In this case, assuming for simplicity that we enter the focusing section with $\alpha = 0$, the general solution to Eq. 1.26 is [32]

$$\beta(s) = \beta^* \cos^2(\sqrt{\kappa}s) + \frac{1}{\beta^*\kappa} \sin^2(\sqrt{\kappa}s), \quad (1.29)$$

where we can identify a particular case, $\beta(s) = 1/\sqrt{\kappa} = 1/k_\beta = \text{constant} = \beta_m$, i.e. a solution where the beam envelope remains constant with $\beta(s) = \beta_m$. At these conditions, the beam is said to be *matched* to the focusing strength [33] - the beam’s intrinsic tendency to diverge, in this case because of the finite emittance, is perfectly balanced by the external focusing. An example of what matched and non-matched beams can look like is shown in Fig. 1.5b. The envelope oscillations shown in Fig. 1.5b, coming from a mismatch in focusing, are often referred to as *beam scalloping*. We saw in Section 1.1.3 that constant focusing yielded harmonically oscillating trajectories, see Eq. 1.19, which together with this result means that under matched conditions, individual particles perform harmonic oscillations while their ensemble remains at a fixed size. The betatron phase advance through the system between the points s_0 and s can be calculated using the β function as [22, 25]

$$\Delta\Psi = \int_{s_0}^s \frac{1}{\beta(s)} ds. \quad (1.30)$$

Particle bunches and their fields

As alluded to above, the particle beams which are relevant in this thesis are not continuous, but confined transversely and longitudinally into *bunches*. As the bunches propagate near the speed of light, their longitudinal distribution is related to their temporal distribution as $\sigma_z \approx c\sigma_t$. In Chapter 2 and 3, it will become clear how and why these bunches become bunched in the first place, but here we will show how the fields of a bunch of electrons behave, as compared to a single electron.

Although not usually the case in linear accelerators, we will make the assumption that the longitudinal charge distribution is Gaussian. Gaussian distributions are

easy to treat analytically, and will allow us to simplify our analysis. The fields from the bunch have a straightforward dependence on the longitudinal distribution - in principle, the longitudinal portion can be swapped out for any appropriate analytical function. Assuming a radially symmetric charge distribution with rms radius σ_ρ and length σ_z , we then have the bunch charge density

$$\rho_b(\rho, z) = -e n_{b,0} \exp\left[-\frac{z^2}{2\sigma_z^2}\right] \exp\left[-\frac{\rho^2}{2\sigma_\rho^2}\right], \quad n_{b,0} = \frac{1}{(2\pi)^{3/2}} \frac{N_b}{\sigma_\rho^2 \sigma_z}, \quad (1.31)$$

where N_b is the total number of electrons and $n_{b,0}$ is the peak electron number density. With this expression, we can calculate the fields from such a bunch of electrons. While Eq. 1.10 gives the single-electron electric field, the total field from a distribution of electrons is somewhat different, and is given by [22]

$$E_\rho(\rho, z) = -\frac{1}{(2\pi)^{3/2} \epsilon_0} \frac{Q_b}{\sigma_z} \exp\left[-\frac{z^2}{2\sigma_z^2}\right] \frac{1}{\rho} \left(1 - \exp\left[-\frac{\rho^2}{2\sigma_\rho^2}\right]\right). \quad (1.32)$$

where $Q_b = N_b e$ is the bunch charge. The longitudinal field distribution follows the longitudinal charge profile, or current profile (in time), while the radial field first grows almost linearly, peaks at $\rho_0 \approx 1.585 \sigma_\rho$,¹ and then decays as roughly $1/\rho$, see Fig. 1.6. Note that for a single electron, the transverse field decays as $1/\rho^2$ and depends explicitly on γ_b . The $1/\rho$ -scaling in Eq. 1.32 of a relativistic bunch is identical to that of a long, cylindrical charge distribution, which is indeed the shape of the bunch in the bunch rest frame. Inserting ρ_0 into Eq. 1.32 at $z = 0$ yields the peak radial

¹Obtained with WolframAlpha [34]

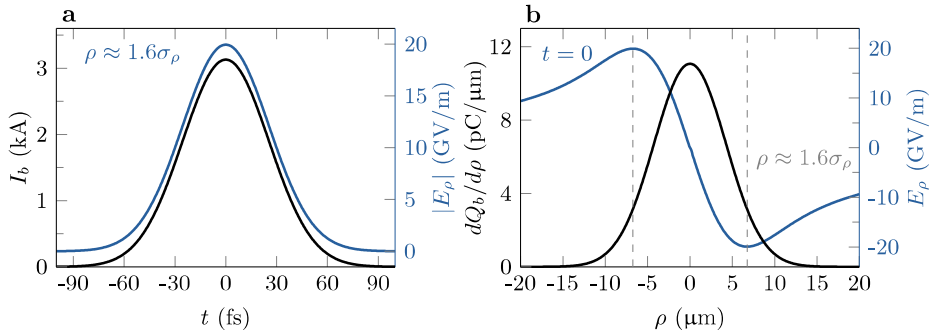


Figure 1.6: Temporal and transverse distributions of the transverse electric fields of a Gaussian electron bunch. **a**, Maximum radial electric field (blue) at $\rho \approx 1.6\sigma_\rho$ (see **b**), as a function of time, for a 200 pC, 60 fs fwhm bunch (note that 0 is offset). The field amplitude follows the bunch current profile (black) exactly. **b**, Radial electric field distribution (blue) of the bunch in **a**, at $t = 0$, when focused to a charge distribution spot size of 10 μm fwhm (black). The field amplitude is maximal at a distance of $\rho \approx 1.6\sigma_\rho$ from the axis, and 0 on the axis. Note that ρ cannot actually have negative values: it has been plotted this way to show the symmetric, Gaussian charge distribution, and the sign of the electric field represents it pointing towards the axis.

electric field

$$E_{\rho,0}(\sigma_\rho, \sigma_z) \approx -\frac{1}{2.216 \times (2\pi)^{3/2} \epsilon_0} \frac{Q_b}{\sigma_\rho \sigma_z}. \quad (1.33)$$

Evidently, the peak radial field scales equally with the longitudinal and transverse sizes of the bunch. The transverse field behaves slightly differently near the bunch symmetry axis, where the exponential in Eq. 1.32 can be expanded around $\rho = 0$, such that the radial electric field near the radial bunch center increases linearly as

$$E_\rho(\rho, z) = -\frac{1}{2(2\pi)^{3/2} \epsilon_0} \frac{Q_b}{\sigma_\rho^2 \sigma_z} \exp\left[-\frac{z^2}{2\sigma_z^2}\right] \rho. \quad (1.34)$$

The magnetic field can, as before, be calculated easily using Eq. 1.11.

Brightness and emittance preservation

The quality of a bunch is not always judged only by its transverse emittance, there is also the longitudinal phase space to consider. When including the bunch charge, Q_b , duration, σ_t , and relative energy spread, σ_δ , one often talks about the 6D brightness [35]

$$B_{6D} = \frac{Q_b}{\epsilon_{n,x} \epsilon_{n,y} \epsilon_{n,z}} \approx \frac{2I_b}{\epsilon_{n,x} \epsilon_{n,y} \sigma_\delta} \approx \frac{B_{5D}}{\sigma_\delta}, \quad (1.35)$$

where $\epsilon_{n,z}$ is the longitudinal normalized emittance and B_{5D} is the 5D brightness, another common figure of merit. For most applications, including particle colliders and FELs, the brightness should be as high as possible.

According to Liouville's theorem, the emittance can only be constant or increase, and it is therefore of great interest to preserve it to the greatest possible extent. Luckily, the linear focusing elements that we have looked at so far preserve the emittance - indeed any element that applies a transversely constant and/or linear, conservative force will preserve the it. Non-linear forces, such as space-charge or transverse restoring forces which vary non-linearly with the transverse position, can increase the emittance in ways which are difficult or impossible to reverse. However, if the bunch has a finite energy spread, a linear electromagnetic force will not act on all particles the same way, such that also a linear force contributes to the emittance degradation. This variation in focusing strength is visible in Eq. 1.16 and 1.17, where it is explicitly evident that the focusing strength depends on the relativistic particle momentum. In the case of a non-zero energy spread, we can write Eq. 1.17 as

$$\kappa(\delta) = \frac{eg_q}{p(1+\delta)} = \kappa (1 - \delta + \delta^2 - \delta^3 + \dots). \quad (1.36)$$

This expression can be inserted in the transfer matrices of focusing elements to yield their chromatic response to a given order. The effect of the varying focusing strength is *chromatic aberrations*, and from Eq. 1.36 and 1.20, we can see that higher energies ($\delta > 0$) lengthen the focal length, and vice versa. The propagation of the CS parameters is affected by these aberrations, and emittance increase can occur. However, by properly combining focusing elements, a system can be made *apochromatic* [17], such that its impact on a bunch with finite energy spread is mitigated. Minimization of chromatic effects in certain electron transport systems were part of the work presented in **Paper I, II and IV**.

1.2 Electromagnetic waves and their propagation

Electromagnetic waves are of paramount importance to the types of accelerators discussed in this thesis. They are both directly and indirectly involved in the acceleration of electrons in several accelerator types. Within the scope of this thesis, the most common forms of these electromagnetic waves are ultra-short and intense laser pulses, high-power RF pulses and X-rays. Laser pulses are for example often used for generating and manipulating the electron bunches, and can also be the desired end “product” of an accelerator in the case of a free-electron laser. The (incoherent) X-rays are another example of a desirable product of the bunches. The RF fields are most commonly used for acceleration and characterization. The special case of electromagnetic waves in plasmas is treated in Section 1.3. Before examining the actual wave description, we will use the simpler ray optics model to introduce some fundamental concepts of electromagnetic wave propagation.

1.2.1 Ray optics

Ray optics, or geometrical optics, is the simplest theory of how light propagates and interacts with objects. Most of us come in contact with this model in school, and have probably come across its uses for e.g. imaging using simple lenses or mirrors. Just as in single-particle propagation, each light ray is associated with a position and an angle, x and x' , for each position s along the propagation axis, which can be controlled using for example focusing devices. However, in photon optics the control of the photons is not governed directly by forces on the photons (as is the case for massive particles, see e.g. Eq. 1.6), but typically by the refractive index, η , of optical media. The refractive index is defined by how much it slows the phase velocity down compared to vacuum, i.e. $c \rightarrow c/\eta$. Therefore, during passage through such a medium, the ray will appear to cover a greater path length, $d = \eta L$, called the optical path length, where L is the physical length of the medium. η can vary in space as $\eta = \eta(\mathbf{r})$, such that the coordinates of the ray must follow the ray equation [36]

$$\frac{d}{ds} \left(\eta(\mathbf{r}) \frac{d\mathbf{r}}{ds} \right) = \nabla \eta(\mathbf{r}). \quad (1.37)$$

The ray equation simplifies greatly for constant η , such that $d^2\mathbf{r}/ds^2 = 0$, which describes a straight line, $x(s) = x_0 + x'_0 s$. $x'_0 = dx_0/ds_0$ is the propagation angle with respect to the axis, which is assumed to be small such that $ds \approx dz$, which is called the paraxial approximation. Just as in Section 1.1.3, we can describe the ray propagation with a matrix equation,

$$\begin{pmatrix} x_f \\ x'_f \end{pmatrix} = \mathbf{x}_f = \begin{bmatrix} 1 & s \\ 0 & 1 \end{bmatrix} \begin{pmatrix} x_0 \\ x'_0 \end{pmatrix} = \mathbf{M}_d \mathbf{x}_0, \quad (1.38)$$

c.f. Eq. 1.18. In the context of Eq. 1.37, constant focusing for an optical ray is not quite as intuitive to envision as for a particle, see Eq. 1.16. Focusing can be achieved either by a geometrically tailored medium of constant η , for example a lens, or a geometrically continuous medium with transversely varying η , such as *graded-index* (GRIN) materials. Constant focusing can be achieved in GRIN materials, but not in lenses which are geometrically confined. By borrowing from Eq. 1.16, we can see that

constant focusing suggests the condition

$$\frac{d^2x}{ds^2} = -C^2x, \quad (1.39)$$

where C is some constant. If we also assume that η only varies transversely (along x , for simplicity) and is independent of s , Eq. 1.37 becomes

$$\frac{d^2x}{ds^2} = \frac{1}{\eta(x)} \frac{d\eta(x)}{dx} \frac{dx}{ds} \iff \frac{d\eta(x)}{dx} + C^2\eta(x)x = 0, \quad (1.40)$$

which has the solution

$$\eta(x) = \eta_0 \exp\left[-\frac{1}{2}C^2x^2\right], \quad (1.41)$$

which is a Gaussian distribution with a peak refractive index η_0 at $x = 0$ and standard deviation $\sigma_\eta = C^{-1}$. If we assume that we are close to the axis, e.g. within $|x| < C^{-1}$, the Gaussian curve is well approximated by the parabola

$$\eta(x) \approx \eta_0 \left(1 - \frac{1}{2}C^2x^2\right). \quad (1.42)$$

Since we have imposed the focusing condition in Eq. 1.39, we obtain the same type of solution as for a particle in Section 1.1.3, namely [36]

$$\mathbf{M}_F = \begin{bmatrix} \cos(Cs) & \frac{1}{C} \sin(Cs) \\ -C \sin(Cs) & \cos(Cs) \end{bmatrix}, \quad (1.43)$$

c.f. Eq. 1.19. These matrices can be concatenated in the same practical way as is done in particle beam optics, see Eq. 1.21.

As optical components, GRIN materials are not very common within the context of particle accelerators, but as we will see in Section 1.2.4 and 1.3.3, a similar refractive index gradient can arise from the interaction between intense laser pulses and media, causing the pulses to self-focus. A more common optical component is the (geometric) lens, where a constant refractive index material of varying thickness is used to provide focusing. The two focusing techniques have in common that the optical path length along the optical axis is longer than away from it.

Another set of common elements are reflective mirrors. Flat mirrors merely redirect the beam, while curved mirrors are used for (de-)focusing. Just as for lenses, a spherical mirror can be used without introducing aberrations provided that the beam is small compared to the mirror's radius of curvature, and that it propagates along the mirror's symmetry axis. However, the optimal shape for a mirror is in fact a *paraboloid*, which will redirect all incoming rays, parallel to the optical axis, towards the focal point [36]. If one uses a spherical mirror or lens for a beam which is too large or divergent, such that the paraxial approximation does not hold, or if the beam is sent onto a focusing parabolic mirror at the wrong angle, so-called *geometric aberrations* will arise. These aberrations distort imaging done with these optics, and show up as for example asymmetric focal spots.

At this point, one might think that defining a ray beam distribution and emittance might be the next step, since this is quite natural for particles and was done in Section 1.1.4. While defining such distributions is effective for some things, such as ray-tracing

calculations, this relatively simple ray optics model misses out on many effects inherent to the propagation of light, such as diffraction and interference. To capture more of the actual behaviour, one must describe the light as an electromagnetic wave instead. These waves are governed by Maxwell's equations, which we will introduce before proceeding to the wave description.

1.2.2 Maxwell's equations and potentials

As seen in Section 1.1.2, moving charges give rise to electromagnetic fields, some of which are in the form of radiation. Electromagnetic radiation can also arise from photoemission (such as the stimulated emission in laser amplification [36]), where changes in quantum states of atoms and molecules are accompanied by the emissions of photons of corresponding energy to the energy difference between these states. Once emitted, the radiation is indifferent of its source and follows the same set of rules that governs all non-quantum electromagnetic fields. This set of rules is called the Maxwell's equations [21]:

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad (1.44a)$$

$$\nabla \times \mathbf{B} = \mu_0 \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right), \quad (1.44b)$$

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0}, \quad (1.44c)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (1.44d)$$

where μ_0 is the vacuum magnetic permeability, \mathbf{J} is the electric current density (i.e. the moving charges, our radiation source term), and ρ is the static charge density. In a vacuum with no charges or currents, i.e. $\rho = 0$ and $\mathbf{J} = \mathbf{0}$, Eq. 1.44a and 1.44b can be combined to yield the free-space wave equation

$$\nabla^2 \mathbf{E} - \frac{1}{c^2} \frac{\partial^2 \mathbf{E}}{\partial t^2} = 0, \quad (1.45)$$

where we see that the phase velocity of these waves is $c = 1/\sqrt{\mu_0\epsilon_0}$. One solution of Eq. 1.45 is plane waves,

$$\mathbf{E}(\mathbf{r}, t) = \mathbf{E}_0(\mathbf{r}, t) \sin(\mathbf{k} \cdot \mathbf{r} - \omega t), \quad (1.46)$$

where \mathbf{k} is the wave vector with $|\mathbf{k}| = k = 2\pi/\lambda = \omega/c$, where λ is the wavelength and $\omega = 2\pi f$ is the angular frequency of the wave, and f is the frequency. Note that \mathbf{E} has a direction in space which might change in time, called its *polarization*. \mathbf{B} can be readily found when \mathbf{E} is known (see Eq. 1.11), so usually it is enough to find the solution for \mathbf{E} . The electric and magnetic fields are usually perpendicular to each other and \mathbf{k} , and in phase.

When formulating theory, introducing the scalar and vector potentials Φ and \mathbf{A} can simplify the description of many problems. These potentials are defined through [21, 23],

$$\begin{aligned} \mathbf{B} &= \nabla \times \mathbf{A}, \\ \mathbf{E} &= -\left(\nabla \Phi + \frac{\partial \mathbf{A}}{\partial t} \right), \end{aligned} \quad (1.47)$$

where Φ is the scalar potential and \mathbf{A} is the vector potential. There is some freedom in choosing these potentials' relation to each other, but for this application the Lorentz gauge $\nabla \cdot \mathbf{A} = -c^{-2} \partial \Phi / \partial t$ is often used [23], as the potentials can then also form wave equations. Φ is the electric potential where voltage, U , is derived from. In vacuum, $\Phi = 0$, and Eq. 1.46 can be expressed as

$$\mathbf{A}(\mathbf{r}, t) = -\mathbf{A}_0(\mathbf{r}, t) \cos(\mathbf{k} \cdot \mathbf{r} - \omega t), \quad (1.48)$$

where $\mathbf{A}_0 = \mathbf{E}_0 / \omega$.

In wakefield acceleration theory, which will be examined further in Section 1.3.3 and 1.3.4, it is most common to use the normalized vector and scalar potentials

$$\begin{aligned} \varphi &= \frac{e}{m_e c^2} \Phi, \\ \mathbf{a} &= \frac{e}{m_e c} \mathbf{A}, \end{aligned} \quad (1.49)$$

where φ is associated with the charge-density variation of particle bunches and electrostatic-like plasma waves (which will be treated in Section 1.3.3 and 1.3.4), and \mathbf{a} the laser pulses. Regarding the intricacies relating to potential composition, the interested reader is referred to for example Ref. [37] for more information. A measure of the electric field strength which is commonly used in the laser-wakefield acceleration community is the parameter

$$a_0 = |\mathbf{a}|_{max} = \frac{e}{m_e c \omega} E_0, \quad (1.50)$$

which is the *peak normalized vector potential*. a_0 quantifies the strength of the electric field relative to the electron mass - if $a_0 > 1$, then the transverse momentum gained by an electron in the electric field of the wave over one cycle is larger than $m_e c$, i.e. $\gamma_b \gg 1$ and relativistic effects must be considered.

1.2.3 Electromagnetic beam propagation

The electromagnetic waves are described by a real, vectorial wavefunction $\mathbf{E}(\mathbf{r}, t)$, but it is usually convenient to treat them as a complex-valued, scalar function, such that [36]

$$E(\mathbf{r}, t) = E_0(\mathbf{r}, t) \Re \{ \exp [i\phi(\mathbf{r})] \exp [-i\omega t] \}, \quad (1.51)$$

where E_0 is the amplitude, or envelope, of the wave and ϕ is the spatial phase and $\Re \{ \cdot \}$ denotes taking the real part of the enclosed quantity. $E(\mathbf{r}, t)$ can take a few different forms in space, two of which are the spherical wave and the plane wave (as in Eq. 1.46), but of most interest to our application for plasma-based accelerators is the Gaussian beam [36]. Temporarily omitting the time dependence, which will be treated below in Section 1.2.4, a Gaussian beam in space can be described as

$$E(\rho, s) = E_0 \frac{W_0}{W(s)} \exp \left[-\frac{\rho^2}{W^2(s)} \right] \Re \left\{ \exp \left[iks + ik \frac{\rho^2}{2\mathcal{R}(s)} - i\zeta(s) \right] \right\}, \quad (1.52)$$

where $W(s)$ is the beam radius at $1/e$ of the peak amplitude (“ $1/e$ radius”), $\mathcal{R}(s)$ is the wavefront radius of curvature, $\zeta(s)$ is an additional phase factor called the Gouy

phase, and we let \mathbf{k} point along the s -axis. $W_0 = W(0) = \sqrt{s_R \lambda / \pi}$ is the beam waist (or focal spot) radius, and s_R is the diffraction length (normally called Rayleigh length), and is related to the increase of the beam area at either side of a focal plane. If we omit the Gouy phase, we see that the spatial phase, $\phi(\mathbf{r})$, is directly proportional to k . The lines that are drawn by $\phi(\mathbf{r})$ are called wavefronts, and represent regions in space where the complex wavefunction has the same phase. They are spaced by 2π in phase, which is equivalent to one wavelength, λ . When entering a medium with refractive index $\eta > 1$, the frequency is unchanged, but the phase velocity $c \rightarrow c/\eta$ decreases, and so the wavelength becomes shorter. The normal to the wavefronts in space is given by $\mathbf{k}(\mathbf{r})$, and we can see that the rays from Section 1.2.1 actually represent \mathbf{k} [36, 38].

Compared to plane waves, which exist (mathematically) in all of space and therefore have infinite energy, Gaussian beams are better anchored in reality. For light, the wave oscillates incredibly fast ($f \approx 6 \times 10^{14}$ Hz for green light at 500 nm), and one does typically not measure the waveform directly. Rather, if we put up a detector and intercept the beam, it is the *intensity*, $I(\mathbf{r}) = c\epsilon_0 |\mathbf{E}(\mathbf{r})|^2 / 2$, that we detect. The intensity is the energy per unit time and unit area flowing through a cross section of the beam. Averaging over the rapidly oscillatory terms in Eq. 1.52 causes them to disappear, and the intensity obtains the simpler form

$$I(\rho, s) = I_0 \left(\frac{W_0}{W(s)} \right)^2 \exp \left(-\frac{2\rho^2}{W^2(s)} \right), \quad (1.53)$$

where I_0 is the peak intensity. Comparing to the rms radius (σ_ρ) commonly used for particle distributions, we see that W is the $1/e^2 = 2\sigma_\rho$ -radius of the intensity ($1/e$ -radius for the field). For a Gaussian distribution, the rms relates to the full-width-at-half-maximum (fwhm) as $\text{fwhm} = 2\sqrt{2 \ln 2} \times \text{rms}$. Knowing the intensity allows us to write a_0 in a simplified way, as $a_0 \approx 0.855 \lambda [\mu\text{m}] \sqrt{I_0 [10^{18} \text{ W/cm}^2]}$. Intensities where $a_0 > 1$ are called relativistic intensities.

Gaussian beam propagation

Determining how the Gaussian beam radius propagates is not straight-forward, but it can be shown that the beam radius propagation, somewhat akin to the β function propagation (see Section 1.1.4), is governed by [39, 40]

$$\frac{d^2 W}{ds^2} = \frac{4}{k^2 W^3} \left(1 - \frac{\eta_0^2 k^2 C^2}{4} W^4 \right), \quad (1.54)$$

where the leftmost term on the right-hand side is related to propagation in vacuum, and the rightmost term is related to the parabolic refractive index in Eq. 1.42. In vacuum, i.e. $C = 0$, Eq. 1.54 has the solution

$$W(s) = W_0 \sqrt{1 + \left(\frac{s}{s_R} \right)^2}, \quad (1.55)$$

which bears a striking resemblance to Eq. 1.28. β^* is evidently equivalent to the diffraction/Rayleigh length for a particle beam, and for both light and (massive) particles, the physical waist radius depends on the square root of this diffraction length.

Comparing to Eq. 1.28, we can see that for free-space propagation of Gaussian beams, there is an equivalence between particle the beam emittance and the wavelength of light, $\varepsilon \leftrightarrow \lambda/\pi$.

If we, similarly to what we did in Section 1.1.4 (c.f. Eq. 1.29 and Fig. 1.5), focus a Gaussian beam onto a medium with parabolic refractive index of a specific steepness (see Eq. 1.42), $C^2 = 4/(\eta_0^2 k^2 W^4)$, the right-hand side of Eq. 1.54 vanishes, and the solution is $W = W_m = \sqrt{2/(\eta_0 k C)} = \text{constant}$, since $dW/ds = 0$ at the focus. Here, we have defined the matched beam radius, W_m , which depends on the refractive index and the steepness of the parabolic channel, as well as the wavelength of the light. Analytical solutions for non-matched cases are more difficult to calculate. Qualitatively, the scalloping of a mismatched laser beam looks similar to a mismatched particle beam, see Fig. 1.5.

1.2.4 Electromagnetic pulse propagation

The analysis so far has been for a continuous beam, but in the topics treated in this thesis, ultra-short *pulses* are also a common topic for both light and electrons. We will now consider an electric field envelope that is time dependent. Assuming a Gaussian pulse in time, the electric field envelope $E(\mathbf{r}, t)$ now explicitly depends on time as $E(t) = E_0 \exp[-2 \ln 2 t^2 / \tau_\ell^2]$, where τ_ℓ is the fwhm pulse duration (of the intensity), such that the pulse intensity becomes

$$I(\rho, s, t) = I_0 \left(\frac{W_0}{W(s)} \right)^2 \exp \left[-\frac{2\rho^2}{W^2(s)} \right] \exp \left[-\frac{4 \ln 2 t^2}{\tau_\ell^2} \right]. \quad (1.56)$$

The peak intensity of this Gaussian pulse in time and radius is given by, c.f. Eq. 1.31,

$$I_0 = \frac{1}{(2\pi)^{3/2}} \frac{\mathcal{E}_\ell}{\sigma_\rho^2 \sigma_t} = \frac{4\sqrt{\ln(2)}}{\pi^{3/2}} \frac{\mathcal{E}_\ell}{W_0^2 \tau_\ell} = \frac{2}{\pi} \frac{P_0}{W_0^2}, \quad (1.57)$$

where $\tau_\ell = 2\sqrt{2 \ln(2)} \sigma_t$ is the fwhm pulse duration, $P_0 = \mathcal{E}_\ell / (\sqrt{2\pi} \sigma_t) = 2\sqrt{\ln 2} \mathcal{E}_\ell / (\sqrt{\pi} \tau_\ell)$ is the pulse peak power, and \mathcal{E}_ℓ is the total pulse energy. Within the scope of this thesis, pulse energies of $\mathcal{E}_\ell \approx 1$ J and durations of ~ 30 fs were used, which brings the power to around 30 TW. With relatively tight focusing, yielding focal spot $1/e^2$ radii around $W_0 \approx 11$ μm , this brings the intensity up to $I_0 \approx 1 \times 10^{19}$ W/cm². With a typical laser wavelength of $\lambda = 800$ nm, Eq. 1.50 gives $a_0 \approx 2.5$. The intensity, power and electric field of such a laser pulse are shown in Fig. 1.7.

An infinitely long, continuous beam of light has one single frequency (or wavelength), but a pulse has a spectrum of frequencies. These different frequencies can be regarded as many “small” waves with different frequencies which add up in time to form a larger wave packet, or pulse. The shorter the pulse, the more frequencies are required, i.e. the broader the spectrum. The width of the spectrum, $\Delta\omega$, is referred to as the *bandwidth* of the pulse, and the frequency around which the spectrum is centered, ω_0 , is called the *carrier frequency*. The increase in spectral width can appear quite subtle. The spectrum of the pulse, $S(\omega)$, also called the spectral intensity, is obtained through a Fourier transformation of the temporal waveform, $S(\omega) \propto |\mathcal{F}\{E(t)\}|^2$. Since the temporal shape is here assumed Gaussian, the spectrum is also Gaussian. The relation between the fwhm spectral width and the temporal minimum, or *Fourier*

transform-limited, fwhm pulse duration is [36]

$$\tau_\ell \Delta f = \frac{2 \ln 2}{\pi} = 0.441 \quad \leftrightarrow \quad \tau_\ell \Delta \omega = 4 \ln 2 = 2.77. \quad (1.58)$$

This is a fundamental limitation for electromagnetic pulses which does not exist (at least practically) for relativistic particle bunches, because the particle energy is too high for the particles to display any wave-like behavior. If the spectrum of a transform-limited pulse of light is made narrower, its duration automatically increases. For a transform-limited pulse, the phase of all the frequencies line up in time, i.e. they are in phase, and the waveform will appear as having the carrier frequency $\omega = \omega_0$ everywhere, only that the pulse amplitude changes. By changing the phase between the waves, the frequencies will begin to change with time along the pulse, which is called a *chirp*. Just like when removing frequencies from the spectrum, chirping a transform-limited pulse makes it longer in time.

Temporal dispersion and non-linear effects

An interesting effect that can arise when the intensity of the pulse becomes high enough is that the pulse itself can start to affect the refractive index of the material it is propagating in. In certain media, called Kerr media, the refractive index then gets an additional term, $\eta(I) = \eta + \eta_2 I(\rho, t)$, where η_2 is a material-specific constant [36, 41]. The intensity has a Gaussian shape, which as we saw in Eq. 1.42 can be approximated by a parabola, so that we again get a transversely parabolic refractive index that has a maximum on axis. As before, this means that the phase velocity is higher in the wings of the pulse than in the center, which will focus the pulse. The pulse creates its own focusing GRIN medium, and the effect is called *self-focusing*.

However, the refractive index is generally also frequency dependent, $\eta = \eta(\omega)$, which means that a pulse like the one described above will see the different spectral

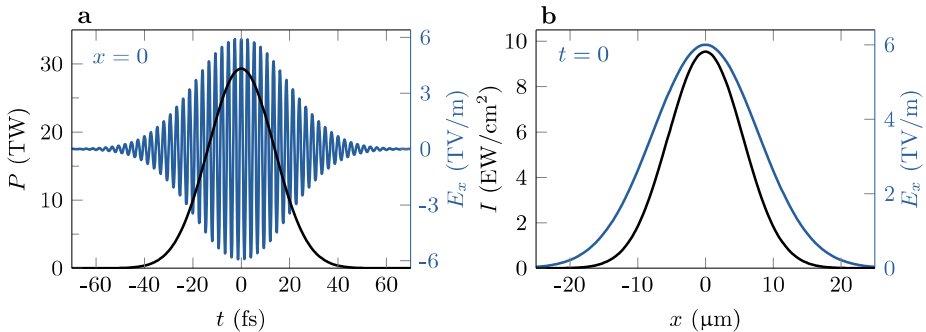


Figure 1.7: Temporal and transverse distributions of the transverse electric fields of a Gaussian laser pulse. **a**, Laser pulse average power (black) and corresponding electric field (blue) on the axis for a 1 J, 32 fs fwhm laser pulse at 800 nm central wavelength, as a function of time. **b**, Transverse distribution of the field (blue) of the pulse in **a** at $t = 0$, when focused to an intensity (black) $1/e^2$ radius of 11 μm . Half a cycle later (and before), this field is negative with almost the same amplitude. The peak intensity is almost $1 \times 10^{19} \text{ W/cm}^2$.

components behave differently. A medium with such a refractive index is called a *dispersive* medium. Dispersion is also evident in ray optics, and causes different refraction angles for different wavelengths, which causes chromatic aberrations in for example lenses. If the lens material is dispersive, the focal length will vary with the wavelength, similar to particle optics (see Section 1.1.4). Furthermore, since the wavenumber (and wavelength) is related to the phase velocity, we get that $k = \omega\eta/c \rightarrow \omega\eta(\omega)/c$, which has great implications.

The first implication is that the entire pulse slows down when propagating. It can be shown that the envelope will travel with a reduced velocity of [36, 42]

$$\frac{1}{v_g} = \frac{dk}{d\omega} = \frac{1}{c} \frac{d}{d\omega} \omega\eta(\omega), \quad (1.59)$$

where v_g is called the *group velocity*. It is therefore warranted to define a *group refractive index*, $\eta_g = c/v_g$. The slow-down of the group velocity means that there will be a time delay between a pulse propagating in vacuum and a pulse propagating in a dispersive medium, and that the electric field amplitude will move underneath the envelope (c.f. Fig. 1.7). The second notable effect is called *group velocity dispersion*, where the envelope not only slows down, but also stretches in time, i.e. the pulse acquires a chirp. The chirp comes from the fact that the different frequency components of the pulse move with different group velocities, which means that the phase between them will change. Group velocity dispersion is typically characterized by the group velocity dispersion coefficient [36, 43]

$$D_\omega = \frac{d}{d\omega} \frac{1}{v_g} = \frac{1}{c} \frac{d^2}{d\omega^2} \omega\eta(\omega), \quad (1.60)$$

which will yield a final rms pulse duration of $\sigma_t^2 = \sigma_{t,0}^2 + \sigma_D^2$, where $\sigma_{t,0}$ is the original, Fourier transform-limited rms pulse duration and $\sigma_D = |D_\omega| \sigma_\omega L$ is the rms temporal increase from group velocity dispersion. The bunch will become chirped, with the chirp depending on the slope of $\eta(\omega)$, but its frequency content, i.e. the spectrum, will not change. If $D_\omega > 0$, the frequency under the envelope will increase with time (longer wavelengths will be at the head of the pulse), and vice versa for $D_\omega < 0$. If a chirped pulse is sent into a dispersing medium with the appropriate sign of D_ω , the effect will be to compress the pulse in time. There are higher-order dispersion effects than these, depending on e.g. $d/d\omega(D_\omega)$, but they are outside the scope of this thesis.

Similarly to the transverse self-focusing effect, a Kerr medium also has an effect on the longitudinal part of intense enough pulses, with an effect called self-phase modulation [36, 44]. The refractive index is modified in the same way, $\eta(I) = \eta + \eta_2 I(\rho, t)$, which also in the longitudinal dimension adds a phase component, here $\delta\phi(t) = -\eta_2 I(t) k_0 s$, which in fact corresponds to the creation of additional frequencies in the pulse. The frequency shift depends on the slope of the intensity, and so is different for the front and the back of the pulse - if η_2 is positive, the head of the pulse will exhibit lower frequencies (red-shift) and the tail higher frequencies (blue-shift). This anti-symmetric frequency shift means that the pulse will become chirped without the pulse lengthening. If paired with a medium that exhibits the appropriate sign of $D_\omega < 0$, the pulse will therefore also self-compress.




Figure 1.8: Ionization intensities. Ionization appearance intensities (calculated with Eq. 1.61), together with corresponding electric field and a_0 , for the most commonly used elements throughout the work for this thesis. The laser pulse parameters are the same as for Fig. 1.7: 1 J, 32 fs fwhm, 13 μm fwhm.

1.3 Plasmas

1.3.1 What is a plasma?

The final topic that we need to examine before moving on to describe the actual electron accelerators, is plasmas. A plasma is a state of matter [45], just like the usual three - solid, liquid and gas - which are typically referred to as “the” states of matter in for example high-school textbooks. Plasma is in fact the most common state of (ordinary) matter in the universe, as it is what stars are made of. Other states of matter exist, but these are not directly relevant for this work.

Starting with a solid material, the atoms or molecules are held together in tight proximity by interactions between the electrons in the material and are relatively stationary. When applying energy (e.g. heat) to a solid material it will eventually melt and become liquid. In a liquid, the attractive forces between the molecules are not strong enough to keep them completely stationary with respect to each other, and they become more free to move around. When applying more energy, the liquid starts to boil and the material becomes a gas, and the molecules become completely free to move with respect to each other. Applying even more energy instigates a rather dramatic change to the material; where it was previously microscopically charge-neutral, with an equal amount of positive charges (atomic nuclei) and negative charges (electrons) bound tightly together into molecules or atoms, the additional energy content in the material is now high enough that the nuclei cannot keep their electrons bound, i.e. the atoms are ionized. This ionization forms a “soup” of free electrons and positive ions, which is called a plasma. The plasma as a whole is often still neutral, at least on average, but there can now be large, local separations of charge inside the plasma itself. As the medium is already ionized, it cannot really sustain any damage - in a sense it is already “destroyed”, which, as we will see later, is a large benefit for electron acceleration.

In most cases considered within this thesis, the medium is assumed to be already

ionized when the laser or beam interacts with it. For laser pulses of the intensities relevant herein, with gases of the relevant elements (see below), this assumption is often good, since the rising edge of the pulse creates the plasma before the peak of the pulse arrives. There are a few different ionization mechanisms, of which two are of most interest in this thesis - *single-photon ionization* (photoelectric effect), exploited in photocathode electron guns (see Section 2.2), and *field ionization*, which is typically what generates the plasmas we are considering with laser pulses. In particular, the “flavor” of field ionization called over-the-barrier ionization is of great importance. Each electron in an atom or molecule is bound to its ion with some ionization energy, \mathcal{E}_{ion} . Using hydrogen as an example, the electron is bound in a spherically symmetric Coulomb potential (corresponding to the field in Eq. 1.8 with opposite sign). It is possible to calculate how strong an external electric field must be to skew this symmetric potential enough that it dips below the electron binding energy, which releases the electron. The corresponding applied intensity is called the *appearance intensity* and is given by [46]

$$I_{\text{app}} = \frac{c}{128\pi e^6} \frac{\mathcal{E}_{\text{ion}}^4}{Z^2}, \quad (1.61)$$

where Z is the charge state of the created ion. Figure 1.8 shows the time-averaged laser peak intensity, with corresponding electric field amplitude and a_0 , versus time, with the appearance intensity for some ions of commonly used elements within the work of this thesis marked with dashed lines. The main constituent of the plasma is often helium or hydrogen, such that the plasma is fully ionized at 50-60 fs before the peak of the laser pulse. In contrast, the lowest-lying states in nitrogen, which is often used as a dopant on the few-% level, are only ionized near the very peak of the pulse - this effect can be used for electron bunch generation and will be covered in Section 3.2.1. A comprehensive table of different ionization energies can be found in Ref. [47].

1.3.2 Electromagnetic waves in plasmas

Plasmas interact differently with electromagnetic waves than “normal” media. A plasma is normally characterized by its electron number density, n_p ,² which is given as the number of free electrons per unit volume. Since these electrons are free to move, they will be accelerated along applied electric fields, which creates a charge density separation between the electrons and the background ions. The ions are at least 1836 times heavier than the electrons [48] (for hydrogen ions), and will typically be regarded as stationary on the time scales relevant in this thesis. The charge density separation is akin to a capacitor, with one positively charged and one negatively charged end with an electric field in between. The electric field causes the electrons to be attracted back towards the ions, but because of the energy they have gained, they will overshoot the ions when they reach them, thus setting up an oscillation about some equilibrium position. This causes a coupled wave of electromagnetic fields and electron motion, i.e. a plasma wave, to propagate in the plasma.

The description of the wave above is in fact closest to one particular type of plasma wave, of which many exist [49]. It is an electrostatic wave called a Langmuir wave, and

²We are only considering electron densities in this thesis. In sources where an ambiguity might occur, notations such as n_e or n_{pe} are often used.

its electric field points along the direction of the wave's motion, unlike for example a laser field in vacuum. The longitudinal motion of the electrons looks qualitatively similar to the motion of molecules in an acoustic pressure wave, although the phenomenon is different and the frequency of oscillation is much higher in a plasma oscillation. Langmuir waves are drastically different from electromagnetic waves in vacuum or dielectric media, but are perhaps the easiest to analyze analytically. For relatively gentle electron displacements, the collective, non-relativistic plasma oscillation behaves as a harmonic oscillator with a frequency [21, 50],

$$\omega_p = \sqrt{\frac{n_p e^2}{m_e \epsilon_0}}, \quad (1.62)$$

called the plasma (electron) frequency. The phase velocity $v_p = \omega_p/k$ of the plasma wave can have any value, and the group velocity is $v_g = 0$.

So far, we have not mentioned what might “trigger”, or *drive*, a plasma wave in the first place. Within this thesis, two methods are used: ultra-short laser pulses and relativistic electron bunches with durations of a few 10's of fs. There are several different regimes in which plasma waves can be driven, see Ref. [51], but the focus of this thesis is on laser-wakefield acceleration (LWFA) and (beam-driven) plasma-wakefield acceleration (PWFA), where the driving pulse/bunch is shorter than the plasma wavelength (which will be defined in Section 1.3.3). The temporal and transverse distributions of the transverse electric field of a laser pulse and a relativistic electron bunch are shown in Fig. 1.6 and 1.7. That ultra-relativistic electrons, with their transverse, disk-like space-charge electric field (see Fig. 1.2), can expel the plasma electrons from the axis and drive a plasma wave feels (at least to me) rather intuitive. That a laser pulse with its oscillating electric field can also drive a wake is perhaps less intuitive, so we will start by examining the propagation of a laser pulse in a plasma and the impact it has on the ambient electrons.

Laser propagation

Like vacuum and dielectrics, plasmas allow propagation of waves with transverse electric and magnetic fields, such as the laser pulses treated in Section 1.2.4. The frequency and wavenumber of these waves must fulfil the dispersion relation [21]

$$\omega^2 = \omega_p^2 + c^2 k^2. \quad (1.63)$$

This relation tells us that if the plasma frequency, which only depends on the plasma density, is higher than the carrier frequency of the laser pulse, $ck = (\omega^2 - \omega_p^2)^{1/2}$ becomes imaginary, which means that the pulse cannot propagate (spatially) inside the plasma. Such a high-density plasma is called *over-dense*. When the electron density is high enough, the collective electron motion can cancel the fields of the incoming pulse, instead mostly reflecting it. The collective electron motion makes the plasma behave like a reflective metal, which can in fact be regarded as housing a very high-density electron plasma because of the many free electrons in the metallic crystal structure. Part of the pulse still penetrates into the surface as an evanescent wave [36], which is exponentially decaying with depth, and these waves can under the right conditions propagate along the material surface [52].

Since they do not allow for a laser pulse to propagate, over-dense plasmas are not particularly useful for electron acceleration. However, plasma densities with a corresponding plasma frequency smaller than the pulse carrier frequency do allow for propagating waves. In such *under-dense* plasmas, the collective electron motion is too slow with respect to the laser pulse frequency, so that it does not have time to rearrange and cancel the incoming fields. Equation 1.63 then yields the laser phase velocity as

$$v_{ph} = \frac{\omega}{k} = \frac{c}{\sqrt{1 - \frac{\omega_p^2}{\omega^2}}}, \quad (1.64)$$

which together with Eq. 1.59 and 1.60 yields the laser group velocity and group velocity dispersion

$$\begin{aligned} v_g &= \left(\frac{dk}{d\omega} \right)^{-1} = c \sqrt{1 - \frac{\omega_p^2}{\omega^2}}, \\ D_\omega &= \frac{d^2k}{d\omega^2} = -\frac{1}{c} \frac{\omega_p^2}{(\omega^2 - \omega_p^2)^{3/2}}. \end{aligned} \quad (1.65)$$

From these relations it is evident that the phase velocity of the pulse in a plasma is always higher than the phase velocity in vacuum, c , but that the group velocity is lower than its vacuum counterpart. The same situation arises in for example hollow, metallic microwave waveguides [53], where the phase velocity of given modes can exceed c , while the group velocity does not. It is therefore warranted to introduce the refractive indices

$$\begin{aligned} \eta &= \left(1 - \frac{\omega_p^2}{\omega^2} \right)^{\frac{1}{2}}, \\ \eta_g &= \left(1 - \frac{\omega_p^2}{\omega^2} \right)^{-\frac{1}{2}}, \end{aligned} \quad (1.66)$$

which only depend on the plasma density and the frequency of the incoming wave. For $n_p = 5 \times 10^{18} \text{ cm}^{-3}$ and $\lambda = 800 \text{ nm}$, $(1 - \omega_p^2/\omega^2)^{1/2} \approx 0.9986$, which is a small but significant difference from the in-vacuum values.

Single-electron motion in the laser fields

Just as for “regular” media, propagation of intense laser pulses through plasmas introduces other effects. These non-linear effects are the foundation for most of the research on plasma-based accelerators, as they are the basis of the accelerating structures. To exemplify this, we will first look at the response of a single electron as it encounters an intense electromagnetic plane wave. Figure 1.9 shows the transverse (along the polarization) and longitudinal trajectory of an electron for three wave amplitudes, both in the laboratory frame (1.9a) and the average electron drift frame (1.9b). For amplitudes corresponding to relativistic intensities, the motion is no longer harmonic, but becomes “spiky”, and there is significant lengthening of the spatial period. For $a_0 \approx 1$, the longitudinal motion begins to dominate since the electric field makes the electron relativistic in less than half a cycle, such that the wave’s magnetic field is no

longer negligible. The contribution of the magnetic field is a forward force on the electron (see Eq. 1.1, 1.11 and 1.46). Solving Eq. 1.6 for the relativistic case, for a laser electric (and corresponding magnetic) field on the form $\mathbf{E}(s, t) = E_0 \cos(ks - \omega t) \hat{\mathbf{x}}$, and with the local coordinate $z = ct - s$, the electron trajectory becomes [46]

$$\begin{aligned} x &= -\frac{a_0}{k} \cos(kz), \\ s &= \frac{a_0^2}{8k} [2kz - \sin(2kz)], \end{aligned} \quad (1.67)$$

which is shown in Fig. 1.9.

The electron motion in Eq. 1.67 and Fig. 1.9 is for an infinite plane wave, and while it illuminates some important effects on the electron motion, intense laser pulses (see Section 1.2) are more relevant to study in the context of the thesis. That the wave is temporally and spatially confined has some important consequences for the plasma response. If we instead insert an electromagnetic field corresponding to Eq. 1.56 into Eq. 1.1 and average the force in time over several laser cycles, the resulting force can be cast on the form [54]

$$\mathbf{F}_p = -\frac{e^2}{m_e c \epsilon_0 \omega_0^2} \nabla \langle I \rangle = -\frac{e^2}{2m_e \omega_0^2} \nabla \langle |\mathbf{E}|^2 \rangle = -\frac{m_e c^2}{2} \nabla \langle |\mathbf{a}|^2 \rangle, \quad (1.68)$$

where $\langle \cdot \rangle$ here denotes the central moment in time. We can see three notable features of this so-called *ponderomotive force*. Firstly, the force itself is inversely proportional to mass, $\mathbf{F}_p \propto m_e^{-1}$, which means that the acceleration due to this force will be proportional to m_e^{-2} . For the laser parameters relevant within this thesis, we are safe in considering the ions to be stationary during a laser cycle, and also during the laser pulse duration. Secondly, the force is proportional to the negative gradient of the intensity, $\mathbf{F}_p \propto -\nabla \langle I \rangle$, so the force acts away from regions of high intensity, and regardless of the sign of the particle. This gradient force means that the force acts




Figure 1.9: Trajectories of a free electron in a strong laser field. **a**, Single-electron trajectories in the $x - s$ plane in a strong laser field for three different laser strengths. The trajectories are shown in the laboratory frame. **b**, Corresponding trajectories in the average electron frame. Stronger laser fields yield larger excursions in the longitudinal direction - a slow, non-relativistic electron would only be a vertical line.

as a form of light pressure, and it springs up from that the fields of the laser are not identical from cycle to cycle in time or radially in space. The change in applied field from cycle to cycle causes the electrons to not be accelerated in the same way from one half-cycle to the next. The ponderomotive force exists both transversely and longitudinally. Thirdly, the force increases with the laser wavelength, $\mathbf{F}_p \propto \lambda^2$, so for a strong response, we want a long-wavelength laser (seen in Eq. 1.67). This is caused by the slower oscillations imparting more transverse momentum on the electrons within one half-cycle.

1.3.3 Laser-driven plasma waves

The collective response of the plasma electrons, i.e. the “actual” plasma wave, to the ponderomotive force is trickier to describe analytically than the single-electron response to a plane wave (Eq. 1.67 and Fig. 1.9). It also becomes increasingly difficult with the dimensionality of the problem, i.e. if the problem is approached in one, two or three spatial dimensions, and for non-linear waves. It is for these theoretical descriptions where the normalized potentials (Eq. 1.49) become useful, and the apparent detour of introducing them pays off.

Linear plasma waves

For a weak laser pulse, $a_0 \ll 1$, a relatively small fraction of the background plasma electrons are displaced by the pulse, and the modulation of the plasma electron density can therefore be considered as a perturbation, $\delta n_p = n_p - n_0$, where n_0 is the uniform background plasma density and n_p is the local plasma density. The scalar potential for this density perturbation, in three dimensions, is described by [37, 54]

$$\left(\frac{\partial^2}{\partial z^2} + k_p^2 \right) \varphi = \frac{k_p^2}{2} \langle \mathbf{a}^2 \rangle, \quad (1.69)$$

where $\langle \mathbf{a}^2 \rangle = a^2/2$ for a linearly polarized laser. Equation 1.69 has the form of a driven harmonic oscillator, where the driving term is dependent on the laser intensity, and has the general analytical solution [37],

$$\varphi = \frac{k_p}{2} \int_z^\infty \sin [\omega_p(z - z')] \langle \mathbf{a}^2(z') \rangle dz', \quad (1.70)$$

which allows for arbitrary laser pulse shapes. Equation 1.70 shows that the plasma wave potential follows in the wake of the laser, and is hence normally called the *wake potential*. Knowing φ gives easy access to both electric fields and plasma electron density variations in the wake, as (see Eq. 1.44c, 1.47 and 1.49)

$$\begin{aligned} \mathbf{E} &= -\nabla\Phi = -\frac{m_e c^2}{e} \nabla\varphi, \\ \frac{\delta n_p}{n_0} &= -\frac{\epsilon_0}{e} \nabla \cdot \mathbf{E} = \frac{1}{k_p^2} \nabla^2 \varphi. \end{aligned} \quad (1.71)$$

Not all laser pulse shapes give analytical solutions, but one that does is a Gaussian pulse in space and time (see Eq. 1.56), which (after the pulse has passed) gives the




Figure 1.10: Plasma density perturbation and electric fields in a linear plasma wake. **a**, Laser a envelope (red) and plasma electron density perturbation (black) for $a_0 = 0.5$, $\tau_\ell = 32$ fs, $\lambda = 800$ nm and $n_p = 5 \times 10^{18}$ cm^{-3} . The relative change in density is on the order of a few percent. The shaded area is extended in 2D in Figs. **b-d**. **b**, Plasma electron density perturbation as function of longitudinal and radial coordinate, within the longitudinal bounds shown in **a**. Darker shade shows areas of higher electron density, and lighter areas show lower density. As in Fig. 1.6, note that ρ does not actually have negative values. **c**, Longitudinal electric field strength in the wake. Positive field (red) is directed to the right and negative (blue) is directed to the left. The shaded area (see also **d**) marks a region which is both accelerating and focusing for an electron placed therein. Remember that the force on an electron in an electric field is directed opposite to the direction of the field (Eq. 1.1). **d**, Radial electric field strength in the wake. The field in the red areas is directed upward in the figure and in blue areas directed downward.

wakefields [37]

$$\begin{aligned}
 E_s &= \frac{\chi_\tau a_0^2 E_{wb}}{4} \exp\left[-\frac{2\rho^2}{W^2}\right] \cos(k_p z), \\
 E_\rho &= -\frac{\chi_\tau a_0^2 c E_{wb}}{\omega_p W^2} \rho \exp\left[-\frac{2\rho^2}{W^2}\right] \sin(k_p z), \\
 \chi_\tau &= \sqrt{\frac{\pi}{4 \ln 2}} \omega_p \tau_\ell \exp\left[-\frac{\omega_p^2 \tau_\ell^2}{16 \ln 2}\right],
 \end{aligned} \tag{1.72}$$

and subsequently the plasma density wake

$$\frac{\delta n_p}{n_0} = -\chi_\tau a_0^2 \left(\frac{1}{4} + \frac{2}{k_p^2 W^2} - \frac{4\rho^2}{k_p^2 W^4} \right) \exp\left[-\frac{2\rho^2}{W^2}\right] \sin(k_p z), \tag{1.73}$$

where

$$E_{wb} = \frac{m_e c \omega_p}{e} \approx 96 \sqrt{n_p [10^{18} \text{ cm}^{-3}]} \text{ [GV/m]} \tag{1.74}$$

is the maximum achievable electric field strength in a linear wake, also called the *wavebreaking field*, and χ_τ is a dimensionless laser pulse-duration dependent factor for the Gaussian pulse. The electron density and the electric fields are plotted in Fig. 1.10. The first term in brackets in Eq. 1.73 comes from the longitudinal field, and the second two from the transverse field. Longitudinally, the wake is simply a sinusoidal oscillation, while it obtains roughly the laser profile radially. The main contribution to the density wake is longitudinal ($2/(k_p^2 W^2) \approx 0.09 < 1/4$ for the laser parameters given in Fig. 1.10), similar to a Langmuir wave. It is now clear that the phase velocity of this plasma wave is governed by the group velocity of the driving pulse, $v_p \approx v_g$, and as such we can define a plasma wavelength $\lambda_p = 2\pi v_p / \omega_p \approx 2\pi c / \omega_p$, which is a natural scale length of the wake. The plasma wavelength is marked in Fig. 1.10a.

The fields formed in the laser wake are the *wakefields*, and it is these that can be used to perform the electron acceleration that we are interested in. The electric field points from regions of lower electron density to higher electron density, which shows up as an offset of $\pi/2$ in phase between the longitudinal and radial field in the linear wake (see Fig. 1.10c and d). The phase difference between longitudinal and transverse fields means that only small sections of $\lambda_p/4$ are simultaneously accelerating and focusing for a test electron inside the wake, marked as dashed regions in Fig. 1.10c and d. From χ_τ in Eq. 1.72, we can identify a resonant condition when the electric field in the wake is maximized. When the pulse duration fulfills $c\tau_\ell = \sqrt{8 \ln 2} / (2\pi) \lambda_p \approx 0.37 \lambda_p$, the forward and backward push of the laser on the plasma electrons, from the ponderomotive force on the front and rear intensity slope, respectively, is optimally synchronized with the natural motion of the plasma electrons from the restoring space-charge force. In general, it is desirable to have a pulse duration shorter than half the plasma wavelength for all regimes of LWFA, so as to make the best use of the available laser energy. For $n_p = 5 \times 10^{18} \text{ cm}^{-3}$, the optimal pulse duration is quite short, at $\tau_\ell \approx 19 \text{ fs}$ fwhm.

1D non-linear plasma waves

As a_0 increases, so does the amplitude of the plasma waves, and at some point the collective oscillations of the plasma electrons are not simply harmonic anymore, even




Figure 1.11: Plasma density and electric field in a 1D non-linear plasma wake. Plots of the laser a envelope (red), plasma electron density (black) and longitudinal electric field (blue) for $\tau_\ell = 32$ fs, $\lambda = 800$ nm and $n_p = 5 \times 10^{18} \text{ cm}^{-3}$ for three values of a_0 , $a_0 = 0.5, 1$ and 2 , as in Fig. 1.9. **a**, $a_0 = 0.5$, which also corresponds to Fig. 1.10. **b**, $a_0 = 1$. Note that the wake is starting to deviate from a sinusoidal function. **c**, $a_0 = 2$. The wake is now significantly non-linear. Note that the plasma wavelength has increased. **d**, Similar plots to **a-c**, but for a plasma of $n_p = 1 \times 10^{17} \text{ cm}^{-3}$ and with an electron beam driver (see Section 1.3.4) with $Q_b = 40$ pC, $\tau_b = 60$ fs and $\sigma_p = 4.25 \mu\text{m}$, for $n_b/n_p \approx 1.15$. The lack of forward force from a beam driver can be seen in the lack of plasma density buildup at the driver head, c.f. **c**.

after the laser has passed. When $a_0 \gtrsim 1$, the assumption of the plasma response, φ , being a small perturbation does not hold, and the potential Eq. 1.69 takes the form [54]

$$\frac{\partial^2 \varphi}{\partial z^2} = k_p^2 \gamma_p^2 \left[\beta_p \left(1 - \frac{1 + a^2}{\gamma_p^2 (1 + \varphi)^2} \right)^{-\frac{1}{2}} - 1 \right] \approx \frac{k_p^2}{2} \left[\frac{1 + a^2}{(1 + \varphi)^2} - 1 \right], \quad (1.75)$$

where $\beta_p = v_p/c$ and $\gamma_p = (1 - \beta_p^2)^{-1/2}$ are the normalized velocity and relativistic factor of the plasma wave, respectively. The right-hand side of Eq. 1.75 is obtained by assuming $\gamma_p \gg 1 \leftrightarrow \beta_p \lesssim 1$, i.e. the laser pulse moves easily through the sufficiently under-dense plasma, while Eq. 1.69 can be subsequently retrieved by expanding for small φ . Equation 1.75 can be solved analytically in one dimension, but to simply demonstrate the effects, its numerical solutions are shown in Fig. 1.11a-c for different a_0 . It is clear that the wave is no longer sinusoidal for $a_0 > 1$, and because of the relativistic mass increase of the electrons in the strong laser field, it takes them longer to be attracted back to their equilibrium position, which lengthens the plasma wavelength. In fact, one half of the plasma period is stretched and the other is compressed, leading to strong longitudinal asymmetry, as seen in for example Fig. 1.11c. There is a buildup of electrons at the back of each wake period and an almost linear electric field gradient inside the constant electron density between these electron density peaks.

The bubble regime

As the above model is only one-dimensional, the electrons can only move in the longitudinal direction, and so there are still electrons left inside each period (the density is non-zero) as they are forced to move along the wave. They are therefore also forced to move over the peak a_0 of the pulse. In three dimensions, the electrons are also pushed radially outwards by the ponderomotive force, which lowers the electron density inside each wake period because the electrons can then also travel around the sides of the driving pulse. In this case, the electrons will typically not traverse the peak of the pulse, as they largely flow around it. The end result of this 3-dimensional motion is that if we have an intense laser of $a_0 \gtrsim 2$, it will act akin to a “snow plow”, or a bullet being fired into a water tank, and push virtually all the electrons away from its propagation axis [55]. This strong interaction forms a quasi-spherical, bare *ion cavity*, or “bubble”, in the wake of the laser pulse, and hence this regime is called the *bubble regime*, or *blowout regime*. A plot of the plasma electron density in such a bubble, from a particle-in-cell (PIC) simulation, is shown in Fig. 1.12. The bubble is somewhat shorter than the 1D non-linear theory would suggest, since the electrons gain less momentum as they do not see the peak laser field. The electric field inside a fully blown-out bubble is radial nearly everywhere, with a linear gradient outward from the bubble center [55]. This means that the fields are transversely focusing for an electron everywhere inside the bubble, and that rear part of the bubble is also accelerating. As will be discussed further in Chapter 3, the blowout regime is therefore excellent for accelerating electrons. The maximum accelerating electric field in the bubble regime is given by [56]

$$E_{z,\max} = E_{wb} \sqrt{a_0}, \quad (1.76)$$

with E_{wb} from Eq. 1.74. The accelerating field gradient is approximately linear with z and zero at the center of the bubble, such that the longitudinal field can be estimated as $E_z \approx -E_{z,\max}(z - z_0)/R_B$, where R_B is the bubble radius and z_0 marks the center of the bubble.

Apart from having $a_0 \gtrsim 2$, the pulse should ideally also fulfil some other requirements to produce an efficient blowout [55]. Just as in the linear regime, the pulse length can be matched to the dimensions of the plasma wave to more efficiently excite it. However, since the bubble is 3-dimensional, not only longitudinal, but also transverse, matching is possible. Because the laser is so intense, it creates both a longitudinal (see Fig. 1.11c) and transverse gradient of the plasma refractive index, which arises from the increase in relativistic momentum of the plasma electrons and the local gradient in plasma density (c.f. Eq. 1.66). The longitudinal density gradient arises from the forward ponderomotive force and the cavitation behind the pulse, just as in the 1D non-linear case, but in a 3D blowout, the bubble can be completely void of electrons, yielding a larger gradient. The longitudinal refractive index gradient can provide both self-phase modulation and self-compression simultaneously, c.f. Section 1.2.4. The 3D structure also provides a large radial density gradient, and since the laser phase velocity, which increases with n_p (see Eq. 1.64), is then larger towards the sides if the pulse, the net effect is radial focusing of the pulse. This focusing effect is similar to the gradient refractive index channel described in Section 1.2.1 and 1.2.3. The increase in transverse momentum for the plasma electrons is proportional to $|\mathbf{a}|$,




Figure 1.12: Plasma density in a laser-driven bubble. Plot of the plasma electron density in a cut through the center of a laser-driven bubble from a quasi-3D particle-in-cell (PIC) simulation. The laser pulse, propagating from left to right in the figure, is not shown, but can be inferred from the vertical lines in the plasma density, which result from the ponderomotive force of the individual laser cycles. The Gaussian $\lambda = 800$ nm laser pulse is polarized in the x -direction, and has a peak normalized vector potential of $a_0 = 4$, radius of $W_0 = 11.2$ μm at $1/e^2$ of the peak intensity and duration of $\tau_\ell = 19.5$ fs fwhm. The ambient electron density is $1 \times 10^{19} \text{ cm}^{-3}$. Courtesy of H. Ekerfelt (see p. 12-13 in Ref. [57] for more details on for example the bubble fields).

which means that it is largest on axis. The more energetic electrons do not interact as strongly with the laser field (their “ability” to cancel the fields decreases because they appear heavier and thus respond slower), which also contributes to the focusing. The effect is therefore called *relativistic self-focusing*, and is a somewhat similar effect to “regular” self-focusing (see Section 1.2.4), but stronger. If matched correctly with respect to the laser spot size, these mechanisms allow the laser to remain focused for many Rayleigh lengths at small spot sizes.

Assuming that \mathbf{a} is Gaussian, a comparison to the analysis of the focusing channel depth in Section 1.2.3 can be done, where it was shown that a parabolic channel of depth parameter $C\eta_0 = 2/(kW^2)$ can guide a Gaussian beam. Similarly to non-relativistic self-focusing, relativistic self-focusing channels are also Gaussian in shape, as they are also dependent on the transverse laser profile, and can thus be approximated by parabolas around the peak. A more thorough analysis gives that $\Delta n/n_p = 4/(k_p W^2)$ [51], which is a rather similar result to that obtained in Section 1.2.3. A full blowout means that $\Delta n/n_p \approx 1$ between the inside of the bubble and the background plasma, so that the condition for self-guiding becomes $W_m \gtrsim 2/k_p = \lambda_p/\pi$. Pulses which are too narrow can therefore not be focused by the channel, but wide pulses can - provided that a_0 is high enough. The relation $a_0^2/8 \gtrsim 4/(k_p W)^2$ between the normalized intensity and area of the pulse gives a condition on the minimum laser power required for relativistic self-focusing to occur, called the *critical power for relativistic self-focusing* [51, 56]

$$P_0 \gtrsim P_c = 2c \left(\frac{e}{r_e} \right)^2 \left(\frac{\omega_0}{\omega_p} \right)^2 \approx 17 \left(\frac{\omega_0}{\omega_p} \right)^2 \text{ GW}, \quad (1.77)$$

where r_e is the classical electron radius.

Similarly to the linear case (c.f. Eq. 1.72), the bubble will be of a similar transverse size to the laser spot size [55], such that $R_B \approx W$. The laser pulse length is also taken to be roughly equal to the bubble radius, $c\tau_\ell \lesssim R_B$ (a more specific optimal duration will be shown in Section 3.1). With these two assumptions, the bubble radius, and thus the matched laser radius, can be shown to be [55]

$$k_p R_B \approx 2\sqrt{a_0} \quad \leftrightarrow \quad W_m = R_B \approx \frac{1}{\pi} \sqrt{a_0} \lambda_p, \quad (1.78)$$

which is similar to the guiding condition above. Recent work [58] has emphasized that the geometry of the bubble is not only dependent on a_0 , but also explicitly on the laser spot radius, W . It is possible to achieve blowout with laser pulses of various aspect ratios (W/τ_ℓ), and the bubble length scales also with W , essentially because of the same reason that the 1D non-linear model overestimates the bubble size - that the laser pulse aspect ratio dictates how many electrons pass over the intensity peak.

Since the self-focusing depends on the local laser intensity and plasma density gradient, the entire pulse is not focused. In particular the front of the pulse is subject to *erosion* (also referred to as *etching*, or simply diffraction), such that the front of the pulse is lost and the pulse appears to erode from the front towards the back. This erosion occurs because while there can be some relativistic self-focusing in the head of the pulse, the increase in plasma density caused by the forward ponderomotive force (see Fig. 1.11c) locally decreases the refractive index on axis, which counteracts the focusing effect. The focusing refractive index profile takes on the order of ω_p^{-1} to

build up [56]. As the laser drives a wake, it also deposits energy into the wake itself. The loss of available laser energy that the diffraction and energy deposition lead to is usually referred to as *pump depletion*, and the (non-linear) *pump depletion length* over which the depletion occurs is given by [56]

$$L_{PD}^{NL} = \left(\frac{\omega_0}{\omega_p} \right)^2 c\tau_\ell. \quad (1.79)$$

If the pulse is long or the bubble radius is very large (e.g. following an ultra-intense laser pulse), the tail of the pulse will find itself in a vacuum, and begin to diffract until it fills the bubble.

1.3.4 Particle beam-driven plasma waves

As alluded to in Section 1.3.2, plasma waves can also be driven by relativistic particle bunches. Within the scope of this thesis, these bunches consist of electrons, but it is also possible to drive wakes using positively charged particles such as protons [59–62] or positrons [63–65]. There are some key differences between laser and particle drivers, some of which can be seen in Fig. 1.6 and 1.7. Most notably, the electric field structure is fundamentally different, with the particle drivers typically having a much weaker transverse electric and magnetic field, making both transverse and longitudinal push on the plasma electrons significantly weaker. However, the field is not time harmonic, and is always directed inward towards the driver bunch center (for an electron driver). Because of this “rectified” electric field, the force on the plasma electrons is always directed away from the driver bunch, and one can efficiently excite high-amplitude wakes also with beam drivers. Another important feature of particle drivers is that their velocity, $c\beta_b$, is independent of the plasma density, which influences the phase velocity of the wake, $v_p = c\beta_b$. The phase velocities of electron beam-driven wakes are therefore usually higher than laser-driven wakes, for the relevant electron energies, plasma electron densities and laser wavelengths. As will be discussed in Chapter 3, this is one of the main reasons of using a particle-beam driver.

For the beam-driven case, we will only examine the non-linear waves. Like in the laser-driven case, there is a parameter which can be used to distinguish between the different regimes. For beam drivers, it is the dimensionless (normalized) bunch charge [66, 67]

$$\tilde{Q} = \frac{N_b k_p^3}{n_p}, \quad (1.80)$$

where $\tilde{Q} \ll 1$ indicates the linear regime, $\tilde{Q} \approx 1$ quasi-linear and $\tilde{Q} > 1$ non-linear, similar to the case for a_0 . \tilde{Q} gives the ratio between the number of bunch electrons and the number of plasma electrons within one cubic plasma skin depth. The response of the plasma to a particle bunch thus differs from the response to a laser pulse in that a_0 is independent of the plasma density.

Similarly to laser drivers, we can describe the wake excitation in the 1D non-linear, beam-driven case, but a small modification to Eq. 1.75 must be made. As no laser is present, $a = 0$, and instead we have a bunch of electrons acting on the plasma. Knowing the electron bunch parameters, it is straightforward to calculate the beam electron density, see Eq. 1.31. From Eq. 1.71 we see that an additional charge density

term (the driver) will show up as a second derivative of the wake potential. The wake potential is thus governed by [51]

$$\frac{\partial^2 \varphi}{\partial z^2} = k_p^2 \left\{ \frac{n_b}{n_0} + \gamma_p^2 \left[\beta_p \left(1 - \frac{1}{\gamma_p^2 (1 + \varphi)^2} \right)^{-\frac{1}{2}} - 1 \right] \right\} \approx k_p^2 \left\{ \frac{n_b}{n_0} + \frac{1}{2} \left[\frac{1}{(1 + \varphi)^2} - 1 \right] \right\}, \quad (1.81)$$

where the right-hand-side is typically a good approximation, as v , and thus v_p , are even closer to c than for a laser driver. The solution of Eq. 1.81 for a 40 pC driver of 60 fs fwhm duration and 10 μm fwhm diameter is shown in Fig. 1.11d. The charge was lowered compared to for example Fig. 1.6 because unrealistic wake elongation is evident at this low bunch electron density of $n_b/n_0 \approx 1.15$, corresponding to $\bar{Q} \approx 0.5$. The 1D non-linear model vastly overestimates the effect of the factor n_b/n_0 , which is accounted for in 3D non-linear theory, shown below.

Beam-driven bubble regime

Also in the beam-driven case, the wakes can be driven into blowout. An example of a beam-driven blowout obtained from a PIC simulation is shown in Fig. 1.13. To reach the 3D blowout regime, the driver electron density must fulfill $n_b \gtrsim 1.8n_p$ [55]. However, from a beam-driver point of view, indefinitely increasing the electron density by decreasing the beam radius after blowout has occurred will lead to no gain in wakefield amplitude: at some point, the driver impulse on the plasma becomes similar to that of a line current density, and the sheath electrons do not receive any additional transverse momentum. The onset of this wake saturation is approximately $n_b/n_p \approx 10$ [55]. Only increasing the driver current by increasing its total charge or decreasing its duration will increase the wake amplitude in this case. A smaller driver radius does excite a blowout further out in the driver head, though, which causes a larger part of the beam to be located inside the focusing wake.

Just as the laser vector potential is normalized to yield \mathbf{a} (see Eq. 1.44, 1.47 and 1.49), the beam's charge density per unit length can be normalized as [55, 68]

$$\Lambda(z) = \frac{k_p^2}{n_p} \int_0^\infty \rho n_b(\rho, z) d\rho. \quad (1.82)$$

Using a bi-Gaussian driver shape as, in Eq. 1.31, yields the *normalized peak current* [55, 68]

$$\Lambda_0 = \frac{n_{b,0}}{n_p} k_p^2 \sigma_\rho^2 = \frac{e^2}{m_e c^2 \epsilon_0} \frac{n_{b,0}}{n_p} n_p \sigma_\rho^2 = \frac{e}{m_e c^2} \frac{1}{(2\pi)^{3/2} \epsilon_0} \frac{Q_b}{\sigma_z}, \quad (1.83)$$

which, c.f. Eq. 1.32, is closely related to the transverse electric field of the bunch. It is usually assumed that $k_p \sigma_z \lesssim \sqrt{2}$ and $k_p \sigma_\rho \ll 1$, meaning that the bunch is smaller than half a plasma wavelength in all dimensions. Just as for a laser driver where $k_p R_B \approx 2\sqrt{a_0}$, a beam driver in the blowout regime gives a bubble radius according to $k_p R_B \approx 2\sqrt{\Lambda_0}$. Even if the blowout regime is reached, beam drivers do not always produce transversely relativistic plasma electrons - for this, one must approach $\Lambda_0 \approx 1$. For the parameters used in Fig. 1.6, $\Lambda_0 \approx 0.37$, and in Fig. 1.13, $\Lambda_0 \approx 0.39$. The quasi-spherical shape, common with laser drivers, is not necessarily obtained - it is possible to have a blown-out ion cavity which has a length near the linear plasma wavelength in a non-relativistic blowout, as is the case in Fig. 1.13 and

Paper IV. When the driver current is increased and the plasma electrons start to become relativistic, the “bubble” is lengthened and the rear portion starts rounding, such that the whole bubble takes the form of a teardrop [55, 68]. Only with very high-current drivers of $\Lambda_0 > 1$ (for example, $\Lambda_0 = 1.1$ for $Q_b = 500$ pC and $\tau_b = 50$ fs fwhm, i.e. $I_b \approx 10$ kA) can quasi-spherical cavities be obtained. Even if the bubble is not spherical, R_B predicts the maximum radius rather well.

When a clear blowout has been reached (relativistic or non-relativistic), the maximum electric field inside the bubble is [55]

$$E_{z,\max} \approx E_{wb} 1.3 \Lambda_0 \ln \left(\frac{10}{\Lambda_0} \right). \quad (1.84)$$

For a non-relativistic blowout, the ion cavity half length is $L_h \approx \lambda_p / (2\sqrt{2})$, so that the longitudinal electric field is approximately $E_z \approx -E_{z,\max} 2\sqrt{2}(z - z_0) / \lambda_p$, where z_0 is the location of the bubble center, where $E_z = 0$. For an ultra-relativistic blowout, the back of the bubble becomes nearly spherical and the cavity half length is approximately equal to the bubble radius. Therefore, the electric field for an ultra-relativistic blowout scales similarly to the laser-driven case for the bubble regime in Section 1.3.3. For the region in between non-relativistic and ultra-relativistic the bubble is approximately elliptical.

As we saw in Section 1.1.4, transversely matched conditions exist for relativistic particle bunches, but not in the same way or through the same mechanism as for laser pulses - transverse matching in the particle beam case is related only to the focusing fields inside the wake and not, for example, the plasma density gradients. Particle beam matching with respect to the wakefield will be discussed further in Section 3.3.2. Nonetheless, both laser pulses and electron bunches share the attribute that the head of the driver, which is outside or in the very front edge of the wake, is subject to erosion due to the lack of wake focusing. However, the part of the driver




Figure 1.13: Plasma density in a beam-driven bubble. Plot of the plasma and bunch electron density in a cut through the center of a non-relativistic beam-driven blowout from a quasi-3D PIC simulation. The driver bunch electron density, propagating from left to right in the figure, is shown in orange-black, where orange is higher, and the plasma density is shown in white-blue, where blue is higher. The driver bunch has a charge of 150 pC, energy of 3 GeV, duration of 43 fs fwhm and diameter of 10 μm fwhm, yielding a peak electron density of approximately $6 \times 10^{17} \text{ cm}^{-3}$. The ambient electron density is $1.1 \times 10^{17} \text{ cm}^{-3}$. Courtesy of H. Ekerfelt (see Ref. [57] for more information).

which is inside the wake is subject to strong focusing and deceleration, so the effective depletion length is in this case a situation-dependent question of which effect is the strongest, energy loss or head erosion. Also, as shown in Eq. 1.32, the space-charge field from the bunch extends far beyond the extents of the bunch itself, so the bubble radius can be significantly larger than the driver radius. Conversely, also drivers which are much smaller than the wake radius are focused, unlike in the laser-driven case.

As the driver only imparts a weak longitudinal force on the plasma electrons, there is no resonant driver length as such, but since the plasma electrons naturally start to return to the beam axis because of the focusing force of the ion cavity, it is more efficient to confine the driver to the first half of the bubble. Assuming that blowout is reached, the driver current profile can be tailored to more efficiently deliver its energy to the wake. It turns out that a linear up-ramped current profile, essentially a triangle with a small “step” at the head, is an optimal shape for a bream driver [69, 70]. This current profile can make the decelerating wakefield seen by the driver constant, such that all the driver particles lose energy to the wake at the same rate. It also increases the *transformer ratio* [71], i.e. the ratio between the decelerating field at the driver and the accelerating field at an accelerating bunch in the rear of the bubble.

Hosing

An effect which is not usually discussed for laser drivers (although it can certainly occur, and has been used in theoretical schemes for X-ray generation [72]) is the *hose instability* (HI), or *hosing*, for short [73]. Here, a variation in spatial and/or angular offset is present along the driver, i.e. the bunch is “tilted” as viewed from the side. For laser pulses, this is usually referred to as *pulse-front tilt* [74]. Because the driver is not fully aligned with its propagation axis, it will push the wake predominantly to one side. The strong focusing inside the wake will push the driver towards the wake center, which is typically located closer to the original propagation axis than the offset driver slice is. As the driver offset starts to be pushed over to the opposite side of the wake, it pushes the wake with it, and so a resonance can appear between the oscillation of the driver and the wake, leading to an instability that can disrupt the wake excitation. Such wake oscillations are also detrimental to any acceleration happening at the back of the wake. It is therefore important to initially reduce such spatial and angular *beam tilts* as much as possible. Reducing beam tilts was a primary objective of the study in **Paper I**, the results of which were implemented in the work of **Paper IV** for that purpose.

Some passive features can be implemented which mitigate the impact of beam tilts in the case that they are present [75, 76]. All of the mitigation methods listed below rely on somehow changing the oscillation frequency along the driver beam, such that not all the driver electrons push to the same side simultaneously. For one, if the beam radius can be made large with respect to the bubble radius, hosing is reduced because of the longitudinal variation in focusing strength along the driver before full blowout is reached. A large-diameter driver could for example be realized by increasing its emittance. A large chirp or slice (uncorrelated) energy spread also mitigates hosing effects, since then not all the beam particles oscillate with the same frequency. The spread in oscillation frequencies makes the oscillations decohere along the beam. Any amount of energy spread improves the situation over no energy spread, but to efficiently mitigate hosing, it must be in the several-percent-range, which is

typically difficult to handle in a controlled way. Another method, which also has other benefits and will be covered in Section 3.3.2, is to use long plasma density ramps going into the accelerating plasma section. Here, the focusing strength, and thus the oscillation frequency, changes during propagation, such that the resonant oscillations are initially damped. The damping in such density ramps adds to the damping effects coming from the other two methods. In essence, it is desirable to have a high-current, high-emittance driver with a large energy spread and high-enough electron density to drive an intrinsically stabilized blowout.

1.4 Summary

We now have the basic tools we need to investigate the actual electron acceleration that this thesis is about. It is clear that electrodynamics is the foundation upon which the entire field is based (and there are, in fact, still more related topics which have not been treated for which electrodynamics are needed). The topic is complex to study in detail, but there are some fundamental points which are important to remember.

To summarize this Chapter, we can note again that to increase the kinetic energy of charged particles, which is the usual meaning of the term “acceleration”, electric fields are required. These electric fields can be created in different ways, the main method of interest here being non-linear plasma waves driven by ultra-intense laser pulses and/or relativistic electron bunches. To keep the particles from simply diffracting away in whichever direction they happen to be going, focusing forces must be applied, and forces which grow linearly with the distance from the propagation axis are the most important, as they do not spoil the emittance of the particle ensembles. As the trajectories of relativistic particles are bent (technically also acceleration, although the kinetic energy is not changed), they emit radiation within a forward-facing cone which gets more narrow the higher the particle energy is. The particle energy increases the number of photons emitted per unit time as well as their energy. The emitted radiation can be both beneficial and detrimental, depending on what the application is. Relativistic effects are responsible for the strong transverse electromagnetic fields of rapidly moving charged particles, and the high particle velocities enhance the particle interaction with magnetic fields, such as in a laser pulse or a dipole magnet.

We saw that plasmas, “soups” of free electrons and positive ions, can under certain conditions support the propagation of electromagnetic waves, and that intense laser pulses and dense particle bunches can generate charge-density waves in these plasmas. These waves follow in the wake of the pulse or bunch, and are also referred to as the wakes. The electromagnetic fields in the wakes, the wakefields, can be very strong compared to for example radio-frequency fields, and highly non-linear wakes have excellent properties for accelerating negatively charged particles.

In the next Chapter, we will apply many of these concepts to electron acceleration in RF structures, which is currently the most common way of accelerating particles to high energy, and in Chapter 3 the concepts will be used in the context of electron acceleration in plasma wakes.

RADIO-FREQUENCY-BASED ACCELERATORS

Before getting into plasma-based acceleration in Chapter 3, we will go through some important aspects of RF-based acceleration, sometimes also referred to as “conventional” acceleration by people in the plasma-accelerator community. More specifically, we will focus on linear accelerators, since circular accelerators are usually not used in conjunction with plasma-based schemes (with one notable exception, the project on proton-driven PWFA at CERN, AWAKE [77]).

2.1 Acceleration in radio-frequency waves

2.1.1 Radio-frequency cavities

As mentioned in the Introduction, RF-based accelerators are (logically) based on RF technology, where electromagnetic fields are fed into hollow, metallic cavities via waveguides and/or antennae. These fields reflect off the conducting inside walls of the cavities and build up so-called modes, each mode associated with a particular electric and magnetic field distribution and resonance frequency. For acceleration, modes with an electric field in the longitudinal direction, $\mathbf{E} = |\mathbf{E}|\hat{\mathbf{s}}$ (such as the TM_{01m} modes), are of biggest interest. However, there are uses for mode types with transverse electric fields as well, for example in beam diagnostics, see Chapter 4. The accelerating structures are usually made cylindrically symmetric, such that their TM modes can also be cylindrically symmetric, indicated by the first index 0 in TM_{01m} .

2.1.2 Different types of radio-frequency structures

Traditionally, these RF structures have been separated into two regimes - *normal-conducting* (or warm, “room-temperature”) or *superconducting* (cryogenic). Room-temperature structures are typically made out of copper, which is a very good electrical conductor at around room temperature and is therefore highly reflective for the RF-frequencies used. The most common RF bands in linear accelerators are L-band

(1-2 GHz), S-band (2-4 GHz), C-band (4-8 GHz) and X-band (8-12 GHz) [78]. The normal-conducting RF structures are usually water cooled to dissipate the excess heat from the RF power. Superconducting structures are often made out of niobium alloys, which become superconductive at temperatures below 18.3 K, and are therefore cooled by liquid helium ($\lesssim 4.2$ K) inside a cryostat [79, 80].

Superconducting RF structures are not primarily used to get higher accelerating fields or to save energy (very little power is dissipated as resistive losses in the cavity walls), but rather for other reasons: The first reason is that the very low resistive losses makes it possible to operate at high duty cycle (i.e. large “on-time”) of the cavity, or even at continuous-wave (CW), making it possible to run the accelerator at high repetition rates, up to the MHz regime [81–83]. This is considered an important feature for many FELs and future colliders. In contrast, copper cavities are usually operated in pulsed mode (low duty cycle) at a maximum repetition rate of just above 100 Hz, because cooling becomes unfeasible at higher resistive heat loads. The second reason is that the low RF losses allow the geometry to be significantly different, having a much larger internal diameter, which decreases the RF wakefields which the bunches generate inside the cavities. These wakefields (as we will discuss in the next Section), can act back on the bunches and cause potential harm (such as the beam breakup instability [84]). For high repetition rates, the wakefields from a previous bunch can linger in the cavity and affect subsequent bunches too. The third reason is that, because the resistive losses are minimal, in principle all RF power goes into the beam, which increases the total RF efficiency.

However, superconducting RF structures are difficult to manufacture and are more expensive than normal-conducting ones. Furthermore, if the superconductive state is lost, irreversible damage and melting of the structure can occur [79]. Recently, a new drive of research into cryogenic copper structures and uses thereof have been spearheaded at UCLA and SLAC [85–89]. Copper does not become superconducting, but cooling it down significantly increases the electric and thermal conductivity and material hardness, and decreases the RF breakdown rate (e.g. spontaneous arcs triggered by field emission inside the structure, destroying the RF fields). These effects allow much higher accelerating fields to be reached with a given structure, which enables a reduction of the acceleration length required to reach a certain particle energy or lowering the emittance from photoinjectors [88, 89]. It also allows for a higher RF-to-beam efficiency, as the higher shunt impedance of the RF structures enables higher beam loading (described below in Section 2.1.3) [87]. As no critical point of superconductivity must be reached, potential loss of a superconducting state is not an issue, and gains can be seen already at liquid nitrogen temperatures (77 K) [90]. Operating with liquid nitrogen is much cheaper than for example liquid helium [90]. The cryogenic normal-conducting regime is less explored than superconducting and warm normal-conducting.

2.1.3 Energy gain in a time-harmonic voltage

Regardless of the type of RF structure being used, the accelerating electric field, and voltage, U , is time-harmonic and varies in time as

$$U = U_0 \sin(kz + \omega t) = U_0 \sin(\psi_{RF}), \quad (2.1)$$

and evidently placing the particles at the correct phase, ψ_{RF} , is paramount since the energy gain varies with the voltage as (see Eq. 1.2)

$$\Delta\mathcal{E} = qU = -eU_0 \sin(\psi_{RF}) = \mathcal{E}_0 \sin(\psi_{RF}). \quad (2.2)$$

As shown in Fig. 2.1, different RF phases will yield different energy gain to the particles, and since the particle bunches have some finite time duration, the bunches will typically acquire different change in energy along the length of the bunch. Indeed, if the particles are not already bunched coming into such an accelerating voltage, it is evident that they will be afterwards, since some of the particles will be accelerated in the opposite direction.

This change in energy over the length/duration of the bunch is, in analogy to its photonic counterpart in Section 1.2.4, called a chirp. A positive chirp has lower energies located towards the front (smaller z and t), and vice versa for a negative chirp, see Fig. 2.1. Figure 2.2a shows the *longitudinal phase space* (LPS) distribution of a beam at the same RF phase as in Fig. 2.1. The LPS corresponds to the transverse phase spaces shown Section 1.1.4 (see Fig. 1.4), but for the $z - \delta$ plane. Note that the longitudinal spatial coordinate, z , is in some sources defined in the opposite direction, with larger values at the head and smaller at the tail.

RF wakefields

While they have not been studied in detail as part of the work of this thesis, the effects of wakefields in RF structures, within this work referred to as *RF wakefields* to avoid any ambiguities with respect to plasma wakefields, have been considered during for example linac tracking simulations. Therefore, a brief treatment is given here. More comprehensive introductions can be found in for example Ref. [91] and [92]. In analogy with beam-driven plasma-wakefields, the RF wakefields follow in the wake of a particle bunch travelling through an RF structure. The discussion in this section is




Figure 2.1: Energy gain in a time-harmonic voltage. Illustration of the energy gain of a bunch (green) of finite length in a sinusoidal RF voltage (black). The bunch is centered at -30 degrees ($-\pi/6$) off-crest. The finite length imparts a correlated energy spread (chirp) along the bunch. The dashed, grey lines show the polynomial expansions of the voltage around $\psi_{RF} = 0$ and $\pi/2$, dominated by a linear term and quadratic term, respectively. Remember that ψ_{RF} depends on time and space as $\psi_{RF} = kz + \omega t$.

concerning *short-range wakefields*, which act within the bunch, unlike the *long-range wakefields* which linger in the RF structures for a longer time and potentially act upon subsequent bunches, as mentioned in Section 2.1.2. The focus will be on the longitudinal component of the wakefields.

As an electron bunch traverses an RF cavity, it crosses a discontinuity of the conductive, cylindrical boundary where the beam pipe radius suddenly changes. While the bunch fields themselves are nearly transverse (see Section 1.1.2), transitioning the boundary excites also longitudinal field components as the bunch fields are reflected off of the boundary walls. Apart from the bunch exciting various higher-order modes (long-range wakefields) in the cavity, causing the bunch to lose some energy as a whole, the effects of short-range wakefields change over the bunch length. Particles toward the tail of the bunch will see a larger decelerating field, as fields from particles before them add up, causing a decrease in particle energy towards the tail. The short-range wakefields therefore affect the chirp of the bunch, increasing or decreasing it depending on which side of the RF voltage crest the bunch is accelerated on. The effect on the bunch chirp is important for e.g. bunch compression (see Section 2.3.1 below) and was accounted for during the particle tracking simulations for **Paper IV** and **V**.

Beam loading

The concept of *beam loading* [22] has not been actively considered for RF-based accelerators during the work this thesis, but is worth a brief mention here as its counterpart in plasma-based acceleration (see Section 3.3.1) has been, as it is considered an imperative effect for that class of accelerator.

When a bunch is accelerated by the electric field in an RF cavity, it gains energy, which must inevitably come from the accelerating field. As a result, if the bunch current is high enough, a significant amount of energy can be extracted from the accelerating field. The situation is often modeled with an electric parallel resonant circuit, with an external voltage generator providing power, and the cavity modeled by a parallel inductance, capacitance and shunt resistance with the bunch current passing through it [22]. If the power source parameters, cavity parameters and bunch current are matched, an optimal load situation can be found such that no RF power is reflected from the cavity, and the maximum amount is transferred to the beam. The matching of parameters also means that the bunch current should not be too high, which could add energy to the accelerating fields instead. Optimal beam loading increases the energy efficiency of the accelerator, and can reach many tens of percent.

2.2 Electron sources

The electrons must be sourced from somewhere, and for most high-brightness electron linear accelerators, or *linacs* for short, that source is a *photocathode electron gun* (PG). In such a device, a laser with a photon energy higher than the work function of the electrons in the cathode is used to liberate electrons near the cathode's surface via single-photon ionization, i.e. the photoelectric effect. The photocathode is placed inside an RF cavity, and with proper timing between the laser and the RF phase, the liberated electrons are accelerated away from the cathode.

Although the accelerating electric field is usually on the order of 100 MV/m at the cathode, a considerable accelerating distance is required to make the electrons

relativistic enough that space-charge (Coulomb) forces can be neglected (see Section 1.1.2). The low electron energy means that they will initially repel each other, and that if left alone the bunch will blow itself up. Furthermore, the Coulomb force is non-linear with the distance between particles, see Section 1.1.2, which can degrade the emittance of the beam. The latter problem can be mitigated by, for example, making the bunch charge density uniform [8], see also Eq. 1.34, which makes the repulsive force linear inside the charge distribution. To handle the repulsion itself, a method called *emittance compensation* has been developed [93]. In analogy with the matching described in Section 1.1.4, where the *emittance dominated* beams' divergence could be matched by an external focusing force, so can *space-charge dominated* beams' repulsion at low particle energy be controlled with proper focusing. By allowing the beam to first expand and then using a radially focusing solenoid magnet to focus the low-energy beam into a linac section immediately downstream, the emittance can be preserved to a large extent. Accelerating the beam further transitions it from being space-charge dominated to being emittance dominated. The PG and the first accelerating linac are typically denoted together as the photoinjector.

2.3 Beam control

After the beam has exited the injector, the beam energy is higher and space charge forces are less significant than in the gun. However, this does not mean that the beam can propagate indefinitely without expanding - the finite emittance of a realistic beam means that there is always some un-correlation between the particles' position and angle. Because of the finite emittance, the ensemble divergence cannot be perfectly compensated, and the emittance will act as a kind of "pressure" to make the beam diverge. Therefore, throughout any accelerator, transverse electron beam focusing is required. This, as most of the beam control in general in "conventional" accelerators, is typically done using transverse magnetic fields, as mentioned in Section 1.1.1. By utilizing the RF-induced $z - \delta$ correlations in the longitudinal phase space, these transverse magnetic fields can also be used to manipulate the longitudinal charge distribution, i.e. the current profile, which we will examine first.

2.3.1 Longitudinal properties

For applications such as FELs, linear colliders, and, of course, PWFAs, high beam current is often desirable. However, because of space charge, short bunches cannot usually be created directly in a PG without severe beam degradation, and so one needs to compress the bunches longitudinally. This is normally done using magnetic *bunch compressors* (although it is also possible to use ballistic compression, also called velocity bunching [94], which will not be discussed further here). Different types of magnetic bunch compressors exist, two of which will be examined below, but they all rely on a rather simple principle which can be understood from Eq. 1.7 and Fig. 1.1: when passing through a homogeneous magnetic field, such as inside a dipole magnet, electrons of different energy will take different paths and exit with different angles. By combining sets of dipole magnets (and other magnets, which we will examine in Section 2.3.2), these differences in path and angle can be made to add up (or cancel). This can make the electrons in the head and tail move closer to the position of the

reference energy, which is usually taken to be the mean energy, thus shortening the bunch.

In the following, we will rely on the notation in Eq. 1.22 and 1.23. A particle with an initial longitudinal position z_0 and a relative energy deviation δ which passes through a compressor with a first-order longitudinal dispersion, or *momentum compaction*, R_{56} , will end up at a final longitudinal position of

$$z_f = z_0 + \Delta z = z_0 + R_{56}\delta + \mathcal{O}(\delta^2), \quad (2.3)$$

where Δz is the longitudinal displacement after the compressor relative to its initial position in the co-moving frame $z = ct - s$. This assumes that other terms, such as R_{51} and R_{52} , etc., are negligible, which is usually the case in magnetic bunch compressors. R_{56} arises inside the dipole magnets, and can be calculated as [26, 95, 96]

$$R_{56} = \int \frac{R_{16}}{R} ds, \quad (2.4)$$

where R_{16} is the first-order horizontal dispersion, see Eq. 1.23, and R is the nominal bending radius of the dipole magnet. A bunch which has a linear chirp (linear distribution in $z - \delta$ space) of the correct sign, can be compressed by this transformation.

However, as is evident from Eq. 2.2 and Fig. 2.1, unless we place the bunch at $\psi_{RF} = \nu\pi$, where ν is an integer, the chirp will not be linear. This can be seen also by Taylor expanding Eq. 2.2 around $\psi_{RF} = 0$ and $\psi_{RF} = \pi/2$, yielding

$$\begin{aligned} \Delta\mathcal{E}(0) &= \mathcal{E}_0\psi_{RF} - \frac{1}{6}\mathcal{E}_0(\psi_{RF})^3 + \dots \\ \Delta\mathcal{E}\left(\frac{\pi}{2}\right) &= \mathcal{E}_0 - \frac{1}{2}\mathcal{E}_0\left(\psi_{RF} - \frac{\pi}{2}\right)^2 + \dots \end{aligned} \quad (2.5)$$

At the voltage zero-crossing ($\psi_{RF} = 0$), we do not have any gain of the average energy, which makes it useless for acceleration purposes. Therefore, one must accept some amount of non-linear chirp and rather figure out how to manage it. This is done differently for different compression schemes, two of which are explained below. Near the voltage peak ($\psi_{RF} = \pi/2$), the energy gain has a constant term and a quadratic term. Acceleration is often done relatively close to, but not exactly on, the peak voltage to maximize the energy gain and still allow for compression, which introduces all orders of non-linearities. At 30 degrees before the crest, i.e. $\psi_{RF} = \pi/2 - \pi/6 = \pi/3$, an expansion of Eq. 2.2 yields

$$2\Delta\mathcal{E}\left(\frac{\pi}{3}\right) = \sqrt{3}\mathcal{E}_0 + \mathcal{E}_0\left(\psi_{RF} - \frac{\pi}{3}\right) - \frac{\sqrt{3}}{2}\mathcal{E}_0\left(\psi_{RF} - \frac{\pi}{3}\right)^2 - \frac{1}{6}\mathcal{E}_0\left(\psi_{RF} - \frac{\pi}{3}\right)^3 \dots \quad (2.6)$$

At the other side of the crest, on the falling slope, the pre-factors of even orders (0, 2, ...) of ψ_{RF} remain unchanged, while for odd orders (1, 3, ...) they change sign.

For a perfectly linear compressor, compressing a beam with a non-linear chirp as in Eq. 2.6 would impose a ‘‘banana’’-shape in the LPS after compression, see Fig. 2.2b, which is often undesirable. However, just as we have a first-order term, R_{56} , describing the linear compression, we can extend Eq. 2.3 and make an equivalent to Eq. 1.23 for longitudinal dispersion,

$$z_f = z_0 + \Delta z = z_0 + R_{56}\delta + T_{566}\delta^2 + U_{5666}\delta^3 + \mathcal{O}(\delta^4), \quad (2.7)$$

where T_{566} and U_{5666} are the second- and third-order longitudinal dispersion, respectively. Normally, higher order terms than this are not treated analytically. These higher-order terms can be used to manipulate the LPS distribution of the bunch after compression, and used to shape its current profile. Figure 2.2c and d show the removal of the second-order curvature by tuning T_{566} , and tuning the current for PWFA using T_{566} and U_{5666} , respectively. T_{566} can be calculated as [95, 96]

$$T_{566} = \int \left(\frac{R_{16}^2}{R^2} + \frac{R_{26}^2}{2} + \frac{T_{166}}{R} \right) ds, \quad (2.8)$$

where R_{26} is the first-order horizontal angular dispersion and T_{166} is the second-

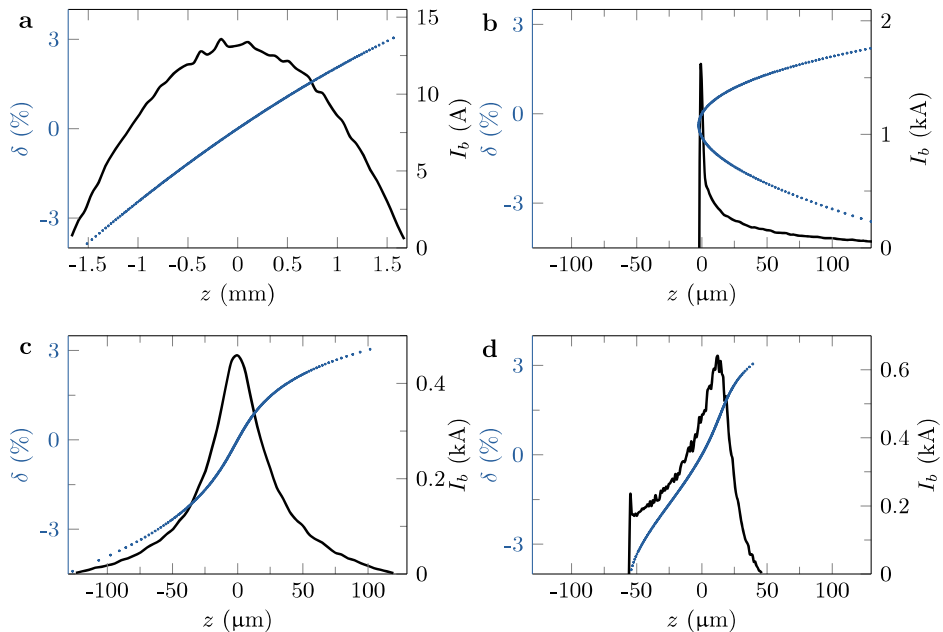


Figure 2.2: Longitudinal phase space distributions after compression. The effects on the LPS distribution under compression in a hypothetical compressor calculated using Eq. 2.7 for three different scenarios. **a**, Initial LPS distribution (blue dots) and the current (projection onto the z -axis, black) for a randomly generated parabolic current distribution of 6 ps fwhm placed at -30 degrees ($-\pi/6$ rad) off-crest, as in Fig. 2.1. The distribution consists of 10^5 particles, of which every 500^{th} is shown, totalling 100 pC. **b**, The LPS distribution after passage through a compressor with $R_{56} = -43$ mm and $T_{566} = U_{5666} = 0$ mm. The small curvature in **a** has been exacerbated. **c**, The same compressor has been tuned to $T_{566} = -165$ mm, canceling the second-order curvature and clearly showing the third-order residual from the RF voltage. **d**, $T_{566} = -190$ and $U_{5666} = -1400$ mm. The LPS distribution is more linear and the peak current has increased as a result of bringing the “wings” from **c** in, but some curvature remains: because the energy spread is large, there is some fourth-order chirp visible. However, the current profile has here been tuned to yield something close to the optimum for driving a wake, albeit with a relatively low peak current. Experimentally obtained current profiles similar to that in **d** are presented in **Paper III**.

order horizontal spatial dispersion. It is evident that magnetic bunch compressors intrinsically connect the transverse and longitudinal phase spaces.

Chicanes

The bunch compression principle is shown schematically in Fig. 2.3 for a *chicane* bunch compressor, which is probably the most common compressor type. As shown in Fig. 2.3, high-energy electrons take a shorter path through the chicane compared to the low-energy ones since their trajectories are bent less in the dipole magnets. With $z = 0$ and $\delta = 0$ defined as the bunch center, we need a positive Δz for particles at the head ($z < 0$) and negative Δz for particles at the tail ($z > 0$) to compress the bunch. Following Eq. 2.3, we define R_{56} to be negative for a chicane. In terms of Eq. 2.3, to decrease the magnitude of z_f using a system with $R_{56} < 0$, we need $z_0 < 0$ for $\delta < 0$ and vice versa, which is defined as a positive chirp. Comparing to Section 2.1.3, we can see that we therefore need to accelerate the beam on the rising slope of the RF voltage (so that the head arrives at a lower voltage) if we want to be able to compress it. Since the chirp non-linearity is predominantly of second degree, it would make sense to use T_{566} to negate this by imposing the opposite, second-degree curvature, as was done in Fig. 2.2c. However, T_{566} is not free to choose arbitrarily in a 4-dipole chicane: it is $T_{566} \approx -3/2R_{56}$ [97], i.e. it has the opposite sign of R_{56} and is thus always positive.

It is not necessarily intuitive that the opposite sign of R_{56} and T_{566} poses a problem, and it cannot be directly seen from the above equations, but can be easily shown. Starting from Eq. 2.6, we examine the spatial coordinate at a single moment in time and set $\psi_{RF} - \pi/3 = kz_0$, and recognize that $\langle \mathcal{E} \rangle \approx \mathcal{E}_0$, such that

$$\delta = \frac{\mathcal{E} - \langle \mathcal{E} \rangle}{\langle \mathcal{E} \rangle} = h_1 z_0 - h_2 z_0^2 - h_3 z_0^3 \dots, \quad (2.9)$$

where $h_n > 0$ are numerical constants, often denoted as the n^{th} -order chirp, depending on k^n . Inserting this δ in Eq. 2.7, and keeping only terms up to z_0^3 , we get that

$$\Delta z = R_{56} h_1 z_0 + (T_{566} h_1^2 - R_{56} h_2) z_0^2 + (U_{5666} h_1^3 - 2T_{566} h_1 h_2 - R_{56} h_3) z_0^3. \quad (2.10)$$

Since $T_{566} = -3/2R_{56}$ in a chicane, the z_0^2 -term cannot be directly canceled. Therefore, something about the setup must be changed to get a linear LPS distribution after compression.

In chicane-based linacs, this “something” is typically another RF system that operates at the 3rd or 4th harmonic, which modifies the RF curvature and linearizes the chirp when relative amplitudes and phases are matched. These harmonic cavities are usually placed before the first bunch compressor (two sets of compressors per linac is not uncommon) where the particle energy is lower, and so the cavity length (i.e. total voltage) can be decreased for the same chirp modulation. For example, in an S-band linac at around $f_{RF} \approx 3$ GHz, which has been a rather common frequency range for the past 2-3 decades or so, the harmonic cavity might be in the X-band. This means that one also needs another specific RF infrastructure, such as power amplifiers and waveguides, which infers a significant increase in cost over just having one set of RF equipment. While much research on X-band components have been done in the last decade, e.g. for the CLIC collaboration [11], making it more efficient and potentially

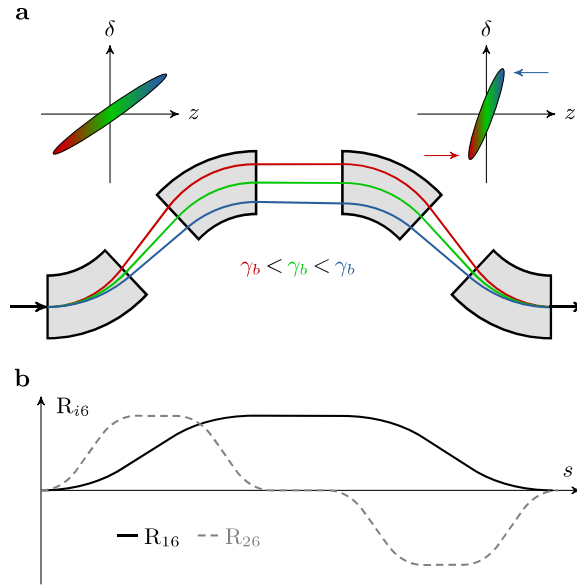


Figure 2.3: The principle of bunch compression in a chicane. **a**, Schematic trajectories of particles of different energies as they pass through a chicane bunch compressors. Higher energies (blue) take a shorter path and lower energies (red) take a longer path compared to the nominal energy (green). Schematic longitudinal phase spaces at the beginning and end of the compressor are also shown, where the z -motion of the off-momentum particles is illustrated. **b**, Schematic plots of R_{16} and R_{26} through the chicane, c.f. Eq. 1.22 and 1.23.

more affordable, designs using normal-conducting RF are also looking towards increasing the fundamental RF frequency (and with it higher accelerating voltages and smaller structures) already at the injector to C-band or X-band [87, 89, 98], pushing the harmonic frequency well into the 10s of GHz. These high frequencies become more sensitive to timing jitter, as a jitter of the same magnitude in time means a larger shift in phase at a higher frequency, and amplifiers are expensive and have lower efficiencies. However, there are other linearization techniques which do not require harmonic cavities, which could be more practical and cost-effective to use.

Double-achromats

Such a linearization technique is used in for example double-achromat bunch compressors, which currently only exists in the S-shaped format at MAX IV (detailed in Section 2.4 below), and which are treated in **Paper I** and **II**. These compressors have $R_{56} > 0$, which means that the higher energies take a longer, rather than a shorter, path. This is also the case for the somewhat similar FODO- and dogleg-compressors. This, in turn, means that they are operated with electrons accelerated on the opposite side of the RF voltage crest compared to chicanes, i.e. the falling slope, such that the bunch chirp is negative. T_{566} is positive also for these compressors. Doing the same

exercise as for the chicane, and placing the bunch at $+\pi/6$ off-crest, yields

$$\Delta z = -R_{56}h_1z_0 + (T_{566}h_1^2 - R_{56}h_2)z_0^2 + (-U_{5666}h_1^3 + 2T_{566}h_1h_2 + R_{56}h_3)z_0^3. \quad (2.11)$$

Since now R_{56} and T_{566} have the same sign, these compressors can be designed (and tuned during operation) to cancel the second-degree chirp in the longitudinal phase space, as in Fig. 2.2c. This way of removing the non-linear chirp is called *optical linearization*. Furthermore, T_{566} can be tuned using sextupole magnets (see Section 2.3.2 below), which do not affect R_{56} , such that these terms can be individually adjusted during design and operation. A potential downside of the original design [99] is that R_{56} is constant and the degree of compression can only be tuned using ψ_{RF} , which can introduce a large energy spread in the beam and limit the maximum energy gain. However, a possible solution to this is offered in **Paper II**, where additional dipole magnets of the opposite sign are introduced to give tuneability also of R_{56} .

Because the double-achromats are optically linearizing, they do not require a harmonic cavity, but nothing prevents such a device from being used for additional flexibility in tuning the longitudinal phase space. The increased number of magnets compared to a chicane is a smaller cost compared to the harmonic RF system. Since people are increasingly looking towards higher frequency linacs, this compressor type might possibly be adapted to such systems to reduce the cost significantly, but no research has been done to this end yet, to the author's knowledge.

One downside for some use-cases (while an upside for others) is that, unlike chicanes, the current, dogleg-like format of these compressors gives a lateral shift of the beam propagation axis. However, as shown in for example Ref. [100], two such systems can be put back-to-back, and return the beam to its original axis. Such a system has the same overall shape as a chicane, although the dimensions differ, but the functionality remains like that of achromats. Such systems might, in fact, be beneficial in another way, which has also not been studied yet to the author's knowledge - namely by reducing coherent synchrotron radiation kicks in the compressors. Using two back-to-back chicanes with opposite bending directions, arranged in a "zig-zag" manner, has been shown to reduce detrimental effects of coherent synchrotron radiation [101].

2.3.2 Transverse properties

In a conventional accelerator, accelerating electric fields and confining magnetic fields are usually separated, with accelerating structures interleaved with focusing magnets. However, after the injector, the electron energy is typically too high for solenoid magnets to be practical, and therefore quadrupole magnets are usually used for beam focusing. A quadrupole magnet, with its 4 poles and specific pole shape [22], gives a transverse magnetic field, and thus a transverse force, that changes linearly with the distance from the center of the magnet. In an iron-yoke magnet (unlike superconducting magnets [102] and Halbach arrays [103, 104]), the magnetic field distribution is given only by the shape of the pole tips. The strength of the field is in electromagnets tuned by changing the current through the magnet's coils, while geometrical alterations are necessary for permanent magnet-based setups. An intrinsic downside of quadrupole magnets, regardless of type, is that while they are focusing along one transverse direction, they are defocusing along the other - this is a fundamental limitation. Because of this limitation, they must be placed at least two and two, such

that the combined effect is focusing. This is a natural reason to choose a Cartesian coordinate system for beam transport, as alluded to in Chapter 1.

There also exists a category of so-called higher-order magnets. While a quadrupole magnet has 4-fold symmetry, the first higher-order magnet, the *sextupole*, has 6-fold symmetry, and *octupoles* have 8-fold symmetry, etc. A magnet of a given order, for example a quadrupole (order 1), will affect all the matrix elements of that order and higher, i.e. of first (R), second (T) and third order (U), etc. A sextupole affects T-terms and up, etc. These magnet types are often used for corrections of chromatic aberrations arising in the quadrupoles [22, 105, 106]. Because of their impact on higher-order terms, they can also be used to tune bunch compressors without affecting the overall compression strength, c.f. Eq. 2.4, 2.8 and 2.7. The common trait of higher-order magnets is that their fields do not increase linearly with the distance from the axis, as in a quadrupole magnet. The transverse magnetic field, and thus the transverse force, in sextupoles increases as x^2 and in octupoles as x^3 . While these non-linear gradients are precisely what allow for these magnets to be used for optics correction, they can also introduce aberrations of their own, even in a beam with zero energy spread. Because their gradients are non-linear, these magnets can degrade the beam emittance via *geometric* aberrations (see Section 1.1.4) - beams which pass through a non-linear focusing object will receive a non-linear angular kick, and if the transverse beam dimensions are too large, this non-linear kick will have significant effects on the emittance. However, it is possible to amend this by arranging the magnets in the proper way, for example by placing a second, equally strong magnet of the same type in such a way that the second magnet imparts a kick of the same magnitude but opposite sign further down the line from the first magnet [107].

Because of realistic geometries and manufacturing imperfections, any real quadrupole magnet will not have a perfectly linear gradient, but have small contributions of higher-order components, which is called the *multipole content* [102] of the magnet. This is also true for the higher-order magnets. Iron-yoke and superconducting magnets can be made very precisely, and typically have higher-quality fields than for example Halbach-type permanent magnet quadrupoles (PMQs). However, it is possible to make the PMQs very small compared to other magnet types, and Halbach PMQs can have gradients of up to 500 T/m [108] (compared to $\lesssim 100$ T/m for iron-yoke magnets and low 100's for superconducting magnets [102]). Since permanent magnets are passive, they do not need to be fed power or cooling water and can therefore easily be placed in vacuum. Because of this, they have received a deal of attention from the LWFA community, but as will be discussed in Section 3.3.2, there are options which are likely better for that specific application.

Tilted bunches

Another point of interest is to control the transverse spatial and angular *slice centroids* of the beam. In this context, the slice centroids are the “center of mass” of the spatial transverse coordinates and their angles, in thin sections along the beam. If the beam is not well aligned with its propagation axis, such *beam tilts* can reduce the output power and efficiency of FELs [109], but can also give rise to beam hosing in a plasma, as discussed in Section 1.3.4. While the FEL efficiency might be decreased, beam tilts can render a PWFA completely useless. There are two main sources of such beam tilts, which we will cover here.

In Section 2.3.1 the higher-order longitudinal dispersion, and its connection to the transverse dispersion, were discussed, see for example Eq. 2.4 and 2.8. From there, it is evident that the higher-order transverse terms influence the longitudinal compression directly. However, while R_{56} and T_{566} should be non-zero after a compressor (otherwise no compression and linacization can be achieved), the transverse terms should typically be cancelled. Otherwise, the result would be a beam with some sort of tilt coming out of the compressor. An example of this is shown with dark crosses in Fig. 2.4 - the relationship between δ and z is linear here, so this is essentially a banana flying through space. Removing this type of *leaking dispersion* is the specific topic of **Paper I**, and also a large part of the studies in **Paper II** and **IV**.

However, leaking dispersion is not the only source of beam tilt - another main source is *coherent synchrotron radiation* (CSR) [110, 111]. While dispersion might typically be a problem in the beginning of the machine, where the bunch energy spread is larger, CSR often becomes an issue towards the end, when the bunches are compressed more. CSR is emitted in the dipoles of for example bunch compressors or beam transfer lines, and consists of wavelengths which are longer than twice the bunch length, $\lambda_{CSR} \gtrsim 2c\sigma_t$, such that the bunch electrons emit in unison. Since the radiation is emitted within a cone and the electrons travel on an arc, the radiation from the tail of the bunch irradiates the head. Because the bunch has some momentum in the direction of the radiated electric field, there can be a net transfer of energy between radiation and bunch. This change in energy along the bunch alters the local bending radius, such that different parts of the bunch gain a non-linear kick in angle. The induced angular spread increases the projected emittance, but also produces a non-linear tilt in the beam which will propagate through the system. An example is shown by the light blue circles in Fig. 2.4. CSR is unavoidable in magnetic compression schemes, since the particles are being accelerated laterally, but there are some ways to design systems to minimize the effects.




Figure 2.4: Beam tilt/slice offset examples. Examples of the two major sources of beam tilt for two independent tracking simulations - CSR (light blue circles) and leaking dispersion (dark blue crosses), with the red line showing the offset prediction from Eq. 1.23 using T_{166} and U_{5666} for the system in question. The centroid position was calculated for each slice, and δ has a linear relation to z . The CSR-induced offset is in this case predominantly third-order. This is not always the case, and varies with degree of compression, bunch charge, optics, etc.

As an example of the impact of CSR, the projected normalized emittance increase in a dipole with a short, Gaussian bunch can be given as [111]

$$\Delta\varepsilon_n = 7.5 \times 10^{-3} \frac{\beta_x}{\gamma_b} \left(\frac{N_b r_e L_D^2}{R^{5/3} \sigma_z^{4/3}} \right)^2, \quad (2.12)$$

where L_D is the length of the dipole magnet. Equation 2.12 shows the main scaling factors of CSR-induced emittance growth. There is also a transverse CSR effect which can yield emittance degradation of a similar order [111]. When designing a dispersive system, it is thus an advantage to keep the β function small and the dipole short and of large radius. Due to the non-linear scaling, it can be advantageous to use multiple dipoles which each of them are shorter and with a smaller bending angle (larger radius). This is assuming that a certain current, $I_b \sim N_b/\sigma_z$, at a certain energy is required. If σ_z is constant through propagation in the system, through for example a small chirp and/or R_{56} , the beam optics can be designed such that the CSR kicks virtually cancel each other [112]. If σ_z changes, optics design alone might not be enough, and a second compressor, with a different R_{56} , opposite bending direction and separated by the first by an optimized betatron phase advance, can be employed [101]. Yet another, complementary method is given in **Paper I**, where the non-linear transverse dispersion is de-tuned from its nominal values near zero, to partially counteract the CSR-induced beam offsets. This technique was used to tune the bunches for the study in **Paper IV**. A similar technique was developed earlier, in Ref. [109].

2.4 The MAX IV linear accelerator

As an example of these concepts in practice, we will examine the MAX IV linear accelerator [113] - the subject of study presented in **Paper III, IV** and **V**. A schematic overview of the MAX IV linac is shown in Fig. 2.5. The linac is a warm electron accelerator which serves two purposes, both as a full-energy injector for the 1.5 and 3.0 GeV storage rings [114] and for the Short-Pulse Facility (SPF) [115], located downstream of the injection point for the 3 GeV ring. There is also an ongoing investigation into using the linac as a driver for a soft X-ray FEL [116].

The linac has two electron sources - one thermionic electron gun (TG) used for ring injection, and one PG used mainly for SPF delivery [117]. Together with a solenoid magnet and the first linear accelerating structure, the guns make up the injector. The laser driving the photoinjector [118] is a frequency-tripled Ti:Sapphire system at 263 nm wavelength, corresponding to a photon energy high enough to overcome the




Figure 2.5: Schematic of the MAX IV linac. Yellow denotes radiofrequency (RF) structures, red are quadrupole magnets, blue are dipole magnets, green are sextupole magnets and gray are special magnets and extraction points for ring injection. PG and TG denote photocathode and thermionic RF guns, respectively, L is the linac structures and BC1/2 are the first/second bunch compressor. SPF denotes the short-pulse facility.

work function of the copper cathode. To achieve frequency-tripling, the laser pulses must be rather short, which is not always desirable for PG operation. Therefore, the tripled, sub-100 fs laser pulses are stretched in a double passed prism stretcher and then stacked using α -BBO crystals to a nominal length of 3 or 6 ps fwhm. α -BBO crystals are birefringent, which is used to separate the polarization components of the pulses and delay them with respect to each other, creating duplicate pulses with perpendicular polarization. Proper configuration produces several pulse-pairs which are delayed just enough to make an almost flat-top temporal profile, which is called *pulse stacking*. There are also plans for other types of longitudinal pulse shaping, using a grating-based stretcher without the need for pulse stacking [119].

The 39 linear accelerating structures, like the guns, operate at 2.9985 GHz (in the S-band), with a unit length of 5.2 m and an average accelerating gradient of about 18 MV/m. The power is fed to these structures in pulses of 4 μ s fwhm duration, generated by one klystron (power amplifier) and modulator (frequency generator), via a SLED cavity (RF pulse compressor) and RF power divider. The phases of the RF systems for the different linac sections (L0 + gun, L1 and L2-L19, cf. Fig. 2.5) are set for each group separately, each pair of linac structures can be individually tuned by (comparatively slow) mechanical phase shifters in the RF waveguide system.

The SPF is based around producing short, spontaneous, undulator X-ray pulses, requiring short electron bunches. Before delivery to the SPF, the bunches are therefore compressed longitudinally to yield shorter, higher current bunches than what is used for ring injection. In the MAX IV linac, this is done using double-achromat bunch compressors [99], and not chicanes, see Section 2.3.1 above. This means that no harmonic linearizing cavity is required. The nominal bunch parameters for SPF delivery are a duration of 100 fs, an emittance of <1 mm mrad and a charge of 100 pC. The facility layout of the SPF is such that there are three parallel slots to which the linac can supply the electron beam, see Fig. 2.6. The only currently populated slot is SP02, which houses the FemtoMAX beamline [120]. Previous simulations indicate that it is possible to compress the bunches to well below 100 fs fwhm and keep the emittance low [99]. The simulated bunch durations in the study of **Paper IV** were 60-200 fs fwhm, while in the study of **Paper V** they were 10-20 fs fwhm. The experimental measurements shown in **Paper III** were a few hundred fs fwhm, but only the first bunch compressor was used for that study. During work parallel to the presented study, single bunches were compressed to 30-40 fs using only the first bunch compressor.




Figure 2.6: Schematic of the SPF. Schematic of the SPF, with the 3 possible beamlines, of which FemtoMAX is the only populated to date. U denotes undulator. Color conventions are the same as for Fig. 2.5. Labels marked with asterisks (*) are proposed for the PWEA project, outlined in **Paper IV**. The blue field marks the diagnostics beamline which is currently under construction.

PLASMA-BASED ACCELERATORS

In Section 1.3.3 and 1.3.4 we discussed how plasma wakes can be generated, and in this Chapter we will examine how to actually use these structures for accelerating particles. What are the advantages and disadvantages? The focus will mostly be on the non-linear bubble regime, but some comparisons with the linear regime will also be made.

Central in plasma-based accelerators is that there are generally at least two components to keep track of, the *driver* and the *witness*. The driver, be it a laser pulse or particle bunch, drives the wake within which the witness is accelerated. In LWFA, the accelerated electrons always belong to a witness bunch, since the driver is a laser pulse. Therefore, the term “witness” is often not used, as there is no ambiguity. For many previous PWFAs experiments, such as the one described in Ref. [121], the tail of the driver is accelerated, with no distinct witness bunch. However, to reach high efficiencies and a higher degree of control over the final bunch parameters, a distinct witness bunch is useful, as will be discussed in Section 3.3.1. Within this thesis, the witness bunch signifies the bunch being accelerated, regardless of driver type.

3.1 Acceleration in plasma waves

The radial electric field in a linear (laser) wake is given in Eq. 1.72, and varies as $E_\rho(\rho, z) \propto \rho \exp[-2\rho^2/W^2] \sin(k_p z)$. This means that the focusing gradient is approximately linear only for small-diameter electron beams, and that it changes in strength depending on the distance behind driver. A bunch of some finite length (assumed to not drive its own wake) would then see different focusing fields along its axis. These are not good conditions for particle acceleration, and emittance growth can occur over short distances.

One of the clear advantages of acceleration in the bubble regime is that the focusing fields have linear gradients which do not depend on the position behind the driver. The radial force on an electron inside a bare ion cavity is $F_\rho = -m_e \omega_p^2 \rho / 2$ [55]. This focusing force comes from the background ions, which have the same charge density as that of the background electrons (but of the opposite sign), together with the radial plasma current, i.e. the radial movement of the sheath electrons [55]. To put this force

on the same format as for a focusing magnet, to be able to use the theory we built in Section 1.1.3 and 1.1.4, we compare to the focusing force in Eq. 1.14, $F_x = -ecg_q x$. By setting $\rho = x$ and equating these two forces, we obtain the quadrupole-equivalent focusing gradient for a plasma channel,

$$g_q = \frac{n_p e}{2c\epsilon_0} = \frac{m_e c}{e} \gamma_b \kappa. \quad (3.1)$$

For a plasma density of $n_p = 1 \times 10^{17} \text{ cm}^{-3}$, $g_q = 3 \text{ MT/m}$, roughly 5 orders of magnitude higher than an electromagnetic quadrupole and 4 orders of magnitude higher than a permanent magnet quadrupole (PMQ). We can also immediately obtain the betatron wavenumber, as $k_\beta = \sqrt{\kappa}$, given that the energy is known. For 1 GeV electrons, $\lambda_\beta = 2\pi/k_\beta \approx 6.6 \text{ mm}$ at the above density. These conditions are excellent and in principle allow for long acceleration distances without focusing-induced emittance growth, but the strong fields incur some difficulties, as will be shown later in this Chapter. Firstly, we shall examine some other effects which do limit the acceleration length.

3.1.1 Limitations to the acceleration length

Dephasing

Unlike in RF accelerators, the particles within plasma-based accelerators move along with the ion cavities as they accelerate, but it is still important to control the accelerating phase of the bunches. For LWFA, one major limitation to this control is the fundamental limits on the laser group velocity which were introduced in Section 1.3.2, see Eq. 1.65. To recapitulate - the higher the plasma density, the lower the group velocity of the laser pulse, and thus the lower the phase velocity of the plasma wave. As mentioned in Section 1.3.2, For $n_p = 5 \times 10^{18} \text{ cm}^{-3}$ and $\lambda = 800 \text{ nm}$ we get $v_g = c(1 - \omega_p^2/\omega^2)^{1/2} \approx 0.9986c$. While this seems like “almost” the speed of light, a 100 MeV electron moves at $v \approx 0.999987c$, which is even closer to c . This small difference is significant, even over short distances. Because the electron quickly becomes faster than the plasma wave, it begins to move forward relative to the laser driver, and its phase in the wave changes. Eventually, the electron enters a decelerating phase (such as the front half of the bubble in the bubble regime), and begins to lose energy - this is called *dephasing*. The distance that the electron can be accelerated before this happens is called the *dephasing length*, L_D . In the 1D, linear regime, the dephasing length is [51]

$$L_D^{1D} = \frac{\lambda_p}{2} \frac{1}{v - v_g} \approx \frac{\lambda_p}{2} \frac{1}{c - v_g} \approx \left(\frac{\omega_0}{\omega_p}\right)^2 \lambda_p. \quad (3.2)$$

For $n_p = 5 \times 10^{18} \text{ cm}^{-3}$ and $\lambda = 800 \text{ nm}$, $L_D^{1D} \approx 5.2 \text{ mm}$. In the bubble regime, there is a stronger dependence on a_0 , and the dephasing length becomes [56]

$$L_D^{3D} = \frac{R_B}{1 - \beta_p} \approx \frac{2}{3} \left(\frac{\omega_0}{\omega_p}\right)^2 R_B = \frac{4}{3} \left(\frac{\omega_0}{\omega_p}\right)^2 \frac{\sqrt{a_0}}{k_p} = \frac{2}{3\pi} \left(\frac{\omega_0}{\omega_p}\right)^2 \lambda_p \sqrt{a_0}. \quad (3.3)$$

For the plasma density and laser wavelength above, $a_0 = 2 \rightarrow L_D^{3D} \approx 1.6 \text{ mm}$. Evidently, there is a large difference between these two regimes. A 1.6 mm acceleration distance does not sound like much, but considering the magnitude of the average

accelerating field before dephasing [56], $E_{s,avg} \approx \sqrt{a_0} E_{wb}/2 \approx 152$ GV/m (see Eq. 1.76), one would still expect maximum electron energies of roughly

$$\Delta\mathcal{E} = eE_{s,avg}L_D^{3D} = \frac{2e}{3} \left(\frac{\omega_0}{\omega_p}\right)^2 \frac{a_0 E_{wb}}{k_p} = \frac{2m_e c^2}{3} \left(\frac{\omega_0}{\omega_p}\right)^2 a_0 \lesssim 240 \text{ MeV}. \quad (3.4)$$

240 MeV is near the maximum electron energies of 225–235 MeV that were reached over approximately 1.6 mm during the study presented in **Paper VI**. Casting Eq. 3.4 in practical quantities yields [56]

$$\Delta\mathcal{E} [\text{GeV}] = 1.7 \left(\frac{P_0 [\text{TW}]}{100}\right)^{1/3} \left(\frac{0.8}{\lambda_0 [\mu\text{m}]}\right)^{4/3} \left(\frac{1}{n_p [10^{18}\text{cm}^{-3}]}\right)^{2/3}. \quad (3.5)$$

This expression shows that to reach higher electron energies, it is advantageous to use a higher-power laser, with a shorter wavelength and at lower plasma densities. Available laser technologies typically fix a feasible wavelength (800 nm Ti:Sapphire is the workhorse technology), but going towards higher laser power and lower densities, with a near-matched spot size, has indeed been the strategy for achieving the current world record of about 8 GeV from a 20 cm long capillary discharge waveguide at $n_p \approx 3 \times 10^{17} \text{ cm}^{-3}$ using 850 TW laser pulses at 800 nm [122]. At these (for LWFA) low plasma densities, self-focusing is significantly less pronounced, and additional schemes, such as the capillary discharge waveguide, must be implemented to enable proper guiding.

Recently, several groups have proposed schemes for circumventing the dephasing limit [123–126]. While the above treatment is valid for Gaussian pulses, the new concepts proposed in Ref. [123–125] are based on introducing spatio-temporal coupling in the pulse front by using special optics [124, 125] or using the overlap between two obliquely incident laser pulses [123] to tune the group velocity of the laser wake. These concepts promise accessing higher plasma densities, and therefore higher accelerating fields, while circumventing the lower laser group velocities normally associated with higher plasma densities (see Eq. 1.65). These novel techniques also promise a minimal reduction in intensity over an enhanced length compared to Gaussian pulses. In the coming years, it is possible that new record energies in single-stage LWFA might be achieved using advanced pulse-tuning techniques. However, because of the relative ease of generating Gaussian or near-Gaussian laser pulses, and good knowledge of how they interact with plasmas, they will likely persist for the foreseeable future.

As mentioned in Section 1.3.4, the velocity of a particle driver does not depend on the plasma density. A notable effect of this is that for electron bunch drivers, where $\beta_b \lesssim 1$, the driver and witness bunches usually have similar velocities, so the dephasing length is very long. Other effects, such as depletion or diffraction, limit the maximum attainable energy gain before dephasing becomes a concern. The smaller impact of dephasing is one of the major benefits of PWFA.

Matching the dephasing and laser depletion lengths

The dephasing length is not the only constraint on the acceleration distance. As we saw in Section 1.3.3, the laser driver will deplete after propagating for one depletion length (see Eq. 1.79). Matching the dephasing and depletion lengths allows reaching

the maximum possible electron energy while making use of the largest possible fraction of laser energy, i.e. maximizing the efficiency. This matching yields

$$L_D^{3D} = \frac{4}{3} \left(\frac{\omega_0}{\omega_p} \right)^2 \frac{\sqrt{a_0}}{k_p} = \left(\frac{\omega_0}{\omega_p} \right)^2 c\tau_\ell = L_{PD}^{NL} \leftrightarrow \tau_{\ell,m} = \frac{4}{3} \frac{\sqrt{a_0}}{ck_p} = \frac{2}{3} \frac{R_B}{c}, \quad (3.6)$$

where $\tau_{\ell,m}$ is the matched laser pulse fwhm duration with respect to dephasing and depletion. This length is within the range $c\tau_\ell \lesssim R_B$ that was assumed for matched blowout conditions in Section 1.3.3.

Staging

Regardless of driver type and whether the driver is perfectly matched with respect to dephasing and diffraction, it will eventually run out of energy and the acceleration will stop. For this reason, *staging*, i.e. concatenating multiple accelerating stages, is necessary if one wants to reach higher witness electron energies. An accelerating stage should not be run to full depletion, because the wakefields would then weaken towards the end and the witness might end up in a decelerating or defocusing phase, and might even start to drive a wake of its own. Therefore, whatever is left of the previous driver, which can contain a significant amount of energy, must be safely discarded before entry to the next stage, and the witness must also be properly transported between the two. This will be discussed further in Section 3.3 below.

3.1.2 Betatron radiation

One effect of great interest for plasma-based accelerators in general, and laser-plasma accelerators in particular, is *betatron radiation* [127, 128]. As was shown in Section 1.1.3 and 1.1.4, particles in a constantly focusing field will perform transverse oscillations - *betatron oscillations* - under a common envelope which might be constant. Because the focusing fields in ion cavities are so strong, any particle which finds itself off-axis will be rapidly pushed back towards the axis. In fact, because the beam has a finite emittance, the majority of the particles have a non-zero transverse position and/or angle, such that effectively all the particles will oscillate to some extent. As was discussed in Section 1.1.2, relativistic charged particles undergoing transverse acceleration radiate strongly in the forward direction, with the radiation from betatron oscillations being named betatron radiation. Particularly in high-density LWFA, these oscillations are often exploited specifically for the generation of betatron radiation, which is typically in the soft X-ray regime. The principle is shown in Fig. 3.1. As a side note, radiation from focusing beams is also generated in for example quadrupole magnets in conventional accelerators: in the final focus system of particle colliders (i.e. before the particle collision point, or *interaction point*), where the beam energy is very high, the beam size is large and the magnets are at peak strength, the focusing forces are strong enough to generate X-rays which must be shielded.

As shown in Section 1.1.2, see Fig. 1.3, each electron radiates within a cone of opening half-angle $\theta_0 = 1/\gamma_b$. As outlined in the Supplementary Information of **Paper VI**, the angular distribution of the radiation can be approximated with a Gaussian of rms radius $\sigma_\theta = 1/(\Upsilon\gamma_b)$, where $2.2 \lesssim \Upsilon \lesssim 3.7$ is a numerical factor depending on observation direction and observed part of the radiation spectrum. When oscillating back and forth on a sinusoidal path in the bubble (c.f. Eq. 1.19), the radiation




Figure 3.1: The principle of betatron radiation. Schematic showing the principle of betatron radiation in a LWFA: An electron which is injected off-axis finds itself in a strong, radially focusing electric field which causes it to perform betatron oscillations as it accelerates. Since the acceleration is the strongest near the trajectory peaks, that is where the radiation will preferentially be emitted. The oscillation amplitude is exaggerated for demonstration purposes. If the electrons interact with the laser field inside the bubble, their oscillation amplitude can become enhanced in the laser polarization direction.

cone also sweeps back and forth, somewhat like the headlights of a car, or a search light, making alternating left and right turns. As evident from Eq. 1.12, the radiated power is maximum when the transverse acceleration is the greatest, at the peaks of the oscillation trajectories. More transverse acceleration means more radiation, but also that the angles of the electron trajectories increase, and thus the total angle spanned over by the X-ray “search light”. Assuming Gaussian electron and X-ray beams allow adding the divergences in quadrature, such that the effective divergence can be estimated as [129]

$$\Theta_x^2 = \sigma_{x'}^2 + \sigma_\theta^2, \quad (3.7)$$

where we assume oscillations along the x -axis. Corresponding expressions can be found for y .

The amplitude of the angular excursion can be quantified using the K -parameter, as $K = \gamma_b x'_0$ [128]. The K -parameter is also used, and defined in the same way, in “conventional” light sources, where it is used to quantify the strength of the radiation-generating *insertion devices* called *undulators* and *wigglers* [8, 22]. Therefore, the parameter is also called the undulator or wiggler parameter. The distinction is made that an undulator has $K \ll 1$ and a wiggler has $K > 1$. Because of spatial coherence effects, the photons from an undulator have an rms divergence of $\sigma_\theta \approx \sqrt{\lambda/L_U}$, where λ is the radiation wavelength and L_U is the undulator length [29], so $\Theta \approx \sigma_\theta$ when assuming small $\sigma_{x'}$ and K . With the fundamental radiated wavelength from an undulator being $\lambda \approx \lambda_U/(2\gamma_b^2)$, the effective divergence becomes $\Theta \approx \sigma_\theta \approx 1/(\gamma_b \sqrt{2N_U})$, where $N_U = L_U/\lambda_U$ is the number of undulator periods. Using the format from the paragraph above, we see that $\Upsilon = \sqrt{2N_U}$ for this case. In a wiggler, the transverse oscillations are much larger, because of the larger K , and the spatial coherence decreases. This gives an effective divergence of $\Theta \approx K/\gamma_b$. Laser-plasma

accelerator-based betatron sources are usually operated in the range $10 \lesssim K \lesssim 20$, and γ_b is a few hundred, such that the X-ray rms divergence is often on the order of a few tens of mrad. This is very large compared to the few μrad divergences produced by undulators ($\gamma_b \gtrsim 1000$ and $N_U \approx 100$), and makes the betatron radiation difficult to focus and transport to applications. However, some applications (see for example **Paper VII**, **VIII** and **IX**) can make good use of the divergent beams. K can also be related directly to the plasma and electron parameters through [128]

$$K = \frac{1}{c} \omega_\beta x_\beta \gamma_b = \frac{1}{\sqrt{2}c} \omega_p x_\beta \sqrt{\gamma_b}, \quad (3.8)$$

where x_β is the peak betatron oscillation amplitude of the horizontal electron motion. Again, a corresponding expression can be found for y .¹

The sinusoidal oscillations have considerable effects on the spectrum. The piecewise circular trajectory causes the spectrum to be synchrotron-like with a critical energy \mathcal{E}_c , see Eq. 1.13. The critical energy can be related to the plasma and electron parameters using Eq. 3.8 as [128]

$$\mathcal{E}_c = \hbar\omega_c = 3\pi c\hbar \frac{K\gamma_b^2}{\lambda_\beta} = \frac{3}{2}\hbar\omega_\beta K\gamma_b^2 = \frac{3\hbar}{4c} \omega_p^2 \gamma_b^2 x_\beta = \frac{3}{4} \hbar\omega_p k_p \gamma_b^2 x_\beta. \quad (3.9)$$

The critical energy from betatron sources is often in the keV photon energy range. Finally, using Eq. 1.12, the radiated power can be shown to be [127, 128]

$$P_{\text{rad}} = \frac{\pi e^2 c}{2\epsilon_0} \frac{\gamma_b^2 K^2}{\lambda_\beta^2} = \frac{r_e m_e}{3c} \gamma_b^4 \omega_\beta^4 x_\beta^2 = \frac{4\pi^2 r_e^3 m_e c^3}{3} \gamma_b^2 n_p^2 x_\beta^2. \quad (3.10)$$

During the work for this thesis, betatron radiation was used for different applications, as described in **Paper VII**, **VIII** and **IX**. Generation of betatron radiation was also studied and described in **Paper X**. **Paper VI** also concerns radiation generation with plasma-based accelerators, but in a new regime, somewhat different from the betatron regime, and is discussed further in Section 3.3.2 below. Apart from being useful for a diverse range of X-ray applications, betatron radiation can also be used as an electron beam diagnostic. This will be covered in Section 4.2.2.

Direct laser acceleration

An effect which can influence the electron beam inside the bubble is *direct laser acceleration* (DLA) [130, 131], which occurs when the electron bunch is overlapped with part of the laser pulse. The effect is similar to inverse free-electron laser, where an oscillating external electric field delivers energy to likewise oscillating electrons if their oscillation phases are properly matched. Since the laser field is usually linearly polarized, the component of the betatron oscillation which is along the laser polarization can become enhanced, and the electrons can gain (or lose, if the phase relation is opposite) energy. The increased oscillation amplitude and energy is easily exploited for producing more and higher-energy betatron X-rays (see Eq. 3.9 and 3.10), but the X-ray divergence is also increased along the laser polarization, making the beam profile elliptical.

¹The spatial transverse coordinate is sometimes given as the radial coordinate $\rho_\beta^2 = x_\beta^2 + y_\beta^2$ [128].

3.2 Electron sources

The electron source for plasma-based accelerators can be either an external source, for example a short bunch from an RF accelerator or another plasma-based accelerator, or an internal source, i.e. the plasma itself. In the former case, one typically talks about *external injection*, and in the latter *internal injection*, and relating to both is the phenomenon of *trapping*. Trapping occurs when the electrons have the right position, propagation direction and momentum to be trapped inside the wake and subsequently accelerated. Obviously, the electrons must be placed spatially and temporally inside the wake, but they must also be made to travel with the correct velocity with respect to the wake phase velocity. These constraints are clearly also present in RF accelerators, but are there easier to control and relatively less emphasis is put on it in the literature. External injection is conceptually rather simple - you have a bunch and you need to shoot it into the wake - but can be very difficult to execute properly. External injection is discussed in Section 3.3. For internal injection, several different methods exist - these are conceptually more complex than external injection, but are less demanding to implement experimentally, and we will pay some extra attention to those who have been used during the work for this thesis. External and internal injection can be used with both laser- and beam-driven schemes, but external injection is by definition required for staging. Either technique requires a plasma source, which will be discussed next.

Plasma sources

As discussed in Section 1.3.1, the plasma precursors used in the work for this thesis were gases, most commonly (nitrogen-doped) helium, but hydrogen was also used. In LWFA, the gases are typically ionized by the rising edge of the laser pulses, while in PWFA, it is common to pre-ionize the gas using an electric discharge or a laser pulse [77, 132, 133], although field-ionization using the driver itself is sometimes possible [133–136]. To control the injection and acceleration of electrons, the plasma must be somehow confined, with a length and electron density appropriate for the driver. To this end, two classes of sources are employed: *gas jets* and *gas cells*, see Fig. 3.2a and b, respectively.




Figure 3.2: Schematic drawing of two types of plasma sources. **a**, A gas jet, with the gas flowing unhindered, at high velocity directly into the vacuum chamber. The example shows an LWFA, with the gas marked in yellow, with darker color meaning higher density, laser in red, electrons in green and X-rays in orange. **b**, A gas cell, with a slower gas flow filling a confined volume (X-rays not depicted).

Gas jets are allowed to flow unrestricted into the surrounding vacuum chamber, and are often produced using high-pressure nozzles with a nozzle opening geometry such that the emerging gas flow is supersonic, so called *de Laval nozzles* [137, 138]. The supersonic character allows the gas flow to remain relatively collimated within a few mm from the nozzle opening, with only slow spatial changes to jet density and width when moving away from the nozzle. In contrast, subsonic gas flows, such as those from simple tubes, expand rapidly in the lateral direction relative to the gas flow, causing the jet density and width to change rapidly in the direction away from the nozzle. The gas jet nozzle geometries are fixed, and the plasma parameters are varied by regulating the gas backing pressure and changing the distance between driver propagation axis and nozzle opening.

Unlike gas jets, gas cells confine the gas to some volume with an opening in each end along the beam propagation axis to allow driver and witness to enter and exit. The restrictive geometry can allow for a significantly longer plasma in the beam propagation direction, while keeping the transverse direction narrow. Whereas gas jet nozzles have fixed geometries, some gas cells have variable length. Electric discharge systems can sometimes be integrated into gas cells, which allow both ionization of the gas and generation of laser-guiding structures.

Gas cells are used for both laser- and beam-driven acceleration, whereas gas jets are usually only used with laser drivers. Both classes of gas source exist in pulsed and constant-flow varieties: the repetition rate of the experiment and the pump capacity of the vacuum pumps determine the operational mode. Because of the lack of obstructive material, gas jets, and the plasma structures therein, are easier to characterize than for gas cells [139]. Gas jets were used for the studies in **Paper VI**, **IX** and **X**, while gas cells were used for the studies in **Paper VII** and **VIII**.

3.2.1 Internal injection mechanisms

In internal injection schemes, electrons from the background plasma are placed inside the wake by different mechanisms. These different mechanisms will typically yield different bunch characteristics - such as charge, spectrum and duration - and the application might dictate the preferred method. Below, the most common injection methods used for this thesis will be listed, together with what type of beams they tend to produce.

Self-injection

Perhaps the most fundamental internal injection method, although possibly not the most conceptually simple, is *self-injection* [140–143]. This mechanism relies only on the interaction between a uniform plasma and the laser pulse, and was the mechanism used for the landmark experiments in 2004 where quasi-monoenergetic electron spectra from LWFA were first observed [144–146]. There are two varieties of self-injection: longitudinal and transverse. Longitudinal injection can occur when the laser aspect ratio, $W/(c\tau_\ell)$, is large, i.e. the pulse is much wider than it is long. This makes the wake take on a more 1D-nonlinear shape, and forces the plasma electrons to move across the peak of the laser pulse to a greater extent. Because the pulse is wide, the transverse ponderomotive force is relatively weak near the peak, and thus the passing electrons mainly gain some longitudinal momentum, such that their velocity

can become comparable to that of the wake, which promotes trapping. Such bunches can be relatively reliable and of good quality, but low charge and energy [147].

In contrast, transverse self-injection only happens when self-focusing and self-compression has increased the intensity of the laser pulse significantly, such that $a_0 \gtrsim 4$. This intense and radially smaller pulse now causes the plasma electrons to gain significant transverse momentum, such that some of them, when returning towards the laser axis, do not follow the bubble sheath but obtain enough longitudinal momentum to enter the wake and get trapped. The self-evolution of the laser pulse can also cause the bubble to expand, which triggers injection in a similar way. The process where part of the wave collapses in on itself in this way is usually called *wavebreaking*, in an analogy with when for example water waves under the right circumstances become peaked (i.e. non-linear), start to roll in on themselves and break. The transverse ponderomotive force responsible for the transverse momentum is very sensitive to fluctuations in the laser parameters, such that injection is often transversely asymmetric and varies from shot to shot [147]. The bunches can be of high charge, narrow spectrum and short duration, but are not always reproducible. During the experimental LWFA work for this thesis, self-injection was usually not desired, and measures, such as lowering the plasma density to limit the laser evolution, were taken to suppress it. Experiments today are often going toward better beam quality and reliability. To this end, other injection schemes, where the injection is controlled by other means, tend to be used. A collection of the ones relevant for this thesis is discussed briefly below.

In PWFA, the only significant bunch self-modulation is transverse, as (part of) the bunch focuses inside the wake. As discussed in Section 1.3.4, the wake excitation amplitude saturates once the driver reaches $n_b \approx 10n_p$, after which the wake does not evolve significantly. Also, the large longitudinal component of the electron velocities that arises in intense laser fields, as seen in Fig. 1.9b, is small for most PWFA-cases. The wake phase velocity is also higher for particle beam drivers than for laser drivers, and PWFA usually operates at lower plasma densities providing lower accelerating fields than in LWFA, which decreases the probability that electrons reach the required forward velocity to enable trapping, compared to LWFA. For these reasons, other methods of internal injection (discussed below) are typically relied upon in PWFA schemes, as the bunch requirements for self-injection are too high. With the high-current bunches projected at the FACET-II facility [133, 135], injection during wake lengthening caused by transverse driver evolution, similar to the laser case above, is theoretically possible [148], but such parameters are not achievable for most PWFA experiments and has thus far not been verified experimentally.

Density down-ramp injection

One method for controlling the wavebreaking is to use a plasma density transition. The transition occurs between a region of higher density and lower density along the driver propagation direction. Propagating from higher to lower density makes the bubble size increase by changing the phase velocity of the rear of the wake and decreasing the restoring force on the sheath electrons. The process is depicted in Fig. 3.3. Two regimes are usually considered, depending on if the transition length, L_T , is longer than the initial plasma wavelength, $L_T \gg \lambda_p$, or not. If $L_T \gg \lambda_p$, the injection method is usually referred to as *density down-ramp injection* [149–153], and

if $L_T \lesssim \lambda_p$, it is called *shock-front injection* [154–156], but the distinction is not always made. Density down-ramp injection is here, for simplicity, chosen as a general denotation of this class of injection methods using plasma density transitions.

In general, the electrons are injected by transverse wavebreaking: electrons initially moving as part of the bubble sheath can find themselves located inside the bubble because the phase velocity of the rear portion of the bubble (behind the driver) decreases. This happens because as the density decreases, so does the plasma frequency, so the response of the plasma is slower, which causes the bubble to “elongate”. For e.g. a beam driver, the driver velocity is constant, such that a transition from a small bubble to a large bubble can only happen if the rear portion slows down. For a laser driver, the relative increase in group velocity between the two regions might only be on the per-mille level, which is not enough to offset the slowing of the wake response. Because of the lower phase velocity at the rear of the bubble, the sheath electrons passing inside it can gain enough energy that their velocity becomes higher than the local phase velocity, such that they begin to follow the driver, and are hence trapped.

Density down-ramp injection can be implemented for both laser and beam drivers, although it is easier to implement in LWFA since wavebreaking is generally easier to achieve than in PWFA to begin with, as discussed above. The bunches produced can have normalized (projected and slice) emittances below 0.5 mm mrad and charges in the few-100 pC range [152, 153]. The duration depends to a large extent on the difference between the bubble sizes in the two regions, and so varies with several parameters (see Section 1.3.3 and 1.3.4). The electron spectra can be peaked and sometimes with few-percent fwhm energy spread. Density down-ramp injection was observed, as a side effect, in the study shown in **Paper VI**, but shock-front injection was used in combination with ionization injection (see below) in the same study as the primary injection technique.

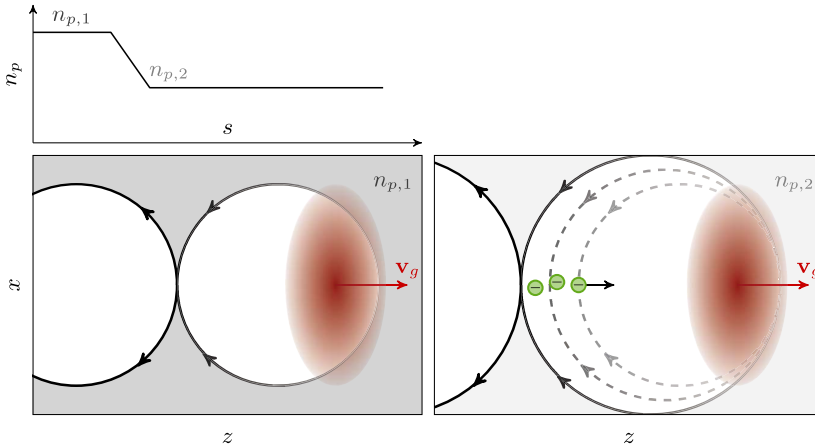


Figure 3.3: Schematic depiction of density down-ramp injection. As the laser pulse propagates from the higher density, $n_{p,1}$, to the lower density, $n_{p,2}$, the bubble expands, which can trigger injection and trapping of electrons.

Ionization injection

In *ionization injection* [157, 158], the different ionization energies of atoms shown in Fig. 1.8 are utilized. The fully ionized H or He plasma (common in LWFA) is here doped with a small amount, usually around 1% or so, of an element with more tightly bound electrons, for example nitrogen. Nitrogen atoms have seven electrons each, of which the five first are significantly easier to ionize than the last two, see Fig. 1.8. These first five electrons are liberated in the rising slope of the laser pulse, together with the electrons from the carrier gas molecules, making them part of the ambient plasma. However, the N^{5+} ions remain stationary, and when the laser pulse peak reaches $a_0 \geq 2.20$ and $a_0 \geq 2.75$, it begins to ionize the 6th and 7th state, respectively, as it passes over the ions. This releases some fraction of the bound electrons around the peak of the laser pulse, inside the first bubble, which then begin to drift backward with respect to the laser. If the difference between the wake potential (see for example Eq. 1.49) at the ionization location, φ_{ion} , and the potential minimum in the bubble, φ_{min} , is such that $\varphi_{\text{min}} - \varphi_{\text{ion}} \geq 1$, i.e. the potential energy difference is larger than the electron rest energy, the electron gains enough forward momentum to be caught in the wake. The process is depicted in Fig. 3.4.

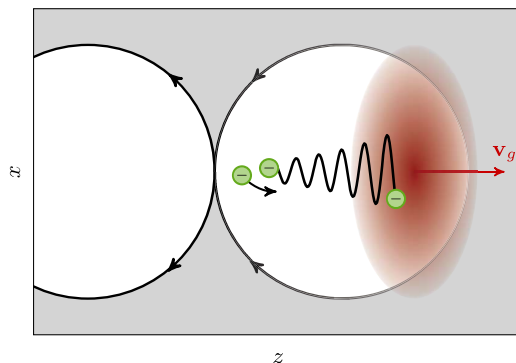


Figure 3.4: Schematic depiction of ionization injection. As the laser pulse propagates over the not fully ionized ions, some still bound electrons can be ejected near the laser peak. While first drifting back with respect to the laser pulse, they can gain enough momentum to become trapped if the wake potential difference between ionization point and the rear of the bubble is large enough.

Ionization injection can produce stable and reproducible bunches of high charge, but often with a broad, continuous spectrum since injection typically occurs continuously over the acceleration length. However, in certain variations of the scheme, injection can be relatively confined, which produces bunches with more peaked spectra and estimated bunch currents of several 10s of kA [159]. The bunch emittances are relatively low, but not as low as for example shock-front injected bunches [160]. Because of the inherent overlap with the laser, DLA effects in the electron beam are pronounced, so because of the experimentally simple implementation, high charge and enhanced electron oscillation amplitudes [131], ionization injection is favorable to use for laser-driven betatron X-ray generation. Ionization injection was used in the studies described in **Paper VII**, **VIII**, **IX** and **X**. The study in **Paper VI** also employed ion-

ization injection, but it was there combined with shock-front injection in *shock-assisted ionization injection* [161].

3.2.2 Hybrid plasma-based accelerators

Since extremely high-current bunches can be produced by LWFA, one alternative is to use these incredibly dense, but relatively high-emittance, bunches for driving a PWFA immediately after the LWFA, in a scheme called hybrid L-PWFA [67, 162]. Such a scheme could then be used for example as a “brightness transformer”, where the high-density, lower-brightness bunch drives a high-amplitude PWFA in a high-density plasma, wherein any PWFA-appropriate injection scheme could be used to generate secondary electron bunches of very high brightness, all within a cm-scale setup. Since PWFA is nearly dephasing-free, much longer acceleration distances could be achieved than with a similar LWFA setup, and thus higher energies could be reached, possibly in the multi-GeV range using current laser technology. As mentioned in Section 3.1.1, guiding laser pulses over many cm to reach multi-GeV electrons can be a significant challenge.

As a side note, an injection scheme such as *plasma photocathode injection* [163, 164] might be easier to realize in hybrid L-PWFA conditions, because part of the main LWFA-driving laser pulse can be split off and independently focused onto the PWFA stage, similar to what is often done in *colliding-pulse injection* [165–167]. These injection methods were not considered during the work for this thesis, but are interesting to notice because they, like a photocathode injector in an RF accelerator, decouple the generation of the accelerating field from the generation of the electron bunch. The decoupling provides additional tuneability with respect to for example injected charge, which can be adjusted independently of the wake excitation. Because the injection pulse in the suggested configuration above originates from the driver pulse, and the total path length difference after splitting is typically short (~ 1 m), the relative timing between the two is intrinsically as stable as it could ever get ($\ll 1$ fs).

Such hybrid, dual-stage schemes could also be used for photon generation, where the first stage is the injector/accelerator and the second stage is the radiator. This allows for efficient and tunable radiation generation, where the focus of the first stage is to produce a high-brightness electron beam and the second can be used to tune the properties of the radiation to a much greater extent than when both parts are performed in the same stage, as for betatron radiation. The experimental complexity only increases by a small margin. Such schemes have been theoretically proposed for γ -ray production, using radiator stages of high density immediately following the lower-density acceleration stage [168, 169]. Incidentally, this geometry is similar to a density down-ramp injection scheme, but in reverse. The idea is to generate a very dense, high-energy electron beam by LWFA in the first stage, which is then sent into the high-density second stage where it (possibly together with the original laser pulse) drives a highly mismatched PWFA, where the high-amplitude transverse oscillations are induced. The accelerator and radiator can also be separated longitudinally by a few mm, as was presented in **Paper VI**. In this work it was found that the radiator can behave akin to a plasma-lens, and will be detailed more below in Section 3.3.2. The radiated photon energy is lower than for the case in Ref. [168] and [169], but the opening angle of the radiation cone can be reduced significantly. In fact, the discovery of this generation mechanism was not intentional: the original goal of that experiment

was to realize a hybrid L-PWFA scheme, but while the acceleration scheme largely did not succeed, other interesting physics were learned.

3.3 Beam control

Just as in RF accelerators, the electron bunches in and from plasma-based accelerators need to be properly controlled and transported if they are to be used for some application such as staging. Unlike in RF accelerators, the accelerating structures, i.e. the wake, provides considerable focusing forces, and so one does not have to interleave accelerating and focusing elements in the same way. However, the acceleration stages are still of finite length, because the driver eventually runs out of energy or diffracts too much to drive the wake, which means that several stages might be required to reach the desired final energy. In the case of a particle collider, the desired collision energy can be > 1 TeV [11]. Depending on the driver type and plasma density, a large number of accelerating stages might be required, and a new driver must then be provided for each. Specifically because the focusing inside the wake is so strong, the usual methods for focusing the beam using quadrupole magnets (see Section 2.3) can be inadequate, depending on the electron beam parameters and the demands on it after focusing. This has led to complementary techniques for electron beam control, which will be detailed below. However, just as in Chapter 2, before we dive into the transverse effects, we will examine the longitudinal effects.

3.3.1 Longitudinal properties

One of the main attractions of plasma-based accelerators is that they can deliver intrinsically short bunches of high quality. Because of the strong focusing and accelerating fields, high charges of 100s of pC can typically be injected without space-charge ruining the beam emittance, yielding very high peak currents of several tens of kA [159]. Therefore, to use the beam for e.g. an FEL, no additional compression would (in general) be required, unlike in an RF accelerator. However, even ultra-short bunches can become strongly chirped, with large energy spreads, because of the large longitudinal electric field gradient (see Section 1.3.3 and 1.3.4), which is not desirable for most applications, let alone FELs. To this end, depending on the sign of the chirp, the bunch can be de-chirped in an RF-based or plasma-based [170] dechirper, or the chirp can be exploited for post-compression, yielding even higher peak currents, or flipped (by over-compression) to cancel the chirp in a second acceleration stage [171]. Energy spread control is one of the main challenges of plasma-based accelerators, together with transverse matching (see Section 3.3.2 below), and large research efforts are invested in solving or circumventing these issues.

Beam loading

Although it also occurs in RF accelerators (see Section 2.1.3), beam loading is a more prominent, and possibly also more important, effect in plasma-based accelerators. Closely related to optimizing the beam-driver current profile for efficient wake excitation (see Section 1.3.4), beam loading using the witness bunch aims to flatten the longitudinal wakefield gradient along the witness bunch to minimize the resulting bunch energy spread. Optimal beam loading also allows the witness bunch to extract

the maximum amount of energy from the wake, such that the efficiency increases. In fact, the efficiency (from wake to witness) can, at least theoretically, approach 100% in the bubble regime [15, 68]. Phenomenologically, loading a linear wake is relatively intuitive: by placing a witness bunch of the appropriate current profile at an appropriate phase, it will excite a linear wake of its own which simply adds to the existing wake, making it possible to locally flatten the longitudinal wakefield (c.f. Fig. 1.10). However, as beam loading in linear wakes has severe limitations [68], we will focus on non-linear wakes also in this matter.

In the rear portion of the bubble, the electrons in the bubble sheath are moving back towards the axis, gaining significant transverse momentum as they return from the peak of their trajectory. By placing a witness bunch with the appropriate current profile in the back of the bubble at an appropriate phase, the space-charge fields from the bunch can push back on the returning sheath electrons and decrease their transverse momentum. Perfect beam loading thus means that these sheath electrons have their original momentum perfectly canceled by the bunch fields, such that all the wake energy ends up in the bunch and there are no further wake periods (c.f. Fig. 1.10 and 1.11). Because the sheath electrons slow down as they are returning to the axis, the wake, assumed to move with constant velocity, will outrun them slightly, which gives the appearance that the bubble is “stretched” or “tapered” backwards. Putting too little charge in the witness bunch (*under-loading*) will cause the sheath electrons to retain some energy, such that the wake still has some energy left and the field is not perfectly flat, and putting too much (*over-loading*) will cause the sheath electrons to be pushed outwards again, which reverses the field gradient and can cause the bunch to lose energy. Both under- and over-loading generally lead to an increase in energy spread. Beam loading influences the trapping conditions, and typically works




Figure 3.5: Beam loading in a 1D non-linear plasma wake. Plots of the bunch electron density envelope (green), plasma electron density (black) and longitudinal electric field (blue) for an electron beam driver with $\tau_b = 60$ fs, $Q_b = 40$ pC, $\sigma_p = 4.25$ μm and $n_p = 1 \times 10^{17}$ cm^{-3} ($n_b/n_p \approx 1.15$), with and without a witness beam with $Q_{b,w} = 13.5$ pC and otherwise identical parameters ($n_b/n_p \approx 0.39$). The electric field and wake density were obtained by numerically solving Eq. 1.81. The light colors correspond to no witness bunch (identical to Fig. 1.11d) and dark colors are with the witness. Here, the witness is Gaussian in time, and does not have the optimal, trapezoidal shape, but the field is visibly flattened within the fwhm, regardless. The remaining wake has a significantly lower amplitude.

to suppress further trapping when using internal injection methods.

As given in Ref. [68], the optimal current shape is a linearly down-ramped trapezoid - similar to the optimal beam-driver shape (see Section 1.3.4), but with the opposite slope. Therein, it is also shown that flat-top bunches and Gaussian bunches can load the wake well. Also illustrated is the fact that there is a trade-off between the total amount of charge that can be loaded and the accelerating field [68]:

$$\frac{Q_b}{en_p k_p^{-3}} \frac{eE_i}{m_e c \omega_p} = \frac{e^2}{\epsilon_0 (m_e c^2)^2} Q_b E_i = \frac{\pi}{16} (k_p R_B)^4, \quad (3.11)$$

where E_i is the longitudinal electric field at the head of the bunch. From Section 1.3.3 and 1.3.4, we remind that $k_p R_B = 2\sqrt{a_0}$ for laser drivers (Eq. 1.78) and $k_p R_B = 2\sqrt{\Lambda_0}$ for beam drivers (Eq. 1.83). The trade-off in Eq. 3.11 comes from the finite energy available in the wake: it is possible to either accelerate few particles to high energy or many particles to low energy, but the total energy must be conserved. An example of beam loading in the 1D non-linear regime using a Gaussian witness bunch is shown in Fig. 3.5. The loading is not exact (and the 1D model is also limited) as the field gradient is only flattened over part of the bunch, but the main results of the load are clearly visible - a flattened electric field and a subsequently lower wake amplitude. Beam loading was one of the considerations of the study presented in **Paper IV**.

Optimizing beam loading is one of the advantages of using external injection or tunable, and preferably (from this point of view) decoupled, internal injection, see Section 3.2.2. Being able to maintain the relative energy spread, or even decrease it, without losing any charge is an important feature to achieve for plasma-based accelerators. It should also be noted that the wake can be loaded by an externally injected laser pulse [172]. This is maybe not a very relevant laser amplification scheme, but such a pulse placed for example in the second wake period can serve to remove remaining energy from the wake - it will “accelerate” the laser pulse by amplifying it and thus the energy can be removed from the plasma. Both of these mechanisms could be important for running experiments at high repetition rate, where thermal effects from the deposited energy might disrupt the acceleration or even destroy the adjacent equipment, such as the capillary.

3.3.2 Transverse properties

The additional requirements of transverse control on beams from plasma-based accelerators, compared to RF accelerated beams, can be illustrated by a few key points. For one, simply sending an ultra-short (few-fs [173, 174]) and divergent (few mrad) electron bunch, from for example a high-density LWFA, into vacuum will lengthen it, even if the bunch is in principle perfectly monoenergetic ($\sigma_\delta = 0$). This bunch lengthening, unlike the longitudinal manipulations done in bunch compressors (see Section 2.3.1), is not a chromatic effect, but a geometric effect, coming from the particles with large angles (u') simply travelling a longer distance than the ones close to the propagation axis. The path length difference between the two paths is $\Delta L \approx u'^2 L_1 / (2 - u'^2)$, for small u' . After only $L_1 = 25$ cm of propagation, an electron with a $u' = 4$ mrad angle will have gained an on-axis shift corresponding to $\Delta L/c \approx 6.7$ fs, on the order of the bunch length, compared to an on-axis particle. Clearly, if one cares about preserving the short bunch duration, and thus the peak beam brightness, the outgoing beam

divergence must be minimized near the source. One way of collimating the beam immediately after the source is by using plasma-lenses, which was done in the study shown in **Paper VI** and will be covered below.

LWFA-generated electron bunches are rarely monoenergetic, and finite energy spread exacerbates the above effects and introduces additional detrimental effects, but can, as shown in Section 2.3.1, also be used to an advantage in some circumstances. For example, the bunch lengthening could possibly be (partially) reversed by using a bunch compressor. However, the combination of finite energy spread and large propagation angles can directly impact the projected normalized emittance [31] (see Eq. 1.25), such that the brightness (see Section 1.1.4) is decreased. This situation normally does not occur in conventional accelerators, since neither energy spread nor divergence are large enough in most cases for the normalized emittance to degrade. We saw above that the bunch energy spread is affected by a combination of bunch length and charge, but why are the plasma-accelerated bunches so divergent to begin with and can we do something about it? To answer this, we must return to the concept of matching.

Matching

As discussed in Section 1.1.4, matching is about correctly pairing a focusing force with the natural divergence of a beam. If the two are not properly balanced, the beam envelope will perform oscillations, as shown in Fig. 1.5. However, the term “matching” is often also used to describe the matching of beam parameters between different sections of an accelerator - conventional accelerators often have *matching sections*, typically made up of a number of quadrupoles, devoted only to this. These two concepts are more or less equivalent - both are about matching the transverse focusing with the incoming beam parameters to control how the beam behaves downstream. Because of the strong focusing in plasma channels, the matched β functions are generally small, which inevitably makes the divergences large since $\gamma_m = 1/\beta_m$ and $\alpha_m = 0$ (see Eq. 1.24). Using Eq. 1.29 and 3.1, the matched β function in a bare ion cavity can be calculated as

$$\beta_m = \frac{1}{\sqrt{\kappa}} = \frac{1}{k_\beta} = \sqrt{\frac{2m_e c^2 \epsilon_0}{e^2}} \sqrt{\frac{\gamma_b}{n_p}} \approx 7.52 \times 10^6 \text{ m}^{-1/2} \sqrt{\frac{\gamma_b}{n_p}}. \quad (3.12)$$

For example, Eq. 3.12 gives $\beta_m \approx 67 \text{ }\mu\text{m}$ for $n_p = 5 \times 10^{18} \text{ cm}^{-3}$ and $\gamma_b = 400$, and $\beta_m \approx 1.8 \text{ mm}$ for $n_p = 1 \times 10^{17} \text{ cm}^{-3}$ and $\gamma_b = 6000$. The large divergences make it more difficult to inject and extract bunches from the plasma channel using electron optics, as we shall see below. The small β functions also decrease the betatron wavelength, so detrimental effects such as emittance growth from a mismatched envelope show up after a short propagation distance. There is an intrinsic hurdle here; we want the bunches to be matched to the channel to avoid emittance degradation, but the matched parameters yield divergent beams which are difficult to control.

The emittance-deteriorating effects in a constant focusing channel, such as a bare ion cavity, come from differences in focusing force within the bunch. In a linear wake, the focusing gradient varies along a bunch of finite length (see Eq. 1.72), which causes the betatron period to vary between the head and the tail as $k_\beta = \sqrt{\kappa}$. In contrast, in a bare ion cavity, the focusing gradients are constant inside most of the cavity.

However, for a bunch with some chirp, the focusing strength will vary between the head and the tail because of the different particle energies, see Eq. 3.1 and 1.36. This type of mismatch leads to *betatron decoherence*. This occurs to some (sometimes recoverable) extent even when the beam is matched [175], but is greatly exacerbated if the beam is mismatched [176, 177], and for a mismatched beam, lowering the energy spread only serves to delay the emittance growth, and does not prevent it fully.

If a chirped bunch with finite energy spread enters the plasma with the same, matched, CS parameters along the beam, different slices, and thus different energies, will oscillate with different betatron frequencies. In transverse phase space, see Fig. 1.4, this will appear as the slice ellipses of different energies rotating apart, or shearing, moving the ensemble from an ellipse into a “butterfly” shape, and eventually into another ellipse with larger area and thus larger emittance. However, if the rms energy spread is relatively small, on the order of 1% or below, these oscillations can eventually move back into phase and overlap as an ellipse with the initial area again, restoring the emittance. Over many periods, the emittance will therefore oscillate between a minimum and maximum value. Ideally, since the matched β function varies with γ_b (see Eq. 3.12), one would want to send the beam into the plasma with a slight variation in the CS slice parameters, so as to match each slice of the beam, which in principle removes all emittance degradation. However, this would require a particular type of chromatic focusing system, and for example an apochromatic [17] system might be practically more feasible to implement.

The strong focusing in plasma channels also makes matching the beam between plasma and vacuum using standard electron-optical components non-trivial. One approach is to put a strong set of magnets near the plasma, which focuses or collimates the beam. This can work relatively well with well-behaved RF accelerated beams going into a PWFA. However, because superconducting magnets are large (and thus incompatible with a small LWFA setup) and electromagnets are weak and impractical to put in vacuum, the only choice for magnetic collimation very near the exit of a LWFA is a PMQ-based solution. As mentioned in Section 2.3, multiple quadrupole magnets must be used because they are de-focusing in one plane while focusing in the other. This makes the setup relatively long (10-20 cm with 2-3 PMQs), and typically blows up the beam size in one of the planes inside the setup, producing large chromatic aberrations in that plane. Furthermore, they are difficult to tune in strength (although efforts have been relatively successful in this direction, see e.g. Ref. [178]), and are prone to permanent radiation damage [179] from stray electrons. Therefore, using magnets is not necessarily a practical solution to this particular problem: as it turns out, *plasma-lenses*, which will be discussed below, can be much more practical in this situation.

Another way of matching between vacuum and plasma, which can be combined with other focusing elements to great effect, is to use smooth *plasma density transitions*, or *ramps* [32, 175, 177, 180, 181], which gradually (de-)magnify the beam transversely. This way, the large beam divergence that comes with a small spot size does not degrade the normalized emittance as the beam does not propagate freely in vacuum, and significantly relaxes the requirements for any optics up- or downstream of the plasma. By increasing or decreasing the focusing strength slowly (*adiabatically*), over lengths $L \gg \lambda_\beta$, one can ensure that if the beam is initially matched at the start of the ramp, it stays matched, which also minimizes the chromatic emittance growth described above. Adiabatic ramps tend to be quite long, particularly the low-density

ends extend far, but are not sensitive to the exact shape [175]. Any realistic (underdense) plasma source will have some type of gradient, on account of the natural expansion of the gas and plasma from a jet or capillary. Future plasma sources will likely seek to exploit this and provide optimized and/or tunable gradients. As mentioned in Section 1.3.4, density ramps can sometimes also be used to mitigate beam hosing.

Plasma-lenses

As alluded to above, plasma-lenses are a strong contender for electron-optical component of choice around plasma-based accelerators. Here, two different classes of plasma-lens will be discussed - *active lenses* [182, 183], which use a strong discharge current to create an azimuthal magnetic field, and *passive lenses* [184–187], which use the wakefields created by the driver, which is “passively” present anyway. Both flavours are *radially focusing*, unlike the magnetic structures.

Active plasma-lenses (APLs) use a capillary with electrodes attached to each end to contain the gas. A high-voltage pulse (few 10’s of kV over ~ 100 ns) is applied over the electrodes, using for example a Marx bank [188], causing an electric discharge current between them. The few-hundred-A current then creates a radially increasing azimuthal magnetic field with a linear gradient of $g_r = \mu_0 I_0 / (2\pi \rho_c^2)$, where ρ_c is the capillary radius, provided that the current density is uniform and that thermal effects are (made) negligible [183]. This allows such a lens to achieve gradients in the several kT/m-range, as compared to a strong PMQ gradient of 500 T/m. However, to be able to use such a lens, the beam electron density must be significantly below that of the plasma, so as not to excite a wake. In conjunction with a LWFA, such a lens is also at risk of laser-induced damage. As for example a final-focus lens in a collider or beam-capture lens in conjunction with a PWFA, the APL is possibly a better fit. As a side note, the thermal expansion of the plasma from heating by the discharge current lowers the density on-axis, which can actually be exploited to create a parabolic guiding structure for laser pulses [122, 188].

Where high-density electron bunches and/or high-intensity laser pulses are involved, passive plasma lenses might be a better option for electron focusing. These, in turn, exist in two main regimes, linear [184–186] and blow-out [187], where the latter is probably the strongest focusing lens that can be produced for a negatively charged particle. As shown in Eq. 3.1, the gradient in such a lens can be more than a factor of 1000 higher than in an APL. Therefore, to work as a thin lens (see Section 1.1.3), such a lens would have to be very short, even for a high-energy beam [187]. Over such a short length, and at the relatively low plasma densities in which such a lens would operate, only a small change in beam energy would occur. Passive plasma-lenses would, on account of their short focal lengths, be placed very near the accelerating section, and possibly paired with plasma density ramps on the accelerating section for the greatest effect. Because they originally only consist of gas, and that both electron and laser spot sizes are small so close to the accelerating stage, they are not prone to damage. In particular downstream of high-density LWFAs, where the electron beam is highly divergent, passive plasma lenses can reduce the divergence significantly only a few mm after the accelerating stage and thus decrease several of the detrimental effects discussed above.

Passive plasma lenses can also be used for radiation generation, as mentioned above in Section 3.2.2. This is the topic of **Paper VI**. The scheme described therein is somewhat akin to a hybrid L-PWFA setup, but the function of the second stage is more like a plasma-lens than an accelerator. After the bunch has been generated in the first LWFA stage, both electron bunch and laser pulse rapidly diffract in the vacuum (or at least low plasma density) between the stages, which expands their transverse sizes. As they enter the second stage, the laser and/or bunch drive a wake inside which the bunch starts to rapidly focus. When tuned properly, the bunch β function will go from divergent to convergent inside the lens, and thus reach a crest where its slope, i.e. the CS α , is zero, see Eq. 1.24. Since $\alpha = 0$, $\beta = \gamma^{-1}$, and thus the electron bunch divergence

$$\sigma_{u'} = \frac{\varepsilon_{n,u}}{\gamma_b \sigma_u}. \quad (3.13)$$

Because of the expanded beam size, the rms divergence is decreased at the crest. As shown in Eq. 1.12, the radiated power depends strongly on the transverse force, as $F = dp/dt$ (c.f. Eq. 1.4 and 1.5). Since we have a linear, quadrupole-like gradient (Eq. 3.1), we can write for the average particle $F = ec\sigma_u g_q$, c.f. Eq. 1.14. This allows us to cast Eq. 1.12 on the form

$$P_{\text{rad}} = \frac{e^3}{6\pi m_e^2 \epsilon_0 c} Q_b \gamma_b^2 g_q^2 (\sigma_x^2 + \sigma_y^2), \quad (3.14)$$

where we have also multiplied Eq. 1.12 with the number of electrons, $N_b = Q_b/e$, and added the contribution from both transverse planes, as the emission is incoherent. The case here is similar to a harmonic oscillator where the transverse acceleration is, for the average particle, maximum near the envelope crest, such that σ_u is maximized for a given gradient, which causes P_{rad} to also be maximized near the crest. That the radiated power maximum and electron bunch divergence minimum (nearly) coincide means that in this scheme, the maximum amount of photons is emitted at or near the location of smallest electron bunch divergence, and since the X-ray divergence is dependent on the bunch divergence (see Eq. 3.7), also the effective X-ray beam divergence is minimized.

3.4 The laser-wakefield accelerator at the Lund High-Power Laser Facility

The majority of the LWFA experiments during the present thesis work were carried out at the the Lund High-Power Laser Facility at Lund University. This laser is based on chirped-pulse amplification (CPA) [36, 189, 190] with Ti:Sapphire at 800 nm central wavelength. After the final amplification stage, the pulses are compressed in vacuum, yielding pulses of $\lesssim 1$ J energy and 35 fs fwhm duration after compression. These pulses are subsequently transported to the experimental chamber in vacuum via a deformable mirror (DM, see Section 4.1.3) used for wavefront correction to optimize the focal spot. A sketch of the experimental chamber is shown in Fig. 3.6a. The ~ 60 mm diameter top-hat shaped beam is focused in the experimental chamber by a 775 mm focal length ($f/13$), 14.3° off-axis parabolic mirror (OAP). A typical focal spot of around $13 \mu\text{m}$ fwhm ($W_0 \approx 11 \mu\text{m}$) is shown in Fig. 3.6b. The gas source is usually a supersonic gas jet with a high Mach-number, used for the studies presented

in **Paper VI** and **IX**, but a gas cell was used for the studies in **Paper VII** and **VIII**. The most common gas is helium, often doped with nitrogen for ionization injection.

Before the OAP, a small part, roughly 10 mJ (i.e. 1%), is picked off from the main beam using a small 1/2-inch diameter, 45-degree mirror and used for transverse shadowgraphic imaging of the laser-plasma interaction, see Fig. 3.6a and c. After the OAP, an attenuated beam can be sent out of the chamber to a Phasix SID4 wavefront sensor (see Section 4.1.3), which is operated in conjunction with the DM, and a camera imaging the focal plane (which is how Fig. 3.6b was acquired). After the LWFA, the accelerated electrons enter the dipole spectrometer (described in Section 4.2.1), and are dispersed onto a Lanex Regular scintillating screen which is imaged by an Andor Zyla 16-bit sCMOS camera. Example electron spectra for three different injection methods are shown in Fig. 3.6d. The majority of the remaining laser pulse is transmitted through the magnet gap, and blocked by a few- μm thick Aluminium foil. The generated X-rays are mostly transmitted through the thin foil and the downstream Kapton vacuum window (although some absorption does occur), before being detected on an X-ray CCD camera. When used for single-photon counting, the camera is moved further away from the source. Two X-ray beam images are shown in Fig. 3.6e and the X-ray diagnostics in general are described further in Section 4.3. Typical electron beam parameters for this setup are energies up to roughly 230 MeV, with continuous spectra (from ionization injection) down to single-percent relative energy spread (using shock-assisted ionization injection), charges up to 150 pC and divergences in the 3-8 mrad fwhm range. The generated betatron X-ray beams have $10^8 - 10^9$ photons per shot with a critical energy of 2-2.5 keV and a divergence of the order of $30 \text{ mrad} \times 15 \text{ mrad}$ fwhm.

On occasion, the X-ray diagnostics were removed and the spectrometer magnet moved aside, to allow the beam to propagate freely through the chamber. Instead, an electron transport beamline was docked, which allowed for transporting and refocusing the electrons further downstream. The electron transport beamline is presented in Section 4.2.2.

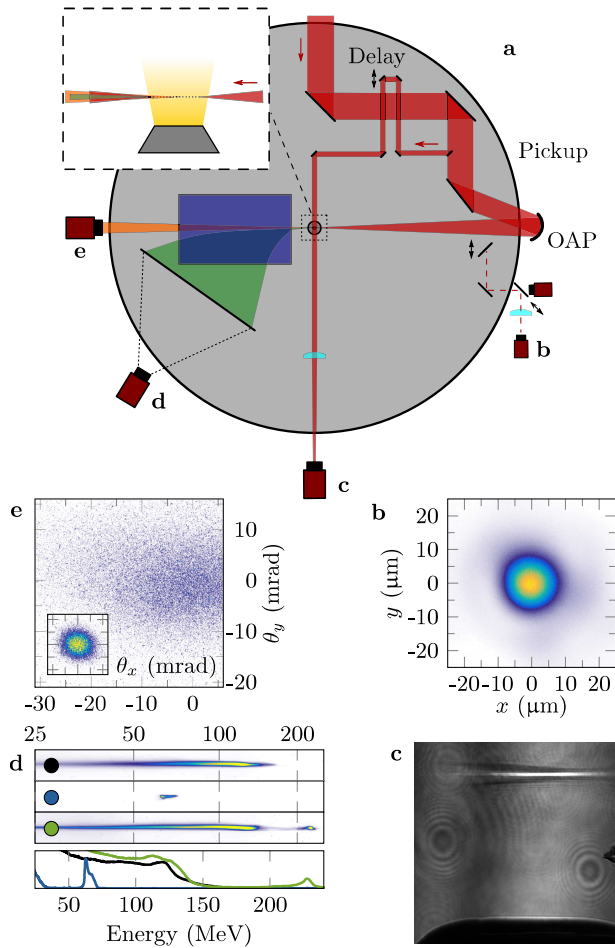


Figure 3.6: Typical LWFA setup and some diagnostics examples. **a**, Schematic of a typical, minimum LWFA setup. The laser pulse enters from the top of the figure, and is focused onto the gas cell in the center of the chamber. Before the OAP, part of the beam is picked up for transverse diagnostics (**c**), and during setup the beam can be intercepted before the focus to diagnose the focal spot (**b**). The accelerated electrons (green) pass through the spectrometer dipole magnet (blue) and onto the Lanex screen. The X-rays (orange) travel straight onto the X-ray CCD camera. **b**, Zoomed image of a good-quality focal spot (false color). **c**, Example image from the sideview (same orientation as the inset in **a**). The dark shape at the bottom is the gas nozzle, the bright, horizontal stripe is the center of the plasma channel. The bright, vertical fringe is from diffraction in the probe pulse. **d**, Three example electron spectrometer images (false color, reversed with respect to the orientation in **a**) using ionization injection (black), shock-front injection (blue), and shock-assisted ionization injection (green), and their respective spectra (bottom). **e**, Two X-ray beam profiles (false color). The large figure is a weak, elliptical betatron profile, while the small inset (shown to scale) shows a beam from the new radiation generation method demonstrated in **Paper VI**, which is described partially in Section 3.2.2 and 3.3.2.

BEAM INSTRUMENTATION

Any type of beam needs to be characterized in a multitude of ways. Preferably, one would have access to the full 6D characteristics, to really know all the details, but this is often difficult to obtain. In this Chapter, some diagnostic methods for laser and X-ray pulses, but mainly electron bunches, will be discussed. Diagnostics of the plasma structures themselves are also necessary, but these will not be covered here. A comprehensive review of diagnostics for electron bunches and plasma structures can be found in Ref. [139].

Much of what is described in this Chapter is a more direct account of work that has been carried out for this thesis than the previous Chapters have been. The majority of the described work concerns implementations of well-known diagnostic techniques, which have been used throughout various publications, or will be used for upcoming ones.

4.1 Laser diagnostics

While laser pulse characterization was not a major part of the thesis work itself, the pulses which were used for the different LWFA studies were characterized, and some basics on laser pulse characterization are therefore warranted.

4.1.1 Pulse energy

As seen in Eq. 1.56, knowledge of the pulse energy is required to be able to calculate the peak power and intensity of a pulse. During the work for this thesis, two different ways of measuring the pulse energy were used - a thermoelectric power meter and a pyroelectric energy meter [191, 192]. The thermoelectric (or Seebeck) effect is effectively the Peltier effect in reverse: a difference in temperature over the thermoelectric element gives rise to a voltage between its two sides. The laser pulses are absorbed by one side of such a power meter, depositing its energy as heat, and the temperature with respect to the “cold” side, typically a heat sink at room temperature, generates a voltage which can be measured and related to a certain power. This gives the average deposited power, such that for e.g. a 10 Hz laser the pulse energy is $1/10^{\text{th}}$ of the

power, and the device might heat up over time when using high enough average power, such that the measurement precision decreases.

A more precise measurement of laser pulses can be obtained with pyroelectric detector, where the pyroelectric effect is exploited [193]. Just as for thermoelectric detectors, the laser pulse is absorbed in the front face of the device. When exposed to a change in heat load, a crystal exhibiting the pyroelectric effects acts akin to a capacitor, and a voltage builds up between its two sides. The voltage is only generated as the heat load changes, constant heat loads generate no voltage, and so these detectors can only be used for short pulses. The generated voltage can then be related to a pulse energy. Such a pyroelectric pulse energy meter was used as part of the daily operation of the Lund High-Power Laser.

4.1.2 Longitudinal properties

In analogy to the longitudinal phase space of a particle bunch (see for example Fig. 2.2), the energy (typically displayed as frequency or wavelength) as a function of longitudinal position/time in a laser pulse can also be mapped in a *spectrogram*. Below follows a brief account of the methods used during the work for this thesis to measure the two constituents of the spectrogram, or the full spectrogram.

Spectrum

The simplest part to measure of the spectrogram is the spectrum of the light. To this end, angularly dispersive components are used, with gratings being the most reliable and common for the application. The dispersive element spreads the different energies, i.e. wavelengths, across a detector, and the intensity of each wavelength (the spectral intensity) is measured. The line density of the grating and detector spectral response must be matched to the spectrum, otherwise the dispersive strength or the detection efficiency, respectively, would suffer. Spectral measurements are performed at several locations along the Lund High-Power Laser system as part of the regular operation.

Temporal structure

Measuring the temporal structure of a laser pulse is less trivial to measure than the spectrum. The spectral intensity alone contains no information concerning when in time these frequencies are placed inside the pulse, only that they are there. For ultra-short laser pulses (pulse durations of few 10s of fs or less), no modulators exist which are fast enough to enable direct measurements of the pulse duration. Therefore, non-linear optical effects in materials are usually exploited to gain access to the time domain and/or phase by splitting the pulses and use the two copies to measure themselves. The two methods that were used during the thesis work were single-shot *intensity autocorrelation* [36, 194] and *frequency-resolved optical gating* (FROG) [36, 195].

In intensity autocorrelation, one observes the results of a process which is sensitive to the pulse intensity (as opposed to the field), and use this to correlate the intensity envelope with itself by splitting the pulse in two and scanning one half in time with respect to the other. Mathematically, this can be expressed as a convolution of the pulse intensity envelope with itself, the width of which differs by a known factor from

the width of the pulse envelope for a given pulse shape. However, a pulse shape must be assumed, and the factor between a signal and its autocorrelation can vary significantly depending on the pulse shape. Furthermore, the autocorrelation signal is always symmetric and no information about the phase is obtained, so while it can be practical to use for quick and simple pulse duration estimations, a full pulse shape retrieval cannot be performed in this way. A common way of implementing an intensity autocorrelator is to overlap the two pulse halves in a second-order non-linear medium and generate the second harmonic of the incident pulse frequency [36]. The generated field of the second harmonic pulse depends on the intensity of the generating pulse. As the pulses are scanned over each other, the overlapping intensity will vary, which affects the amount of generated second harmonic. To circumvent the requirement of scanning, the device can be made single-shot by instead allowing the two pulse halves to intersect at an angle in the second-order medium. The intersection angle allows streaking the two halves across each other, generating the second harmonic along the direction between the two pulses.

To obtain information about the full pulse envelope and its phase, additional operations must be performed. One way this can be done is by complementing the second-harmonic generation (SHG) scheme with a spectral measurement, yielding a spectrogram. One such implementation is called SHG-FROG. The additional spectral information at each time delay allows calculating a phase and pulse shape which produces a spectrogram equal to the measured one. Iterative algorithms are typically implemented to this end, which slows the process down compared to autocorrelation, but the benefits are that arbitrary pulse shapes can be obtained, together with the temporal phase.

To characterize the laser pulse duration throughout the work for this thesis, intensity autocorrelation has been the most common method. More thorough characterization using SHG-FROG has been performed more often during the past months, at the time of writing, although not directly by the author.

4.1.3 Transverse properties

Similarly to the CS parameters describing the transverse propagation of a particle beam, as partially seen in Eq. 1.52, there are also several components to describing the transverse propagation of a light pulse. The wavefront curvature and the wavelength together determine how the pulse will propagate, and what shape it will have at a waist. In reality, every mirror and other optic will have a small amount of roughness, which together with other effects during for example amplification and transport will affect the wavefront of a laser pulse. Instead of being simply plane or parabolic, the wavefront will in general be “wrinkled” and perhaps not even azimuthally or radially symmetric. To be able to focus the pulse to a symmetric and reliable spot, these aberrations should be corrected for to as large an extent as possible.

There are several ways of measuring wavefronts, but in general they rely on interference and/or diffraction of the beam. One example is a Shack-Hartmann wavefront sensor [196]. In such a sensor, a mask is introduced, for example a dense matrix of holes or a microlens array, after which the wavefront propagates onto a sensor. Depending on the local curvature of the wavefront at each aperture, the diffracting beamlets will have different sizes and divergences, which end up at different locations and with different sizes on the sensor, which can allow for a reconstruction of the

original wavefront. A wavefront sensor can then be used together with a *deformable mirror* (DM) to correct for certain aberrations. A DM generally consists of a malleable mirror membrane placed on top of a spatially adjustable structure such as an array of piezoelectric or electromechanical actuators. This means that the mirror can be made to change its local curvature by changing the actuators, which can correct for some aberrations. The wavefront sensor and the DM can be connected in a closed loop, such that a control system reading the wavefront sensor can be made to do the correction by actuating the DM. Such a closed loop is implemented at the LWFA in the Lund High-Power Laser Facility, where a Phasics SID4 wavefront sensor is used in connection with a piezoelectric DM as part of the daily facility operation to correct for wavefront aberrations (see Fig. 3.6).

4.2 Electron diagnostics

4.2.1 Longitudinal properties

Magnetic spectrometers

As shown in Section 1.1.1 and Fig. 1.1, particles with different energy will take different trajectories when passing through the same (electro-)magnetic field. This fact is exploited extensively to measure the energies, or (equivalently) momenta, of relativistic charged particles in beams by using magnetic spectrometers. The fundamental component of such a device is a dipole magnet, but charged particle spectrometers can, like their photon-equivalents, be made imaging, so as to improve the energy resolution. By carefully mapping out the \mathbf{B} -field and the geometry of the spectrometer, it is possible to infer the particle's energy from its position on a scintillating screen placed downstream of the dipole magnet. If the dipole magnet is electromagnetic, it can be tuned to optimally disperse different energy ranges. In LWFA setups, it is not uncommon to build magnetostatic dipoles for in-vacuum spectrometers using permanent magnets, see Fig. 3.6. As discussed in Section 2.3.2 and 3.3.2, the different magnet technologies have different advantages and disadvantages.

In principle, the simplest possible spectrometer design consists of a dipole magnet, a drift section, and a screen. However, because the electron beam always has some divergence, at some setup-specific combinations of divergence and energy spread, the divergence starts to limit the minimum measurable energy spread because the beam size on the screen starts to depend more on the divergence than the dispersion. This is particularly true for narrow-band LWFA beams, whose divergence is often considerable. To minimize the effect of the beam divergence, one can construct an imaging spectrometer, where focusing elements are used to image some point onto the spectrometer screen, through the dipole. For narrow-band beams in particular, where chromatic effects on the focusing are not a large concern, the total resolution of the spectrometer can be increased considerably this way. Chromatic focusing can also be used to estimate the beam emittance, as will be briefly discussed in Section 4.2.2 below. However, an imaging spectrometer will reduce the visibility of any angular effects in the beam, for example from DLA in LWFAs or hosing instabilities in PWFA. One way to partially solve this, and still get angular information along the axis perpendicular to the dispersion plane, is to use the dipole *edge focusing* [22], where different entrance positions and/or angles on the dipole will yield different exit angles depending on the

total integrated field that a particle traverses. This effect is very noticeable, though not optimized, in the electron spectrometer that was designed and built for the LWFA setup at the Lund High-Power Laser Facility, described below.

Two other practical resolution limits of a spectrometer are the amount of dispersion at a given energy, and the imaging resolution of the optical system imaging the spectrometer scintillating screen. For a broadband magnetostatic dipole spectrometer, the former typically causes the resolution to decrease with particle energy, because the dipole strength is fixed and thus causes a smaller spread in outgoing angles as the particle energy increases. This dispersion issue can be mitigated, to some extent, by optimizing the spectrometer geometry through placement of the scintillating screen or the spectrometer dipole shape and/or field distribution. The latter limit can be mitigated by using multiple cameras with overlapping fields-of-view. This allows a larger magnification of each section of the spectrometer screen, such that smaller features in the beam can be resolved, but requires a more complex experimental setup and more post-processing to produce a valid spectrometer image. The final resolution limit comes from the scintillating screen itself - any scintillator gives some minimum scintillation signal size which depends on the type of scintillator and its thickness. A thicker scintillator typically has more scintillating material, producing more photons per incident electron. However, care must be taken, since for opaque materials such as Lanex screens, a thicker scintillator also gives a “blurrier” signal because of for example scattering in the material, while for a transparent (usually crystalline) material such as yttrium aluminium garnet (YAG), the length of the scintillating trace can give a false signal if imaged at improper angles [197]. YAG screens can give a higher spatial resolution than Lanex screens, but the latter are significantly cheaper and exist in significantly larger sizes, so the application dictates which material is the best choice.

Knowing the conversion efficiency from charge to photons in the scintillator can be used to simultaneously obtain values for the bunch charge [198]. This is particularly valuable in LWFA experiments, where the noise from low-energy electrons and electromagnetic pulses (EMPs) can prevent, or at least impede, the use of conventional charge-measurement methods such as Faraday cups or integrating current transformers (ICTs). Estimating the charge from the scintillation signal was the charge-measurement method used with the electron spectrometer at the LWFA in Lund, which is described below.

Spectrometer design for the LWFA setup at the Lund High-Power Laser Facility

As an early part of the thesis work, a new electron spectrometer was designed and built for the LWFA experiments at the Lund High-Power Laser Facility. The design goals were to increase the energy resolution and detectable energy range over the then current spectrometer, and decrease the X-ray noise from bremsstrahlung coming from electrons colliding with the inside of the magnet. The previous spectrometer magnet was of O-type, meaning that the magnet yoke, i.e. the magnetically conductive (typically iron (low-carbon steel) -based) girder that helps close the field lines, was O-shaped, with a rectangular slot for the electrons to pass through. To reduce the amount of bremsstrahlung, the new magnet was chosen to be C-type, thus leaving one side open for the electrons to exit through. The open yoke shape also decreased




Figure 4.1: Spectrometer dipole magnet model. Comsol Multiphysics model of the spectrometer dipole cross section. The color map shows the magnitude of the magnetic flux density, $|\mathbf{B}|$, and the lines show the orientation of the field. The model shows that no magnetic flux should be present on the outside of the yoke. The measured, 2D B -field map of the dipole midplane is shown in Fig. 4.3a.



Figure 4.2: Spectrometer dipole magnet render. Rendered CAD model image of the electron spectrometer dipole magnet designed and used for LWFA experiments at the Lund High-Power Laser Facility. The nearest outermost surfaces are aluminium spacers. The shinier magnet blocks are visible inside the gap, and the internal wall is also an aluminium spacer.

the minimum detectable electron energy significantly, from roughly 40 MeV down to approximately 10 MeV. The magnet was to be used in vacuum and close to the source, so a permanent magnet-based solution using NdFeB (“neodymium”) magnets was chosen. The resulting design is shown in Fig. 4.1 and 4.2. In both figures, the electrons would enter the magnet “into” the page and exit to the left.

As shown in Fig. 4.1, the yoke is the largest part of the magnet, with the NdFeB magnet blocks placed as magnet poles above and below the pole gap in the center. The leftmost blocks, and the T-shaped block at the very center of the magnet, are




Figure 4.3: Geometry, trajectories and resolution of the electron spectrometer. **a**, Top-view of the spectrometer geometry, with the B -field map of the dipole magnet (color scale) and some example electron trajectories (black, solid lines). The grey line shows the Lanex screen and the dashed black lines show the aluminium spacers on the outside of the magnet (see Fig. 4.1 and 4.2). The electron source is at $S = 0$, and the different trajectory lines correspond to angles of 0 and ± 5 mrad with respect to the laser axis. It is clear that the highest energies, intersecting the screen at larger S , have a larger angular separation when impacting the screen. However, the angular separation is near zero at the second lowest energy depicted here - at this energy, the natural horizontal focusing of the dipole magnet cancels the divergence of the beam and effectively images the beam in the horizontal plane onto the Lanex screen. Lower energies over-focus, such that the angular separation increases again. **b**, Dispersion curve of the spectrometer, i.e. what position on the spectrometer screen corresponds to what energy. The full range of the spectrometer is approximately 350 MeV for the full screen length of 25 cm. **c**, Energy resolution as a function of angle and energy, given as the observed relative energy deviation per mrad of angular deviation from zero in the dispersion plane. Note the minimum at approximately 20 MeV; this minimum occurs because of the horizontal focusing visible in **a**.

aluminium safety spacers, also visible in Fig. 4.2. One downside of choosing a C-shape yoke is that all of the magnetic flux, the “return flux”, must pass through a single piece of iron, whereas for example O-type yokes provide two paths for the flux. Because of this, the cross section of the new yoke was made significantly larger than the previous yoke. A larger cross section prevents magnetic saturation of the yoke material, which would lower the gap field and force some fraction of the flux out through the surface, making the yoke magnetic. Being placed in a chamber where much work occurs often, typically using steel tools, an externally non-magnetic object is preferable.

The dipole magnet dimensions were designed around the off-the-shelf magnet blocks of $4 \times 2 \times 0.5$ inches, and the COMSOL Multiphysics [199] model shown in Fig. 4.1 used the magnetization given by the manufacturer, together with the $B(H)$ -curve [102] for the yoke steel. The magnet blocks were offset laterally away from the return yoke since the field is of opposite direction in the gap and the yoke and so passes through zero on or near the boundary. Considerations were made as to manual assembly of the magnet, something which influenced the choice of placing the magnets directly as the poles. The pole gap height, i.e. the vertical distance between the poles, directly influences the magnitude of the field and its distribution inside the gap, and a pole gap of 15 mm was chosen. A smaller pole gap favors a larger maximum field, thus increasing the dispersion, but clips parts of vertically divergent beams before they reach the spectrometer screen. Also the laser transmission is clipped vertically. After final assembly, the midplane of the magnet was mapped using a Hall probe, to obtain the field distribution, see Fig. 4.3a.

After its construction, this spectrometer became the main diagnostic for electron beams in the LWFA experimental chamber, see Fig. 3.6 for a sketch of the geometry and Fig. 4.3 for an example of electron trajectories and spectrometer resolution. Both electron and X-ray diagnostic performance were greatly improved over the old setup, with larger electron dispersion, giving better energy resolution and a larger detectable electron energy range, and less bremsstrahlung on the X-ray diagnostics, all owing to the open C-shape and stronger magnetic field of the spectrometer dipole. The spectrometer was used for the work presented in **Paper VI, VII, VIII and IX**.

Permanent magnets are prone to demagnetization by heat and ionizing radiation (such as energetic electrons), both of which are present currently. Over time the magnets will degrade at a rate corresponding to the deposited dose. The dose rate is rather low at a typical experiment using the Lund High-Power Laser, but warrants re-characterization every few years. If the full production and assembly had been outsourced, other designs could have been considered, such as placing the magnet blocks inside what is currently the return yoke, shielding them further from radiation. This could also allow for machining of the pole tips, then made out of steel, to provide other field distributions. One could also opt for using the slightly weaker but more radiation-resistant (and more expensive) samarium-cobalt (SmCo) magnets. The boom in use of permanent magnets for future synchrotron storage ring designs has led to many creative designs over the last few years, where inspiration could be gathered. Also other spectrometer geometries could be considered to improve the energy resolution at the higher energies, potentially together with the use of multiple cameras to increase the imaging resolution of the Lanex screen.

Transverse deflecting structures

Measuring the other part of the longitudinal phase space distribution, the time structure, is a less trivial matter also for electron bunches, particularly for the intrinsically ultra-short and divergent bunches from high-density LWFA. A technique which has been used successfully with high-brightness bunches from conventional accelerators is the *transverse deflecting structure* (TDS) or *cavity* (TDC) [200, 201]. The function is similar to that of a streak camera, which relies on imposing a known, time-dependent transverse kick, to map the temporal structure into a spatial one. A TDS beamline is currently being built at the MAX IV linac, see Section 2.4. This will act as the main diagnostic of the linac after the second bunch compressor. As part of the thesis work, evaluative simulations of the planned beamline were performed. TDSs rely on RF structures similar to those used for acceleration, but instead exciting a field distribution which is not necessarily cylindrically symmetric or belonging to either a TM or TE mode. There are both transverse electric and magnetic fields, for example E_x and B_y , which both contribute to the transverse force on the particle as (see Eq. 1.1) $F_x = -e(E_x + cB_y)$, where we assume that the electron is relativistic. These fields are typically constant over the aperture of the TDS. For simplicity, the fields are often combined into an effective voltage, $V_0 = -(E_x + cB_y)L/2$, where L is the length of the structure.

As in the accelerating structures (see Section 2.1.3), the fields are time-harmonic. If the bunch is placed at the zero-crossing of this transverse field, the head is kicked in one direction and the tail in the other, but the average still propagates straight. After propagation, this kick angle is turned into a change in transverse position - a streaking of the bunch where the spatial position on a scintillating screen depends on the bunch duration and the kick strength - and from this the time structure can be inferred. The temporal streak can also be coupled with a spectrometer dipole magnet bending in the perpendicular plane, for a full, single-shot characterization of the longitudinal phase space. The kick angle imparted by the TDS on a particle at position z is [200]

$$\Delta x'(z) = \frac{eV_0}{pc} \sin(\psi_{RF}) \approx \frac{eV_0}{p_z c} kz, \quad (4.1)$$

where we have expanded to first order for $\psi_{RF} = kz + \omega t$ around $z = 0$. It is evident that at $\psi_{RF} = 0$, we get the streak that we want along z , but no change of the average angle of the bunch. Using Eq. 4.1, we can define the *shear parameter* [201]

$$S = \frac{e}{c} \frac{V_0 k}{p_z} \approx \frac{eV_0 k}{\mathcal{E}}, \quad (4.2)$$

as the peak linear “strength” of the streak. At some point downstream of the TDS, the beam is focused onto a screen. The first-order transfer matrix element which maps x' to x , R_{12} (c.f. Eq. 1.22), can be given as $R_{12} = \sqrt{\beta_d \beta_s} \sin(\Delta\Psi)$ [200], where β_d and β_s are the β functions at the deflector and screen, respectively, and $\Delta\Psi$ is the betatron phase advance between the deflector and the screen, see Eq. 1.30. Evidently, we get a good transfer from angle to position if the β functions are large; at the deflector, the imparted TDS kick is a larger fraction of the intrinsic beam divergence if the beam is large, since $\gamma = (1 + \alpha^2)/\beta$. A phase advance of $\Delta\Psi = \pi/2$ maximizes the sinus term at unity. However, we want a large streak compared to the natural, unstreaked beam

size, $\sigma_{s,0} = \sqrt{\varepsilon\beta_s}$, at the screen to maximize the resolution, so β_s should be small. Assuming $\Delta\Psi = \pi/2$ and $\psi_{RF} = 0$, the streaked rms beam size at the screen is [200]

$$\sigma_s = \sqrt{\sigma_{s,0}^2 + \sigma_z^2 S^2 \beta_d \beta_s}. \quad (4.3)$$

For a streak to be strong enough for the streaked beam to exceed the un-streaked size on the screen, it must then fulfil

$$S > \frac{\sigma_{s,0}}{\sigma_z \sqrt{\beta_d \beta_s}} = \frac{1}{\sigma_z} \sqrt{\frac{\varepsilon}{\beta_d}} = \frac{\varepsilon}{\sigma_z \sigma_d}, \quad (4.4)$$

which is determined by a combination of RF frequency, streak voltage and beam energy. Typically the total voltage is a free parameter. Equation 4.3 can finally be used to calculate the bunch length from the measured beam size on the screen as

$$\sigma_z = \frac{1}{S} \sqrt{\frac{\sigma_s^2 - \sigma_{s,0}^2}{\beta_d \beta_s}}. \quad (4.5)$$

Apart from the transverse streaking field, there is also a radially varying longitudinal electric field which is important to consider, as it is fundamentally linked to the streaking component. At $\psi_{RF} = 0$, this field is 0 on axis and has a linear gradient in the streaking direction. This means that particles on transversely opposite sides of the beam will gain or lose energy, respectively, depending on which side and which zero-crossing they are on. This increases the energy spread in each slice (the uncorrelated energy spread), which potentially distorts the spectrum and decreases the measurement precision of the longitudinal phase space. The increased rms slice energy spread increases with the streak strength as [201]

$$\sigma_{\mathcal{E}} = ekV_0 \sigma_d \iff \frac{\sigma_{\mathcal{E}}}{\mathcal{E}} = \sigma_{\delta} = S \sigma_d, \quad (4.6)$$

and so for a higher temporal resolution, one must accept that the slice energy spread increases.

For the MAX IV linac diagnostics beamline, a polarizable TDS, meaning that its streaking axis is variable, was designed [202] under the (partial) constraints that 1. the same RF frequency as the linac, 2.9985 GHz, and existing RF infrastructure was to be used, and 2. the total streaking voltage should be more than 100 MV to reach the desired resolution of 1 fs. With a nominal energy of 3 GeV and an emittance of $\varepsilon = \varepsilon_n/\gamma_b$, where $\varepsilon_n \approx 1$ mm mrad, the minimum streaking voltage is guided by Eq. 4.4 such that $V_0 \sqrt{\beta_d} > \mathcal{E} \sqrt{\varepsilon}/(ek\sigma_z) = \mathcal{E}_e \sqrt{\gamma_b \varepsilon_n}/(ek\sigma_z) \approx 2 \times 10^9$ Vm^{1/2}. To get sub-fs resolution, $\beta_d = 1000$ m \leftrightarrow $V_0 \gtrsim 66$ MV. To this end, two 3-m TDSs, each with $V_0 \approx 61$ MV, were designed. As part of the thesis work, particle tracking simulations of this beamline were carried out for internal evaluation. A result from these simulations is shown in Fig. 4.4, where the beamline layout and a virtual diagnostics screen are shown, together with the reference simulation traces at the TDS position. The TDS screen data with the analytical dt/dx scaling (from Eq. 4.1) and the ‘‘calibrated’’ scaling, obtained by varying the RF phase in small increments and observing the centroid displacement, were nearly identical, but both differ a little from the reference current trace. The resolution is just above 1 fs, possibly due to

smearing from chromatic aberrations caused by the induced slice energy spread which is increased by approximately a factor 10-30 in the TDS, varying along the bunch. Equation 4.6 projects nearly 1 per mille of induced slice energy spread.

For comparison, another polarizable TDS type, which has been designed as a collaboration between PSI, CERN and DESY, is the PolariX-TDS [204, 205]. This device is based on X-band technology, with an RF frequency of ~ 12 GHz, paired with a voltage of ~ 25 MV for a final resolution of about 5 fs. The first such TDS is in use at the FLASHForward beamline at DESY. Experimental results using the device are currently under review [206].

While TDSs are useful, they are difficult (although not impossible) to use in conjunction with an LWFA setup, and typically expensive to manufacture and operate. The required RF infrastructure, which typically exists naturally at “conventional” accelerator facilities, is usually not present in laser laboratories. As discussed in Section 3.3.2, the bunch length from an LWFA quickly increases for several reasons. Since

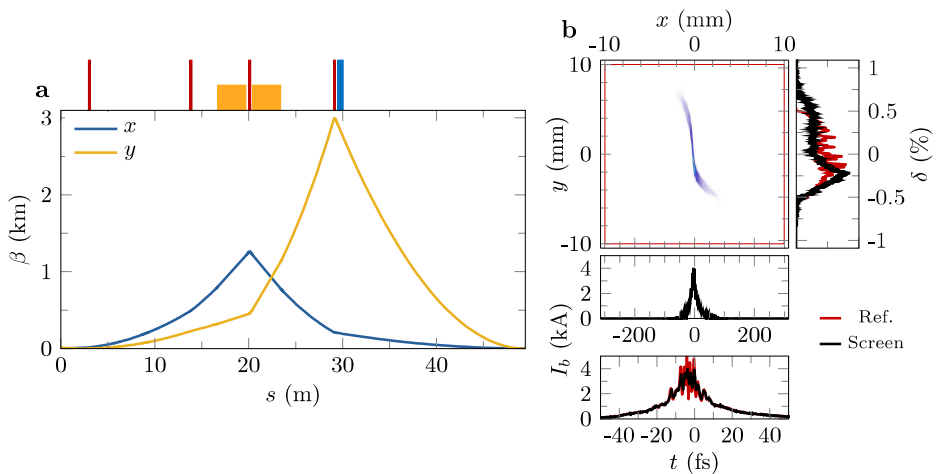


Figure 4.4: MAX IV TDC beamline and simulation example. **a**, Magnet layout (top) and β functions (bottom) for the TDC diagnostic beamline. Quadrupoles are red, dipoles blue and TDC units orange. **b**, Simulation result from tracking simulations through the TDC beamline using Elegant [203]. The bunch charge is 100 pC. The 2D histogram is a virtual diagnostics screen in the transverse plane, placed at the end of the beamline, such as it might look in an experiment. The grid has been set to 1024×1024 to be in the same range as a camera imaging the screen, and the red square indicates the edges of the scintillating screen. To the right, the spectrum obtained from the screen (black) and the reference (red, simulated beam before the TDS) are shown, where the slice energy spread induced by the TDS (and possibly also coherent synchrotron radiation (see Section 2.3.2)) is clearly seen. Below, the corresponding current trace is shown. The current trace is not limited by the grid resolution. The scaling is obtained from a virtual calibration, by shifting the RF phase to not correspond to zero-crossing, and thus moving the beam in space. At the bottom, a zoomed-in current trace is shown together with the reference. Some details are partially lost, such as the coherent synchrotron radiation-induced microbunching spikes, which are around 1 fs long in the reference trace. Thus, the resolution appears to be just above 1 fs.

the bunch duration which is inferred from a TDS trace is the duration of the bunch inside the TDS, when the angular kick is applied, and the resolution also relies on proper imaging of the beam (see Eq. 4.5), the bunch would have to be carefully collimated and refocused to enable a good temporal measurement. With the fs-durations of LWFA bunches, and with percent-level energy bandwidths in a good case, this requires a very powerful streak and complicated electron optics to mitigate detrimental effects. Placing the TDS close to the source would solve some problems, but operating a high-power RF structure requires good vacuum (RF breakdown from arcing quickly becomes a problem otherwise), which is difficult to achieve close to a gas source. For these reasons, TDSs are normally impractical to use in combination with LWFA experiments.

Coherent transition radiation

Data of the actual time structure is a very valuable metric, and so alternative ways of measuring this have been explored. One of these, which was also attempted during the thesis work, uses *coherent transition radiation* (CTR) [21, 207], which has been used by several groups to this end. Just as radiation is generated when relativistic charges accelerate and their electromagnetic fields are forced to rearrange, see Section 1.1.2, so can radiation be generated when charged particles traverse the boundary between two spatial regions with different dielectric properties. A sharp transition between two strongly differing media is preferred. Practically, the transition between a thin metal foil and vacuum is easiest to implement, as (for LWFA) the remaining laser pulse energy can then be blocked simultaneously. By measuring the spectrum of the CTR, i.e. specifically the coherent part of the transition radiation (TR), one can infer the duration of the electron bunch.

As mentioned in Section 1.3, metals and plasmas respond similarly to electromagnetic fields, as both are macroscopically neutral and contain a quasi-stationary ion background and a “gas” of free electrons. Therefore, one can deduce a plasma frequency for the metal as well. Because of their high electron density, on account of being solid materials, only high frequencies can propagate in metals at all, c.f. Eq. 1.62 and 1.63. A practical choice of metallic foil is aluminium, as it is cheap and easy to work with. At room temperature, aluminium has a free-electron density of roughly $6 \times 10^{22} \text{ cm}^{-3}$, ~ 4 orders of magnitude higher than the typical LWFA plasmas used in the experiments presented in this thesis. This electron density gives a plasma frequency of $\omega_p \approx 1.4 \times 10^{16} \text{ rad/s}$, as compared to $\omega_0 = 2.35 \times 10^{15} \text{ rad/s}$ for an 800 nm laser. Since the TR can also be seen as being generated by the oscillating charges on the medium surface, triggered by the passage of a short electron bunch, the shortest wavelength that can be generated corresponds roughly to the plasma frequency, $\lambda_{\min} \gtrsim 2\pi c/\omega_p \approx 136 \text{ nm}$. However, similar to CSR, the coherent part of the TR is only generated when the medium response from individual bunch electrons are roughly in phase, which happens when the bunch length is shorter than half the radiated wavelength, $\lambda_{\text{CTR}} \gtrsim 2c\sigma_t$. For $\lambda_{\text{CTR}} = \lambda_{\min}$, the bunch duration corresponds to $\sigma_t \approx 0.23 \text{ fs}$, shorter than what would be expected from typical LWFA parameters used during the thesis work, and so these short wavelengths are not typically generated to a large extent. Conversely, a bunch duration of $\sigma_t = 1.3 \text{ fs rms}$ (3 fs fwhm) would suggest $\lambda_{\text{CTR}} \gtrsim 0.8 \mu\text{m}$.

The maximum wavelength could be estimated by a similar argument regarding

the dimensions of the coherence volume [21]. One would expect that wavelengths longer than the coherence length, D , are not efficiently generated by a bunch which is significantly shorter, i.e. $\lambda_{\max} \lesssim D$, where $D = \gamma_b c / \omega_p \approx 8.6 \mu\text{m}$ for $\gamma_b = 400$. This means that one would expect the CTR to be emitted within a very broad wavelength range from 136 nm to roughly 9 μm . To properly characterize the bunch length, an advanced spectroscopic setup is required, since truncating the spectrum means losing information about the CTR pulse. However, just as for a laser pulse (see Section 4.1), the spectral intensity of the CTR alone is not enough to characterize the pulse duration; without information about the phase, there is no information about the relationships between the different wavelengths, and one would have to make strong assumptions for the bunch length calculations. For more exact characterization, phase-retrieval algorithms need to be employed [208], which is also done in for example FROG characterization of laser pulses. In the case when there are multiple electron bunches, the ambiguity is less, as the spectral oscillation period is directly related to the bunch separation [208, 209], but a full characterization would still require the phase to be calculated. However, CTR incurs an additional hurdle compared to a typical laser pulse: the spectral and angular components of the CTR are coupled, so considerations as to the transverse distribution of the radiation must be made.

The TR energy radiated per solid angle and unit frequency by a single electron is [207]

$$\frac{d^2 W_e}{d\omega d\Omega} = \frac{e^2}{\pi^2 c} \frac{\beta_b^2 \sin(\theta)}{(1 - \beta_b \cos^2(\theta))^2}, \quad (4.7)$$

where θ is the angle between the observation direction and the propagation axis of the particle. The radiated TR energy is, unlike synchrotron radiation (see Section 1.1.2), maximal near $\theta = 1/\gamma_b$ and zero on axis. This radiation angular profile thus has the shape of a hollow cone. Assuming that the motion of electrons in a bunch is roughly normal to the transition surface, and that their momentum distribution is cylindrically symmetric, the total radiated energy is given by adding the incoherent (ITR) and coherent parts as [207]

$$\frac{d^2 W_{\text{TR}}}{d\omega d\Omega} = \frac{d^2 W_{\text{ITR}}}{d\omega d\Omega} + \frac{d^2 W_{\text{CTR}}}{d\omega d\Omega} = \frac{d^2 W_e}{d\omega d\Omega} (N_b + N_b^2 |F|^2), \quad (4.8)$$

where $F = F(\omega, \theta)$ is the spatial form factor of the electron distribution. A bunch charge of 10 pC corresponds to $N_b = 6 \times 10^7$ electrons, so the coherent part (right-hand side) will clearly dominate for wavelengths that emit coherently. By assuming that the bunch distribution is Gaussian in space (see Eq. 1.31), and that there are no correlations between positions and momenta, the form factor can be expressed as [207]

$$F = F_{\perp} F_{\parallel} = \exp\left[-\frac{1}{2} \left(\frac{\omega \sigma_{\rho}}{c}\right)^2 \sin^2(\theta)\right] \exp\left[-\frac{1}{2} \left(\frac{\omega \sigma_z}{c}\right)^2\right]. \quad (4.9)$$

Since the goal of this technique is to get a value for σ_z , F_{\perp} must be suppressed ($F_{\perp} \approx 1$) or individually characterized, as it is F which is measured. Because of the assumption of no correlation between position and momentum, the radiator must be placed as close to the source as possible (unless focusing optics are used, which might also distort the bunch length), so that σ_{ρ} is relatively small. Also the radiation collection angle should be kept small to minimize θ while still collecting enough energy.

If these requirements are met, σ_z can be calculated using inverse Fourier transform, ideally with the phase retrieved.

As part of the work leading up to the study in **Paper VI**, an effort was made at the LWFA in Lund to generate double-bunches and characterize their separation using CTR. As mentioned above, two temporally separated bunches would yield spectral interference, which contains information about the bunch separation [209]. However, to evaluate the setup parameters, the single-bunch CTR should be considered first. As an example of the generated CTR properties, we show calculations similar to the geometry used in the double-bunch pre-study, as shown in Fig. 4.6. Using a typical measured bunch divergence of $\sigma_{\rho'} = 1.5$ mrad and a distance to the foil of 10 mm, the electron beam size at the foil would be roughly $\sigma_{\rho} \approx 15$ μm , which is a significant increase from the $\lesssim 1$ μm at the source. If we want to measure frequencies down to for example $\omega_p/4 \leftrightarrow 410$ nm with decent precision, we need the exponent of F_{\perp} to be less than $1/2$, and so $\sin(\theta) < c/(\omega\sigma_{\rho}) = 4c/(\omega_p\sigma_{\rho})$, which means that $\theta \lesssim 5.8$ mrad. This is well outside of the maximum of the radiation cone at $\theta \approx 1/\gamma_b \approx 2.5$ mrad ($\gamma_b = 400$), such that a decent amount of energy would still be collected. The angular distribution is shown in Fig. 4.5a. Narrowing the collected angle further would yield further improvements in precision, but might be limited by the capture efficiency of the imaging system and detectors. For example, the same calculation for $F_{\perp} > 0.9$ gives $\theta \lesssim 2.6$ mrad, around the peak radiated energy.

To obtain the expected CTR spectrum, Eq. 4.8 can be integrated over the angle, $d\Omega$, using Eq. 4.9. As the captured spectra are typically given in wavelength, it can be more useful to visualize the spectrum as a function of λ , i.e. $dW/d\lambda$. A numerical integration is the most straight-forward approach, and can provide solutions for a large variety of bunch shapes (also non-Gaussian). Figure 4.5b shows the spectra from a bunch with the parameters given above and a duration of 3 fs fwhm, for integration intervals of $0 \leq \theta \leq 2.6$ mrad and $0 \leq \theta \leq \pi/2$. It is evident that the spectrum of the radiation captured within a smaller angle is shifted with respect to a using the full




Figure 4.5: Calculated CTR angular distribution and spectrum. **a**, Angular distribution of the CTR with $\gamma_b = 400$. The distribution peaks at $\theta = 1/\gamma_b \approx 2.5$ mrad. The dark grey line marks $\theta = 2.6$ mrad, the upper limit for one of the spectral calculations in **b**. **b**, Normalized, calculated CTR spectrum when integrating over all angles and over a limited region to minimize F_{\perp} . The smaller integration limit is shown by the dark grey line in **a**. The spectrum is noticeably shifted towards shorter wavelengths. The full-angle integration has been reduced by a factor of ~ 15.5 compared to the limited integration.

emission angle and that this would give the wrong values of σ_z .

A schematic of the setup used in the pre-study is shown in Fig. 4.6a. A replenishable aluminium tape, 40 μm thick, was placed 9 mm from the source, as close as the gas nozzle allowed. The tape blocked the remaining laser energy and was used to generate CTR. A silver-coated off-axis parabola and plane mirror were used to collimate and redirect the generated radiation and send it to a ZnSe wedge placed at ~ 45 degrees, where the high refractive index of ZnSe caused much of the (S-polarized) visible light to be specularly reflected, see Fig. 4.6b. Accounting for polarization is important, because CTR is radially polarized. The visible light exited the vacuum chamber through an ultra-violet-enhanced fused silica (UVFS) wedged window and was directed via two plane silver mirrors to an imaging spectrometer for 500-1000 nm, onto which it was imaged with a UVFS planoconvex lens. The visible arm was reliable and worked well. Figure 4.6e shows a visible CTR spectrum and inferred time structure (Fig. 4.6d) generated by the electron bunch whose spectrum is shown in Fig. 4.6c. The near-IR CTR was mostly transmitted through the ZnSe wedge and exited the vacuum chamber through a CaF_2 wedged window and imaged onto a StellarNet cooled InGaAs array spectrometer for 900-2300 nm using a CaF_2 planoconvex lens. The NIR spectrometer often malfunctioned and the data that was taken turned out to not be usable. As seen in Fig. 4.6e, the visible spectrometer alone gave too little spectral information to draw any conclusion, as only partial oscillations could be measured. There were indications of time separations on the order of 6 fs (Fig. 4.6d), which were corroborated by simple analytical calculations of the CTR spectrum

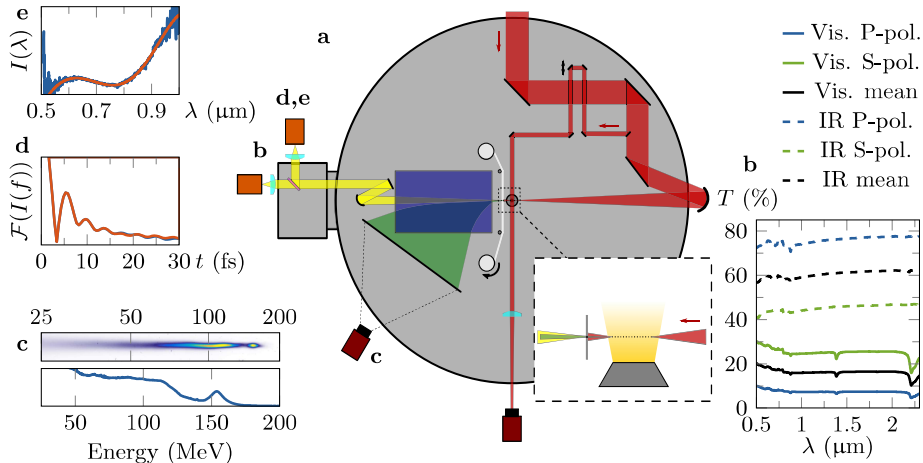


Figure 4.6: LWFA CTR setup and data examples. a, Setup schematic, c.f. Fig. 3.6. The X-ray camera is replaced with two spectrometers, one for visible light (uppermost, d) and one for near-IR (leftmost). There is also a replenishable aluminium foil placed near the electron source to block the laser and generate CTR. b, calculated optical transmission for S- and P- polarization for the optical components (except lenses) up to the entrance of each spectrometer. Transmission was calculated for the full range for both arms. c, Electron spectrometer image and spectrum for the electron beam that generated the CTR spectrum and inferred temporal structure shown in e and d, respectively.

using Gaussian bunch current shapes and the experimentally measured charges, but the data was inconclusive and ended up not being used.

4.2.2 Transverse properties

While measurements of the longitudinal phase space distribution are important and the subject of much research, the full 6D brightness cannot be characterized without also measuring the transverse phase space distribution. This distribution is not typically measured directly, as the longitudinal phase space distribution can be with a TDS - one instead typically measures the emittance of the beam in the two transverse directions. Emittance measurements were not performed as part of the present thesis work, but is mentioned briefly here because emittance has been a recurring topic through this text, and that the LWFA electron transport beamline (described below), which was designed as part of the thesis work, could be used for e.g. emittance measurements in the future.

Emittance measurements

The workhorse of emittance measurements is the *quad scan*, where the strength of an electromagnetic quadrupole magnet is varied and its effect on the transverse beam dimensions is observed on a downstream diagnostics screen [29]. Apart from the emittance, such a procedure also determines the CS parameters going into the (first) quadrupole used in the measurement. The quad scan technique is intrinsically multi-shot, and while it works well at conventional accelerators where the beams are generally stable, single-shot techniques are preferred for PBAs, which can have large shot-to-shot fluctuations. A technique which borrows some aspects from the quad scan is the “butterfly technique”¹ [160, 210]. Here, an imaging spectrometer is used to focus a given energy onto the spectrometer screen. Because of the tight focusing, chromatic effects are noticeable, and the beam size grows quickly outside of the nominal energy because lower energies are over-focused and higher energies are under-focused at the position of the screen. This gives the trace on the screen a butterfly-like waist which depends on the energy, the beam transport system, and, of course, the emittance. The beam width as function of energy can be fitted using the emittance (provided the beam transport system is well known), and provides single-shot measurements. However, it only gives a value for the direction perpendicular to the dispersion plane of the spectrometer magnet. Similar techniques could be used with a TDS to diagnose the slice emittance of the bunch. Recently, a compact imaging spectrometer setup for spectrometry and emittance measurements was realized using a single active plasma-lens as the focusing optic [211].

Single-shot measurements can also be obtained by using betatron [212, 213] or inverse Compton scattered [214] X-ray photons. Like the above methods, these methods are destructive, since when the electron spectrum is measured the beam is lost. An important difference to the methods described above is that the X-ray-based measurements provide information about the emittance inside (or very near) the plasma channel, whereas all other measurements are performed after a significant propagation distance (which might incur significant emittance increase). While the latter case is

¹This technique has not, to my knowledge, been officially named, so I this is what I decided to call it.

arguably the more interesting metric when it comes to applications, the former gives information on the *intrinsic* emittance, in principle the lowest attainable value and thus what could potentially be reached post-plasma if proper density ramps and optics are used. While the above references were specifically for LWFA experiments, similar and adapted techniques are interesting also for PWFA experiments [132, 215, 216] for the same reasons.

Transport beamline

The very first project that was started during the work of this thesis was the design of a short transport beamline for LWFA electrons. Because of budget reasons, it was not finished until several years later. Several different constellations were envisioned, all based around magnets - some paired a primary permanent magnet quadrupole triplet with a secondary electromagnet quadrupole triplet, while others employed one or the other. The final version is based around an electromagnet quadrupole triplet, placed outside the main vacuum chamber, with each magnet individually controlled to provide the maximum possible flexibility with respect to imaging and collimation. The maximum specified gradient of each magnet is about 40 T/m (with some margin for higher gradient), with a pole aperture diameter of 15 mm and an effective magnetic length of 100 mm. Each magnet has a supporting structure with a large adjustment range of ± 5 mm in x and y , to be able to counteract transverse misalignment of supporting tables, etc. All three magnets are geometrically and electrically identical,




Figure 4.7: LWFA beamline quadrupole in the lab. One of the three identical electromagnetic quadrupole magnets used for electron beam transport in the lab at the Lund High-Power Laser Facility. This is the same magnet that is shown on the thesis cover.

and have the same (nominal) gradient range of 0-40 T/m. A picture of one of the magnets is shown in Fig. 4.7.

A 1-to-1 imaging constellation of the triplet is shown in Fig. 4.8a. The initial CS parameters were $\beta^* \approx 0.5$ mm and $\alpha^* = 0$, as estimated from the divergence of experimental beams of 4 mrad fwhm, a normalized emittance of 0.5-1 mm mrad and energy of 200 MeV. As the quadrupoles have a relatively low gradient, they must be placed some distance away from the source to be effective at focusing in this energy range. The front face of the first magnet in this setup is 65 cm from the source, right outside the main vacuum chamber. The magnets are placed with 10 cm spacing, followed by another drift of 65 cm, focusing to a nominal β function equal to that of the source. Because of for example normalized emittance growth (see Section 3.3.2) and energy spread (yielding large chromatic aberrations from the focusing, see Fig. 4.8b), the beam radius at the downstream focus should in principle never be expected to be equal to that at the source. Since the strength of every magnet can be individually controlled, completely different constellations could also be envisioned, but the first drift length has a minimum determined by the size of the experimental chamber. The center quadrupole, which is the strongest in a symmetric imaging triplet configuration, has $\kappa \approx 56$ m⁻² for the configuration in Fig. 4.8, which gives a maximum refocusable energy of approximately 215 MeV for the nominal gradient of 40 T/m. However, the magnets were tested up to a maximum of 52 T/m by the manufacturer, corresponding to a maximum energy of 280 MeV. Another option is to use a plasma lens near the source, to collimate the electrons before magnetic transport, which has several benefits (see Section 3.3.2).

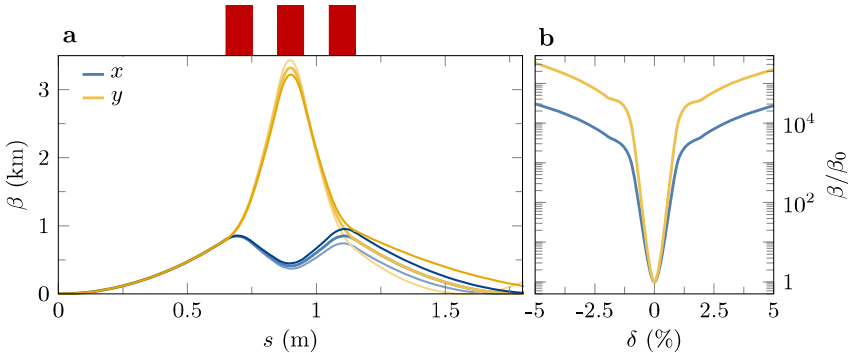


Figure 4.8: LWFA electron beamline optics. **a**, Element positions (top) and β functions along the electron transport beamline (bottom), for nominal energy (thicker line) and $\delta = \pm 5\%$ (thinner lines). Darker shade shows higher energy and vice versa. The first drift is determined by the size of the experimental chamber, since the quadrupoles are placed in air. β_y becomes largest in the center quadrupole and beam collision with the walls of the vacuum pipe will be an issue. The chromatic aberrations are also larger in this direction. **b**, β functions as a function of δ at the end of the beamline. Evidently, the chromaticity of this system is enormous, as should be expected from the tight focusing and large β functions. The chromaticity is largest in the y -direction, such that if paired with a spectrometer dipole, the dipole should disperse in x . However, the dipole would alter the optics functions somewhat, which must be taken into account.

Not all applications are equally sensitive to the chromatic effects in the beam. The transport beamline was recently used in a study towards using LWFA electrons for medical radiotherapy. High-energy electrons have better properties than X-rays when it comes to depositing energy in tissue, and the requirements of preserved bunch lengths, etc., are significantly relaxed (although other demands do arise instead). Therefore, a quadrupole magnet setup can work well, also for LWFA electrons. The study used a tissue phantom, emulating human muscle tissue, and the electrons were focused into this phantom, which was rotated to simulate a radiotherapy treatment session. Unfortunately, the manuscript is in a too early stage to include in this thesis.

This quadrupole magnet triplet can also be paired with for example a downstream dipole magnet, provided that a vacuum chamber with an electromagnetic dipole magnet is designed and procured. The electromagnetic dipole (as opposed to a permanent magnetic dipole) would, together with the electromagnetic quadrupoles, allow for tunable imaging of specific energies onto a given point. Such a setup would allow high-resolution spectrometry of narrow-band bunches, such as the ones shown in **Paper VI**, and also emittance measurements with the butterfly technique, as described above in Section 4.2.1. To avoid emittance degradation, the bunch must be collimated near the LWFA (see Section 3.3.2), which could be done with for example a passive laser-plasma lens for low-charge bunches (so as not to drive a beam-wake), or an active plasma-lens placed further from the source for higher-charge bunches.

4.3 X-ray diagnostics

When producing X-rays for applications or diagnostics, the X-rays must also be detected somehow. Although not the main focus of this thesis, different X-ray diagnostics were used extensively. Below will follow a brief summary of the techniques used herein.

Direct X-ray detection

Complementary metal-oxide semiconductor (CMOS) technology is today the most common technology for producing camera chips for the visible part of the spectrum, largely on account of being cheaper and easier to manufacture than their charge-coupled device (CCD)-based counterparts. However, CCD-based photodetectors are still useful in some applications, such as *direct X-ray detection* [217]. The CCDs' low noise, good spatial resolution and the possibility to detect and read out the X-ray photons in the same device (*direct detection*) have made them popular for e.g. X-ray detection. Another detection method is based on scintillators, where the X-ray energy deposition and signal detection happens in more than one step (indirect detection). In *deep depletion* CCD X-ray detectors, the incident X-ray photons deposit their energy in a thick semiconductor layer. Electron-hole pairs are created where each pair requires some relatively small amount of energy to be created, and thus many pairs are created per X-ray photon. These electrons can then be detected by an integrated circuit, which requires some number of electrons to register a count. When properly cooled, such that thermal excitation is only responsible for a small number of spurious pair-creations, this method typically gives a very good ratio between detected and false counts, i.e. a good signal-to-noise ratio. Since the number of counts is proportional to the deposited photon energy, such a detector can also be used directly for X-ray spectroscopy, under some conditions. Modern chips are sensitive to a wide range of

photon energies, for example $\sim 400 - 10\,000$ eV for the camera used during the thesis work.

The main X-ray diagnostic used for the LWFA experiments at the Lund High-Power Laser Facility (see Fig. 3.6) was an Andor Ikon-L SO back illuminated deep-depletion X-ray CCD [218] placed directly in-line with the source. The sensor has 2048×2048 pixels of $13.5\ \mu\text{m}$ side length, typically cooled to $-70\ \text{°C}$ to reduce the electronic noise. This type of detection setup was used for **Paper VI, X, IX** and **VIII** where single-photon counting was used for **Paper VI** and Ross-filters were used for **Paper VIII** and **IX**. These techniques are explained below.

Single-photon counting

One method for doing X-ray spectroscopy is *single-photon counting* [219]. Here, the energy deposited per X-ray photon is used directly to calculate the spectrum of the source. However, this approach would yield a false spectrum if several photons were to hit the same or close-lying pixels, so the photon flux must be reduced enough that there is at least a few pixels between each detection event. Typically, one also discriminates between events based on their shape. A single photon might trigger electron-hole pair creation in several pixels, by for example a high-energy photon creating enough pairs that they spill over to neighboring pixels, or the photon being incident with an angle such that it deposits its energy partially in several pixels. Some shapes are less likely than others to be created by a single photon, and are therefore removed from the analysis. To get enough statistics to recreate a reliable spectrum, it is desirable to detect many events, so a large CCD chip with a relatively high pixel density is useful. The X-ray flux can be decreased by moving the CCD further away from the source, letting natural diffraction spread out the photons, or by introducing characterized absorbers. The divergent betatron X-ray beams are a good match for this technique.

Ross filters

Well characterized absorbers can also be used in another spectroscopic method, using a *Ross filter* setup [220, 221]. Here, the well-known spectral response and thickness-dependent X-ray absorption of different elements is used to characterize the spectrum. By letting the X-ray pulse illuminate several different pieces of material, and possibly overlapping combinations thereof, the transmission through the different filter pieces can be used to fit a spectrum. In such a setup, the X-ray photon flux must be much higher than for single-photon counting, and so the CCD is placed closer to the source, and other attenuators are avoided. A downside is that a spectral shape must be assumed - it does not provide an arbitrary spectrum as single-photon counting does.

Dispersive spectrometers

Another method for spectroscopy is to use an angularly dispersive element such as a grating or, as in the case of **Paper VII**, a diffractive crystal. Such setups are not restricted in terms of photon density on the detector, like single-photon counting is, but are typically quite lossy because of for example photon absorption in the optics. However, they can yield a much higher spectral resolution than single-photon counting, and can be placed in an imaging configuration. This method (at least the grating-based implementation) is similar to spectroscopic techniques used for visible light.

SUMMARY AND OUTLOOK

In summary, I hope that these previous Chapters have provided an interesting read, and that those who want to read through the Papers now have enough background information to do so. After all, the purpose of the thesis introduction of a Swedish thesis is precisely to provide a background for the Papers.

In Chapter 1, we went through what is required to increase the kinetic energy of a charged particle, how its fields transform when it becomes relativistic, and that relativistic particles radiate when accelerated perpendicular to their direction of motion. We then went on to describe how just a single particle propagates in vacuum and in the presence of a linearly restoring force. From there, particle ensembles and their propagation were described, as well as the fields of such an ensemble. We also described a common figure of merit of a particle beam, the beam brightness, and that because it can never improve again once it has deteriorated, one must make sure to preserve it by handling the beams properly.

After particles, electromagnetic wave propagation was described by beginning with the simple ray optics model and going via Maxwell's equations to the description of electromagnetic waves and their potentials. It was described how rays, beams and pulses propagate in media, and that intense pulses can alter the media they are propagating in. What plasmas are and how they are generated was described next, followed by how electromagnetic waves propagate in them. The response of the plasma to intense pulses was shown, and that sufficiently intense pulses can drive wakes in the plasma. The wakefields in such wakes are very high compared to fields in "conventional" radio-frequency accelerating structures. It was also shown that dense particle beams can drive such wakes.

Chapter 2 outlined the principles of conventional accelerators, how electron bunches can be generated in them and how the bunches can be accelerated and controlled. Problems arising from bunch compressors were discussed in the context of plasma-based accelerators, and the "conventional" accelerator which was studied during for this thesis, the MAX IV linear accelerator, was described. In Chapter 3, some limitations to the acceleration distance of mainly LWFAs were shown, where the accelerating particles can outrun the wake and dephase. The limitation does not apply in the same way if the wake is driven by an electron bunch. If the driver and witness have nearly the same velocity, the strong focusing inside the wake can keep them

propagating for long distances without loss of brightness. It was also shown that the strong focusing in the wake can generate betatron X-ray pulses, enabling generation of keV-energy photons from small sources. These sources are typically gas jets or gas cells of mm to cm lengths, depending on the experimental parameters. The different source types are also preferentially used with different types of injection methods to generate the electron bunches inside the plasma.

It was also discussed that the plasma-generated bunches can be used in a second stage as a driver in a hybrid configuration, and that the bunches are very divergent, such that the means for beam control that were discussed in Chapter 2 are sometimes difficult to apply. To control plasma-generated bunches transversely, one might instead use plasma-lenses of various types, as these are much stronger than the typical magnet types commonly used for focusing. The plasma lenses are also radially focusing, unlike the magnets. The strong focusing can be a problem if the energy spread of the bunch is large, which is one reason why beam loading is important for minimizing the induced energy spread and increasing efficiency of plasma-based accelerators. An example of a plasma-based accelerator was shown with the LWFA at the Lund High-Power Laser Facility.

The final Chapter, Chapter 4, discussed how to characterize laser pulses and electron bunches, and outlined some of the differences and similarities between these methods. The design of an electron spectrometer for use in the LWFA setup in Lund was outlined, and evaluative simulations for a transverse deflecting structure diagnostics beamline at the MAX IV Laboratory were shown. The principles of using coherent transition radiation for temporal bunch measurements were explained and some preliminary measurements and calculations thereof were shown. Some characterization techniques for the beam emittance were outlined, as was the design of an electron transport beamline for the LWFA in Lund. The final Section concerned some detection principles for X-rays and their application to spectral measurements.

Outlook

Below follows an outlook based on the different Papers included in this thesis, located in the Appendix. For the outlook to make sense, I would advise to at least read through the “Article summaries and the author’s contributions” on page 103, if not the Papers themselves.

Building on the work presented in this thesis is in many cases possible to do almost directly, and there are some promising leads. Starting with LWFA, a main trajectory would be to follow up on the new radiation generation scheme shown in **Paper VI**. To increase the generated brightness, higher bunch charges, and possibly also electron energies, are needed. Bunches such as those presented in Ref. [159] might be ideal to this end. This might be possible to do in Lund, by for example increasing the total acceleration distance by using a different gas source, but higher laser powers would be very advantageous (as shown in Ref. [159]). Therefore, experiments at a higher-power laser facility, guided by an extended set of simulations and theory, could really show the potential of this source, which I believe to be very high. Future experiments should include, among other things, spectral and source-size measurements of the X-rays, and their scaling with the parameters of electron beam and plasma lens. Because of the experimentally simple implementation of this scheme, it is my hope that it could be a go-to method for compact X-ray generation in the future.

One of the key benefits of the collimated X-ray pulses demonstrated in **Paper VI** is that they are easier to re-focus and transport for applications, such as was done for betatron X-rays for the study in **Paper VII**. Following a (potential) source development as per above, applications benefiting from using X-ray optics could be considered. It would be useful to involve other researchers here, with expertise in for example X-ray optics and applications of X-ray imaging/spectroscopy. It might also be possible to use the collimated source for phase-contrast imaging, such as in **Paper VIII**. Having a more collimated beam might also be advantageous for a wide-field imaging technique such as the one in **Paper IX**. In that experiment, the imaged sprays were very close to the acceleration chamber, because of the large divergence of the betatron source - moving the studied objects further away might facilitate the experimental conditions in some cases.

Learning from the work on radiation generation with plasma lenses, these lenses could be used also with the specific goal of collimating the electron beams. This has several applications, for example together with the quadrupole magnet transport beamline implemented at the LWFA in Lund. Collimating the beam close to the source could allow for better transport using the conventional magnet optics, leading to improved performance in for example further radiotherapy experiments or beam characterization. The plasma lenses must be matched with respect to the bunch charges, such that the possibility of a beam-driven wake is considered for the transport. The plasma density should most likely be lower than in **Paper VI**, maybe by 2 orders of magnitude, to yield better control over both laser- and beam-driven wakes. Showing matching between an LPA and conventional transport optics using a passive plasma-lens has not been shown yet, to the author's knowledge, and might be an important step for applications of LPA-generated electron beams.

Regarding the tuning of bunch compressor and accelerator performance, as was covered in **Paper I, II, III, IV, and V**, there are multiple ways forward also here. As was shown in **Paper I, II and IV**, the performance of double-achromat dogleg compressors can be improved, in some cases with relatively minor upgrades. However, CSR in particular can still be a problem, as in **Paper IV**, and so further improvements to optically cancel these effects are necessary to investigate, for any high-brightness application. The possibility for on-line, beam-based tuning is extremely important, and should be included in any design. Borrowing from the concept outlined in for example Ref. [101], where an additional, oppositely deflecting sub-compressor with different R_{56} can be added downstream of the first, could be an alternative. The first dogleg would then be designed to overshoot the axis of the downstream section, and a smaller, second dogleg would return it (similar to the bypass line in Ref. [100], but with a shorter second half). This could also increase the total R_{56} , requiring a smaller energy spread for the same degree of compression, which could then alleviate chromatic aberrations as well. A small straight-section between the two compressors could be used for tuning the phase advance on a beam-by-beam basis. Also, making the two achromats of a single such dogleg compressor asymmetric might improve performance with respect to CSR.

In the meantime, lower-charge double-bunches for FEL, as in **Paper V**, could be studied, as the CSR effects are somewhat suppressed in those cases. Some initial double-bunch results were shown in **Paper III**, but further studies are definitely warranted. An interferometer was in fact designed for the purpose of creating continuously controllable double-pulses in the laser system, but remains to be constructed.

The lower-charge bunches can also be placed closer together in time in the PG without severe space-charge degradation, such that the total energy spread of the beam is lower around BC1 and thus more controllable. Both the lower-charge and higher-charge beams would also benefit greatly from stronger PG fields, which can decrease space-charge effects, which recent research indicates could be achieved “simply” by cooling normal-conducting RF structures to cryogenic temperatures, also at S-band [88]. There are improvements even at liquid nitrogen temperatures, which greatly reduces the cost of the cryostats. A cryogenic injector could enable much shorter beams, while still lowering the emittance, which would also serve to reduce chromatic effects down the line for the same compression degree. Such a cryogenic injector could therefore work well together with an optimized compressor layout - it could in principle improve beams for any application.

In summary, there are many interesting future projects to undertake, which build at least partially upon work done within the frame of this thesis. Hopefully, some of the work herein will also be of use to the field in general.

ARTICLE SUMMARIES AND THE AUTHOR'S CONTRIBUTIONS

For all articles, the author has read and reviewed the manuscript.

I Third-order double-achromat bunch compressors for broadband beams

This paper describes a design and optimization strategy for the minimization of chromatic aberrations in dogleg-shaped double-achromat bunch compressors, with specific application to the second bunch compressor in the MAX IV linac. The goal was originally to eliminate higher-order transverse dispersion at the end of the compressor, but the strategy was also found to yield less chromatic aberrations in the slice CS parameters, as well as a decrease in geometric aberrations from a tuned betatron phase advance between non-linear magnets in the achromat centers. The non-linear dispersion could also be tuned to mitigate transverse effects from CSR. The lessons learned in this study were valuable for the study described in **Paper II**.

I performed all simulations and calculations and wrote the manuscript with help and feedback from the co-authors.

II Arc-like variable bunch compressors

The study described in this paper was aimed towards the design of double-achromat, in the paper described as “arc-like”, bunch compressors with a variable R_{56} . The study started from the first bunch compressor in the MAX IV linac, and sought to make the R_{56} variable while keeping the footprint of the compressor, as well as improve on its chromatic performance in the same vein as **Paper I**. Two different routes were explored, where one emerged as clearly superior to the other.

I was involved in the discussions and continuous evaluations during the study and provided feedback on for example strategies for handling the various chromatic aberrations of the compressor optics.

III **Beamline design for plasma-wakefield acceleration experiments at MAX IV**

This paper describes both a status update of the (then) current state of the diagnostics beamline being built for the MAX IV linac, as well as experimental results from the linac with potential PWFA applications. The experimental results were on both single- and double-bunches, where the single-bunch current profiles could be tuned using only the first bunch compressor to yield linear ramps. The double-bunches were a preliminary test of such an operation mode, both for potential PWFA and FEL, by a modification to the pulse-stacking system in the PG laser.

I co-performed the experiments, did the data analysis, partook in part of the diagnostics beamline planning and wrote the manuscript with feedback from the co-authors.

IV **Start-to-end simulations of the MAX IV linear accelerator towards plasma-wakefield acceleration experiments**

The study presented in this paper is a start-to-end simulation study of the MAX IV linac towards the application of PWFA, with the PG, linac and beam-plasma interaction being simulated as three separate blocks, one feeding the particle data to the next. While some relatively straight-forward “upgrades” were made to the magnet layout, such as an improved version of the second bunch compressor as shown in **Paper I**, the parameters of for example the PG and the PG laser system were matched with what is currently available in the machine. Both single-bunch and double-bunch operating modes were simulated, for use with internal and external injection, respectively. As only a relatively small amount of charge was feasible (200 pC in total), low-emittance operation and high compression was chosen to yield a high charge density at the plasma entrance. While the study does show that there clearly is potential for doing this type of experiment with the MAX IV linac, it also elucidates some of the difficulties of using low-emittance bunches for PWFA, and higher-charge (if possible), higher-emittance operation should be advantageous.

I coordinated the study, implemented all machine upgrades in the code, performed all linac particle-tracking simulations and wrote the manuscript with feedback from the co-authors.

V Double-bunches for two-color soft X-ray free-electron laser at the MAX IV laboratory

This paper shows a short study on the acceleration of double-bunches with parameters suitable for soft X-ray FEL at MAX IV. In this study, only the longitudinal phase space distributions were considered, but these distributions were promising as the achievable peak current was relatively high, considering the low-charge operation that was chosen to mitigate CSR-induced emittance growth, and the energy spread of the bunches was relatively low. Apart from the PWFA studies, this application was a motivation behind the preliminary double-bunch experiment shown in **Paper III**.

I performed all tracking simulations and wrote the manuscript with feedback from the co-authors.

VI Low-divergence, femtosecond X-ray pulses from a passive plasma lens

The study presented in this paper shows a novel method for generating X-ray pulses with unprecedented low divergence using a compact, two-stage LWFA approach. The first stage is the injector and accelerator, while the second stage is the radiator, in effect a strong, passive plasma-lens. The two stages are separated by a small distance, allowing the electron bunches and laser pulse to diverge between the stages, which increases their size. When the electron bunch is focused inside the second stage, its expanded size yields a decreased bunch divergence at the location of highest radiated power, such that the radiated X-ray divergence is also decreased. The number of radiated photons per electron is nearly the same as for betatron radiation, such that the angular X-ray brilliance could be significantly higher if using high-charge bunches generated using higher-power lasers. The method is general, more tuneable than betatron radiation, and the experimental implementation is straight-forward, showing potential for being a future method of choice for LWFA X-ray applications.

I co-planned and co-performed the experiment, analyzed the electron spectrometer data, developed the numerical electron beam propagation model and wrote the manuscript with feedback from the co-authors.

VII Highly efficient angularly resolving X-ray spectrometer optimized for absorption measurements with collimated sources

This paper presents the design and implementation of a compact X-ray spectrometer for time-resolved X-ray absorption, applied to the study of copper in the warm dense matter state. The spectrometer was found to have a high efficiency and angular resolution, and it allowed taking shot-by-shot, self-referenced data. The X-ray source in the study was betatron X-rays from the LWFA in Lund.

I assisted in preparing the experimental setup, performing the experiment, and the analysis of electron spectrometer data.

VIII Optimization of soft X-ray phase-contrast tomography using a laser wakefield accelerator

This paper details a study of applying betatron X-rays generated with a LWFA using ionization injection for X-ray phase-contrast tomography of a small insect. The resolution of the final 3D-tomographic image was in the few 10s of μm , and the study also outlined the optimization method and considerations.

I assisted in preparing the experimental setup and the analysis of electron spectrometer data.

IX Simultaneous laser-driven X-ray and two-photon fluorescence imaging of atomizing sprays

The study performed for this paper was a collaboration with the Division of Combustion Physics towards combining the two previously separate characterization methods of X-ray absorption and two-photon fluorescence imaging for atomizing sprays, such as liquid fuel injectors. Yielding complementary information, of the integrated droplet mass and droplet size for X-rays and two-photon fluorescence, respectively, this combined diagnostic provides a more complete set of data than either method on its own. The X-ray source was betatron radiation from the LWFA in Lund. The study also showed the advantage of the divergent betatron X-rays for this applications, which allowed capturing the entire spray in a single image, instead of doing for example a raster scan.

I assisted in preparing the experimental setup and performing the experiment.

X Electron acceleration in merging laser wakefields

In this study, two separate laser pulses, split from a single pulse from the VEGA II laser at CLPU in Salamanca, Spain, were focused into a gas jet with a small focal spot separation and relative angle. This excited two separate wakes, each showing electron injection, whose interaction was studied by varying the temporal delay and angle (affecting the intersection position) between the two pulses. The merging of two separate wakes, which was observed, has potential applications in for example enhancing betatron radiation from LWFA, or as a means of driving wakes in staged LWFA schemes.

I contributed to planning, setting up and performing the experiment, in particular concerning the electron spectrometer.

ACKNOWLEDGEMENTS

In the introduction of this thesis, I mentioned that I'd had the privilege to work with a lot of different people, from a lot of different places. While I might not have worked directly with everyone mentioned below, everyone listed here (and others that I have probably forgotten to mention) is someone who has impacted this thesis, and the work gone into it, in one way or another.

First and foremost, I would like to extend my sincere thanks to my supervisors: Olle Lundh, Claes-Göran Wahlström, Sara Thorin and Francesca Curbis. I am very grateful that all of you have taken me under your respective wing, and under your mentorship I have learned a great deal of things and grown both on a professional and personal level. It would not have been possible to complete this body of work without your guidance. You have always been encouraging and made me feel welcome.

I am also incredibly grateful to the people who have worked with me directly for so many hours on various projects. To Henrik, for always being there (both for me and others) when help, motivation or general discussion and ventilation was needed. I, too, have enjoyed your company, in Lund and on the many excursions we have done together over the years. To Isabel, for always being good-humored, willing to help me realign my misalignments, and sharing my enthusiasm for “very important meetings” in the coffee room. To Diego, for all the time spent together in the salt mines, for teaching me many different things about optics, and for providing regular raclette to keep the spirit up. To Tessa, for not getting tired of my endless questions (or at least not showing it!), for teaching me so much about particle transport, and for all your motivation and huge support, especially now at the end. I hope I can somehow return the favor. To Joel, for your enthusiasm about everything from beer to accelerator projects, it is good to know that I am not the only one who can get carried away with cool physics. To Martin, for your pragmatic approach to things and for helping me get into the topic when I first started, and for being always reliable when “very important meetings” were needed. To Kristoffer, for your good company and admirably relaxed attitude, I do occasionally get jealous of that. To Anders, for always knowing (or figuring out) what's where in the lab, and for keeping the machine running, I would have lost a great deal of hair if not for you. I hope you find your *skrädmjöl* eventually. To Julien, for managing to handle both my many simulation questions/requests and a baby (two babies, then?) at the same time, and for all the good discussions.

To Giada, for sharing my enthusiasm for some things and enthusiastically not sharing it for others. Don't forget to give yourself a break every once in a while, and

I'll try to do the same. Now, about that YouTube channel...

And to all my other colleagues who, even though we might not have worked as extensively on the same projects, have also been important during these past years: Kristoffer, Lovisa, Malay, Alexander and everyone else at the Division of Atomic Physics, for all the conversations and discussions in the lunch room and corridors around the division. To Anne, Åke and Jakob for your help with a great big pile of practical stuff. To the old Maxlab gang: Galina, Jonas, Joel, Teresia, Olivia, Walan, Alan and Christian for coffee breaks and summer lunches behind the old lab when I was very new to everything, and for Martin, Magnus, Pedro and Åke for supervising me on various projects throughout that time. And thanks to Magnus and Francesca for sharing their not-so-gigantic office with me. Also thanks to the mostly-current gang: Joel, Jonas, Francis, David, Weilun, Mihai and Billy, for dragging me out of the office every once in a while, for a lunch or a beer or a few slices of pizza. I hope there will be more such occasions. Thanks to Sara, Erik and the rest of the MAX IV linac team for letting me be part of your club, with breakfast meetings, barbecues and everything else.

To all the collaborators over the years: everyone in PLIONA, the CLPU team, Katerina and Michal, Edouard, all the participants of the 2019-2020 EuPRAXIA experiments here in Lund, the folks at Daresbury and Cockroft, and László. It is always interesting to meet and work with new people, and you have all enriched my time during the work going into this thesis.

To Galina, Ilya and Kuma, for always being hospitable, and for taking Amalia and myself to places we would otherwise not have visited. We would have been many experiences poorer without you. And to Galina, for encouraging me to apply for the PhD position that has finally lead me here.

To Jonas and Rob, for all the cycling adventures and coffee breaks on the countryside which have helped clear my mind when there has been too much on it. It is always a pleasure (even after tens of kilometers when stuff starts to hurt a bit), and I hope we get to ride many more kilometers together. And thanks to Francis and Matilda for keeping an eye on your respective boy when I can't!

A big thanks also to Carl Lindstrøm and Richard D'Arcy, for inviting me to FLASHForward - your encouragement has motivated me throughout the past year, during times when I have very much needed a mental boost.

To all my friends, for New Year's eves, midsummers, trips to France and Latvia, beers and Vietnamese food in Malmö, board game nights, barbecues, shuffleboard, computer gaming, sci-fi conventions, crayfish fishing, concerts and a whole host of other good times with good people, far away from work.

To my family, for their ability to instantly derail dinner conversations, their compassion and honesty, and their endless love and support, both for me and each other. And to my extended family, for accepting me with open arms and letting me be part of their lives too.

To Amalia, for her unconditional love which I often wonder what I do to deserve. I don't need to worry about the future when I know that you'll be there to keep me safe.

I don't think that I appreciate often enough that I am in fact, in many regards, very, very lucky.

REFERENCES

1. J. D. Cockcroft and E. T. S. Walton. *Experiments with high velocity positive ions*. Proc. R. Soc. Lond. A **129**, 477–489 (1930).
2. R. Widerøe. *Über ein neues prinzip zur herstellung hoher spannungen*. Arch. Elektrotech. **21**, 387–406 (1928).
3. E. O. Lawrence and N. E. Edlefsen. *On the production of high speed protons*. Science **72**, 376 (1930).
4. CERN. LHC Homepage (2020-08-13).
<https://home.cern/science/accelerators/large-hadron-collider>.
5. SLAC. LCLS Homepage (2020-08-13).
<https://lcls.slac.stanford.edu/>.
6. V. Shiltsev. *Particle beams behind physics discoveries*. Phys. Today **73**, 32–39 (2020).
7. D. H. Bilderback, P. Elleaume and E. Weckert. *Review of third and next generation synchrotron light sources*. J. Phys. B **38**, S773–S797 (2005).
8. P. Schmüser, M. Dohlus, J. Rossbach and C. Behrens. *Free-Electron Lasers in the Ultraviolet and X-Ray Regime*. Springer International Publishing, Switzerland, 2nd edition (2014). ISBN 978-3-319-04080-6.
9. C. Pellegrini, A. Marinelli and S. Reiche. *The physics of x-ray free-electron lasers*. Rev. Mod. Phys. **88**, 015006 (2016).
10. T. K. Charles. *Synchrotrons on the coronavirus frontline*. Cern Courier (2020).
11. M. Aicheler, P. Burrows, M. Draper, T. Garvey et al. *A Multi-TeV Linear Collider based on CLIC Technology: CLIC Conceptual Design Report*. Technical Report CERN-2012-007 CERN (2012).
12. T. Behnke, J. E. Brau, B. Foster, J. Fuster et al. *The International Linear Collider Technical Design Report*. Technical report CERN (2013).
13. M. Benedikt, A. Blondel, O. Brunner, M. Capeans Garrido et al. *FCC-ee: The Lepton Collider: Future Circular Collider Conceptual Design Report Volume 2. Future Circular Collider*. Technical Report CERN-ACC-2018-0057. 2 CERN Geneva (2018).

14. European Strategy Group. 2020 Update of the European Strategy for Particle Physics. Technical Report CERN-ESU-015 Geneva (2020).
15. M. Tzoufras, W. Lu, F. S. Tsung, C. Huang et al. *Beam loading in the nonlinear regime of plasma-based acceleration*. Phys. Rev. Lett. **101**, 145002 (2008).
16. G. G. Manahan, A. F. Habib, P. Scherkl, P. Delinikolas et al. *Single-stage plasma-based correlated energy spread compensation for ultrahigh 6D brightness electron beams*. Nature Communications **8**, 15705 (2017).
17. C. A. Lindström and E. Adli. *Design of general apochromatic drift-quadrupole beam lines*. Phys. Rev. Accel. Beams **19**, 071002 (2016).
18. T. Liu, T. Zhang, D. Wang and Z. Huang. *Compact beam transport system for free-electron lasers driven by a laser plasma accelerator*. Phys. Rev. Accel. Beams **20**, 020701 (2017).
19. M. E. Couprie, M. Labat, C. Evain, F. Marteau et al. *An application of laser-plasma acceleration: towards a free-electron laser amplification*. Plasma Phys. Control. Fusion **58**, 034020 (2016).
20. A. Bernhard, V. Afonso Rodríguez, S. Kuschel, M. Leier et al. *Progress on experiments towards LWFA-driven transverse gradient undulator-based FELs*. Nucl. Instrum. Methods Phys. Res. A **909**, 391–397 (2018).
21. J. D. Jackson. *Classical Electrodynamics*. Wiley, New York, NY, USA 3rd edition (1999). ISBN 978-047-1309-32-1.
22. K. Wille. *The Physics of Particle Accelerators: An Introduction*. Oxford University Press, Oxford, UK (2000). ISBN 0-19-850549-3.
23. D. J. Griffiths. *Introduction to Electrodynamics*. Pearson Education, Inc., San Francisco, CA, USA, 3rd edition (2008). ISBN 0-13-919960-8.
24. H. Goldstein, C. P. Poole and J. L. Safko. *Classical Mechanics*. Pearson Education Ltd., Harlow, UK, 3rd edition (2014). ISBN 978-1-292-02655-8. Pearson New International Edition.
25. S. Y. Lee. *Accelerator Physics*. World Scientific Publishing Co. Pte. Ltd., Singapore, 3rd edition (2012). ISBN 978-981-4374-94-1.
26. K. L. Brown. A first- and second-order matrix theory for the design of beam transport systems and charged particle spectrometers. Report SLAC No. 75 SLAC (1982).
27. D. C. Carey, K. L. Brown and F. Rothacker. Third-order TRANSPORT: A computer program for designing charged particle beam transport systems. Report SLAC-R-95-462 FERMILAB, SLAC (1995).
28. Wikipedia. Moment (mathematics) (2020-08-13). [https://en.wikipedia.org/wiki/Moment_\(mathematics\)](https://en.wikipedia.org/wiki/Moment_(mathematics)).

-
29. H. Wiedemann. *Particle Accelerator Physics*. Springer-Verlag, Berlin, Germany, 3rd edition (2007). ISBN 978-3-540-49045-6.
 30. K. Floettmann. *Some basic features of the beam emittance*. Phys. Rev. ST Accel. Beams **6**, 034202 (2003).
 31. M. Migliorati, A. Bacci, C. Benedetti, E. Chiadroni et al. *Intrinsic normalized emittance growth in laser-driven electron accelerators*. Phys. Rev. ST Accel. Beams **16**, 011302 (2013).
 32. K. Floettmann. *Adiabatic matching section for plasma accelerated beams*. Phys. Rev. ST Accel. Beams **17**, 054402 (2014).
 33. K. A. Marsh, C. E. Clayton, D. K. Johnson, C. Huang et al. *Beam matching to a plasma wake field accelerator using a ramped density profile at the plasma boundary*. Proc. PAC'05, Knoxville, TN, USA (2005).
 34. WolframAlpha Homepage (2020-08-13).
<https://www.wolframalpha.com/>.
 35. S. Di Mitri and S. Spampinati. *Estimate of free electron laser gain length in the presence of electron beam collective effects*. Phys. Rev. ST Accel. Beams **17**, 110702 (2014).
 36. B. E. A. Saleh and M. C. Teich. *Fundamentals of Photonics*. Wiley Series in Pure and Applied Optics. John Wiley & Sons, Inc., Hoboken, NJ, USA, 2nd edition (2007). ISBN 978-0-471-35832-9.
 37. R. Lehe. *Improvement of Laser-Wakefield Accelerators: Towards a Compact Free Electron Laser*. PhD thesis Palaiseau, Lab. Opt. Appl. (2014).
 38. Wikipedia. Wave vector (2020-08-13).
https://en.wikipedia.org/wiki/Wave_vector.
 39. E. Esarey, P. Sprangle, J. Krall and A. Ting. *Self-focusing and guiding of short laser pulses in ionizing gases and plasmas*. IEEE J. Quantum Electron. **33**, 1879–1914 (1997).
 40. D. A. Jaroszynski, R. A. Bingham and R. A. Cairns, editors. *Laser-Plasma Interactions*. CRC Press (2017). ISBN 9781138114364.
 41. Wikipedia. Self-focusing (2020-08-13).
<https://en.wikipedia.org/wiki/Self-focusing>.
 42. Wikipedia. Group Velocity (2020-08-13).
https://en.wikipedia.org/wiki/Group_velocity.
 43. Wikipedia. Group Velocity Dispersion (2020-08-13).
https://en.wikipedia.org/wiki/Group_velocity_dispersion.
 44. Wikipedia. Self-phase modulation (2020-08-13).
https://en.wikipedia.org/wiki/Self-phase_modulation.

45. Wikipedia. State of matter (2020-05-13).
https://en.wikipedia.org/wiki/State_of_matter.
46. P. Gibbon. *Short Pulse Laser Interactions with Matter*. Imperial College Press (2005). ISBN 1-86094-135-4.
47. Wikipedia. Ionization energies of the elements (data page) (2020-08-13).
[https://en.wikipedia.org/wiki/Ionization_energies_of_the_elements_\(data_page\)](https://en.wikipedia.org/wiki/Ionization_energies_of_the_elements_(data_page)).
48. NIST. Proton-electron mass ratio (2020-08-13).
<https://physics.nist.gov/cgi-bin/cuu/Value?mpsme>.
49. Wikipedia. Waves in plasmas (2020-08-13).
https://en.wikipedia.org/wiki/Waves_in_plasmas.
50. Wikipedia. Plasma oscillation (2020-08-13).
https://en.wikipedia.org/wiki/Plasma_oscillation.
51. E. Esarey, P. Sprangle, J. Krall and A. Ting. *Overview of plasma-based accelerator concepts*. IEEE Trans. Plasma Sci. **24**, 252–288 (1996).
52. G. Cantono. *Relativistic Plasmonics for Ultra-Short Radiation Sources*. PhD thesis (2017).
53. A. Karlsson and G. Kristensson. *Microwave Theory*. Department of Electrical and Information Technology, Faculty of Engineering, LTH, Lund University (2015).
54. E. Esarey, C. B. Schroeder and W. P. Leemans. *Physics of laser-driven plasma-based electron accelerators*. Rev. Mod. Phys. **81**, 1229–1285 (2009).
55. W. Lu, C. Huang, M. Zhou, M. Tzoufras et al. *A nonlinear theory for multidimensional relativistic plasma wave wakefields*. Phys. Plasmas **13**, 056709 (2006).
56. W. Lu, M. Tzoufras, C. Joshi, F. S. Tsung et al. *Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime*. Phys. Rev. ST Accel. Beams **10**, 061301 (2007).
57. H. Ekerfelt. *Numerical and Experimental Studies of Wakefield Accelerators*. PhD thesis Lund University (2019).
58. H. Ding, A. Döpp, M. Gilljohann, J. Götzfried et al. *Nonlinear plasma wavelength scalings in a laser wakefield accelerator*. Phys. Rev. E **101**, 023209 (2020).
59. A. Caldwell, K. Lotov, A. Pukhov and F. Simon. *Proton-driven plasma-wakefield acceleration*. Nat. Phys. **5**, 363–367 (2009).
60. A. Caldwell, E. Adli, L. Amorim, R. Apsimon et al. *Path to AWAKE: Evolution of the concept*. Nucl. Instrum. Methods Phys. Res. A **829**, 3–16 (2016).
61. E. Adli, A. Ahuja, O. Apsimon, R. Apsimon et al. *Acceleration of electrons in the plasma wakefield of a proton bunch*. Nature **561**, 363–367 (2018).

-
62. M. Turner, P. Muggli, E. Adli, R. Agnello et al. *Experimental study of wakefields driven by a self-modulating proton bunch in plasma*. Phys. Rev. Accel. Beams **23**, 081302 (2020).
 63. B. E. Blue, C. E. Clayton, C. L. O'Connell, F.-J. Decker et al. *Plasma-wakefield acceleration of an intense positron beam*. Phys. Rev. Lett. **90**, 214801 (2003).
 64. S. Corde, E. Adli, J. M. Allen, W. An et al. *Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield*. Nature **524**, 442–445 (2015).
 65. S. Gessner, E. Adli, J. M. Allen, W. An et al. *Demonstration of a positron beam-driven hollow channel plasma wakefield accelerator*. Nat. Commun. **7**, 11785 (2016).
 66. N. Barov, J. B. Rosenzweig, M. C. Thompson and R. B. Yoder. *Energy loss of a high-charge bunched electron beam in plasma: Analysis*. Phys. Rev. ST Accel. Beams **7**, 061301 (2004).
 67. B. Hidding, A. Beaton, L. Boulton, S. Corde et al. *Fundamentals and applications of hybrid LWFA-PWFA*. Appl. Sci. **9**, 2626 (2019).
 68. M. Tzoufras, W. Lu, F. S. Tsung, C. Huang et al. *Beam loading by electrons in nonlinear plasma wakes*. Phys. Plasmas **16**, 056705 (2009).
 69. J. B. Rosenzweig, B. Breizman, T. Katsouleas and J. J. Su. *Acceleration and focusing of electrons in two-dimensional nonlinear plasma wake fields*. Phys. Rev. A **44**, R6189–R6192 (1991).
 70. K. V. Lotov. *Efficient operating mode of the plasma wakefield accelerator*. Phys. Plasmas **12** (2005).
 71. Pisin Chen, J. J. Su, J. M. Dawson, K. L. F. Bane and P. B. Wilson. *Energy transfer in the plasma wake-field accelerator*. Phys. Rev. Lett. **56**, 1252–1255 (1986).
 72. H. Xiao, M. Chen, S. M. Weng, L. M. Chen et al. *Control of transverse motion and x-ray emission of electrons accelerated in laser-driven wakefields by tuning laser spatial chirp*. Plasma Phys. Control. Fusion **62**, 024002 (2019).
 73. D. H. Whittum, W. M. Sharp, S. S. Yu, M. Lampe and G. Joyce. *Electron-hose instability in the ion-focused regime*. Phys. Rev. Lett. **67**, 991–994 (1991).
 74. S. Akturk, X. Gu, E. Zeek and R. Trebino. *Pulse-front tilt caused by spatial and temporal chirp*. Opt. Express **12**, 4399–4410 (2004).
 75. T. J. Mehrling, R. A. Fonseca, A. Martinez de la Ossa and J. Vieira. *Mitigation of the hose instability in plasma-wakefield accelerators*. Phys. Rev. Lett. **118**, 174801 (2017).
 76. A. Martinez de la Ossa, T. J. Mehrling and J. Osterhoff. *Intrinsic stabilization of the drive beam in plasma-wakefield accelerators*. Phys. Rev. Lett. **121**, 064803 (2018).

77. E. Gschwendtner, E. Adli, L. Amorim, R. Apsimon et al. *AWAKE, the advanced proton driven plasma wakefield acceleration experiment at CERN*. Nucl. Instrum. Methods Phys. Res. A **829**, 76–82 (2016).
78. Microwaves 101. Frequency Letter Bands (2020-08-13).
<https://www.microwaves101.com/encyclopedias/frequency-letter-bands>.
79. P. Schmüser. Basic principles of RF superconductivity and superconducting cavities. Technical report Universität Hamburg, Institut für Experimentalphysik (2006).
<http://srf.desy.de/fap/paper/MoT01.pdf>.
80. Wikipedia. Superconducting radio frequency (2020-08-13).
https://en.wikipedia.org/wiki/Superconducting_radio_frequency.
81. SLAC. LCLS-II: A world-class discovery machine (2020-08-13).
<https://lcls.slac.stanford.edu/lcls-ii>.
82. J. N. Galayda. *The LCLS-II: a high power upgrade to the LCLS*. Proc. IPAC'18, Vancouver, BC, Canada (2018).
83. W. Decking, S. Abeghyan, P. Abramian, A. Abramsky et al. *A MHz-repetition-rate hard X-ray free-electron laser driven by a superconducting linear accelerator*. Nat. Photon. **14**, 391–397 (2020).
84. Y. Y. Lau and D. G. Colombant. *A general theory of beam breakup instabilities in linear accelerators*. AIP Conf. Proc. **193**, 244–273 (1989).
85. A. D. Cahill, J. B. Rosenzweig, V. A. Dolgashev, Z. Li et al. *rf losses in a high gradient cryogenic copper cavity*. Phys. Rev. Accel. Beams **21**, 061301 (2018).
86. A. D. Cahill, J. B. Rosenzweig, V. A. Dolgashev, S. G. Tantawi and S. Weathersby. *High gradient experiments with X-band cryogenic copper accelerating cavities*. Phys. Rev. Accel. Beams **21**, 102002 (2018).
87. K. L. Bane, T. L. Barklow, M. Breidenbach, C. P. Burkhart et al. *An advanced NCRF linac concept for a high energy e^+e^- linear collider*. arXiv preprint arXiv:1807.10195v2 (2019).
88. J. B. Rosenzweig, A. Cahill, V. Dolgashev, C. Emma et al. *Next generation high brightness electron beams from ultrahigh field cryogenic rf photocathode sources*. Phys. Rev. Accel. Beams **22**, 023403 (2019).
89. J. B. Rosenzweig, N. Majernik, R. R. Robles, G. Andonian et al. *An ultra-compact X-ray free-electron laser*. arXiv preprint arXiv:2003.06083v2 (2020).
90. M. Nasr. *Cryogenic-copper accelerating structures: new frontier for beam brightness, efficiency and cost-capability*. In *Online presentation at IPAC'20* (2020).
<https://www.ipac20.org/mandouh-nasr-slac/>.
91. P. B. Wilson. Introduction to wakefields and wake potentials. resreport PUB-4547 SLAC (1989).

-
92. T. P. Wrangler. *RF Linear Accelerators*. Wiley-VCH Verlag GmbH, Weinheim, Germany, 2nd edition (2008). ISBN 978-3-527-40680-7.
 93. L. Serafini and J. B. Rosenzweig. *Envelope analysis of intense relativistic quasilinear beams in rf photoinjectors: a theory of emittance compensation*. Phys. Rev. E **55**, 7565–7590 (1997).
 94. S. G. Anderson, P. Musumeci, J. B. Rosenzweig, W. J. Brown et al. *Velocity bunching of high-brightness electron beams*. Phys. Rev. ST Accel. Beams **8**, 014401 (2005).
 95. J. R. Thompson. *Bunch compression for a short-pulse mode in Cornell's ERL* (2009).
<https://www.classe.cornell.edu/~hoff/hoff/papers/09ThompsonHonorsThesis.pdf>.
 96. A. Sharma. *Optics Design and Optimization of Electron Bunch Compressor Transfer Line*. PhD thesis Indore, Raja Ramanna Centre for Advanced Technology (2013).
 97. M. Schreck and P. Wesolowski. *Analytical bunch compression studies for a linac-based electron accelerator*. Phys. Rev. ST Accel. Beams **18**, 100101 (2015).
 98. D. Alesini, G. Castorina, M. Croia, M. Ferrario et al. *Design of a full C-band injector for ultra-high brightness electron beam*. Proc. IPAC'19, Melbourne, Australia (2019).
 99. S. Thorin, M. Eriksson, S. Werin, D. Angal-Kalinin et al. *Bunch compression by linearising achromats for the MAX IV injector*. Proc. FEL'10, Malmö, Sweden (2010).
 100. Y. Sun. *Second-order achromat design based on FODO cell*. Phys. Rev. ST Accel. Beams **14**, 060703 (2011).
 101. Y. Jing, Y. Hao and V. N. Litvinenko. *Compensating effect of the coherent synchrotron radiation in bunch compressors*. Phys. Rev. ST Accel. Beams **16**, 060704 (2013).
 102. S. Russenschuck. *Field computation for accelerator magnets: Analytical and numerical methods for electromagnetic design and optimization*. Wiley-VCH Weinheim (2010). ISBN 9783527407699.
 103. K. Halbach. *Application of permanent magnets in accelerators and electron storage rings*. J. Appl. Phys. **57**, 3605–3608 (1985).
 104. Wikipedia. Halbach array (2020-08-13).
https://en.wikipedia.org/wiki/Halbach_array.
 105. P. Raimondi and A. Seryi. *Novel final focus design for future linear colliders*. Phys. Rev. Lett. **86**, 3779–3782 (2001).

106. R. Brinkmann, P. Raimondi and A. Seryi. *Halo reduction by means of non linear optical elements in the NLC final focus system*. Proc. PAC'01, Chicago, IL, USA (2001).
107. R. V. Servranckx and K. L. Brown. *Circular machine design techniques and tools*. Nucl. Instrum. Methods Phys. Res. A **258**, 525–535 (1987).
108. A. Ghaith, D. Oumbarek, C. Kitégi, M. Valléau et al. *Permanent magnet-based quadrupoles for plasma acceleration sources*. Instruments **3**, 27 (2019).
109. M. W. Guetg, B. Beutner, E. Prat and S. Reiche. *Optimization of free electron laser performance by dispersion-based beam-tilt correction*. Phys. Rev. ST Accel. Beams **18**, 030701 (2015).
110. E. L. Saldin, E. A. Schneidmiller and M. V. Yurkov. *On the coherent radiation of an electron bunch moving in an arc of a circle*. Nucl. Instrum. Methods Phys. Res. A **398**, 373–394 (1997).
111. G. V. Stupakov. *Effect of centrifugal transverse wakefield for microbunch in bend*. resreport SLAC-PUB 8028 SLAC (2006).
112. S. Di Mitri, M. Cornacchia and S. Spampinati. *Cancellation of coherent synchrotron radiation kicks with optics balance*. Phys. Rev. Lett. , 014801 (2013).
113. S. Thorin, J. Andersson, F. Curbis, M. Eriksson et al. *The MAX IV linac*. Proc. LINAC'14, Geneva, Switzerland (2014).
114. S. C. Leemann, Å. Andersson, M. Eriksson, L.-J. Lindgren et al. *Beam dynamics and expected performance of Sweden's new storage-ring light source: MAX IV*. Phys. Rev. ST Accel. Beams **12**, 120701 (2009).
115. S. Werin, S. Thorin, M. Eriksson and J. Larsson. *Short pulse facility for MAX-lab*. Nucl. Instrum. Methods Phys. Res. A **601**, 98 (2009).
116. S. Werin, J. Andersen, P. Fernandes Tavares, L. Isaksson et al. *The soft X-ray laser project at MAX IV*. Proc. IPAC'17, Copenhagen, Denmark (2017).
117. J. Andersson, F. Curbis, D. Kumbaro, F. Lindau and S. Werin. *Beam performance of the photocathode gun for the MAX IV linac*. Proc. FEL'14, Basel, Switzerland (2014).
118. F. Lindau, J. Andersson, J. Björklund Svensson, M. Brandin et al. *MAX IV photocathode gun laser system specification and diagnostics*. Proc. IPAC'17, Copenhagen, Denmark (2017).
119. M. Kotur, J. Andersson, J. Björklund Svensson, M. Brandin et al. *Pulse shaping at the MAX IV photoelectron gun laser*. Proc. IPAC'17, Copenhagen, Denmark (2017).
120. FemtoMAX (2020-08-13).
<https://www.maxiv.lu.se/accelerators-beamlines/beamlines/femtomax/>.

-
121. I. Blumenfeld, C. E. Clayton, F.-J. Decker, M. J. Hogan et al. *Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator*. *Nature* **445**, 741–744 (2007).
 122. A. J. Gonsalves, K. Nakamura, J. Daniels, C. Benedetti et al. *Petawatt laser guiding and electron beam acceleration to 8 GeV in a laser-heated capillary discharge waveguide*. *Phys. Rev. Lett.* **122**, 084801 (2019).
 123. A. Debus, R. Pausch, A. Huebl, K. Steiniger et al. *Circumventing the dephasing and depletion limits of laser-wakefield acceleration*. *Phys. Rev. X* **9**, 031044 (2019).
 124. C. Caizergues, S. Smartsev, V. Malka and C. Thauray. *Phase-locked laser-wakefield electron acceleration*. *Nat. Photon.* **14**, 475–479 (2020).
 125. J. P. Palastro, J. L. Shaw, P. Franke, D. Ramsey et al. *Dephasingless laser wakefield acceleration*. *Phys. Rev. Lett.* **124**, 134802 (2020).
 126. R. J. Shalloo and S. P. D. Mangles. *Faster laser pulses boost plasma accelerators*. *Nat. Photon.* **14**, 470–471 (2020).
 127. E. Esarey, B. A. Shadwick, P. Catravas and W. P. Leemans. *Synchrotron radiation from electron beams in plasma-focusing channels*. *Phys. Rev. E* **65**, 056505 (2002).
 128. S. Corde, K. Ta Phuoc, G. Lambert, R. Fitour et al. *Femtosecond x rays from laser-plasma accelerators*. *Rev. Mod. Phys.* **85**, 1–48 (2013).
 129. R. P. Walker. *Undulator radiation brightness and coherence near the diffraction limit*. *Phys. Rev. Accel. Beams* **22**, 050704 (2019).
 130. J. L. Shaw, N. Lemos, L. D. Amorim, N. Vafaei-Najafabadi et al. *Role of direct laser acceleration of electrons in a laser wakefield accelerator with Ionization Injection*. *Phys. Rev. Lett.* **118**, 064801 (2017).
 131. I. Gallardo González, H. Ekerfelt, M. Hansson, T. L. Audet et al. *Effects of the dopant concentration in laser wakefield and direct laser acceleration of electrons*. *New J. Phys.* **20**, 053011 (2018).
 132. A. Aschikhin, C. Behrens, S. Bohlen, J. Dale et al. *The FLASHForward facility at DESY*. *Nucl. Instrum. Methods Phys. Res. A* **806**, 175–183 (2016).
 133. V. Yakimenko, L. Alsberg, E. Bong, G. Bouchard et al. *FACET-II facility for advanced accelerator experimental tests*. *Phys. Rev. Accel. Beams* **22**, 101301 (2019).
 134. M. J. Hogan, T. O. Raubenheimer, A. Seryi, P. Muggli et al. *Plasma wakefield acceleration experiments at FACET*. *New J. Phys.* **12**, 055030 (2010).
 135. C. Joshi, E. Adli, W. An, C. E. Clayton et al. *Plasma wakefield acceleration experiments at FACET II*. *Plasma Phys. Control. Fusion* **60**, 034001 (2018).

136. M. F. Gilljohann, H. Ding, A. Döpp, J. Götzfried et al. *Direct Observation of Plasma Waves and Dynamics Induced by Laser-Accelerated Electron Beams*. Phys. Rev. X **9**, 011046 (2019).
137. S. Lorenz, G. Grittani, E. Chacon-Golcher, C. M. Lazzarini et al. *Characterization of supersonic and subsonic gas targets for laser wakefield electron acceleration experiments*. Matter Radiat. Extremes **4**, 015401 (2019).
138. Wikipedia. de Laval nozzle (2020-08-12). https://en.wikipedia.org/wiki/De_Laval_nozzle.
139. M. C. Downer, R. Zgad Zaj, A. Debus, U. Schramm and M. C. Kaluza. *Diagnostics for plasma-based electron accelerators*. Rev. Mod. Phys. **90**, 035002 (2018).
140. S. Bulanov, N. Naumova, F. Pegoraro and J. Sakai. *Particle injection into the wave acceleration phase due to nonlinear wake wave breaking*. Phys. Rev. E **58**, R5257–R5260 (1998).
141. S. Kalmykov, S. A. Yi, V. Khudik and G. Shvets. *Electron self-injection and trapping into an evolving plasma bubble*. Phys. Rev. Lett. **103**, 135004 (2009).
142. I. Kostyukov, E. Nerush, A. Pukhov and V. Seredov. *A multidimensional theory for electron trapping by a plasma wake generated in the bubble regime*. New J. Phys. **12**, 045009 (2010).
143. S. A. Yi, V. Khudik, S. Y. Kalmykov and G. Shvets. *Hamiltonian analysis of electron self-injection and acceleration into an evolving plasma bubble*. Plasma Phys. Control. Fusion **53**, 014012 (2010).
144. S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas et al. *Monoenergetic beams of relativistic electrons from intense laser-plasma interactions*. Nature **431**, 535–538 (2004).
145. C. G. R. Geddes, Cs. Tóth, J. van Tilborg, E. Esarey et al. *High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding*. Nature **431**, 538–431 (2004).
146. J. Faure, Y. Glinec, A. Pukhov, S. Kiselev et al. *A laser-plasma accelerator producing monoenergetic electron beams*. Nature **431**, 541–544 (2004).
147. S. Corde, C. Thaury, A. Lifschitz, G. Lambert et al. *Observation of longitudinal and transverse self-injections in laser-plasma accelerators*. Nat. Commun. **4**, 1501 (2013).
148. T. N. Dalichaouch, X. L. Xu, F. Li, A. Tableman et al. *Generating high quality ultrarelativistic electron beams using an evolving electron beam driver*. Phys. Rev. Accel. Beams **23**, 021304 (2020).
149. C. G. R. Geddes, K. Nakamura, G. R. Plateau, Cs. Tóth et al. *Plasma-density-gradient injection of low absolute-momentum-spread electron bunches*. Phys. Rev. Lett. **100**, 215004 (2008).

-
150. J. Faure, C. Rechatin, O. Lundh, L. Ammoura and V. Malka. *Injection and acceleration of quasimonoeenergetic relativistic electron beams using density gradients at the edges of a plasma channel*. Phys. Plasmas **17**, 083107 (2010).
151. M. Hansson, B. Aurand, X. Davoine, H. Ekerfelt et al. *Down-ramp injection and independently controlled acceleration of electrons in a tailored laser wakefield accelerator*. Phys. Rev. ST Accel. Beams **18**, 071303 (2015).
152. A. Martinez de la Ossa, Z. Hu, M. J. V. Streeter, T. J. Mehrling et al. *Optimizing density down-ramp injection for beam-driven plasma wakefield accelerators*. Phys. Rev. Accel. Beams **20**, 091301 (2017).
153. H. Ekerfelt, M. Hansson, I. Gallardo González, X. Davoine and O. Lundh. *A tunable electron beam source using trapping of electrons in a density down-ramp in laser wakefield acceleration*. Sci. Rep. **7**, 12229 (2017).
154. H. Suk, N. Barov, J. B. Rosenzweig and E. Esarey. *Plasma electron trapping and acceleration in a plasma wake field using a density transition*. Phys. Rev. Lett. **86**, 1011–1014 (2001).
155. K. Schmid, A. Buck, C. M. S. Sears, J. M. Mikhailova et al. *Density-transition based electron injector for laser driven wakefield accelerators*. Phys. Rev. ST Accel. Beams **13**, 091301 (2010).
156. A. Buck, J. Wenz, J. Xu, K. Khrennikov et al. *Shock-front injector for high-quality laser-plasma acceleration*. Phys. Rev. Lett. **110**, 185006 (2013).
157. M. Chen, Z.-M. Sheng, Y.-Y. Ma and J. Zhang. *Electron injection and trapping in a laser wakefield by field ionization to high-charge states of gases*. J. Appl. Phys. **99**, 056109 (2006).
158. C. McGuffey, A. G. R. Thomas, W. Schumaker, T. Matsuoaka et al. *Ionization induced trapping in a laser wakefield accelerator*. Phys. Rev. Lett. **104**, 025004 (2010).
159. J. P. Couperus, R. Pausch, A. Köhler, O. Zarini et al. *Demonstration of a beam loaded nanocoulomb-class laser wakefield accelerator*. Nat. Commun. **8**, 487 (2017).
160. S. K. Barber, J. van Tilborg, C. B. Schroeder, R. Lehe et al. *Measured emittance dependence on the injection method in laser plasma accelerators*. Phys. Rev. Lett. **119**, 104801 (2017).
161. C. Thaury, E. Guillaume, A. Lifschitz, K. Ta Phuoc et al. *Shock assisted ionization injection in laser-plasma accelerators*. Sci. Rep. **5** (2015).
162. B. Hidding, T. Königstein, J. Osterholz, S. Karsch et al. *Monoenergetic energy doubling in a hybrid laser-plasma wakefield accelerator*. Phys. Rev. Lett. **104**, 195002 (2010).
163. B. Hidding, G. Pretzler, J. B. Rosenzweig, T. Königstein et al. *Ultracold electron bunch generation via plasma photocathode emission and acceleration in a beam-driven plasma blowout*. Phys. Rev. Lett. **108**, 035001 (2012).

164. A. Deng, O. S. Karger, T. Heinemann, A. Knetsch et al. *Generation and acceleration of electron bunches from a plasma photocathode*. Nat. Phys. **15**, 1156–1160 (2019).
165. E. Esarey, R. F. Hubbard, W. P. Leemans, A. Ting and P. Sprangle. *Electron injection into plasma wakefields by colliding laser pulses*. Phys. Rev. Lett. **79**, 2682–2685 (1997).
166. J. Faure, C. Rechatin, A. Norlin, A. Lifschitz et al. *Controlled injection and acceleration of electrons in plasma wakefields by colliding laser pulses*. Nature **444**, 737–739 (2006).
167. M. Hansson, B. Aurand, H. Ekerfelt, A. Persson and O. Lundh. *Injection of electrons by colliding laser pulses in a laser wakefield accelerator*. Nucl. Instrum. Methods Phys. Res. A **829**, 99–103 (2016).
168. J. Ferri, S. Corde, A. Döpp, A. Lifschitz et al. *High-brilliance betatron γ -ray source powered by laser-accelerated electrons*. Phys. Rev. Lett. **120**, 254802 (2018).
169. X.-L. Zhu, M. Chen, S.-M. Weng, T.-P. Yu et al. *Extremely brilliant GeV γ -rays from a two-stage laser-plasma accelerator*. Sci. Adv. **6**, eaaz7240 (2020).
170. R. D’Arcy, S. Wesch, A. Aschikhin, S. Bohlen et al. *Tunable plasma-based energy dechirper*. Phys. Rev. Lett. **122**, 034801 (2019).
171. A. Ferran Pousa, A. Martinez de la Ossa, R. Brinkmann and R. W. Assmann. *Compact multistage plasma-based accelerator design for correlated energy spread compensation*. Phys. Rev. Lett. **123**, 054801 (2019).
172. C. B. Schroeder, E. Esarey, C. Benedetti and W. P. Leemans. *Efficiency considerations for high-energy physics applications of laser-plasma accelerators*. AIP Conference Proceedings **1777**, 020001 (2016).
173. O. Lundh, J. Lim, C. Rechatin, L. Ammoura et al. *Few femtosecond, few kilo-ampere electron bunch produced by a laser-plasma accelerator*. Nat. Phys. **7**, 219–222 (2011).
174. A. Buck, M. Nicolai, K. Schmid, C. M. S. Sears et al. *Real-time observation of laser-driven electron acceleration*. Nat. Phys. **7**, 543–548 (2011).
175. R. Ariniello, C. E. Doss, K. Hunt-Stone, J. R. Cary et al. *Transverse beam dynamics in a plasma density ramp*. Phys. Rev. Acc. Beams **22**, 041304 (2019).
176. T. Mehrling, J. Grebenyuk, F. S. Tsung, K. Floettmann and J. Osterhoff. *Transverse emittance growth in staged laser-wakefield acceleration*. Phys. Rev. ST Accel. Beams **15**, 111303 (2012).
177. X. L. Xu, J. F. Hua, Y. P. Wu, C. J. Zhang et al. *Physics of phase space matching for staging plasma and traditional accelerator components using longitudinally tailored plasma profiles*. Phys. Rev. Lett. **116**, 124801 (2016).

-
178. F. Marteau, A. Ghaith, P. N'Gotta, C. Benabderrahmane et al. *Variable high gradient permanent magnet quadrupole (QUAPEVA)*. Appl. Phys. Lett. **111**, 253503 (2017).
179. C. Benabderrahmane. *Review of permanent magnet technology for accelerators*. Proc. IPAC'17, Copenhagen, Denmark (2017).
180. M. D. Litos, R. Ariniello, C. E. Doss, K. Hunt-Stone and J. R. Cary. *Beam emittance preservation using Gaussian density ramps in a beam-driven plasma wakefield accelerator*. Philos. Trans. R. Soc. A **377**, 20180181 (2019).
181. Y. Zhao, W. An, X. Xu, F. Li et al. *Emittance preservation through density ramp matching sections in a plasma wakefield accelerator*. Phys. Rev. Accel. Beams **23**, 011302 (2020).
182. J. van Tilborg, S. Steinke, C. G. R. Geddes, N. H. Matlis et al. *Active plasma lensing for relativistic laser-plasma-accelerated electron beams*. Phys. Rev. Lett. **115**, 184802 (2015).
183. C. A. Lindstrøm, E. Adli, G. Boyle, R. Corsini et al. *Emittance preservation in an aberration-free active plasma lens*. Phys. Rev. Lett. **121**, 194801 (2018).
184. R. Lehe, C. Thauray, E. Guillaume, A. Lifschitz and V. Malka. *Laser-plasma lens for laser-wakefield accelerators*. Phys. Rev. ST Accel. Beams **17**, 121301 (2014).
185. C. Thauray, E. Guillaume, A. Döpp, R. Lehe et al. *Demonstration of relativistic electron beam focusing by a laser-plasma lens*. Nat. Comm. **6**, 6860 (2015).
186. R. Lehe, C. Thauray, A. Lifschitz, J.-M. Rax and V. Malka. *Transverse dynamics of an intense electron bunch traveling through a pre-ionized plasma*. Phys. Plasmas **21**, 043104 (2014).
187. C. E. Doss, E. Adli, R. Ariniello, J. Cary et al. *Laser-ionized, beam-driven, underdense, passive thin plasma lens*. Phys. Rev. Accel. Beams **22**, 111001 (2019).
188. A. E. Dyson, C. Thornton and Hooker S. M. *A compact, low cost Marx bank for generating capillary discharge plasmas*. Rev. Sci. Instrum. **87**, 093302 (2016).
189. D. Strickland and G. Mourou. *Compression of amplified chirped optical pulses*. Optics Communications **56**, 219–221 (1985).
190. Wikipedia. Chirped pulse amplification (2020-08-13).
https://en.wikipedia.org/wiki/Chirped_pulse_amplification.
191. R. Paschotta. *Optical energy meters*. RP Photonics Encyclopedia (2020-08-13).
192. Thorlabs. Power Meter and Sensor Tutorial (2020-08-13).
https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=6188.
193. R. Paschotta. *Pyroelectric detectors*. RP Photonics Encyclopedia (2020-08-13).
194. Wikipedia. Optical autocorrelation (2020-08-13).
https://en.wikipedia.org/wiki/Optical_autocorrelation.

195. Wikipedia. Frequency-resolved optical gating (2020-08-13).
https://en.wikipedia.org/wiki/Frequency-resolved_optical_gating.
196. B. C. Platt and R. B. Shack. *History and principles of Shack-Hartmann wavefront sensing*. J. Refract. Surg. **17**, S573–S577 (2001).
197. R. Ischebeck, E. Prat, V. Thominet and C. Ozkan Loch. *Transverse profile imager for ultrabright electron beams*. Phys. Rev. ST Accel. Beams **18**, 082802 (2015).
198. T. Kurz, J. P. Couperus, J. M. Krämer, H. Ding et al. *Calibration and cross-laboratory implementation of scintillating screens for electron bunch charge determination*. Rev. Sci. Instrum. **89**, 093303 (2018).
199. COMSOL AB, Stockholm, Sweden. COMSOL Multiphysics Homepage (2020-08-13).
<https://www.comsol.com/>. v5.5.
200. P. Emma, J. Frisch and P. Krejcik. A transverse RF deflecting structure for bunch length and phase space diagnostics. Technical Report LCLS-TN-00-12 SLAC (2001).
201. K. Floettmann and V. V. Paramonov. *Beam dynamics in transverse deflecting rf structures*. Phys. Rev. ST Accel. Beams **17**, 024001 (2014).
202. D. Olsson, F. Curbis, E. Mansten, S. Thorin and S. Werin. *Transverse RF deflecting structures for the MAX IV linac*. Proc. IPAC'18, Vancouver, BC, Canada (2018).
203. M. Borland. Elegant: A flexible SDDS-compliant code for accelerator simulation. Technical Report LS-287 Argonne National Lab., IL (US) (2000).
204. D. Marx, R. Assmann, R. D'Arcy and B. Marchetti. *Simulations of 3D charge density measurements for commissioning of the PolariX-TDS*. J. Phys. Conf. Ser. **1067**, 072012 (2018).
205. D. Marx, R. W. Assmann, P. Craievich, K. Floettmann et al. *Simulation studies for characterizing ultrashort bunches using novel polarizable X-band transverse deflection structures*. Sci. Rep. **9**, 19912 (2019).
206. R. D'Arcy, Deutsches Elektronen-Synchrotron DESY. *Private communication* (2020).
207. C. B. Schroeder, E. Esarey, J. van Tilborg and W. P. Leemans. *Theory of coherent transition radiation generated at a plasma-vacuum interface*. Phys. Rev. E **69**, 016501 (2004).
208. S. I. Bajlekov, M. Heigoldt, A. Popp, J. Wenz et al. *Longitudinal electron bunch profile reconstruction by performing phase retrieval on coherent transition radiation spectra*. Phys. Rev. ST Accel. Beams **16**, 040701 (2013).

-
209. O. Lundh, C. Rechatin, J. Lim, V. Malka and J. Faure. *Experimental Measurements of Electron-Bunch Trains in a Laser-Plasma Accelerator*. Phys. Rev. Lett. **110**, 065005 (2013).
210. R. Weingartner, S. Raith, A. Popp, S. Chou et al. *Ultralow emittance electron beams from a laser-wakefield accelerator*. Phys. Rev. ST Accel. Beams **15**, 111302 (2012).
211. S. K. Barber, J. H. Bin, A. J. Gonsalves, F. Isono et al. *A compact, high resolution energy, and emittance diagnostic for electron beams using active plasma lenses*. Appl. Phys. Lett. **116**, 234108 (2020).
212. S. Kneip, C. McGuffey, J. L. Martins, M. S. Bloom et al. *Characterization of transverse beam emittance of electrons from a laser-plasma wakefield accelerator in the bubble regime using betatron x-ray radiation*. Phys. Rev. ST Accel. Beams **15**, 021302 (2012).
213. G. R. Plateau, C. G. R. Geddes, D. B. Thorn, M. Chen et al. *Low-emittance electron bunches from a laser-plasma accelerator measured using single-shot X-ray spectroscopy*. Phys. Rev. Lett. **109**, 064802 (2012).
214. G. Golovin, S. Banerjee, C. Liu, S. Chen et al. *Intrinsic beam emittance of laser-accelerated electrons measured by x-ray spectroscopic imaging*. Sci. Rep. **6**, 24622 (2016).
215. P. San Miguel Claveria, E. Adli, L. D. Amorim, W. An et al. *Betatron radiation and emittance growth in plasma wakefield accelerators*. Philos. Trans. R. Soc. A **377**, 20180173 (2019).
216. B. Williamson, G. Xia, S. Gessner, A. Petrenko et al. *Betatron radiation diagnostics for AWAKE Run 2*. Nucl. Instrum. Methods Phys. Res. A **971**, 164076 (2020).
217. S. M. Gruner, M. W. Tate and E. F. Eikenberry. *Charge-coupled device area x-ray detectors*. Rev. Sci. Instrum. **73**, 2815–2842 (2002).
218. Andor Technology Ltd., Belfast, UK. Andor iKon-L SO (2020-08-13). <https://andor.oxinst.com/products/high-energy-detection/ikon-l-so>.
219. C. Fourment, C. Arazam, C. Bonte, T. Caillaud et al. *Broadband, high dynamics and high resolution charge coupled device-based spectrometer in dynamic mode for multi-keV repetitive x-ray sources*. Rev. Sci. Instrum. **80**, 083505 (2009).
220. P. Kirkpatrick. *On the theory and use of Ross filters*. Rev. Sci. Instrum. **10**, 186–191 (1939).
221. P. Kirkpatrick. *Theory and use of Ross filters II*. Rev. Sci. Instrum. **15**, 223–229 (1944).