Towards a Free-Electron Laser Driven by Electrons from a Laser-Wakefield Accelerator: Simulations and Bunch Diagnostics

Svetoslav Bajlekov
Merton College, Oxford

Thesis submitted in fulfillment of the requirements for the degree of
Doctor of Philosophy at the University of Oxford

Trinity Term, 2011

Supervised by
Prof. Simon Hooker
Dr. Riccardo Bartolini

Abstract

This thesis presents results from two strands of work towards realizing a free-electron laser (FEL) driven by electron bunches generated by a laser-wakefield accelerator (LWFA). The first strand focuses on selecting operating parameters for such a light source, on the basis of currently achievable bunch parameters as well as near-term projections. The viability of LWFA-driven incoherent undulator sources producing nanojoule-level pulses of femtosecond duration at wavelengths of 5 nm and 0.5 nm is demonstrated. A study on the prospective operation of an FEL at 32 nm is carried out, on the basis of scaling laws and full 3-D time-dependent simulations. A working point is selected, based on realistic bunch parameters. At that working point saturation is expected to occur within a length of 1.6 m with peak power at the 0.1 GW-level. This level, as well as the stability of the amplification process, can be improved significantly by seeding the FEL with an external radiation source.

In the context of FEL seeding, we study the ability of conventional simulation codes to correctly handle seeds from high-harmonic generation (HHG) sources, which have a broad bandwidth and temporal structure on the attosecond scale. Namely, they violate the slowly-varying envelope approximation (SVEA) that underpins the governing equations in conventional codes. For this purpose we develop a 1-D simulation code that works outside the SVEA. We carry out a set of benchmarks that lead us to conclude that conventional codes are adequately capable of simulating seeding with broadband radiation, which is in line with an analytical treatment of the interaction.

The second strand of work is experimental, and focuses on on the use of coherent transition radiation (CTR) as an electron bunch diagnostic. The thesis presents results from two experimental campaigns at the MPI für Quantenoptik in Garching, Germany. We present the first set of single-shot measurements of CTR over a continuous wavelength range from 420 nm to 7 µm. Data over such a broad spectral range allows for the first reconstruction of the longitudinal profiles of electron bunches from a laser-wakefield accelerator, indicating full-width at half-maximum bunch lengths around 1.4 µm (4.7 fs), corresponding to peak currents of several kiloampères. The bunch profiles are reconstructed through the application of phase reconstruction algorithms that were initially developed for studying x-ray diffraction data, and are adapted here for the first time to the analysis of CTR data. The measurements allow for an analysis of acceleration dynamics, and suggest that upon depletion of the driving laser the accelerated bunch can itself drive a wake in which electrons are injected. High levels of coherence at optical wavelengths indicate the presence of an interaction between the bunch and the driving laser pulse.

To my parents

Acknowledgements

I remember getting off the Oxford Tube for the first time, reflecting on the four years of undergrad that I had just waved goodbye to, and grumbling under my breath as I strapped up my backpack:

> *Oh God, here we go again!*

I thought I'd seen it all.

Nearly four years have passed since that moment, and of the multitude of things that present-day me would like to say to the me of back then, none seems as apt as

> *Little do you know. . .*

These last four years have been a rollercoaster ride through the weird, the whacky, the wonderful, and laser-plasma interactions. Many people have helped shape my life during that time, and I feel indebted to all of them for the various ways in which they have contributed to where I am now.

Simon, I think there are few supervisors who would tolerate a student that kept making mistakes with the same frequency and flair as myself. I am grateful for your faith that one day[1] I would learn from them, and for the boundless support you have provided me with. Riccardo, you have introduced me to a world of big machines that few graduate students have the fortune to experience — thank you for all your help and guidance!

Andi, I distinctly remember the start of our time at the Clarendon — how at the beginning we'd finish work at 5pm, then 6pm. . . then a couple of weeks later we were dropping into the Wok'n'Roll for dinner at 11. Thankfully we both realized that was untenable in the long run, and finally settled on 9.30pm — at least until 3rd year hit, that is. Thank you for all the camaraderie! Tom I, I have a related first-day memory of you, where I was trying to ascertain what the normal working hours of a physics DPhil student are; you tried to paint things in a positive light and said "sometimes I get bored and leave at 6." The rest, of course, is history (see above). Thank you for all the computing wisdom you've imparted over the years.

Nicolas, thank you for bringing some Caribbean flair to dreary England — it's been a pleasure working with you! Tom R-R, I never thought I'd have a colleague who was the star of YouTube videos with multiple thousands of hits, *and* was an absolute gun at laser physics! I remember those few nights towards the end of my first year when we were both in the lab at 3am — you because you had to finish your thesis; me because I thought making a conference poster pretty was the most important thing in the world — and secretly felt like the experience would bring me closer to knowing what it's like to finish a thesis.

---

[1] Hopefully sometime in the near future.

Charlotte, thank you for all the newbie advice, and for helping me find my feet at the Clarendon. Thank you also for consistently providing me with a distraction from work around the hours of 11am and 4pm[2], and for the copious quantities of cake over the years! Kevin, thank you for persevering with pilgrimages to the Clarendon tea room in Charlotte's absence, and for keeping spirits high in a distinctly Irish way. Thank you also for being my marathon expert-in-residence. Tom R — thank you for single-handedly ensuring that I always knew at least one person (you) whichever party or bop I went to in Oxford! Wolf, Dave, Lewis, James C: thank you for injecting fresh blood and enthusiasm into the Hooker collective — keep up the good work!

Huge thanks to Nicky, Jesse, Simon W, Andreas D, Gianni, and Daniele for showing that physicists are some of the coolest people around — and for making Oxford as a whole a far more fun place to be! Thanks to Chris, Sam, James M, Katja, Matt, Andy H, Gabriele and the rest of the Simon room occupants for always providing me a good reason to leave my desk for a cup of coffee or a pint. I also have a great deal of appreciation for the hard work put in by the staff at the Clarendon — whether it was Sue's readiness to help whenever I needed *any* paperwork sorted out, or Rob's expert advice in the workshop. Transcending physics and the real world, I would like to thank the *Accelerate!* outreach team — in particular Suzie, Andrew, and Rosalind — for all the good times, and for turning me irreversibly onto science communication!

Simon and Riccardo, of all the countless contributions to my personal and professional development, one stands out in the current context: your unyielding efforts to give me exposure to top-notch research groups around the world. I am grateful for all the doors you opened for me! Correspondingly, I am grateful to all the people I have had the pleasure to collaborate with, of whom I would like to mention a few.

Prof. Suk, Dr. Kotaki, and the team at GIST — thank you not just for the hospitality and all the hard work you put in during the day, but also for all the great Korean food in the evenings! I am also grateful to Todor and the surprisingly numerous eastern European contingent at GIST for being my home away from home.

One of the most enduring and fruitful collaborations I had the fortune to be part of was with Bill Fawley and Carl Schroeder at Berkeley — I am grateful for the time and energy you invested in me, and for the wisdom you imparted! Bill, you have an unparalleled ability to spot the mistakes in my work — which in the end I'd like to think proved a match for my ability to introduce them. You have taught me a great deal simply by setting such a high standard for me to work to. Carl, I have little doubt that you are the cleverest man alive, and that quality is only surpassed by your humility! I am indebted as well to the rest of the Berkeley team for their hospitality — not just Wim for his leadership and for facilitating my visits — but to Jeroen, Tony, Jens, Kei and co. for showing me the ropes and making me feel at home.

Having the chance to work with one team closer to home stands out in particular for me. I may be inclined to say this partly because this collaboration was instrumental in me actually *getting* a DPhil, but all things considered I cannot overstate my gratitude and warm feelings for the fantastic team at MPQ! Stefan, to this day I am not quite sure whether your team works so well because of your leadership — or in spite of it — but there is no doubt

---

[2] The rest of the time, I had to rely on the internet to distract me.

that you are an experimental genius. And a top-notch guy!

Antonia, your hospitality and care for me as a guest were astounding — you made me feel right at home in what initially was a strange land (and a strange lab) — and I will always be grateful for this. Beyond care for my wellbeing, though, you went out of your way to supply me with boundless scientific insight, and were always ready to answer all the absurd questions I came up with. Matthias, as you well know spending endless hours in a laser lab together (often between the hours of midnight and 7am) builds the kinds of bonds that are difficult to replicate by other means. I suspect ours would have been even stronger had I taken up smoking, but someone I resisted the temptation. It was a pleasure working with you! Raphael, thank you for helping me settle in, for sharing the enjoyment of watching Germany meticulously thrash Argentina in the World Cup, and for treating me to your kiwi barbecue skills! Johannes, Kostya, Shao-Wei, and Alex: likewise, thank you for the great science as well as the awesome times together! And last but not least, Frau Wild, your boundless energy and enthusiasm leave me speechless — even if they resulted in me living in a house with 6 dogs for a while.

I don't think anyone can live through a DPhil and retain a shred of sanity when they come out the other side without the support of friends. There are many people who I would like to thank for keeping me company, and for supplying me with food, drink, and banter through the last four years — a few particular individuals deserve a mention here.

Ben: it's been a long road, and it seems we've made it through unscathed! You are due thanks for a large portion of the sanity — and of the great memories — that I've retained during the last 4 years. You are also responsible for a large portion of the memories I haven't retained. Thank you for making it such an enjoyable ride! Dave: sometimes I think I understand you, but most of the time I resign myself to the reality that nobody really does. It's amazing how you keep your intelligence, wit, and all-round gun-baller-ness such closely-guarded secrets. Martin: I remember meeting you in the capacity of a stern ex-Polish-military rowing coach. Then I realized that you have a softer side to you. It's been great seeing you work on that! Sriram: one day my name will be tarnished and I'll need contacts in media to minimize the damage — so we'll be in touch. You guys are a riot!

Carmen: I occasionally feel I owe you an apology for (so you claim?) nudging you into doing a PhD. But now that you're on the verge of figuring out how the Universe works, I don't feel so guilty. Apart from Hamiltonian Monte Carlo, you have taught me a great deal about life, and I am fortunate to have met you! Amy, Barney, Christian, Vic, Patricia, Rachel, Marcus, Felix, Christina, Rick, Toni, Janina, Elie, and several others — both inside and outside Merton's ominous walls — your company has been a real treat!

Mari, thank you first of all for your kindness and for keeping me sane during the final stages of the DPhil experience. Thank you also for giving me a great deal to look forward to afterwards!

Finally, I am grateful to my family for all the support they have provided me throughout, especially since much of the time they did not seem to completely understand why I'm putting myself through this. Above all else, Emilia and Ivan, I owe this to you. On the one hand, you have dedicated immense effort to providing me the opportunities I've enjoyed throughout my life. But more importantly, on the other, you have engendered the ambition and curiosity that have lead me to where I am today — and on which I can rely on to lead me into the future. Thank you!

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Since the dawn of mankind, one of the hallmarks of sentient thought has been the desire to explore our surrounding environment. While this desire is not unique to us as a species, humans are the only ones who have created instruments specifically for the purpose of exploration. In many cases, these instruments have allowed us to overcome the intrinsic limitations of our sensory organs — and none of our senses has been as crucial for exploration as our sight. In first-century Rome, emperor Nero's tutor observed that [1]

> *Letters, however small and indistinct, are seen enlarged and more clearly through a globe of glass filled with water.*

During the two millennia since then we have engineered increasingly sophisticated optical devices, which have allowed us to overcome the limits of our vision and observe objects on smaller and smaller scales.

Fundamentally, however, the spatial resolution of optical microscopes using visible light is limited by the wavelength of the light to a few hundred nanometers — corresponding to photon energies of a few electronvolts. The first half of the 20^{th} century saw the development of electron microscopy, which instead of using photons to probe the samples, bombarded them with electrons accelerated to energies of many kiloelectronvolts. The effective (de Broglie) wavelength of these electrons brought the resolution down to mere fractions of a nanometer, allowing the observation of single atoms. During the second half of the century, the desire to find out still more about the basic constituents of our world led us to the development of neutron and muon sources at even higher energies, allowing us to probe matter further. The Large Hadron Collider, of course, is currently at the frontier of these developments: promising to illuminate the fundamental laws of our Universe by smashing together protons at teraelectronvolt energies.

1.1 Light sources

Returning to the nanoscale, transmission electron microscopy has seen broad use in determining the atomic structure of materials, however in order for the samples to remain trans-parent to the electrons, they must be prepared in slices of thickness $\sim 100\,\mathrm{nm}$ [2]. Moreover, in many cases the electrons can induce sample damage, making electron microscopy unsuitable. These limitations can to some extent be overcome by replacing the keV electrons with keV photons — in other words x-rays — that interact with the atoms more weakly and can penetrate deeper into the sample. A resounding success in this respect has been the application of x-ray crystallography to find the molecular structures of a wealth of materials — including the double-helix structure of DNA in 1953 [3] — thereby providing invaluable knowledge to many spheres of science [4]. X-ray crystallography relies on the scattering of x-rays by a crystal composed of a large number of the target molecules, arranged along the regular crystal lattice. The resulting diffraction pattern contains a series of interference peaks (Bragg peaks), from which the structure of a single molecule can be retrieved.

The "workhorses" behind some of the most important (and certainly most challenging) x-ray crystallography studies have been synchrotron light sources. These are large facilities where bunches of electrons are accelerated to ultra-relativistic gigaelectronvolt-scale energies, and deflected along a pseudo-circular path around a storage ring with a circumference of typically several hundred meters[1]. Radiation is generated at several points around the ring: either in bending magnets, which literally force the electron along a curved path, causing it to emit broadband synchrotron radiation; or in so-called "insertion devices". Insertion devices — typically regular lattices of magnets of alternating polarity — produce a periodically varying magnetic field along the path of the electron bunch, introducing periodic transverse deflections in its motion. The period of the resulting transverse oscillations is usually several millimeters or centimeters, however due to relativistic contraction and Doppler-shifting, the emitted radiation has a wavelength that is scaled by $\sim \gamma^{-2}$. Here $\gamma = E/mc^2$ is the relativistic Lorentz factor of a bunch of energy E, and m and c are the electron rest mass and speed of light in vacuum, respectively. Insertion devices are broadly classified by their strength, where weaker "undulators" produce quasi-monochromatic radiation, while in stronger "wigglers" considerable spectral broadening and harmonic emission take place. For a more detailed discussion of undulator and wiggler radiation see Section 2.3.1.

While synchrotrons produce intense, tightly-focused beams of tunable short-wavelength radiation, they still normally require the materials under investigation be crystallized — containing many regularly-arranged copies of the sample molecule — in order for the diffraction pattern to have sufficient contrast. At the same time, only a small fraction of materials are amenable to crystallization to sufficiently large sample sizes. Powder diffraction can offer a middle ground in this respect, where instead of a single crystal the sample comprises a large number of randomly oriented micro-crystals, still allowing for a certain degree of structural reconstruction [5]. In some sense, though, the holy grail is obtaining x-ray diffraction patterns from single molecules, through small-angle x-ray scattering (SAXS) [6]. Such single-molecule diffraction requires the delivery of a high dose of radiation — however since delivering this dose within a short time would damage or destroy the sample, the diffraction pattern needs to be recorded over an exposure time of roughly a second. On the atomic time-

---

[1] For an up-to-date list of facilities and their parameters see `http://www.lightsources.org/cms/?pid=1000098`.

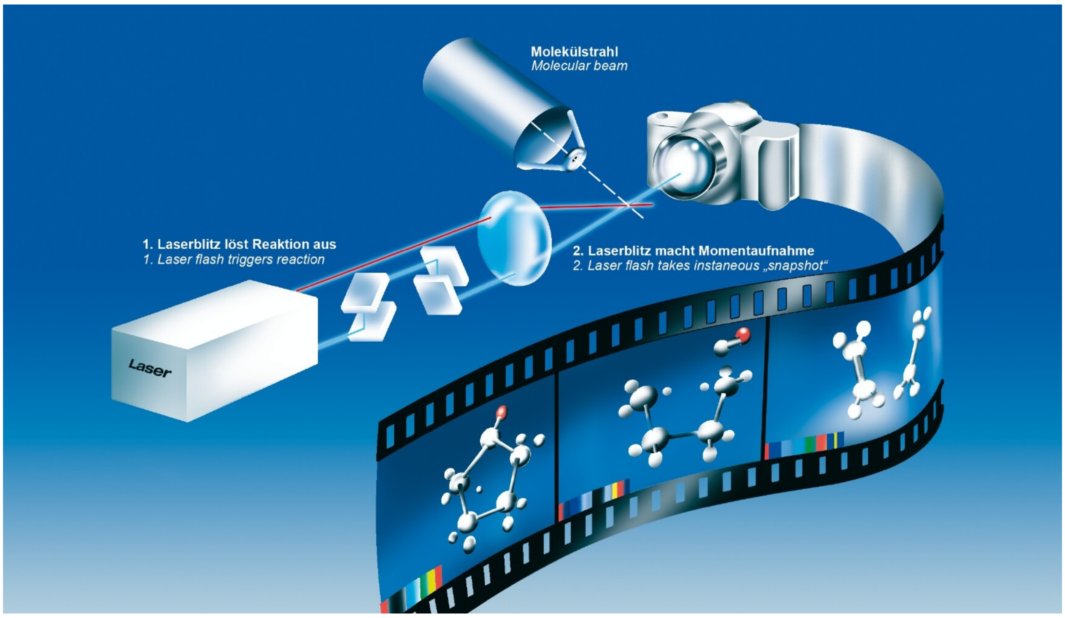

Reproduced with permission from [7]. Copyright © 2011 DESY.

Figure 1.1: Creating a "movie" of a chemical reaction with a femtosecond x-ray free-electron laser. 1. Laser flash triggers reaction; 2. Laser flash makes instantaneous "snapshot".

scale a second is eternity, and only the average atomic positions (over, e.g., thermal motion) are recorded, leading to low resolution.

The solution to this is presented by x-ray *free-electron lasers*. A free-electron laser (FEL) is akin to a synchrotron source, in that high-energy electrons propagate through a set of undulators and generate radiation due to the induced periodic transverse motion. However, in an FEL the coupling between the electrons and the radiation is strong enough that the emitted radiation feeds back on the electrons, causing them to micro-bunch at the radiation wavelength and emit coherently. The gain process is similar to that in a normal laser, with the ultra-relativistic electron bunch constituting the gain medium. The electron bunch is accelerated in a linear accelerator (linac) prior to entering the undulators, and is discarded after it has emitted the x-ray pulse. The process of FEL amplification is treated in detail in Section 2.3.2.

The radiation from an x-ray FEL has a peak brightness[2] that can be 10 orders of magnitude greater than that from the best synchrotron sources [8] — partly due to the higher number of coherently emitted x-ray photons, partly due to the much shorter timescale of emission, and partly due to the better transverse characteristics of the emitting bunch. In particular, while the minimum x-ray pulse duration at synchrotron sources is already a mere ~ 50 picoseconds, at FELs it is below 100 femtoseconds, almost three orders of magnitude shorter. This allows a dose of x-rays that is sufficient for the measurement of single-molecule diffraction to be delivered — and for the pattern to be recorded — before the sample has exploded due to radiation damage [9, 10].

X-ray FELs can generate radiation that has an unprecedented combination of short

---

[2] The number of photons emitted per second per mm^2 crosssection per mrad^2 divergence per 0.1% bandwidth.

wavelength, extremely high brightness, ultra-short pulse duration and if desired a very small — sub-micron scale — focal spot with transverse coherence. This opens up myriad possible applications [8, 11, 12], including the tantalizing prospect of femtosecond-resolution "molecular movies" of intermolecular or photon-molecule interactions (Fig. 1.1). The first (and so far only) hard x-ray FEL in operation is the Linac Coherent Light Source at SLAC, which achieved lasing in 2009 [13, 14]. Already, results from it have successfully been used for the reconstruction of a membrane protein structure [15], as well as that of a non-crystalline virus [16]. Two other x-ray FELs: SPring-8 in Japan [17, 18] and the European X-ray Free Electron Laser (XFEL) at DESY, Hamburg [19] are also set to begin operation during the coming years, broadening the available parameter space. The latter, in particular, builds on the success of FLASH [20] — the first operational soft x-ray FEL and initially a testbed for the XFEL, but currently a user facility in its own right, operating at wavelengths down to 4.12 nm.

1.2 Laser-driven acceleration of electrons in plasma

While synchrotron sources and free-electron lasers are invaluable to scientific research, the availability of experimental time on them is very limited due to the substantial expenditure involved in designing, constructing, and running such large facilities — and the correspondingly small number of sources in operation. A significant part of this expenditure is associated with the equipment required to generate electron bunches and accelerate them to GeV-scale energies. Since conventional radio-frequency (RF) accelerators are limited by electrical breakdown to accelerating fields of around 50 MV/m, the required energies call for linear accelerators that are hundreds of meters long[3] — or a kilometer in the case of the LCLS linac, for an energy of 14.35 GeV.

The advent of the plasma accelerator, however, allows us to overcome this upper limit on the accelerating field, and build smaller accelerators. The concept was first proposed by Tajima and Dawson in 1979 [21]: the ponderomotive force of an intense laser beam can drive a density wake in a plasma, in which electrons are accelerated. The plasma mediates the transfer of energy from the laser to the electron bunch. The wake is analogous to that set up behind a motorboat in water, and the electrons have been said to "surf" down the crests [22]. This concept was extended by Chen *et al*. [23] to include wakefields driven by particle beams instead of lasers. Because plasma is already ionized, it is not subject to the same electrical breakdown limits that afflict conventional accelerators: instead, a plasma can sustain accelerating gradients of several hundred GV/m — over three orders of magnitude greater than an RF accelerator.

For many years, experimental research into plasma accelerators failed to turn Tajima and Dawson's proposal into reality. A plasma wake can only be driven efficiently if the driving pulse has a length shorter than the plasma period. For typical plasma densities

---

[3] The case is somewhat different for synchrotron sources, where a bunch from a linear pre-accelerator experiences gradual acceleration over a number of revolutions through a "booster synchrotron", before it is energetic enough for injection into the main storage ring. While the full acceleration does not take place over a single linear stretch, the size of the facility is still of order hundreds of meters across.

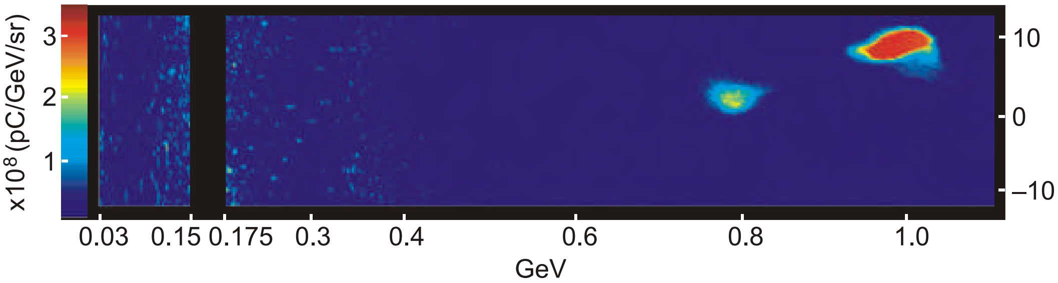

Reproduced with permission from Macmillan Publishers Ltd: *Nature Physics* [36]. Copyright © 2006.

Figure 1.2: Electron spectrometer image showing the first electrons from a laser-wakefield accelerator to reach an energy of 1 GeV. A second bunch at ~ 800 MeV is also visible.

($n_e \sim 10^{18}\,\mathrm{cm}^{-3}$), this period is of order 10 microns, requiring sub-100-femtosecond drivers. As these were not available, through the late ’80s and ’90s early successes included the plasma beat-wave accelerator (PBWA) [24, 25] and the self-modulated laser-wakefield accelerator (SMLWFA) [26–28]; in both schemes a relatively long (picosecond) laser pulse interacts respectively either with another pulse to form a beat-wave, or with the plasma — which induces intensity modulations via Raman scattering — in order to generate a periodic ponderomotive force with spatial modulations that are sufficiently short to drive a wake [29]. However, most of these experiments demonstrated only modest increases in electron energy (albeit over just millimeters of acceleration) and the resulting electrons had very broad thermal energy spectra.

It was the advent of chirped pulse amplification (CPA) of laser pulses that transformed the playing field. First developed by Mourou and Strickland in the mid-1980s [30, 31], CPA involves stretching the laser pulse in time before it is amplified in the gain medium, then compressing it back to its original length. This eliminates the previous limits on intensity that were imposed by the damage threshold of the gain medium, and nowadays allows the generation of pulses with peak powers of hundreds of terawatts and durations below 50 fs [32]. During the first years of this century, developments in high-power CPA-based laser systems led to a watershed moment in laser-wakefield acceleration, as the available intensities and pulse durations edged into the parameter space originally envisioned by Tajima and Dawson. This allowed a single pulse to efficiently drive a wake, into which electrons from the background plasma were trapped and accelerated — culminating in the “dream-beam” experiments of 2004 [33–35]. These experiments demonstrated for the first time the ability of a laser-wakefield accelerator to generate quasi-monoenergetic bunches at the 100-MeV level. These bunches contained a high (100-pC level) charge within a narrow (few percent) energy spread and a narrow angle of divergence.

The dream-beam experiments set off unprecedented growth in the field of laser-wakefield acceleration, with many experimental successes taking place in the years that followed. Notably, in 2006 plasma channeling of the laser pulse was used to extend the acceleration length to 33 mm, allowing the acceleration of bunches to an energy of 1 GeV [36] — a celebrated result that is depicted in Figure 1.2. While this concludes our historical overview of laser-wakefield acceleration, Section 2.1 provides some details on the underlying theory, as well as a summary of the diverse experimental achievements in the field since the dream-beam experiments.

1.3 Light sources driven by laser-wakefield accelerators

The ability of laser-wakefield accelerators to produce high-energy electron bunches with qualities approaching (or exceeding) those from conventional accelerators has brought up the possibility of using them as drivers for light sources. This would lead to dramatic reductions in both the size and cost of these facilities, meaning they could be housed within university, industrial or medical laboratories. Access to intense femtosecond pulses of short-wavelength radiation would become much more readily available to researchers, at a much lower cost to cash-strapped funding bodies. This sort of commoditization would have a transformational effect on countless fields of science.

The production of synchrotron radiation from laser-accelerated electrons has already been demonstrated down to soft x-ray wavelengths [37, 38], and is likely to be pushed to higher energies in the near future. Moreover, since the electron bunches intrinsically are of femtosecond duration due to the size of the plasma wake in which they are accelerated, the corresponding kiloampère peak currents suggest that operation in the free-electron regime could be a possibility [39, 40]. As described previously, this would result in the emission of exceptionally bright femtosecond pulses, which would in addition be intrinsically synchronized in time to the drive laser, making them particularly appealing for time-resolved pump-probe studies.

1.4 Outline of the thesis

The prospect of driving a free-electron laser using electrons from a laser-wakefield accelerator is the subject of theoretical, computational and experimental work presented in this thesis. The primary aims are to establish viable machine parameters for a LWFA-driven FEL source and to characterize its projected performance; as well as to develop and carry out LWFA diagnostics that would afford the community a better understanding of the properties of the electron bunches and the acceleration process. The topics covered are as follows:

Chapter 2 An overview of the components of a “compact light source”, which we take to mean a short-wavelength (XUV or x-ray) light source driven by electrons from a laser-wakefield accelerator. The background theory of LWFA is presented, as well as a summary of experimental results, on the basis of which sample parameter sets are established for electron bunches produced currently, and in the near future. The concept of emittance is introduced, and some details of beam transport from the accelerator to the undulator are given, including focusing and space-charge effects. The basic theory of undulator radiation is presented, together with an overview of the free-electron laser regime of operation, and of electron bunch dynamics in the undulator. Relevant undulator technology is also summarized.

Chapter 3 A selection of parameter scans and simulations to characterize the operation of a compact light source. Preliminary studies on the production of incoherent undulator radiation at wavelengths of 5 nm and 0.5 nm are presented. A comprehensive study on the prospective operation of an FEL at 32 nm follows. A set of machine parameters are selected on the basis of scaling laws, and the results are complemented with a set of full 3-D time-dependent simulations for a wide range of bunch parameters. These demonstrate the need for comprehensive and reliable electron bunch diagnostics. The effect of seeding with a high-harmonic generation (HHG) source on the output level and stability of FEL operation is investigated.

Chapter 4 Results from preliminary diagnostics experiments based on coherent transition radiation (CTR). The theory of CTR is given, and the evolution of the experimental setup is described. Several different measurements are presented, which as a broad theme illustrate the presence of sub-structure in the bunch that leads to the emission of coherent transition radiation at optical wavelengths. Preliminary tests demonstrate the potential for CTR to be used as a non-destructive transverse profile diagnostic.

Chapter 5 Results from broadband CTR measurements allowing longitudinal bunch profile reconstruction. The first set of single-shot measurements of CTR spectra over a continuous wavelength range from 420 nm to 7 µm is presented. A method for reconstructing the longitudinal bunch profiles from these spectra is adapted from phase reconstruction algorithms that were initially developed for studying x-ray diffraction data. The reconstructed profiles — in conjunction with the other diagnostics — allow for a comprehensive analysis of acceleration dynamics including dephasing, pump depletion, injection into multiple buckets and the existence of beam-driven wakes. This bunch profile diagnostic also allows a more accurate evaluation of the projected performance of LWFA-driven light sources. Finally, the seemingly anomalous levels of coherence at certain optical wavelengths are analyzed.

Chapter 6 A treatment of issues related to the simulation of FELs seeded with broadband radiation. A theoretical representation of seeding in the 1-D linearized regime is presented. In the context of FEL seeding, the ability of conventional simulation codes to correctly handle seeds from high-harmonic generation (HHG) sources, which have a broad bandwidth and temporal structure on the attosecond scale, is studied. Namely, these seeds violate the slowly-varying envelope approximation (SVEA) that conventional codes rely on. A 1-D simulation code that was developed for this purpose and works outside the SVEA is described, and FEL simulations of broadband seeding are benchmarked against conventional codes. Finally, the possibility of obtaining resonant FEL gain due to the interaction of non-resonant seed components and non-resonant longitudinal modulations on the electron bunch is highlighted.

Chapter 7 A summary of results, together with an overview of future work towards realizing a free-electron laser driven by electron from a laser-wakefield accelerator.

1.5 Role of the author

The data presented in Chapter 4 was collected during an experimental campaign at the Atlas laser facility at the MPI für Quantenoptik in Garching, Germany in June 2010. This was led by the author, and carried out in collaboration with the research group of Stefan Karsch at MPQ, the members of which made sizable contributions to the experiments. The equipment required to generate and record the transition radiation — including appropriate shielding from stray radiation — was designed and produced by the author prior to the experiment. Data acquisition for all non-CTR related diagnostics was carried out by a semi-automated system in place at MPQ; CTR data acquisition was carried out by the author. All data analysis was carried out by the author, using codes written by the author for this purpose.

The data presented in Chapter 5 was collected during an experimental campaign at MPQ in March/April 2011. It was again carried out in collaboration with the research group of Stefan Karsch and was led by one of its members, Matthias Heigoldt, who had constructed and commissioned the terahertz spectrometer. The author was responsible for the set-up and operation of the other two (visible and near-infrared) spectrometers, with data acquisition tasks being divided likewise. The author contributed to setting up the experiment and calibrating the visible and near-infrared spectrometers, and was present during two of the three data acquisition runs. All data analysis — except for theoretical calibration of the THz data — was carried out by the author. During the course of analysis, the author designed thirty-six different metrics that were evaluated (usually, but not always, automatically) on the basis of the raw data for each of the 2974 recorded shots, in order to identify meaningful trends in the large dataset. The longitudinal profile reconstruction algorithm was designed and coded by the author, prompted by a suggestion (and some preliminary analysis) from Simon Hooker to use phase retrieval algorithms to analyze the CTR data. These results are to be submitted for publication, pending the absolute calibration of the terahertz spectrometer at a light source later this year.

The theoretical work presented in Chapter 6 was carried out in collaboration with Bill Fawley and Carl Schroeder at Lawrence Berkeley National Laboratory. The code AURORA was designed and written by the author and is not based on pre-existing codes. All simulations were carried out by the author. The work on resonant gain due to the interaction of non-resonant longitudinal modes of the seed radiation and electron bunch is entirely the author's own.

The author also participated in a number of experiments that are not presented in this thesis. These include a LWFA campaign at the Gwangju Institute of Science and Technology in Korea in May 2008; a LWFA campaign at the Astra-Gemini laser in June/July 2008 [41]; preparations for a campaign to measure LWFA electron bunch emittance using a pepper-pot at MPQ in November/December 2009 that never took place due to laser malfunction; and during April–June 2010 a set of emittance measurements using transition radiation at the conventional accelerator driving the Diamond Light Source [42].

All figures are the author's own work, unless credited otherwise.

Chapter 2

Overview of Components and Design of a Compact Light Source

As stated in the introduction, one of the most alluring uses of laser-wakefield accelerators is as drivers for short-wavelength undulator-based light sources. This chapter and the next are dedicated to an analysis of how these light sources — hereafter referred to generally as ‘compact light sources’ — might be achieved in practice. Broadly speaking, such a source comprises three stages: the laser-wakefield accelerator, the beam transport, and the undulator where radiation is generated. These are pictured in Figure 2.1, and are the subject of the current chapter. Details on the production of the driving laser pulse and on the potential applications of such a light source can be found in e.g. Refs. [30–32, 43–45] and Refs. [9, 10, 15, 46], respectively.

We begin by introducing the theory of laser-wakefield acceleration, together with a review of recent experimental achievements and a summary of future projections. From there we move to a discussion of issues concerning beam transport from the accelerator to the undulator — including emittance, focusing, and space-charge effects. Finally, in Sec. 2.3 we present the basic theory of undulator radiation, including a discussion of operation in the free-electron laser regime. This allows us to, in Chapter 3, present proposals for concrete machine parameters, and analyze the expected performance of these future light sources.

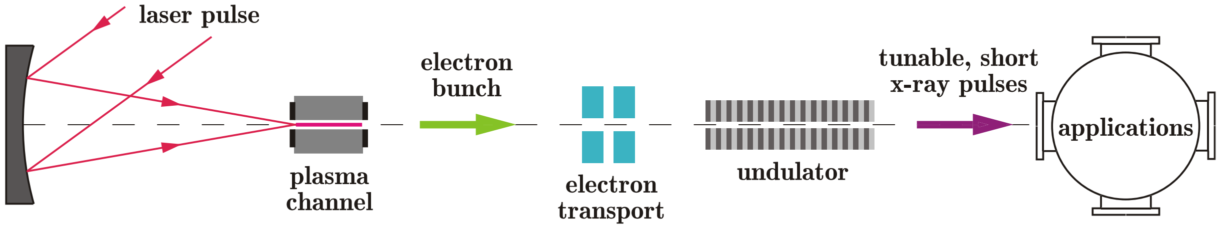

Figure 2.1: Layout of a compact light source driven by electrons from a laser-wakefield accelerator. After injection and acceleration of an electron bunch in the plasma (described in Sec. 2.1) and appropriate transport (Sec. 2.2), the bunch drives the generation of short-wavelength femtosecond pulses of radiation in an undulator (Sec. 2.3) that can be used for various applications (Sec. 1.1). Image credit: Simon Hooker.

2.1 Laser-wakefield accelerator

While Sec. 1.2 introduced the development of plasma-based acceleration as an experimental field in a historical context, here we present the elementary theory of laser-driven plasma accelerators. Here we focus on the case of the laser-wakefield accelerator, where the pulse length of the driving laser is similar to the plasma wavelength — as opposed to, e.g., the self-modulated laser-wakefield accelerator [26, 27] or the plasma beat-wave accelerator [24, 25], which can be driven by longer pulses — though much of the theory will be applicable in other contexts. Following this, we review recent experimental results and novel approaches, in order to establish a viable parameter space for electrons accelerated by a laser-wakefield accelerator.

2.1.1 Theory of laser-wakefield acceleration

The basic concept behind the proposal by Tajima and Dawson [21] is the creation of a plasma density wake behind a driving laser pulse, in which the electrons are accelerated. We therefore begin with a discussion of the ponderomotive force of the laser that leads to the generation of this wake, and of the plasma oscillations that govern the wake's behavior.

Ponderomotive force

The processes associated with intense laser-plasma interactions are often represented in terms of the normalized electrostatic potential, $\phi = e\Phi/mc^2$, and the normalized electromagnetic vector potential, $\mathbf{a} = e\mathbf{A}/mc$, where Φ and $\mathbf{A}$ are the corresponding non-normalized potentials in the Coulomb gauge. The constants e, m, and c are respectively the elementary unit of charge, the electron rest mass, and the speed of light in vacuum. SI units are used here and throughout this thesis.

A linearly polarized laser pulse with peak intensity I_0 and peak power P_0, and a Gaussian transverse profile of spot size w_0, has a peak normalized vector potential [47]

$$\begin{aligned} a_0 &= \sqrt{\frac{e^2\lambda^2 I}{2\pi^2\epsilon_0 m^2 c^5}} \\ &\approx 8.6\times10^{-10}\,\lambda\,[\mu\text{m}]\sqrt{I\,[\text{W/cm}^2]} \\ &\approx \frac{\lambda}{w_0}\sqrt{\frac{P\,[\text{GW}]}{21.5}}, \end{aligned} \tag{2.1}$$

where λ is the laser wavelength and ϵ_0 is the vacuum permittivity. This is based on a vector potential of the form

$$\mathbf{a} = a_0\, e^{-r^2/w_0^2}\cos(kz-\omega t)\,\mathbf{e}_x, \tag{2.2}$$

where $k = 2\pi/\lambda$ and $\omega = ck$; r and z are the radial and longitudinal coordinate, respectively; and $\mathbf{e}_x$ is the unit vector along the direction of laser polarization, here chosen to be x. The corresponding electric field, in the absence of a space-charge contribution, $\nabla\phi = 0$, can be expressed in terms of the potential as

$$\mathbf{E} = -\mathrm{d}\mathbf{A}/\mathrm{d}t = -E_0\, e^{-r^2/w_0^2} \sin(kz - \omega t)\,\mathbf{e}_x,$$

where $E_0 = (mc/e)\omega a_0$. An electron in this field undergoes transverse oscillatory 'quiver motion' at the driving laser frequency.

Conservation of canonical momentum implies $\mathbf{u}_\perp = \mathbf{a}_\perp$, where $\mathbf{u}_\perp = \mathbf{p}_\perp/mc$ is the normalized transverse momentum of the electron. This is true in 1-D, as well as in 3-D for the case of a sufficiently large spot size, $w_0 > \lambda_\mathrm{p} \gg \lambda$, in which case the quiver motion is still the leading order motion in the transverse dimension [47, 48]. Here λ_p is the plasma wavelength, which we introduce shortly.

Let us now take a potential with a more general spatial envelope, $\hat{a}(r,z)$, such that $\mathbf{a}_\perp = \hat{a}\cos(\omega t)\mathbf{e}_x$. Inserting this into the momentum conservation equation, we obtain the time-average $\langle u_\perp^2 \rangle = \hat{a}^2/2$, and the corresponding kinetic energy

$$\left\langle E^2 \right\rangle = mc^2(1 + \hat{a}^2/2).$$

This indicates that a gradient in laser intensity would correspond to an energy gradient, and therefore a force due to $\mathbf{F} = -\nabla E$,

$$\mathbf{F}_\mathrm{p} = -\frac{mc^2}{4\sqrt{1+\hat{a}^2/2}}\nabla\hat{a}^2. \tag{2.3}$$

Equation (2.3) is the ponderomotive force, also known as radiation pressure, which pushes electrons away from areas of high intensity. Ponderomotive motion of heavier charged particles falls off quadratically with particle mass[1], e.g. by a factor $\sim 3\times10^{-7}$ for protons, therefore the effect of $\mathbf{F}_\mathrm{p}$ on the plasma ions can be neglected [49].

Finally, we note that electrons undergoing quiver motion will have a time-averaged squared Lorentz factor of $\langle\gamma^2\rangle = 1 + \hat{a}^2/2$, and can thus become significantly relativistic for $a \gtrsim 1$, which corresponds to laser intensities $\gtrsim 10^{18}\,\mathrm{W/cm^2}$. Most recent LWFA experiments have taken place in this relativistic intensity regime, and relativistic corrections have to be taken into account in many contexts, in particular as relating to the change in effective electron mass.

Plasma waves

The ponderomotive force of a laser propagating through a plasma perturbs the plasma density, and can drive a plasma density wake that trails the laser pulse. The behavior of this wake is governed by plasma oscillations, which we go on to describe. Here we assume the plasma is cold and unmagnetized, i.e. the thermal motion of electrons is negligible compared

---

[1] Equation (2.3) in its current form is only valid for electrons. To find the appropriate relation for particles of arbitrary mass m_p we need to express the force in terms of the gradient in the electric field amplitude, $\nabla\hat{E}^2$, which cancels the m in the numerator and introduces a factor m_p in the denominator. The corresponding acceleration is then a factor $(m_\mathrm{p}/m)^2$ less than that caused by the same field in the case of electrons.

to the quiver motion. Furthermore, the much heavier ions are assumed to remain stationary throughout the interaction.

The equations governing the propagation of an electromagnetic wave are the Maxwell-Faraday equation and Ampère's law,

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}; \tag{2.4}$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}, \tag{2.5}$$

where $\mathbf{J}$ is the current density and μ_0 is the permeability of free space. For small perturbations in plasma density, general solutions to these coupled equations can be decomposed as sums over the wavevectors, $\mathbf{k}$, of the corresponding plane wave solutions:

$$\mathbf{E_k} e^{i(\mathbf{k}\cdot\mathbf{r}-\omega_\mathbf{k} t)}; \qquad \mathbf{B_k} e^{i(\mathbf{k}\cdot\mathbf{r}-\omega_\mathbf{k} t)}; \qquad \mathbf{J_k} e^{i(\mathbf{k}\cdot\mathbf{r}-\omega_\mathbf{k} t)}. \tag{2.6}$$

Additionally, from the Lorentz force we obtain the electron velocity corresponding to a wave of vector $\mathbf{k}$ to be $\mathbf{v_k} = -ie\mathbf{E_k}/m\omega_\mathbf{k}$, omitting the oscillating exponentials for clarity. This is valid for non-relativistic motion.

Substituting this velocity in the current density, $\mathbf{J} = -n_\mathrm{e} e \mathbf{v}$, where n_e is the plasma number density, we obtain $\mathbf{J_k} = i n_\mathrm{e} e^2 \mathbf{E_k}/m\omega_\mathbf{k}$. This expression can be substituted into equations (2.4) and (2.5), for a plane wave of particular vector $\mathbf{k}$,

$$\mathbf{k} \times \mathbf{E_k} = \omega \mathbf{B_k}; \tag{2.7}$$

$$\mathbf{k} \times \mathbf{B_k} = \mu_0 \frac{n_\mathrm{e} e^2}{m\omega_\mathbf{k}} \mathbf{E_k} - \omega_\mathbf{k} \mu_0 \epsilon_0 \mathbf{E_k}. \tag{2.8}$$

Two classes of wave solutions exist for equations (2.7) and (2.8): those where $\mathbf{k} \parallel \mathbf{E_k}$, and those with $\mathbf{k} \perp \mathbf{E_k}$. The first class of equations are longitudinal electrostatic plasma oscillations, or Langmuir waves, for which from Eq. (2.7) we have $\mathbf{B_k} = 0$. Substituting this into equation (2.8) we obtain the plasma frequency in the linear regime,

$$\omega_\mathrm{p} \equiv \omega_\mathbf{k} = \sqrt{\frac{e^2 n_\mathrm{e}}{m\epsilon_0}}. \tag{2.9}$$

The frequency of these oscillations is only dependent on the plasma density, n_e, and due to the independence of ω_p on $\mathbf{k}$, their phase velocity is undefined. Important related quantities are the wavelength of the oscillations, $\lambda_\mathrm{p} = 2\pi c/\omega_\mathrm{p}$, also known as the plasma period, and the associated wavenumber $k_\mathrm{p} = 2\pi/\lambda_\mathrm{p}$. Evaluating the constants,

$$\lambda_\mathrm{p}[\mu\mathrm{m}] \approx \frac{3.3 \times 10^{10}}{\sqrt{n_\mathrm{e}[\mathrm{cm}^{-3}]}}. \tag{2.10}$$

It is these plasma oscillations that provide the longitudinal accelerating fields in laser-wakefield accelerators.

For the other class of waves, where $\mathbf{k} \parallel \mathbf{E_k}$, we can substitute equation (2.7) into equation (2.8) to obtain the dispersion relation

$$ck^2 = \omega_k^2 - \omega_\mathrm{p}^2, \tag{2.11}$$

that governs the propagation of electromagnetic waves through plasma. We will designate the frequency of these electromagnetic waves $\omega \equiv \omega_k$. Since k is imaginary for $\omega < \omega_\mathrm{p}$, the plasma only allows the unattenuated propagation of waves with frequency $\omega > \omega_\mathrm{p}$. Their phase velocity is $v_\mathrm{p} = \omega/k = c/\sqrt{1-\omega_\mathrm{p}^2/\omega^2}$, and their group velocity is $v_\mathrm{g} = \partial\omega/\partial k = c\sqrt{1-\omega_\mathrm{p}^2/\omega^2}$. For typical laser-wakefield accelerator parameters where the driving laser wavelength is $\lambda \sim 1\,\mu\mathrm{m} \to \omega \sim 2\times10^{15}\,\mathrm{s}^{-1}$, and the plasma density is $n_\mathrm{e} \sim 5\times10^{18}\,\mathrm{cm}^{-3} \to \omega_\mathrm{p} \sim 10^{14}\,\mathrm{s}^{-1}$ ($\lambda_\mathrm{p} \sim 15\,\mu\mathrm{m}$), we have $\omega_\mathrm{p}^2 \ll \omega^2$, known as the "underdense regime". Under these conditions the laser pulse propagates through the plasma at a group velocity

$$v_\mathrm{g} \approx c\left(1-\omega_\mathrm{p}^2/2\omega^2\right). \tag{2.12}$$

More general, longitudinal plasma waves in an initially uniform plasma driven by the ponderomotive force of a laser pulse obey the relations [48, 50]

$$\left(\frac{\partial^2}{\partial t^2}+\omega_\mathrm{p}^2\right)\frac{\delta n_\mathrm{e}}{n_{\mathrm{e},0}} = c^2\nabla^2 a^2/2,$$
$$\left(\frac{\partial^2}{\partial t^2}+\omega_\mathrm{p}^2\right)\phi = \omega_\mathrm{p}^2 a^2/2,$$

where $n_{\mathrm{e},0}$ is the background plasma density and $|\delta n_\mathrm{e}/n_{\mathrm{e},0}| \ll 1$ is the normalized density perturbation corresponding to the electrostatic potential ϕ. The solutions for this density perturbation and the associated electric field are [48]

$$\frac{\delta n_\mathrm{e}}{n_{\mathrm{e},0}} = \frac{c^2}{\omega_\mathrm{p}}\int_0^t \mathrm{d}t'\,\sin\left[\omega_\mathrm{p}(t-t')\right]\,\nabla^2 a^2(\mathbf{r},t')/2, \tag{2.13}$$
$$\frac{\mathbf{E}}{E_0} = -c\int_0^t \mathrm{d}t'\,\sin\left[\omega_\mathrm{p}(t-t')\right]\,\nabla a^2(\mathbf{r},t')/2, \tag{2.14}$$

where the latter is normalized to

$$E_0 = \frac{mc\omega_\mathrm{p}}{e}, \tag{2.15}$$

the cold non-relativistic wave-breaking field [51]. This is the peak field for waves in the linear regime; the field that can be sustained by nonlinear waves can be significantly more, for reasons that we outline shortly. Evaluating the constants, $E_0\,[\mathrm{V/m}] \approx 96\sqrt{n_\mathrm{e}\,[\mathrm{cm}^{-3}]}$, indicating the possibility for plasma waves to sustain field gradients of over 200 GV/m for $n_\mathrm{e} \sim 5\times10^{18}\,\mathrm{cm}^{-3}$. The phase velocity of the plasma wave is determined by the group velocity of the driving laser pulse, $v_{\mathrm{p,wake}} \approx c(1-\omega_\mathrm{p}^2/2\omega^2)$.

Solutions to equations (2.13) and (2.14) indicate that for optimal plasma wake generation, the pulse length of the driving laser, i.e. the extent of the envelope a^2 in time, should be close to the plasma wavelength, λ_p. This is a key reason for the historical connection between the development of ultrashort-pulse lasers (lengths ~ 100 fs) and the proliferation of LWFA experiments.

Figure 2.2a illustrates the (sinusoidal) normalized density perturbation and electric field of the wake trailing a pulse with peak $a_0 = 0.5$ and rms length $\lambda_\mathrm{p}/2\pi$. As noted earlier, however, recent LWFA experiments have in general employed relativistic laser intensities,

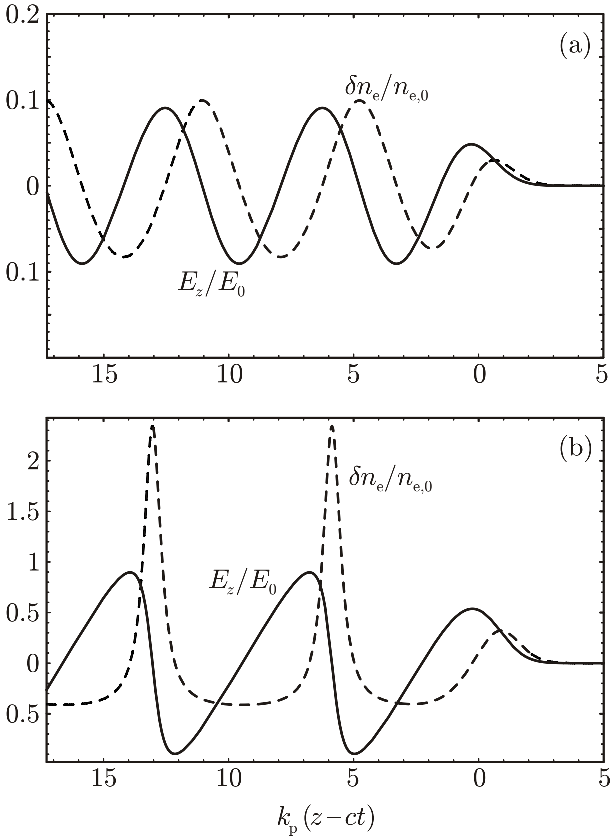

Reproduced with permission from [48]. Copyright © 2009 American Physical Society.

Figure 2.2: Plasma wakes in 1-D driven by laser pulses in (a) the linear regime, $a_0 = 0.5$ and (b) the nonlinear regime, $a_0 = 2.0$. The driving pulse is centered at $(z - ct) = 0$ and has an intensity envelope with a Gaussian profile of rms length $\sigma_{ct} = \lambda_p/2\pi$. Note the wavefront steepening and elongation of the period in the nonlinear regime.

$a_0 \gtrsim 1$. This regime calls for a nonlinear treatment of the problem: one that no longer assumes the normalized density perturbation to be small, and includes relativistic effects. Nonlinear models have been developed in the 1-D limit and can be found in e.g. Refs. [47, 48], where the case of a non-evolving drive pulse is treated. Figure 2.2b illustrates the key features of a nonlinear wake: the density perturbations are steeper, and the plasma period is elongated with respect to λ_p. The non-linear plasma wavelength is [48]

$$\lambda_{p,n} = \frac{2}{\pi}\left(\tilde{E}_{max} + \tilde{E}_{max}^{-1}\right), \tag{2.16}$$

where $\tilde{E}_{max}$ is the peak electric field of the wake, normalized to the peak field in the linear regime, E_0. For a drive pulse with a square temporal profile and an optimum length $\tilde{E}_{max} = \left(a_0^2/2\right)\left(1 + a_0^2/2\right)^{-1/2}$, though for other pulse shapes it usually needs to be evaluated numerically. For a Gaussian pulse it is around 0.7–0.8 of that value.

The nonlinear 1-D regime illustrated in Fig. 2.2b is valid in the limit of broad drive pulses, $k_p w_0 \gg 1$, where the laser and plasma characteristics have only a weak dependence on the transverse dimension. For relativistically intense laser pulses that are more tightly focused, in the regime $k_p w_0 \lesssim 1$, the nonlinear (intensity-dependent) aspects of the wake characteristics acquire an additional transverse dependence. For instance, due to the nonlinear dependence of plasma period on wake amplitude, the wavefronts of the wake become curved. Effects such as this lead to more complex interplays that can no longer be modeled analytically, and require numerical simulation. Due to this complexity, the remainder of this discussion will be of a more qualitative nature.

The bubble regime

In the extreme nonlinear case, a relativistically intense ($a_0^2 \gg 1$) laser pulse that is focused to a sufficiently small spot size, $w_0 \leq (2/k_\mathrm{p})\sqrt{a_0}$, can ponderomotively expel all plasma electrons from its axis of propagation in what is known as the "bubble", "blow-out", or "cavitation" regime [52–56]. The column of positively charged background ions that is left behind the laser pulse attracts the radially expelled electrons — they swing back and overshoot, forming a wake. The end result is illustrated in Figure 2.3a: here the acceleration parameters have been matched in a way that produces a spherical cavity of radius $R \approx \sqrt{a_0}\lambda_\mathrm{p}/\pi$, corresponding to an effective plasma wavelength in the bubble regime of

$$\lambda_\mathrm{p,b} = \frac{2\sqrt{a_0}}{\pi}\lambda_\mathrm{p}. \tag{2.17}$$

This matched regime was identified by Lu *et al.* [55, 56]; it corresponds to optimum wake generation and occurs for a laser spot size[2] $w_0 \approx \sqrt{a_0}\lambda_\mathrm{p}/\pi$ and pulse length $c\tau < \sqrt{a_0}\lambda_\mathrm{p}/\pi$ at normalized laser potentials $a_0 \gtrsim 4$. In this case, the two-way interaction between the laser pulse and the wake is such that transverse focusing of the pulse is maintained and this accelerating structure can in principle propagate stably over many Rayleigh lengths.

At the rear edge of the bubble, the longitudinal accelerating field reaches its peak value of $E_\mathrm{m} \approx \sqrt{a_0}E_0$, where E_0 is the cold non-relativistic wave-breaking field (cf. Eq. (2.15)). The field falls off approximately linearly with z for locations further forward in the bubble; it is zero at the center, and is decelerating in the front hemisphere. Additionally, the bubble exerts a transverse focusing force on electrons inside it that is approximately linear in r and broadly independent of z. This leads off-axis electrons to perform characteristic transverse "betatron" motion, often leading to the emission of betatron radiation in the x-ray spectral region [57–59].

Injection of electrons into the wake

A key step towards the acceleration of an electron bunch to high energy is its injection into the accelerating structure — either a plasma bubble or a regular plasma wave. This can either take place spontaneously — e.g. due to electrons that were perturbed by the laser entering the wake and being trapped in it — or in a controlled fashion. According to simulations [56, 60], in the bubble regime self-injection can take place with high probability for $a_0 \gtrsim 3$, e.g. as seen in Fig. 2.3a–b; experiments support this assertion [41, 61–63]. The

---

[2] Note that for these conditions the radius of the bubble is equal to the laser spot size.

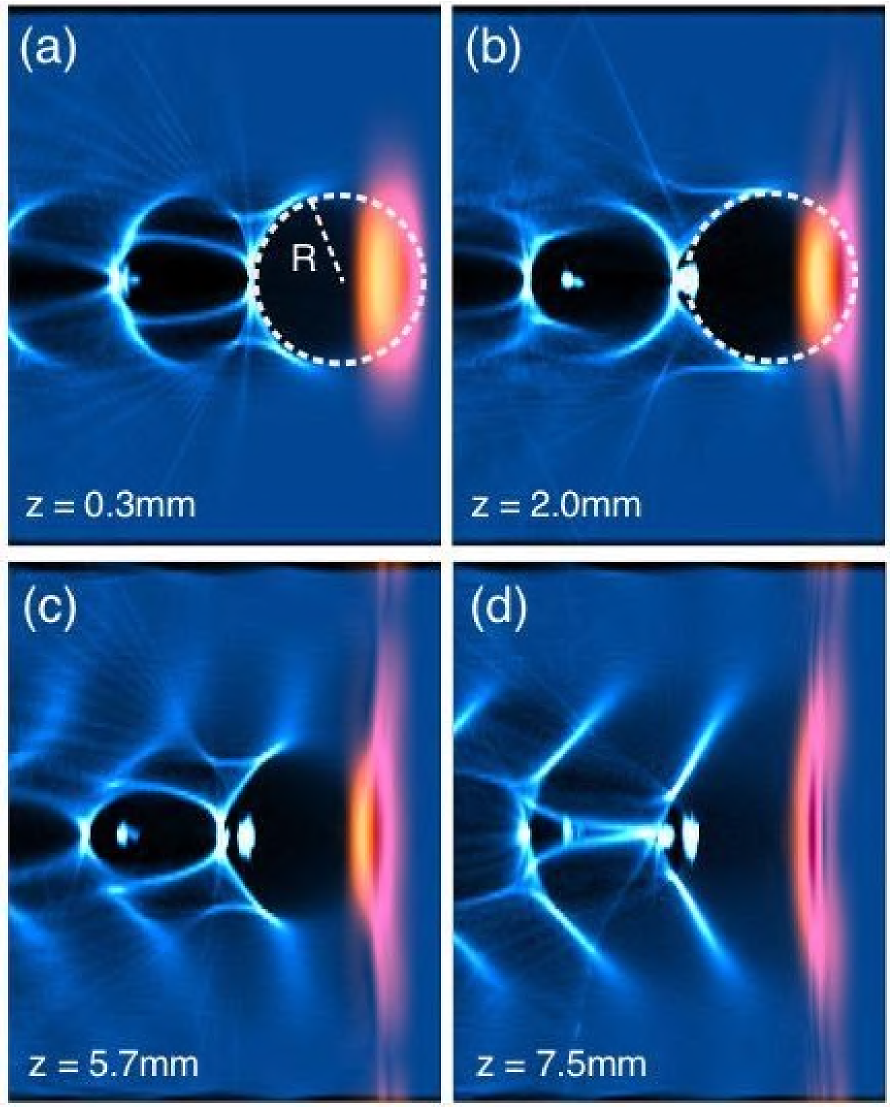

Reproduced from [56] under Creative Commons license.

Figure 2.3: A sequence of 2-D slices from a particle-in-cell simulation of laser-wakefield acceleration in the bubble regime. Electron density is represented in blue, while laser intensity is orange. The near-spherical shape of the cavitated region is clearly visible in slices (a)–(b), emphasized by a dashed white circle. Self trapped electrons in the first bucket are visible at $z = 2.0\,\mathrm{mm}$ (b) and beyond, and beam loading due to their space-charge is observed to distort the spherical shape of the cavity. As the front of the laser pulse diffracts and is etched away, the back is self-guided and the wakefield remains relatively robust even after a propagation distance of 7.5 mm (5 Rayleigh lengths). After this the depleted laser pulse can no longer sustain self-focusing and diffracts, terminating the acceleration process. Each plot is a rectangle of size $z = 101.7\,\mu\mathrm{m}$ in the longitudinal direction and $x = 129.3\,\mu\mathrm{m}$ in the transverse direction.

drawbacks of self-injection are the lack of shot-to-shot stability and a decreased bunch quality. Despite the existence of certain conditions which electrons need to meet for them to be trapped in the wake [48], injection can take place during different stages of the acceleration process, and the electrons are injected at a range of locations within the wake, with a range of initial (transverse and longitudinal) velocities. This can lead to fluctuations in the final energy, energy spread, bunch length, charge, transverse bunch size, pointing, and divergence.

A number of schemes for controlling injection have been implemented in recent years: here we present a brief overview of the most successful ones. So-called "colliding-pulse injection" was one of the first schemes to be proposed [64, 65], and was realized soon after the initial generation of quasi-monoenergetic beams, see e.g. Refs. [66–70]. In this scheme an intense 'pump' laser pulse drives the wake, and a weaker counter-propagating 'injection' pulse collides with it at the desired injection point. The interference of the two pulses creates a standing beatwave that exerts a strong ponderomotive force on the background plasma electrons, and under certain conditions has the potential to inject them into the wake.

Another proposal by Bulanov *et al.* [71] involves a downward transition in plasma density. Such a transition leads to an increase of the plasma period, $\lambda_{\mathrm{p}} \propto n_{\mathrm{e}}^{-1/2}$, cf. Eq. (2.10), and an increase in the group velocity of the laser (cf. Eq. (2.12)) and correspondingly the velocity of the wake. Under certain conditions this can lead to wave-breaking: namely, despite the acceleration of the front of the wake, the lengthening of the plasma period slows down the phase velocity of the actual wake crests, thus briefly lowering the threshold velocity for electrons to be trapped in the wake. The injection process can be controlled by tailoring the plasma density profile. Experimentally, this has been demonstrated by Geddes *et al.* [72], Faure *et al.* [73], and Gonsalves *et al.* [74]. In the latter case a gas jet down-ramp is used purely for injection, while acceleration takes place in a separate stage. Finally, Schmid *et al.* [75] demonstrated a related injection mechanism that relies on a sharp shock-front-induced density transition.

Finally, recent experiments demonstrated the possibility of injecting electrons by ionization of additional species present in a partially-ionized plasma: this was done by Oz *et al.* in 2007 [76] for the case of beam-driven acceleration, and by McGuffey *et al.* [77] and Pak *el al.* [78] in 2010 for laser-driven acceleration. In the laser-driven case, this mechanism relies on the fact that these additional species (dopants) have higher optical field ionization thresholds for certain electron shells. While the background plasma (typically hydrogen or helium) would be fully ionized by the front of the laser pulse, the interaction can be tailored so that some of the electrons in the dopants would only ionize at a certain point near the peak of intensity of the driving laser, thereby giving control over the injection position and velocity. Several groups have recently employed this method for inducing injections in regimes where it does not normally occur, and controlling it [77–81]. In particular, Pollock *et al.* [81] demonstrated the operation of a two-stage accelerator, where the injection — from ionization of 0.5 % N_2 gas in a 99.5 % He background — and acceleration stages are separate from each other.

We conclude our brief overview of injection with a mention of beam loading. The charge that has been injected in the wake exerts its own electrostatic repulsion, which alters the shape of the wake and thus degrades its accelerating properties. This can for instance be observed in Fig. 2.3b as the deviation of the wake shape from an ideal sphere in the region behind the injected bunch. In the simplest sense, beam loading sets a limit on the amount of bunch charge that can be sustained by a wake, but it also imposes a trade-off between the injected charge and the energy to which it can be accelerated. The effects of beam loading have been studied theoretically and computationally [82, 83], and clear signatures have been observed in experiments [68, 84].

Laser evolution

A further aspect that affects acceleration is the evolution of the laser pulse through nonlinear interactions with the plasma. Two important phenomena in this respect are relativistic self-focusing, and pulse compression [85–87]. Self-focusing occurs due to the fact that the relativistic quiver motion of the electrons increases their effective (relativistic) mass, and correspondingly the local refractive index of the plasma. This leads to a gradient in refractive index that scales with the intensity gradient of the laser — with a peak on-axis — that can lead to self-focusing. The critical laser power at which self-focusing can counteract diffraction is [85]

$$P_\mathrm{c} = \frac{8\pi\epsilon_0 m^2 c^5}{e^2}\frac{\omega^2}{\omega_\mathrm{p}^2} \approx 17.4\frac{\omega^2}{\omega_\mathrm{p}^2}\,[\mathrm{GW}], \tag{2.18}$$

where ω is the angular frequency of the laser. In particular, while a pulse in vacuum undergoes Rayleigh diffraction with its spot size w evolving in z as $w^2/w_0^2 = 1 + z^2/Z_\mathrm{R}^2$, where Z_R is the Rayleigh length, in the case of self-focusing for a laser pulse with peak power P_0 this is modified to $w^2/w_0^2 = 1 + (1 - P_0/P_\mathrm{c})\, z^2/Z_\mathrm{R}^2$ [48]. In theory, this will cause a sufficiently powerful laser to self-focus to a spot size $w_\mathrm{sf} = \left(\sqrt{a_0}/\pi\right)\lambda_\mathrm{p}$, the matched spot size in the bubble regime [56] — that is the point at which full cavitation occurs and the radial ponderomotive force of the laser is balanced by the radial attractive force of the ion channel, leading to no further self-focusing. Substituting for a_0 from Eq. (2.1), we can re-express w_sf in terms of the critical power for self-focusing,

$$w_\mathrm{sf} = \frac{\sqrt{2}}{\pi}\lambda_\mathrm{p}\left(\frac{P_0}{P_\mathrm{c}}\right)^{1/6}. \tag{2.19}$$

This matched self-focused spot size corresponds to a peak normalized potential of

$$a_0^\mathrm{sf} \approx 2\left(\frac{P_0}{P_\mathrm{c}}\right)^{1/3}. \tag{2.20}$$

While this is the ideal case, in practice the evolution of the laser pulse may proceed differently, and analysis through simulations is advisable. In particular, if electrons are injected before the laser has self-focused to the matched spot size, their space-charge could alter the profile — and thus focusing effect — of the channel (see below), leading to different evolution.

Despite the effect of relativistic self-focusing, outward ponderomotive pressure causes the build-up of a plasma density gradient at the front of the pulse — the resulting gradient in refractive index serves to almost exactly negate the effect of relativistic self-focusing.

In the absence of an external guiding structure, therefore, the front of the laser pulse will always diffract, while the rest of the pulse is guided behind it. This is less problematic in the case where the rate at which the front of the pulse etches away due to energy transfer to the wake is faster than the rate at which it diffracts, which in the bubble regime happens for powers above

$$P_\mathrm{c}^\mathrm{diff} = (1/8)(\omega/\omega_\mathrm{p})^{6/5} P_\mathrm{c}, \tag{2.21}$$

known as the critical power for self-guiding [56, 88].

These conditions can be relaxed through the use of external guiding structures, such as preformed plasma channels. A plasma channel is a plasma column with a radial density profile — corresponding to a gradient in the index of refraction — that can guide a laser pulse of a certain matched spot size without diffraction. These channels can be created either by laser-induced hydrodynamic expansion (the “ignitor-heater technique”) [35, 89, 90], or by a capillary discharge [36, 41, 49, 91–93]. Use of the latter for LWFA was pioneered by Simon Hooker’s group at Oxford, which the author belongs to, in collaboration with the group of Wim Leemans at LBNL. It was key to the first acceleration of electrons to GeV energies in a laser-wakefield accelerator [36].

Pulse compression occurs due to similar considerations. The longitudinal density gradient created by ponderomotive pressure at the front of the pulse results in a gradient in the index of refraction, which in turn results in a gradient in the laser group velocity. This leads to steepening of the leading edge of the pulse and shortening of its overall duration. This happens in conjunction with a redshifting of the photons at the front of the pulse — a natural consequence of energy transfer to the wake [56].

Limits to acceleration

So far we described a process where a laser pulse drives a wake in which electrons are trapped and accelerated, although we gave few details of the dynamic evolution — and indeed termination — of this process. Here we briefly cover the key points in this respect. Broadly speaking, electron acceleration comes to an end upon one of three possible limiting events: the diffraction of the laser pulse, the depletion of its energy, or the dephasing of the electron bunch within the wake so that it no longer experiences an accelerating (and focusing) field.

Above we briefly discuss laser diffraction, and in particular ways that it can be counteracted. Virtually all contemporary acceleration experiments ensure guiding of the laser pulse over many Rayleigh lengths by either operating in the parameter range for self-guiding, or by employing external guiding structures. Termination of the acceleration process due to laser diffraction is therefore rarely a limiting factor for well-designed systems.

The regime where acceleration is terminated due to depletion of the pump laser (or ‘pump depletion’), on the other hand, is more common. This depletion process can be thought of as taking place somewhat differently in the 1-D ($w_0 \gg \lambda_\mathrm{p}$) nonlinear ($a_0^2 \gg 1$) regime, and in the bubble regime. In the 1-D case we can consider the overall transfer of energy from the drive pulse to the wake, and by equating the energy in the wake to that of the pulse we obtain a pump depletion length of [48]

$$L_{\mathrm{pd,n}} = \frac{\sqrt{2}a_0}{\pi}\frac{\lambda_{\mathrm{p}}^3}{\lambda^2}. \tag{2.22}$$

In the bubble regime, only the front of the laser actually interacts with the plasma, since all plasma electrons behind it have been expelled. This causes the front to etch at a rate $v_{\mathrm{etch}} \approx c\omega_{\mathrm{p}}^2/\omega^2$ [94], which for a laser pulse of length τ_{d} results in a bubble-regime pump depletion length of [56]

$$L_{\mathrm{pd,b}} = \frac{\lambda_{\mathrm{p}}^2}{\lambda^2}c\tau_{\mathrm{d}}. \tag{2.23}$$

The final limiting factor is electron dephasing. Electrons are only accelerated towards $+z$ in the part of the wake that has a negative electrostatic potential gradient along z, i.e. $E_z < 0$. This roughly corresponds to the rear half of each "bucket" of the wake, see Fig. 2.2. When it reaches the middle of the bucket the bunch is no longer accelerated, and in the front of the bucket it is decelerated. This movement of the bunch towards the front of the wake during the acceleration process is inevitable, since under typical conditions the group velocity of the driving laser — and thus the phase velocity of the wake — corresponds to a relativistic Lorentz factor $\gamma(v_{\mathrm{g}}) = \omega/\omega_{\mathrm{p}} \sim 20$, while electron bunches are frequently accelerated to energies corresponding to Lorentz factors $\gamma \sim 1000$. In other words, the bunch velocity (after some acceleration) can be far closer to c than the wake velocity, causing it to outpace the wake. The dephasing length in the nonlinear 1-D regime [48],

$$L_{\mathrm{d,n}} = \frac{\sqrt{2}a_0}{\pi}\frac{\lambda_{\mathrm{p}}^3}{\lambda^2}, \tag{2.24}$$

again differs from that in the bubble regime [56],

$$L_{\mathrm{d,b}} = \frac{2\sqrt{a_0}}{3\pi}\frac{\lambda_{\mathrm{p}}^3}{\lambda^2}, \tag{2.25}$$

although only by a small constant factor. Here we have neglected the requirement for transverse focusing of the bunch throughout its acceleration, which in the nonlinear 1-D regime[3] would have reduced $L_{\mathrm{d,n}}$ by a further factor of 2. Note that in this regime the dephasing length is equal to the pump depletion length, $L_{\mathrm{d,n}} = L_{\mathrm{pd,n}}$. For typical parameters, acceleration takes place over a distance of a few to a few tens of millimeters.

For the case where acceleration is limited by dephasing, the maximum energy gains are given in literature as [48]

$$\Delta W_{\mathrm{d,n}}\,[\mathrm{MeV}] \approx \frac{400\,I\,[\mathrm{W/cm^2}]}{n_{\mathrm{e}}\,[\mathrm{cm^{-3}}]} \tag{2.26}$$

for the 1-D nonlinear regime, and as [56]

$$\Delta W_{\mathrm{d,b}}\,[\mathrm{MeV}] \approx 0.25\left(\frac{\lambda_{\mathrm{p}}}{\lambda}\right)^{4/3}\sqrt{P\,[\mathrm{GW}]} \tag{2.27}$$

for the bubble regime. Note that even though E_0, the peak accelerating field in the linear regime, scales with the square root of plasma density (cf. Eq. (2.15)), the dephasing length

---

[3] In the bubble regime transverse focusing occurs at all longitudinal positions within the bubble.

has an inverse dependence on $n_e^{3/2}$, and therefore the total energy gain achievable grows with decreasing plasma density. Both of these idealized relations give energies in the range of several 100 MeV to several GeV for typical acceleration parameters — particle-in-cell simulations are generally necessary to capture the full dynamics of the acceleration (including, e.g., laser evolution and beam loading effects) and obtain more reliable estimates.

This subsection aimed to give an overview of the key concepts behind laser-wakefield acceleration. Many details were omitted here: theoretical treatments of the nonlinear regimes; details of the criteria for self-injection; a more detailed theory of laser evolution; or indeed scaling laws for acceleration in the linear regime, which is not frequently used in practice. For further information we direct the reader to the many excellent publications on the topic, and in particular the comprehensive review articles by Esarey, Sprangle, Krall and Ting [47], and by Esarey, Schroeder and Leemans [48]. We go on to discuss recent experimental achievements in the field.

2.1.2 Review of recent experiments

Since the initial generation and acceleration of quasi-monoenergetic "dream beams" in the experiments by Mangles *et al.* [33], Geddes *et al.* [35], and Faure *et al.* [34], the field of laser-wakefield acceleration has practically ballooned: in terms of the number of research groups involved in experiments, as well as theory and simulations; the number of laser facilities dedicated to LWFA experiments; and the number of publications. Here we give a brief overview of recent experimental trends, on the basis of which we select viable parameter sets for electron bunch production, now and in the near future.

Table 2.1 summarizes the approximate acceleration (laser and plasma) parameters for which recent laser-wakefield acceleration experiments have been carried out, and lists figures of merit for the quality of the produced electron bunches. The parameters P and I are respectively the peak power and intensity of the driving laser, and τ_d is the full-width at half-maximum (FWHM) pulse duration; E is the energy of the best quasi-monoenergetic bunches that were produced, and σ_E/E is their FWHM energy spread; Q is the charge, which is sometimes quoted as the charge over all detected energies — explaining the very high values given in some cases — though more often the charge in the quasi-monoenergetic peak; $\sigma_{r'}$ is the FWHM divergence; ε_n is the normalized emittance (only where it was deduced experimentally — many publications quote figures of a few mm mrad based on simulation results); and σ_z is the FWHM bunch duration, again only where this was measured experimentally. The experiments are listed in roughly chronological order. The years are approximate.

A simple skim down the table reveals a number of general trends: the frequency of LWFA experiments — or at least the frequency of those that yielded publishable results — has increased significantly, as has the diversity of the high-power laser facilities where the experiments are taking place. As the experiments mature the application of controlled injection techniques is becoming more commonplace, with ionization injection in particular gaining recent popularity. More sophisticated bunch diagnostics are also now being carried out, such as measurements of emittance using pepper pots [95, 96], or of bunch length via coherent transition radiation [97] (also Chapter 5 of the present thesis) and direct imaging inside the wake [98].

A trend that is less clear from the presented figures is the fact that the stability of acceleration has increased dramatically: while earlier publications were based around data taken for a small number of successful shots, during more recent experiments regimes have been identified where quasi-monoenergetic electron bunches are produced for $> 90\%$ (sometimes $> 99\%$) of laser shots [99, 100], with corresponding stability improvements in bunch energy and pointing (typically $< 10\%$ and a few mrad, respectively). The contribution of better pulse quality and stability of the driving lasers in this respect cannot be overstated.

Nonetheless, it seems that some bunch parameters have not improved as significantly during the seven years since the initial "dream-beam" experiments. One of these is the bunch energy spread, which appears to have remained stubbornly at a level $\gtrsim 2\%$ FWHM. A closer look reveals that for typical experiments this is close to the resolution limit of the electron spectrometer [35, 36, 101, 102], which clearly calls for more sophisticated diagnostics, for instance bunch focusing in the spectrometer as reported by Weingartner *et al.* [100].

Bunch charge measurement is also an issue of some contention. One reason for this is the lack of definition of what the "bunch charge" actually comprises: whether it is the charge in the entire bunch, or above some cut-off energy, or in the quasi-monoenergetic peak — though in the last case, there is no universal definition of the bounds on the peak. This is particularly an issue in the (common) cases where the quasi-monoenergetic peak in the electron spectrum is superimposed on a background with a broad quasi-thermal spectrum, which can itself contain a lot of charge. Recent experiments where charge at the 10 pC level was consistently observed in the high-energy quasi-monoenergetic peak [97, 99] have called into question the reliability of earlier measurements of nanocoulomb-level bunches [34, 35]; the latter imply $\sim 10\%$ conversion efficiency of laser into electron energy. In either case, reliable calibration of the scintillating screens used in electron spectrometers — as recently carried out by e.g. Nakamura *et al.* [103] and Buck *et al.* [104] — and its consistent application across experiments are of crucial importance.

Based on the results in Table 2.1 and the above considerations, in Table 2.2 we present a summary of current experimental capabilities for stable operation. While the experimental values quoted in Table 2.1 correspond to full-width at half-maximum measures of spread (e.g. in energy, divergence and bunch length), the values given in Table 2.2 have been converted to root-mean-square, assuming a Gaussian distribution, i.e. by dividing by $2\sqrt{2\log 2}$. This reflects the fact that root-mean-square measures are more common when dealing with topics pertinent to conventional accelerators, and will facilitate the discussion in the remainder of this chapter, and the next.

In Table 2.2 we also present projected parameters for stable operation in the near future. These projections are in the context of a number of concepts that are currently being developed and tested, that are expected to lead to notable improvements in the near term. Of these, controlled injection is probably the most important, but certainly not the only one.

Table 2.1: Overview of recent LWFA experimental parameters and results. Quoted spreads (σ) are full-width at half-maximum.

Laser			**Plasma**			**Electron bunch**								
P (TW)	τ_d (fs)	$I/10^{18}$ ($\mathrm{W/cm^{-2}}$)	$n_\mathrm{e}/10^{18}$ ($\mathrm{cm^{-3}}$)	Target	Inj.	E (MeV)	σ_E/E (%)	Q (pC)	$\sigma_{r'}$ (mrad)	ε_n (µm)	σ_z (µm)	Ref.	Facility	Year (~)
12	40	2.5	20.0	$\veebar$		72	3	22	87.0			[33]	ASTRA (RAL)	2004
30	33	3.2	6.0	$\veebar$		170	24	500	5.0			[34]	LOA	2004
9	55	11.0	19.0	Υ		86	2	320	3.0			[35]	LOASIS (LBNL)	2004
30	30	3.2	6.0	$\veebar$		170	24					[106]	LOA	2005
30	30	3.0	6.0	$\veebar$		200	5		10.0		13.0	[107]	LOA	2006
24	30	3.4	7.5	$\veebar$	C	125	9	25	4.0			[66]	LOA	2006
35	35	4.0	20.0	$\veebar$		150	7					[58, 61]	Lund	2006
8	80	50.0	40.0	$\veebar$		47	4		10.0			[108]	Jena	2006
40	38	1.0	4.3	$\circledcirc$		1000	2	30	1.6			[36, 109]	LOASIS (LBNL)	2006
11	45		15.0	$\veebar$		40	30	30			13.6	[110]	ASTRA (RAL)	2006
14	50	2.2	7.5	$\veebar$		80	6		5.0			[111]	ASTRA (RAL)	2007
18	42	1.5	3.2	$\circledcirc$		300	3	45	1.3			[112]	ATLAS (MPQ)	2007
8	40	19.0	10.0	$\veebar$	C	9	7		1.3			[113]	ASTRA (RAL)	2007
13	45	0.8	8.0	$\circledcirc$	I	200	5	100	1.5			[93]	ASTRA (RAL)	2007
10	47	16.0	22.0	$\veebar$	D	1	22	500	20.0			[72]	LOASIS (LBNL)	2007
5	80	5.0	20.0	$\veebar$		64	2	28	1.6			[37, 101]	Jena	2008
20	42	1.7	7.3	$\boxdot$		198	8	10	1.8			[99]	ATLAS (MPQ)	2008
11	40	3.0	20.0	$\veebar$		134	8	9	4.0			[70]	J-KAREN (JAEA)	2008
	30	4.6	5.7	$\veebar$	C	180	5	20				[68]	LOA	2008
30	80		2.5	$\circledcirc$		540	10	250				[41, 88]	GEMINI (RAL)	2008
200	55	19.0	5.7	$\veebar$		800		550	3.6			[114]	GEMINI (RAL)	2008
5	8	12.0	20.0	$\veebar$		25	3	3	6.3			[115]	LWS-10 (MPQ)	2009
20	37	1.7	8.0	$\boxdot$		210		7	0.7			[38]	ATLAS (MPQ)	2009
30	30	3.3		$\circledast$		80						[116]	HERCULES (MI)	2009
	30	4.6	5.7	$\veebar$	C	206	7	13	4.5			[67, 84]	LOA	2009

Continued on next page

Table 2.1: Overview of recent LWFA experimental parameters and results. *(cont.)*

P	τ_d	$I/10^{18}$	$n_\mathrm{e}/10^{18}$	Target	Inj.	E	σ_E/E	Q	$\sigma_{r'}$	ε_n	σ_z	Ref.	Facility	Year
65	60		3.0	⊻		720	7	100	2.9			[63]	Callisto (LLNL)	2009
72	32	47.0	8.0	⊻		230	11	200	3.8			[59]	Hercules (MI)	2009
8	8	2.5	19.0	⊻	D	24	4	2	0.5			[75]	LWS-10 (MPQ)	2010
23	37		7.7	⊡		170		23				[117]	Atlas (MPQ)	2010
24	30	35.0	25.0	⊻	I	110	10	15	3.0			[77]	Hercules (MI)	2010
11	45		14.0	⊻	I	92						[78]	UCLA	2010
30	30	2.0	10.0	⊻		125			3.0	2.2		[95]	Alpha-X (Strath.)	2010
10	8		50.0	⊻	D	20		1	6.0	2.3		[96]	LWS-10 (MPQ)	2010
110	60	31.0	1.3	⊻	I	1450	50	4	25.0			[79]	Callisto (LLNL)	2010
30	30	3.6	8.5	⊻	C	100	10	100	4.3			[73]	LOA	2010
30	30	3.6	10.0	⊻	C	84	25	15	6.0		1.1	[97]	LOA	2010
16	8	6.0	32.0	⊻		19	12	2	11.0		5.3	[98]	LWS-20 (MPQ)	2010
28	30	40.0	14.0	⊻		190						[118]	Hercules (MI)	2011
23	37			⊡		190	2		3.2			[100]	Atlas (MPQ)	2011
45	40		5.7	⊻	I	800	15	4	2.6			[80]	Shanghai	2011
40	60		3.0	⊻	I	460	5	35	2.3			[81]	Callisto (LLNL)	2011

Key:

⊻	gas jet	C	colliding-pulse injection (collinear or colliding)
⊡	gas cell	I	ionization-induced injection
◎	discharge capillary waveguide	D	density-ramp or density-transition injection
⊛	ablated capillary waveguide		
⋎	plasma channel by laser pre-ionization		

Table 2.2: Sample electron bunch parameters from a LWFA source: current state and near-term estimate. All spreads are root-mean-square.

		Current	Future	
Energy	E_e	< 800	< 1500	MeV
Energy spread	σ_E/E	1	0.5	%
Bunch charge	Q	20	100	pC
Bunch length	σ_z	0.5	0.5	µm
Transverse size	σ_r	1	1	µm
Divergence	$\sigma_{r'}$	1	0.5	mrad
Geometric emittance	ε	1	0.5	µm mrad
Normalized emittance	ε_n	1.5	1.5	mm mrad

The separation of the injection from the acceleration stages (demonstrated recently by Pollock *et al.* [81] and Gonsalves *et al.* [74]) presents another significant advance, since the conditions required for injection are often very different from those desired for acceleration. Namely, while injection typically requires high plasma densities[4], the optimum conditions for acceleration to high energies are obtained in lower density plasmas (see e.g. Ref. [105], and the scalings in Sec. 2.1.1). Further target optimization, such as tailoring of the plasma density profile, as well as improvements in laser stability and reliability, are likely to lead to continued improvements in the observed bunch characteristics. At the same time, improved and more advanced diagnostics (see e.g. [96–98] and Chapters 4–5 of the current thesis) will continue to shed light not just on these characteristics, but on the acceleration process itself.

2.2 Beam transport

Upon exiting the plasma accelerator, the electron bunch will propagate a certain distance in vacuum before reaching its intended target, in this case the undulator. In this section we discuss the dynamics of the bunch during free-space propagation, and highlight ways in which it can be manipulated — using so-called beam optics — in order to bring it to the state required for optimum generation of radiation. We also discuss the issue of repulsion due to space-charge forces.

2.2.1 Introduction to beam emittance

The electrons that comprise a bunch will have a certain distribution of positions, x, y, z, and corresponding momenta, p_x, p_y, p_z, that are all functions of time. We may introduce a function, $f(x, p_x, y, p_y, z, p_z, t)$, that describes this distribution in 6-D phase space, representing it as a continuous particle density rather than a collection of discrete particles. According to Liouville's theorem [119], this density — and hence the volume in phase space occupied

---

[4] Even though self-injection is also observed at low densities, higher densities almost universally lead to higher injected bunch charges.

by f — will remain constant under the action of conservative forces. This fact underpins the definition of emittance, denoted ε, as a measure of the phase-space volume occupied by f [120, 121]. Emittance, in the simplest sense, is a measure of the bunch quality: incorporating in it the bunch length and energy spread, as well as the transverse size and divergence. It is conserved under propagation of the bunch through drift space, or a linear focusing system (see Sec. 2.2.2).

Typically, emittance along the spatial dimensions x, y and z is considered separately, by projecting the volume of f onto the orthogonal planes $x - p_x$, $y - p_y$, and $z - p_z$. In fact, the unqualified term "emittance" is frequently used to refer to transverse emittance (along $x - p_x; y - p_y$) alone. Furthermore, the "trace-space" quantities $x' = p_x/p_z$ and $y' = p_y/p_z$ are often used as a more convenient (angular) measure of deviation from straight motion along the z-axis, than the momenta p_x, p_y [122]. This assumes $p_x, p_y \ll p_z$, which is true for the relativistic bunches considered here.

Here we focus on transverse emittance, and for illustrative purposes limit ourselves to the x-dimension. While definitions vary, a common measure of emittance is the root-mean-square emittance, which in 1-D is

$$\varepsilon_{x,\mathrm{rms}} = \sqrt{\left\langle x^2 \right\rangle \left\langle x'^2 \right\rangle - \langle x x' \rangle^2}. \tag{2.28}$$

Transverse emittance is commonly normalized by the mean relativistic Lorentz factor of the bunch, γ, to obtain the normalized emittance, $\varepsilon_\mathrm{n} = \gamma\varepsilon$. As well as being conserved under the above-mentioned conditions, normalized emittance is also conserved under perfect acceleration or deceleration of the bunch. For clarity, non-normalized emittance is sometimes referred to as 'geometric emittance'.

One of the advantages of representing a bunch by the collective phase-space distribution of the particles — rather than by their individual positions in phase space — is the simplifications of calculations of bunch evolution along the beamline: the tracking of individual particle trajectories is no longer necessary, as this representation allows for the propagation of the overall distribution. Such calculations are normally carried out by representing the bunch distribution in each spatial dimension as an ellipse, e.g. in $x - x'$ space, which in particular allows for analytical propagation. Figure 2.4 illustrates this phase ellipse representation.

The Twiss parameters[5], also known as transport parameters, α_T, β_T, and γ_T, together with the emittance, completely describe the phase ellipse boundary via the relation

$$\gamma_\mathrm{T} x^2 + 2\alpha_\mathrm{T} x x' + \beta_\mathrm{T} x'^2 = \varepsilon, \tag{2.29}$$

e.g. as shown in Fig. 2.4. In particular, the projected bunch size — i.e. the size measured experimentally by a device, such as a scintillating screen, that only records particle positions and not directions of motion — is $X = \sqrt{\beta_\mathrm{T}\varepsilon}$, while the spread in propagation angles is $X' = \sqrt{\gamma_\mathrm{T}\varepsilon}$. Correspondingly, the Twiss parameters β_T and γ_T have dimensions of length

---

[5] We use the subscript $_\mathrm{T}$ to denote Twiss parameters in order to avoid confusion, in particular with the normalized particle velocity β or Lorentz factor γ.

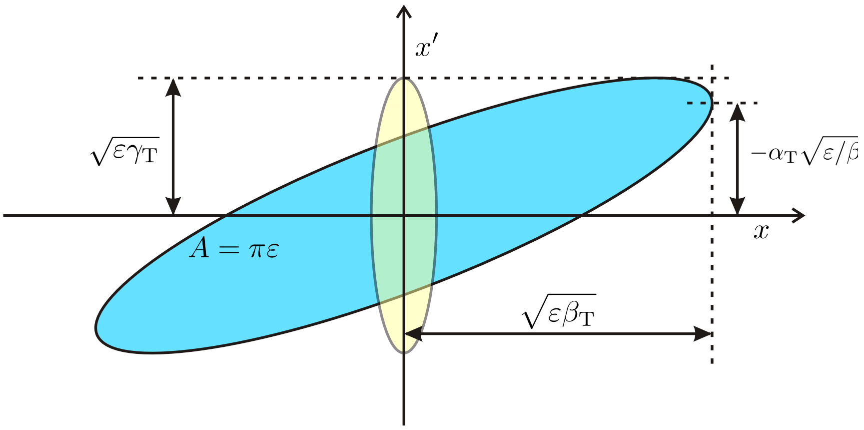

Figure 2.4: Phase ellipse representing the distribution of particles in x–x' space, i.e., particle location along the horizontal axis and divergence (angle of propagation with respect to the z-axis) along the vertical axis. The area of the ellipse is the emittance, ε_x, multiplied by π. The Twiss parameters β_{T} and γ_{T} represent the projected bunch size and divergence, respectively, while α_{T} is a measure of whether the bunch is converging towards a waist in x ($\alpha_{\mathrm{T}} > 0$), or diverging away from one ($\alpha_{\mathrm{T}} < 0$; blue ellipse). The yellow ellipse represents a bunch at waist ($\alpha_{\mathrm{T}} = 0$).

and length^{-1}. During propagation through a free space of length L, the three parameters evolve according to the transformation

$$\begin{pmatrix} \alpha'_{\mathrm{T}} \\ \beta'_{\mathrm{T}} \\ \gamma'_{\mathrm{T}} \end{pmatrix} = \begin{pmatrix} 1 & 0 & -L \\ -2L & 1 & L^2 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \alpha_{\mathrm{T}} \\ \beta_{\mathrm{T}} \\ \gamma_{\mathrm{T}} \end{pmatrix}, \tag{2.30}$$

where the primed quantities are those after the drift space.

Experimentally, transverse emittance is a useful quantity because it encapsulates both the size and the angular divergence of the bunch. It represents a limit on the ability to keep the bunch transversely focused to a narrow cross-section over a long distance — a topic that is the subject of the next subsection.

2.2.2 Focusing optics

An electron bunch exiting the laser-wakefield accelerator typically has a transverse size on the micrometer scale, in part imposed by the transverse plasma bubble dimension, and transverse divergence on the milliradian scale. While the combination of these parameters results in a low normalized emittance of order 1 mm mrad (see e.g. Sec. 2.1.2), the relatively large divergence leads the bunch to expand transversely in size to several millimeters over just a meter or two of drift space. A bunch of this size would be useless for most radiation generation applications.

Ideally, we would like to exploit the low emittance and achieve a better balance between transverse size and divergence. In practice this means going from an ellipse that is narrow (in x) but tall (in x') — such as the one colored yellow in Fig. 2.4 — to one that is wide but

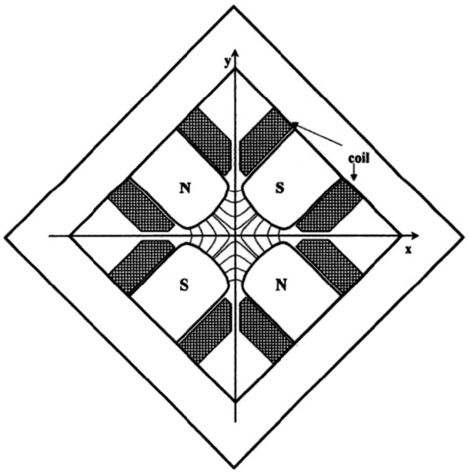

Reproduced with permission from World Scientific Publishing Co [123]. Copyright © 2000.

Figure 2.5: Cross-section of a quadrupole magnet lens, with magnetic field lines shown emanating from the electromagnets.

short. The latter represents a bunch of the same emittance that has a larger transverse size (at focus), but a lower divergence. Key to achieving this transformation is the use of a focusing system. The basic component of such a system is the quadrupole lens, a schematic of which is shown in Fig. 2.5. The field strengths in x and y are proportional to the distance from the axis [124]:

$$B_x = \frac{\partial B_x}{\partial y} y = k_\mathrm{q} y, \qquad B_y = \frac{\partial B_y}{\partial x} x = k_\mathrm{q} x, \tag{2.31}$$

where k_q is the quadrupole gradient. Via the Lorentz force, this results in transverse acceleration [123]

$$\frac{\mathrm{d}^2 x}{\mathrm{d}z^2} = k_\mathrm{q} x, \qquad \frac{\mathrm{d}^2 y}{\mathrm{d}z^2} = -k_\mathrm{q} y, \tag{2.32}$$

i.e., while particles in one of the dimensions are pushed back towards the axis[6] in the other dimension they are accelerated away from it. Due to this, a quadrupole is either horizontally focusing and vertically defocusing, or horizontally defocusing and vertically focusing.

For quadrupoles that are relatively short, so that particles undergo negligible change in transverse position during the time they take to propagate through them, we may employ the thin-lens approximation, yielding the changes in angular direction as functions of transverse position,

$$\Delta x' \approx -K_\mathrm{f} l x = -\frac{x}{f}, \qquad \Delta y' \approx K_\mathrm{f} l y = \frac{y}{f}, \tag{2.33}$$

where $K_\mathrm{f} = e k_\mathrm{q}/(\gamma m c)$ is the focusing parameter, l is the quadrupole length, and f is the effective focal length. The thin-lens approximation is valid for focal lengths $f \gg l$. In the nomenclature of Twiss parameters, a thin lens leads to the transformation

$$\begin{pmatrix} \alpha'_\mathrm{T} \\ \beta'_\mathrm{T} \\ \gamma'_\mathrm{T} \end{pmatrix} = \begin{pmatrix} 1 & 1/f & 0 \\ 0 & 1 & 0 \\ 2/f & 1/f^2 & 1 \end{pmatrix} \begin{pmatrix} \alpha_\mathrm{T} \\ \beta_\mathrm{T} \\ \gamma_\mathrm{T} \end{pmatrix}, \tag{2.34}$$

---

[6] For instance, for $k_\mathrm{q} < 0$, particles at $x < 0$ are accelerated towards positive x and particles at $x > 0$ are accelerated towards negative x.

cf. equation (2.30).

In Section 2.2.1 we noted that Liouville's theorem, which underpins emittance conservation, is applicable under the action of conservative forces. The forces that result from an ideal quadrupole magnet's linear field gradients, cf. Eq. (2.31), are indeed conservative, as verified by Wiedemann in Ref. [120], where he also provides a generalized quantitative criterion for the class of transformations that preserve the phase-space volume of a bunch. On the other hand, nonlinearities in these fields may increase emittance — an effect that can be compensated using beamline optics with higher-order fields, such as sextupole magnets [124].

A related effect — chromaticity — arises when an electron bunch with a finite energy spread passes through the quadrupole. Since the focusing strength of a quadrupole, K_f, depends on the electron energy, electrons of different energies are deflected in somewhat different directions. After passing through a focusing system, the transverse bunch size will therefore be larger than the focused size of a monochromatic bunch at the design energy. This can especially be an issue for electrons from a laser-wakefield accelerator, with energy spreads on the 1-% level. Chromatic effects should therefore be taken into account — and if possible minimized — when designing quadrupole focusing systems for light sources driven by LWFA-generated electrons.

Synchrotron radiation sources commonly employ lattices comprising quadrupoles with alternating focusing directions that are interspersed with drift spaces, known as "FODO lattices", that provides continuous transverse focusing [124]. Moreover, this focusing is periodic, whereby the Twiss parameters at the end of each FODO cell[7] are equal to those at the beginning, provided they are initially matched. This introduces the notion of a 'matched β-function' for a particular focusing configuration. In a storage ring the characteristics of an initially matched bunch can be preserved over large numbers of circulations around the ring.

In the case of a compact light source driven by bunches from a laser-wakefield accelerator, however, FODO lattices are unlikely to be employed, at least in the first instance. All the proposals outlined in Chapter 3 assume that radiation is generated in a single undulator, with a single set of beam optics (a quadrupole doublet or triplet) present between the accelerator and the undulator — rather than a series of undulators interspersed with beam optics. There are several reasons for this, apart from the obvious savings of cost and space. One of the arguments for a single-undulator setup is that the transverse bunch emittance is predicted to be low enough to allow a single set of focusing optics to focus the bunch to a sufficiently narrow transverse size ($\sim 30\,\mu\mathrm{m}$) that can be maintained over a sufficiently long distance ($\sim 2\,\mathrm{m}$) for FEL saturation to occur (cf. Table 3.3). Another argument is that in light of shot-to-shot fluctuations in a LWFA, it would be difficult to consistently produce bunches with parameters — such as transverse size and divergence, but also energy — that comply with the more restrictive matching conditions along the full length of a multi-stage focusing system.

---

[7] A single cell comprises a lens that is focusing in x (defocusing in y), a drift space, a lens that is defocusing in x (focusing in y), followed by another drift space.

The research group of Florian Grüner at the Max-Planck-Institut für Quantenoptik in Garching, Germany has in recent years carried out research towards the realization of a "table-top" FEL source [39]. A significant outcome of that research has been the development — together with external collaborations — of miniature permanent-magnet quadrupoles with very high field gradients, $k_\mathrm{q} \approx 500\,\mathrm{T/m}$ [125, 126]. Refs. [39, 125] describe particular focusing schemes employing pairs of these quadrupoles, while Ref. [100] presents an experimental analysis of this focusing in practice. We direct the reader to these references for details of concrete quadrupole arrangements that may be used for focusing in a compact light source. For a wealth of further discussion on emittance, focusing, and bunch dynamics in general, we moreover direct the reader to the excellent texts by Wiedemann [120] and Humphries [121].

2.2.3 Space-charge considerations

Electron bunches produced in a laser-wakefield accelerator can have charge of order 100 pC confined within a volume with micrometer-scale dimensions both longitudinally and transversely. This raises the question of possible degrading effects due to space-charge forces. These effects have been the subject of several recent studies [39, 127–130], where a range of theoretical and computational approaches has been applied to the problem. Space-charge forces were also the subject of an internal study which was carried out under the guidance of the author [131], the key conclusions of which are also published in Ref. [132]. Here we present a brief summary of the main findings in these publications.

In the longitudinal dimension, in its rest frame the bunch size is relativistically expanded by its Lorentz factor to $\sigma_z' = \gamma\sigma_z$, whereby the space-charge force is greatly reduced [39]. Under these circumstances, even for ultra-high-current bunches with $I_\mathrm{p} \sim 50\,\mathrm{kA}$, as long as the bunch energy is moderately high, $\gamma \sim 600$, the induced energy chirp and the increase in bunch length after approx. 1 m of propagation are both at just the few-percent level [129]. While this level of bunch expansion is not notable, the energy chirp may be an issue in the context of driving a free-electron laser, where energy spreads below 1 % are normally required (see Secs. 2.3.2, 3.2.2). Refs. [39, 129] offer some suggestions on how to counter this chirp, for example by driving the laser-wakefield accelerator just past dephasing, so that the bunch emerges with a negative energy chirp that is then negated by the positive chirp induced by the space-charge forces.

In the transverse dimension, on the other hand, the bunch is not relativistically expanded and the contribution of space-charge effects to transverse growth may be more significant. It has been demonstrated, however, that space-charge driven expansion is only significant during the very earliest stages of bunch evolution after exiting the plasma, before it has expanded to a transverse size of $\sim 100\,\mu\mathrm{m}$ at which point space-charge forces again become negligible [129, 131]. This is again only an issue for bunches of extremely high charge, $\sim 0.5\,\mathrm{nC}$. Moreover, any transverse Coulomb expansion has already taken place at the point at which the divergence of the bunch is experimentally measured, therefore current diagnostics already account for space-charge induced divergence.

With these considerations in mind, and on the basis of theoretical evidence in the currently available body of work on the subject, it appears that space-charge effects should be insignificant for the regimes that we go on to consider. In particular, they are not expected to strongly manifest for bunch charges below approx. 100 pC and bunch energies $\gtrsim 200\,\mathrm{MeV}$. We will therefore not treat space-charge effects further in our analysis of bunch evolution. We now go on to a discussion of the final components of a compact light source.

2.3 The generation of undulator radiation

Having followed the path from injection and acceleration through beam transport, the electron bunch now reaches its ultimate goal — the generation of radiation in an undulator. The present section provides a theoretical description of the generation of undulator radiation, extending to free-electron lasers, in order to allow us to make informed predictions and suggest optimized parameters for a future lights source driven by electrons from a laser-wakefield accelerator. In the following subsection, we describe the motion of electrons through an undulator, and derive the spatial and spectral characteristics of the emitted incoherent radiation. Subsequently, we assess the possibility of feedback of the emitted radiation on the electrons, and describe how this may lead to exponential gain in the radiation intensity, in the process harnessed by free-electron lasers.

The final subsections give a treatment of transverse bunch dynamics in the undulator, as well as a brief overview of presently available undulator technology. Here we use the term “undulator” to refer to periodic magnetic lattices of any strength, despite a distinction in some literature between “undulators” and “wigglers”, where the latter term is used for lattices with stronger fields.

2.3.1 Incoherent undulator radiation

Figure 2.6 shows schematically the propagation of an electron through an undulator. Here we consider an undulator with a plane-polarized magnetic field that has an on-axis magnitude

$$B_y(z) = B_0 \cos(k_\mathrm{u} z),$$

where $k_\mathrm{u} = 2\pi/\lambda_\mathrm{u}$, λ_u is the undulator period, and B_0 is the peak field strength. The treatment that follows is similar to that for a helical undulator (which is carried out in e.g. [133]), and the main differences will be highlighted in the text.

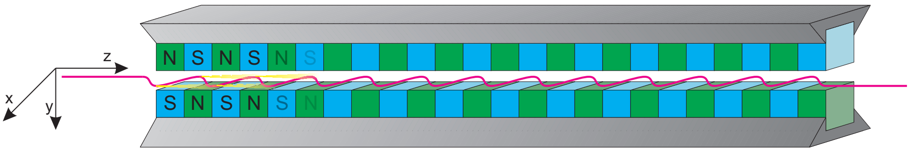

Figure 2.6: The path of an electron (red) as it traverses the alternating magnetic fields of an undulator assembly. The electron propagates along the z-axis, while a sinusoidal magnetic field oriented along the y-axis deflects it alternatingly towards $+x$ and $-x$.

If, as a first approximation, we take the speed of the electron along z to be c, we can write the Lorentz equation for acceleration along x as

$$\gamma m \frac{\mathrm{d}v_x}{\mathrm{d}t} = -ev_z B_y \approx -ecB_0 \cos(k_\mathrm{u} z),$$

which can be integrated to give

$$v_x = -\frac{\sqrt{2} c a_\mathrm{u}}{\gamma} \sin(k_\mathrm{u} z). \tag{2.35}$$

Here we have introduced the normalized rms undulator strength parameter,

$$a_\mathrm{u} = \frac{eB_0}{\sqrt{2} m c k_\mathrm{u}}. \tag{2.36}$$

It should be noted that in a significant part of the literature on the topic, the peak undulator parameter, $K = \sqrt{2} a_\mathrm{u}$, is used instead. As alluded to previously, when this parameter is much greater than unity the term “wiggler” is sometimes used instead of “undulator”. This corresponds to a regime where the maximum deflection angle of the electrons away from the axis, $\sim \sqrt{2} a_\mathrm{u}/\gamma$, is greater than the emission angle of the radiation, $\sim 1/\gamma$ [134]. While some of the radiation characteristics do indeed differ between the two regimes, the fundamental emission processes are equivalent, and therefore we simply use the designation “undulator” throughout.

Assuming no significant deflection along y, from the definition of $\gamma = 1/\sqrt{1-|\mathbf{v}|^2/c^2}$ we obtain

$$v_z = c\sqrt{1 - \frac{1}{\gamma^2} - \frac{v_x^2}{c^2}}.$$

This relation stems from the fact that a static magnetic field cannot do work on a charged particle, therefore γ is conserved. Substituting v_x from Eq. (2.35) and carrying out a binomial expansion of the square root, the expression becomes

$$v_z \approx c\left(1 - \frac{1+a_\mathrm{u}^2}{2\gamma^2}\right) + \frac{a_\mathrm{u}^2 c}{2\gamma^2} \cos(2k_\mathrm{u} z). \tag{2.37}$$

The first (bracketed) term corresponds to the overall speed of propagation of the electron along z, which is reduced by about $(a_\mathrm{u}^2/2\gamma^2)c$ with respect to the speed of direct propagation in the absence a deflecting structure. Based on this, we introduce the normalized mean longitudinal velocity

$$\beta_{z0} = \langle v_z \rangle / c = 1 - \frac{1+a_\mathrm{u}^2}{2\gamma^2}.$$

The second term is an oscillation in the speed along z that is coupled to the magnitude of v_x, corresponding to the characteristic “figure-8” motion of the electron in x-z space. This oscillation is the source of higher odd harmonics of the fundamental emission wavelength, and is not present in the case of a helical undulator.

As we established, an electron of nominal energy $\gamma m c^2$ is moving through the undulator at a mean speed along z of $\beta_{z0} c$. The corresponding frame of reference moving at this mean speed has a Lorentz factor of

$$\gamma^* \approx \frac{\gamma}{\sqrt{1+a_\mathrm{u}^2}},$$

see e.g. [124, 134]. In this frame, the electron will experience a magnetic lattice that is relativistically contracted in the longitudinal direction by a factor γ^*. The transverse motion of the electron will therefore have a period of $\lambda^* = \lambda_\mathrm{u}/\gamma^*$, corresponding to the emission of dipole radiation of frequency

$$f^* = \frac{c\gamma^*}{\lambda_\mathrm{u}}.$$

Transforming back to the lab frame, an on-axis observer downstream from the undulator will see this emission with an additional relativistic Doppler shift:

$$f = \frac{f^*}{\gamma^*(1-\beta_{z0})} = \frac{c}{\lambda_\mathrm{u}(1-\beta_{z0})}. \tag{2.38}$$

Finally, since for ultra-relativistic electrons, $(1-\beta) \approx 1/(2\gamma^2)$, this frequency corresponds to a wavelength of emission of

$$\lambda_\mathrm{R} = \frac{c}{f} = \frac{\lambda_\mathrm{u}}{2\gamma_{z0}^2} = \frac{\lambda_\mathrm{u}}{2\gamma^2}\left(1+a_\mathrm{u}^2\right). \tag{2.39}$$

This is the undulator equation in its most basic form, describing the emission by a relativistic electron traveling on-axis through a quasi-infinite undulator. It can be extended to account for off-axis observation by multiplying β_{z0} in equation (2.38) by a factor $\cos\theta$, where θ is the angle of observation. This yields

$$\lambda_\mathrm{R} = \frac{\lambda_\mathrm{u}}{2\gamma^2}\left(1+a_\mathrm{u}^2+\gamma^2\theta^2\right), \tag{2.40}$$

i.e. photons emitted off-axis are less energetic. We go on to discuss further modifications that adapt the undulator equation to more realistic operating conditions.

Due to the finite time of travel of the electron through the undulator, the emitted radiation pulse will have a finite duration and thus a finite bandwidth. An electron making N_u oscillations in an undulator of length $N_\mathrm{u}\lambda_\mathrm{u}$ will emit a wavetrain with that same number of periods, corresponding to a Fourier-transform-limited spectrum with a sinc^2 distribution around the central frequency. The full width at half-maximum of this spectrum is $\Delta\omega/\omega_\mathrm{R} \approx 0.89/N_\mathrm{u}$, where $\omega_\mathrm{R} = 2\pi c/\lambda_\mathrm{R}$ is the central angular frequency of emission.

Furthermore, while equation (2.39) is limited to a single electron of energy γmc^2 nominally propagating along the undulator axis, a realistic electron bunch will not be purely monoenergetic (or “cold”), nor mono-directional. Both the finite energy spread and the finite angular divergence of the bunch will affect the radiated spectrum in similar ways, as they both cause a spread in the rates of propagation along z. In other words, β_{z0} and correspondingly γ^* now take a range of values. For any given electron, a deviation from the central energy by $\Delta\gamma$ leads to a radiation frequency deviation of $\Delta\omega/\omega_\mathrm{R} = 2\Delta\gamma/\gamma$, while an angular deviation of α with respect to the z-axis leads to a reduction in radiation frequency of $\Delta\omega/\omega_\mathrm{R} = -\gamma^{*2}\alpha^2$.

The true linewidth of on-axis undulator emission at the fundamental wavelength will be the result of convolving the spectral broadening effects outlined here, as well as others that may result from e.g. nonuniformities in the undulator magnetic field. In particular, we

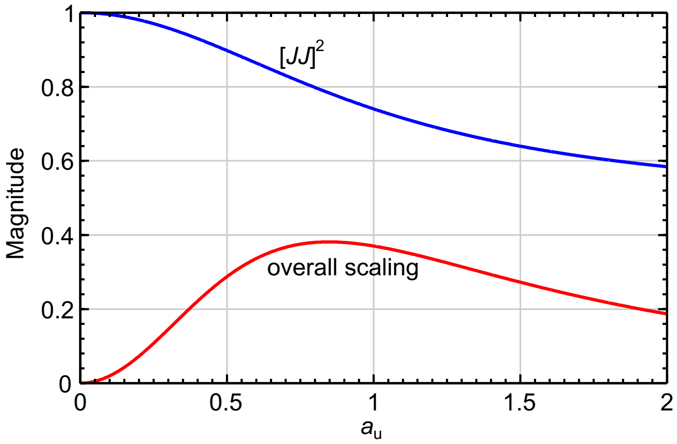

Figure 2.7: Scaling of $[JJ]^2$ with a_u, as well as overall dependence on a_u of energy emitted into the central cone and natural bandwidth at the first harmonic (that is, $[JJ]^2 \cdot 2a_\mathrm{u}^2/(1+a_\mathrm{u}^2)^2$). For details refer to text, in particular Eq. (2.43).

have assumed that the magnetic field has only a z dependence, while in reality (even in the ideal case), it will vary with transverse position as well. This means electrons entering the undulator at different transverse positions will emit differently, leading to additional modifications in the spectral and spatial characteristics of the radiation. For further information on transverse effects, including transverse focusing effects of the undulator, we direct the reader to Reiche's treatment in Ref. [133, Sec. 2.1], and references therein.

As well as emission at the fundamental, electrons traveling through a planar undulator will emit higher odd harmonics due to the z-dependent variation of v_z described in equation (2.37). This can be illustrated by integrating equation (2.35) to first order (i.e., assuming $z = \beta_{z0}ct$) in order to obtain the transverse electron position as a function of time, t,

$$x(t) = \frac{\sqrt{2}a_\mathrm{u}}{\gamma k_\mathrm{u}\beta_{z0}}\cos(k_\mathrm{u}v_z t). \tag{2.41}$$

Substituting the expression for v_z from Eq. (2.37) here, and again only retaining first order terms in t, we obtain

$$x(t) = \frac{\sqrt{2}a_\mathrm{u}}{\gamma k_\mathrm{u}\beta_{z0}}\cos\left[\zeta\left(1+\frac{a_\mathrm{u}^2}{2\gamma^2\beta_{z0}}\cos(2\zeta)\right)\right],$$

where $\zeta = k_\mathrm{u}\beta_{z0}ct$ is the oscillation phase of the bunch along the undulator. One of the approaches to analyzing this expression is to express the cosines as sums of exponentials and expand the outer ones into Taylor series. One thus obtains a series of exponentials of the form $e^{i(2n+1)\zeta}$, where n is an integer. These are the odd harmonics of the fundamental that are emitted due to the oscillations of v_z in the planar undulator, and their emission amplitudes are strongly dependent on a_u (for a quantitative treatment see Kim's 1989 monograph [135]). Even harmonics are also emitted off-axis, and are again significantly more intense for higher a_u. In a helical undulator the longitudinal velocity is constant, therefore the odd on-axis harmonics are generally not present, though harmonic emission can still be observed off-axis, especially for larger a_u.

The total energy radiated over all frequencies and into all directions by an electron bunch with N_e electrons propagating through an undulator with N_u periods is

$$E_\mathrm{rad} = \frac{N_\mathrm{e} e^2 \gamma^2 N_\mathrm{u} k_\mathrm{u}^2 a_\mathrm{u}^2}{3\epsilon_0}, \tag{2.42}$$

as presented in the detailed work by Hofmann [124]. More specifically, the total energy emitted into the central radiation cone, $\theta_\mathrm{cen} = 1/(\gamma^* \sqrt{N_\mathrm{u}})$, within the natural relative spectral bandwidth, $\Delta\omega/\omega_0 = 1/N_\mathrm{u}$ at just the 1st harmonic, is [134]

$$E_\mathrm{rad,1} = \frac{N_\mathrm{e} e^2 \gamma^2 k_\mathrm{u}}{\epsilon_0} \frac{a_\mathrm{u}^2}{(1 + a_\mathrm{u}^2)^2} [JJ]^2. \tag{2.43}$$

Here $[JJ]$ is a coupling factor that quantifies the fraction of radiation emitted at the fundamental, versus the higher harmonics. For a planar undulator, $[JJ] = J_0(\xi) - J_1(\xi)$, where $\xi = a_\mathrm{u}^2/(2 + 2a_\mathrm{u}^2)$ and $J_{0,1}$ are Bessel functions. Figure 2.7 shows a plot of $[JJ]^2$ as well as the overall scaling of the emission for $0 < a_\mathrm{u} < 2$. Since for the case of a helical undulator there is no on-axis emission at the high harmonics, there $[JJ] = 1$.

In Ref. [124], Hofmann also presents full analytical derivations of the spatial, angular and spectral characteristics of undulator radiation, complementing previous work by Kim [135]. For the purposes of the present discussion, we will use numerical means to further determine these characteristics. In particular, it is possible to obtain the emission properties for an arbitrary electron bunch propagating through any undulator by integrating the Liénard-Wiechert potentials of the oscillating electrons. This method is implemented in the SPECTRA code [136].

Figure 2.8, for instance, shows the spectrum from a bunch of 500 MeV electrons through a 2 cm-period ($a_\mathrm{u} = 1.06$) undulator, both on-axis and integrated over a 1 mrad acceptance half-angle. The difference between the two curves clearly highlights (a) off-axis spectral shifting; and (b) rich off- and on-axis harmonic content. Having described the emission principle of incoherent undulator radiation, we continue with a discussion of the free-electron laser regime.

2.3.2 Free-electron lasers

Undulator radiation is incoherent by nature. Barring special manipulation, the longitudinal distribution of electrons in the bunch will be more or less homogeneous, therefore the phases of radiation emitted by all electrons will not be correlated. Correspondingly, the output energy scales linearly with electron bunch charge. When the system meets certain criteria, however, the emission can become coherent in a regime known as a free-electron laser. This can significantly enhance output energy level.

In a free-electron laser the spontaneous undulator radiation impacts a degree of longitudinal *microbunching* on the electrons, on the scale of the radiation wavelength. The microbunching imparts a degree of coherence on the newly emitted photons, leading to an increase in radiation intensity, which in turn increases the degree of microbunching. This positive feedback process leads to an exponential gain in power output as the electrons travel along the undulator.

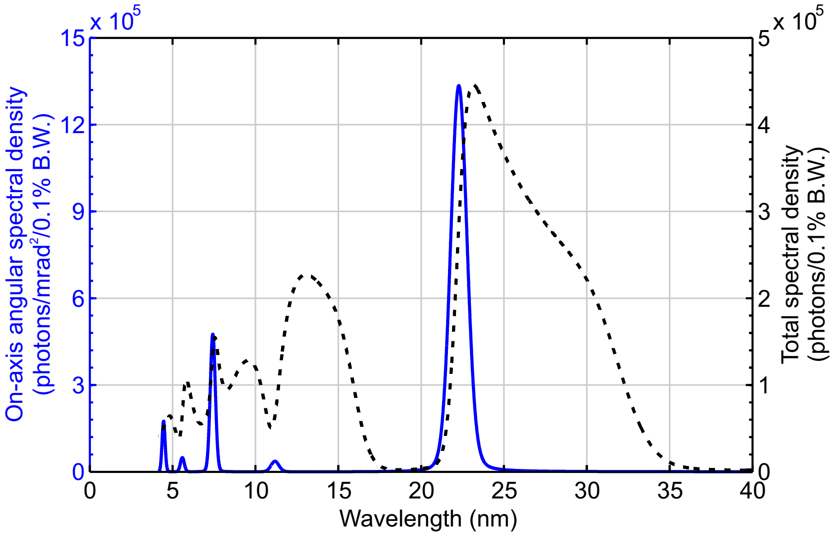

Figure 2.8: Undulator spectra from a 100 pC bunch of 500 MeV electrons (1% rms energy spread) with 1 mm · mrad normalized emittance, focused at the center of a 1 m long undulator of period 2 cm with $a_\mathrm{u} = 1.06$. Blue line represents on-axis photon count per unit solid angle (per 0.1% bandwidth); dashed line represents photon count integrated over a solid angle of $1\pi\,\mathrm{mrad}^2$.

Microbunching due to a resonant electric field

To illustrate the basic physics, let us consider the way an electric field affects the energy of electrons propagating through the undulator. From the zeroth component of the covariant form of the Lorentz force, we obtain

$$\frac{\mathrm{d}\gamma}{\mathrm{d}t} = \frac{e}{mc^2}\mathbf{E}\cdot\mathbf{v}, \tag{2.44}$$

where $\mathbf{E}$ is the electric field vector [133, p. 20]. Limiting ourselves to electron motion in the x-z plane and an electric field polarized along x, the electrons will gain energy whenever E_x and v_x are parallel, or lose energy whenever they are antiparallel. In addition, let us remind ourselves that the electrons nominally propagate along z at a speed of $\beta_{z0}c = \left[1-(1+a_\mathrm{u}^2)/(2\gamma^2)\right]c$, while the radiation field[8] will propagate at c. The electrons therefore slip behind the radiation at a rate of exactly one resonant wavelength, λ_R, for each undulator period, $\Delta z = \lambda_\mathrm{u}$ (cf. Eq. (2.39)). These considerations underpin the microbunching process, as illustrated in Figure 2.9.

Let us consider the sample electrons in Fig. 2.9. In frame (1) of panel (a), both electrons have $v_x > 0$, however they interact with different phases of the radiation field. This causes the one colored cyan to lose energy, while the one colored magenta gains energy — corresponding to the relative directions of E_x and v_x for each. In frame (2) the electron phases with respect to the undulator field have advanced by π and they both now have $v_x < 0$, however the electric field has slipped over them by π and reversed sign. This means the cyan-colored electron is *still* losing energy, while the magenta-colored one continues to gain energy. In frame (3) the phases have advanced by π once again and the situation of frame

---

[8] For the moment we ignore diffraction effects.

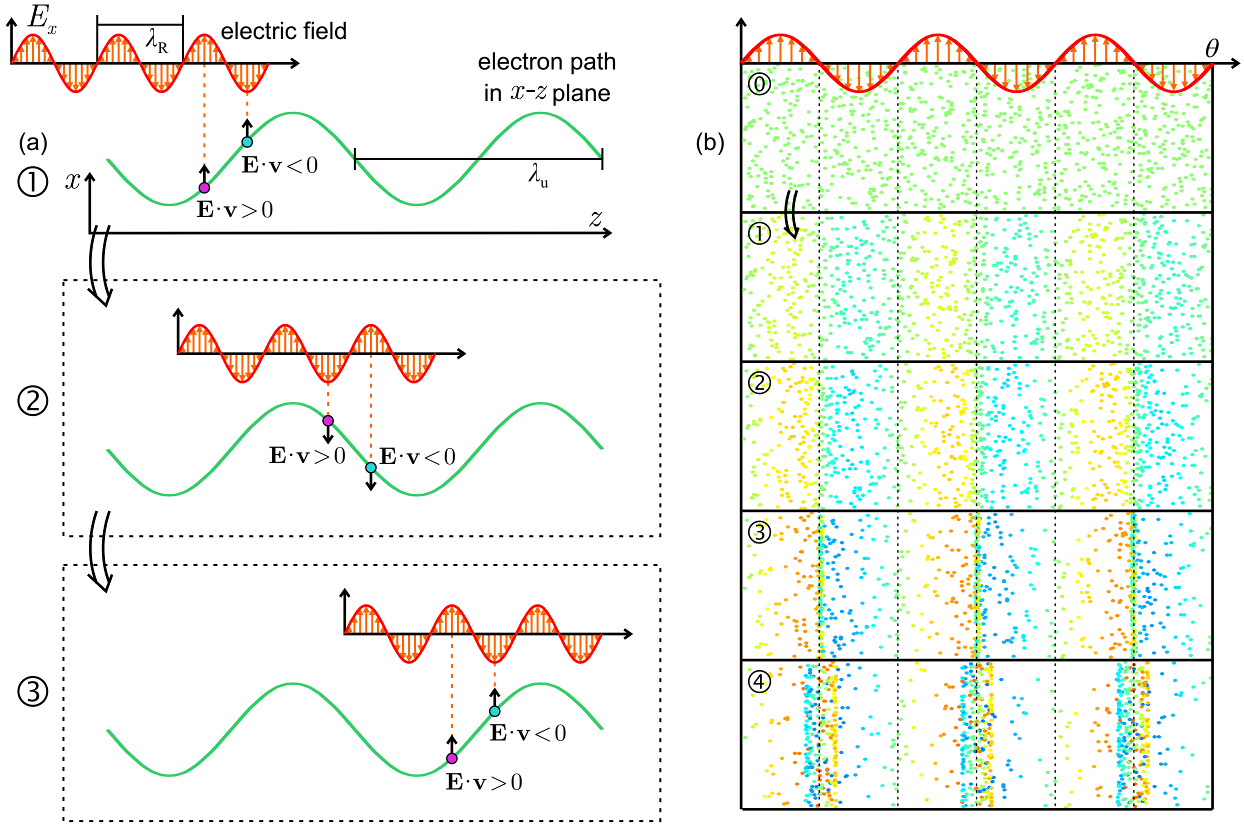

Figure 2.9: (a) The left panel illustrates qualitatively how an electric field of wavelength λ_R leads to an energy modulation. The electrons, represented by magenta and cyan circles, follow a quasi-sinusoidal path in x (green line) as they propagate along z. The radiation field, pictured in red, propagates along with them, but slips forward of the electrons by one radiation wavelength (λ_R) for each advance of the electrons by an undulator period (λ_u). (b) The right panel illustrates the evolution of the relative energies and positions of sample electrons at different initial phases with respect to the radiation. Electrons at the initial energy are colored green, while red signifies higher energy and blue signifies lower. Detailed explanations of both panels are given in the text.

(1) is repeated, though the electrons are now interacting with the next cycle of the radiation field.

Consequently, for a radiation field of wavelength λ_R with constant amplitude and phase, the energy modulation is enhanced over subsequent undulator periods. The accumulated modulation is sinusoidal, with a period equal to that of the radiation. This is illustrated in panel (b) of Fig. 2.9, where sample electron positions are shown with respect to the frame moving at $\beta_{z0}c$. The evolution of the energy modulation of each electron (represented by its color), and the corresponding position changes are shown at different points along the undulator in frames (0)–(4). A snapshot of the electric field that interacts with the bunch is shown for reference. In this case the snapshot is of the field overlapping the bunch when it is at $z = \lambda_\mathrm{u}/4$, i.e., when its velocity along x is maximum (cf. Eq. (2.35)), which is the situation depicted in frame (1) of panel (a).

The energy modulation can be thought of as arising from a ponderomotive potential that is the result of the combined effects of the undulator and the radiation. Since this potential is constant in the reference frame moving together with the electron bunch (i.e., at $\beta_{z0}c$), we introduce the ponderomotive phase

$$\theta = (k_\mathrm{R} + k_\mathrm{u})z - k_\mathrm{R}ct. \tag{2.45}$$

The coordinates commonly used when characterizing the free-electron laser interaction are z, the distance the electron bunch has traveled along the undulator, and θ, the position within the electron bunch at which an event takes place. Somewhat counter-intuitively, then, the z coordinate indicates the stage of evolution of the system.

The interaction described above can be quantified by substituting the appropriate values of E_x and v_x in equation (2.44). Expressing the radiation field in terms of an envelope and a carrier wave of wavenumber k and phase ψ,

$$\begin{aligned} E_x &= E_0 \cos\left[k(z-ct)+\psi\right] \\ &= \frac{mc^2k}{e} a_{\mathrm{r}0} \cos\left[k(z-ct)+\psi\right], \end{aligned}$$

where $a_\mathrm{r} = eE_x/(mc^2k)$ is the dimensionless magnetic vector potential of amplitude $a_{\mathrm{r}0}$, we obtain

$$\frac{\mathrm{d}\gamma}{\mathrm{d}t} = \frac{\sqrt{2}cka_\mathrm{u}a_{\mathrm{r}0}}{\gamma} \cos\left[k(z-ct)+\psi\right] \sin(k_\mathrm{u}z). \tag{2.46}$$

Ignoring the constant phase ψ, the product of the trigonometric functions is oscillating as $(k \pm k_\mathrm{u})z - kct$. For a resonant build-up of energy modulation this, therefore, needs to be zero. Setting $z = \beta_{z0}ct$, the resonance condition becomes

$$\beta_{z0} = \frac{k}{k \pm k_\mathrm{u}}.$$

Ignoring the '−' sign in the denominator since it requires $v_z > c$, we note that the resonance condition $\beta_{z0} = k/(k+k_\mathrm{u})$ is exactly matched when $k = k_\mathrm{R} = 2\pi/\lambda_\mathrm{R}$ (cf. Eq. (2.39)). That is, the radiation wavelength required to induce a resonant energy modulation is the same as the resonant wavelength of the undulator radiation emitted by the electrons. This reinforces the pictorial treatment provided earlier.

FEL performance metrics

Since radiation emitted by monoenergetic electrons in an undulator can induce microbunching of the electrons with a period equal to the radiation wavelength, the degree of coherence of the emitted radiation can increase, thereby leading to a power output that grows exponentially in z. While this process is analyzed in detail in Chapter 6 (and further literature, such as Refs. [133, 137]), here we provide some simple metrics that can be used to characterize the performance of an FEL. These are also outlined elsewhere, e.g., by Pellegrini in Ref. [138] in relation to the proposed (at the time) operation of the LCLS.

The high-gain FEL collective instability that leads to coupled exponential enhancement in electron microbunching and radiation intensity was identified in the late 70s [139] and developed theoretically through the early 80s [140–142], with Haus [140] providing the first treatment of start-up from noise. The 1-D model of Bonifacio, Pellegrini, and Narducci [142] has proved particularly valuable in underpinning further theoretical developments. I this model the Pierce parameter,

$$\rho = \frac{1}{2\gamma}\left[\frac{I_\mathrm{p}}{I_\mathrm{A}}\left(\frac{a_\mathrm{u}[JJ]}{k_\mathrm{u}\sigma_r}\right)^2\right]^{1/3}, \tag{2.47}$$

plays a central role. Broadly speaking, this parameter characterizes the strength of the coupling of the electron beam to the radiation field in the undulator. In the context of this model, we will refer to electron *beam* rather than *bunch*, since the model assumes uniformity in the longitudinal dimension. Here I_p is the peak electron beam current; $I_\mathrm{A} = mc^3/e \approx 17\,\mathrm{kA}$ is the Alfvén current; σ_r is the transverse rms electron beam size; and $[JJ]$ is a factor that was defined after equation (2.43). Within this model the radiation intensity grows as

$$I \approx \frac{I_0}{9}\exp(z/L_\mathrm{G}) \tag{2.48}$$

where z is the distance along the undulator and L_G is the exponential gain length, given by

$$L_\mathrm{G} = \frac{\lambda_\mathrm{u}}{4\sqrt{3}\pi\rho}, \tag{2.49}$$

and I_0 is the spontaneous coherent radiation intensity from an undulator of length L_G. Even if the electrons have not been explicitly manipulated to induce microbunching prior to their entry into the undulator, the "shot noise" — the random fluctuations in their discrete positions — will lead to a small degree of coherent emission that will then be amplified by the FEL in what is known as *self-amplification of spontaneous emission* (SASE). The number of such coherent photons spontaneously emitted by N_e electrons is $N_\mathrm{ph} = \pi\alpha N_\mathrm{e} a_\mathrm{u}^2/(1+a_\mathrm{u}^2)$, where α is the fine structure constant [138].

Saturation is reached when the energy radiated by the electrons causes them to lose sufficient energy so as to no longer radiate within the FEL bandwidth[9] $\Delta\lambda_\mathrm{R} = \rho\lambda_\mathrm{R}$. This results in a saturation power of

$$P_\mathrm{sat} = \rho P_\mathrm{beam}, \tag{2.50}$$

where $P_\mathrm{beam} = \langle\gamma\rangle I_\mathrm{p} mc^2/e$ is the electron beam power. The saturation length, L_sat, is then the distance at which the emitted power is P_sat, normally $L_\mathrm{sat} \approx 20L_\mathrm{G}$.

While the radiation field slips over the electrons by one resonant wavelength, λ_R, for each undulator period, λ_u, the total slippage over one gain length is the so-called cooperation length [146],

$$L_\mathrm{c} = \frac{\lambda}{\lambda_\mathrm{u}}L_\mathrm{G}, \tag{2.51}$$

which is the distance over which a phase correlation between the electrons and field is established. The temporal structure of the FEL output thus comprises a sequence of longitudinally coherent spikes of duration $\Delta t = L_\mathrm{c}/c$. Transverse coherence is similarly established

---

[9] It is possible to extract more energy by tapering the undulator period to shift the resonance to lower electron energies, see e.g. Fawley *et al.* [143] for a proposal, and Ratner *et al.* [144] or Giannessi *et al.* [145] for discussions of experimental implementations.

through the diffraction of the radiation field over a gain length. Coherence-related aspects of spontaneous FEL radiation are further discussed in Ref. [147].

While radiation diffraction enhances transverse coherence, efficient gain cannot occur if the diffraction is too great, imposing the condition

$$L_{\mathrm{G}} < Z_{\mathrm{R}}. \tag{2.52a}$$

Here Z_{R} is the Rayleigh range, $Z_{\mathrm{R}} = \pi w_0^2/\lambda_{\mathrm{R}}$, where the radiation width, w_0, can be taken to equal the beam size, σ_r, under equilibrium. Additionally, the electron beam should not diverge faster than the radiation, setting the following condition on its transverse emittance:

$$\varepsilon < \frac{\lambda_{\mathrm{R}}}{4\pi}. \tag{2.52b}$$

The final condition that delineates the validity of this 1-D model concerns the energy spread of the electrons,

$$\frac{\sigma_\gamma}{\gamma} < \rho, \tag{2.52c}$$

since only electrons that have energies within a factor ρ of the resonant energy will emit at wavelengths that undergo exponential FEL gain.

While the model described above captures the key physics, it can give misleading results if the parameters of operation are not all well within the limits set out above. A full three-dimensional treatment was carried out by Ming Xie [148] by solving the eigenmode equations and carrying out a variational analysis to obtain a set of scaling laws[10]. The model allows the derivation of a scaling constant Λ that aims to capture the degrading effects on FEL gain of radiation diffraction, transverse electron bunch emittance, and energy spread. This constant relates the 3-D solutions to the 1-D ones as $\rho_{3\mathrm{D}} = \rho/(1+\Lambda)$, and correspondingly $L_{\mathrm{G,3D}} = (1+\Lambda)L_{\mathrm{G}}$. In the next chapter we will use Xie's scaling in order to quickly scan a large parameter space and identify the most favorable working points for a proposed FEL, before carrying out more detailed simulations.

Seeding with high harmonics

The longitudinal coherence length of FEL radiation can be extended beyond L_{c} via *seeding* by an appropriately coherent external radiation source. The presence of a sufficiently intense[11] seed ensures that the phase of microbunching is the same along the entire electron bunch — or at least to the extent to which the seeding source is temporally coherent. However, conventional lasers operating at wavelengths shorter than ∼ 200 nm wavelength are not readily available for use as seed sources. Recently there has been great interest at numerous laboratories in seeding FELs directly with coherent radiation produced by high harmonic generation (HHG) from intense visible laser pulses interacting with a gas. This scheme has been tested experimentally in single pass FEL amplifiers to wavelengths as

---

[10] Also see references therein for other FEL analyses in three dimensions.

[11] The seed needs to be more intense than the shot noise emission by the electrons. Giannessi derives a quantitative criterion in Ref. [149].

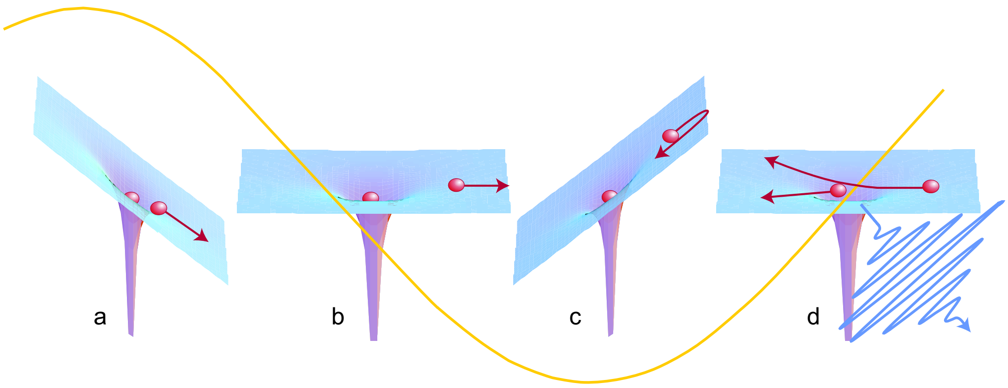

Adapted with permission from Macmillan Publishers Ltd: *Nature Physics* [152]. Copyright © 2007.

Figure 2.10: (a) An intense drive laser field (yellow) extracts an electron wavepacket (red circle) from an atom or molecule (blue potential). The electron is pulled away from the atom and accelerated in the continuum, then driven back towards the atom when the field reverses (b,c). During a small fraction of the laser oscillation cycle, the electron can recollide (d). During this recollision the kinetic energy gained by the electron while in the continuum can be converted to an attosecond optical pulse, which adopts the amplitude and phase imparted on the electron wavepacket during interaction with the laser.

short as ~ 60 nm (13th harmonic of Ti:sapphire) [150, 151] and, in principle, can work down to the few-nm regime. We will further consider the implications of HHG seeding on FEL performance in Chapter 6, here we briefly introduce the high harmonic generation process and the characteristics of the generated short-wavelength radiation.

In HHG an intense drive laser pulse ionizes atoms in a low-density gas medium. Each ionized electron is accelerated in the continuum by the laser's electric field and can recollide with the parent ion. Upon recombination, the energy gained by the electron due to the laser field is emitted in the form of a high-energy photon. This process is described in considerable detail in a review article by Corkum and Krausz [152], from which we have adapted Figure 2.10.

During recombination, the most energetic photons are emitted by electrons ionized at a fixed offset from the extrema of the driving pulse. Due to this condition, HHG emission nominally comprises a series of sub-femtosecond spikes, one for each half-cycle of the driving laser. This corresponds to a spectrum containing all odd harmonics of the driving frequency, up to a cut-off determined by the wavelength and intensity of the driving laser, as well as the ionization potential of the gas [152, 153].

Apart from seeding with high harmonics of an optical driving laser, it is also possible to directly use an optical laser seed to modulate the electron bunch energy in one undulator (the "modulator"), followed by a dogleg in which the induced energy modulation is converted to a spatial modulation — i.e., microbunching — and a second undulator (the "radiator") where the electrons radiate at a high harmonic of the initial optical laser [154, 155]. This process is known as high-gain harmonic generation (HGHG), and can be used to efficiently produce emission at wavelengths shorter than the resonant, although it clearly requires a more complex setup than simple HHG seeding.

2.3.3 Transverse electron motion in an undulator

Apart from the quasi-sinusoidal transverse oscillations of the electron bunch as it passes through the undulator, it is also subject to separate "bulk" transverse evolution. This evolution was not considered thus far, partly because its effect can generally be treated as a higher-order correction to the resonant oscillations, and partly because a full treatment can introduce considerable complexity and detract from the fundamental picture.

In Section 2.3.1 we discussed the fact that angular divergence of the electron bunch shifts emission to longer wavelengths, due to the decrease in the axial velocity of the electrons. Here we elaborate on the dynamics of the bunch itself, without reference to the effect on emission — this is treated by Reiche in Ref. [133]. In particular, we discuss focusing of the beam due to the variation in magnetic field strength away from $x = y = 0$, referred to as "natural" undulator focusing.

A realistic magnetic field inside the undulator gap must satisfy Maxwell's equations in a source-free environment, namely $\nabla \cdot \mathbf{B} = 0$ and $\nabla \times \mathbf{B} = 0$. This means the field can be represented in terms of a scalar potential, $\mathbf{B} = -\nabla\phi$, which itself must be a solution of the Laplace equation, $\nabla^2\phi = 0$. A realistic potential that satisfies this for an undulator geometry is

$$\phi = -\frac{B_0}{k_y}\cosh(k_x x)\sinh(k_y y)\cos(k_\mathrm{u} z),$$

where k_x and k_y satisfy $k_x^2 + k_y^2 = k_\mathrm{u}^2$ (see e.g., Refs. [120, 133]). For flat magnetic pole faces $k_x = 0$, for poles bent towards each other it is > 0, and for poles that are bent outwards it is imaginary. In the limit $k_x x, k_y y \ll 1$, i.e., for small excursions of the electrons from the z-axis compared to the undulator period — as is the case generally — the hyperbolic functions can be expanded into Taylor series around $x = y = 0$. To 2nd order in x and y this gives the components of the magnetic field as

$$\mathbf{B} = B_0 \begin{pmatrix} k_x^2 xy & \cos(k_\mathrm{u} z) \\ \left(1 + \frac{k_x^2 x^2}{2} + \frac{k_y^2 y^2}{2}\right) & \cos(k_\mathrm{u} z) \\ -k_\mathrm{u} y & \sin(k_\mathrm{u} z) \end{pmatrix}. \tag{2.53}$$

Working out the trajectories of motion in x and y (see e.g., Ref. [133]), and averaging out any components with variation on the scale of an undulator period, we obtain,

$$\frac{\mathrm{d}x'}{\mathrm{d}z} = -\left(\frac{a_\mathrm{u} k_x}{\gamma}\right)^2 x; \tag{2.54}$$

$$\frac{\mathrm{d}y'}{\mathrm{d}z} = -\left(\frac{a_\mathrm{u} k_y}{\gamma}\right)^2 y, \tag{2.55}$$

where the primed quantities are the angular directions with respect to the z-axis. This leads to harmonic oscillations in x and y at the frequencies in brackets, corresponding to a beam envelope with a β-function of $\beta_\mathrm{T} = (\gamma/a_\mathrm{u})\lambda_\mathrm{u}$ (cf. Sec. 2.2.1). The focusing can be along both dimensions, with e.g., $k_x = k_y = k_\mathrm{u}/\sqrt{2}$ for appropriately curved undulator poles. For flat pole faces the focusing is only along y, with $k_x = 0$ and $k_y = k_\mathrm{u}$. In the next chapter we will exploit this as part of optimizing the radiation generation process.

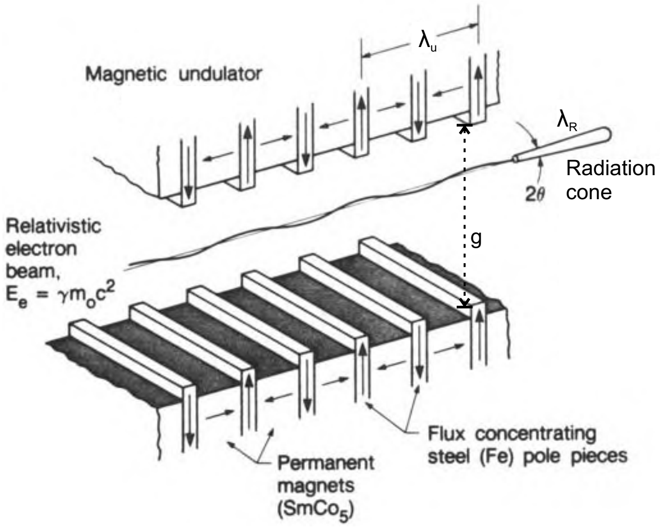

Reproduced with permission from [135]. Copyright © 2000, American Institute of Physics.

Figure 2.11: Schematic of a hybrid undulator based on permanent magnets (e.g., NdFeB or $SmCo_5$) with a period of λ_u. The path of the electron bunch and radiation cone are also depicted.

2.3.4 Undulator technology

Finally, we would like to develop a sense of the range of undulator parameters that are available based on current technologies. Broadly speaking, undulators can be constructed in the following ways: [156, 157]

(a) a pure sequence of permanent magnets of alternating polarity;

(b) a sequence of permanent magnets separated by flux-concentrating poles made of steel (so-called "hybrid undulators", see Fig. 2.11);

(c) same as (a) or (b), with cryogenic cooling applied to the magnets in order to increase the magnetic flux;

(d) a sequence of electromagnets;

(e) a sequence of superconducting electromagnets.

Despite significant recent progress — especially in superconducting technologies — electromagnetic undulators are generally limited to longer periods (> 10 mm) due to the space required by the windings. Due to this, and their more complex nature, we will not further consider options (d) or (e). Cryogenically-cooled undulators, option (c), can offer somewhat improved performance over ones operating at room temperature [156], however the technique is relatively unproven, and as such not optimal for our needs.

Currently the best trade-off between performance, reliability, and simplicity is generally achieved through the use of hybrid undulators, and these will be the basis for discussion here and in the following chapter. Nearly 30 years ago, Halbach expressed the optimum peak magnetic field that can be obtained from a hybrid undulator in the form [158]

$$B_0 = a \exp\left[b \frac{g}{\lambda_\mathrm{u}} + c \left(\frac{g}{\lambda_\mathrm{u}} \right)^2 \right], \tag{2.56}$$

where g is the gap between opposing poles of the undulator, a is a constant with units of Tesla, and b and c are dimensionless constants. In Ref. [157] Elleaume *et al*. give the values of these constants for several different undulator types. Of those listed, the hybrid undulator that can deliver the highest on-axis magnetic field consists of NdFeB magnet blocks and vanadium permendur poles, for which

$$a = 3.694\,\mathrm{T}, \quad b = -5.068, \quad c = 1.520.$$

The peak on-axis field as a function of g/λ_u is plotted for this as well as several other configurations in Figure 2.12. For the pure permanent magnet and the hybrid undulators (cases A–E), the gap can vary in the range $0.1\lambda_\mathrm{u} < g < \lambda_\mathrm{u}$. Further aspects of undulator and beamline technology are covered by Ciocci in Ref. [123].

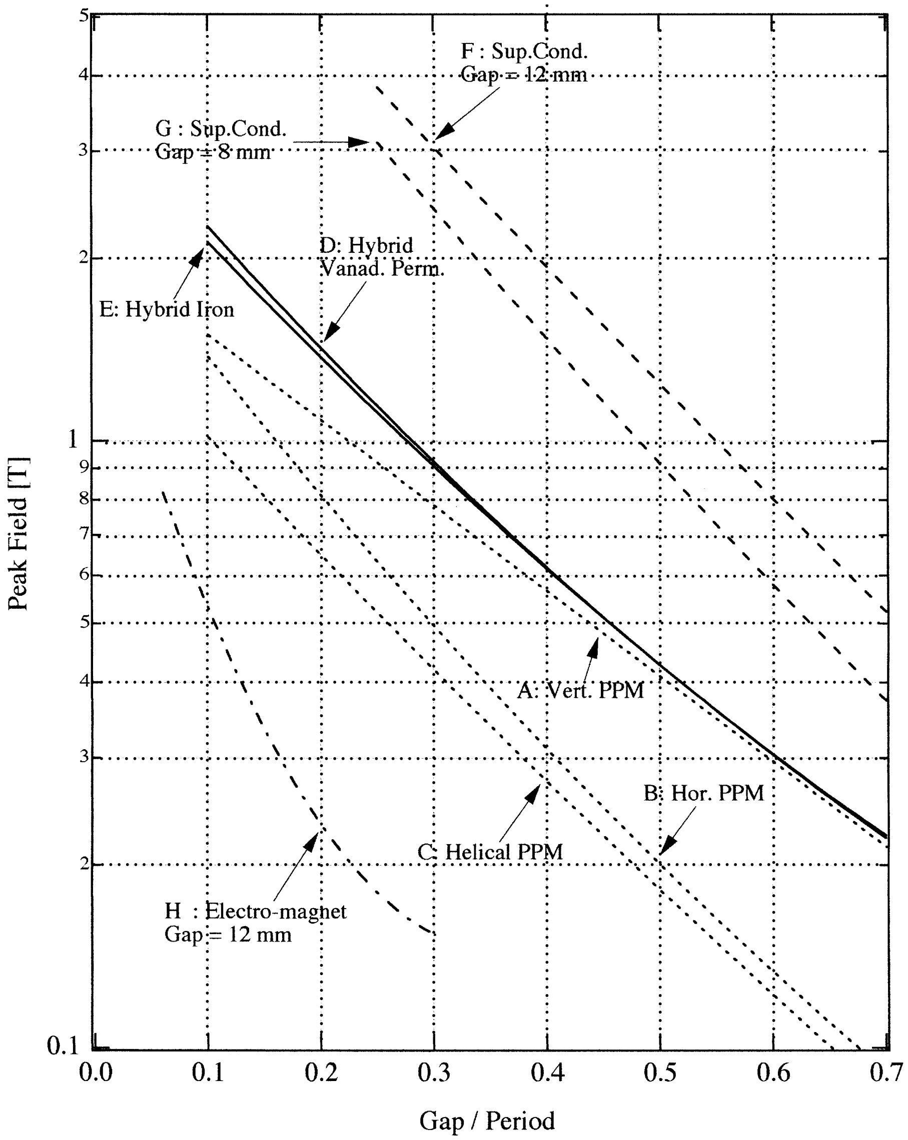

Reproduced with permission from [157]. Copyright © 2000 Elsevier.

Figure 2.12: Peak magnetic fields versus g/λ_u for different undulator types. A–C: pure permanent magnet undulators; D–E: hybrid undulators; F–G: superconducting undulators; H: electromagnetic undulator.

Chapter 3

Parameter Optimization and Simulations for a Compact Light Source

In this chapter we build on the background developed previously in order to establish a suitable parameter range for operation of a light source driven by a laser plasma accelerator, which we refer to here as a “compact light source”. We describe proposals for compact light sources producing incoherent soft x-ray radiation, as well as a free-electron laser. We will carry out initial parameter studies based on the established analytical scaling laws, and subsequently carry out more detailed simulations in order to firm up the proposals.

3.1 Production of incoherent soft x-ray radiation

The scalings for producing incoherent radiation are really rather simple. As shown is Section 2.3.1, the main parameters that determine the characteristics of the output radiation are the electron bunch energy and charge, and the undulator period. Most others — such as electron bunch energy spread, length, and transverse emittance — lead to corrections to the base performance that can be considered separately. This is in contrast to the case of the free-electron laser that will be considered in the following section, where the performance can have strong non-linear dependence on the operating parameters.

Common sense dictates that best performance is achieved for electron bunches with maximum charge, minimum energy spread and minimum transverse emittance — allowing the emission to be intense, have a narrow bandwidth, a small source size, and small divergence. Increases in the transverse emittance or energy spread will generally lead to a deterioration in emission quality, therefore they would ideally be minimized during the injection and acceleration processes. For discussions on how to achieve this, we refer the reader to Section 2.1 and references therein.

Besides these parameters, however, selecting a combination of the optimal undulator parameters (a_u and λ_u) and the optimum electron energy in order generate a given radiation wavelength is not an entirely trivial problem. In the following subsection we present the scaling of output power in the fundamental for different parameter combinations. Subsequently we present simulations of the full output characteristics for the chosen working points.

3.1.1 Scaling of output with undulator parameters and electron energy

We carry out parameter scans for target (fundamental) emission wavelengths in the x-ray part of the spectrum, at 5 nm (248 eV) and 0.5 nm (2.48 keV). The free parameters are the undulator period, the undulator strength and the electron energy. If any two of these are fixed, the third can be varied in order to ensure emission takes place at a particular wavelength (to within some limits).

In order to better assess the physical feasibility of a given set of parameters, instead of directly scanning the undulator strength, a_u, we scan the undulator gap, g. We then employ the relation $B_0 = 3.694[\mathrm{T}]\exp\left[-5.068g/\lambda_\mathrm{u} + 1.520(g/\lambda_\mathrm{u})^2\right]$ (see Sec. 2.3.4, in particular Eq. (2.56)) to find the peak on-axis magnetic field for a state-of-the art hybrid undulator, and obtain the corresponding value of a_u using equation (2.36).

For each of the two target emission wavelengths we perform a wider scan to choose a set of realistic undulator periods, λ_u. For each of the chosen periods we perform scans over ranges of electron energies, $E_\mathrm{e} = \gamma mc^2$, and undulator gaps, g, that lead to emission of the desired wavelength. We then compute the output energy at the fundamental within the central emission cone and natural bandwidth (see Eq. 2.43) for each set of parameters. Throughout, we choose a bunch charge of 100 pC. The contributions of further factors such as bunch energy spread and divergence will be considered in the next subsection.

A source of 5 nm radiation

Figure 3.1a shows the scaling of emitted energy at 5 nm for undulators with $\lambda_\mathrm{u} = 8\,\mathrm{mm}$, 10 mm or 12 mm, electron energies in the energy range 500–1000 MeV and an undulator gap between 1 and 3 mm. The scalings show that higher electron energies coupled with longer undulator periods and smaller gaps (i.e., higher a_u) generally lead to higher emission levels, as a rough examination of equation (2.43) would suggest. However, the dependence we observe here is relatively weak for $\lambda_\mathrm{u} \geq 10\,\mathrm{mm}$ and $E_\mathrm{e} \geq 700\,\mathrm{MeV}$.

Examining the scalings, a good working point appears to be an undulator of 10 mm period and 2 mm gap — for which $a_\mathrm{u} = 0.936$ — with a driving electron bunch of energy 700 MeV ($\gamma = 1370$). We note that this value of a_u is close to the optimum in terms of emission scaling with a_u alone, see Fig. 2.7.

These parameters are a compromise in terms of achieving an emission level close to the maximum available within the range, while not excessively challenging the machine parameters in terms of electron energy or undulator gap. In particular, a 2 mm gap would provide some scope for fluctuations in electron pointing and divergence — as well as for energy fluctuations that could contribute to divergence due to the chromaticity of the electron focusing optics. If a variable gap undulator is used, Figure 3.1a also illustrates the ability to tune the fundamental emission to 5 nm for a range of electron energies[1].

A source of 0.5 nm radiation

Figure 3.1b shows the scaling of emitted energy at 0.5 nm for undulators with $\lambda_\mathrm{u} = 5\,\mathrm{mm}$, 6 mm, and 7 mm, electron energies in the energy range 1000–1800 MeV and an undulator gap between 1 and 3 mm. The trends are broadly in line with those observed in Figure 3.1a, however it is evident that in this case the output level is much more sensitive to the parameters. This is partly because the shorter undulator period — which is necessary for the emission of short-wavelength radiation using realistic electron energies — leads to a lower[2] a_u, and thus weaker emission (see Fig. 2.7).

In this case it is evident that closing the undulator gap is advantageous in terms of output energy. In order to allow reasonable propagation of an electron bunch from a laser-wakefield accelerator through the undulator it would, however, be unwise to go below $g = 1\,\mathrm{mm}$. Even though 1 mm is an aggressive gap size, in recent experiments an undulator with a gap of 1.2 mm was used to generate radiation from laser-accelerated electrons — and led to the detection of undulator radiation on over 70 % of shots [38]. Here, the 6 mm-period undulator appears to offer a significant advantage over the 5 mm-period one in terms of output level, however the case for going to a 7 mm-period undulator is not as clear. We will thus select $\lambda_\mathrm{u} = 6\,\mathrm{mm}$, $g = 1\,\mathrm{mm}$ — therefore $a_\mathrm{u} = 0.661$ — and $E_\mathrm{e} = 1500\,\mathrm{MeV}$.

3.1.2 Simulations

Table 3.1 provides a summary of the working points chosen in the previous subsection. Based on the discussion in Sec. 2.1.2, we complement them with further projected electron bunch parameters. In both cases the total length of the undulator is chosen to be 0.6 m, in line with this being a “compact” light source. In this section we do not consider interference effects, therefore the bunch length has no effect on the total output energy, only on the radiation pulse length.

In both cases the electron bunch is focused in the center of the undulator. We have selected focusing strengths that maintain transverse sizes $\lesssim 40\,\mu\mathrm{m}$ — i.e., significantly smaller than the undulator gap — while ensuring the divergence does not contribute excessively to the spectral and angular spread of the radiation. This is especially important in the case of 0.5 nm radiation, where the central wavelength is more sensitive to divergence effects. This is because divergence causes a broadening of the emission towards longer

---

[1] Conversely, it is of course also possible to tune the output wavelength by changing the electron energy and/or undulator gap (see Eq. (2.39)).

[2] The effect of reducing λ_u on a_u is twofold: it enters linearly in the basic equation for a_u, Eq. (2.36), but reducing λ_u also increases the ratio g/λ_u in Eq. (2.56), thereby generally reducing the peak field, B_0, for a given undulator gap.

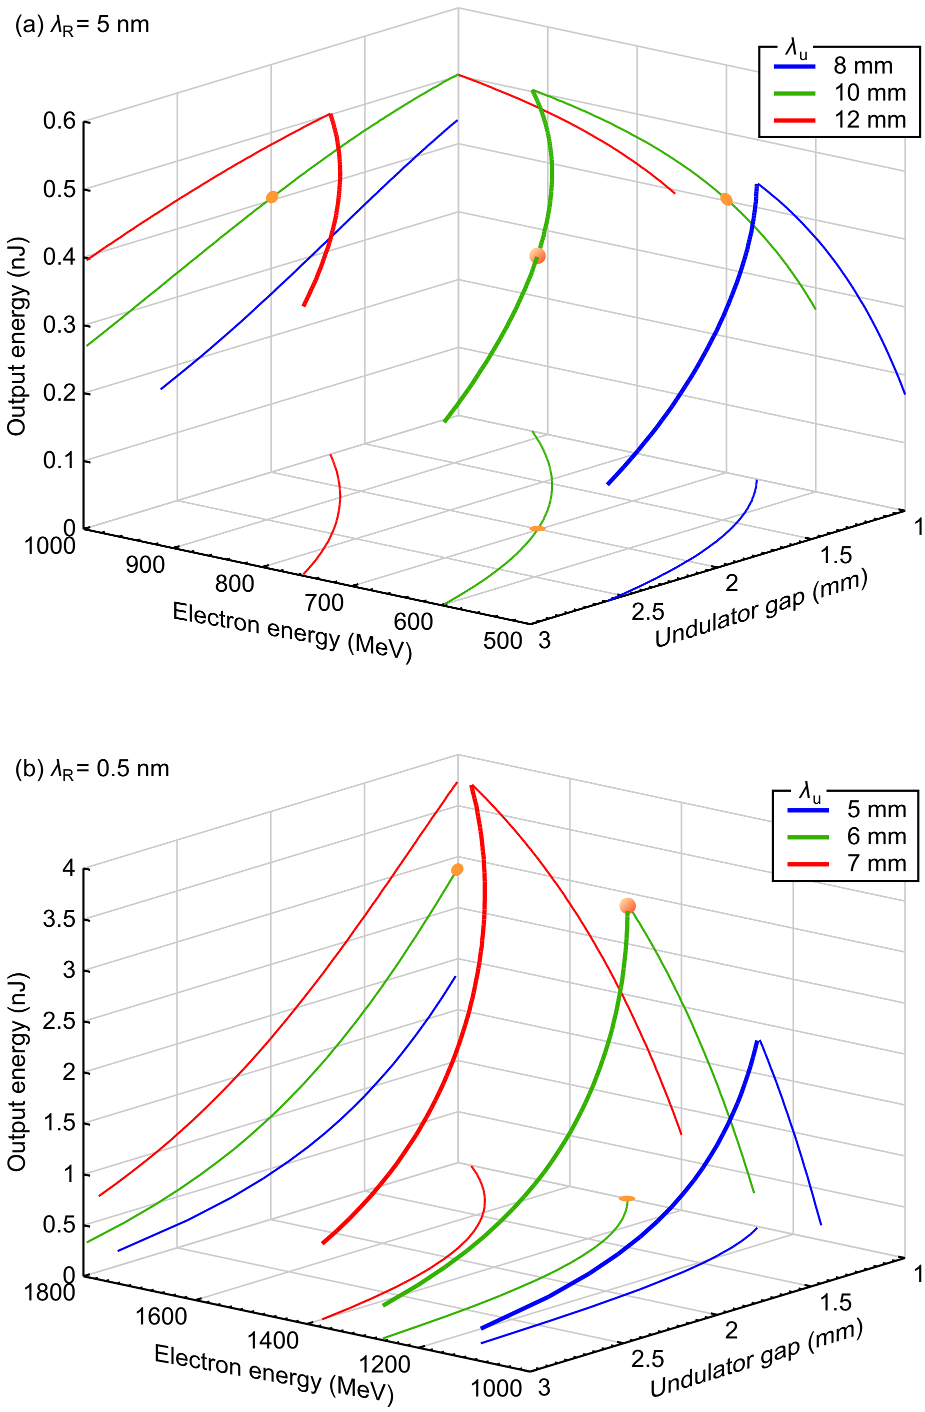

Figure 3.1: Scaling of energy emitted at the fundamental wavelength within the central emission cone and natural bandwidth versus undulator gap and electron energy, for three different undulator periods. The target fundamental wavelengths are (a) 5 nm and (b) 0.5 nm. Thick lines are representations in 3-D, while the thin lines are the corresponding projections onto each set of axes. The orange dots represents the chosen working points for both cases.

wavelengths that scales with γ^2 — this can be seen for on-axis emission in Figure 3.3.

We use the SPECTRA code [136] to complement the analytical scalings and to obtain more

Table 3.1: Parameters for incoherent light sources at 5 nm and 0.5 nm.

Radiation				
Resonant wavelength	λ_R	5	0.5	nm
Electrons				
Energy	E_e	700	1500	MeV
Lorentz factor	γ	1370	2935	
Energy spread (rms)	σ_γ/γ	1	0.5	%
Bunch charge	Q	100	100	pC
Normalized emittance	ε_n	1.5	2	mm mrad
Geometric emittance	ε	1.09	0.68	µm mrad
Undulator				
Period	λ_u	10	6	mm
Gap	g	2	1	mm
Strength (rms)	a_u	0.936	0.661	
Number of periods	N_u	60	100	
Total length	L_u	0.6	0.6	m
Focusing				
β-function at waist	β_T	0.8	2.4	m
Waist size	$\sigma_{r,\mathrm{min}}$	29.6	40.4	µm
Size at und. entrance	$\sigma_{r,\mathrm{max}}$	31.6	40.7	µm

detailed performance indicators. The code calculates the radiation emission by integrating the Liénard-Wiechert potentials for a single (idealized) electron trajectory, and then applies the corresponding convolutions with transverse size and energy spread to obtain the actual emission characteristics. It has the capability to calculate emission in the near as well as the far field — the present simulations are all carried out in the far field.

Figure 3.2 shows the results of simulations carried out using SPECTRA for the 5 nm case. The spectral density plots clearly illustrate the broadening that occurs for radiation emitted off-axis, as characterized by e.g., Eq. (2.40). The angular dependence of the emission also exhibits notable characteristics: while emission at the fundamental, (b), is broader along the y-axis than along the x-axis (as expected, see e.g., Ref. [124]), emission over all wavelengths, (a), is broader along the x-axis — due to the fact that the higher harmonics have significant off-axis components along that direction. Emission within the natural bandwidth of the fundamental, (c), is as expected almost exclusively on-axis, with the ring at ~ 1 mrad arising due to shifting of the second harmonic wavelength at that emission angle (cf. Eq. (2.40)).

The characteristics of emission at 0.5 nm, shown in Fig. 3.3, follow a broadly similar pattern. It is notable that due to the lower a_u there is less emission at higher harmonics. This is evident both from the spectral flux density plot, as well as from the relative heights of the peaks in the angular energy density plots, (a) and (b).

The parameters proposed here, and the output characteristics projected by the simu-

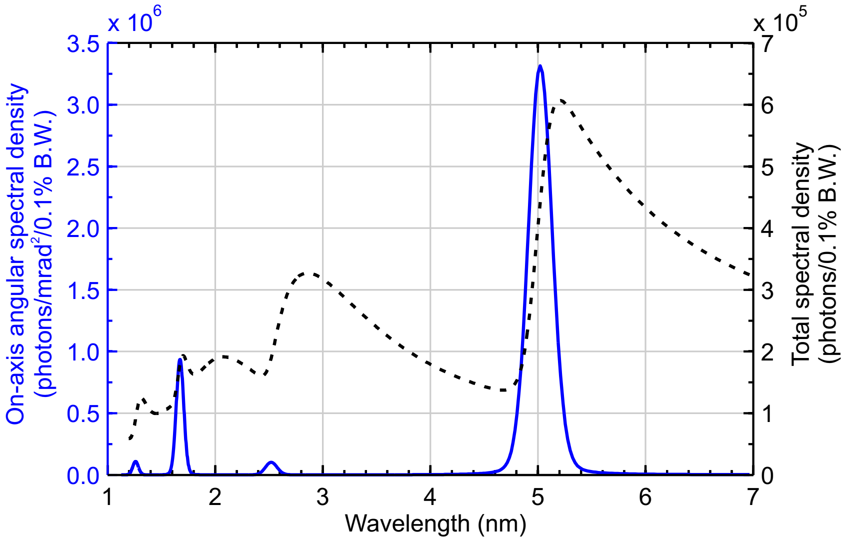

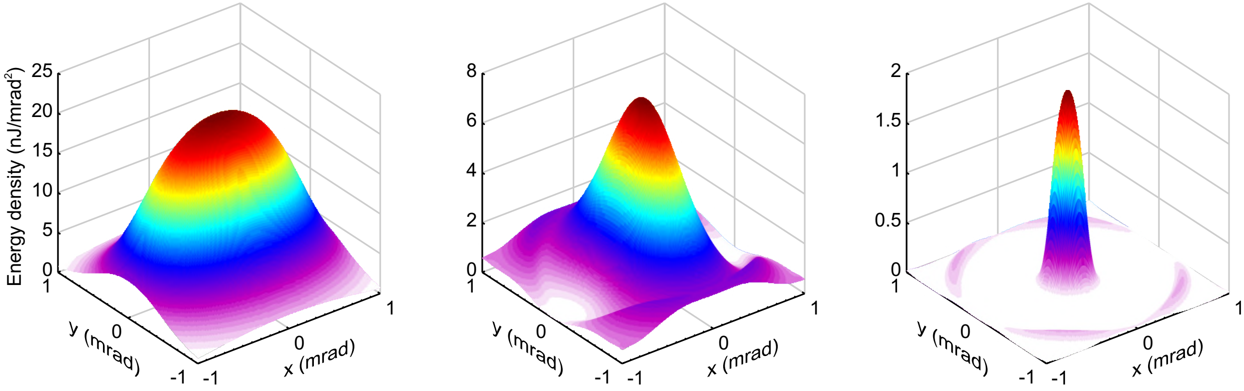

Figure 3.2: Results from simulation of emission at 5 nm for the parameters in Table 3.1. Top: On-axis (blue) and total (dotted black) spectral flux density. Bottom: angular energy density: (a) total over all wavelengths; (b) at wavelengths longer than 4.68 nm, i.e., mostly fundamental; (c) within natural bandwidth of fundamental emission, i.e., Gaussian filter of FWHM 0.074 nm centered on 5 nm.

lations, can be used as guidelines on the performance of future compact light sources of incoherent x-rays driven by electrons from a laser-wakefield accelerator. This chapter now turns to a similar — but more detailed — analysis of the feasibility and performance of future FELs driven by such electrons.

3.2 Production of coherent XUV radiation in a free-electron laser

While undulator radiation is relatively straightforward to produce and has many uses, it lacks coherence and its intensity scales only linearly with the charge in the driving bunch. This limits the amount of emission, especially in the case of bunches from laser-wakefield accelerators, where the bunch charge is of order 10–100 pC. Despite this relatively low

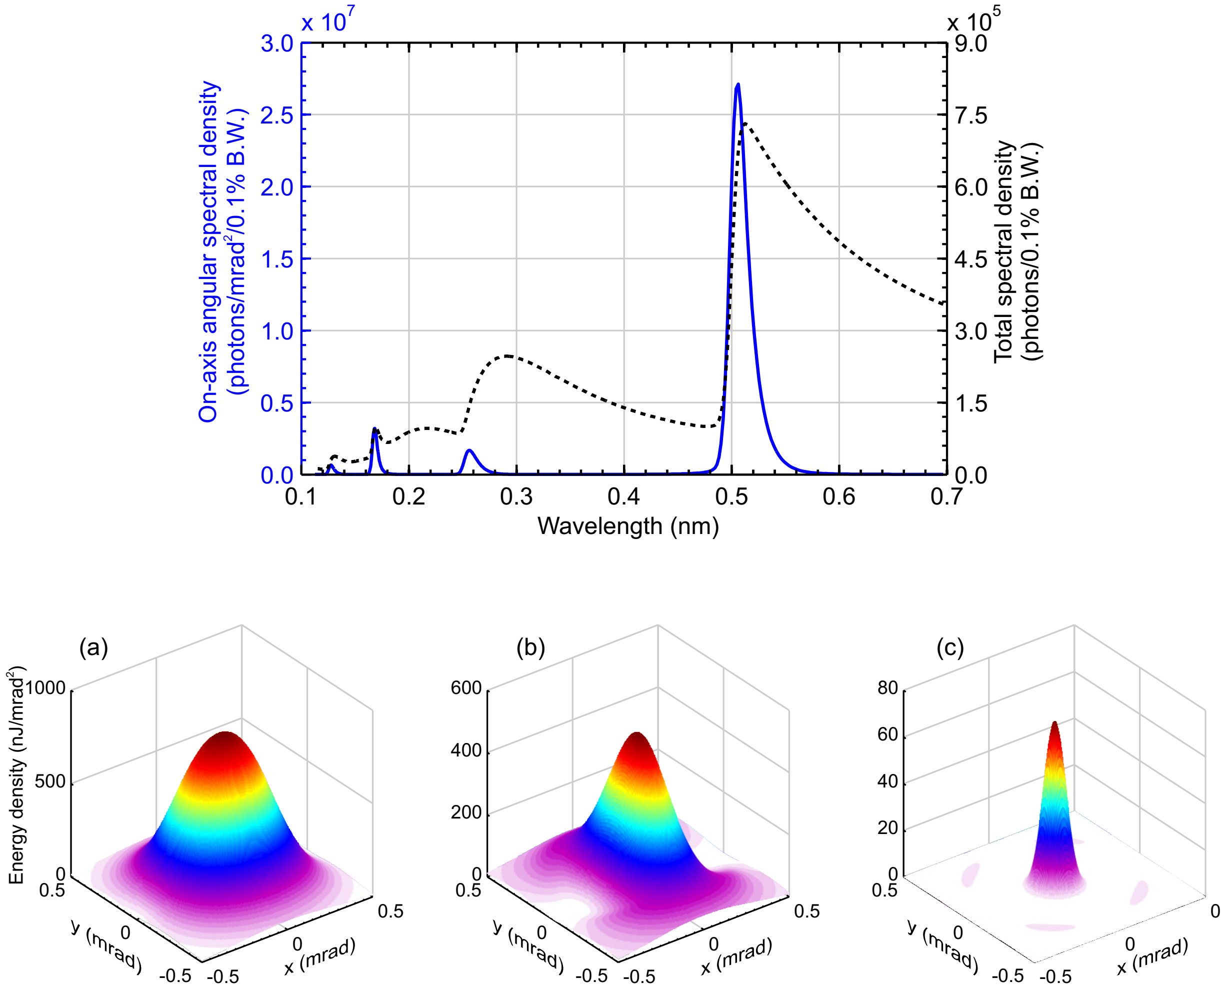

Figure 3.3: Results from simulation of emission at 0.5 nm for the parameters in Table 3.1. Top: On-axis (blue) and total (dotted black) spectral flux density. Bottom: angular energy density: (a) total over all wavelengths; (b) at wavelengths longer than 0.479 nm, i.e., mostly fundamental; (c) within natural bandwidth of fundamental emission, i.e., Gaussian filter of FWHM 0.0045 nm centered on 0.5 nm.

charge, however, laser-accelerated electron bunches can be very short — of order 1 µm rms — meaning their peak current is at the kiloampère level. This makes operation in the free-electron laser regime — where the output can scale with the square of the bunch charge — a tantalizing prospect.

Here we investigate the possibility of driving a free-electron laser in the XUV spectral range with electrons from a laser-wakefield accelerator; see Sec. 2.3.2 for background on FELs. We first discuss some of the scaling laws that can be applied to the problem. We then use them to select a viable working point, and finally carry out full simulations of the FEL emission process.

3.2.1 Scaling laws

As indicated in Section 2.3.2, a fully analytical characterization of FEL performance beyond the 1-D regime is prohibitively complex. Nevertheless, in Ref. [148] Ming Xie has carried out a variational analysis of the growth of the eigenmodes in an FEL in order to allow factors such as electron energy spread, emittance and radiation diffraction to be incorporated as a simple scaling of 1-D growth.

An additional effect not accounted for by these scalings is the reduction in gain due to the short bunch length. When the bunch length is on the order of a few cooperation lengths, L_c (c.f. Eq. (2.51)), radiation slips over the bunch at a rate that is too fast to allow full amplification to take place.

This subsection gives a summary of Xie's 3-D scaling laws, and outlines empirically derived scaling laws that allow for the characterization of the impact of short bunch length.

3-D scaling laws by Ming Xie

Xie introduces four parameters on the basis of which his analysis is carried out. These are as follows:

- The radiation diffraction contribution,

  $$\eta_\mathrm{d} = \frac{L_\mathrm{G}\lambda_\mathrm{R}}{4\pi\sigma_r^2},$$

  where the transverse radiation size is assumed to be matched to the electron bunch size, σ_r.

- The emittance (and thus electron divergence) contribution,

  $$\eta_\varepsilon = 4\pi\frac{L_\mathrm{G}}{\lambda_\mathrm{R}}\frac{\varepsilon_\mathrm{n}}{\gamma\beta_\mathrm{T}},$$

  where ε_n is the normalized transverse emittance and β_T is the Twiss β-function (see Sec. 2.2.1). Xie's model assumes continuous transverse focusing that maintains a constant β-function in the undulator.

- The electron energy spread contribution,

  $$\eta_\gamma = 4\pi\frac{L_\mathrm{G}}{\lambda_\mathrm{u}}\frac{\sigma_\gamma}{\gamma},$$

  where σ_γ/γ is the relative energy spread.

- The frequency detuning parameter,

  $$\eta_\omega = 4\pi\frac{L_\mathrm{G}}{\lambda_\mathrm{u}}\frac{\omega-\omega_\mathrm{R}}{\omega_\mathrm{R}}.$$

The 1-D limit described in section 2.3.2 corresponds to the case where all these parameters are zero, i.e., the gain is the maximum achievable for a given set of γ, a_u, λ_u, I_p, σ_r. For non-zero electron energy spread, emittance, or radiation diffraction, the gain length is greater with respect to that in the 1-D limit by

$$L_{\mathrm{G,3D}} = (1+\Lambda)L_\mathrm{G}, \tag{3.1}$$

where L_G is the 1-D gain length as defined in Eq. (2.49), and Λ is a function of η_d, η_ε and η_γ at the optimum detuning, η_ω^*. Xie gives a fit for Λ as [148]

$$\begin{aligned}\Lambda = & a_1\eta_\mathrm{d}^{a_2} + a_3\eta_\varepsilon^{a_4} + a_5\eta_\gamma^{a_6} & (3.2)\\ & + a_7\eta_\varepsilon^{a_8}\eta_\gamma^{a_9} + a_{10}\eta_\mathrm{d}^{a_{11}}\eta_\gamma^{a_{12}} + a_{13}\eta_\mathrm{d}^{a_{14}}\eta_\varepsilon^{a_{15}} & (3.3)\\ & + a_{16}\eta_\mathrm{d}^{a_{17}}\eta_\varepsilon^{a_{18}}\eta_\gamma^{a_{19}}. & (3.4)\end{aligned}$$

The 19 fitting parameters $\{a\}$ are given in Table 3.2. We will use this scaling in Section 3.2.2 in order to facilitate the choice of a realistic working point.

Table 3.2: Xie fitting parameters for gain length scaling.

$a_1 = 0.45$	$a_2 = 0.57$	$a_3 = 0.55$	$a_4 = 1.6$
$a_5 = 3$	$a_6 = 2$	$a_7 = 0.35$	$a_8 = 2.9$
$a_9 = 2.4$	$a_{10} = 51$	$a_{11} = 0.95$	$a_{12} = 3$
$a_{13} = 5.4$	$a_{14} = 0.7$	$a_{15} = 1.9$	$a_{16} = 1140$
$a_{17} = 2.2$	$a_{18} = 2.9$	$a_{19} = 3.2$	

Gain degradation due to short bunch length

All the characteristics of FEL gain described hitherto — in 1-D as well as in 3-D — have been derived under the assumption that the electron beam current is constant in time; this is the so-called “coasting beam”, or steady-state regime. It is broadly applicable to the majority of conventional short-wavelength FELs currently in operation or under construction, since the electron bunches driving them are much longer than the total slippage of the radiation over the electron bunch during the gain process[3] [13, 17, 19, 20]. This long bunch length allows radiation “born” at the tail of the bunch to reach saturation while it is still interacting with the bunch, i.e., before it has slipped over it.

However, in the case of electron bunches from laser-wakefield accelerators the scenario is likely to be different: for a large range of FEL parameters radiation originating at the tail of the bunch would slip over the entire bunch well before reaching saturation. This is also the case for certain short-bunch modes of operation at conventional facilities such as LCLS [159] and SPARX [160, 161]. The regime where the bunch length is short compared to the FEL cooperation length is known as the (weak[4]) “superradiant” regime of FEL operation

---

[3] This total slippage is the number of gain lengths required to reach saturation, normally ~ 20, times the cooperation length, L_c, which is the slippage over a single gain length.

[4] The term “weak superradiance” is often used to distinguish this from the “strong” superradiant regime, where the bunch length is much longer than the cooperation length, and spikes in the temporal profile of the output radiation occur specifically for radiation originating in the tail of the bunch.

[146, 162, 163], leading to emission characteristics that are different from the steady-state regime. In particular, the total emitted radiation energy scales as N_e^2, where N_e is the number of electrons in the bunch, rather than $\propto N_\mathrm{e}^{4/3}$, as is the case in the steady-state regime. Gain in the superradiant regime is slower, saturation is reached later, and the pulse length at saturation is shorter than in the steady-state regime.

In Ref. [162] Bonifacio, Maroli and Piovella present an analytical treatment of FEL gain in the short-bunch limit, based on a top-hat temporal bunch profile. They derive the exponentially growing component[5] of the peak power to be

$$P_\mathrm{peak} \propto \frac{\bar{L}_\mathrm{b}}{3\pi(\bar{z}-\bar{L}_\mathrm{b})^2} \exp\left[3\sqrt{3}\bar{L}_\mathrm{b}^{1/3}\left(\frac{\bar{z}-\bar{L}_\mathrm{b}}{2}\right)^{2/3}\right], \tag{3.5}$$

for $\bar{z} \geq \bar{L}_\mathrm{b}$, where $\bar{z} = (4\pi\rho/\lambda_\mathrm{u})z$ is the normalized distance along the undulator and $\bar{L}_\mathrm{b} = (4\pi\rho/\lambda_\mathrm{R})L_\mathrm{b} = L_\mathrm{b}/(\sqrt{3}L_\mathrm{c})$ is the normalized bunch length. For reference, in the steady-state case the power grows as $P \propto \exp\left[\sqrt{3}\bar{z}\right]/3$.

Ideally, we would like to obtain a scaling law that describes the effect of electron bunch length on gain length, and potentially on saturation length. Due to the fairly involved dependence on $\bar{z}$ and $\bar{L}_\mathrm{b}$ of the expression in Eq. (3.5), and the lack of characterization of saturation dynamics in the linear regime in which the expression is derived, it is not trivial to obtain such a scaling based on Ref. [162]. We face similar challenges in the context of other analytical treatments, such as the result for an arbitrary longitudinal bunch shape obtained by Dattoli *et al.* in Ref. [164].

Due to these considerations, here we present an empirical scaling law that has been obtained from a set of simulations of FEL gain driven by bunches with Gaussian temporal profiles. These simulations were carried out in a 1-D time-dependent environment with no transverse effects (e.g. diffraction or emittance) and no electron energy spread, using the code described in Appendix B. Although a range of values of $\rho = \{1,2,5\}\times 10^{-3}$ was used in the simulations, the final results for the measured gain and saturation lengths were only significantly dependent on the ratio σ_z/L_c, where σ_z is the rms bunch length and L_c is the cooperation length. In all simulations an initial microbunching level of 10^{-4} was introduced at the same phase along the bunch, in order to avoid ambiguities associated with start-up from shot noise.

The deviation in mean gain length and the saturation lengths for different values of σ_z/L_c are plotted in Figure 3.4. Since the rate of growth in peak power is dependent on z (as Eq. (3.5) suggests), the gain lengths are averaged over z between the onset of exponential gain and the point at which the gain drops below 75% of the median gain, which we take to be the saturation point. The parameter η_z plotted in Figure 3.4 is the fraction by which the gain length observed in the short-bunch regime is greater than the steady-state gain length, L_G, that is,

$$L_\mathrm{G,eff} = (1+\eta_z)L_\mathrm{G}.$$

---

[5] As opposed to the oscillating or exponentially decaying components, see e.g., Ref. [142].

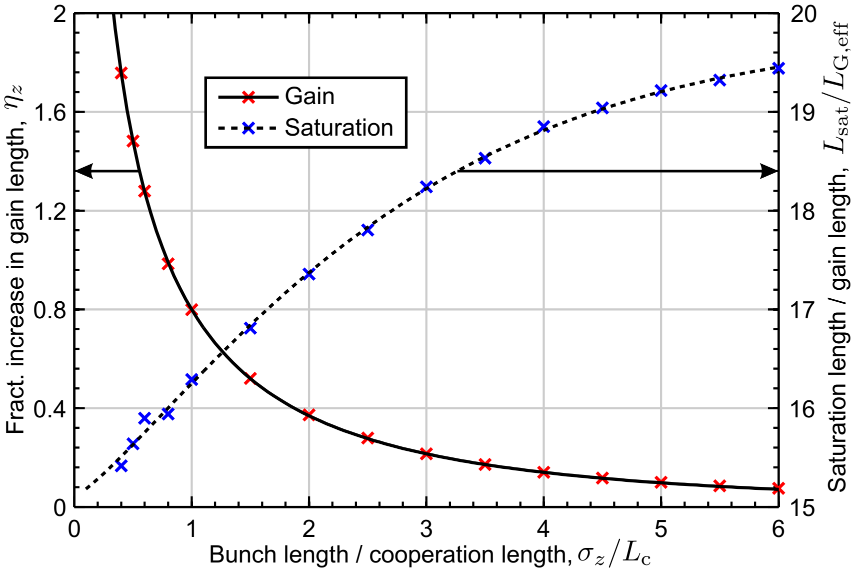

Figure 3.4: Empirical scaling of FEL gain length and saturation length versus rms length of driving Gaussian bunch relative to the FEL coherence length. Red (blue) crosses and solid (dotted) curve represent simulation results and fit for scaling of gain length (saturation length). The gain length is expressed as a fractional increase over the steady state gain length. The saturation length is expressed as number of effective gain lengths for that bunch length. The functions for the fit curves are given in the text.

The choice of functions for fitting was such that

$$\lim_{\sigma_z/L_{\rm c}\to\infty} \eta_z = 0$$

and

$$\lim_{\sigma_z/L_{\rm c}\to\infty} L_{\rm sat}/L_{\rm G,eff} = c_0$$

for some constant c_0. The best fits were obtained using the functional forms

$$\eta_z = \frac{L_{\rm G,eff}}{L_{\rm G}} - 1 = b_1 \exp\left[b_2\left(\frac{\sigma_z}{L_{\rm c}}\right)^{b_3}\right], \tag{3.6}$$

with fit parameters

$$\begin{aligned} b_1 &= 16.7512, \\ b_2 &= -3.0420, \\ b_3 &= 0.3267; \end{aligned}$$

and

$$\frac{L_{\rm sat}}{L_{\rm G,eff}} = c_0 + c_1 \exp\left[c_2\left(\frac{\sigma_z}{L_{\rm c}}\right)^{c_3}\right], \tag{3.7}$$

with fit parameters

$$\begin{aligned} c_0 &= 19.8202, \\ c_1 &= -4.7141, \\ c_2 &= -0.2778, \\ c_3 &= 1.2367. \end{aligned}$$

These are the fits shown in Figure 3.4.

Finally, we would like to combine the 3-D Xie scaling and the bunch length scaling in order to fully take into consideration the presence of radiation diffraction, bunch emittance, energy spread, and finite bunch length. This can be done by first finding the Xie scaling parameter Λ from Eq. (3.4), and the corresponding 3-D cooperation length $L_{\mathrm{c,3D}} = (\lambda_{\mathrm{R}}/\lambda_{\mathrm{u}})L_{\mathrm{G,3D}} = (1+\Lambda)L_{\mathrm{c}}$. On the basis of this we can find η_z for the bunch length scaling, and thus the effective overall gain length,

$$L_{\mathrm{G,eff}} = (1+\eta_z)(1+\Lambda)L_{\mathrm{G}}, \tag{3.8}$$

where η_z is a function of σ_z and $(1+\Lambda)L_{\mathrm{c}}$. This is based on the assumption that the 3-D steady-state scaling is not substantively altered in the case of a finite bunch length, and that the bunch length scaling is only dependent on the effective 3-D steady-state gain length, rather than the underlying effects. I.e., gain degradation due to finite bunch length is an effect that is secondary to the 3-D effects characterized by the Xie scaling.

Having found the effective gain length, the saturation length can be approximated using equation (3.7). It should be noted that while the σ_z scaling was derived for a bunch with constant-phase prebunching, start-up in the SASE regime is likely to proceed differently, with different longitudinal modes initially in competition for gain until a dominant one emerges. The signatures of this are likely to be delayed saturation, a fluctuation in saturation power, and possibly multiple saturation points (along z) at different peak powers [147].

Using the scaling laws established here, we proceed to select a working point for a proposed free-electron laser operating in the XUV.

3.2.2 Selection of working point

A graphical user interface was developed in order to carry out FEL parameter optimization. This is described in more detail in Appendix A. It can be used to find suitable electron bunch and undulator parameters for a given target emission wavelength. For the purposes of designing a proof-of-principle free-electron laser driven by electrons from a laser-wakefield accelerator, we select an operating wavelength of 32 nm. As we will later demonstrate, this should be a reasonable choice given currently available electron bunch parameters, see e.g., Sec. 2.1.2. Moreover, this wavelength is amenable for seeding with the 25th harmonic of an 800 nm driving laser. Other proposals for initial operation of FELs driven by laser-accelerated electrons have selected similar wavelengths — see e.g. Refs [39, 165, 166] — though our choice was made independently of those.

The optimal undulator parameters and electron bunch energy selected to minimize the gain length were $\lambda_{\mathrm{u}} = 8\,\mathrm{mm}$, $g = 1.2\,\mathrm{mm}$ — thus $a_{\mathrm{u}} = 0.878$ — and $E_{\mathrm{e}} = 240.4\,\mathrm{MeV}$ ($\gamma = 470.5$). These, as well as the further electron bunch parameters selected are given in Table 3.3. The table also gives a sample set of parameters for operation at a resonant wavelength of 2 nm, which we will not discuss in detail. It is, however, notable that the optimized undulator parameters for operation at 2 nm are the same as for 32 nm, and the only change is a quadrupling of the electron energy. The undulator parameter, $a_{\mathrm{u}} = 0.878$, is close to the optimal for incoherent emission (cf. Fig. 2.7). Acceptable operation in the 2 nm case also requires lower electron energy spread and transverse emittance, due to the smaller ρ and correspondingly slower gain and longer undulator. These criteria can be relaxed somewhat if a greater bunch charge or shorter bunch length were achieved, though ρ only scales as $\sim I_{\mathrm{p}}^{1/3}$. Back to the 32 nm case that is the topic of discussion of this section, we can confirm that the simple criteria for FEL operation set out in equations (2.52), namely

$$\begin{gathered}(L_{\mathrm{G,eff}} = 7.7\,\mathrm{cm}) < (Z_{\mathrm{R}} = 8.4\,\mathrm{cm});\\ (\varepsilon = 1.1\,\mathrm{nm}) < (\lambda_{\mathrm{R}}/4\pi = 2.6\,\mathrm{nm});\\ (\sigma_\gamma/\gamma = 5\times10^{-3}) < (\rho_{\mathrm{eff}} = 7.4\times10^{-3}),\end{gathered}$$

are met.

Figure 3.5 shows contour plots of the dependence of saturation length on many of the parameters listed above. Plots (a) and (b) are for the undulator parameters and electron energy — i.e., the parameters that need to be balanced to obtain the desired output radiation wavelengths — while plots (c) and (d) illustrate how other electron bunch parameters affect performance. Figure 3.5a shows that a narrower undulator gap would have decreased the saturation length, although such a gap is unlikely to be achievable in practice.

Clearly, in the context of recent experimental achievements (cf. Sec. 2.1.2), achieving an electron energy above the target 240.4 MeV would not be challenging, and at the same time Fig. 3.5b indicates that for a given undulator period a higher electron energy would improve performance. However, this would again require decreasing the undulator gap in order to maintain the same resonant emission wavelength. If we instead kept the undulator period and gap constant while increasing the electron energy, we would achieve emission at a shorter wavelength. While this may at first appear desirable, it means that the electrons would also need to be microbunched at a shorter wavelength, imposing more stringent criteria on their other characteristics.

Figure 3.6 complements Figs. 3.5c–d and shows how the gain length scales when electron bunch charge, bunch length, energy spread, and normalized emittance are varied around the selected working point. This allows a preliminary assessment of the sensitivity of gain on these essentially variable parameters.

3.2.3 Full simulations of FEL operation

In order to verify the viability of the parameters selected in the previous subsection, we carry out a number of full 3-D time-dependent FEL simulations using the GENESIS 1.3 code [167]. For this purpose the author created a set of computer scripts to automatically compose input files for the GENESIS simulations, run the simulations on a remote computer, and retrieve and synthesize the results from multiple simulations.

The bulk of the simulations are carried out in the SASE regime, with electron shot noise enabled but without an initial radiation seed. Since in the SASE regime the emission

Table 3.3: Parameters for free-electron lasers at 32 nm and 2 nm.

Radiation				
Resonant wavelength	λ_R	32	2	nm
Electrons				
Energy	E_e	240.4	961.6	MeV
Lorentz factor	γ	470.5	1882	
Energy spread (rms)	σ_γ/γ	0.5	0.2	%
Bunch charge	Q	50	50	pC
Bunch length	σ_z	0.5	0.5	µm
Peak current	I_p	12	12	kA
Normalized emittance	ε_n	0.5	0.2	mm mrad
Geometric emittance	ε	1.06	0.106	µm mrad
Undulator				
Period	λ_u	8	8	mm
Gap	g	1.2	1.2	mm
Peak on-axis magnetic field	B_0	1.66	1.66	T
Strength (rms)	a_u	0.878	0.878	
Number of periods	N_u	200	350	
Total length	L_u	1.6	2.8	m
Focusing				
Mean β-function in x	$\beta_{\mathrm{T},x}$	0.924	1.617	m
Mean β-function in y	$\beta_{\mathrm{T},y}$	0.682	1.617	m
Mean transverse size	σ_r	29.2	13.1	µm
FEL gain				
1-D gain length	L_G	3.73	8.75	cm
Xie 3-D scaling parameter	Λ	0.547	0.435	
3-D gain length	$L_\mathrm{G,3D}$	5.77	12.55	cm
Bunch length scaling parameter	η_z	0.334	0.009	
Effective gain length	$L_\mathrm{G,eff}$	7.70	12.67	cm
Cooperation length	$L_\mathrm{c,eff}$	308	31.7	nm
Effective Pierce parameter	ρ_eff	7.39	4.16	$\times 10^{-3}$
Saturation length	L_sat	1.35	2.51	m

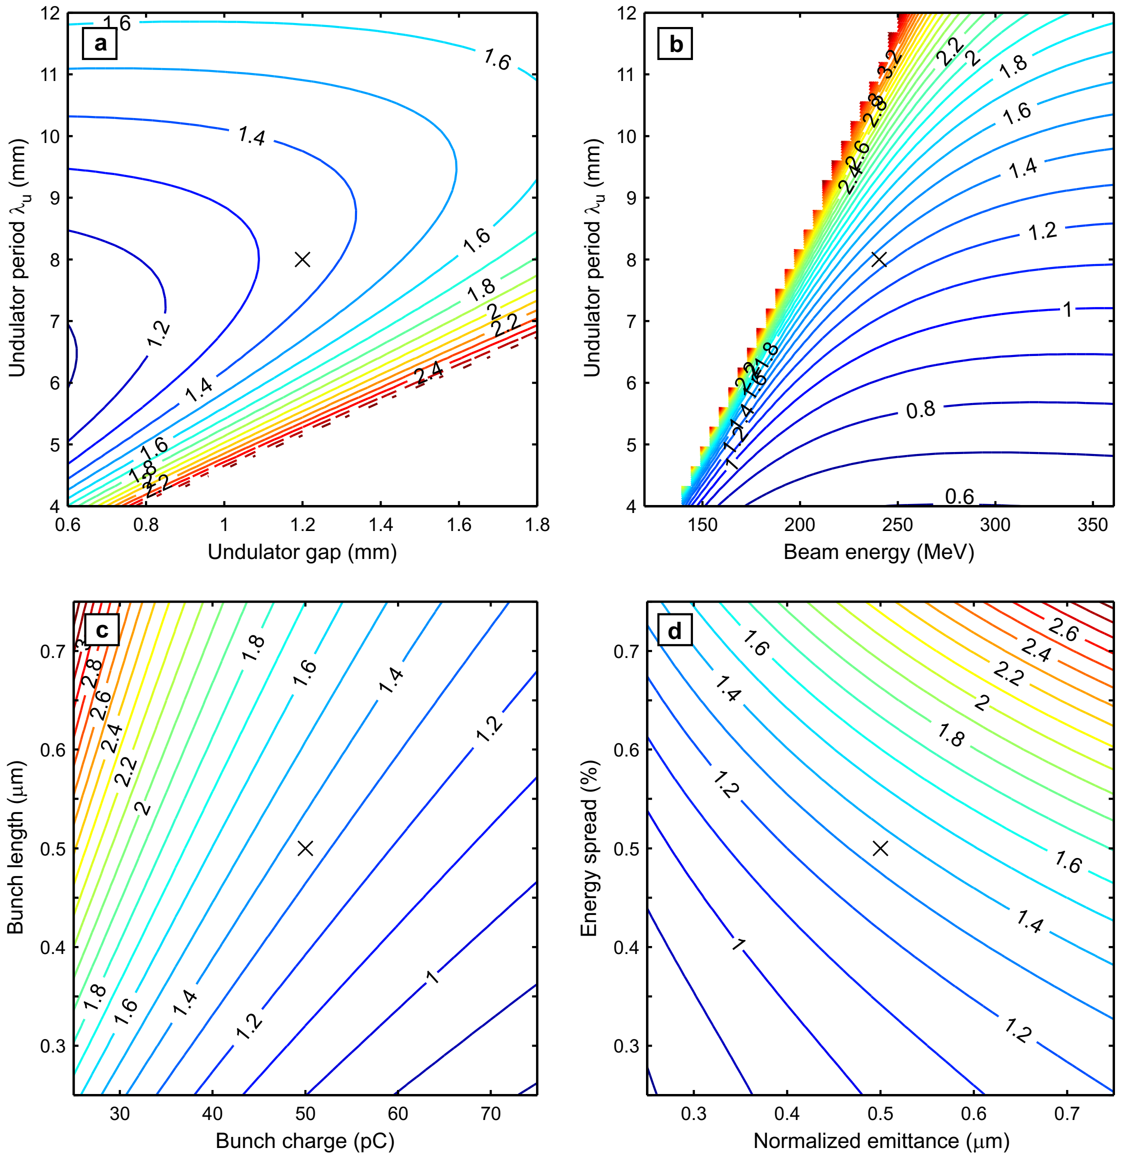

Figure 3.5: Contour plots of saturation length versus (a) undulator period and gap; (b) undulator period and electron energy; (c) bunch charge and bunch length; (d) energy spread and normalized emittance. In each case the '×' symbol marks the selected working point, and each plot shows variation in L_{sat} around that point for changes in the corresponding parameters.

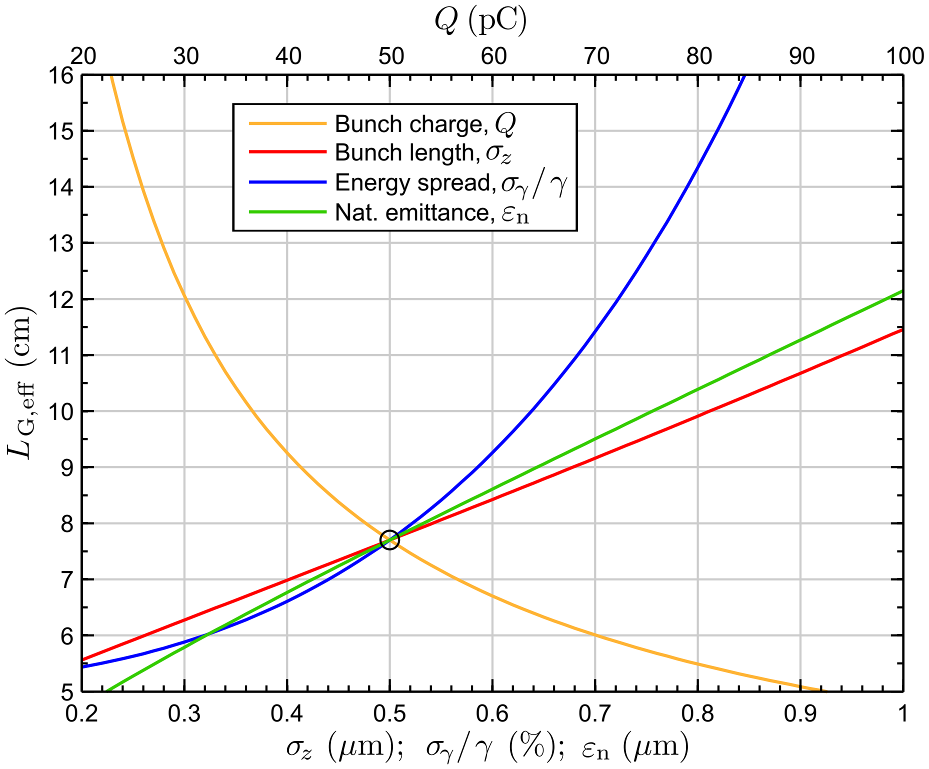

Figure 3.6: Scaling of FEL gain length for variations in bunch charge, bunch length, energy spread, and normalized emittance around the working point (marked ‘∘’).

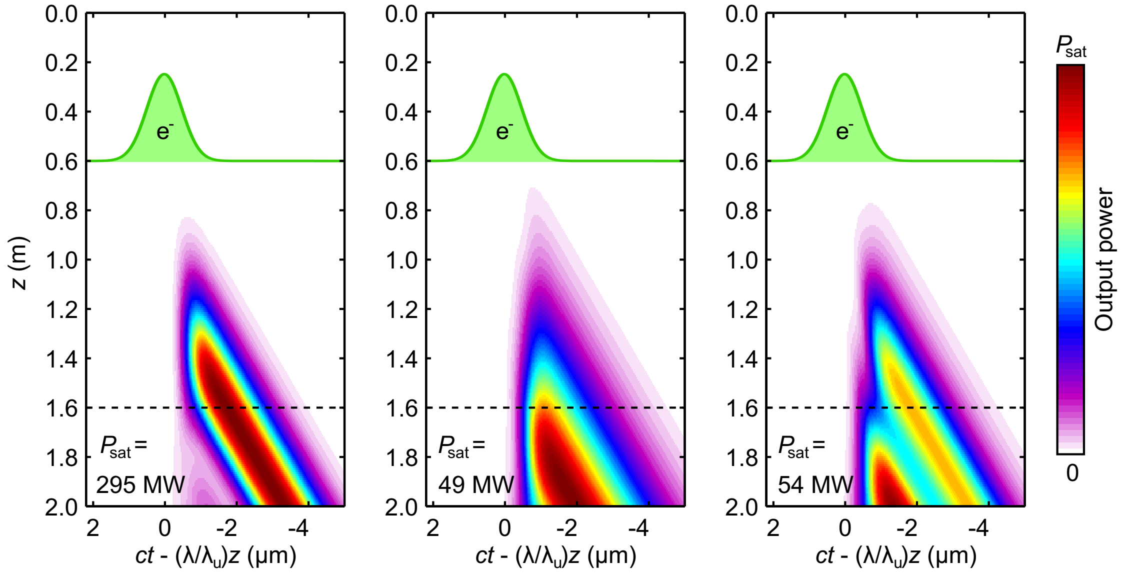

Figure 3.7: Samples of FEL power evolution with z, against ct. The color scale is linear, therefore the emission only appears close to saturation. The bunch current profiles — whose position within the ponderomotive window, $ct-(\lambda_\mathrm{R}/\lambda_\mathrm{u})z$, remains constant during evolution in z — are shows for reference.

process is dependent on the initial phases of the shot noise microbunching — which are random for each bunch — the gain length and saturation power differ from shot to shot. This is documented in e.g., Refs. [147, 168, 169]. In particular, Bonifactio *et al.* point out that shot-to-shot fluctuations in SASE emission increase as the bunch length decreases relative to the coherence length. In order to therefore more completely capture the predicted FEL performance, for each set of conditions we carry out 10 simulations, where the random shot noise is initialized differently in each simulation.

The electron bunch is focused in x to a waist half-way along the undulator, at $z = 0.8\,\mathrm{m}$, and natural undulator focusing is used to maintain a constant size in y. The simulation is actually run over an undulator distance somewhat greater than that specified in Table 3.3, in order to obtain a clearer picture of the saturation dynamics. Wakefields and magnetic field errors are not considered, but short-range longitudinal space-charge effects are.

Emission characteristics

Figure 3.7 shows the simulated evolution of emitted power versus undulator location, z, and location along the ponderomotive coordinate, $\theta/k_\mathrm{R} = ct-(\lambda_\mathrm{R}/\lambda_\mathrm{u})z$, for three different shots at the selected working point parameters. Additionally, Figure 3.8 shows the output profiles and spectra obtained at $z = 1.6\,\mathrm{m}$ for all 10 cases. These plots serve to illustrate the shot-to-shot fluctuations in emission in the SASE regime for such short electron bunches.

Figures of merit for the simulated radiation characteristics near saturation (around $z = 1.5$–$1.6\,\mathrm{m}$) are given in Table 3.4. The bandwidth, pulse length, source size and divergence quoted in the table are root-mean-square quantities. It is clear that the majority of performance metrics exhibit low shot-to-shot fluctuations, with the exception of the actual

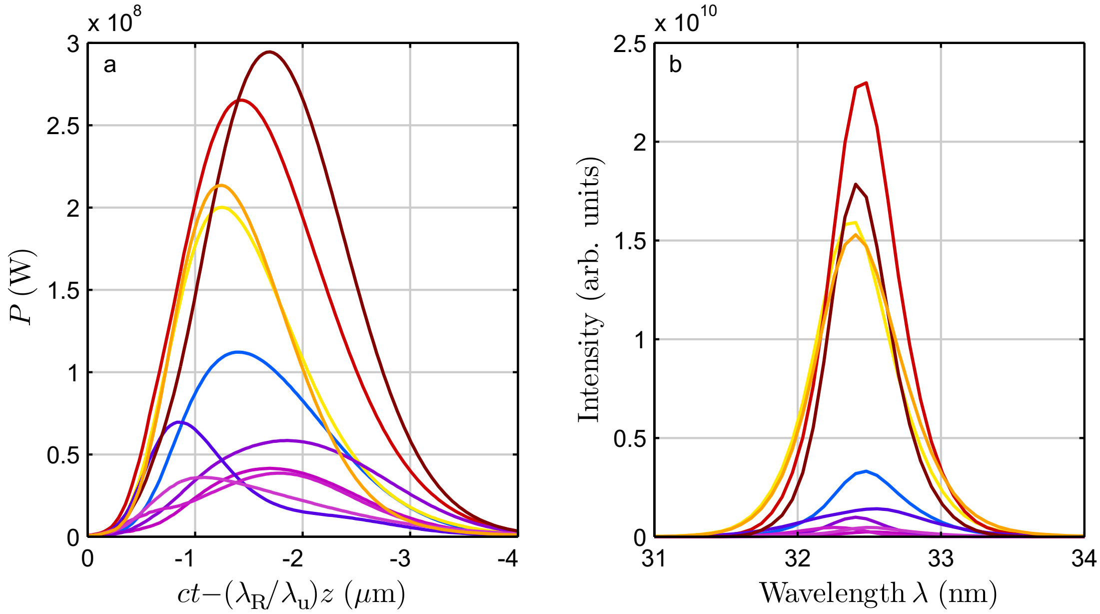

Figure 3.8: (a) Simulated temporal emission profiles for 10 shots at the parameters in Table 3.3 at $z = 1.6\,\mathrm{m}$, i.e., at or near saturation. (b) Corresponding spectra.

emission level (number of photons), which varies by approximately an order of magnitude for different initial shot noise. The range of peak brilliance is given based on the brilliances calculated independently for each shot, rather than superficially from the metrics given in the table.

Table 3.4: Simulated emission characteristics at FEL saturation. Ranges observed over 10 SASE simulations with randomized initial noise.

Number of photons	N_{ph}	3 – 20	$\times 10^{10}$
Peak wavelength	λ	32.3 – 32.5	nm
Bandwidth	σ_λ/λ	3 – 5	%
Pulse length (rms)	σ_z	0.42 – 0.48	µm
		↪1.4 – 1.6	fs
Source size	σ_r	50 – 60	µm
Divergence at source	$\sigma_{r'}$	0.55 – 0.70	mrad
Peak brilliance	B	$\sim 10^{26} - 10^{27}$	photons /second /mm^2 /mrad^2 /0.1% B.W.

Sensitivity to electron bunch characteristics

Following our discussion of SASE fluctuations, Fig. 3.9a illustrates the evolution of peak power along z for the selected working point parameters, in terms of the mean and rms levels from the 10 simulations. It also shows how this evolution differs for different bunch charges.

Furthermore, Figures 3.9b, 3.9c and 3.9d show the evolution of peak power for different values of bunch length, energy spread and normalized emittance. Note in particular the high sensitivity of gain to bunch charge and energy spread, and the relatively low sensitivity to bunch length. These are all in line with the predictions of the analytical and empirical scalings presented in Fig. 3.6. In particular, the bunch length is involved in a trade-off where although shorter bunches have higher peak currents, they lead to partial gain degradation as described in Sec. 3.2.1.

Another striking feature is the surge in output by a factor of ~ 60 when the transverse emittance is decreased by 50% — this exceeds the change expected on the basis of the gain length scaling in Fig. 3.6, and suggests that emittance effects may play a greater role than the Xie scaling accounts for.

These sets of simulations aim to elucidate some of the different ways in which FEL performance may be affected by fluctuations in bunch characteristics. Even in state of the art laser-wakefield acceleration systems, such fluctuations are unfortunately all too common. Parameters such as energy spread and charge will clearly have to be kept to the specified levels, as Fig. 3.9 illustrates. This emphasizes the need for accurate diagnostics.

Seeded operation

Finally, we simulate the amplification of radiation seeds of relatively low mean powers, 10 kW and 50 kW. From the results in Figure 3.10 it is clear that these seeds not only have the potential to significantly enhance gain, but they dramatically decrease the shot-to-shot fluctuations in output that are inherent to start-up from shot noise.

Seeding levels of 10 or 50 kW at 32 nm would be readily achievable with present high-harmonic generation technology — see e.g., Refs. [150, 170] where μJ-level harmonic energies are reported within pulse durations of order 50 fs. As described previously, the temporal profile of HHG emission is not smooth but instead comprises a series of sub-femtosecond spikes. We show in Chapter 6 that in the general case this temporal structure does not impact the actual gain process. In Section 6.5 we discuss a possible interplay between the temporal (and spectral) structure of the seed radiation and that of the electron bunch.

It should be noted that while temporal and spatial synchronization of the HHG pulse and the electron bunch — both of which have fs-order lengths — may be challenging, it is facilitated by the fact that both the seed and bunch can be generated using intrinsically synchronized laser pulses originating from a single laser system.

In terms of peak brilliance in the seeded case, the 10 kW seed leads to $7.7 \pm 0.3 \times 10^{27}$ photons /second /mm^2 /mrad^2 /0.1% B.W., where the error given is standard deviation. This is almost an order of magnitude greater than the unseeded case, and as Fig. 3.10 suggests the most significant reason for this is the increase in output level itself. The only other significant change in the output characteristics with respect to those listed in Table 3.4 is an approximate halving of the bandwidth of emission, though this is partly attributable to the fact the idealized seed used is monochromatic. For the 50 kW seed the peak brilliance is yet higher, at $1.8 \pm 0.2 \times 10^{28}$ photons /second /mm^2 /mrad^2 /0.1% B.W.

We conclude this chapter by presenting the brilliance figures obtained here, as well as those for the unseeded case, alongside those for existing conventional light sources in Figure 3.11. Clearly, while the photon energies discussed here are not as high as those achievable at conventional free-electron lasers (in particular FLASH, which is closest in performance), the projected brilliance is in a comparable range. The operation of a proof-of-principle XUV FEL driven by laser-accelerated electrons in the near future seems increasingly likely. This would be a major milestone along the way to building a compact and affordable LWFA-based light source for real-world applications.

(a) Charge dependence

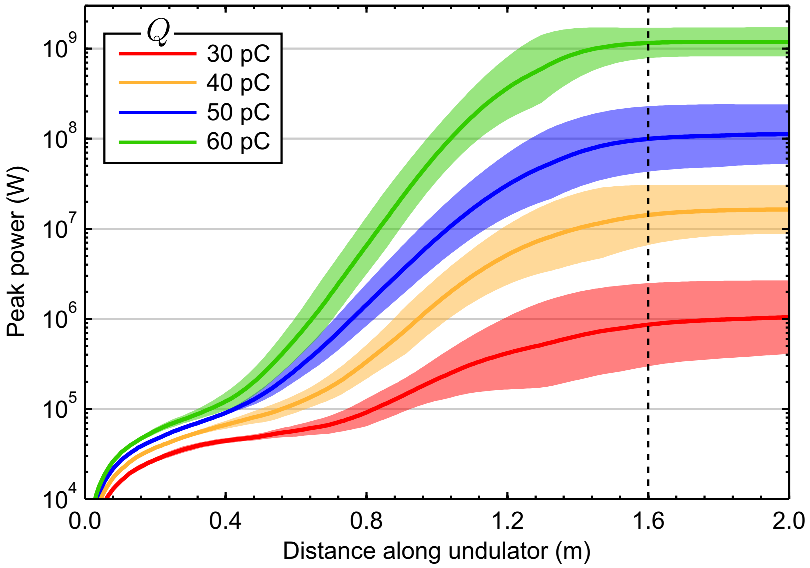

(b) Energy spread dependence

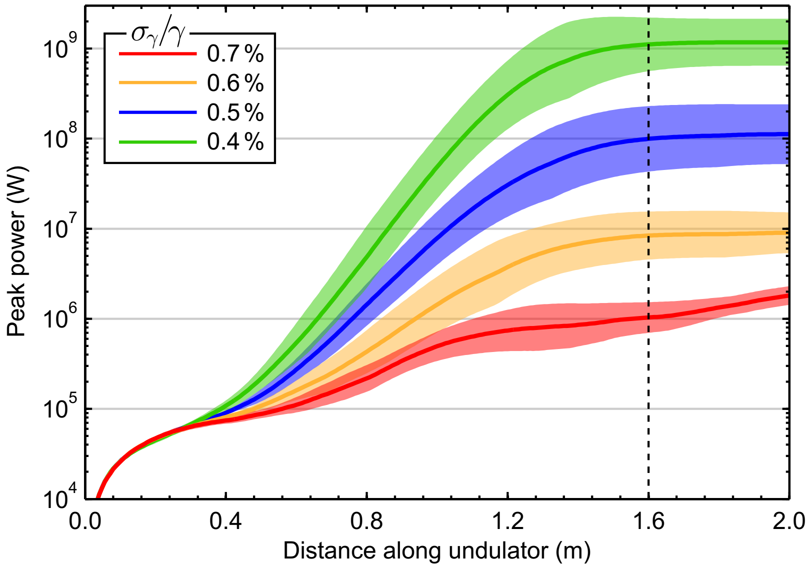

Figure 3.9: Evolution of FEL peak power for different bunch parameters. In each case one parameter is varied: (a) bunch charge, Q; (b) energy spread, σ_γ/γ; (c) bunch length, σ_z; (d) normalized emittance, ε_n. All parameters other than the one that is varied are according to the selected working point, Table 3.3. Solid lines represent the mean level observed over 10 simulations starting from randomized shot noise; shaded areas represent the rms variation.

(c) Bunch length dependence

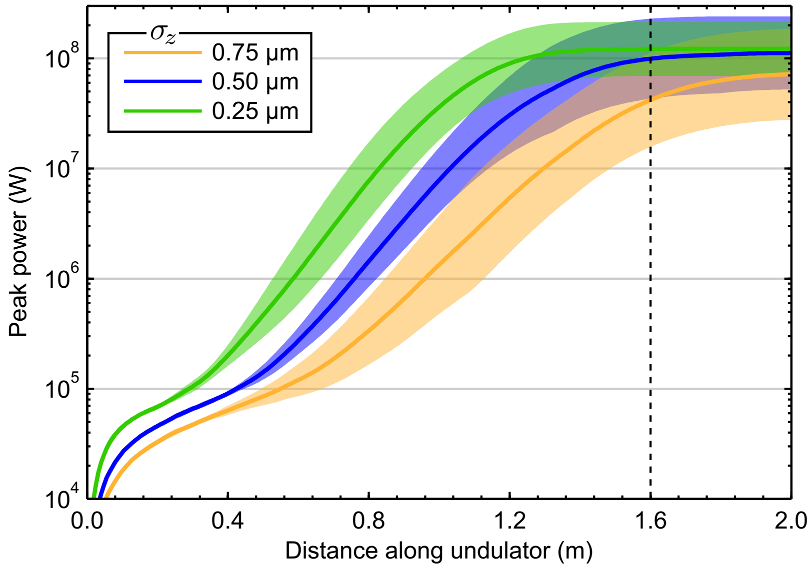

(d) Emittance dependence

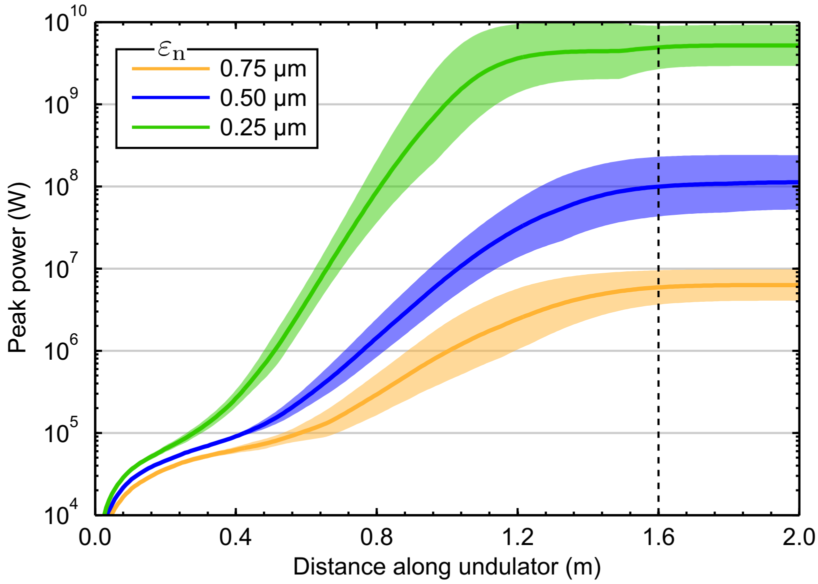

Figure 3.9: Evolution of FEL peak power for different bunch parameters. In each case one parameter is varied: (a) bunch charge, Q; (b) energy spread, σ_γ/γ; (c) bunch length, σ_z; (d) normalized emittance, ε_n. All parameters other than the one that is varied are according to the selected working point, Table 3.3. Solid lines represent the mean level observed over 10 simulations starting from randomized shot noise; shaded areas represent the rms variation.

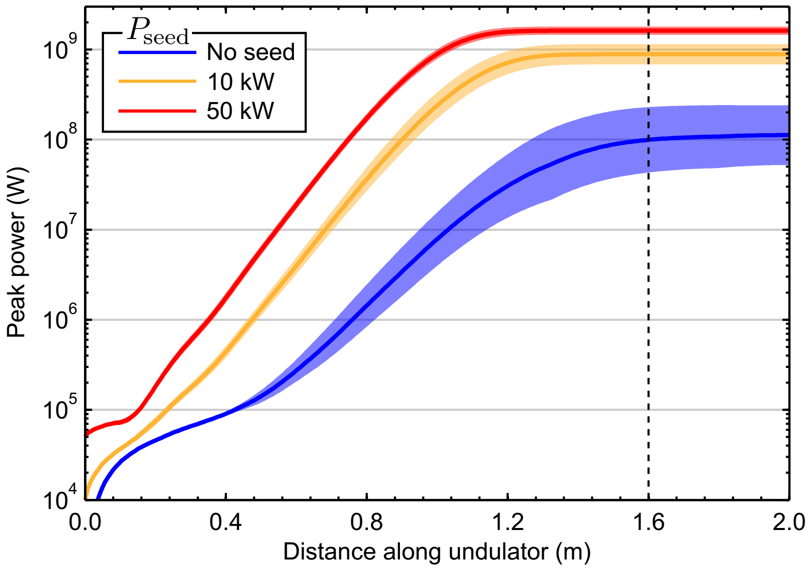

Figure 3.10: Evolution of FEL peak power for SASE and seeded operation, where the seeds have constant powers of 10 kW and 50 kW. Shot noise is included in all cases, leading to some observable shot-to-shot variation. All parameters are according to the selected working point, Table 3.3. Solid lines represent the mean level observed over 10 simulations starting from randomized shot noise; shaded areas represent the rms variation.

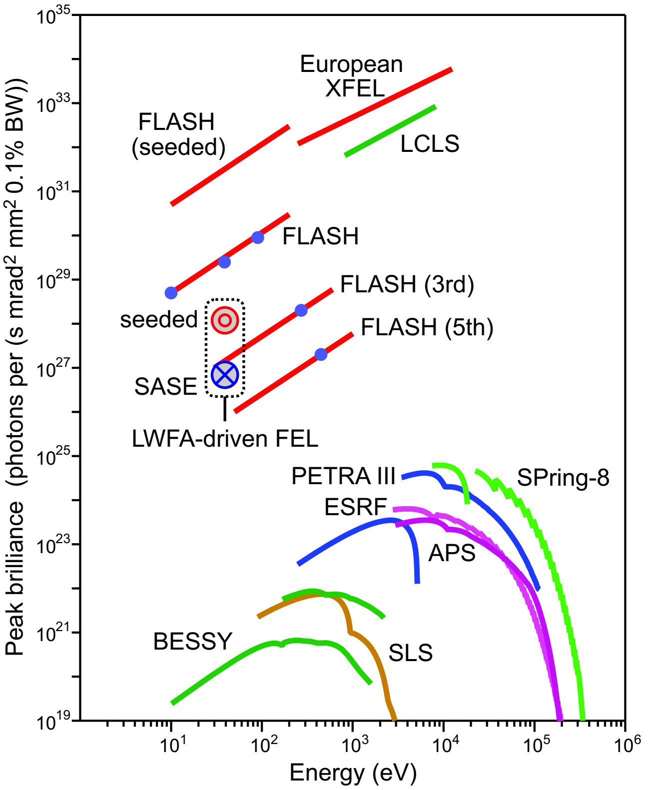

Figure 3.11: Plot of peak brilliance versus photon energy for various conventional synchrotron and free-electron laser light sources currently in operation or under construction (adapted from [20]). The projected performance of an FEL driven by laser-accelerated electrons at the working point described in Sec. 3.2 is also included, for SASE as well as seeded (50 kW seed) operation.

Chapter 4

Preliminary Diagnostics Based on Coherent Transition Radiation

Having established a suitable parameter space for operating a free-electron laser, it is important to determine whether the electron bunches produced by a laser-wakefield accelerator are of sufficiently high quality. The bunch charge, energy, and energy spread can be measured to a high degree of accuracy, however there have been few reliable measurements of transverse and longitudinal emittance[1]. On a simplistic level, to find the transverse emittance one needs to simultaneously measure the transverse bunch size and angular divergence; while the longitudinal emittance is a product of the bunch length and energy spread.

In this chapter we present initial experiments towards reliably and repeatably carrying out a selection of these measurements. These experiments are based on the detection of transition radiation (TR) that is emitted by the electron bunch as it passes through a screen, and we therefore begin with a discussion of the theory of transition radiation. The experiments presented here were carried out at the Max-Planck-Institut für Quantenoptik in Garching, Germany, where coherent transition radiation from laser-accelerated electron bunches was observed at optical frequencies. In the next chapter we discuss a second experimental campaign where the electron bunch and acceleration process were more fully characterized through a broadband CTR measurement.

4.1 Theory of transition radiation

Transition radiation is emitted whenever a charged particle traverses a boundary between two regions with different dielectric constants. It was first postulated by Ginzburg and Frank in 1945 [171], when they — following Cherenkov [172] — dispensed with the notion that a charged particle does not radiate unless it is accelerating. Indeed, in a later work Ginzburg and Tsytovich give a generalized definition of transition radiation as "radiation emitted in the case of uniform and rectilinear motion of a charge [. . .] under inhomogeneous

---

[1] For an introduction to emittance see Section 2.2.1.

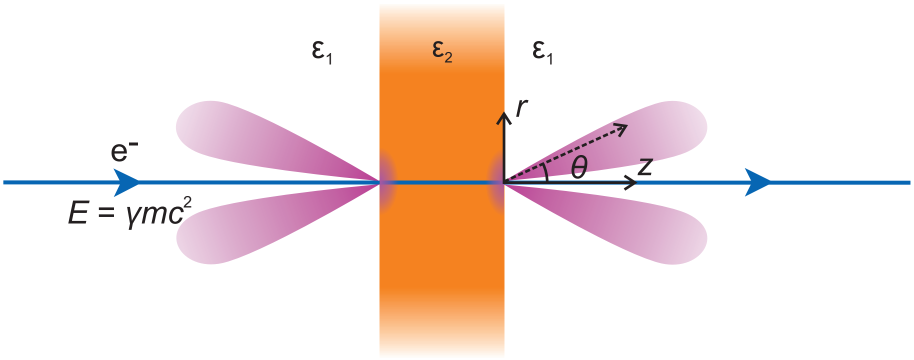

Figure 4.1: Basic setup for emission of transition radiation from relativistic electrons. In the text, the material (orange) is assumed to be a perfect conductor ($\varepsilon_2 = \varepsilon_0 - i\infty$), while the background medium is vacuum ($\varepsilon_1 = \varepsilon_0$). As well as forward transition radiation, backward TR is emitted for reflective materials.

conditions" [173].

Conceptually, transition radiation is emitted due to the *self-field* of the particle[2] being different for media of different dielectric constants: therefore when transiting between media, the change in this self-field leads to emission of radiation. Mathematically, the process can be approached in one of several ways. Ginzburg and Tsytovich offer a simple treatment in terms of the radiation emitted at the interface between vacuum and a perfect conductor, upon the effective disappearance of the charge (and its corresponding image) at the moment when it enters the conductor [173, p. 14]. When the medium is not a perfect conductor, emission can still result from a qualitatively similar process, with characteristics corresponding to the particular alteration in the self-field caused by the medium's dielectric constant.

Rather than working in terms of the self-fields, a more versatile model can be developed by following Jackson's approach [174] and considering the polarization induced in a material's surface by a transiting charge. We adopt this method in the following subsection, in order to evaluate the angular, spatial and spectral characteristics of transition radiation from ultra-relativistic electron bunches. Subsequently, in Section 4.1.2 we consider how these characteristics differ in the case where a number of electrons emit coherently at the observed wavelengths.

4.1.1 Incoherent transition radiation

An electron crossing the surface of a perfect conductor (Figure 4.1) leads to the following time-dependent radial and longitudinal fields at the surface [174, p. 650]:

---

[2] The self-field is the solution to Maxwell's equations in the vicinity of the particle.

$$\begin{aligned} \mathbf{E}_r(\mathbf{r},k) &= -\frac{e}{4\pi^2\varepsilon_0\beta^2 c}\frac{k}{\gamma}K_1\left(\frac{kr}{\beta\gamma}\right)\frac{\mathbf{r}}{r}, & (4.1)\\ E_z(\mathbf{r},k) &= i\frac{e}{4\pi^2\varepsilon_0\beta^2 c}\frac{k}{\gamma^2}K_0\left(\frac{kr}{\beta\gamma}\right), & (4.2) \end{aligned}$$

where βc and γ are respectively the velocity and Lorentz factor of the electron. Here $\mathbf{r}$ is the vector in the transverse plane and z is the longitudinal coordinate (see Figure 4.1), while $k = \omega/c$ is the angular wavenumber corresponding to the frequency-space representation of the field's temporal variation, and $K_{0,1}$ are modified Bessel functions of the second kind. Since the strength of the longitudinal field is a factor γ less than the radial field, it can be neglected for ultra-relativistic electrons.

By carrying out the far-field diffraction integral on $\mathbf{E}_r$ (see e.g. Casalbuoni [175]), we can obtain the Ginzburg-Frank formula for the spectral energy as a function of emission angle θ for a single electron,

$$\frac{\mathrm{d}^2 U_{\mathrm{GF}}}{\mathrm{d}\omega\mathrm{d}\Omega} = \frac{e^2}{4\pi^3\varepsilon_0 c}\frac{\beta^2\sin^2\theta}{(1-\beta^2\cos^2\theta)^2}, \tag{4.3}$$

where Ω is the solid angle subtended at the observation screen. Investigation of Eq. (4.3) shows that the angular distribution of the emission peaks at $\theta = 1/\gamma$, with no emission on-axis (as crudely illustrated in Figure 4.1). Furthermore, the intensity of emission is independent of frequency, meaning TR is extremely broadband. This latter consideration may appear to imply that the total emitted energy is infinite, which would, of course, not be physical: in reality, TR has a high-frequency cut-off around $\omega = \gamma\omega_{\mathrm{p}}$, the plasma frequency of the screen material multiplied by the Lorentz factor of the electron; additionally, equation (4.1) implies an effective source size on the screen of $r_{\mathrm{eff}} \approx \gamma\lambda$, where λ is the emission wavelength — therefore a finite screen size, a, introduces an additional term on the r.h.s. of Eq. (4.3) that suppresses emission at frequencies $\lesssim \omega = 2\pi c\gamma/a$.

In the current treatment, we focus on transition radiation in the wavelength range $\lambda = 0.3-10\,\mu\mathrm{m}$ from electron bunches with $\gamma \approx 600$ crossing a steel screen of diameter $a \approx 10\,\mathrm{mm}$. Under these conditions, equation (4.3) is a good representation of the spectral response of incoherent radiation from single electrons when imaged in the far field. However, in the experiments described in Sections 4.3 and 4.4, the radiating screen is imaged in the near field rather than the far field. A different treatment of the radiation propagation is thus required.

For this purpose it is instructive to take the Fourier transform of the radiating field, equation (4.1), in the transverse plane:

$$\tilde{\mathbf{E}}(\boldsymbol{\kappa},k) = \frac{1}{2\pi}\frac{-ie}{4\pi^2\varepsilon_0\beta c}\frac{\boldsymbol{\kappa}}{\kappa^2+(k/\beta\gamma)^2}. \tag{4.4}$$

We proceed with the approach taken by Loos *et al.* [176] (or see Castellano & Verzilov [177] for a more detailed treatment), integrating over a lens of radius r_{l} and focal length f_{l} at a distance $z = z_{\mathrm{l}} = 2f_{\mathrm{l}}$ from the screen, which focuses the radiation onto a detector at

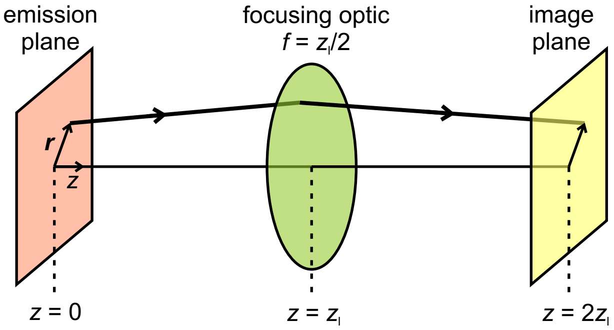

Figure 4.2: The norminal arrangement in a 4-f imaging system of the source plane, lens, and image plane.

$z = 2z_l$, as illustrated in Fig. 4.2. The field at the image plane is therefore

$$\mathbf{E}_1(\mathbf{r},k) = \int^{kr_l/z_l} \mathrm{d}^2\kappa\, e^{i\boldsymbol{\kappa}\cdot\mathbf{r}}\tilde{\mathbf{E}}(\boldsymbol{\kappa}),$$

where we work in the Fraunhofer regime and ignore variation of phase with $\mathbf{r}$. In Ref. [177] this approximation is justified for the imaging at visible wavelengths of the field due to a single electron, whenever $\lambda\gamma^2 \ll z_l$. The justification is based on the fact that the extent of the field profile due to a single electron — the so-called point spread function (PSF) — is small enough, $\sim \gamma\lambda$, that there will be no appreciable difference in the phase acquired by different parts of the PSF during propagation through the imaging system. This assumption can be extended to emission due to multiple electrons, since the field at any given point on the screen is the combined emission from sources only within a transverse extent of $\sim \gamma\lambda$ at the source plane.

The above expression for $\mathbf{E}_1(\mathbf{r},k)$ can be approximated as

$$\mathbf{E}_1(\mathbf{r},k) = \frac{-e}{4\pi^2\varepsilon_0\beta c}\left[\frac{k}{\beta\gamma}K_1\left(\frac{kr}{\beta\gamma}\right) - \frac{1}{r}J_0\left(\frac{r_l}{z_l}kr\right)\right]\frac{\mathbf{r}}{r}, \tag{4.5}$$

for $r_l/z_l \gg 1/\gamma$. Figure 4.3 shows a comparison of the squared moduli of the source function (Eq. (4.1)) and of the focused fields (Eq. (4.5)) for a lens aperture of $r_l/z_l = 0.075$ and $\gamma = 500$ or ∞. We can see that a realistic imaging system suppresses the central part of the field, resulting in a ring with a peak at around $0.4\lambda z_l/r_l$, with emission extending out to $\sim \gamma\lambda$.

For N_e electrons, the total field at the image plane will be the sum of the fields from each individual electron,

$$\mathbf{E}_N(\mathbf{r},k) = \sum_{j=1}^{N_\mathrm{e}} e^{-ikz_j}\mathbf{E}_1(\mathbf{r}-\mathbf{r}_j,k),$$

with the phases corresponding to the longitudinal positions of the electrons, z_i, taken into account. In the limit of an electron bunch that is long compared to the emission wavelength, the radiation phases, kz_j, are evenly distributed and the emission is incoherent. In this case the square modulus of the field on the screen is

$$\left|\mathbf{E}_{N,\mathrm{incoh}}(\mathbf{r},k)\right|^2 = N_\mathrm{e}\int \mathrm{d}^2r'\,\mathrm{d}z\,\rho(\mathbf{r}',z)\left|\mathbf{E}_1(\mathbf{r}-\mathbf{r}',k)\right|^2, \tag{4.6}$$

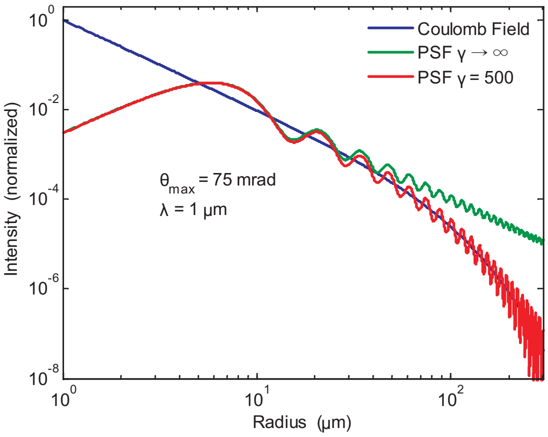

Figure 4.3: Source field on screen for electron with $\gamma = 500$ at wavelength $\lambda = 1\,\mu\text{m}$ (blue), and point spread functions at detection plane after lens of acceptance angle $r_l/z_l = 75\,\text{mrad}$ for $\gamma = 500$ (red) and $\gamma = \infty$ (green). From Loos *et al.* [176].

where we have represented the sum over the $N_\text{e} \gg 1$ electrons as an integral over the normalized electron distribution, $\rho(\mathbf{r}, z)$. For short bunches, however, the emission phases, kz_j, may not be evenly distributed, and coherence effects come into play. These can be quantified by means of the full integral representation of the sum,

$$|\mathbf{E}_N(\mathbf{r},k)|^2 = N_\text{e}^2 \left| \int \mathrm{d}^2 r' \, \mathrm{d}z \, e^{-ikz} \rho(\mathbf{r}', z) \mathbf{E}_1(\mathbf{r} - \mathbf{r}', k) \right|^2 , \tag{4.7}$$

where the emitted intensity scales with N_e^2 rather than with N_e. We, and others [178], have observed transition radiation from laser-accelerated electron bunches to exhibit coherence even at optical wavelengths — we thus go on to study this aspect of it.

4.1.2 Coherent transition radiation

Substituting in $\mathbf{E}_1$ from Eq. (4.5), we obtain

$$\mathbf{E}_N(\mathbf{r},k) = \frac{N_\text{e} e}{2i\pi\varepsilon_0 \beta c} \int^{kr_l/z_l} \mathrm{d}^2\kappa \, \tilde{\rho}(\boldsymbol{\kappa}, k) \frac{\boldsymbol{\kappa} \, e^{i\boldsymbol{\kappa}\cdot\mathbf{r}}}{\kappa^2 + (k/\beta\gamma)^2} . \tag{4.8}$$

We split the Fourier transform of the electron density into independent transverse and longitudinal components, $\tilde{\rho}(\boldsymbol{\kappa}, k) = \tilde{\rho}_\perp(\boldsymbol{\kappa})\tilde{\rho}_\parallel(k)$. Each of these is independently the Fourier transform of the real-space distributions:

$$\tilde{\rho}_\parallel(k) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathrm{d}z \, \rho_\parallel(z) e^{-ikz},$$
$$\tilde{\rho}_\perp(\boldsymbol{\kappa}) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathrm{d}^2 r \, \rho_\perp(\mathbf{r}) e^{-i\boldsymbol{\kappa}\cdot\mathbf{r}}$$

where for the case of a radially symmetric transverse electron distribution, the transform can be expressed as

$$\tilde{\rho}_\perp^\text{s}(\kappa) = \frac{1}{2\pi} \int_0^\infty \mathrm{d}r \, \rho_\perp^\text{s}(r) r J_0(\kappa r).$$

With these considerations in mind, Eq. (4.8) becomes

$$\mathbf{E}_N^{\mathrm{s}}(r,k) = \frac{\tilde{\rho}_{\parallel}(k) N_{\mathrm{e}} e}{\varepsilon_0 \beta c} \int_0^{k r_1/z_1} \mathrm{d}\kappa \, \tilde{\rho}_{\perp}^{\mathrm{s}}(\kappa) \frac{\kappa^2 J_1(\kappa r)}{\kappa^2 + (k/\beta\gamma)^2} \hat{\mathbf{r}}, \tag{4.9}$$

where $\hat{\mathbf{r}}$ is a radial unit vector. Clearly, the Fourier transform of the longitudinal electron distribution, $\tilde{\rho}_{\parallel}(k)$, plays a key role in determining the spectrum of the emitted radiation — and this spectrum therefore contains much information about the longitudinal bunch profile. We later evaluate this integral numerically, in order to model source distributions for experimentally recorded spatially resolved TR spectra.

As a final analytical consideration, provided the Fourier transform of the transverse distribution, $\tilde{\rho}_{\perp}(\kappa)$, does not have significant frequency components above $\kappa_{\max} \approx k r_1/z_1$ or $\kappa_{\max} \approx k/\gamma$, the term κ^2 can be dropped in the denominator of Eq. (4.8) and the integral can be evaluated to

$$\mathbf{E}_N(\mathbf{r},k) = -\frac{\tilde{\rho}_{\parallel}(k) N_{\mathrm{e}} e}{2\pi\varepsilon_0 \beta c} \frac{\gamma^2 \beta^2}{k^2} \nabla \rho_{\perp}(\mathbf{r}). \tag{4.10}$$

Therefore the transverse distribution of coherent transition radiation is proportional to the *gradient* of the transverse distribution of the electron bunch, leading to characteristic ring-shaped images for regular electron bunch profiles. Qualitatively, this is due to the fact that the field arising from a single electron, $\mathbf{E}_1$, has opposite signs on different sides of the electron. This means that fields from neighboring electrons that are emitting in phase almost completely cancel out, except where there is a gradient in the electron density.

4.2 Introduction to the experimental setup and preliminary considerations

As discussed in the previous section, and defined in Eq. (4.9), the spectrum of coherent transition radiation has a strong dependence on the longitudinal profile of the emitting bunch. The transverse distribution of the emission is also related in a known way to the transverse distribution of the coherently emitting electrons (c.f. Eq. (4.10)). Experimental measurements of the spatial and spectral characteristics of CTR emitted from laser-accelerated electrons were carried out at the Max-Planck-Institut für Quantenoptik (MPQ) in Garching, Germany.

MPQ houses the ATLAS Ti:sapphire CPA laser system, which can deliver pulses of energy 1.2 J with a 28 fs FWHM duration on target at a wavelength of approximately 800 nm. These pulses were focused by an $f/22$ off-axis parabola to a spot of transverse size $w_0 = 18.7\,\mu\mathrm{m}$. This resulted in a vacuum intensity of $7.7 \times 10^{18}\,\mathrm{W\,cm^{-2}}$, corresponding to a normalized laser vector potential of $a_0 = 1.9$. In the experiments discussed here the target was a steady-state-flow gas cell. The gas cell was constructed in the form of a miniature cylinder with one open end, where a piston was inserted, as depicted in Figure 4.4. The laser entered through a hole in the piston face, propagating along the cylinder's axis, where it interacted with gas constrained by the piston to an interaction length in the range 2–14 mm. The density of the hydrogen gas in the cell was varied in the range 5–$9 \times 10^{18}\,\mathrm{cm^{-3}}$.

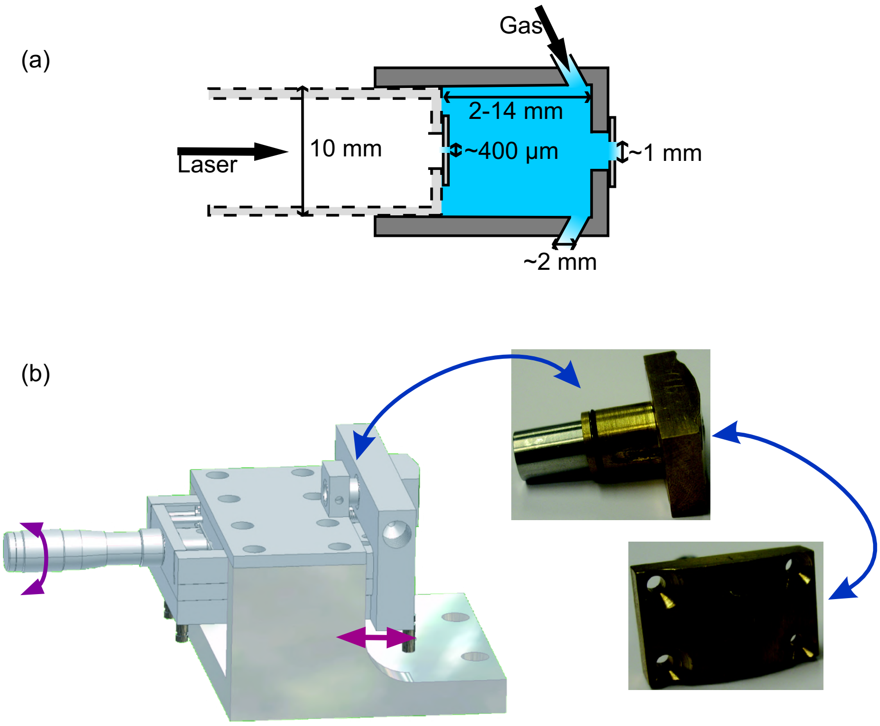

Figure 4.4: Variable length gas cell used as a plasma target in laser-wakefield acceleration experiments at MPQ. (a) Cross-section of the gas cell, showing the the volume occupied by gas (blue shading) and the direction of laser incidence. (b) 3-D schematic of the assembly, with photographs of the gas cell. Image credit: Antonia Popp, MPQ.

This (or a similar) setup has been used successfully for a number of laser-wakefield acceleration experiments at MPQ, the results of some of which are documented in Refs. [38, 99, 100, 112, 117]. The details of the acceleration process will therefore be omitted here. The electron bunches obtained here had characteristics broadly similar to those documented in previous experiments.

Over the course of the experimental campaign we observed quasi-monoenergetic bunches with energies in the range 150–400 MeV; energy spreads of $\geq 10\,\%$ FWHM; bunch charges of 5–20 pC; and transverse pointing fluctuations and divergences of a few mrad. In the following sections we present details of particular shots, as well as multi-shot statistics, side-by-side with the corresponding TR data. However before presenting the experimental setup and data for the spectral distribution measurements that form the crux of this chapter, we first explore some of the considerations and preliminary measurements that underpin them.

4.2.1 Separation of electrons and transition radiation

Due to the large quantities of x-rays generated during scattering and the potential for nuclear activation, it is undesirable for the ultra-relativistic electron bunches to enter any bulk[3] material other than a designated beam dump at the end of the beamline. Therefore,

---

[3] I.e. more than just a thin screen.

detecting the transition radiation requires separating it from the electron bunch.

As Figure 4.1 illustrates, in the forward direction TR is emitted along the propagation direction of the electrons. Although strictly speaking it is emitted into a cone, its opening angle of $1/\gamma$ is of order 1 mrad, which is comparable to the angular divergence of the electron bunch itself. Spatial separation of the forward TR from the electrons is therefore not practical without deflecting the electron bunch. This can be achieved by use of a bending magnet, although this may introduce undesirable constraints on the setup.

In the case of backward TR emission[4], however, the path of the radiation overlaps the path of electron propagation only in the case of an electron bunch that is normally incident on the radiator surface. When the bunch is incident at an angle α to the surface normal, the radiation is reflected at an angle α away from the normal: thus a screen oriented at 45° to the incident electron bunch will lead to TR emission at 90°. In this context, using a tilted screen and observing the backward rather than the forward TR appears to be a straightforward way to avoid interference of the electron bunch with the radiation detector. We carried out such a measurement using the experimental setup outlined in Figure 4.5.

This method, however, has a potential drawback when used in conjunction with a laser-wakefield accelerator. If the laser pulse driving the acceleration is not blocked, it will co-propagate with the electron bunch and be reflected into the TR detector — drowning out the TR signal, or even damaging the detector. On the other hand, introducing a screen to block the laser pulse as it exits the accelerator means the electrons too will cross this screen and generate transition radiation at it. This additional radiation, co-propagating with the electrons, will then be reflected off the tilted screen and enter the detector together with the radiation from the tilted screen, potentially leading to interference or other undesirable effects. Despite these concerns this was the setup used initially, as described below.

4.3 Transverse profile measurements in two-screen experiment

Figure 4.5 illustrates the two-screen version of the setup, where the purpose of the first screen is to block the driving laser beam after it exits the gas cell, and we aim to observe radiation from the second screen, which is tilted at 45° w.r.t. the electron beam path. The radiation is imaged outside the vacuum chamber using an $f/2.5$ Navitar Zoom 7000 lens mounted on a Point Grey Flea2 12-bit CCD camera of resolution 800×600. Table 4.1 gives the absolute positions of the components shown in Figure 4.5.

Due to the fact that upon interaction and exit from the gas cell the intense driving laser beam is scattered throughout the vacuum chamber, the TR screens had to be built into a light-tight enclosure. This consisted of a modular array of aluminium tubes and connectors. Care was also taken to ensure no light enters the imaging setup at the point where the radiation exits the vacuum chamber. The light-tight enclosure, together with the screens, was mounted on a motorized vertical translation stage and could be withdrawn from the

---

[4] Note that backward emission occurs only for reflective screen materials, to which we limit ourselves here.

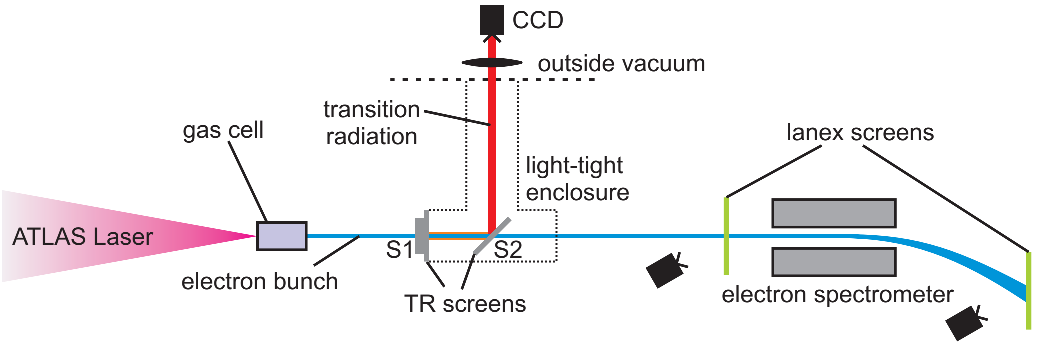

Figure 4.5: Schematic of two-screen TR experiment at MPQ. The laser beam, entering from the left, accelerates an electron bunch in the gas cell. The bunch (blue) enters the light-tight enclosure through one screen (S1) and crosses another (S2), oriented at 45° to its direction of propagation. A camera outside the vacuum chamber captures the transition radiation (red). The electron bunch continues through a lanex screen where its pointing and transverse profile are measured and is deflected by a dipole magnet onto a second screen, before being discarded into the beam dump.

beamline — e.g. during initial set-up of the accelerator.

A custom mount was designed by the author to ensure the TR screens were as flat as possible, while allowing them to be very thin. Pictured in Figure 4.6, it consists of a ring-shaped base with an inside diameter of 23.5 mm, with a thin groove etched into one face. The screen material is placed against this face, and a rubber O-ring matching the groove is overlaid. When the cap is placed over the O-ring and screwed into the base, the O-ring recedes into the groove, gripping the screen material and stretching it out evenly in all directions. Such a mount can either be attached directly to a post; or to a custom-made foot that allows better control over orientation angle; or be mounted flush against the open end of the light-tight tubing.

The most common screen materials used — depending on the requirements — were aluminium foil of thickness 5–10 µm, aluminized Mylar foil of thickness 5 µm, or 10 µm-thick steel. The choice of screen material was made on the basis of a compromise between reliable blocking of the laser beam (esp. at S1), and minimum scattering of electrons. For

Table 4.1: Beamline component locations for two-screen experiment.

Component	z (mm)
Gas cell exit	0 ± 1[a]
TR screen S1	561 ± 2
TR screen S2	676 ± 2
Flea2 CCD	1217 ± 5[b]
Profile lanex	1461 ± 5

---

[a] The gas cell was moved in z during the run in order to optimize acceleration, hence the uncertainty.

[b] This is the sum of the distances from the gas cell exit to the center of S2, and from there to the CCD.

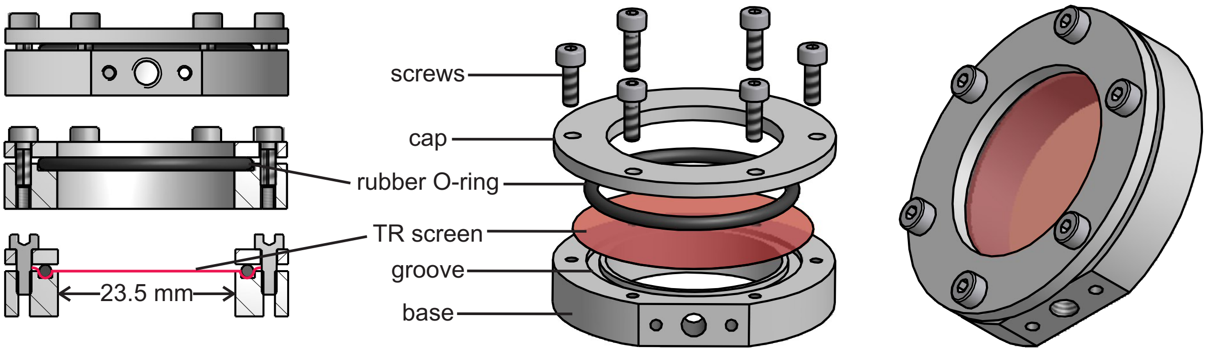

Figure 4.6: Scale drawing of mount for transition radiation screens. The red color and transparency are used solely for illustrative purposes and are not representative of the screen's real-life appearance.

this two-screen setup, the main set of data was taken with 10 µm-thick Al foil blocking the laser at S1 and 5 µm-thick aluminized Mylar foil at S2.

The zoom of the Navitar lens was adjusted so that screen S2 filled the field of view, and stopped down to $\approx f/3$ in order for the depth of field to be enough to have a sufficient part of the tilted screen in focus. Since this depth of field was still much narrower than the distance between S1 and S2, we assumed the radiation generated at S1 would have a minimal contribution to the transverse profile measurements.

4.3.1 Results from two-screen experiment

Using the setup described above, after tuning the accelerator to obtain stable electron beams and making minor adjustments to alignment, transition radiation was observed fairly easily. Part of this ease is attributable to the brightness of the radiation, which far exceeded initial projections based on the bunch charges expected from a laser-wakefield accelerator.

The primary aims of this experiment were an initial detection of optical transition radiation, and a proof-of-principle demonstration of the use of OTR as a transverse profile diagnostic for laser-accelerated electron bunches. However, the transverse profiles of the emitted radiation could not conclusively demonstrate the setup's fitness for this purpose.

Figure 4.7 shows a side-by-side comparison of bunch profiles measured at the lanex before the electron spectrometer (column (a)), and OTR emission at screen S2 (column (b)). The lanex profiles are comparable to those measured without the OTR setup in place, demonstrating that the OTR screens do not scatter the electrons significantly. The emitted OTR appears at the same angular position with relation to the gas cell exit, and it has a transverse size that is strongly correlated with the size of emission from the lanex. The OTR emission, however, exhibits very strong spatial modulations in the majority of cases. Moreover, the intensity of the emitted OTR was not observed to reliably correlate with the fluorescence intensity at the lanex screen, which nominally scales linearly with charge.

These features, together with the brightness of the optical transition radiation, led us to conclude that the observed radiation was most likely coherent. As such, it is possible

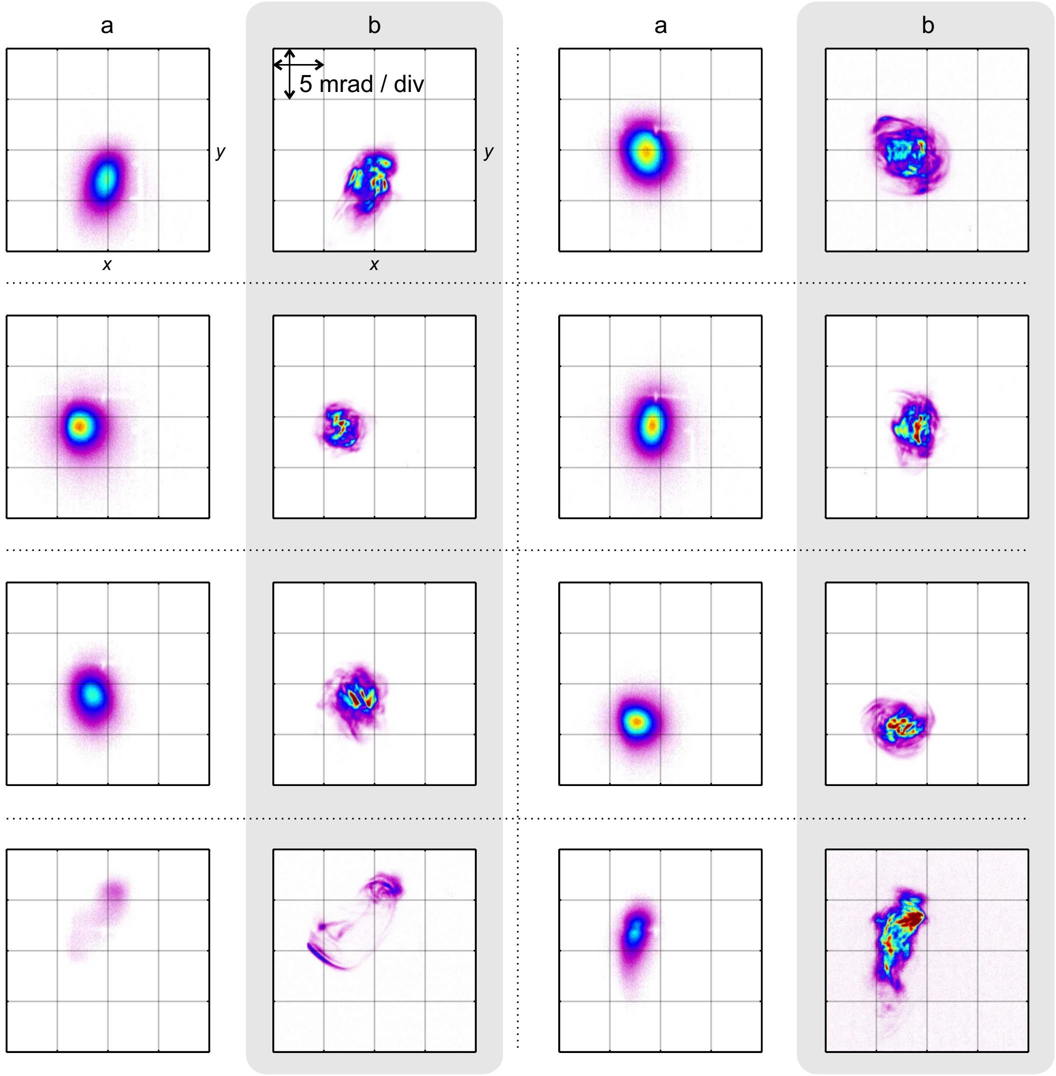

Figure 4.7: Measurements of transverse electron bunch profile at a lanex screen (columns (a)), and corresponding optical transition radiation measurements (columns (b)) for a selection of shots. All plots have an angular scale of 5 mrad per division with respect to the gas cell exit, and the false color scaling is kept constant for each shot.

that the observed modulations are due to interference of radiation from *both* screens S1 and S2, rather than intrinsic features in the transverse profile of the electron bunch. On the other hand, similar “hot spots” in the transverse profile of COTR emission have been observed several meters downstream from the accelerator by Lin *et al.* [179], therefore we can also not exclude the possibility that the observed features correspond to the presence of microbunching at optical wavelengths at certain transverse locations. In order to eliminate any ambiguity, however, we instead reconfigured the setup so as to only observe radiation from a single screen — as we go on to describe next.

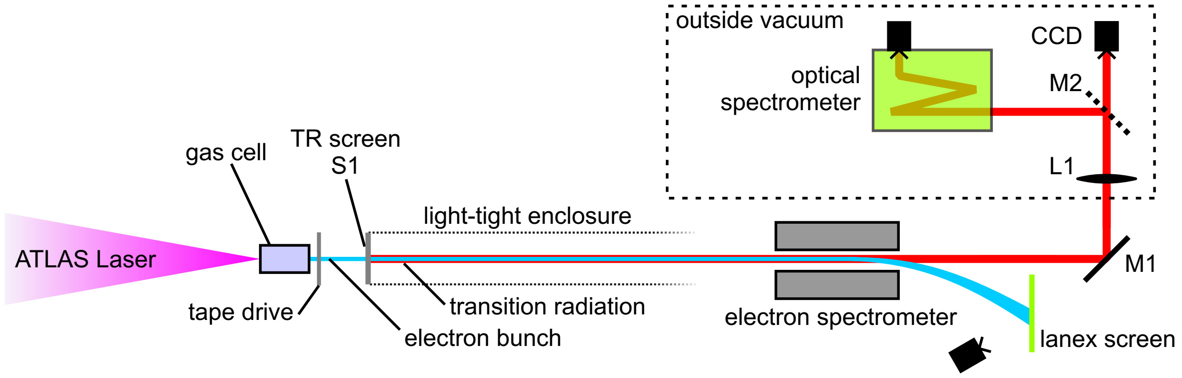

Figure 4.8: Schematic of single-screen TR experiment at MPQ. The laser beam, entering from the left, accelerates an electron bunch in the gas cell. The bunch (blue) passes through a steel tape and through a TR screen. It copropagates with the transition radiation and is then deflected by a dipole magnet spectrometer onto a lanex screen. The transition radiation from S1 (red) is either imaged onto a CCD camera or spectrally dispersed in an imaging spectrometer.

4.4 Transverse and spectral measurements from single TR screen

In order to be certain we are not observing interference between coherent transition radiation from two separate screens, it is necessary to observe forward TR from a single screen. We achieved this by placing a screen close to the gas cell exit and imaging the transition radiation from it after the electron bunch had been deflected by the dipole magnet spectrometer. Figure 4.8 shows a schematic of this setup.

Here we aimed to detect and characterize TR generated very close to the gas cell exit, in order to measure the electron bunch properties at an early stage of its evolution through the drift space. For this purpose, a “tape drive” consisting of two rollers with 20 µm-thick steel tape stretched between them was placed directly behind the gas cell exit. This allowed significant attenuation of the driving laser, but did not block it completely. A TR screen like the one pictured in Figure 4.6 was then placed at $z = 72$ mm from the gas cell exit. It was mounted at one end of a light-tight enclosure that extended along the electron beam path, into the next chamber that is itself shielded from the laser beam.

The TR screen and enclosure were mounted on a vertical translation stage, which allowed the screen to either be retracted to allow free-space propagation of the electrons, or to be exchanged for a pinhole. The pinhole was placed at the same position in z as the screen (± 100 µm), and it was used for setting up the imaging system. The imaging was carried out using a 100 mm-diameter achromatic lens with a focal length of 1000 mm placed just outside the vacuum chamber, approx. 2 m after the gas cell exit (L1 in Fig. 4.8). Meanwhile, the electron bunch was deflected by a permanent magnet dipole spectrometer onto an absolutely calibrated lanex screen.

The transition radiation was directed to another optical table where it was either imaged by a 12-bit PCO pixelfly CCD camera or, using a kinematic flip mirror (M2 in Fig. 4.8),

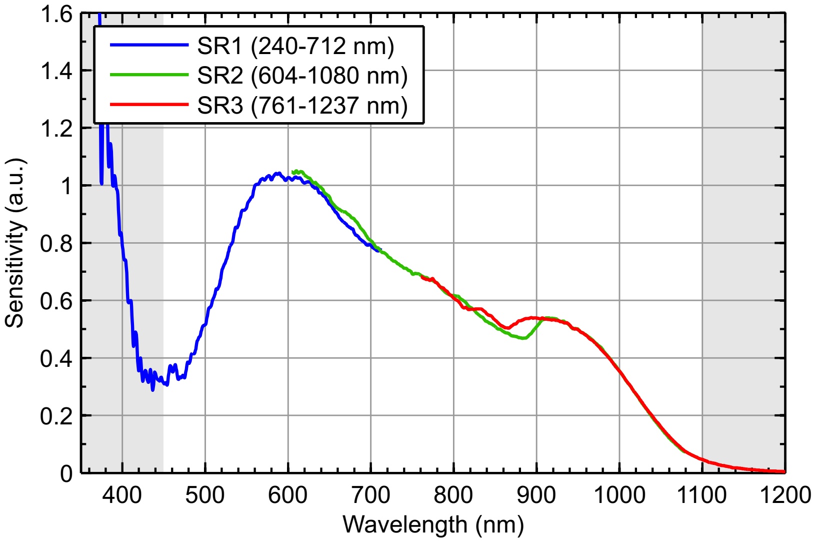

Figure 4.9: Relative spectrometer sensitivity over different spectral ranges, as determined using a blackbody radiation source. The sensitivity is scaled to unity at 632 nm for the SR2 spectral range, where the absolute calibration was carried out. In the grayed out ranges, $\lambda <$ 480 nm and $\lambda >$ 1100 nm, either the sensitivity of the CCD or the intensity of the blackbody source were too low to carry out a reliable calibration.

directed into an Oriel MS 127i Czerny-Turner imaging spectrograph and dispersed onto an Andor DV420 16-bit CCD camera. The pixelfly CCD and the spectrometer's entrance were both ≈ 2 m from the achromat[5] L1, and positioned so that radiation from screen S1 is in focus. The aperture of this imaging system was limited by the size of the exit of the light-tight enclosure and by the electron spectrometer to an estimated half-angle of 20 mrad, corresponding to an aperture of $f/25$.

The grating used in the spectrometer had 400 lines/cm and was suitable for measurements across the operating wavelengths of the silicon CCD, $\approx$ 450–1100 nm, however, the spectral range at any given micrometer setting[6] was limited to about 450 nm. Calibration spectra from a mercury-neon gas-discharge lamp were taken for each of the four different micrometer settings at which TR spectra were measured.

Subsequently, an Ocean Optics HL-2000-CAL light source was used for calibrating the spectrometer sensitivity. The calibration spectra were not recorded for the exact same micrometer settings (i.e. spectral ranges) as those used for taking the data, but pairing them with Hg-Ar discharge spectra allowed them to be adapted for application to the data. A further feature that allowed for this adaptation is the fact that the sensitivity at any given wavelength appeared to be largely independent of the particular micrometer setting, as illustrated in Figure 4.9. A HeNe laser with a known intensity at 632 nm was shone through a 100 µm pinhole and used to obtain an absolute calibration of 7.0×10^{-17} J/count at that wavelength, which was then scaled using the relative calibration to cover the full spectral

---

[5] This arrangement corresponds to a $4f$ imaging system, since the total distance from the object plane to the image plane is four times the focal length of the lens."

[6] The Oriel MS127i is equipped with a micrometer that rotates the grating, thus allowing for different wavelength ranges to be dispersed onto the CCD.

range.

Most of the spectral data was taken with the spectrograph's entrance slit removed, as it was found to be prohibitively difficult to match the transverse location of the imaged radiation with the slit location. This introduced an uncertainty in the absolute wavelength range of the overall spectrum, however since the transverse size of the imaged TR emission was itself limited to $\approx 250\,\mu\text{m}$ by the electron bunch size (see e.g., Fig. 4.10), a relative resolution of $\approx 15\,\text{nm}$ was maintained.

4.4.1 Transverse profile measurements

Before recording transition radiation spectra, we recorded a number of images of the transverse profile of the emitted radiation. The primary purpose of these was to verify the emission of transition radiation with the single-screen setup, but they could also be useful as a diagnosis of the transverse size of the electron bunch. Here we present and analyze the results from these measurements.

Experimental measurements

Figure 4.10a shows a representative sample of transverse profile measurements. The ring-like shape was prevalent among the intensity profiles measured, and is consistent with theoretical predictions for coherent emission in the near field, cf. Eq. (4.10) and Ref. [176].

Note that since this was a near-field measurement, the ring-like shape is not related to the (conical) angular emission profile of TR (Eq. (4.3)), but is rather a result of the transverse coherence of the emitted radiation. Since the transverse electric field at the screen due to a single electron at $\mathbf{r} = 0$ is $\mathbf{E}_r(\mathbf{r}) \sim \mathbf{r} f(r)$ (Eq. (4.1)), the field polarized along either x or y will be positive on one side of the electron and negative on the other (e.g., $E_x(\mathbf{r}) \sim x f(r)$). Therefore the fields due to a number of evenly (and sufficiently densely) distributed electrons will cancel out everywhere, except at the distribution's edge. Correspondingly, the fields due to a general transverse distribution will be proportional to the gradient of the distribution, as we concluded in equation (4.10).

Analysis

According to equation (4.10), $\mathbf{E}_N(\mathbf{r},k) \sim N_e \tilde{\rho}_\parallel(k)\, \gamma^2/k^2\, \nabla\rho_\perp(\mathbf{r})$. As outlined in Sec. 4.1.2, this is a simplification of the general integral, Eq. (4.9), based on the assumptions that the transverse electron distribution does not have spatial frequency components κ_{max} greater than kr_1/z_1 or k/γ, where $k = 2\pi/\lambda$ is the wavenumber of the observed radiation. While the first condition limits the accuracy of our 1^{st} order approximation to $\kappa_{\text{max}} \approx 2\pi/30\,\mu\text{m}$, the second condition imposes a more stringent limit of $\kappa_{\text{max}} \approx 2\pi/300\,\mu\text{m}$.

Further limitations include the fact that we lack direct knowledge of the wavelength of the observed TR, though based on the camera's quantum efficiency curves and observations

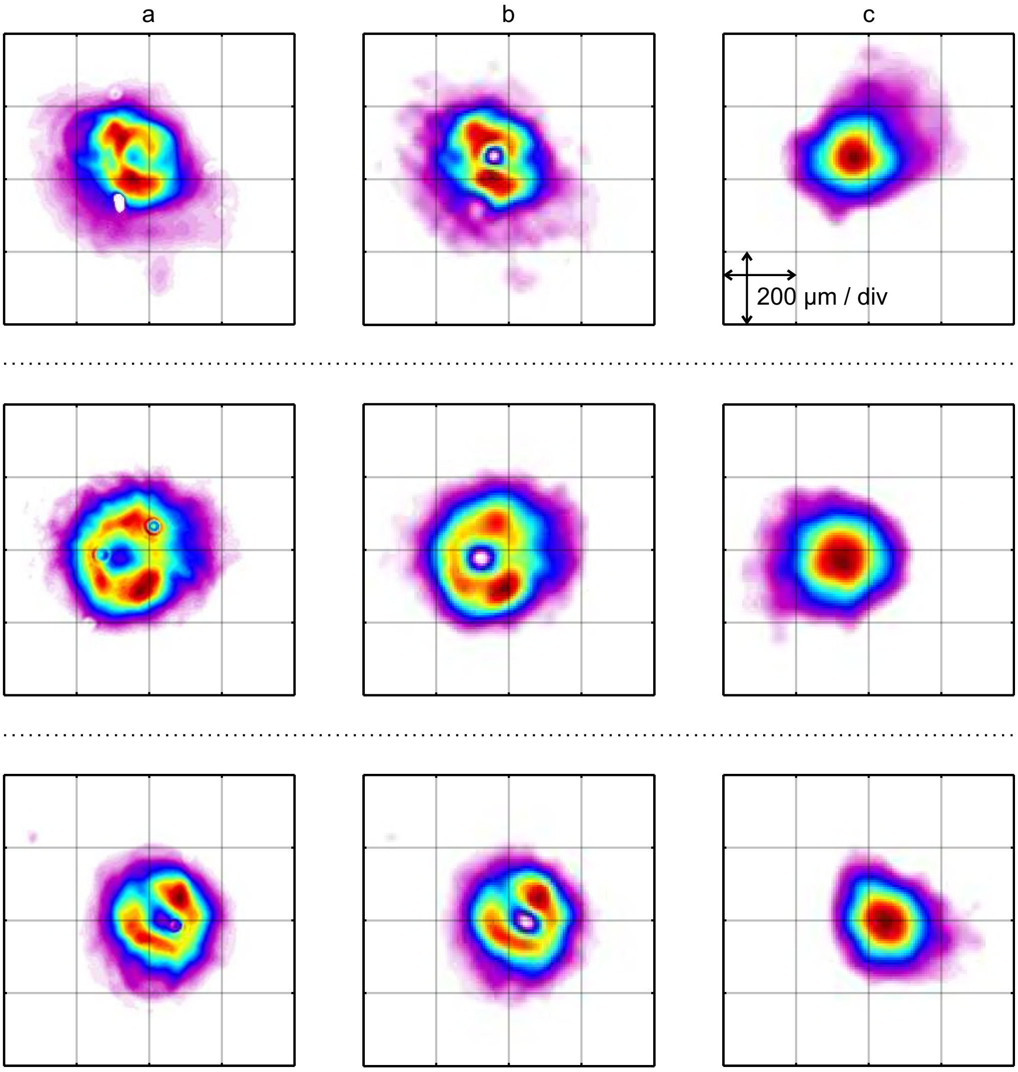

Figure 4.10: Three measurements of transverse profiles of TR intensity are shown in column *a*. Column *b* contains theoretical fits to the data, with the corresponding transverse electron bunch profiles shown in column *c*. The color scales are relative, and normalized to the maximum intensity for each shot, with the same color scale used for each pair of images in columns *a* and *b*.

presented in the following subsection we estimate it to be 500–700 nm. Furthermore, although the electron energy spectra of these bunches exhibit quasi-monoenergetic peaks in the range 200–300 MeV, the energy spreads are rather large (~ 50 MeV FWHM) and in some cases there are multiple peaks, therefore it is difficult to determine the precise energy of the electrons that are emitting the observed coherent transition radiation.

With these considerations in mind, it is nonetheless illustrative to carry out a first order reconstruction of the transverse electron bunch profile by employing the simplified relation $\mathbf{E}_N(\mathbf{r},k) \sim \nabla\rho_\perp(\mathbf{r})$. Since this corresponds to an observed intensity profile $I_{\rm coh}(x,y) \sim |\partial_x \rho_\perp(x,y)|^2 + |\partial_y \rho_\perp(x,y)|^2$, the direct reconstruction of $\rho_\perp(x,y)$ is strictly speaking an ill-posed problem. However, the knowledge that the charge density is non-zero within only a limited range of (x,y), and has a positive peak near the middle of the "ring" of emission, allows us to satisfactorily reconstruct its profile.

We proceed with the reconstruction by generating a sample initial distribution $\tilde{\rho}_\perp(\kappa_x, \kappa_y)$ on a 18×18 grid in transverse Fourier space. Since $\partial_{(x,y)}\rho_\perp = \mathscr{F}^{-1}[i\kappa_{(x,y)}\tilde{\rho}_\perp]$, this distribution results in a real-space intensity profile of $I_{\mathrm{coh}}(x,y) \sim \left(\mathscr{F}^{-1}[i\kappa_x\tilde{\rho}_\perp]\right)^2 + \left(\mathscr{F}^{-1}[i\kappa_y\tilde{\rho}_\perp]\right)^2$. Using the trust-region-reflective algorithm of MATLAB's `lsqnonlin` function, we carry out a least-squares minimization on $I_{\mathrm{coh}}(x,y)$, in order to obtain the target distribution $\tilde{\rho}_\perp(\kappa_x, \kappa_y)$ that most closely corresponds to the observed intensity profile. Note that since $\rho_\perp(x,y) = \mathscr{F}^{-1}[\tilde{\rho}_\perp(\kappa_x, \kappa_y)]$ must be a real function, and since its absolute offset from 0 is unimportant, the total number of parameters to fit is reduced to $18^2/2 + 1$.

Sample results of this process are shown in Figure 4.10, where column (b) shows the reconstructed transverse intensity profiles, and column (c) shows the corresponding density distributions. Clearly, despite the limitations outlined above, it is possible to obtain reconstructed intensity profiles that match the observations very closely.

With further improvements, such as the introduction of spectral filters into the imaging setup, and the inclusion of higher-order contributions in the model used for fitting, it may be possible to employ COTR measurements for accurate reconstruction of electron bunch transverse profiles. It should be reiterated that this treatment pertains to electrons emitting coherently at the observed wavelengths, therefore further studies are required to establish how their transverse profile relates to that of the electron bunch overall.

4.4.2 Spectral measurements at visible wavelengths

Generally, the primary objective of measuring coherent transition radiation is its use as a bunch length diagnostic. This application has been demonstrated several times previously (see e.g. Refs. [110, 180–182]), with recent measurements by Lundh *et al.* at the Laboratoire d'Optique Appliquée in Paris, France, indicating bunch lengths of approx. 1.5 fs rms [97] for laser-accelerated electrons. Most of these measurements were carried out at near-infrared or terahertz frequencies, corresponding to the wavelength range 1–50 µm.

The measurements at visible wavelengths (approx. 480–1100 nm, cf. Fig. 4.9) that we carried out are not dissimilar to those carried out by Glinec *et al.* [178] where the coherence of the observed transition radiation was to a large extent attributed to fine longitudinal structure that the laser pulse imprints on the bunch during the acceleration interaction. Here we present the results of experiments undertaken at MPQ, and after analyzing them reach broadly similar conclusions.

General trend in CTR emission

A total of 126 OTR spectra were recorded during the course of the experiment, with corresponding electron energy spectra. The conditions for electron acceleration remained broadly the same throughout, the main change occurring to gas pressure in the cell, which was varied over the range 110–180 mbar (n_e = 5.3–8.7 × $10^{18}\,\mathrm{cm}^{-3}$) in order to optimize electron acceleration. A simplified analysis was carried out on the data from each shot in order to obtain a summary and look for statistical correlations.

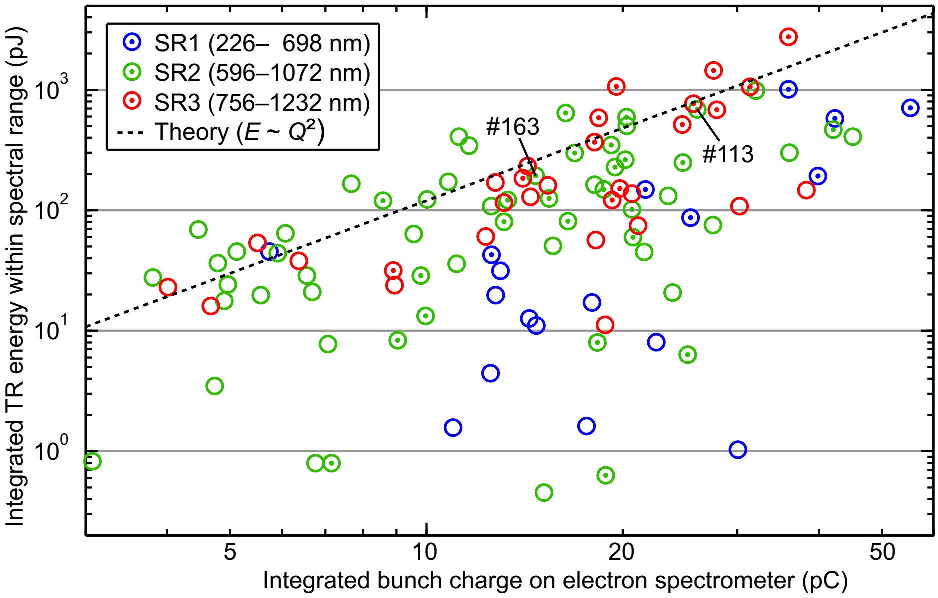

Figure 4.11: Plot of total transition radiation energy (integrated over different spectral ranges) over total electron bunch charge (integrated above 183 MeV). Dotted circles indicate bunches whose energy spectra exhibit clear quasi-monoenergetic features. Due to the optical spectrometer's response (Fig. 4.9), the actual integration ranges for SR1 and SR3 were limited to 450–698 nm and 756–1150 nm respectively. The "theory" line is empirical and corresponds to $E\,[\mathrm{pJ}] = 1.2(Q\,[\mathrm{pC}])^2$.

A plot of total detected OTR energy versus total electron charge, as shown in Figure 4.11, reveals evidence of the quadratic dependence typical of coherent emission. The plotted integrated charge only accounts for electrons of energy above 183 MeV, as the dipole spectrometer deflects those with lower energies outside the imaged lanex. The upper energy limit was 587 MeV, and does not appear to have been exceeded by any electrons. Furthermore, data points with dots inside them indicate that the electron energy spectrum contains one or more clear quasi-monoenergetic peaks. Such data points appear to adhere to the quadratic trend more closely.

Analysis of individual shots

Due to the significant shot-to-shot variation in electron bunch parameters such as charge, energy, and transverse size, and the strong dependence of CTR emission on these, it is difficult to draw general conclusions based on limited data in such a multi-dimensional parameter space. It therefore seems beneficial to more closely analyze individual data points. Here we select a few representative ones for which we have good measurements of both the transition radiation spectrum and the electron energy spectrum.

The analysis is performed by fitting a theoretical prediction derived from equation (4.9),

$$\mathbf{E}^{\mathrm{s}}_{N}(r,k) = \frac{\tilde{\rho}_{\parallel}(k)N_{\mathrm{e}}e}{\varepsilon_0 \beta c} \int_0^{kr_1/z_1} \mathrm{d}\kappa\, \tilde{\rho}^{\mathrm{s}}_{\perp}(\kappa)\, \frac{\kappa^2 J_1(\kappa r)}{\kappa^2 + (k/\beta\gamma)^2}\, \hat{\mathbf{r}},$$

to the observed data. The charge, $N_{\mathrm{e}}e$, is taken to be the total charge detected on the electron spectrometer, while γ is assumed to be the Lorentz factor corresponding to the most prominent peak in the electron energy spectrum; $\gamma \gg 1$ in all cases, therefore we approximate $\beta = 1$. Since we only have transverse resolution along one direction, the transverse form factor, $\tilde{\rho}^{\mathrm{s}}_{\perp}(\kappa)$, is assumed to correspond to a radially symmetric Gaussian transverse electron density distribution of rms size σ_r, i.e., $\tilde{\rho}^{\mathrm{s}}_{\perp}(\kappa) = \exp[-\kappa^2\sigma_r^2/2]/(2\pi)^2$.

Based on knowledge of N_{e} and γ as well as the parameters σ_r and $\tilde{\rho}_{\parallel}(k)$ (the longitudinal form factor of the bunch), we can obtain the transverse and spectral dependence of the emission profile $\mathbf{E}^{\mathrm{s}}_{N}(r,k)$. Since we used the spectrometer without an entrance slit, the transverse extent of this profile in x — i.e. the direction that is *not* spatially resolved — will be preserved and in fact cover an area on the CCD that in principle corresponds to multiple wavelengths. To account for this, the fitting procedure is as follows:

1. Construct a blank map of the actual CCD's full pixel count: 1024 in x (corresponding to different wavelengths) and 255 in y (spatially resolved).

2. For each pixel along the spectrally dispersed direction, x, calculate the transverse intensity profile $\left|\mathbf{E}^{\mathrm{s}}_{N}(\sqrt{x^2+y^2}, k_x)\right|^2$ for the corresponding $k_x = 2\pi/\lambda_x$ and add it to the appropriate pixels on the CCD map, taking into account the spectral bandwidth.

3. Convert the thereby obtained full intensity map into predicted energy incident on each CCD pixel.

MATLAB's `lsqcurvefit` function was used to vary σ_r and a discretized $\tilde{\rho}_{\parallel}(k)$ in order to obtain an intensity distribution that most closely resembled the recorded spectrometer data (in the sense of minimum sum of squared differences over all pixels). Two additional parameters were used to account for absolute vertical offset of the emission, and tilt angle due to grating misalignment in the spectrometer.

In order to illustrate the degree of success of the fitting procedure, Figure 4.12 shows the results for data from two shots, while Fig. 4.13 shows the corresponding electron energy spectra. These two datasets are representative in that they yield a coherence level $\tilde{\rho}_{\parallel}$ that is greater for longer wavelengths, yet exhibits a rather gentle increase. This is consistent with the hypothesis that we are not observing an onset of coherence due to the actual bunch shape — which would be much steeper, and would only occur in the visible part of the spectrum for bunches of rms length $\lesssim 0.3\,\mu\mathrm{m}$ — but are instead seeing coherent emission due to sub-structure in the bunch.

Longitudinal form factors

To illustrate this for the dataset as a whole, Figure 4.14 shows a plot of observed coherence levels (i.e. form factors) obtained from fits for a larger selection of shots, measured at different spectrometer ranges. While there are significant shot-to-shot fluctuations, the coherence levels are mostly within an order of magnitude. Moreover the dependence on wavelength of $\tilde{\rho}_{\parallel}$ exhibits a trend that many of the shots appear to have in common.

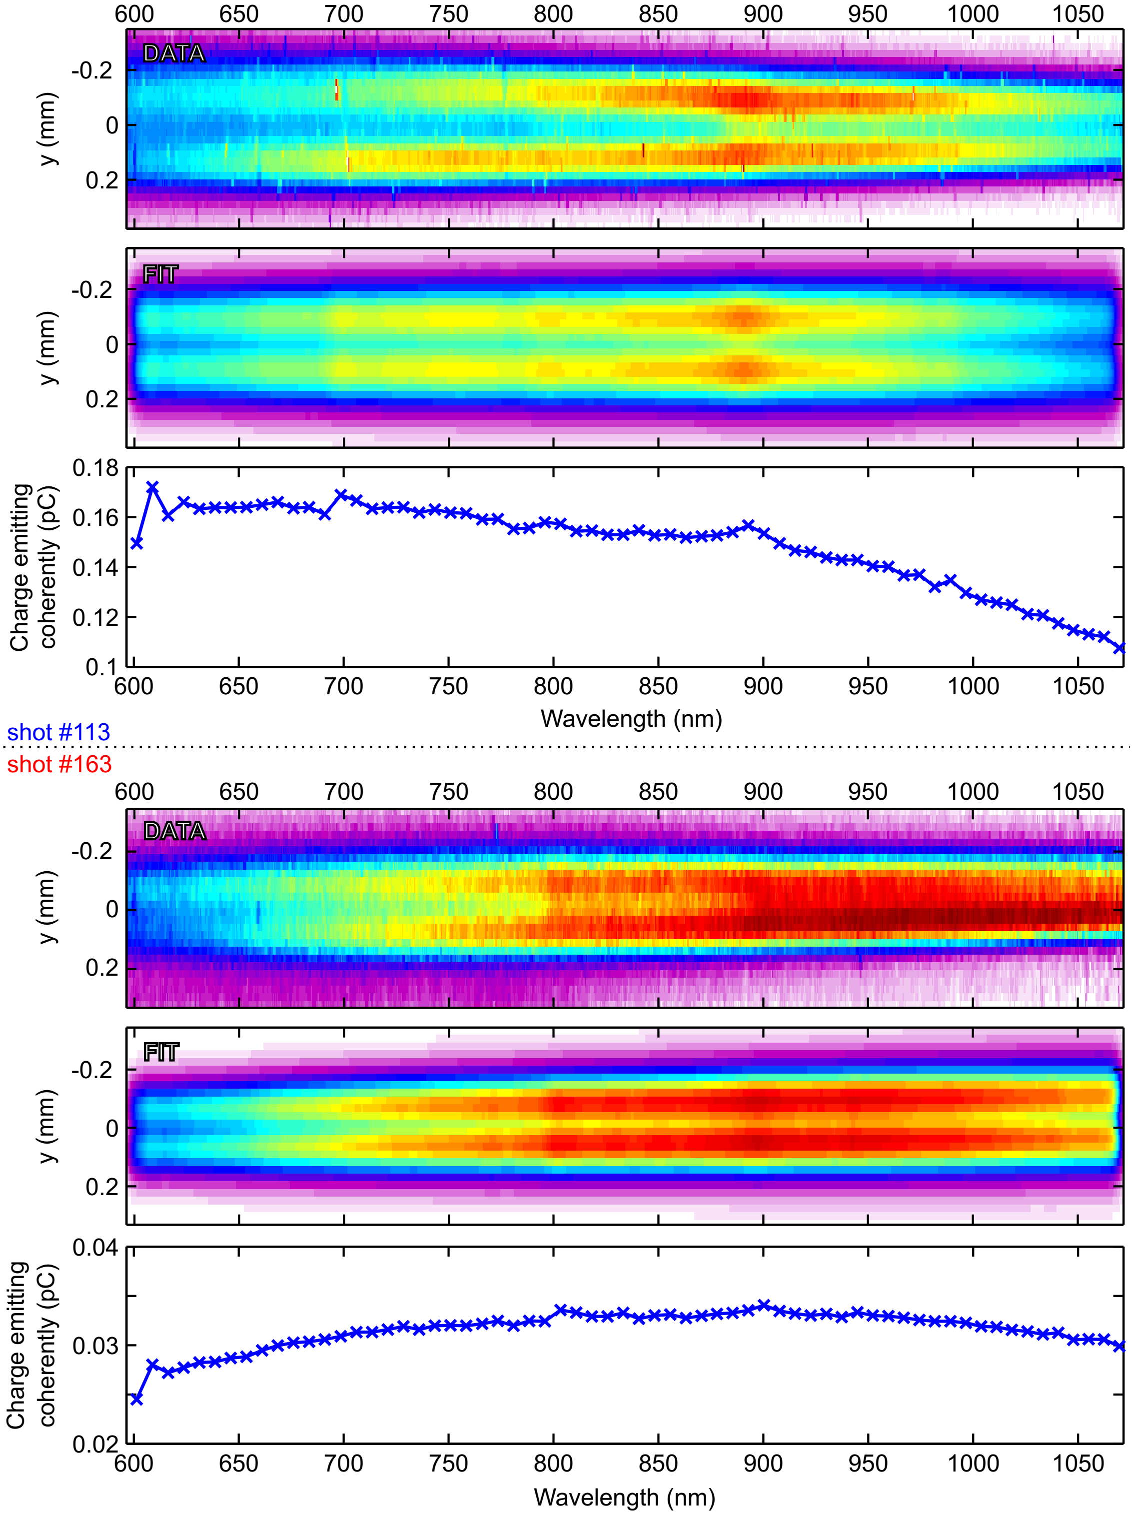

Figure 4.12: False color plots of spatially resolved, sensitivity-corrected transition radiation spectra for two shots (#113: top; #163: bottom). The theoretical fits are shown below the data, together with the corresponding plots of $N_\mathrm{e} e \tilde{\rho}_\parallel(\lambda)$ obtained from the fit. The beam sizes, σ_r, obtained from the fit were 137 µm and 89.5 µm for shots #113 and #163, respectively. The bunch charges and electron energy spectra are given in Fig. 4.13.

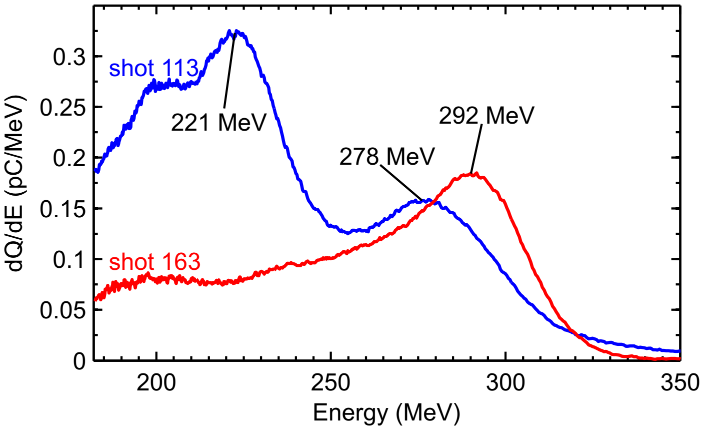

Figure 4.13: Electron energy spectra for shots #113 and #163. The integrated bunch charges were 26.0 pC and 14.7 pC, respectively.

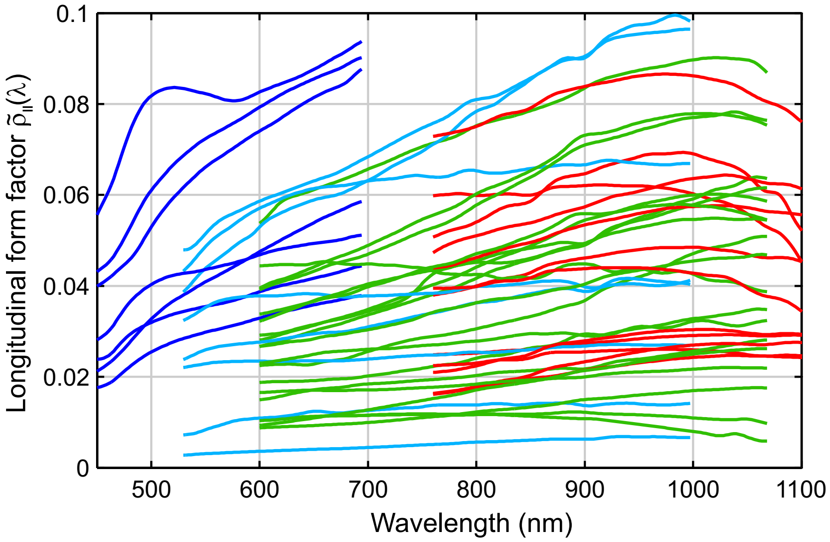

Figure 4.14: Deduced longitudinal form factors (from fitting) for a range of shots where both good OTR signal and good electron energy spectra were observed.

To reiterate, the relatively homogeneous coherence levels that we observed over this spectral range do not appear to be consistent with coherent emission due to the overall bunch shape. The latter would be expected to lead to a difference in coherence level of one or more orders of magnitude over the selected spectral range. This leads us to believe that the observed coherence is due to fine structure in the bunch, rather than a signature of the overall longitudinal bunch profile.

The fine structure that appears to be the cause of the coherence here has been observed before, in particular by Glinec *et al.* [178], while Németh *et al.* [183] present a theoretical justification in terms of the interaction of the electron bunch with the laser pulse during acceleration. Evidence for this interaction is presented by, among others, Mangles *et al.* [58], and in Chapter 5 of the present thesis. Although the distinct shapes of the longitudinal form factors shown in Figure 4.14 hint at a rather more systematic source for this fine structure, it is difficult to identify one beyond the effects described in the preceding references.

In the next chapter we describe an experimental campaign where CTR was measured over a significantly broader spectral range. In that case the results obtained in the visible part of the spectrum were broadly consistent with those presented here in terms of coherence level — though they exhibited a more structured dependence on wavelength. Measurements at longer wavelengths provided context for the coherence levels observed in the visible part of the spectrum, and allowed reconstruction of the full bunch profile. Before continuing to a discussion of that experiment, we take a moment to consider whether the radiation detected in the experiment described here could in fact have been something other than transition radiation.

A note on coherent synchrotron radiation

Since in this setup we used a dipole spectrometer to bend the electron bunch away from the path of the transition radiation, it is possible that some of the observed emission is bending magnet radiation (also known as synchrotron radiation, SR) that co-propagated together with the transition radiation. It is indeed even possible that this radiation was emitted coherently, thereby enhancing its intensity to levels comparable to those of the CTR.

This possibility was considered in some detail, especially noting the fact that the general wavelength dependence of the observed spectra bears some resemblance to the predicted wavelength dependence of SR emitted by $\sim 100\,\mathrm{MeV}$ electrons going through a dipole magnet of field $\sim 0.2\,\mathrm{T}$ (see e.g. Refs. [134, 135]), which are in line with our operating conditions. This suggestion, however, is subject to a number of inconsistencies. The greatest of these stems from a mismatch between the expected transverse profile of the synchrotron radiation on the CCD, and what is observed.

Since this synchrotron radiation would be generated in the bending magnet, around $0.5\,\mathrm{m}$ upstream from the focusing optic (L1 in Fig. 4.8), two rays that are divergent by just $2\,\mathrm{mrad}$ at that point would end up around $3\,\mathrm{mm}$ apart in the image plane. Taking into account the fact that at that point the electron bunch has a transverse size of $\geq 2\,\mathrm{mm}$ and a divergence of at least $\sim 2\,\mathrm{mrad}$ rms, and that the natural cone of emission of SR for our conditions has a half-angle of approx. 10 mrad (Eq. (3.28) in Ref. [135]), we would expect this to result in a diffuse image that covers several millimeters in the image plane. Instead, we consistently observe well-defined images of transverse size on the order of several 100 µm, consistent with radiation originating from the TR screen.

Moreover, the expected synchrotron radiation intensity — assuming incoherent emission — is several orders of magnitude less than what we observe. This difference *could* be accounted for by allowing for reasonable levels of coherent emission[7], however in that case the spectral profile of the radiation will again be primarily dictated by the longitudinal form factor of the bunch. This largely invalidates the initial reason for suggesting that we may be detecting synchrotron radiation — namely the characteristic spectral response — but it also means that even if we *were* looking at SR, the conclusions we drew above regarding the "flat" dependence of coherence level on wavelength still hold true.

---

[7] Although this is unlikely, in part due to the relatively large transverse bunch size at that point.

Chapter 5

Bunch Characterization via Broadband Measurement of Coherent Transition Radiation

In the previous chapter we described some preliminary experiments to characterize the emission of coherent transition radiation by electron bunches from a laser-wakefield accelerator. We observed coherence in the visible part of the emitted spectrum, with indications that the degree of coherence increased at longer wavelengths, cf. Fig. 4.14. This is also consistent with the fact that a significant fraction of the electrons in the bunch will be emitting coherently at wavelengths approaching (or longer than) the bunch length. Quantitatively, ignoring transverse effects, the spectral intensity of the emitted radiation is proportional to the Fourier transform of the longitudinal bunch profile,

$$\frac{\mathrm{d}U^2}{\mathrm{d}\omega\mathrm{d}\Omega} = \frac{\mathrm{d}U_{\mathrm{GF}}^2}{\mathrm{d}\omega\mathrm{d}\Omega} N_{\mathrm{e}}^2 \left|\tilde{\rho}_{\parallel}(\omega)\right|^2, \tag{5.1}$$

where $\mathrm{d}U_{\mathrm{GF}}/(\mathrm{d}\omega\mathrm{d}\Omega)$ is the angular emission profile according to the Ginzburg-Frank formula, Eq. (4.3), N_{e} is the total number of electrons, and $\tilde{\rho}_{\parallel}(\omega) = \mathscr{F}[\rho_{\parallel}(t)]$ is the Fourier transform of the longitudinal bunch profile[1].

Clearly, then, for a comprehensive longitudinal profile diagnostic for femtosecond-scale bunches it is necessary to characterize the emission of coherent transition radiation over a range of wavelengths that extends to the near-infrared and terahertz parts of the spectrum. The 1–10 µm spectral range is indeed where a significant onset of coherence is expected for electron bunch lengths on the order of a micrometer (i.e., a few femtoseconds). Coherence in this range has been detected previously, see e.g., Refs. [180, 181], however those measurements were not frequency-resolved. A frequency-resolved measurement up to 5.5 µm was carried out recently by Lundh *et al.* [97], although, since the spectrum was recorded by a monochromator, data was only taken at a single frequency for each shot. The full reconstructed spectra in Ref. [97] therefore combine data from multiple shots, with a corresponding loss of fidelity due to the shot-to-shot fluctuations inherent to laser-wakefield

---

[1] In the Chapter 4 we had $\rho_{\parallel}$ as a function of z and $\tilde{\rho}_{\parallel}$ correspondingly a function of k. Since we assume $z = ct$ and $\omega = ck$, the two definitions are equivalent.

accelerators.

In this chapter we present the first single-shot measurements of coherent transition radiation (CTR) spanning the wavelength range 0.4–7 µm. This is achieved by simultaneously measuring the radiation using a conventional visible spectrometer based on a Si CCD sensitive to wavelength below 1.1 µm, a liquid nitrogen-cooled InGaAs diode array for the range 1-2.2 µm, and two custom-built pyroelectric detector arrays for wavelengths longer than 1.7 µm.

We first describe the experimental apparatus and data acquisition procedures. We then go on to discuss the possibility for reconstructing the full bunch profile using a phase-retrieval algorithm derived from ones used in crystallography, which we adapt for the purpose at hand. Based on this analysis of CTR data — to our knowledge the first of its kind — we analyze the dependence of bunch length on the different acceleration parameters. Subsequently, we present and analyze evidence for the acceleration of multiple bunches, and the dependence of this phenomenon on the parameters of the laser-plasma system. Lastly, we consider the possibility that an anomalously high level of coherence at a small range of short wavelengths ($\sim 500 \pm 100$ nm) is the result of interaction between the driving laser and the electron bunch during acceleration.

5.1 Experimental setup

The experimental setup used for these measurements was broadly similar to that used in the first experimental campaign, described in Sec. 4.2, at least insofar as the production of ultra-relativistic electron bunches is concerned. The collection and characterization of transition radiation, on the other hand, were carried out in a more sophisticated way. The overall setup is shown in Fig. 5.1, while Fig. 5.2 illustrates the gas target, CTR collection optics, and THz spectrometer.

The laser energy on target was 1.4–1.6 J within a FWHM pulse length of 28 fs, focused to a waist of spot size of $w_0 = 18.7 \pm 1.2\,\mu$m. This gave a peak intensity of about $1.2 \times 10^{19}\,\mathrm{W\,cm^{-2}}$, corresponding to a peak normalized vector potential in vacuum of $a_0 = 2.4$.

The gas cell pictured in Fig. 4.4 was again used as the plasma target where acceleration took place. The length was varied in the range 2–14 mm, while the H_2 pressure was varied in the range 65–260 mbar, corresponding to plasma densities of 3.1–$12.6 \times 10^{18}\,\mathrm{cm^{-3}}$. Upon exiting the gas cell, the accelerated electron bunch passed through a pair of 20 µm-thick steel tapes immediately behind the cell. The first tape served to block the laser beam, while the second one was the source of transition radiation. The transition radiation was reflected aside by a pellicle placed directly behind the tape drive[2], while the electron bunch continued forward; the bunch was dispersed by a dipole magnet spectrometer, and its charge and energy spectrum were recorded by means of detecting the fluorescence of an absolutely

---

[2] Since the distance between the tape drive and pellicle was much smaller than the CTR formation length at the observed wavelengths, this setup should minimize the possibility of interference between TR from the tape drive and TR from the pellicle.

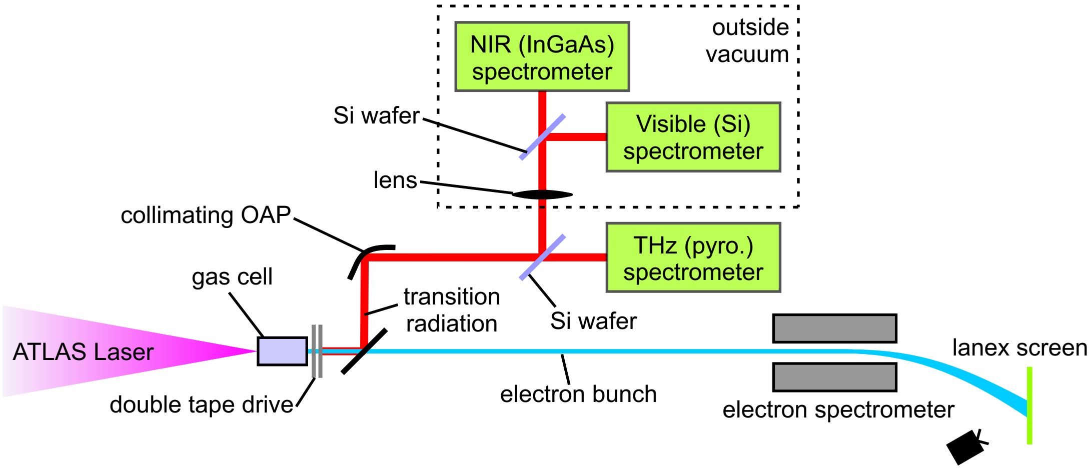

Figure 5.1: Experimental setup for characterization of coherent transition radiation across the visible, near-infrared and terahertz ranges of the electromagnetic spectrum. See text for details.

calibrated lanex screen.

After reflection off the pellicle, the transition radiation was collimated by a 2" off-axis paraboloid of focal length 19 cm. This corresponds to an f-number of $f/7.5$, or an acceptance half-angle of 133 mrad. Part of the collimated radiation propagated forward through a 20 µm-thick silicon wafer, and was reflected into a separate vacuum chamber that housed the THz spectrometer. This spectrometer was based on a design employed for CTR measurements at DESY, which is described in more detail by Delsim-Hashemi [184]. It employs two gratings; the zeroth order reflection from the first one is directed towards the second one. They each dispersed different parts of the spectrum onto two pyroelectric detector arrays, where custom-made gold mirrors were used to refocus the radiation onto the detectors. Two different grating configurations were available, allowing measurements either in the range 1.7–7 µm or 5–25 µm. The raw detector signal was processed and digitized by an array of electronics described in Ref. [184].

The first silicon wafer mentioned above directed part of the radiation out of the vacuum chamber and over to a separate in-air spectrometry setup. There, a 60-cm-focal-length lens was used to refocus the collimated transition radiation onto the entrance slit of an Oriel MS260 imaging spectrometer based on a 1024 × 256 silicon CCD chip, and onto the entrance slit of a Princeton Instruments near-infrared (NIR) spectrometer based on a 1024-element OMA-V liquid nitrogen-cooled InGaAs diode array. The former, hereafter referred to as the "visible" spectrometer, was used for detection in the spectral range 421–1096 nm, while the latter was used over three different ranges — each spanning approx. 700 nm — at wavelengths between 1033 nm and 2135 nm.

Since silicon has good reflectivity ($\gtrsim 0.3$) over both the visible and near-infrared, but only transmits wavelengths longer than ~ 1050 nm [185], a second silicon wafer served as the beamsplitter between the visible and NIR spectrometers, which measured the reflected and transmitted radiation, respectively. By this token, a third silicon wafer (not pictured in Fig. 5.1) was placed directly at the entrance of the NIR spectrometer in order to eliminate

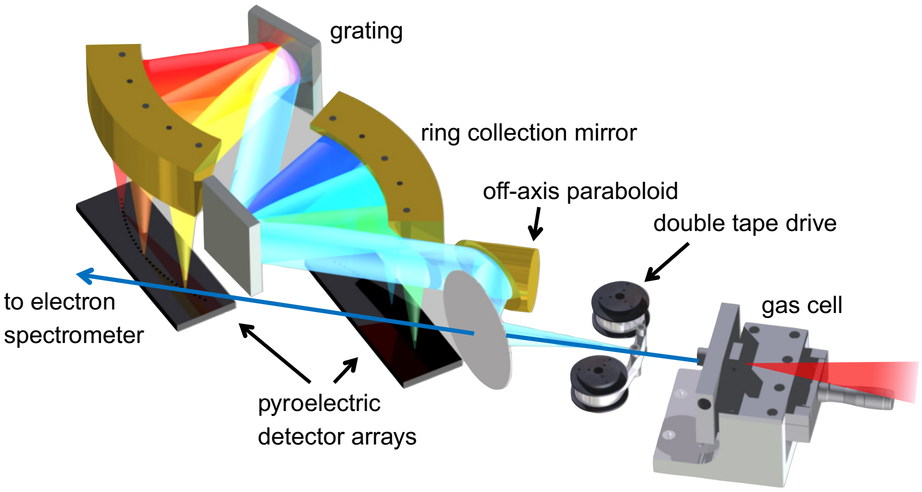

Figure 5.2: Detector setup for CTR measurements in the THz range. The laser-accelerated electrons propagate along the path indicated by the dark blue arrow and generate transition radiation at the double tape drive, which also serves to block the laser. The transition radiation is collimated using an off-axis paraboloid and dispersed by a pair of gratings, after which it is refocused onto two pyroelectric detector arrays so that different wavelengths are incident on different detectors. Image credit: Matthias Heigoldt, MPQ.

stray laser radiation that could manifest as anomalous signal due to higher-order reflection off the NIR spectrometer grating.

Through the use of an 1100 K tungsten blackbody source and an Ocean Optics HL-2000-CAL calibration source, the relative spectral responses of the visible and NIR spectrometers were established. A 50 mW helium-neon laser placed at the CTR source position was used with a chopper in order to obtain an absolute sensitivity calibration for the visible spectrometer, from which an absolute calibration of the NIR spectrometer could also be established. Due to lack of easy access to a well-characterized terahertz source, however, the relative calibration of the THz spectrometer is only based on knowledge of the nominal reflectivities of the gratings and of the spectral range covered by each of the pyroelectric detectors. The absolute level of terahertz emission was determined by means of cross-correlation of readings in the spectral region overlapping that of the NIR spectrometer. The terahertz spectrometer will be absolutely calibrated during an upcoming timeslot at a light source.

5.2 Bunch profile reconstruction

The primary objective of measuring the spectrum of coherent transition radiation from electron bunches is to deduce the bunch length, and potentially the full longitudinal bunch profile. As we showed in Chapter 3, knowledge of the bunch length is crucial to assessing the performance of future free-electron lasers driven by electrons from a laser-wakefield accelerator. In this section we describe a phase-retrieval algorithm and — using synthetic data — demonstrate its ability to closely reconstruct the bunch profiles for many common cases, while also outlining its limitations in the context of the actual spectral range of our experimental measurements. We then carry out the first application of a phase-retrieval algorithm on CTR spectra taken during the experiment in order to reconstruct the profiles of the accelerated bunches. In the following sections we discuss the dependence on acceleration parameters of the bunch profiles deduced in this way.

5.2.1 Introduction to phase-retrieval algorithms

In many situations in experimental science, the measured data is the Fourier transform of the information being sought. Commonly, however, the measurements can only capture the amplitude and not the phase of this transform — thereby not allowing for a simple inverse transform. This is a problem pertinent to disciplines as diverse as x-ray crystallography, astronomy, electron microscopy, holography, particle scattering, as well as many other problems in optics. It is also a key limitation in the context of retrieving the longitudinal electron bunch profile from the transition radiation spectrum, for which we cannot measure phase information.

The notion that phase information can, nonetheless, be retrieved *a posteriori* is underpinned by a 1952 note by Sayre [186], who pointed out that a function with finite support is uniquely defined by knowledge of its Fourier transform at a limited number of points. In particular, a function with support in real space over the range $x \in [-a/2, a/2]$ is completely specified by the values of its Fourier transform at the points $k = 0, \pm 2\pi/a, \pm 4\pi/a, \ldots$. Therefore, in an experimental measurement it is possible to oversample the transform's amplitude by measuring it at a higher resolution than this minimum, and use the additional data to deduce the phase of the limited set of points required for reconstruction of the function. The first practical phase-retrieval algorithm was introduced in 1972 by Gerchberg and Saxton [187], for the purpose of reconstructing the 2-D profile of an object from its diffraction pattern. The Gerchberg-Saxton algorithm is of an iterative nature, as are all others that have come after it; there is no known one-step reconstruction method.

Gerchberg-Saxton algorithm

We first present the Gerchberg-Saxton algorithm, as its operation informs that of the more advanced ones discussed later. Our treatment follows that of Fienup, who presents a comparison of phase retrieval algorithms in Ref. [188]. We introduce the function $f(x)$ that we are trying to reconstruct based on the known modulus of its Fourier transform, $|F(k)| = |\mathscr{F}[f(x)]|$. The function $g_n(x)$ is our estimate of $f(x)$ at the n^{th} iteration, and $G_n(k)$ is its Fourier transform. Figure 5.3 illustrates the four-step iteration cycle.

The algorithm continually switches between Fourier space and real space, and ensures that the candidate estimate g matches the constraints present in each domain. In the Fourier domain the only constraint is that the modulus of G must be equal to the measured amplitude of the spectrum, $|F|$. Since we have no phase information about F, however, we instead retain the current phases of G, $\psi(k)$. The modulus of G is only set to $|F|$ for the points at which we have measured values of $|F|$; the iterative estimate G is preserved at the

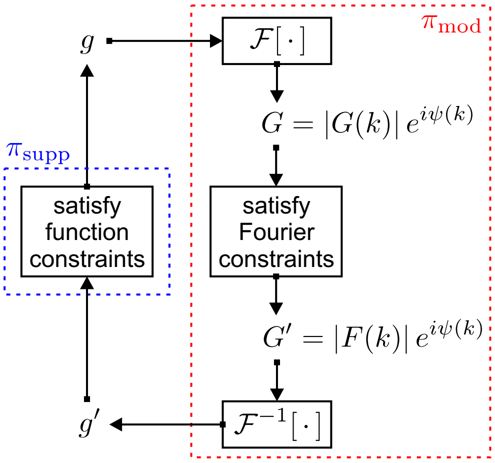

Figure 5.3: Outline of the Gerchberg-Saxton phase-retrieval algorithm. For further details on each step see the text. Dashed boxes indicate the steps that form projections of a candidate function onto the support constraints, π_{supp}, and the Fourier-space modulus constraints, π_{mod}.

remaining points. According to the formalism utilized by Elser [189] this — together with the forward and inverse Fourier transforms — constitutes a “projection”, π_{mod} of g onto the closest function that satisfies the Fourier-space modulus constraint, $g' = \pi_{\text{mod}}[g]$.

The real-space constraints can be more diverse. In particular, we require $f(x)$ to be zero outside a certain support; additionally, we require it to be positive-valued inside the support. We can also require it to be a real function, however this is automatically satisfied for all iterates g_n if the initial estimate g_0 is real-valued[3]. We can therefore define a set of points γ at which $g'_k(x)$ violates the support and non-negativity constraints, for which we set the function's value to zero:

$$g_{k+1}(x) = \begin{cases} g'_k(x), & x \notin \gamma, \\ 0, & x \in \gamma. \end{cases} \tag{5.2}$$

In Elser's formalism, this constitutes the support projection, π_{supp}.

The modifications to g and G described above are the *minimum* changes to the functions that are required to make them compliant with the constraints in each domain. As Fienup points out, each iteration leads to a reduction of the error on the estimate, where the error at the k^{th} iteration is defined as

$$E_k = \sqrt{\sum_x \left| g_{k+1}(x) - g'_k(x) \right|^2}. \tag{5.3}$$

The Gerchberg-Saxton algorithm is therefore also known as the error-reduction algorithm. However, even though at each iteration the difference between the estimate, g_k, and the target function, f, decreases, this algorithm is prone to stagnation. After a number of iterations convergence can become very slow, and sometimes the candidate function can converge towards a local minimum [188]. This is in part a symptom of the fact that the algorithm

---

[3] The Fourier transform of a real-valued $g(x)$ is characterized by $G(-k) = G^*(k)$, correspondingly $\psi(-k) = -\psi(k)$. Since $|F(-k)| = |F(k)|$, we will also have $G'(-k) = G'^*(k)$, and thus a real-valued $g'(x)$. In other words, applying the Fourier-domain constraint to a real-valued input will yield a real-valued output.

has a tendency to “explore” a relatively small set of candidate functions with respect to the starting point $g_0(x)$.

Hybrid input-output algorithm

The Gerchberg-Saxton algorithm can be enhanced in several ways. One of them involves recognizing that $g_k(x)$ does not need to be the current best estimate of $f(x)$, but can instead be thought of as the driving function for the next input, $g'_k(x)$. In this sense the left-hand side of the procedure in Figure 5.3 need not force $g(x)$ to satisfy the actual real-space constraints, but can instead be an arbitrary manipulation, such that the next $g'(x)$ will be closer to convergence. This approach forms a class of algorithms known as “input-output” algorithms.

Fienup compares several such input-output algorithms, as well as a number of others, alongside the Gerchberg-Saxton algorithm. Empirically, the best-performing one appears to be the hybrid input-output (HIO) algorithm, in which the input function $g_{k+1}(x)$ is obtained from the output $g'_k(x)$ via

$$g_{k+1}(x) = \begin{cases} g'_k(x), & x \notin \gamma, \\ g_k(x) - \beta g'_k(x), & x \in \gamma, \end{cases} \tag{5.4}$$

where γ is again the set of points at which $g'_k(x)$ violates the support and non-negativity constraints, and β is a free parameter. Fienup’s numerical tests — where an object is reconstructed from a 2-D diffraction pattern — indicate best performance for $\beta = 1$. The hybrid input-output algorithm has been applied successfully for reconstruction of object data from a range of diffraction patterns, notably by Marchesini *et al.* [190].

Difference map algorithms

Throughout this subsection we referred to the projections π_{supp} and π_{mod}, part of Elser’s nomenclature in Refs. [189, 191]. The Gerchberg-Saxton algorithm consists of alternating application of this pair of projections: $g_k = \pi_{\text{mod}} \circ \pi_{\text{supp}} \ldots \circ \pi_{\text{mod}} \circ \pi_{\text{supp}} [g_0]$. Elser introduces another class of algorithms that utilize so-called “difference maps” to facilitate convergence to a fixed point, i.e., a function f that remains unchanged under either of the two projections.

For further background theory on difference maps and their application to phase retrieval we direct the reader to Ref. [189] and references therein; here we simply present Elser’s result for the optimum difference map:

$$\begin{aligned} g_{k+1} = g_k + \beta_{\text{dm}} \Big\{ &\pi_{\text{supp}} \left[(1 + \beta_{\text{dm}}^{-1}) \pi_{\text{mod}} [g_k] - \beta_{\text{dm}}^{-1} g_k \right] \\ &- \pi_{\text{mod}} \left[(1 - \beta_{\text{dm}}^{-1}) \pi_{\text{supp}} [g_k] + \beta_{\text{dm}}^{-1} g_k \right] \Big\}, \end{aligned} \tag{5.5}$$

where β_{dm} is a free parameter. Elser selects unity as a suitable value for β_{dm}. Incidentally, the projection carried out by the difference map algorithm with $\beta_{\text{dm}} = 1$ is equivalent to the projection carried out by the hybrid input-output algorithm with $\beta = 1$, which Fienup found to be the optimum value for that case. Keeping this equivalence in mind, we later compare the performance of these two algorithms when applied to 1-D CTR data.

Choice of function support and initial estimate

The initial estimate, $g_0(x)$, is generally taken to be a randomly generated real function — not necessarily compliant with the support constraint — in order not to bias the phase retrieval algorithm towards a particular shape. This approach allows an algorithm to be run multiple times for different initial values of g_0 in order to obtain a set of target functions that match all constraints, and to correspondingly verify the uniqueness of the reconstructed function.

In addition to this, an appropriate choice of support is crucial to the successful execution of a phase-retrieval algorithm. It has to be large enough to accommodate the entire reconstructed function, while being sufficiently restrictive for quick convergence to an unambiguous target function to occur. In particular, an excessively large support may yield significantly differing results for runs with different initial conditions.

The support can be chosen on the basis of additional information: for example in the case of x-ray diffraction, a low-resolution image of the object obtained using conventional optics or scanning electron tomography can inform the choice of support; or on the basis of physical considerations: e.g. the length of an electron bunch from a laser-wakefield accelerator is expected to be limited by the bubble size, or plasma wavelength. In addition, the autocorrelation function of the measured spectrum, $\mathscr{F}^{-1}\left[|F|^2\right]$, can be used as an indication of the outer limits of the reconstructed object[4], provided $|F|$ is known over a sufficiently broad range of frequencies. Marchesini *et al.* [190] give further details on possible support choices, with references to more detailed treatments of the relevant methods.

Adaptive support selection: the Shrinkwrap algorithm

In light of the importance of the choice of function support, a key recent development has been the suggestion by Marchesini *et al.* to adaptively change the support during the reconstruction process, shrinking it around subsequent iterates of the candidate function [190]. The initial support is obtained by thresholding the autocorrelation of the diffraction pattern (or spectrum) after applying a Gaussian blur to it. The HIO algorithm is then applied for 20 iterations with $\beta = 0.9$. At that point the candidate function is convolved with a Gaussian of rms σ, and a threshold is applied at 20% of its maximum in order to obtain a new support. After another 20 iterations of the HIO algorithm the support is recalculated again by the same means, and so forth. Initially the blurring Gaussian has width $\sigma = 3$ pixels, though every 20 iterations this is decreased by 1% until it reaches 1.5 pixels — thereby each time the support is recalculated it follows the outline of the candidate function yet more closely. Correspondingly, this is known as the “Shrinkwrap” algorithm.

This algorithm has been applied successfully to a range of diffraction-related problems, demonstrating the possibility for successful reconstruction without *a priori* information [192, 193]. However, after some attempts we found its present form inadequate for the reconstruction of bunch profiles from hypothetical or real CTR spectra, presumably due to

---

[4] As a rule of thumb, the size of the reconstructed object is approximately half the size of the autocorrelation function.

the fact it is optimized for operation on 2-D data. In the following subsection we describe a new algorithm that combines several of the ideas presented here and is tailored to reconstruction of longitudinal bunch shapes from CTR spectra.

5.2.2 Algorithm for reconstruction of bunch shapes from CTR data

A comprehensive survey of the algorithms described previously — the Gerchberg-Saxton algorithm, the hybrid input-output algorithm, and the difference map algorithm — was carried out, where they were applied to hypothetical CTR spectra corresponding to realistic bunch profiles, as well as to experimentally recorded data. For the HIO and the difference map algorithms, a range of values of the corresponding parameters β and β_{dm} were utilized. A range of approaches for adaptive support selection were also tested. Here we present the results of this survey, and outline the approach that appeared to yield the best performance under a broad range of circumstances.

Selection of algorithm

Preliminary tests were carried out with synthetic longitudinal bunch profile data for Gaussian bunches of rms length 0.5 – 3 µm. The simulated spectral data was truncated so as to only include the spectral components actually measured in the experiment, up to 7 µm, in order to emulate realistic reconstruction from experimental data. The Gerchberg-Saxton algorithm was found to perform reasonably well for a number of iterations, but did have a tendency to stagnate or produce incorrect reconstructions — for example multiple bunches — whenever the support allowed for such ambiguity.

The HIO algorithm with $\beta = 1$ did well at probing a broader range of candidate functions. In conjunction with this, though, it induced rather large changes in the candidate function from one iteration to the next, even when the function appeared to be close to a suitable solution. A compromise could be reached by reducing β, which correspondingly reduced the change in the candidate function at each iteration (cf. Eq. (5.4)). Figure 5.4 illustrates the variation in the final error and the rms length of the reconstructed function after 990 iterations of the HIO algorithm followed by 10 iterations of the Gerchberg-Saxton algorithm, performed on input data corresponding to a single Gaussian of rms length 2 µm, on a support of width 15 µm. Ten reconstructions were performed for each value of β, which was varied between 0.1 and 1.2 in steps of 0.1. For these conditions, the HIO algorithm appeared to perform best around $\beta \approx 0.3$–0.4.

Similarly to Fienup [188], we found that performing a small number of iterations of the Gerchberg-Saxton algorithm after a sequence of HIO iterations helps minimize the error. They were particularly good at eliminating non-zero values of the candidate function outside the support — i.e., forcing it to satisfy the support constraint — without significantly affecting the values inside the support.

In the case presented here the reconstructed function generally has an rms length of somewhat less than the original 2 µm. This is due to the fact that we allowed all spectral

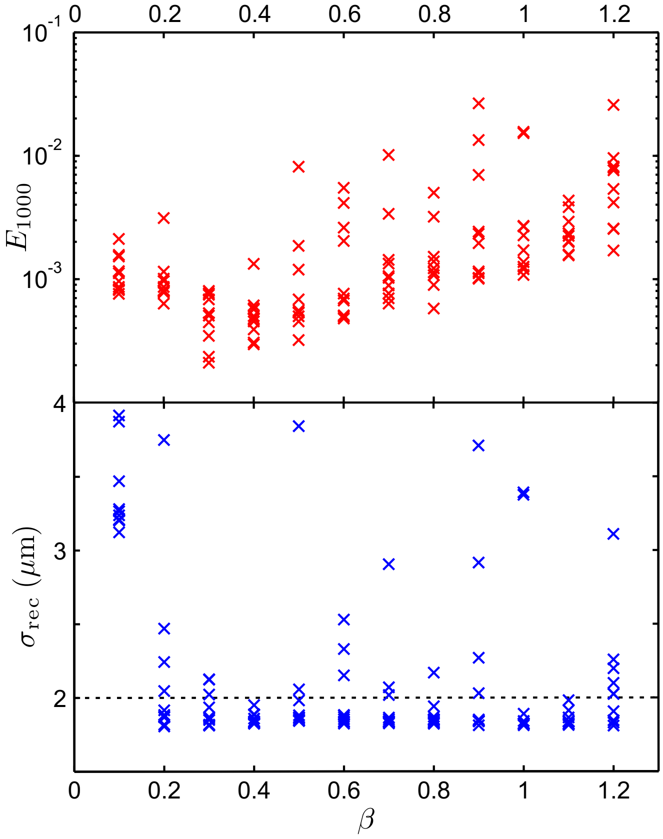

Figure 5.4: Performance of the hybrid input-output algorithm for different values of the parameter β, when reconstructing the spectrum of a Gaussian of 2 µm rms length on a fixed support of length 15 µm. The spectrum is truncated at wavelengths above 7 µm, leading to consistent underestimates of σ_z. Each phase retrieval process consists of 990 iterations of the HIO algorithm followed by 10 iterations of the Gerchberg-Saxton algorithm. Ten reconstructions were carried out for each value of β.

components at wavelengths longer than 7 µm to vary freely, as we would with experimental data. Empirically, for a Gaussian of rms length σ_z to be reconstructed accurately, all spectral components up to a wavelength of ~ $2\pi\sigma_z$ need to be fixed.

Finally, the difference map algorithm suffered from the same issue of large iteration-to-iteration changes that impeded the HIO algorithm's ability to converge on a solution. This is natural, as the two algorithms are equivalent for $\beta_{\rm dm} = \beta = 1$. However, as we showed, in the case of the HIO algorithm this issue could be addressed by reducing β. In the case of the difference map algorithm, on the other hand, changing the parameter $\beta_{\rm dm}$ changed the algorithm's behavior [189], but did not appear to improve its performance with respect to several scenarios involving 1-D data. With these considerations in mind, we select the HIO algorithm as our primary means of phase retrieval, and note the benefit of applying the Gerchberg-Saxton algorithm at the final stages of reconstruction.

Adaptive support selection

A suitable support is crucial for a successful reconstruction. We observe that a Gaussian of rms length 1 µm is reconstructed correctly after 990 iterations of the HIO algorithm followed by 10 iterations of the Gerchberg-Saxton algorithm in virtually every case on a support of length 10 µm, but almost never on a support of length 15 µm, independent of the value of β. Obtaining a support from the spectrum's autocorrelation function is not viable here, since we have no knowledge of the long-wavelength components. In addition, *a priori* knowledge of the rough bunch length is not available. Deducing a support on the basis of physical considerations, i.e., using the theoretical plasma wavelength as a limit on the bunch length, is possible — but unlikely to be restrictive enough for robust reconstruction.

It therefore appears inevitable that the support would have to be calculated adaptively during the course of the reconstruction, as in the Shrinkwrap algorithm. The Shrinkwrap algorithm itself, as described in Ref. [192], tends to lead to reconstructed profiles that are consistently *shorter* than the synthetic profiles used to generate input data. This is a natural outcome of the algorithm's tendency to restrict the size of the support as much as possible. Such a restrictive approach may be valid for reconstructing the images of collections of compact and possibly scattered objects, like those considered in Ref. [192], but it does not appear to be suitable for accurate bunch profile reconstruction in 1-D. We go on to describe potential modifications to this end.

In the original Shrinkwrap algorithm the threshold level applied to the candidate function each time the support was recalculated was fixed at 20 %. In the case of CTR data this level results in a support that consistently cuts off any low-current tails in the electron bunch profile. We find that a good way to address this is by starting at a high threshold level, and subsequently reducing it at each recalculation of the support. This ensures that the function reconstructed during the initial stages of the phase retrieval process is still as compact as possible, however during the latter stages more subtle details are allowed to emerge near the edges. Furthermore, we find it beneficial to initially run the HIO algorithm with a high value of β, to allow probing over a wider range of candidate functions — subsequently, as the reconstruction progresses and the function converges towards a target profile, β is reduced in order to prevent excessive changes occurring.

Parameters

After trialling different combinations of parameters, the following approach appears to yield good results for a range of inputs:

- Initial support of length 20 µm. This can shrink or grow during the process, so the initial value is not crucial.
- Between each recalculation of the support, the HIO algorithm is run for 45 iterations, and the Gerchberg-Saxton algorithm is run for 5 iterations. We will henceforth call this set of 50 iterations an *iteration cycle*.

- The reconstruction process is run for a total of 10,000 iterations, i.e., 200 iteration cycles.
- The initial β for the HIO algorithm is set to 1.0, and decreased by 1.5% after each iteration cycle, bringing it down to a final value of 0.05.
- The initial threshold for support recalculation is set to 25% of the maximum, and decreased by 2.5% after each iteration cycle, bringing it down to a final value of 0.16%.
- The initial blurring width, σ, is set to 3 data points and decreased by 0.5% after each iteration cycle, bringing it down to a final value of 1.1 data points.

These parameters are far from universally applicable — for instance some phase retrievals may require more iterations, and some may reach the target function after much fewer — however they form a good basis set. They can be adjusted further based on observation of the algorithm's performance. The number of data points used in the reconstructions presented here is $2^{13} = 8,192$, however a significantly smaller number can be used without substantively affecting the operation of the phase retrieval process. This is discussed further in Sec. 5.2.4. In the following subsection we present the ability of this algorithm to reconstruct a range of profiles from synthetic data, and outline its limitations.

5.2.3 Reconstruction of profiles from synthetic spectral data

We applied the reconstruction procedure described in the previous subsection to a number of spectra, corresponding to sample bunch shapes. These synthetic spectra were truncated up to a wavelength of 7 μm, i.e., frequencies below 2.7×10^{14} rad s^{-1} were left as free parameters in the reconstruction. Since phase retrieval cannot fix the absolute position or the direction of the reconstructed function, in order to compare the reconstructed profiles they were shifted so that their peaks overlap, and if necessary inverted. Finally, they were rescaled so as to have the same maximum value. Ten reconstructions were carried out for each input spectrum, in order to ascertain the stability of the reconstruction.

For the cases of simple Gaussian functions, we find that the lack of spectral data at long wavelengths impedes the accurate reconstruction of longer profiles, leading to underestimates of their length. Despite the shorter lengths, though, the reconstructed bunch shape

Table 5.1: Phase retrieval for synthetic Gaussian bunch data. The given error margins are standard deviation over the 10 reconstructions (or subsets thereof, in the case of 3 μm).

Original σ_z (μm)	Retrieved σ_z (μm)
0.5	0.500 ± 0.001
1.0	0.978 ± 0.004
1.5	1.414 ± 0.006
2.0	1.736 ± 0.019
3.0	2.066 ± 0.053 (7/10 cases), 4.369 ± 0.270 (3/10 cases)

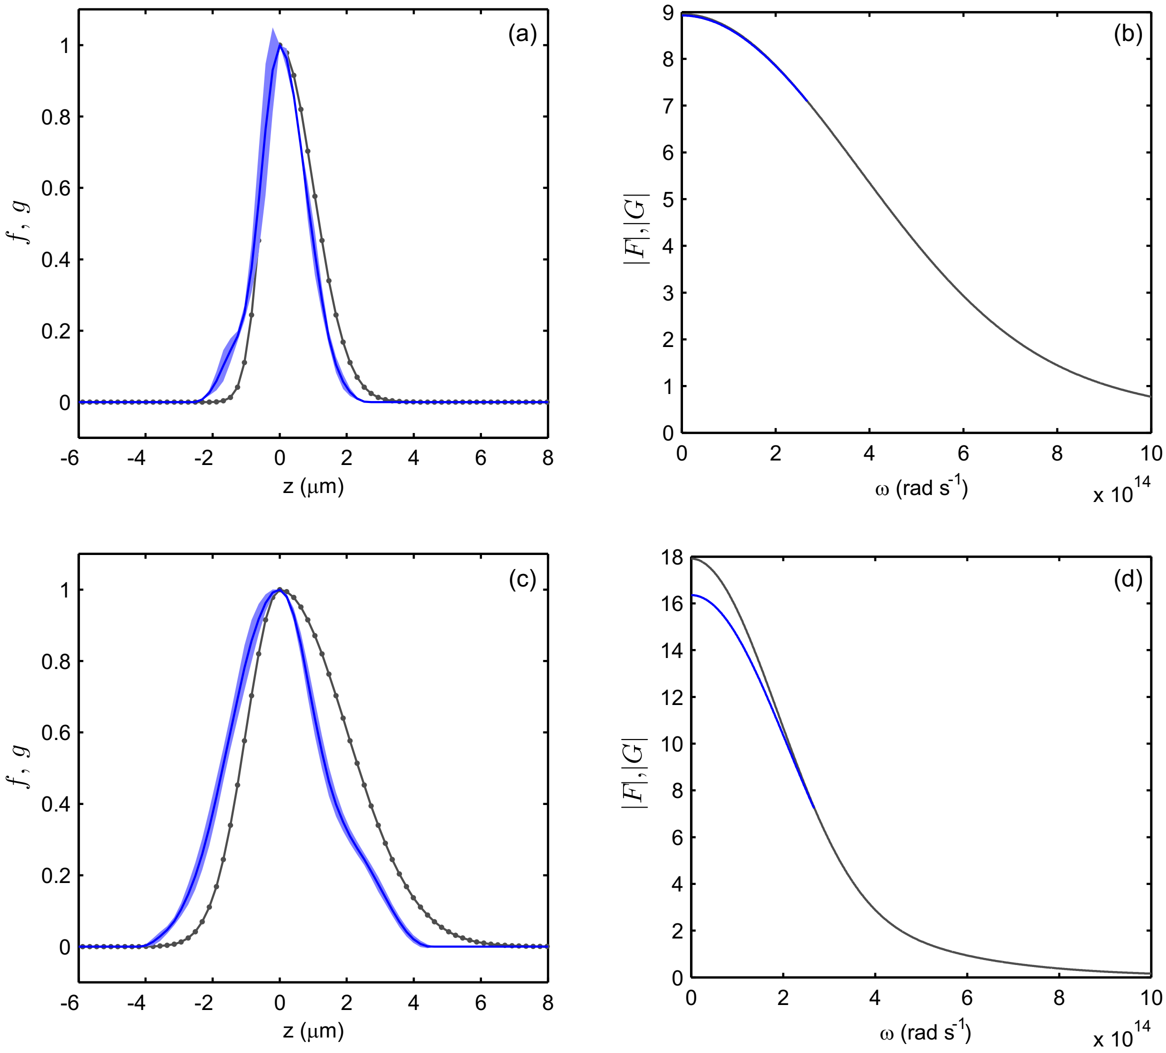

Figure 5.5: Reconstruction of sample asymmetric bunch profiles. Initial profiles are shown as gray lines in (a) and (c), with the data points shown. The mean of the reconstructed profiles — from 10 reconstructions with randomized initial data — are shown as blue lines. The light blue shading represents three-σ deviation from the mean. Profile (a) is a combination of half-Gaussians of rms lengths 0.5 μm (to the left of $z = 0$) and 1 μm (to the right); for profile (c) the lengths are correspondingly 1 μm and 2 μm. The spectral amplitudes corresponding to the original and reconstructed profiles are shown in (b) and (d). The reconstructed spectra (blue) are only shown for $\omega < 2.7 \times 10^{14}$ rad s^{-1}, where the data is withheld from the phase retrieval algorithm.

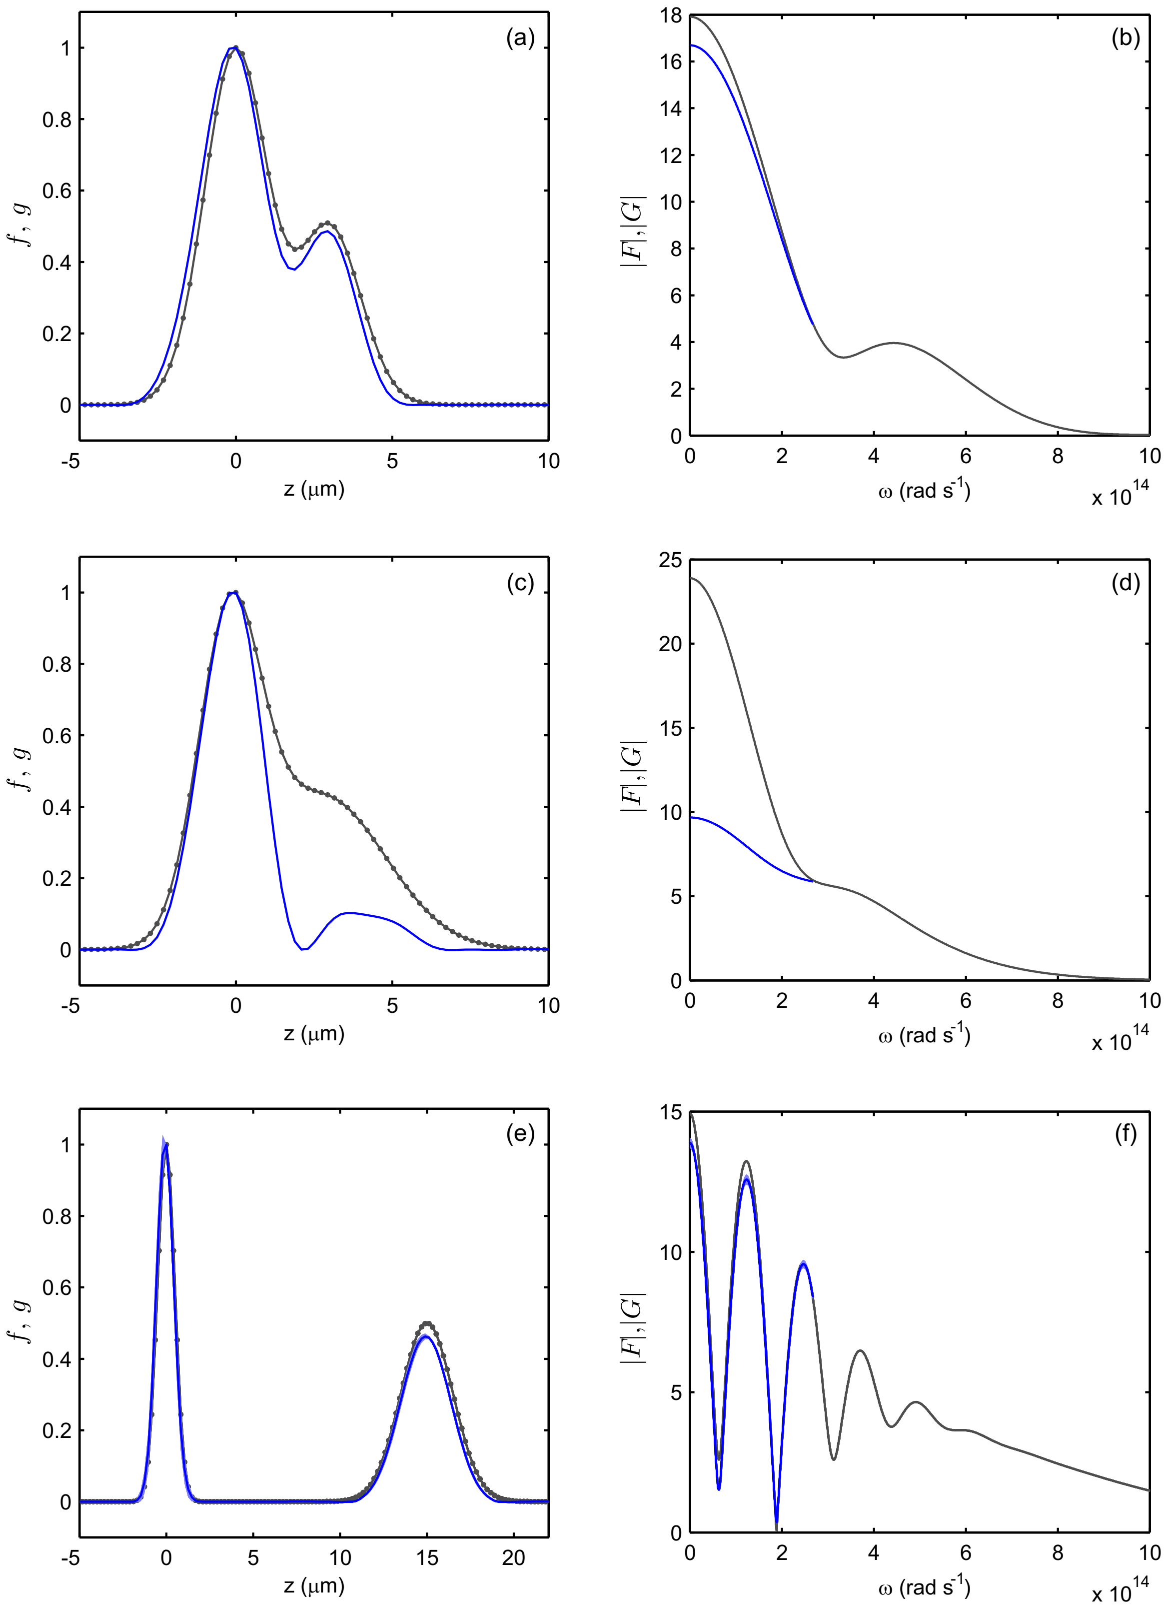

Figure 5.6: Reconstruction of sample bunch profiles consisting of two Gaussians. Initial profiles are shown as gray lines in (a), (c) and (e). The mean of the reconstructed profiles are shown as blue lines, and light blue shading represents three-σ deviation from the mean. For details of the sample profiles see text. The spectral amplitudes corresponding to the original and reconstructed profiles are shown in (b), (d) and (f).

is still approximately Gaussian, occasionally with slight asymmetries. Table 5.1 summarizes the accuracy of reconstruction for several sample cases. We see that for bunches of length $\gtrsim 2\,\mu\mathrm{m}$ the deviation in reconstructed bunch length becomes significant. As the figures in the table indicate, for the case of the 3 µm-rms profile, for some inputs the phase retrieval converges on a profile containing a pair of peaks rather than a single one. This is because for bunches of that length most of the spectral data is at wavelengths longer than 7 µm, and truncating it leads to ambiguity.

Furthermore, we consider reconstruction of asymmetric Gaussian profiles, as illustrated in Figure 5.5. Again, we see that whenever important spectral components of the original profile are left as free parameters during the phase retrieval, the quality of the reconstruction suffers. Finally, Figure 5.6 presents the original and reconstructed profiles of "double-bunch" scenarios, where two Gaussian profiles are either overlapping or separate. In each case, one Gaussian has half the peak height of the other. In profile (a) both have rms length 1 µm, and their centers are separated by 3 µm; in profile (b) the separation is the same, but the smaller Gaussian has an rms length of 2 µm; in profile (c) the rms lengths are 0.5 µm for the larger Gaussian and 1.5 µm for the smaller one, and the separation is 15 µm.

A subset of the above reconstructions were also performed after introducing Gaussian noise to the synthetic data, at levels of 10 % and 20 % rms. For the reconstruction of regular Gaussian profiles, we found the noise to introduce random uncertainties in the reconstructed bunch lengths that are commensurate with, or less than, the applied noise level. Notably, applying such non-systematic high-frequency noise in the Fourier domain results in the anomalous reconstruction of bunch features at long distances from the main bunch in the real-space domain — these can normally be disregarded on the basis of physical considerations. In order for measurement errors to introduce actual variations in the bunch shape — or the appearance of features close to the main bunch — they would need to be systematic, altering the actual shape of the recorded spectrum over a range of frequencies. Such errors may for instance arise from a miscalibration of the detectors' spectral response.

This brief treatment should inform our expectations regarding the algorithm's ability to reconstruct bunch profiles from real-world CTR data. It is notable that in each of the cases considered here the reconstructed profiles exhibit very little to no dependence on the initial randomized candidate function. We observe that this is not as often the case for spectra corresponding to profiles that cannot be reconstructed accurately, such as the 3 µm-rms-length bunch referred to in Table 5.1. In the following subsection we go on to apply the algorithm to a selection of experimentally recorded spectra.

5.2.4 Reconstruction of longitudinal bunch profiles from CTR data

The phase retrieval algorithm described in the previous subsection was applied to each experimental shot where good spectral data was available through the visible, near-infrared, and terahertz, for cases where the 1.7–7 µm grating configuration was used on the THz spectrometer — altogether 816 shots. There were a further 538 shots with good data that were taken using the 5–25 µm grating configuration, however the gap in the spectrum at

2–5 µm precluded reliable application of the phase retrieval algorithm in those cases. Moreover, without an absolute calibration or an overlap region with the calibrated NIR data, it was difficult to determine the absolute level of the THz signal for these long-wavelength measurements.

In this section we present the results of the bunch profile reconstruction from actual CTR data for a small number of representative cases. We use a grid of 2048 points for the reconstruction (hereafter the "reconstruction grid"), which in the Fourier domain spans frequencies from $\omega_{\text{min}} = 0$ up to $\omega_{\text{max}} = 2\pi c/(420\,\text{nm}) = 4.48 \times 10^{15}\,\text{rad s}^{-1}$, where the latter corresponds to the shortest wavelength recorded on the visible spectrometer. This highest frequency also determines the data point spacing in real space as $\Delta z = 210\,\text{nm}$, since by Nyquist's theorem $\omega_{\text{max}} = 2\pi c/(2\Delta z)$. The number of data points then sets the total span of the reconstruction window to approx. 430 µm. Since the total length of the electron bunch — including secondary features — is unlikely to exceed 50 µm, the reconstruction is largely insensitive to the concrete numbers of data points used, as long as this number is greater than ~ 500. In particular, even though the reconstructions of synthetic data in the previous subsection were carried out on grids of 8192 points, the same results are obtained for grids of 2048 points.

Since the data from the visible and NIR spectrometers has a resolution that is higher than (or comparable to) the resolution of the reconstruction grid in Fourier space, we use simple linear interpolation to populate the initial candidate spectrum function, G_0, with the amplitude data (i.e., the square roots of the measured intensities). On the other hand, the terahertz data has a relatively low resolution, with much fewer experimental data points than points on the reconstruction grid. In this case, we populate G_0 by taking each of the terahertz data points and assigning its amplitude to the nearest corresponding point on the reconstruction grid. The rest of the points are left unassigned, and their amplitudes are free to vary during the phase retrieval process.

Having populated G_0 for a particular shot, the profile reconstruction process is carried out 10 times, using different seeds. This allows the stability of the reconstruction to be evaluated, and allows us to pick the most representative profile of the 10. In order to do this, we first align the profiles so that their peaks are centered on the grid; if necessary their direction is then reversed. These steps are carried out since the reconstruction process cannot determine the absolute offset of the profiles with respect to the grid origin, or their left-right orientation (i.e., direction in time). Namely, the offset and direction are solely determined by the spectral phases and have no relation to the spectral amplitudes, leading to a random offset and orientation for each reconstruction.

We then carry out a process of iterative elimination on the set of reconstructed profiles. Starting with all 10 profiles, at each iteration the profile that is most different from the others — in a least squares sense — is removed from the set. This elimination is repeated on the remaining set until only three profiles are left. Of those, the profile that is most similar to the other two is selected at the most representative. This process appears to work satisfactorily for most shots. Even in cases where the reconstruction produces diverging results for some initial seeds, eliminating the outliers appears to ultimately lead to the selection of a representative profile.

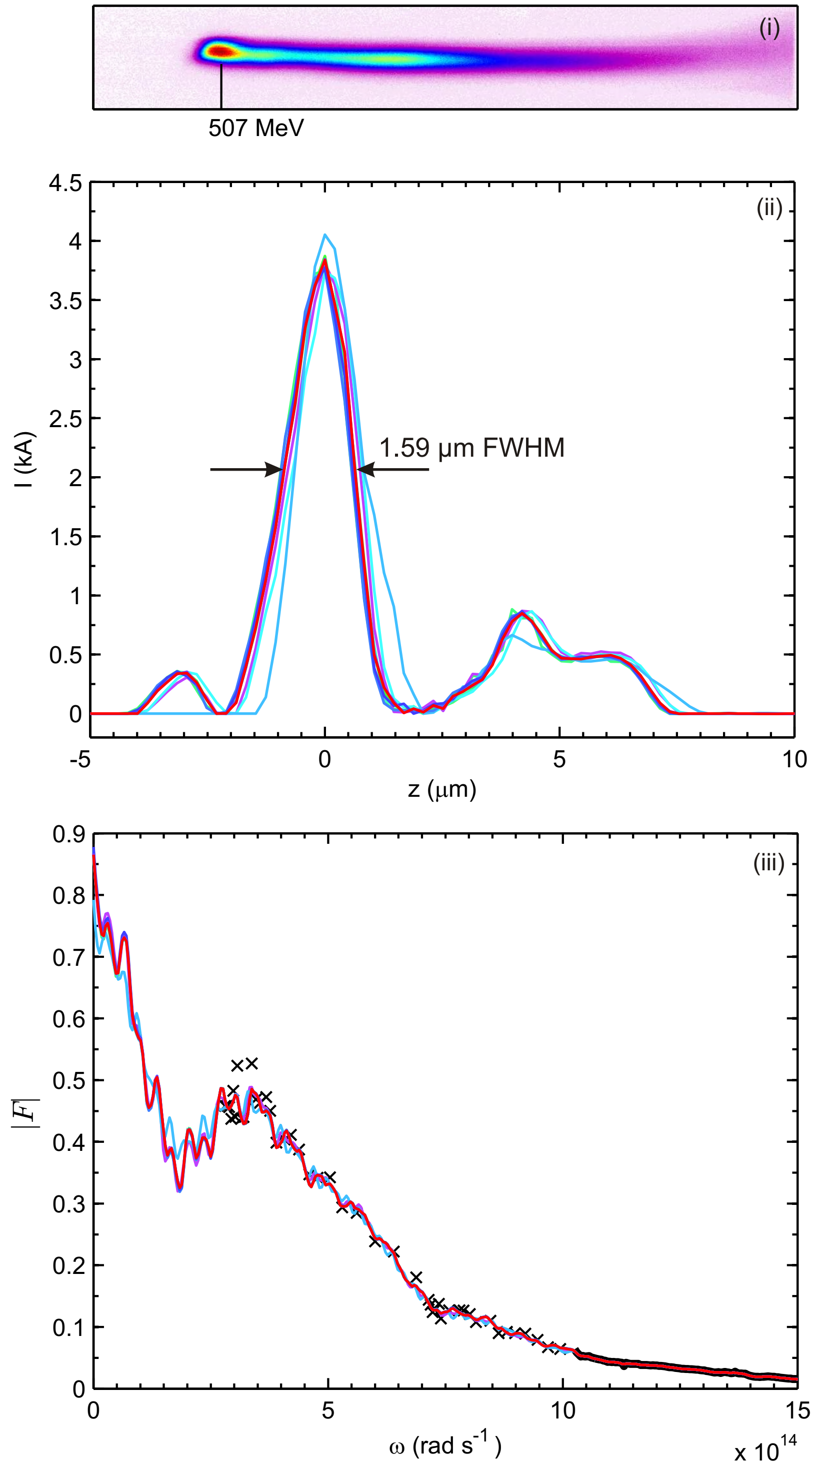
(a) Shot 1930:
$P = 83\,\mathrm{mbar}$ ($n_e = 4.0 \times 10^{18}\,\mathrm{cm}^{-3}$); $L_{\mathrm{cell}} = 6\,\mathrm{mm}$;
$E_{\mathrm{peak}} = 507 \pm 92\,\mathrm{MeV}$ (FWHM); $Q_{\mathrm{peak}} = 8.7\,\mathrm{pC}$; $Q_{\mathrm{total}} = 34.6\,\mathrm{pC}$.

Figure 5.7: Reconstruction of longitudinal bunch profiles from experimentally recorded CTR spectra: (i) electron spectrometer lanex image; (ii) reconstructed profiles, where the most representative one is colored red; (iii) measured spectral data (black line and black crosses), as well as spectra corresponding to the reconstructed profiles (colored lines). A Gaussian smoothing filter with $\sigma = 2.2 \times 10^{12}\,\mathrm{rad\,s^{-1}}$ is applied to the latter.

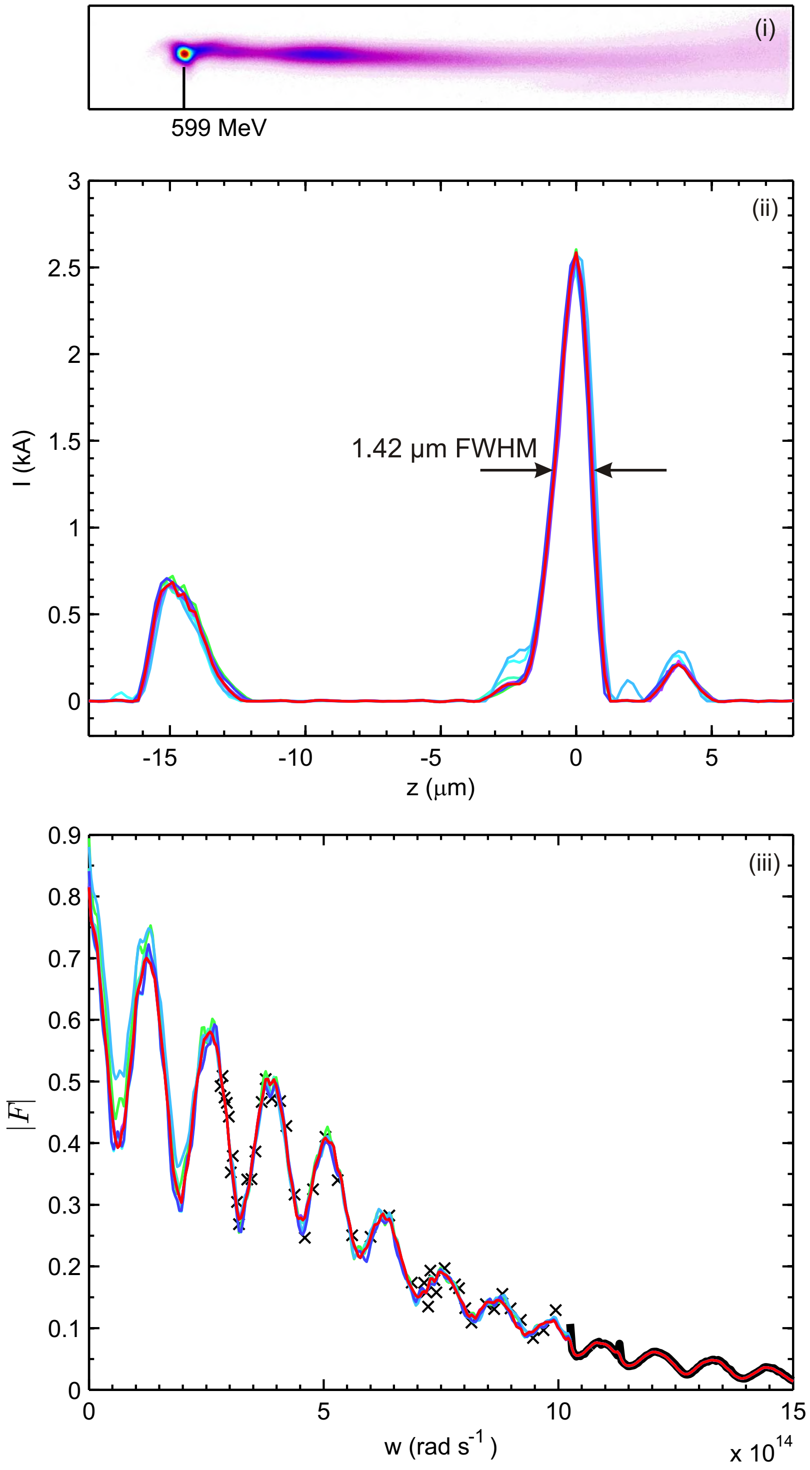

(b) Shot 2168:
$P = 82\,\mathrm{mbar}$ ($n_\mathrm{e} = 4.0 \times 10^{18}\,\mathrm{cm}^{-3}$); $L_\mathrm{cell} = 14\,\mathrm{mm}$;
$E_\mathrm{peak} = 599 \pm 47\,\mathrm{MeV}$; $Q_\mathrm{peak} = 3.8\,\mathrm{pC}$; $Q_\mathrm{total} = 22.5\,\mathrm{pC}$.

Figure 5.7: Reconstruction of longitudinal bunch profiles from experimentally recorded CTR spectra: (i) electron spectrometer lanex image; (ii) reconstructed profiles, where the most representative one is colored red; (iii) measured spectral data (black line and black crosses), as well as spectra corresponding to the reconstructed profiles (colored lines). A Gaussian smoothing filter with $\sigma = 2.2 \times 10^{12}\,\mathrm{rad\,s^{-1}}$ is applied to the latter.

(c) Shot 2147:
$P = 82\,\mathrm{mbar}$ ($n_\mathrm{e} = 4.0 \times 10^{18}\,\mathrm{cm}^{-3}$); $L_\mathrm{cell} = 13\,\mathrm{mm}$;
$E_\mathrm{peak} = 585 \pm 48\,\mathrm{MeV}$; $Q_\mathrm{peak} = 2.8\,\mathrm{pC}$; $Q_\mathrm{total} = 21.0\,\mathrm{pC}$.

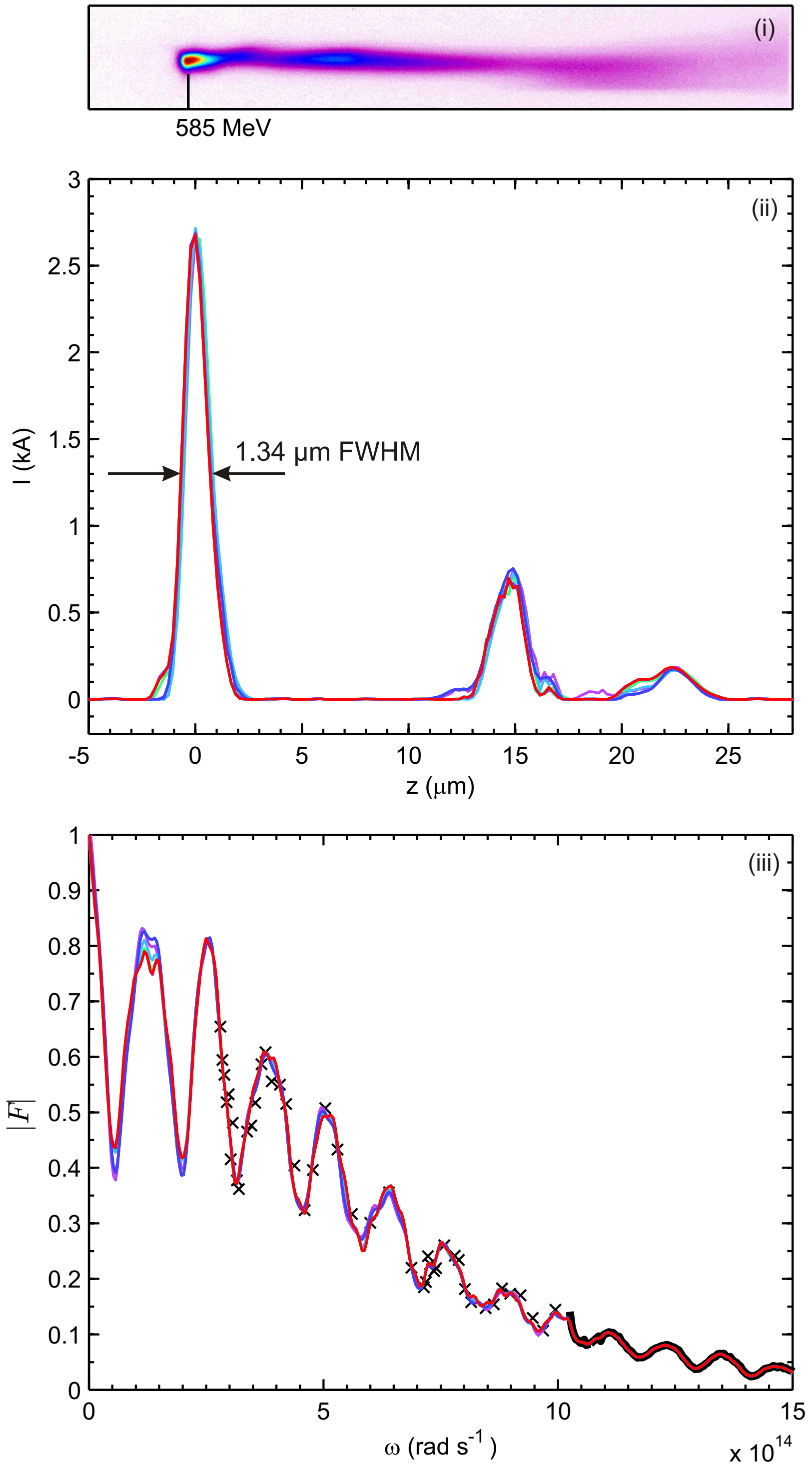

Figure 5.7: Reconstruction of longitudinal bunch profiles from experimentally recorded CTR spectra: (i) electron spectrometer lanex image; (ii) reconstructed profiles, where the most representative one is colored red; (iii) measured spectral data (black line and black crosses), as well as spectra corresponding to the reconstructed profiles (colored lines). A Gaussian smoothing filter with $\sigma = 2.2 \times 10^{12}\,\mathrm{rad\,s}^{-1}$ is applied to the latter.

Figure 5.7 shows the reconstructed profiles for three representative shots (panels (ii) of each subfigure), with a summary of acceleration and electron bunch parameters given for each at the top. The lanex images (panels (i)), and the experimental and reconstructed CTR spectra (panels (iii)) are given for reference. The figure illustrates the general shape of the reconstructed profiles — occasionally with secondary or even tertiary bunches, the interference of which causes characteristic oscillations in the CTR spectrum.

The small pedestal to the left of the main bunch in Fig. 5.7a(ii) is also present in many of the other reconstructed profiles, and occurs due to the lower-than-expected values recorded by the pyroelectric detectors at the longest wavelengths (lowest frequencies), which lead to the large-scale modulation of the reconstructed spectra seen in Fig. 5.7a(iii). Without a proper calibration of the THz spectrometer it is difficult to say whether these spectral modulations — and the corresponding pedestals in the bunch profiles — are a real feature, or artifacts due to a miscalibration in the long-wavelength region.

While when dealing with synthetic data in Section 5.2.3 we consistently used a root-mean-square measure of the bunch lengths, here it is necessary to switch to a full-width at half-maximum measure. This must be done since the presence of noise in the reconstructed profile at large distances from the main bunch — and indeed any pedestals or secondary bunches around it — would interfere with an rms measurements of the overall bunch length. A FWHM measure is more appropriate when dealing with non-ideal experimental data. In the particular case of a Gaussian, the FWHM is a factor of $2\sqrt{2\log 2} \approx 2.35$ greater than the rms width. Therefore if we approximate the bunch profiles reconstructed from experimental data as Gaussians, their corresponding rms widths would be well within the range for which we observe reliable reconstruction in Sec. 5.2.3.

Since both the electron and the CTR spectra are absolutely calibrated[5], we can check that the total amount of observed radiation corresponds to the amount of TR expected from a bunch with the measured energy spectrum and the reconstructed longitudinal profile. This gives us a tentative verification the validity of the construction. To do this, we can take as a starting point the formula for total energy emitted by a single electron, $\mathrm{d}U_1/\mathrm{d}\omega = e^2/(2\pi^2\epsilon_0 c)\log\gamma$, which is the integral of Eq. (4.3) over all angles in the limit $\gamma \gg 1$ (after Ref. [194]), where γ is the Lorentz factor of the electron. Taking into account the total number of electrons, N_e, and the frequency-dependent longitudinal coherence factor, $\tilde{\rho}_\parallel(\omega)$ introduced in Sec. 4.1.2, the total coherent spectrum is given by

$$\frac{\mathrm{d}U_\mathrm{CTR}}{\mathrm{d}\omega} = N_\mathrm{e}^2 \left|\tilde{\rho}_\parallel(\omega)\right|^2 \frac{e^2}{2\pi^2\epsilon_0 c}\langle\log\gamma\rangle, \tag{5.6}$$

where $\langle\log\gamma\rangle = \frac{\int \mathrm{d}E\, \log\gamma(E)(\mathrm{d}Q/\mathrm{d}E)}{\int \mathrm{d}E\,(\mathrm{d}Q/\mathrm{d}E)}$ is the weighted average of $\log\gamma$ over the electron energy spectrum, $\mathrm{d}Q/\mathrm{d}E$. Due to the small variation in $\log\gamma$ ($\approx$ 6–7) over the range of observed electron energies, the use of this weighted average to represent the scaling with γ — instead of separately considering the emission from parts of the electron bunch with different energies — should not lead to a significant discrepancy.

If the reconstructed bunch profile accounts perfectly for the absolute level of observed

---

[5] To within the limits of the cross-calibration of the NIR and THz spectrometers, and of the theoretical relative calibration of the THz spectrometer.

emission from the number of electrons observed on the spectrometer, $N_\mathrm{e} = \int \mathrm{d}E\,(\mathrm{d}Q/\mathrm{d}E)/e$, the corresponding Fourier transform of the profile, or longitudinal coherence factor, $F \equiv \tilde{\rho}_\parallel(\omega)$, should be unity[6] at frequency $\omega = 0$. Within a small error margin, we observe this to be true in most cases, and certainly for those shown in Fig. 5.7.

Note that there is a certain possibility that not all electrons that emit transition radiation are observed on the electron spectrometer. Additionally, the absolute calibration of the THz detectors is only obtained by cross-calibration with the NIR spectrometer at a small number of data points around 2 µm, and may be unreliable. Therefore, the above verification of reconstruction validity should only be used as a rough guide.

Finally, it is worth noting that phase retrieval is far from the most obvious means of determining the bunch length from CTR data — in fact, to our knowledge this is the first time it has applied to this problem. A more rudimentary method of establishing the bunch length involves fitting a Gaussian spectrum (corresponding to a bunch with a Gaussian longitudinal profile) to the measured data. Applying this method to our data and comparing the FWHM lengths obtained by Gaussian fitting to those obtained by phase retrieval, we see that for each shot the length obtained by phase retrieval is almost universally in the range 1.1 ± 0.1 times the length obtained by Gaussian fitting. Such a small discrepancy lends further credibility to the bunch profiles obtained by phase retrieval. The existence of the discrepancy itself highlights the fact that the assumption of a Gaussian profile is likely to limit the accuracy of bunch length estimation by simple fitting, by e.g. not accounting for second or third bunches, or asymmetry.

Based on the reconstructed profiles for the full dataset, we go on to consider multi-shot statistics and evaluate several aspects of the acceleration process.

5.3 Variation of bunch length with acceleration parameters

This single-shot non-destructive diagnostic of electron bunch length can be combined with information from further electron bunch diagnostics in order to carry out a more thorough analysis of the acceleration dynamics. In this section we present a summary of results from runs taken at two different pressures, where in each the gas cell length was varied over a range of values. Setting the gas cell length to a particular value allows us to, in a sense, obtain a snapshot of the bunch at a particular stage of the acceleration process. Combining data obtained at a series of gas cell lengths, then, allows us to build a picture of the evolution of the bunch along the full acceleration process, together with either its dephasing, or the defocusing or depletion of the driving laser.

The two nominal pressures used here are 80 and 170 mbar, corresponding respectively to plasma densities of $n_\mathrm{e} = 3.8 \times 10^{18}\,\mathrm{cm}^{-3}$ and $8.2 \times 10^{18}\,\mathrm{cm}^{-3}$, and to plasma wavelengths of $\lambda_\mathrm{p} = 17.0\,\mu\mathrm{m}$ and $11.7\,\mu\mathrm{m}$. We begin our discussion by presenting statistics for the dependence on gas cell length of the central energy of the main electron bunch (i.e., the high-

---

[6] The zero-frequency Fourier component of a function's transform is the integral of the function in real space.

Table 5.2: Overview of acceleration parameters.

80 mbar				**170 mbar**			
		n_e	$3.8 \times 10^{18}\,\mathrm{cm}^{-3}$			n_e	$8.2 \times 10^{18}\,\mathrm{cm}^{-3}$
		λ_p	$17.0\,\mu\mathrm{m}$			λ_p	$11.7\,\mu\mathrm{m}$
		P_c	7.7 TW			P_c	3.6 TW
		P_c^{diff}	37.6 TW			P_c^{diff}	12.3 TW
		P_0	54.0 TW			P_0	54.0 TW
		a_0	3.8			a_0	4.9
		w_{sf}	$10.6\,\mu\mathrm{m}$			w_{sf}	$7.3\,\mu\mathrm{m}$
1-D nonlin. [48]		bubble [56]		1-D nonlin. [48]		bubble [56]	
$\lambda_{p,n}$	$25.3\,\mu\mathrm{m}$	$\lambda_{p,b}$	$21.6\,\mu\mathrm{m}$	$\lambda_{p,n}$	$18.6\,\mu\mathrm{m}$	$\lambda_{p,b}$	$16.5\,\mu\mathrm{m}$
$L_{d,n}$	13.8 mm	$L_{d,b}$	3.3 mm	$L_{d,n}$	5.5 mm	$L_{d,b}$	1.2 mm
$L_{pd,n}$	13.8 mm	$L_{pd,b}$	3.8 mm	$L_{pd,n}$	5.5 mm	$L_{pd,b}$	1.8 mm

energy peak on the electron spectrometer), and of the reconstructed bunch length (FWHM) of the main bunch. These are shown in Figures 5.8a and 5.8b, respectively.

Looking at the bunch energies alone, Fig. 5.8a, it is clear that for each pressure there is an optimum acceleration length, for which the highest energy is achieved. At 170 mbar this appears to be 5 mm, while for 80 mbar is it somewhat longer, around 9 mm. The notion of an optimum acceleration length also broadly corresponds to the dependence of bunch length on acceleration length shown in Fig. 5.8b, though there are discrepancies between the two figures that we go on to address shortly.

While at 170 mbar the electron energy drops sharply for $L_{\mathrm{cell}} > 5\,\mathrm{mm}$, at 80 mbar it only tails off slightly past $L_{\mathrm{cell}} > 9\,\mathrm{mm}$. This suggests the mechanisms that limit acceleration are different in the two cases. We investigate this in the context of the acceleration parameters presented in Table 5.2, calculated based on the equations in Section 2.1.1. There is some ambiguity regarding whether we are operating in the bubble regime or whether the mode of operation is closer to the 1-D (broad-pulse) nonlinear regime. In some PIC simulations it has been observed that self-injection at early stages causes beam-loading of the wake, which leads to a termination of self-focusing before the matched self-focused spot size, w_{sf}, is actually reached [195]. Due to this ambiguity, Table 5.2 presents the parameters of both regimes.

We see that in the 1-D regime at 80 mbar the dephasing length is $L_{d,n} = 13.8\,\mathrm{mm}$, while at 170 mbar it is $L_{d,n} = 5.5\,\mathrm{mm}$. It therefore appears that at 170 mbar dephasing is the likely cause of the decrease in energy at $L_{\mathrm{cell}} > 5\,\mathrm{mm}$, while it appears that another process stops the acceleration process at lengths $\gtrsim 9\,\mathrm{mm}$ in the case of 80 mbar. This may either be complete depletion of the laser pulse, or its diffraction due to lack of self-guiding after its partial depletion. Namely, at low pressures the critical laser power required for self-guiding is greater, and for our parameters at 80 mbar it is rather close to the initial peak laser power (37.6 TW and 55 TW, respectively). Thus even partial depletion of the laser pulse may lead to diffraction and termination of the acceleration process. We will refer to these processes collectively as “laser depletion”.

These considerations are in line with the dependence of bunch energy on acceleration length at high pressure (170 mbar), which suggests the onset of dephasing — and thus a decrease in energy — while at low pressure (80 mbar) the driving laser diffracts and the electron bunch continues to coast through the plasma, only losing a small amount of energy. This explanation is corroborated by the evolution of bunch charge with gas cell length shown in Figure 5.9. For 80 mbar, at lengths above $L_{\text{cell}} \sim 8\,\text{mm}$ there is a gradual decrease in bunch charge, even though the bunch energy remains the same. This may be due to attrition of the bunch as it propagates through the plasma when there is no longer a laser driving a bubble. On the other hand, at 170 mbar (Fig. 5.9b) the abrupt drop in bunch charge for $L_{\text{cell}} \gtrsim 6\,\text{mm}$ corresponds very closely to the drop in electron energy at those gas cell lengths — behavior that is in line with the electron bunch dephasing within the bubble.

In the context of the above discussion, the evolution of bunch length with L_{cell} can be explained more adequately. At 80 mbar for $2\,\text{mm} < L_{\text{cell}} < 8\,\text{mm}$ the measured bunch length decreases with increasing L. This is consistent with the rear of the bunch experiences a stronger accelerating field than the front, as is expected to be the case. Beyond $L_{\text{cell}} \sim 8\,\text{mm}$, the point at which the driving laser appears to no longer drive a wake, bunch shortening continues at a slower pace, and is more likely to be a result of attrition of the front of the bunch as it propagates through the plasma, which is also in line with the decreasing charge seen in Fig. 5.9a.

At 170 mbar, on the other hand, the acceleration dynamics appear to be more complex with respect to bunch length. While the initial bunch shortening as the electron energy increases with L_{cell} is again clearly present, the bunch remains relatively short even after dephasing, and only appears to become longer for very long gas cell lengths. It is possible that the range of effects that occur during dephasing — such as loss of bunch charge — collectively ensure the short bunch length is maintained.

In any case, even from the small data sample presented here it is clear that an electron bunch length diagnostic can be highly beneficial in analyzing and understanding the acceleration dynamics. Ultimately, bunch length measurements would be crucial for the selection of a suitable working point for a laser-wakefield accelerator operating in a production environment, such as a driver for a compact light source.

5.4 Acceleration of multiple electron bunches

As Figs. 5.7b and 5.7c illustrate, some of the observed transition radiation spectra exhibit high-contrast modulations in frequency. Such modulations are a signature of interference between transition radiation emitted from multiple bunches. The separation of these bunches in z can be inferred from the modulation frequency, and the relative sizes of the bunches can be determined from the contrast. In our case, this is all done automatically as part of the phase retrieval process. As can be seen in Figs. 5.7b and 5.7c, the process also appears to successfully fill in the missing spectral data between the THz data points

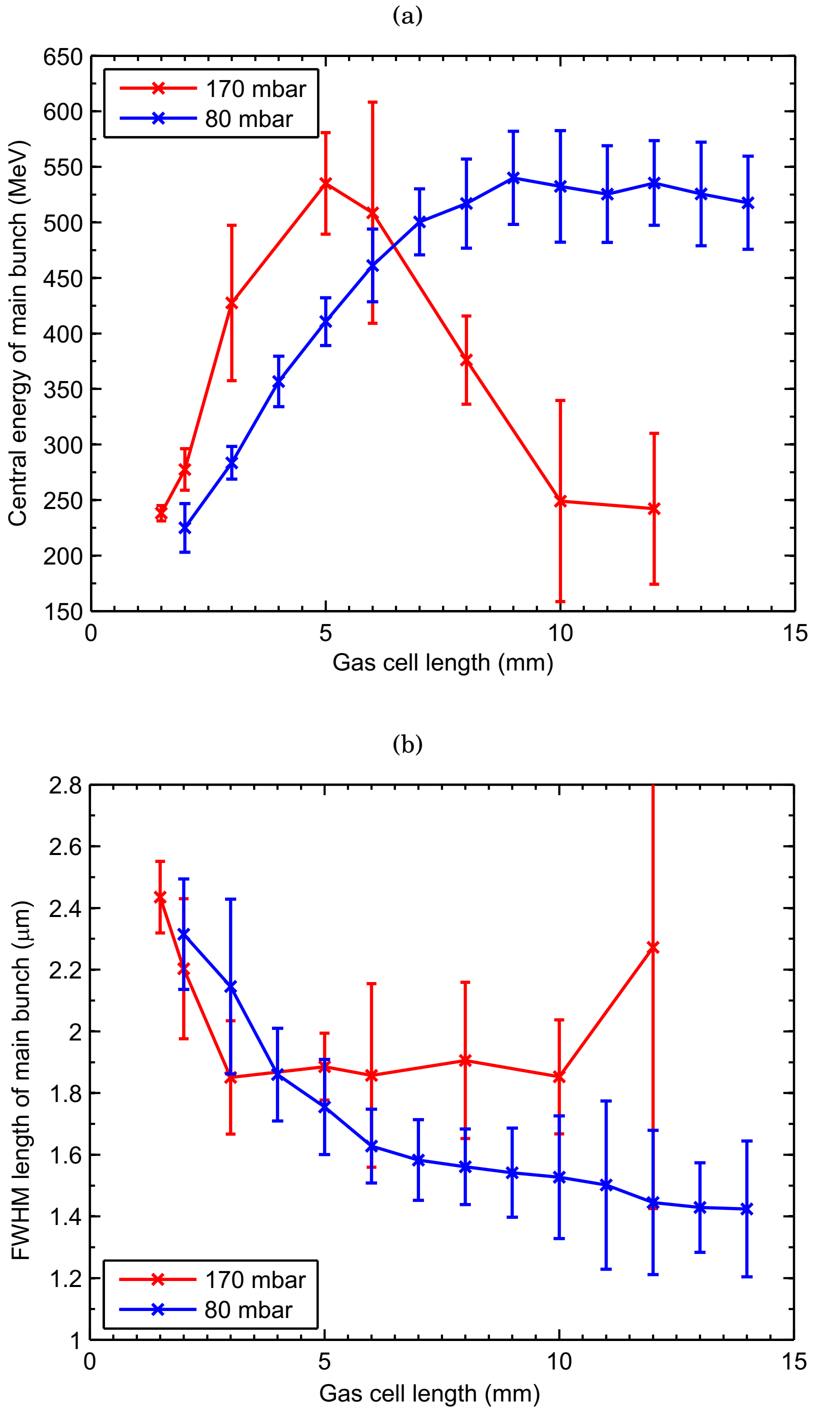

Figure 5.8: Dependence of (a) electron bunch energy and (b) full-width at half maximum of the reconstructed bunch on gas cell length. Statistics are over ~ 30 shots for 80 mbar and ~ 20 shots for 170 mbar, where the data points represent mean values and the error bars are standard deviation.

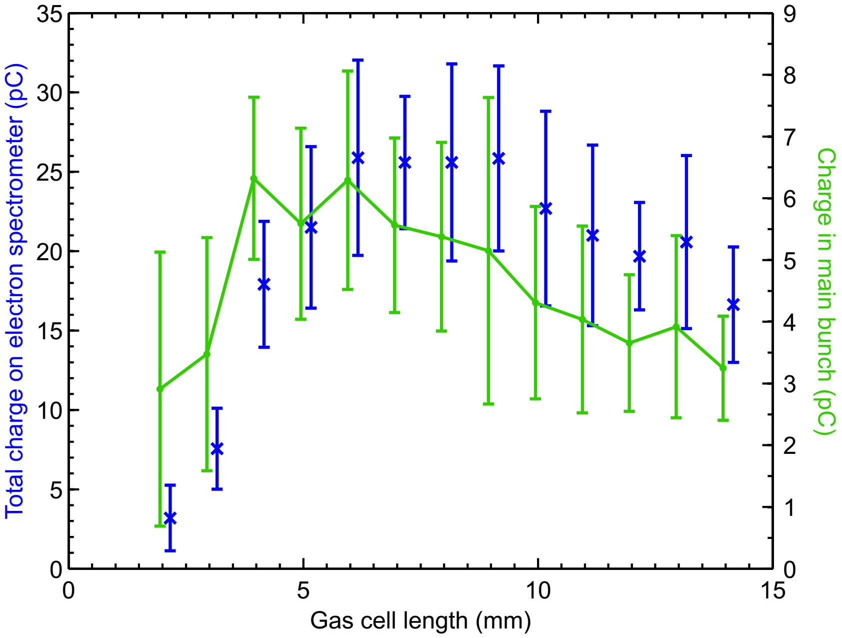
(a) 80 mbar

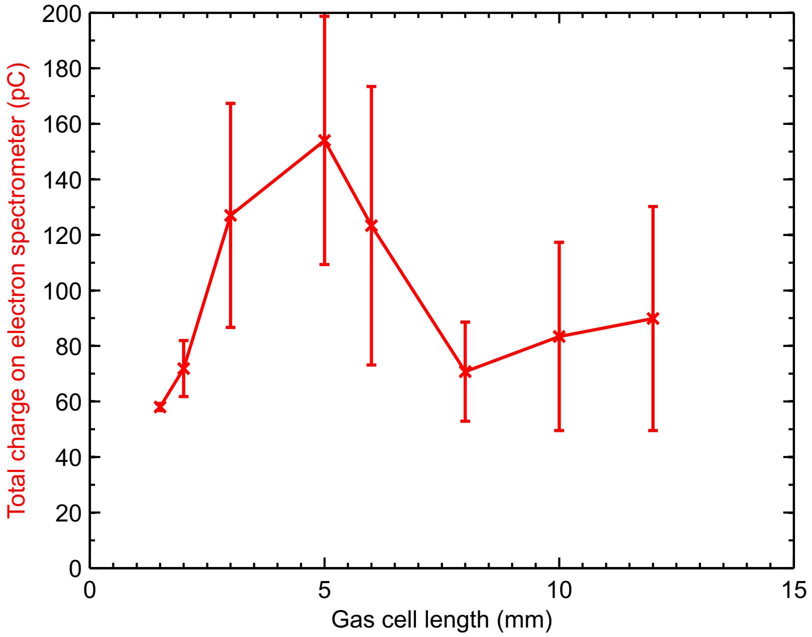
(b) 170 mbar

Figure 5.9: Dependence of accelerated bunch charge on gas cell length at (a) 80 mbar and (b) 170 mbar. The total represents all the charge observed above 180 MeV, the spectrometer cut-off energy. In the 80 mbar case the bunch charge in the quasi-monoenergetic peak is also shown in green, defined as the sum of charge within the FWHM (in terms of dQ/dE) of that peak, and at higher energies. This statistic was omitted at 170 mbar as in many of the spectra the peaks were less well-defined. Statistics are over ~ 30 shots for 80 mbar and ~ 20 shots for 170 mbar, where the data points represent mean values and the error bars are standard deviation.

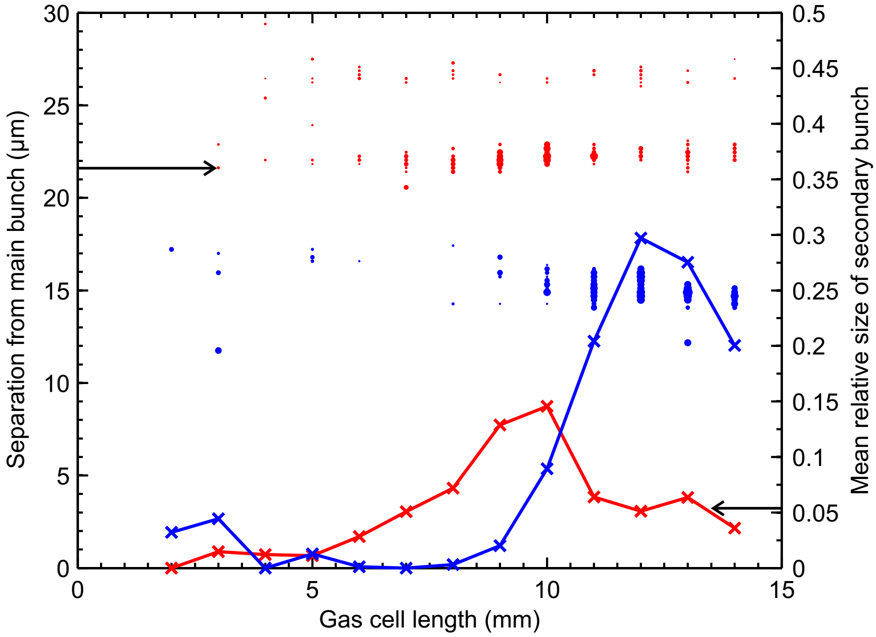

Figure 5.10: The observation of secondary bunches during a gas cell length scan at 80 mbar. Dot positions represent the distance of secondary (or tertiary) bunches from the main one, with dot sizes denoting their peak current as a fraction of the peak current of the main bunch — here also called "relative size". Blue dots represent bunches separated by $\Delta z < 19\,\mu\text{m}$ from the main bunch, while red dots represent separations larger than 19 µm. The blue and red curves show the mean of the relative sizes of bunches observed at $\Delta z < 19\,\mu\text{m}$ and $\Delta z > 19\,\mu\text{m}$, respectivaly. These means are computed over all shots recorded at each gas cell length, including those where no secondary bunches are observed, for which a zero relative size is taken.

(marked with crosses). In the case of 5.7c, it has also picked up simultaneous modulations at two different frequencies, resulting in the reconstruction of three bunches.

In order to understand the conditions required for the acceleration of additional electron bunches — and the properties of those bunches — we carried out an analysis of all secondary and tertiary bunches observed during the 80-mbar run that was presented in more detail in Section 5.3. The key statistics are shown in Figure 5.10. One of the most striking features exhibited by the data is the apparent discreteness of the observed separation between the main bunch and additional bunches. There appear to be at least two clear regimes, with some bunches separated by $\Delta z \sim 15\,\mu\text{m}$ from the main bunch (blue dots in Fig. 5.10), and some by $\Delta z \sim 22\,\mu\text{m}$ (red dots), but practically no bunches seen at intermediate separations. A third mode appears to occur at $\Delta z \sim 26\,\mu\text{m}$, however the bunches seen there are much smaller. We have verified that this "modularity" of bunch separations is not just an artifact of the reconstruction process by carrying out Fourier analyses of the modulations on the NIR data alone.

After observing this modularity, in our analysis we recorded the separation and size[7] of bunches observed at $\Delta z < 19\,\mu\text{m}$ as well as at $\Delta z > 19\,\mu\text{m}$ for each shot, thus being able to account for multiple additional bunches. A prime example where this is required is the

---

[7] Represented as peak current of the secondary bunch as a fraction of the peak current of the main bunch.

profile shown in Figure 5.7c, where additional bunches occur at *both* $\Delta z = 15\,\mu\text{m}$ and $\Delta z = 22.5\,\mu\text{m}$.

From this information we can determine the probability of secondary bunches appearing in either the "small-separation" or "large-separation" regimes. In Figure 5.10 we have represented this probability as the mean relative size of bunches appearing in the two regimes. For the large-separation bunches we see this rise with increasing gas cell length, and peak at 13–15% at $L_\text{cell} = 9$–10 mm, then drop off sharply. On the other hand, small-separation bunches are almost non-existent for $L_\text{cell} \leq 9$, but their mean relative size rises sharply for longer gas cell lengths and peaks at 30% for $L_\text{cell} = 12\,\text{mm}$. Indeed, bunches at $\Delta z \sim 15\,\mu\text{m}$ with peak currents of at least 10% of that of the main bunch are observed in ~ 90% of shots for all gas cell lengths above 11 mm.

The facts that secondary bunches occur only at discrete separations, and that their occurrence has a clear dependence on acceleration length, suggest the presence of two different acceleration mechanisms for bunches at different separations. At this point it is helpful to refer again to Table 5.2. At 80 mbar the non-linear plasma wavelength in the bubble regime is $\lambda_\text{p,b} \approx 21.6\,\mu\text{m}$, and in the 1-D regime it is only slightly longer, $\lambda_\text{p,n} = 25.3\,\mu\text{m}$. Clearly, these values are very close to the separation at which we see the more distant secondary bunches at acceleration distances $L_\text{cell} = 6$–11 mm, suggesting that they are simply accelerated in the second bucket behind the driving laser. This second-bucket acceleration is readily seen in simulations as well as experiments, e.g. in Figs. 2.3b–c, or in Ref. [36].

At $L_\text{cell} \approx 11\,\text{mm}$ a sort of transition occurs, with far fewer (and smaller) bunches seen to be accelerated in the second bucket, while bunches separated from the main one by approx. 15 μm become both common and relatively large. Our interpretation is that these "small-separation" bunches are accelerated in a bubble driven by the main bunch itself, after the driving laser has depleted or diffracted — indeed, the point during the acceleration process at which these bunches appear corresponds very closely to the laser depletion point established in Section 5.3.

Beam-driven plasma acceleration is of course a well-known phenomenon, both theoretically [23] and experimentally [196]. Recently, Hidding *et al.* [197] proposed the idea of using a "hybrid" laser-plasma wakefield accelerator consisting of two stages: in the first stage a laser pulse drives a wake and accelerates two bunches — one in the first and one in the second bucket — then in the second stage the leading bunch drives a wake through a plasma of lower density (and thus longer period), so that the trailing (witness) bunch is trapped at the back of the bubble driven by the first bunch, and is accelerated to double its initial energy. Further simulations were carried out by Pae *et al.* [198], where a similar process is demonstrated to occur in a single (uniform density) plasma column. In that case the witness bunch that is ultimately injected in the beam-driven wake does not exist independently prior to the creation of that wake — as is the case in the scheme by Hidding *et al.* — but appears to be injected from the remnants of the laser-driven wake upon its breakup after laser depletion. The present experimental data strongly indicates that a similar scenario arises here.

Further verification of the above hypothesis regarding the acceleration dynamics will be sought through 3-D particle-in-cell simulations of the exact conditions under which this

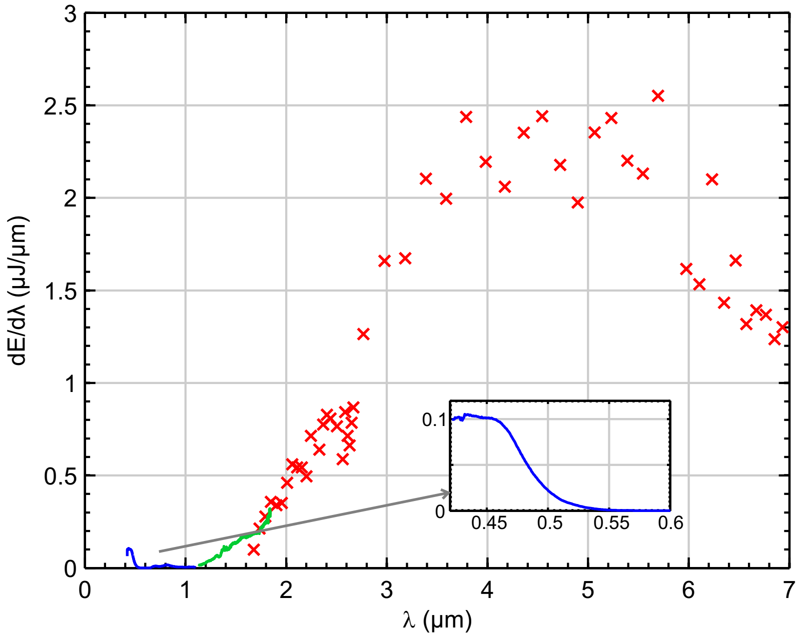

Figure 5.11: Sample CTR spectrum showing data from the visible spectrometer (blue line), NIR spectrometer (green line) and THz spectrometer (red crosses). The "bump" in the spectrum at around 450 nm is clearly visible, and is shown magnified in the inset.

experiment was performed. A clear demonstration that a laser-accelerated electron bunch can itself drive a plasma accelerator in which a witness bunch is accelerated would certainly open up new research areas within the field. Moreover, reliable self-injection of a witness bunch — as we appear to be observing — would remove the requirement for an externally injected bunch from a prior stage (as recommended by Hidding *et al.* [197]) and significantly simplify the operation of such a "hybrid accelerator".

5.5 Short-scale modulations in the longitudinal electron bunch profile

In a number of past experiments [178, 179], including the experimental campaign described in Chapter 4, significant coherent transition radiation was observed at optical frequencies, indicating the presence of short-scale substructure in the longitudinal bunch profile. In the present campaign, however, a particular feature in the visible spectrum was especially prominent. In the majority of cases where electrons were accelerated, a "bump" was seen in the transition radiation spectrum, centered at a wavelength of approx. 430–550 nm and having a width of ~ 100 nm. A sample CTR spectrum showing this bump in the context of the full spectrum is shown in Figure 5.11. In this case it was centered towards the short-wavelength end of the spectral range.

The position and shape of this bump vary from shot to shot, and we observe them to depend on the various acceleration parameters such as gas pressure, gas cell length, and laser pulse length. While it is difficult to completely exclude the possibility, we find it unlikely that the origin of this bump is in fact radiation from the driving laser (possibly from second harmonic generation) or other stray radiation, rather than actual transition radiation from the electron bunch. One of the clearest indications that this bump is in fact CTR stems from the fact that in many of the cases where the overall CTR spectrum is modulated due to interference of multiple bunches, the bump itself also exhibits modulations at the same frequency.

Assuming that this bump in the spectrum is indeed CTR due to short-scale longitudinal modulations in the electron bunch, seeking to explain the origin of these modulations is the obvious next step. While no definitive explanation stands out, interaction with the driving laser pulse is a potential candidate. Glinec *et al.*, in particular, observe coherent structures at very similar wavelengths and perform 3-D PIC simulations that indicate these arise from interactions with the laser [178]. Such interactions have also been observed by others: in PIC simulations [183, 199], as well as experimentally [58, 179, 200, 201]. In theory, interaction of the bunch with the laser pulse alone (in the absence of a plasma wave) can lead to microbunching of the electrons at half the drive laser wavelength [202], however the betatron oscillations of the bunch in the plasma channel can modify the dynamics [183]. With these considerations in mind, and observing that the wavelength of microbunching is indeed close to half the drive laser frequency — which is initially 800 nm, but will be redshifted through interaction with the plasma — interaction with the drive laser appears to be a plausible source of the observed fine structure.

Some of the "bump" data is summarized in Figure 5.12, again for the 80 mbar experimental run described in Sec. 5.3. We see that the central wavelength of the bump varies unambiguously with gas cell length, albeit over a relatively small wavelength range. In a very simple picture, this may be due to the bunch interacting with different parts of the laser pulse — which have different wavelengths due to chirp — as it moves forward through the plasma bubble.

A further striking feature is the fact that modulations of the "bump" are observed for gas cell lengths $\gtrsim 9$ mm, corresponding precisely to the regime where secondary bunches are observed close to the main bunch (cf. Fig. 5.10), but not to the regime where secondary bunches are observed further away (e.g. $L_{\text{cell}} \approx 6$–$9$), as discussed in Sec. 5.4. Indeed, if a bunch is accelerated in the second bucket behind the drive laser, we do not expect it to be modulated by the laser. Therefore only the bunch accelerated in the first bucket will have fine structure, and the "bump" in the spectrum would be present but not modulated due to interference. On the other hand, under the scenario described in Sec. 5.4, where for $L_{\text{cell}} \geq 10$ mm the laser no longer drives a wake but the electron bunch does, it is possible to envisage both electron bunches interacting with the (weakened but still present) laser pulse. Especially so considering the fact that the electron bunches, with relativistic Lorentz factors of order 1000 — would outpace the laser, which has a group velocity corresponding to a Lorentz factor $\gamma_{\text{g}} = \lambda_{\text{p}}/\lambda_{\text{d}} \approx 20$.

This phenomenon certainly warrants further investigation, in part by carrying out a full analysis — and possibly simulations — of the mechanisms that may cause laser-induced microbunching during acceleration, and in part through further experimental investigation, possibly with a spectrometer capable of detecting emission at shorter wavelengths. As we go

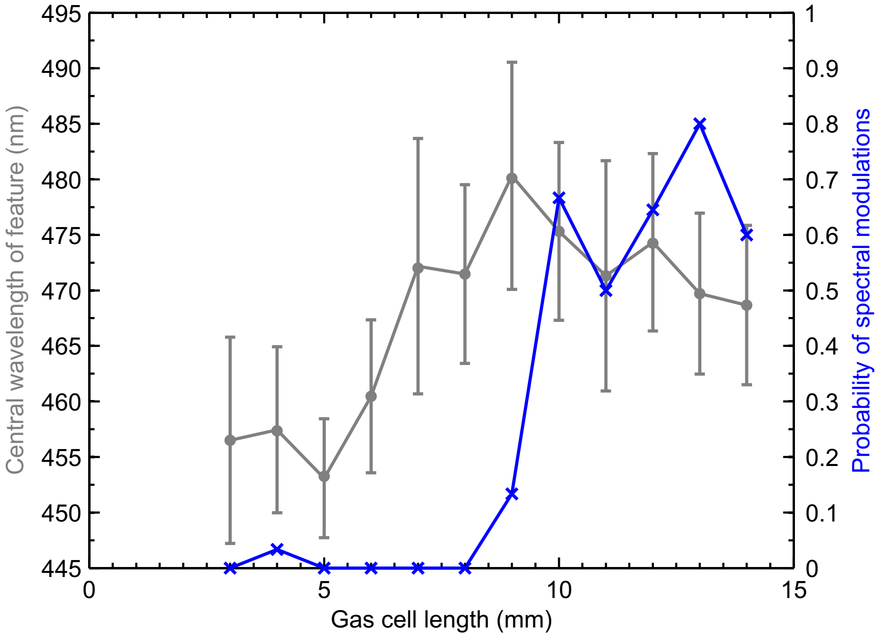

Figure 5.12: The central wavelength of the “bump” in the CTR spectrum, computed by taking the center of mass of the spectrum in the wavelength range 420–640 nm is shown as gray circles (error bars represent rms variation over ~ 30 shots). Blue crosses indicate the probability of the bump exhibiting modulations due to interference of multiple bunches with fine structure. The bump was observed on all shots taken at $L_{\text{cell}} \geq 3\,\text{mm}$ during this 70 mbar run.

on to describe in Section 6.5 of the next chapter, such short-scale longitudinal modulations of the electron bunch may have implications for the operation of future free-electron lasers driven by laser-accelerated bunches.

Chapter 6

Theory and Simulation of Free-Electron Lasers Seeded with Broadband Radiation

In this chapter we return to the subject of free-electron lasers. While in Chapter 3 we presented realistic simulations of a proposed FEL driven by electrons from a laser-wakefield accelerator, in this chapter we explore a related topic — the seeding with high-harmonic radiation — from a more idealized viewpoint. In particular, we aim to determine whether the broad bandwidth of HHG seeds — and the fact they often violates the commonly-used slowly-varying envelope approximation (SVEA) — precludes the use of conventional FEL codes for simulation of FELs seeded with them.

To this end, we begin by introducing a linearized analytical model for the amplification of radiation in an FEL. In Section 6.2 we then flesh out the problem at hand by introducing the slowly-varying envelope approximation, along with undulator-period averaging, and discussing their implications for simulation codes. We also present preliminary arguments for the applicability of the SVEA even in the case of broadband seeds. The non-SVEA, non-undulator-period-averaged code AURORA that we use to further test this hypothesis is presented in Section 6.3, and results from it as well as a conventional SVEA code (described in Appendix B) under various "challenging" conditions are presented in Section 6.4. Finally, in Section 6.5 we discuss the potential for coupling between non-resonant longitudinal modes of the radiation seed and the electron bunch profile that may result in resonant FEL emission.

The bulk of Sections 6.1–6.4 is based on material that has undergone peer review and is published in *Physical Review Special Topics — Accelerators and Beams*, Ref. [203].

6.1 The linearized analytical model of a seeded FEL

Before we begin the discussion of the slowly-varying envelope approximation and its implications for FEL simulations, in this section we present an analytical FEL model that provides another useful point of a reference. It treats the case of gain due to a radiation seed and an initially cold and uniform electron beam, and is based on the linearized Maxwell-Vlasov equations (within the SVEA), similar to the development in Refs. [204, 205].

For a seed field $A(\theta,0)$, represented in terms of magnetic vector potential, the field at a given distance along the undulator $\zeta = k_\mathrm{u} z$, and a given ponderomotive phase, $\theta = k_\mathrm{R}(z - ct) + k_\mathrm{u} z$, can be expressed as

$$A(\zeta,\theta) = g(\zeta,\theta) * A(0,\theta),$$

where $*$ represents convolution with respect to θ. Previous treatments have employed the saddle-point approximation when evaluating an expression for $g(\zeta,\theta)$, thereby precluding the correct modeling of fast-varying seed envelopes. In order to allow such envelopes to be modeled, here we avoid this approximation and derive an expression based on the generalized hypergeometric function ${}_pF_q$ [cf. Eqs. (6.11) and (6.13)]:

$$g(\zeta,\theta) = \delta(\zeta-\theta) + 8i(\zeta-\theta)\theta\rho^3\, {}_0F_2\left[;\frac{3}{2},2;2i(\zeta-\theta)^2\theta\rho^3\right] H(\zeta-\theta), \tag{6.1}$$

where H is the Heaviside step function.

We begin by introducing the function $f(\zeta,\theta,\eta)$ that is used to characterize the electron distribution in phase space, where $\eta = (\gamma-\gamma_\mathrm{R})/\gamma_\mathrm{R}$ is a dimensionless energy-spread variable. The distribution's evolution is governed by the Vlasov equation,

$$\mathrm{d}f = \left(\partial_\zeta + \dot\theta\partial_\theta + \dot\eta\partial_\eta\right) f = 0. \tag{6.2}$$

The dotted terms are derivatives of the phase space variables w.r.t. ζ, correspondingly

$$\dot\theta = 2\eta,$$
$$\dot\eta = -\kappa_B\, a e^{i\theta},$$

where $a \equiv eA/mc$ is the normalized field envelope. The Maxwell equation governing the latter is

$$\left(\partial_\zeta + \partial_\theta\right) a = \kappa_M\, e^{-i\theta} \int \mathrm{d}\eta\, f. \tag{6.3}$$

The two dimensionless constants are

$$\kappa_M = \frac{n_e e^2 a_\mathrm{u}[JJ]}{4c^2\gamma_\mathrm{R} k_\mathrm{R} k_\mathrm{u} m\epsilon_0},$$
$$\kappa_B = \frac{k_\mathrm{R} a_\mathrm{u}[JJ]}{2k_\mathrm{u}\gamma_\mathrm{R}^2},$$

where n_e is the peak electron density and all other parameters are as defined in Section 2.3.2. We note that $\kappa_M\kappa_B = 4\rho^3$, where ρ is the Pierce parameter (cf. Eq. (2.47)).

We rewrite the distribution function as $f = f_0 + f_1$, where f_0 is the initial distribution and f_1 accounts for FEL-induced modulation. The shot-noise contribution is not considered here. Linearizing Eq. (6.2) results in

$$\left(\partial_\zeta + \dot\theta\partial_\theta\right) f_0 = 0 \tag{6.4}$$

and

$$\left(\partial_\zeta + \dot{\theta}\partial_\theta\right) f_1 = \kappa_B a e^{i\theta} \partial_\eta f_0, \tag{6.5}$$

which are valid during the start-up and exponential gain regimes, where $a \ll 1$. Rearranging Eq. (6.5), we obtain

$$\left[\partial_\zeta + \dot{\theta}(i + \partial_\theta)\right](e^{-i\theta} f_1) = \kappa_B a \partial_\eta f_0.$$

We solve this in the Fourier domain, $\mathscr{F}[f(\theta)] = F(\nu)$:

$$\mathscr{F}[e^{-i\theta} f_1] = \int_0^\zeta \mathrm{d}\zeta' \, e^{-i\dot{\theta}(1+\nu)(\zeta-\zeta')} \kappa_B \mathscr{F}[a(\zeta')\partial_\eta f_0(\zeta')].$$

Carrying out the inverse transform and integrating over ν we obtain

$$e^{-i\theta} f_1 = \kappa_B \int_0^\zeta \mathrm{d}\zeta' \, e^{i\dot{\theta}(\zeta'-\zeta)} a(\zeta',\theta') \partial_\eta f_0(\zeta',\theta',\eta),$$

with $\theta' = \theta - \dot{\theta}(\zeta - \zeta')$. Finally, from Eq. (6.4) $f_0(\zeta,\theta,\eta) = f_0(\theta - \dot{\theta}\zeta, \eta)$ and thus also $f_0(\zeta',\theta',\eta) = f_0(\theta - \dot{\theta}\zeta, \eta)$.

Substituting the above into Eq. (6.3) yields

$$\begin{aligned}(\partial_\zeta + \partial_\theta) a &= \kappa_M \int \mathrm{d}\eta \, e^{-i\theta} f_1 && (6.6)\\ &= 4\rho^3 \int_0^\zeta \mathrm{d}\zeta' \, e^{i\dot{\theta}(\zeta'-\zeta)} a(\zeta',\theta') \int \mathrm{d}\eta \, \partial_\eta f_0(\theta - \dot{\theta}\zeta, \eta). && (6.7)\end{aligned}$$

This is amenable to solution via a Laplace transform, $\mathscr{L}[a(\zeta)] = A(s)$,

$$(s + \partial_\theta) A(s,\theta) - a(\zeta = 0, \theta) = 4\rho^3 \int_0^\infty \mathrm{d}\zeta \, e^{-s\zeta} \int_0^\zeta \mathrm{d}\zeta' \, e^{i\dot{\theta}(\zeta'-\zeta)} a(\zeta',\theta') \int \mathrm{d}\eta \, \partial_\eta f_0(\theta - \dot{\theta}\zeta, \eta).$$

Since $a(\theta') f_0(\theta - \dot{\theta}\zeta)$ contains higher-order terms, under linearization it becomes $a(\theta) f_0(\theta)$:

$$\begin{aligned}(s + \partial_\theta) A(s,\theta) - a(\zeta = 0, \theta) &= 4\rho^3 \int \mathrm{d}\eta \, \partial_\eta f_0(\theta,\eta) \int_0^\infty \mathrm{d}\zeta \, e^{-(s+i\dot{\theta})\zeta} \int_0^\zeta \mathrm{d}\zeta' \, e^{i\dot{\theta}\zeta'} a(\zeta',\theta)\\ &= 4\rho^3 \int \mathrm{d}\eta \, \partial_\eta f_0(\theta,\eta) \int_0^\infty \mathrm{d}\zeta' \, e^{-s\zeta'} \frac{a(\zeta',\theta)}{s + i\dot{\theta}}\\ &= 4\rho^3 \int \mathrm{d}\eta \, \partial_\eta f_0(\theta,\eta) \frac{A(s,\theta)}{s + i\dot{\theta}}.\end{aligned}$$

Rearranging and integrating by parts, we obtain

$$(s + \partial_\theta) A(s,\theta) = a_0(\theta) + i(2\rho)^3 A(s,\theta) \int \mathrm{d}\eta \, \frac{f_0(\theta,\eta)}{(s + 2i\eta)^2},$$

where the initial (seed) field is represented by $a_0(\theta) = a(\zeta = 0, \theta)$. Finally, we assume an initially cold, longitudinally uniform electron beam, $f_0(\theta,\eta) = \delta(\eta)$:

$$\left[s + \partial_\theta - \frac{i(2\rho)^3}{s^2}\right] A(s,\theta) = a_0(\theta). \tag{6.8}$$

The solution to this equation is

$$A(s,\theta) = \int_{-\infty}^{\theta} \mathrm{d}\theta' \exp\left[\left(-s + \frac{i(2\rho)^3}{s^2}\right)(\theta-\theta')\right] a_0(\theta'),$$

and carrying out the inverse Laplace transform,

$$a(\zeta,\theta) = \mathscr{L}^{-1}\left\{\exp\left[-s\theta + \frac{i(2\rho)^3}{s^2}\theta\right]\right\} * a_0(\theta), \tag{6.9}$$

where $*$ represents convolution with respect to θ.

Defining $\chi \equiv i(2\rho)^3\theta$, we go on to evaluate the inverse transform in Eq. (6.9). From the power series expansion

$$\frac{1}{s}\exp\left(\frac{\chi}{s^2}\right) = \sum_{n=0}^{\infty} \frac{\chi^n}{s^{2n+1}n!},$$

applying the relation $\mathscr{L}^{-1}\left\{\frac{1}{s^{n+1}}\right\} = \frac{\zeta^n}{n!}$, we obtain the inverse Laplace transform

$$\mathscr{L}^{-1}\left\{\frac{1}{s}\exp\left(\frac{\chi}{s^2}\right)\right\} = \sum_{n=0}^{\infty} \frac{\chi^n \zeta^{2n}}{n!(2n)!} \equiv G(\zeta,\theta).$$

Using

$$\mathscr{L}^{-1}\{sF(s)\} = f'(\zeta) + f(0)\delta(\zeta),$$

and

$$\mathscr{L}^{-1}\{\exp(-s\theta)F(s)\} = f(\zeta-\theta)H(\zeta-\theta),$$

where $\mathscr{L}^{-1}\{F(s)\} = f(\zeta)$, and H is the Heaviside step function, we obtain

$$\mathscr{L}^{-1}\{\exp(-s\theta)\exp(\chi/s^2)\} = \frac{\partial G}{\partial \zeta}(\zeta-\theta,\theta)H(\zeta-\theta) + \delta(\zeta-\theta). \tag{6.10}$$

The solution for $\partial G/\partial\zeta$ can be expressed in terms of the generalized hypergeometric function, ${}_pF_q$,

$$\begin{aligned}\frac{\partial G}{\partial \zeta}(\zeta-\theta,\theta) &= \sum_{n=1}^{\infty} \frac{\left[i(2\rho)^3\theta\right]^n (\zeta-\theta)^{2n-1}}{n!(2n-1)!} \\ &= {}_0F_2\left(;1,1/2;i(2\rho^3)\theta\zeta^2\right),\end{aligned} \tag{6.11}$$

where

$${}_pF_q\left(a_1,\ldots,a_p;b_1,\ldots,b_q;z\right) = \sum_{k=0}^{\infty} \frac{(a_1)_k(a_2)_k\ldots(a_p)_k}{(b_1)_k(b_2)_k\ldots(b_q)_k}\frac{z^k}{k!} \tag{6.12}$$

with $(a)_k \equiv a(a+1)...(a+k-1)$; $(a)_0 \equiv 1$.

The series of Eq. (6.11) is the same as that obtained by Krinsky in Ref. [206]. For values of $\zeta \lesssim \pi/\rho$ (beyond which non-linear effects manifest) and $\theta \in (0,\zeta)$, it is accurately approximated by its first four terms:

$$\begin{aligned}&\frac{\partial G}{\partial \zeta}(\zeta-\theta,\theta) \\ &\approx 8i(\zeta-\theta)\theta\rho^3 - \frac{16}{3}(\zeta-\theta)^3\theta^2\rho^6 \\ &\quad - \frac{32}{45}i(\zeta-\theta)^5\theta^3\rho^9 + \frac{32}{945}(\zeta-\theta)^7\theta^4\rho^{12}.\end{aligned} \tag{6.13}$$

In obtaining this expression we have not used the saddle-point approximation when evaluating the inverse Laplace transform. It is customary to apply this approximation (see e.g. [137, 205]), and indeed it yields the familiar expression for exponential FEL growth in the linear regime — however, it leads to expressions that cannot have short-scale variations in θ, thus precluding the modeling of gain due to fast-varying seeds. The present expression does not suffer from this drawback and can be used to model gain due to seeds with an arbitrary temporal profile.

6.2 The slowly-varying envelope approximation in conventional FEL codes

As discussed in Sec. 2.3.2, FEL seeds produced by high-harmonic generation comprise a series of sub-femtosecond spikes — as short as ≈ 100 attoseconds [207] — and can have a very wide bandwidth. At the same time, conventional SVEA codes such as GENESIS [167] or PERSEO [208] (see also Appendix B) normally represent the seed field in terms of the complex amplitude of the field's resonant component on a grid of spacing $\lambda_0 = n\lambda_{\mathrm{R}}$, for integer $n \geq 1$. This can support a bandwidth of $\frac{2}{3}\lambda_{\mathrm{R}} \leq \lambda < 2\lambda_{\mathrm{R}}$ with a resolution in frequency of $2\pi c/(L\lambda_{\mathrm{R}})$, where L is the total number of data points on the grid.

For a resonant FEL wavelength $\lambda_{\mathrm{R}} = 32\,\mathrm{nm}$, an HHG burst of length 100 as corresponds to just a single data point on this grid (for $n = 1$, $\lambda_0 = \lambda_{\mathrm{R}}$). Correspondingly, the limited bandwidth that an SVEA code supports precludes the inclusion of the broad range of frequencies that are present in these HHG seeds. We would like to investigate the extent to which the omission of these frequencies compromises the ability of single-frequency codes to correctly simulate the seeding process.

The effects of an HHG seed's rich spectral content and sub-femtosecond temporal structure upon FEL gain have been considered previously by Giannessi *et al.* [209], McNeil *et al.* [210] and Wu *et al.* [211], among others. In Refs. [209] and [210] the authors dismiss such effects reasoning that since the FEL gain bandwidth typically encompasses only a single harmonic of the seed, the presence of the other harmonics will therefore negligibly affect FEL dynamics. In both cases, simple filtering of the HHG seed in Fourier space was carried out in order to use numerical codes based on the slowly-varying envelope approximation (SVEA); in general, such codes support only a limited frequency domain around a central frequency. Such filtering can lead to non-causal changes in the temporal structure, e.g., making radiation artificially appear in the valleys between the strong HHG spikes. Wu *et al.*, on the other hand, discuss smearing of the seed's temporal structure within an analytical framework.

A related approximation used by nearly all conventional FEL simulation codes is the so-called undulator-period-averaging in which the particle-field interaction is locally in time and z reduced to only the resonant component. Since a particle slips relative to the radiation field one resonant wavelength in time for each undulator period travelled in z, as with SVEA it is not clear that such averaging will work well for temporal radiation structures with spikes of order a single resonant wavelength in duration. At the other extreme — wavelengths much longer than that corresponding to FEL resonance — it is not clear that the strong radiation components in the HHG spectrum at the low harmonics of the drive laser can be safely neglected in terms of their effects upon the particle dynamics and, if not, whether the combination of spectral filtering and undulator-period averaging accurately capture these effects.

Here we revisit the issue of the use of SVEA and undulator-period averaging in FEL simulation and, more specifically, the applicability of the SVEA when modeling the HHG seed's temporal and spectral structure vis-à-vis FEL amplification. Since the SVEA assumes a radiation amplitude and phase that implicitly varies slowly compared to a central wavelength λ_0 and a radiation spectral content that can be mapped to a relatively narrow region around λ_0, it is not immediately clear that conventional FEL codes operating within the constraints of SVEA will accurately capture all the important details of HHG seeding.

We begin the discussion with an introduction to the slowly-varying envelope approximation in the context of its application in free-electron laser simulation codes, together with some qualitative arguments for why the SVEA in 1-D geometries can handle a surprisingly wide spectral content range for an FEL seed.

6.2.1 Simulation framework and core equations

Simulating gain in a free-electron laser involves the evaluation of the combined action of the undulator and radiation fields on the electrons at each position along the undulator, as well as the evaluation of the emitted field, based on the electron positions and velocities at that point. Generally these interactions are averaged over one or more undulator periods, with all variables propagated from one undulator position to another. At any given interaction point within the classical framework, the effect of the radiation *on* the electrons is independent of the emission of radiation *by* the electrons. From a numerical standpoint the effect of a seed is entirely mediated by the electrons, in the sense that the seed causes energy modulation and then microbunching, which in turn causes emission of radiation. When the microbunched electrons radiate, the presence of the seed radiation has no explicit effect on the newly emitted radiation and the two radiation fields add linearly. This view of the emission process is completely independent of the adoption of the SVEA.

Radiation emission is calculated from Maxwell's wave equation for the transverse components of the field's magnetic vector potential in the Lorenz gauge, $\vec{A}$, and for the source current density, $\vec{J}$,

$$\left(\frac{\partial^2}{\partial z^2} - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)A_\perp = -\frac{1}{\epsilon_0 c^2}J_\perp\,, \tag{6.14}$$

where c is the speed of light and ϵ_0 is the vacuum permittivity. In this paper we restrict ourselves to a 1-D model and have accordingly dropped the diffraction term, $\nabla_\perp^2 A_\perp$.

At this point one may proceed in two different fashions in regards to the eventual applicability of the SVEA in the context of FEL emission. In the strict 1-D case, we can factor the left-hand side of Eq. (6.14) to obtain

$$\left(\frac{\partial}{\partial z}+\frac{1}{c}\frac{\partial}{\partial t}\right)\left(\frac{\partial}{\partial z}-\frac{1}{c}\frac{\partial}{\partial t}\right)A_\perp=-\frac{1}{\epsilon_0 c^2}J_\perp. \tag{6.15}$$

With $B_\perp=\nabla\times A_\perp$ and $E_\perp=-\partial A_\perp/\partial t$, one may separate Eq. (6.15) into two separate equations for the forward-going wave $E_\perp-\vec{z}\times cB_\perp$ and the backward-going wave $E_\perp+\vec{z}\times cB_\perp$. In the FEL context, this result was previously obtained by Maroli [212]; it is commonly used as a means of of advancing Maxwell's equation for electromagnetic codes in different contexts of plasma physics [213].

Given our focus upon short wavelength FELs for which there is no resonant gain in the backward wave, we neglect it and set $cB_\perp=\vec{z}\times E_\perp$ exactly. We then make a Galilean transformation to the frame of reference moving forward in z at velocity c by carrying out the substitutions $\hat{t}=t-z/c$ and $\hat{z}=z$, resulting in the exact first order equation

$$\frac{\partial}{\partial\hat{z}}E_\perp(\hat{t})=-\frac{1}{2\epsilon_0 c}J_\perp \quad . \tag{6.16}$$

This equation can be integrated forward in z in either the time or frequency domain without strictly requiring the application of the SVEA. In this situation the radiation bandwidth that can be supported by the code is of course still dependent on the resolution of the grid in $\hat{t}$ on which $E_\perp$ and $J_\perp$ are recorded.

However, this approach has no equivalent for higher dimensions; including physical effects such as diffraction brings in non-trivial complications, even for the forward-going radiation wave. Accordingly, we proceed along a second path that will illustrate the approximations required for the SVEA. Using the Galilean transformations mentioned above, we rewrite Eq. (6.14) as

$$\left(\frac{\partial^2}{\partial\hat{z}^2}-\frac{2}{c}\frac{\partial^2}{\partial\hat{z}\partial\hat{t}}\right)A_\perp=-\frac{1}{\epsilon_0 c^2}J_\perp. \tag{6.17}$$

The field and source current may be represented as the products of envelopes and oscillations at an angular frequency $\omega_0=2\pi c/\lambda_0$, where λ_0 is typically chosen equal to the resonant wavelength λ_R (and in an SVEA code ω_0 is the center of the frequency window),

$$A(\hat{z},\hat{t})=A_0(\hat{z},\hat{t})e^{-i\omega_0\hat{t}},$$
$$J(\hat{z},\hat{t})=J_0(\hat{z},\hat{t})e^{-i\omega_0\hat{t}},$$

where we have dropped the $_\perp$ subscript to aid readability. Notably, this representation makes no assumptions regarding the envelopes — it still allows for variations on arbitrarily short timescales. Hereby Eq. (6.17) can be recast as

$$\frac{2i\omega_0}{c}\frac{\partial A_0}{\partial\hat{z}}=-\frac{1}{\epsilon_0 c^2}J_0-\frac{\partial^2 A_0}{\partial\hat{z}^2}+\frac{2}{c}\frac{\partial^2 A_0}{\partial\hat{z}\partial\hat{t}}. \tag{6.18}$$

6.2.2 The slowly-varying envelope approximation

The slowly-varying envelope approximation consists of dropping the last two terms on the right-hand side (r.h.s.) of Eq. (6.18). SVEA is applicable provided that

$$\left|\frac{\partial^2 A_0}{\partial\hat{z}^2}\right|\ll 2\frac{\omega_0}{c}\left|\frac{\partial A_0}{\partial\hat{z}}\right| \tag{6.19}$$

and

$$\left|\frac{\partial^2 A_0}{\partial\hat{z}\partial\hat{t}}\right| \ll 2\omega_0 \left|\frac{\partial A_0}{\partial\hat{z}}\right|. \tag{6.20}$$

Condition (6.19) applies strongly for the forward wave component of essentially all realistic FEL configurations, especially those with $\lambda_\mathrm{R} \ll \lambda_\mathrm{u}$ since for $\rho \leq 10^{-2}$ the exponential gain length $L_\mathrm{G} \geq 5\lambda_\mathrm{u} \gg \lambda_\mathrm{R}$, and the rate of change of A_0 along the undulator scales as the gain length, $\partial A_0/\partial\hat{z} \sim A_0/L_\mathrm{G}$. Here ρ is the FEL Pierce parameter introduced in Eq. (2.47), and $L_\mathrm{G} = \lambda_\mathrm{u}/(4\sqrt{3}\pi\rho)$ is the gain length (cf. Eq. (2.49)). Although the backward wave can violate condition (6.19) as mentioned above, there is no resonant gain for this component and its effective strength is a factor $\sim O(\gamma^2\lambda_\mathrm{u}/l_\mathrm{b})$ less than the forward wave where l_b is the electron bunch length.

With regard to condition (6.20) in Eq. (6.18), the $\partial^2 A_0/\partial\hat{z}\partial\hat{t}$ term is a correction that is relevant only when the *change* in the field, $\partial A_0/\partial\hat{z}$, is fast-varying in $\hat{t}$ — therefore, the initial radiation seed, which is invariant with z in the 1-D approximation whatever its spectral content, does not contribute to this term.

Strong initial modulations on the bulk current or externally produced initial microbunching bring about effects such as coherent spontaneous emission (CSE) that might violate condition (6.20); these have been discussed in detail elsewhere [214–218] and are outside the scope of our focus on seeded FELs. A recent paper by Maroli *et al.* [219] discusses FEL simulations outside the SVEA and also highlights CSE as an important effect that is not observed in conventional undulator-period-averaged codes.

On the other hand, whether or not such modulations in the electron beam may develop in z in response to structures in the seed whose characteristic frequency ω is far from ω_0 depends on their evolution with z. Realistic FELs with $\rho \ll 1$ will have resonant growth only for modulations whose frequencies lie within a narrow bandpass $\Delta\omega \sim \rho\omega_0$ centered on ω_0 (ignoring higher harmonics and presuming $\lambda_0 \equiv \lambda_\mathrm{R}$). Non-resonant seed structures are expected to lead at most to short-lived (in z) modulations in electron beam energy and ponderomotive phase. Unless these modulations are sufficiently strong to either affect growth at the central resonant wavelength λ_R (e.g., strongly increasing the coarse-grained energy spread and thus reducing gain) or to shift growth to neighboring wavelengths (e.g., sideband formation), one expects condition (6.20) to remain valid and application of the SVEA to be justifiable in the context of coherent radiation emission even when the initial field envelope is fast-varying (in phase or amplitude) compared to ω_0.

As a final point, here we transformed to a frame of reference moving at c, rather than, as is more common in FEL codes, to a frame moving at $\beta_z c$. In the latter the radiation slips over the electrons at a rate of $(1-\beta_z)c$, thus even in the absence of a source current the field is not invariant. This would pose a problem if the propagation of the field along the wiggler was evaluated by means of the wave equation, which within the SVEA would not handle a fast-varying temporal profile correctly. However, 1-D FEL codes use a simpler approach and advance the radiation field by translating it by the required amount, $\Delta t = \Delta z/(1-\beta_z)c$ [133]. In the absence of a source current this does leave the field invariant, except for the position shift. This procedure deals correctly with any field, regardless of whether it obeys the SVEA or not. Due to the way radiation is advanced between integration steps in FEL codes, therefore, the above discussion, for a reference frame moving at c, will also be valid in the ponderomotive reference frame moving at $(1-\beta_z)c$ that is conventionally used in codes.

In order to check these preliminary indications that the slowly-varying envelope approximation can still be applied in the context of FELs seeded with broadband radiation, in Section 6.4 we compare simulation results from a conventional 1-D SVEA code, in this case the one described in Appendix B, to those from the non-undulator-period-averaged, non-SVEA code AURORA that we introduce next. The results are also compared to those from the linearized model of Section 6.1.

6.3 The 1-D non-undulator-period-averaged, non-SVEA code AURORA

For the purposes of numerically evaluating the effects of the slowly-varying envelope approximation and of undulator-period averaging in FEL simulations, we developed the non-SVEA, non-undulator-period averaged code AURORA [203, 220], which we present here. As a multi-frequency code it shares a number of features with MUFFIN by Piovella [214]. We note that Campbell *et al.* [221] have developed a fully 3D, non-SVEA code that also follows Piovella and advances the wave equation (i.e., our Eq. (6.18) with diffraction terms) in Fourier space.

6.3.1 General framework

In AURORA, the FEL interaction is simulated within a moving time window of duration $L\lambda_\mathrm{R}/c$, where L is an integer, and adopts periodic boundary conditions. As the window propagates down the undulator at a constant speed of $+c$, the evolution of the radiation field and the electron positions and energies are computed at M points per undulator period. The coordinate within the window is the phase measured relative to a resonant, monochromatic plane wave, $\varphi = k_\mathrm{R}(z-ct)$ where $k_\mathrm{R} = 2\pi/\lambda_\mathrm{R}$ is the resonant wavenumber. The usual ponderomotive phase is $\theta = \varphi + \zeta$ where $\zeta = k_\mathrm{u}z$ is the distance along the undulator normalized by the undulator wavenumber $k_\mathrm{u} = 2\pi/\lambda_\mathrm{u}$.

In AURORA the radiation is represented by the transversely-polarized, normalized magnetic vector potential, $a(z,\varphi) = eA_x(z,\varphi)/mc$. Recording this on a grid of M points per wavelength allows a simulation bandwidth of $2\lambda_\mathrm{R}/M < \lambda < L\lambda_\mathrm{R}$. The electron bunch is represented by a set of P macroparticles whose weights, w_i, determine the number of electrons represented by each. The macroparticles are free to move within the window (i.e., are not confined to "bins", as in undulator-period-averaged codes) with any electrons escaping through the back of the window re-entering through the front. Initial shot-noise microbunching can be modeled by introducing a certain microbunching level at the fundamental as suggested by Giannessi [149] or — preferably — by a Poissonian distribution of either macroparticle displacement [143, 222] or weights w_i, as suggested by McNeil *et al.*

[217].

The bunch cross-section Σ_r is assumed to be constant and matched to the radiation cross-section, and, given our 1-D geometry, diffraction effects are not considered. Our studies adopt a linearly-polarized undulator with constant period λ_u and constant rms strength parameter a_u.

6.3.2 FEL equations

The transverse current density at location z, $J_x(z,\varphi)$, is evaluated through

$$J_x(z,\varphi) = \frac{e\sqrt{2}a_\mathrm{u}}{\Sigma_r}\cos(k_\mathrm{u}z)\sum_{i=1}^{P}\frac{w_i}{\gamma_i}C(\varphi-\varphi_i), \tag{6.21}$$

where $k_\mathrm{u} = 2\pi/\lambda_\mathrm{u}$ is the undulator wavenumber, i runs over the macroparticle indices, γ_i are the macroparticle Lorentz factors, and φ_i are the resonant radiation phases at the macroparticle locations. The distribution C represents the particle-in-cell weighting applied to the macroparticles in order to project their distribution onto the radiation grid.

In AURORA, the current density J_x is transformed to Fourier space,

$$\tilde{J}_x(z,\omega) = \int \mathrm{d}t\, J_x(z,t)\, e^{i\omega t},$$

and filtered with a Gaussian of width $\sim 30\omega_0$ centered at $\omega = 0$ in order to reduce high-frequency noise. The result is then inserted into the Fourier-transform of Eq. (6.17), omitting the second derivative with respect to z:

$$\frac{\mathrm{d}}{\mathrm{d}z}\tilde{A}_x(z,\omega) = \frac{1}{2i\omega}\frac{1}{\epsilon_0 c}\tilde{J}_x(z,\omega), \tag{6.22}$$

to compute the field evolution.

Macroparticle energies and positions are propagated using

$$\frac{\mathrm{d}\gamma_i}{\mathrm{d}z} = -\frac{\sqrt{2}ea_\mathrm{u}}{\gamma_i mc^2}\cos(k_\mathrm{u}z)E_x(z,\varphi_i), \tag{6.23}$$

and

$$\frac{\mathrm{d}\varphi_i}{\mathrm{d}z} = -\frac{k_\mathrm{R}}{2\gamma_i^2}\left\{1 + a_\mathrm{u}^2\left[1+\cos(2k_\mathrm{u}z)\right]\right\} \tag{6.24}$$

with $E_x = -\partial A_x/\partial t$. The cosine term in Eq. (6.24) corresponds to the oscillation of the electron's longitudinal velocity that is characteristic of planar undulators, naturally leading to the emission of odd radiation harmonics of the fundamental.

Since AURORA does not explicitly compute the transverse particle trajectories, the above equations omit higher order terms (relating z_i and φ_i) in the longitudinal trajectories. As normally done in FEL simulation codes, the radiation field contribution (i.e., the quiver velocity) to Eq. (6.24) is neglected, and macroparticle motion is assumed to be dominated by the undulator field. We have verified that this is justified since $a \ll a_\mathrm{u}$ for all normalized seed and undulator potentials considered in Sec. 6.4.

Finally the computation of long-scale length space-charge forces is implemented in AURORA, based on the treatment by Grüner *et al*. [129], but was not switched on for the simulations presented here. For several of the cases presented, corresponding simulations including space-charge effects were run in order to verify that they do not introduce any new physics, other than the expected reduction in gain.

6.3.3 Benchmarking

We verified AURORA's performance in benchmarks against PERSEO. The latter was chosen since in the case of a slowly-varying field envelope, small electron energy spread and no coherent spontaneous emission, the output from AURORA should match PERSEO's. Here we present results from a steady-state and a time-dependent SASE benchmark. The parameters used are close to the values discussed in Sec. 3.2.2; they are given in Table 6.1:

Table 6.1: Parameters for AURORA benchmark

Electrons		
Energy	E	320 MeV
	γ	625
Energy spread	σ_γ/γ	0.1 %
Peak current	I_p	10 kA
Transverse bunch size	σ_r	100 µm
Undulator		
Period	λ_u	20 mm
rms strength	a_u	0.5
Resonant wavelength	λ_R	32 nm
Pierce parameter	ρ	0.0041
Gain length	L_G	0.23 m

We found a configuration of $M = 256$ and 800 macroparticles per wavelength (i.e., $P = 800 \times L$) to be sufficient for good agreement with analytical predictions for the gain length, as well as energy conservation (both within $\sim 0.5\,\%$). Figure 6.1 shows a comparison of the results of a steady state simulation, i.e. $L = 1$.

A time-dependent SASE benchmark was also performed in a window of $L = 500$: the output profiles near saturation are shown in Fig. 6.2. The phases of the initial shot noise bunching were random, but identical for both codes. The beam current was constant throughout the window. In order to compare the profile of the power output from AURORA to that from PERSEO, the E-field was split into slices of length λ_0, and the power carried by the fundamental Fourier component of each slice was evaluated, thus reducing the effective bandwidth to match PERSEO's. A similar — but more sophisticated — conversion was applied

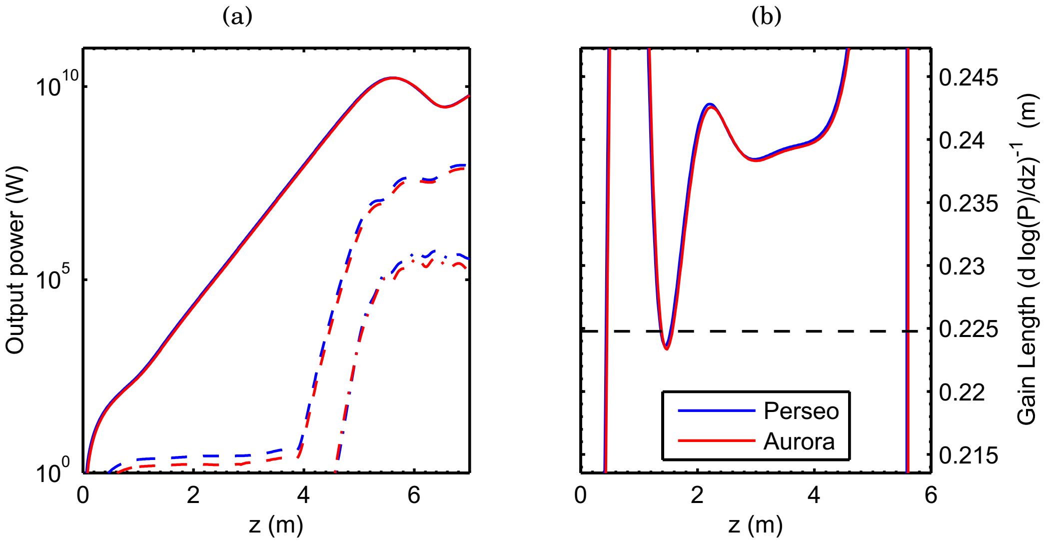

Figure 6.1: Steady state benchmark of output power (a) and gain length (b). In (a) solid lines represent emission in the fundamental, dashed lines are 3rd harmonic and dash-dotted are 5th. The horizontal black line in (b) is the 1-D analytical gain length for an initially cold beam.

for the other results presented in this chapter, as described next. These plots clearly show that the AURORA output matches that of PERSEO in these benchmark cases.

6.3.4 Conversion from a broadband to an SVEA-compatible field

The conventional undulator-period-averaged SVEA code used for the present studies (hereafter "SVEA code") is a 1-D code based on the model of PERSEO [208] that is described in Appendix B. The time grid spacing is typically λ_0/c, where λ_0 is the bin size, which we assume to be equal to the FEL's resonant wavelength $\lambda_0 \equiv \lambda_R$. Hence, there is support for angular frequencies in the range $[\omega_0/2, 3\omega_0/2]$, where $\omega_0 = 2\pi c/\lambda_0$. A key first step towards using the SVEA code to simulate the FEL response to a broadband seed is the requirement to re-bin the seed onto a relatively coarse (relative to that of AURORA) time grid, while ensuring important characteristics of the seed are retained.

The re-binning process consists of two steps. First, a Gaussian spectral filter of root-mean-square (rms) width $\omega_0/(2\pi)$ centered on ω_0 is applied. This spectral filter serves to remove spectral components of extremely high and extremely low frequencies that may otherwise lead to aliasing anomalies; at the same time the filter remains broad enough to preserve the temporal profile of the seed — and therefore of the FEL output during early stages of gain — on scales of $\gtrsim \lambda_0$.

Second, the filtered radiation field is transformed to the time domain and split into sections of length λ_0. The binned field then comprises the complex amplitudes of the fundamental spectral component of the broadband field in each section. The approach of this second step is informed by the discussion in Appendix C, where the energy modulation that a broadband seed induces on the electron bunch is analyzed. The spectral filtering in the

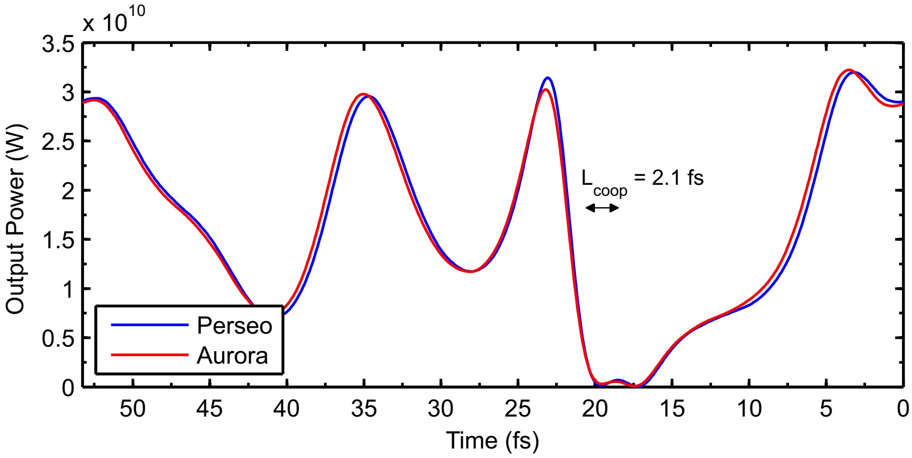

Figure 6.2: SASE output profiles near saturation ($z = 6.9\,\mathrm{m}$) at the fundamental. The cooperation length, $L_\mathrm{c} = \lambda_\mathrm{R}/4\pi\rho$, is also indicated.

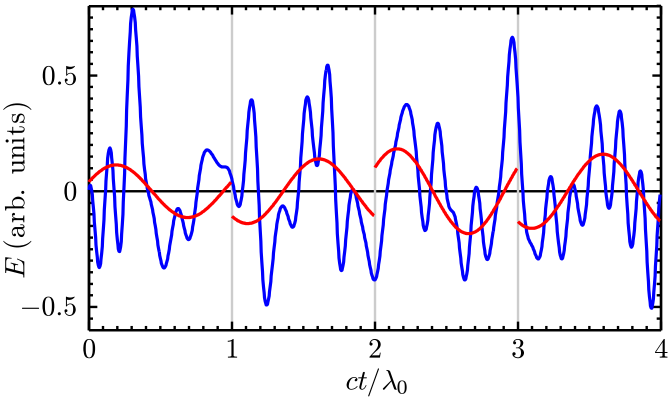

Figure 6.3: The procedure for binning a broadband field to a grid of spacing λ_0, applied to a sample randomly-generated field. Blue is the original electric field; red sinusoids represent the phase and amplitude obtained for the fundamental frequency component within each bin of width λ_0.

first step of the conversion procedure is required to address the caveats of simple binning discussed at the end of Appendix C.

Figure 6.3 illustrates the effects of the procedure described here on a sample seed. Note that this is carried out on the electric field, E, rather than the magnetic vector potential A. In order to allow easier comparison of results from the two simulation codes, the same process is applied to the AURORA radiation output.

In the frequency domain, this filtering process is equivalent to multiplication of the input field $E(\omega)$ with a filter function $W_{\lambda_0}(z_1,\omega)$ that is the product of a Gaussian filter and the Fourier transform of a top-hat of width λ_0, centered at the bin location $z_1 = ct - z$. On a discretized mesh the correct filter to convolve with is

$$W_{\lambda_0}(z_1,\omega) = \frac{\sqrt{2\pi}}{\omega_0} \exp[-2\pi^2(\omega-\omega_0)^2/\omega_0^2] \times \frac{\sin(\pi\omega/\omega_0)}{\sin[\pi\omega/(M\omega_0)]} \exp(-iz_1\omega c), \tag{6.25}$$

where λ_0/M is the data point spacing on the grid on which the full broadband seed is discretized.

6.4 Simulation results

In this section we compare AURORA results with those from the SVEA code and, in one case, from our analytic model. We first examine the FEL response to an ultrashort input seed and then go on to examine several situations involving HHG seeds.

6.4.1 Ultrashort seeds

Our first simulation study involves an FEL seeded with a single ultra-short spike of radiation, whose high frequency components strongly violate the SVEA. The relevant FEL parameters include resonant wavelength $\lambda_{\rm R} = 32\,\text{nm}$, undulator period $\lambda_{\rm u} = 20\,\text{mm}$ and rms strength $a_{\rm u} = 0.5$, and a monoenergetic, zero emittance electron beam with $\gamma = 625$ and a current density corresponding to an FEL parameter $\rho = 0.002$. The following figures use normalized units for space, time, and field measures

$$\bar{z} = 2k_{\rm u}\rho z, \tag{6.26}$$

$$\bar{z}_1 = 2k_{\rm u}\rho\frac{\bar{\beta}_z}{1-\bar{\beta}_z}(ct-z), \tag{6.27}$$

$$\bar{A} = \frac{k_{\rm R}a_{\rm u}[JJ]}{4\gamma_{\rm R}^2 k_{\rm u}\rho^2}\frac{e}{mc}A, \tag{6.28}$$

where $\bar{\beta}_z = \sqrt{1-1/\gamma_{\rm R}^2}$ is the mean normalized electron velocity in z.

Our results show that even for this case where the input seed amplitude exhibits significant variations on the scale of a resonant wavelength, there is excellent agreement between results from the SVEA and non-SVEA codes and the analytical model in the great majority of cases. Figure 6.4 presents a comparison between the radiation output profiles from the non-SVEA code AURORA, the conventional SVEA code presented in Appendix B, and the linearized analytical model of Sec. 6.1 at different stages of gain.

The initial seed has a Gaussian profile with amplitude $A_{0,\text{seed}} \sim e^{-(ct)^2/2\sigma_{\text{seed}}^2}$ and rms duration $\sigma_{\text{seed}} = \lambda_0/8$. In terms of normalized variables, the rms width of the seed spike is $\bar{z}_1 = \pi\cdot 10^{-3}$ (in other words, much shorter than the cooperation length $L_{\rm c}$. In the SVEA code the time grid spacing for the radiation and particle quantities is $\lambda_{\rm R}/c$ and the central wavelength $\lambda_0 = \lambda_{\rm R}$. The evolution of the output radiation profiles illustrate the high degree of correspondence between the results yielded by the different approaches. At $\bar{z} = 10$ the simulations approach saturation and predictably begin to deviate from the analytical model. Simulations of other ultrashort seeds with rms widths as short as $\sim \lambda_0/100$ showed similar excellent agreement between results from the SVEA code initialized from a filtered seed and the non-SVEA AURORA and analytic model.

6.4.2 Seeding with full HHG spectrum

In order to test further the capabilities of the simulation codes we turn to an example that better illustrates actual HHG seeding, with an input seed spectrum comprised of odd har-

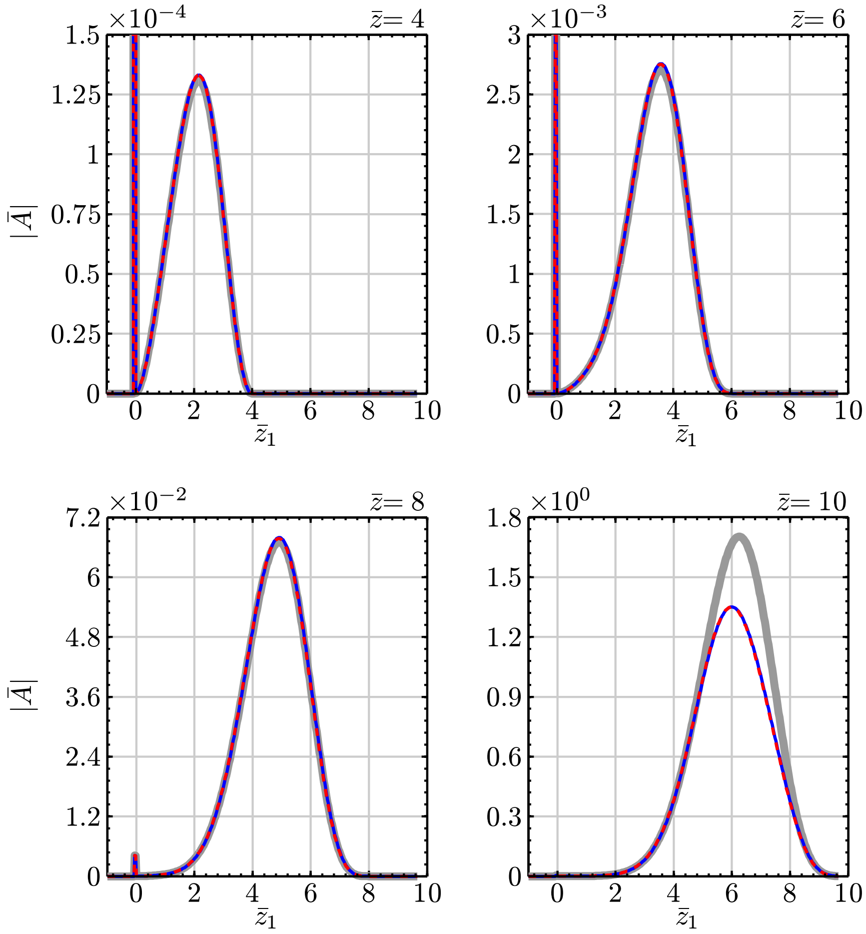

Figure 6.4: Evolution at various $\bar{z}$ locations of the intensity of the radiation A for a system seeded with a single ultra-short spike of rms width $\sigma_{\bar{z}_1} = \pi \cdot 10^{-3}$ (corresponding to $\lambda_{\mathrm{R}}/8$) located at $\bar{z}_1 = 0$. The thick gray curve is from the analytical model, the solid blue from the non-SVEA code Aurora, and the dashed red curve from the conventional SVEA code. The latter almost perfectly overlaps the blue curve.

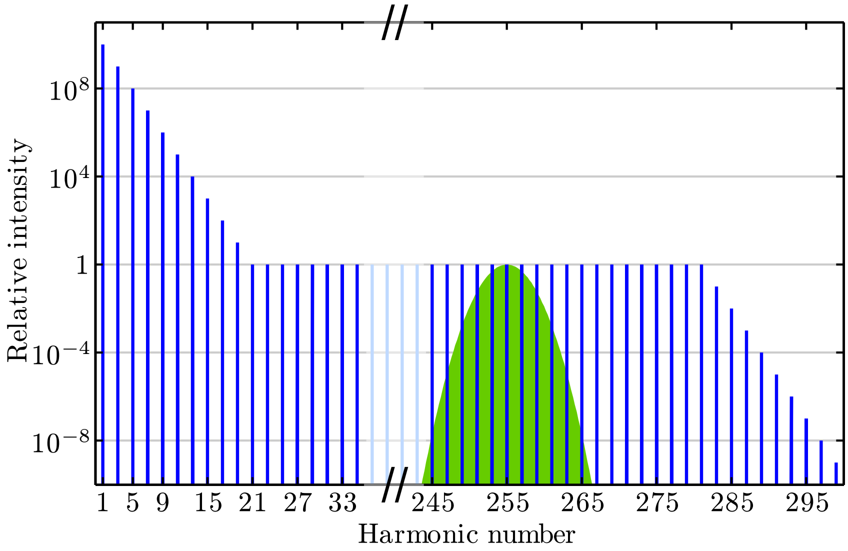

Figure 6.5: Relative intensities of harmonic orders 1–35 and 245–299 of an input HHG seed spectrum. The shaded green Gaussian around the 255th harmonic (also plotted on a logarithmic scale) represents the FEL gain curve for $\rho = 5 \times 10^{-3}$.

monics of a hypothetical drive laser. In the time domain, this corresponds to a sequence of ultra-short spikes, two for each cycle of the driving laser. We simulate seeding with an idealized infinite sequence of such spikes by using a periodic simulation window of size $\lambda_\mathrm{d} = 800\,\mathrm{nm}$, the drive laser wavelength. FEL resonance was set to $\lambda_\mathrm{R} \approx 3.14\,\mathrm{nm}$, the 255th harmonic of the λ_d. Other relevant simulation parameters include $\lambda_\mathrm{u} = 5\,\mathrm{mm}$, $a_\mathrm{u} = 0.5$, $\gamma = 998$, and FEL parameter $\rho = 5 \times 10^{-3}$.

The corresponding input spectrum is shown in Fig. 6.5 where we also presumed that all harmonics are exactly in phase. The adopted spectrum has an exponential drop-off of intensity at the lowest harmonics and a plateau region of constant intensity per harmonic, in line with theoretical predictions and experimental observations [223]. The normalized amplitude of each of the plateau harmonics ($q = 21$ to $q = 281$) is $\bar{A}_q = 2.3 \times 10^{-5}$.

Despite the large electric fields associated with the seed's long-wavelength components, their dimensionless amplitudes $a_\mathrm{seed} = eA_\mathrm{seed}/(mc)$ are still much less than the undulator parameter a_u, therefore we may continue to disregard their direct contribution to electron motion (cf. Eq. (6.24)). Additionally, although in reality these long-wavelength components are likely to diffract significantly over the course of FEL gain (or even during pre-undulator transport), by neglecting diffraction effects we may consider a limiting case in terms of their potential effect on gain.

Figure 6.6 plots simulation results at an early stage of gain, $\bar{z} = 4$, and close to FEL power saturation, $\bar{z} = 14$, where $|\bar{A}| \approx 1$. Comparing the time-dependent output profiles, it is clear that the SVEA code results match those from AURORA very closely — to within the few-percent level — at all stages of gain. The radiation phases exhibit a similar degree of correspondence. For purposes of comparison here, the AURORA results have also been filtered and binned according to the procedure outlined in Sec. 6.3.4. This does not, however, limit their validity, as we have verified that the filter's bandwidth covers all spectral

(a) $\bar{z} = 4$

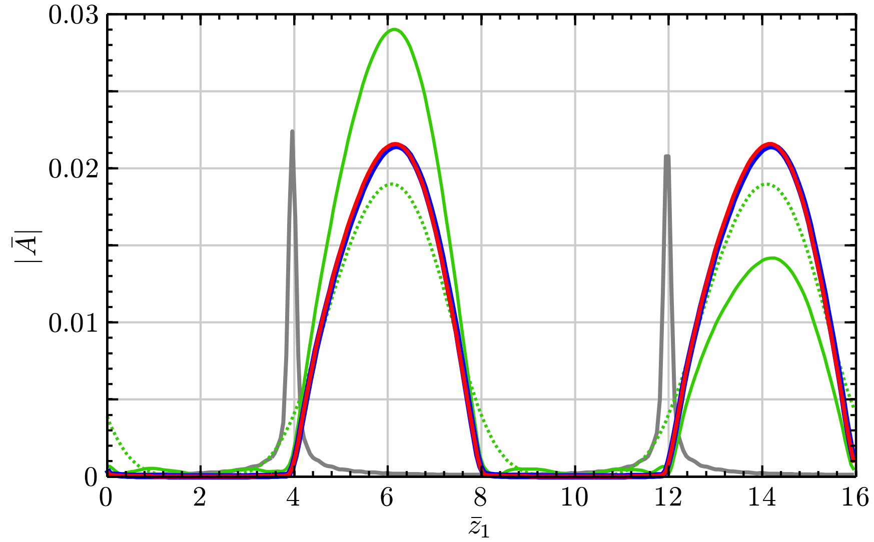

(b) $\bar{z} = 10$

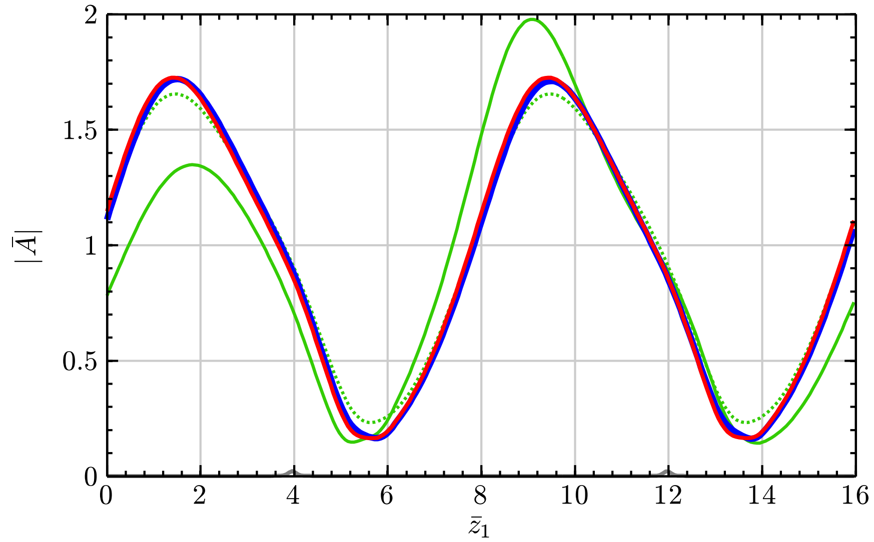

Figure 6.6: Radiation output profiles versus time at $\bar{z} = 4$ (a) and $\bar{z} = 14$ (b) resulting from broadband HHG seeding (cf. seed spectrum in Fig. 6.5). The simulation window has a width of $\lambda_d = 800\,\mathrm{nm}$ with periodic boundary conditions. The gray line represents the initial seed profile, with the characteristic two spikes per driving laser wavelength — this seed profile has been subtracted from the other results, leaving only radiation generated during FEL gain. Non-SVEA AURORA results are in blue while SVEA code results using a seed filtered with a Gaussian of RMS width $\omega_0/2\pi$ are in red. The solid (dotted) green line represents SVEA simulation results where the input seed was filtered with a 1.67× broader (10× narrower) Gaussian bandpass.

components that experience gain.

To demonstrate the sensitivity of the SVEA code results to details of the initial Gaussian spectral filter described in Sec. 6.3.4, we have also plotted results for filters with both a wider bandwidth ($1.67\,\omega_0/2\pi$) and a narrower one ($0.1\,\omega_0/2\pi$). For the case with a broader filter, the high intensities at low harmonics of the drive laser lead to an anomalous contribution to the seed that distorts the gain profile, as well as the total output power. In the narrowband case the seed's initially sharp temporal profile is smeared out, and thus also the profile of the output. This smearing does not, however, generally lead to large deviations in output power — especially at advanced stages of gain. This is because even the narrow filter's bandwidth still covers most of the FEL's gain bandwidth, and there is no interplay between different spectral components.

Based on the analytical arguments of Sec. 6.2 and the numerical analysis presented here, it appears possible to accurately simulate FELs seeded with broadband radiation using conventional single-frequency codes. This is in part a manifestation of the strong mode selection that takes place during FEL gain.

6.4.3 Evolution of HHG harmonics within FEL gain bandwidth

The full harmonic content of the seed is clearly also of importance in cases where the gain bandwidth encompasses more than one harmonic of the seeding laser. This may readily occur at high harmonic orders, where the relative difference in frequency between adjacent harmonics becomes small. We consider again the example from the previous subsection: seeding with the 255th harmonic of an 800 nm driving laser at 3.14 nm. The relative frequency separation between adjacent harmonics of that order is $(\omega_{255} - \omega_{257})/\omega_{255} = 1 - 255/257 \approx 7.8 \times 10^{-3}$, which is comparable to the gain bandwidth of a hypothetical FEL with Pierce parameter $\rho = 5 \times 10^{-3}$. In this regime the presence of harmonics other than the primary seeding one may affect the gain dynamics, particularly near saturation.

We simulate this scenario in AURORA with a full seed comprising all plateau harmonic orders from 21st to 281st at equal intensities, as illustrated in Fig. 6.5, and also with a filtered seed only containing the 255th harmonic. The results are presented in Fig. 6.7. Due to the imposed absence of initial electron bunch shot noise, for the case of the filtered seed the 255th harmonic is the only spectral component that achieves appreciable gain. The additional harmonics in the full seed undergo gain in accordance with the FEL's gain curve (inset). This difference in behavior leads to different saturation dynamics: the full seed enters saturation sooner, but takes longer to reach maximum intensity. Correspondingly, there is a period around $\bar{z} \approx 15$ where the filtered seed actually yields a marginally greater intensity than the full seed. The "tapered saturation" exhibited by the full seed is attributable to the fact that different modes within the gain curve saturate at different times, and in doing so inhibit the saturation of other modes. Thus, for example, in the case of the full seed the 255th harmonic can never reach the same saturation intensity as in the case of the filtered seed.

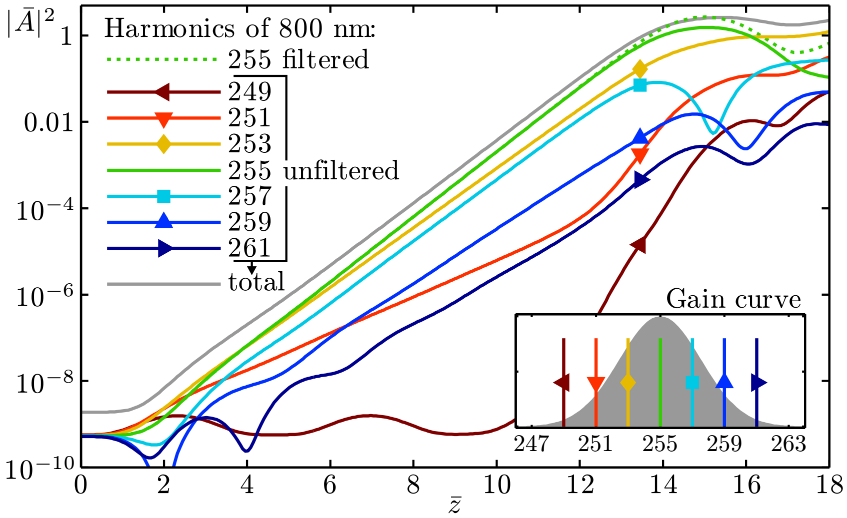

Figure 6.7: Evolution of HHG seed harmonics around the FEL-resonant 255th harmonic of an 800 nm driving laser, for the cases of a full and a filtered seed. The full seed (cf. Fig. 6.5) contains several harmonics that fall within the gain bandwidth ($\rho = 5 \times 10^{-3}$), whereas the filtered seed case contains only the 255th harmonic. Spectral components of FEL output resulting from the full seed are represented by solid lines, where the gray line is total normalized field due to harmonics 249 through 261; the output from the filtered seed is represented by the dotted green line. Inset: harmonic content of the full seed (vertical lines), set against the FEL gain bandwidth (gray).

As demonstrated in the previous subsection, even this case — where seed harmonics other than the primary seeding harmonic have an effect on FEL gain — can in fact be simulated in a conventional SVEA-compliant FEL code, as long as the seed field is appropriately binned. This relates to the arguments of section 6.2, where we showed that dropping the SVEA is not necessary even when the FEL is seeded with radiation that violates the approximation. Indeed, the SVEA only needs to be dropped when the radiation *emitted* by the FEL has a fast-varying profile — the most striking regime where this is the case is when the Pierce parameter approaches unity, i.e. we no longer have $\rho \ll 1$. This corresponds to very broadband emission, as well as significant amplification on the scale of a single undulator period. Initial studies in this regime suggest that working outside the SVEA is indeed necessary in order to accurately simulate it.

These results consider a single spatial dimension. While transverse effects such as radiation diffraction will affect the FEL's overall characteristics, it is unlikely there would be an interplay between these effects and the temporal and spectral characteristics of FEL gain, such that wideband effects are more important in 3-D than 1-D. In fact, the much greater amount of diffraction at long wavelengths suggests that there will be an automatic high passband filtering of HHG seeds, even further limiting non-SVEA effects.

6.5 Coupling between longitudinal modes of bunch current and seed field

The final topic of this chapter focuses on a class of FEL interactions that, to the best of the author's knowledge, have so far been overlooked in literature. This is not surprising, as their effect under normal circumstances is minor, however they may — as we go on to show — have implications for future FELs driven by laser-accelerated electron bunches. These interactions may arise from the coupling between non-resonant frequency components of the electron bunch and the radiation seed that ultimately lead to resonant gain.

In Section 6.1 we derived an analytical 1-D model for the amplification of a radiation seed in an FEL driven by an electron beam with a flat longitudinal profile. Our result, Eq. (6.11), is in agreement with the one obtained by Krinsky in Ref. [206] (cf. Eq. (5.11) therein). Krinsky's work, however, also allows for the case of an electron bunch with an arbitrary longitudinal profile, and includes a treatment of SASE radiation emitted due to variations in the beam current. Here we summarize his main result, which is that in the linear regime the radiation field envelope, $A(\zeta,\theta)$, evolves as

$$\begin{aligned} A(\zeta,\theta) =& A_0(\theta-\zeta) \\ &+\int_{\theta-\zeta}^{\theta} \mathrm{d}\theta'\, g_\mathrm{r}(\zeta,\theta,\theta')A_0(\theta') \\ &+\frac{d_1 k_\mathrm{R}}{\gamma_\mathrm{R}}\int_{\theta-\zeta}^{\theta} \mathrm{d}\theta'\, g_\mathrm{e}(\zeta,\theta,\theta')D(\theta')e^{-i\theta'}, \end{aligned} \tag{6.29}$$

Where $d_1 = n_{e,0}e^2 a_\mathrm{u}/(2c\epsilon_0 m)$ is a scaling factor similar to the κ_M introduced in Sec. 6.1, with $n_{e,0}$ the peak electron line density. Note the absence of the $[JJ]$ factor from the expression for d_1, since here we consider a helical undulator for which $[JJ] = 1$. In Eq. (6.29) $A_0(\theta)$ represents the initial seed field at $\zeta = 0$, while $D(\theta)$ is the electron bunch profile, normalized to 1 so that $n_e(\theta) = D(\theta)n_{e,0}$. The first line of the equation therefore represents the slippage of the original seed field through the ponderomotive window; the second line represents the amplification of that field through the interaction with the electron bunch; and the third line represents SASE radiation due to the electron bunch profile itself. Correspondingly, g_r and g_e are the Green's functions for the coupling of the radiation seed and of the electron bunch profile to FEL emission, respectively. They can be represented as the following series expansions:

$$g_\mathrm{r}(\zeta,\theta,\theta') = \sum_{l=1}^{\infty}\left[i\alpha\int_{\theta'}^{\theta}\mathrm{d}\theta''\,D(\theta'')\right]^l \frac{(\zeta-\theta+\theta')^{2l-1}}{l!\,(2l-1)!}; \tag{6.30}$$

$$g_\mathrm{e}(\zeta,\theta,\theta') = \sum_{l=0}^{\infty}\left[i\alpha\int_{\theta'}^{\theta}\mathrm{d}\theta''\,D(\theta'')\right]^l \frac{(\zeta-\theta+\theta')^{2l}}{l!\,(2l)!}, \tag{6.31}$$

where $\alpha = (2\rho)^3$. The two functions exhibit a significant degree of similarity, representing the fact that the process of radiation growth due to amplification of a seed field is similar to that due to amplification of noise (or modulations) in the bunch profile. In this section we focus on the coupling of seed radiation through g_r. Note that the expression for g_r above is

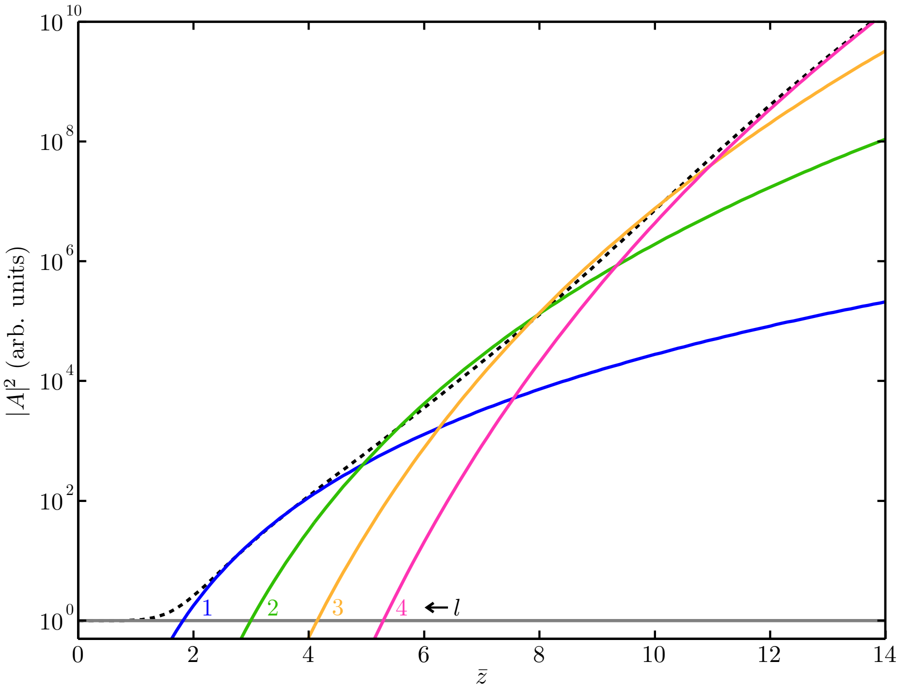

Figure 6.8: Contributions to resonant FEL gain by the first four terms in the summation in equation (6.32). The horizontal gray line is the initial seed intensity. The dotted black line is the sum of the terms — since consecutive terms are π out of phase, the amplitudes of individual terms are occasionally greater than the total. The horizontal axis is normalized distance along the undulator, $\bar{z} = 2\rho\zeta$.

equivalent to the one we obtained in Section 6.1, as presented in Eq. (6.11), for the case of a flat electron beam, $D(\theta) = 1$.

We proceed by taking a closer look at the second line of equation (6.29). To simplify our analysis, we can bring the summation in g_R outside the integral over θ',

$$A_\mathrm{r}(\zeta,\theta) = \sum_{l=1}^{\infty} \left\{ \int_{\theta-\zeta}^{\theta} \mathrm{d}\theta' \left[i\alpha \int_{\theta'}^{\theta} \mathrm{d}\theta'' D(\theta'') \right]^l \frac{(\zeta-\theta+\theta')^{2l-1}}{l!\,(2l-1)!} A_0(\theta') \right\}, \tag{6.32}$$

and consider the integration for successive values of l. Figure 6.8 illustrates the contribution of these different summation terms to conventional FEL gain driven by a flat electron beam and a resonant seed. Furthermore, we will represent the radiation seed envelope and the bunch profile in terms of their Fourier transforms, $\tilde{D}(\nu) = \mathscr{F}[D(\theta)]$, $\tilde{A}_0(\nu) = \mathscr{F}[A_0(\theta)]$, in order to aid the representation of longitudinal modes. The real-space functions are therefore

$$D(\theta) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathrm{d}\nu \, \tilde{D}(\nu) e^{i\theta\nu};$$

$$A_0(\theta) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathrm{d}\nu \, \tilde{A}_0(\nu) e^{i\theta\nu}.$$

Thus, we can reformulate equation (6.32) as

$$A_{\mathrm{r}}(\zeta,\theta)=\sum_{l=1}^{\infty}\int_{\theta-\zeta}^{\theta}\mathrm{d}\theta'\left(\frac{1}{2\pi}\int\mathrm{d}\nu_{\mathrm{r}}\,\tilde{A}_{0}(\nu_{\mathrm{r}})e^{i\theta'\nu_{\mathrm{r}}}\right)$$
$$\left[i\alpha\int_{\theta'}^{\theta}\mathrm{d}\theta''\left(\frac{1}{2\pi}\int\mathrm{d}\nu_{l}\,\tilde{D}(\nu_{l})e^{i\theta''\nu_{l}}\right)\right]^{l}\frac{(\zeta-\theta+\theta')^{2l-1}}{l!\,(2l-1)!},$$

where ν_{r} and ν_{l} are the frequency variables corresponding to the seed field and the electron bunch, respectively. The $-\infty$ to $+\infty$ limits of integration over $\nu_{\mathrm{r},l}$ have been omitted for clarity. Taking only the first term in the series, $l=1$, rearranging the order of integration, and carrying out the integral over θ'', we obtain

$$A_{\mathrm{r},1}(\zeta,\theta)=\frac{1}{4\pi^{2}}\iint\mathrm{d}\nu_{\mathrm{r}}\,\mathrm{d}\nu_{1}\,\tilde{A}_{0}(\nu_{\mathrm{r}})\tilde{D}(\nu_{1})g_{\mathrm{re},1}(\zeta,\theta;\nu_{\mathrm{r}},\nu_{1}),\tag{6.33}$$

where

$$g_{\mathrm{re},1}(\zeta,\theta;\nu_{\mathrm{r}},\nu_{1})=i\alpha\int_{\theta-\zeta}^{\theta}\mathrm{d}\theta'\,e^{i\theta'\nu_{\mathrm{r}}}\,\frac{e^{i\theta\nu_{1}}-e^{i\theta'\nu_{1}}}{i\nu_{1}}(\zeta-\theta+\theta')$$

is the first-order Green's function for the coupling between longitudinal modes of the electron bunch and the radiation. Carrying out the integral over θ'', we obtain the expression

$$g_{\mathrm{re},1}(\zeta,\theta;\nu_{\mathrm{r}},\nu_{1})=\alpha\frac{e^{i\theta(\nu_{\mathrm{r}}+\nu_{1})}}{\nu_{1}}\left[\frac{1}{\nu_{\mathrm{r}}^{2}}\left(1-e^{i\theta\nu_{1}}\right)-\frac{1}{(\nu_{\mathrm{r}}+\nu_{1})^{2}}\left(1-e^{-i\zeta(\nu_{\mathrm{r}}+\nu_{1})}\right)-i\zeta\left(\frac{1}{\nu_{\mathrm{r}}}-\frac{1}{\nu_{\mathrm{r}}+\nu_{1}}\right)\right].\tag{6.34}$$

This function only leads to significant growth for $\nu_{1}=-\nu_{\mathrm{r}}$. For the specific case of $\nu_{1}=\nu_{\mathrm{r}}=0$, i.e. a flat electron bunch and a resonant seed, it evaluates to

$$g_{\mathrm{re},1}(\zeta,\theta;0,0)=\alpha^{3}\frac{i\zeta^{3}}{6},\tag{6.35}$$

represented by the blue line in Fig. 6.8. This constitutes a contribution to the characteristic exponential dependence of FEL output on distance along the undulator, z, that in this case grows as z^{3}.

Notably, though, growth also occurs in the general case of $\nu_{1}=-\nu_{\mathrm{r}}\neq 0$, where the seed is not resonant and the bunch has a modulation at a frequency equal to the detuning of the seed from resonance. In that case the coupling has the following dependence on ζ:

$$g_{\mathrm{re},1}(\zeta,\theta;\nu_{\mathrm{r}},-\nu_{\mathrm{r}})=\alpha^{3}\left[-\frac{1}{\nu_{\mathrm{r}}^{3}}\left(1-e^{-i\zeta\nu_{\mathrm{r}}}\right)+i\frac{\zeta}{\nu_{\mathrm{r}}^{2}}+\frac{\zeta^{2}}{2\nu_{\mathrm{r}}}\right].\tag{6.36}$$

In this case output grows as $\sim z^{2}$, slower than the resonant growth with z^{3}, though the multiplier can be large for small detuning ν_{r}. This contribution will be insignificant for most conventional scenarios, however it can have a noticeable effect (at the few-percent level) in cases where the electron bunch is strongly modulated and the seed has off-resonant frequency components.

One such scenario may arise when an FEL is driven by electron bunches accelerated in a laser-wakefield accelerator that is itself driven by a laser of wavelength λ_{d}, where a

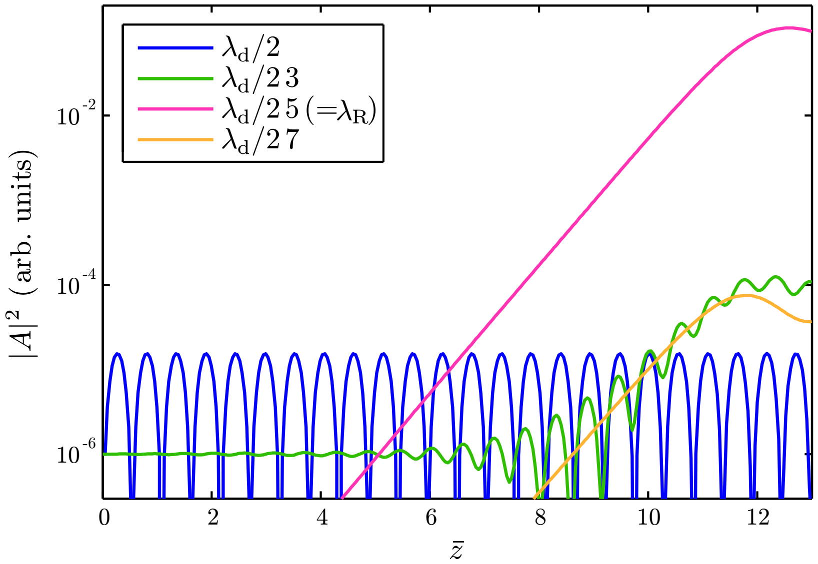

Figure 6.9: FEL harmonic growth for a resonant wavelength equal to the 25th harmonic of a $\lambda_d = 800\,\text{nm}$ drive laser, $\lambda_R = \lambda_d/25 = 32\,\text{nm}$. The FEL is seeded with only the 23th harmonic of the drive laser (green line), and the electron bunch is modulated at the 2nd harmonic (leading to oscillating emission at that harmonic; blue line). Despite the lack of seeding — either directly or through shot noise — at the resonant wavelength, exponential gain occurs due to coupling between the off-resonant radiation seed and the longitudinal bunch modulations.

laser of the same wavelength is used to drive the HHG process used to provide the FEL seed. In Section 5.5 we saw indication that laser-accelerated electron bunches may have longitudinal modulations at approximately half the laser wavelength, $\lambda_d/2$. For a resonant FEL wavelength of λ_R this corresponds to a normalized frequency of $\nu_1 = 2\lambda_R/\lambda_d$. At the same time, the frequency spacing between adjacent HHG harmonics, normalized to the resonant FEL frequency, is also $\nu_r = 2\lambda_R/\lambda_d$.

We simulate this scenario using the code AURORA described in Sec. 6.3. We take $\lambda_d = 800\,\text{nm}$ and $\lambda_R = 32\,\text{nm}$, i.e., the 25th harmonic. The simulation parameters are as given in Table 6.1, except for a lower peak current of $I_p = 6\,\text{kA}$, corresponding to $\rho = 3.4 \times 10^{-3}$. For purposes of illustration, we assume a maximally modulated electron bunch — $D(\theta) = [1 + \cos(\nu_1\theta + \phi)]/2$, where ϕ is a constant phase offset — and shot-noise microbunching is turned off. We consider the case of a seed containing all odd harmonics of the drive laser at equal intensities, as well as the case where only one non-resonant harmonic at $\nu_r = -2\lambda_R/\lambda_d$ is present.

The case of seeding with only one non-resonant harmonic, the 23rd, is illustrated in Figure 6.9. We observe that the bunch modulation induces oscillating emission at the second harmonic of the drive laser (blue line), and couples to the 23rd-harmonic seed (green line), leading to exponential gain in the 25rd harmonic (the resonant FEL wavelength; magenta line). At larger $\bar{z}$ we also observe a higher-order coupling of the modulated bunch current to the seed, which leads to gain in the 27th harmonic. We note that no significant gain is observed at any other harmonic. Additionally, there is no gain at the (resonant) 25th, nor at the 27th harmonic, for the control cases where there is either no seed radiation, or the

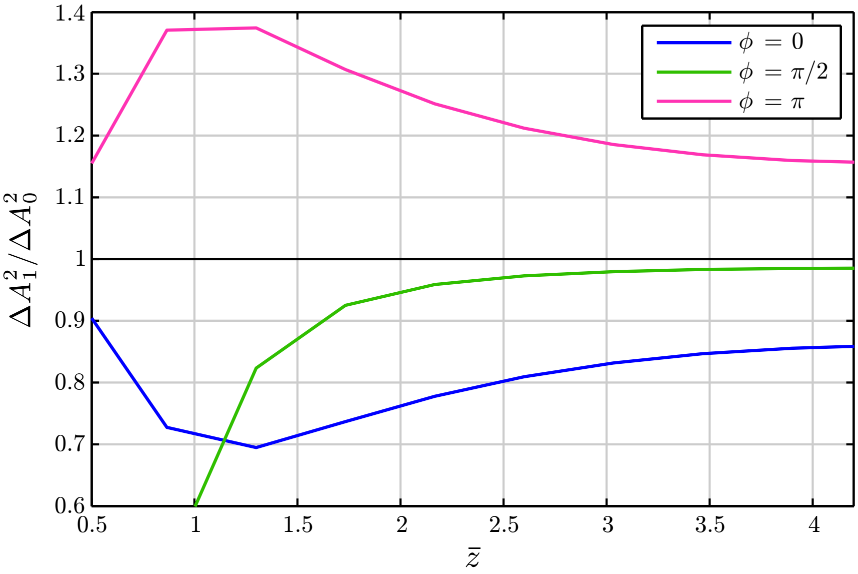

Figure 6.10: Ratio between gain in the fundamental for the case of an FEL seeded with all odd harmonics of an $\lambda_\mathrm{d} = 800\,\mathrm{nm}$ drive laser, ΔA_1^2, and gain in the fundamental of an FEL seeded with only the resonant 25^th harmonic, $\lambda_\mathrm{R} = 32\,\mathrm{nm}$, ΔA_0^2. The driving electron bunch is modulated at $\lambda_\mathrm{d}/2 = 400\,\mathrm{nm}$, thus allowing coupling of the non-resonant harmonics to the resonant, and leading to the observed discrepancy. Different lines represent different phase relations between the seed harmonics and the modulations on the bunch. Sharp bends in the lines are due to the limited number of data points used in the plot, and are not features of the system.

electron beam is not modulated.

Following up on this, Figure 6.10 illustrates the discrepancy in emission in the fundamental induced by the presence of additional harmonics, the effect of which is mediated by the presence of modulations on the electron bunch. The "full seed" case here (ΔA_1^2) is seeded with all odd harmonics of λ_d up to the 35^th at equal levels. Clearly, for the present case of an electron bunch that is maximally modulated at 400 nm, the discrepancy at the fundamental can be rather large (30–40 %) at initial stages of gain — however it becomes less significant later on, when regular gain sets in and the stronger resonant coupling, Eq. (6.35), becomes more important than the non-resonant, Eq. (6.36).

While the present analysis only deals with first-order couplings, i.e. those arising from the $l = 1$ term in the series of Eq. (6.32), higher-order ones may arise as well, where resonant gain can occur due to the coupling of one frequency component of the seed to two or more longitudinal modes of the bunch ($\nu_{2,3,\ldots}$). Additionally, a similar analysis of the last line of equation (6.29) demonstrates that multiple (off-resonant) longitudinal modes of the bunch can couple to each other in order to also produce resonant gain. One key point, however, is that only a single mode of the seed can contribute to these couplings, since terms containing $A_0(\theta)$ to a power greater than 1 never occur in the expansion of Eq. (6.29). Therefore, the only way for an off-resonant seed to lead to resonant gain is via a coupling to electron bunch modulations.

Although the couplings discussed in this section are unlikely to lead to significant corrections in the projected performance of conventional future FELs, they do suggest the potential for tailoring the interaction between the longitudinal modes of the radiation seed and electron bunch. This is of particular interest in future facilities where a single drive laser may be used to generate the radiation seed, as well as accelerate (and modulate) the electron bunch.

Chapter 7

Conclusion

The present thesis discusses a number of theoretical and experimental aspects pertaining to the realization of a free-electron laser driven by electrons from a laser-wakefield accelerator. In this chapter we summarize its findings and assess the overall outlook for the field, as well as suggest further work towards achieving the above-mentioned goal.

7.1 Summary

In Chapter 1 we introduced light sources — in particular synchrotrons and free-electron lasers — and their capabilities, and motivated the desire to use electrons from a laser-wakefield to drive them. Namely, the high accelerating gradients achievable in LWFA allow the electrons to be accelerated to the target energy over a distance that is as much as three orders of magnitude less than the distance required by a conventional accelerator. This has the potential to dramatically decrease the size and cost of light source facilities, meaning they could become more commonplace, and more readily available to scientists. For the sake of brevity, we dub light sources driven by electrons from a laser-wakefield accelerator "compact light sources".

Although not yet proven in this role, laser-wakefield accelerators are anticipated to offer a number of particular advantages as light source drivers. Some of these stem from the fact that the length of the electron bunches corresponds to a few femtoseconds: to a large degree, this sets the length of the output radiation pulses[1], resulting in pulses that can be used to probe matter with exquisite temporal resolution; in addition, it means that even a modest bunch charge of a few picocoulombs leads to peak currents of several kiloampères, thus presenting the potential for operation in the free-electron laser regime. Moreover, since the electron bunches are intrinsically synchronized to the drive laser, a part of that pulse can be split off and either itself used in the light source target area, or it can be used for the generation of secondary radiation sources, such as high-harmonics. This allows pump-probe experiments with pulses across a broad range of the electromagnetic spectrum that are intrinsically synchronized at the femtosecond level.

---

[1] Though in the case of a free-electron laser additional effects come into play, see e.g. Sec. 3.2.

In order to treat the problem in more detail, in Chapter 2 we gave an overview of the key components of a compact light source. We began by covering the background theory of laser-wakefield acceleration and summarized the key experimental results from recent years, in order to produce a set of realistic electron bunch characteristics that can be recreated under current experimental conditions. Recently, at the more advanced laser facilities, bunches have been stably and reproducibly accelerated to energies of approximately 800 MeV with rms energy spreads down to about 1%. Bunch charges of around 20 pC are common, though there is a certain trade-off between the energy to which a bunch is accelerated and its charge. There have been few measurements of bunch length (see e.g. Chap. 5), though they have yielded values around 0.5 µm ($\sim$ 1.7 fs) rms. Emittance measurements are also rare, however a normalized emittance of 1.5 mm mrad appears to be achievable.

We outlined that in the near term, the energy range is expected to be extended to about 1.5 GeV (an energy that has already been demonstrated, though not with high reproducibility), with energy spreads ideally remaining at approximately the same absolute level, i.e. relative energy spread going down to 0.5% rms. The main drivers behind these trends are expected to be further development and implementation of controlled injection as well as separation of the injection and acceleration stages, with recent advances in these fields [67, 74, 78, 81] being consolidated and built upon in the near future. In addition, the practically inevitable improvements in laser performance and stability should translate into direct advances in LWFA performance. These developments are expected to also increase the achievable bunch charges to about 100 pC, and allow the normalized emittance to be maintained at 1.5 mm mrad even as the electron energy is increased. The presented overview highlighted the need for bunch diagnostics that are more accurate and more consistently applied across experiments — in particular where measurements of energy spread and bunch charge are concerned. The recent diagnostic measurements of bunch length and emittance were commended, since, as we saw further on in the thesis, they constitute key parameters when evaluating the potential operation of an FEL.

Chapter 2 also covered the elementary theory pertaining to bunch dynamics between the accelerator and the undulator components of the light source. Focusing systems were discussed, and in particular the possibility of using recently-developed miniature permanent magnet quadrupoles for focusing within a compact system. The final part of the chapter was dedicated to the theory of generation of undulator radiation, and a discussion of operation in the free-electron laser regime, including the key parameters required for attaining it. We concluded with a summary of modern undulator technology, which highlights hybrid undulators based on permanent magnets as the most suitable type to meet the requirements of a compact light source.

After this qualitative overview, in Chapter 3 we presented a collection of parameter optimization studies and simulations in order to assess the performance of future LWFA-driven free-electron lasers as well as incoherent undulator radiation sources. These were based around target wavelengths for operation of 5 nm and 0.5 nm for incoherent radiation and 32 nm and 2 nm for an FEL. For incoherent radiation, optimization studies were carried out based on varying electron energy and undulator gap for a selection of undulator periods, in order to obtain a high value for output energy whilst remaining within realistic parameter bounds. For operation at 5 nm, electrons of energy 700 MeV could be used with an undulator of period $\lambda_{\mathrm{u}} = 10\,\mathrm{mm}$ and gap $g = 2\,\mathrm{mm}$, to yield nominal output in the fundamental of 0.4 nJ (for a bunch of 100 pC charge). Increasing the energy to 1500 MeV and decreasing the period and gap to $\lambda_{\mathrm{u}} = 6\,\mathrm{mm}$ and $\lambda_{\mathrm{u}} = 1\,\mathrm{mm}$, respectively, would lead to 2.9 nJ of output at 0.5 nm. While this undulator gap may appear overly aggressive, in recent experiments an undulator with a gap of 1.2 mm was used to generate radiation from laser-accelerated electrons — leading to the detection of radiation on over 70 % of shots [38]. Via simulations with the SPECTRA code, we gave the predicted spectral and spatial characteristics of the radiation for electron bunches with realistic spread, emittance, and transverse focusing.

In the case of the free-electron laser, the 32 nm wavelength was selected in part because it would be amenable for seeding with the 25th harmonic of an 800 nm laser, and partly because it appears to be a realistic wavelength at which to achieve lasing with current bunch parameters. Performance scaling is far more complex in the case of an FEL due to the highly nonlinear interplay between some of the parameters. For the bulk of the work here we relied on the scaling laws constructed by Ming Xie [148], however we also empirically developed an addition to these that allows for the characterization of gain degradation in the case of electron bunches with a length close to or shorter than the FEL cooperation length. We selected a set of bunch parameters based on modest improvements over those currently achievable, as presented previously: an energy of 240.4 MeV with 0.5 % rms energy spread and a charge of 50 pC; an rms bunch length of 0.5 µm and a normalized emittance of 0.5 mm mrad, corresponding to a transverse size of 1 µm rms and a divergence of 1 mrad rms. It is hoped that by having a target energy reasonably far below the bounds of the attainable, achieving improvements in the other metrics such as bunch charge or emittance would be possible. The selected undulator parameters were a period of 8 mm and a gap of 1.2 mm, corresponding to an rms strength parameter of $a_{\mathrm{u}} = 0.878$. For these parameters, scaling laws predict the gain length to be $L_{\mathrm{G}} = 7.7\,\mathrm{cm}$ and the saturation length to be $L_{\mathrm{sat}} = 1.35\,\mathrm{m}$.

A series of 3-D time-dependent simulations of SASE gain were carried out in order to evaluate performance for these parameters, and to assess the effects of parameter variations. For the selected working point, as an average over 10 SASE simulations, saturation was observed to occur after about 1.6 m with a peak power of 0.1 GW. The total number of photons produced was of order 10^{11}, which together with the other radiation source characteristics corresponded to a peak brightness on the scale of 10^{26}–10^{27} photons /second /mm^2 /mrad^2 /0.1 % B.W.. The range is due to variations resulting from differences in initial shot noise in the SASE simulations.

Carrying out a range of simulations for other bunch parameters revealed some strong dependences, which were already hinted at by the scaling laws. For example, decreasing the bunch charge from 50 pC to 30 pC led to a significant drop in gain, and while saturation still appeared to occur around 1.6 m, it was at a peak power of about 1 MW — a decrease of two orders of magnitude. Increasing the energy spread to 0.7 % had a similar effect. Additionally, increasing or decreasing the emittance by just 0.25 mm mrad caused the saturation power to respectively go down or up by 1.5 orders of magnitude. The one parameter that did not appear to have a significant effect was the bunch length: while a shorter bunch length resulted in a higher peak current, the gain was degraded due to the fact the bunch length was close to the cooperation length. For our proposed parameters these two effects to some extent canceled out, though the benefit of shortening the bunch length and thus increasing the current the current did slightly outweigh the associated degradation in gain. Under other operating regimes, for example at shorter FEL wavelengths where the cooperation length is likely to be shorter, variations in the bunch length would have an effect of similar magnitude to that seen here as a result of variations in the bunch charge. Finally, we evaluated the effect of seeding with HHG radiation, and observed that even a modest seed (10 kW peak power) would lead to almost an order of magnitude increase in output level — though, maybe more importantly, it significantly decreases shot-to-shot fluctuations due to shot noise start-up. A 50 kW seed was seen to practically eliminate these fluctuations, and to bring the saturation point forward to $L_{\mathrm{G}} \approx 1\,\mathrm{m}$. Considering the fact that a portion of the laser that drives LWFA can be used for the production of HHG radiation that is intrinsically synchronized with the accelerated electron bunches, the prospect of seeding appears to be a natural choice for a compact FEL source.

Having observed the strong dependence of FEL performance on the various bunch parameters, in Chapter 4 we switched gear and turned to electron bunch diagnostics in an experimental context. These diagnostics were based on measurements of the transition radiation that is emitted when bunches of electrons (or other charged particles) cross a dielectric boundary. In the case of electrons from a laser-wakefield accelerator, we found the spatial and spectral characteristics of this radiation to contain a wealth of information about the bunches. A set of preliminary measurements at the MPI für Quantenoptik in Garching, Germany, highlighted the presence of high levels of coherence in the radiation emitted at visible wavelengths, indicating that the bunches possess fine structure at those scales. These preliminary measurements allowed us to evaluate suitable experimental setups for measuring the radiation, while eliminating light contamination from the bright drive laser. They also underscored the possibility of using CTR at visible wavelengths as a non-destructive bunch profile measurement — a topic we will further discuss below.

After these preliminary measurements, we focused on performing a single-shot broadband measurement of transition radiation, which had not been realized in the past. The primary purpose of this was to allow us to deduce the electron bunch length, which, while a crucial parameter, had so far evaded reliable measurement[2]. Detection over the spectral range 0.4–7 μm was carried out by simultaneously using three spectrometers, one of which was a custom-built terahertz spectrometer based on pyroelectric detectors. Nearly 3000 shots with data were taken, though not each of them had data from all three spectrometers. A number of gas cell length and plasma density scans were carried out, as well as scans of laser pulse length. Of these, the gas cell length scans proved the most informative.

As part of the data analysis process we developed an algorithm for longitudinal bunch profile reconstruction by phase retrieval on a Fourier spectrum[3], based on phase retrieval algorithms that are frequently used in fields such as x-ray crystallography and astronomy.

---

[2] The multi-shot CTR measurements by Lundh *et al.* [97] were a significant step towards this, though in any case they were performed long after planning for this experiment had begun.

[3] The measured CTR intensities correspond to the amplitudes of the Fourier transform of the bunch profile, however phase data is not available.

We demonstrated the ability of this algorithm to successfully retrieve synthetic bunch profiles from simulated CTR data, and established the bounds on its reliability in the context of the actual spectral range that we observe. This algorithm was applied in order to reconstruct the bunch profiles on all shots for which reliable spectral data was available from all three spectrometers. For successful acceleration of quasi-monoenergetic electron bunches — which was observed in roughly 99 % of cases — the full-width at half-maximum bunch lengths measured were mostly in the range 1.2–2.0 μm, corresponding to Gaussian fits of rms length 0.51–0.85 μm (1.7–2.8 fs).

The retrieved bunch lengths, together with other diagnostics such as bunch energy and charge, allowed us to build a reasonably complete picture of the evolution of the acceleration process for different plasma densities. In particular, this included identification of the onset of dephasing at higher densities, or of laser defocusing at lower densities. In addition, the CTR data often exhibited spectral modulations that indicated interference due to emission from multiple electron bunches, which the profile reconstruction algorithm was able to identify. By analyzing the occurrence of additional bunches — and the distance by which they were separated from the main one, as well as their relative sizes — we were able to identify signatures of acceleration not just in the second bucket of the wake, but also the acceleration of bunches in a wake that appeared to be driven by the main bunch after the laser had defocused and the laser-driven wake had correspondingly disintegrated. This regime would be a first step towards realizing the “hybrid” laser-plasma accelerator proposed by Hidding *et al.* [197]. A final observation we made based on this dataset is that fine structure in the electron bunch appears to often occur in the wavelength range 400–600 nm, with a dependence on acceleration length that suggests it is due to interaction between the bunch and the driving laser.

After this experimental detour, in Chapter 6 we returned to free-electron laser theory in order to discuss the simulation of seeding with broadband radiation, in particular seeds obtained via high-harmonic generation. While this is a topic that appears especially pertinent to potential LWFA-driven light sources, the idea (and implementation) of FEL seeding is gaining momentum at many short-wavelength conventional FELs. Since HHG seeds often have temporal profiles that vary on the attosecond scale, they violate the slowly-varying envelope approximation that underpins the dynamical equations in most conventional FEL codes. We set out to evaluate the validity of this approximation in this context. Based on a number of preliminary theoretical arguments, we concluded that there are no fundamental reasons why conventional FEL codes should not be able to adequately simulate seeding with broadband seeds. We then presented a number of tools for benchmarking, including a linearized theoretical 1-D model for gain in a seeded FEL, for the case of an infinite electron bunch, as well as the non-SVEA non-undulator-period-averaged simulation code AURORA.

Benchmarks using a conventional SVEA code, the theoretical model, and AURORA indicated that for the simple case of seeding with an ultra-short spike of radiation, there are no discrepancies between the results — until saturation sets in, at which point the linear model is no longer correct, although the two codes still give identical results. One of the keys to achieving this level of correspondence appeared to be the application of a filtering and binning procedure to the full broadband seed, in order to make it amenable for use as an input in the conventional SVEA code. We went on to consider the more sophisticated case of seeding with a full idealized HHG seed, in which case we demonstrated how a discrepancy between the results from the SVEA and non-SVEA codes may arise for different initial filters on the broadband seed. In the case of seeding with a full HHG seed, we also considered the effect of filtering out all but the central resonant harmonic of the seed, for the regime where the gain bandwidth covers multiple harmonics. Observing the growth of the harmonics in the filtered and unfiltered cases, we concluded that filtering allows the central harmonic to reach a somewhat higher saturation level.

In the context of seeding with broadband seeds, we also considered the interplay between the spectral components present in seed radiation and longitudinal density modulations in the bunch. Using a linearized theoretical model similar to the one developed at the beginning of Chapter 6, we identify cases where non-resonant components of the seed can couple to non-resonant density modulations in the bunch and lead to resonant FEL gain. We went on to demonstrate this through simulations in AURORA, and evaluated the strength of the coupling with respect to the “regular” FEL coupling for the case of a resonant seed and an electron bunch with a flat profile, which for common FEL parameters can be of order 10 %, provided the electron bunch is maximally modulated. While this is unlikely to have a notable effect in most cases — largely because such large modulations are seldom seen in practice — we pointed out the intriguing perspective of tailoring the bunch profile in a LWFA through interactions with the drive laser during acceleration, in order to induce desired modifications to the FEL interaction upon coupling of the bunch modulations with seed radiation.

7.2 Further work

The material presented in this thesis highlights certain areas with particular scope for further work. Here we give an overview of the most important ones, from the point of view of the author.

With respect to evaluating the performance of FELs driven by laser-accelerated bunches, there are a number of factors that our present simulations did not take into account. One of these is the effect of variations in the transverse position and pointing of the bunches, which would be non-negligible for laser-driven acceleration. Their effects are expected to manifest in a manner similar to the variations in transverse emittance, and can by that token be significant. On the other hand, high reproducibility would not be the primary concern[4] for a proof-of-principle FEL experiment — rather it would be a matter of having some fraction of the bunches propagate sufficiently close to the undulator axis to achieve lasing. Moreover, for a proper study of gain degradation due to pointing fluctuations, the actual magnetic field profile (in 3-D) of the undulator used in the experiments would need to be known, as well as that of the quadrupoles used in the focusing system. Correspondingly, studies of the effect of errors in these field profiles would be appropriate.

---

[4] Of course, sufficiently large pointing fluctuations could lead to undulator damage, in which case it would be a concern.

There are a number of aspects of the propagation of the electron bunch itself that deserve further attention. One of these is the effect of space-charge forces that — while we dismissed based on theoretical studies — should be investigated further. In particular, experimental studies of the bunch size evolution (in all three dimensions) directly after the accelerator exit are required in order to settle any remaining doubts. A related effect is the generation of wakefields [224] that may result from propagation of the bunch through a narrow-gap undulator, and a study of their potential to degrade gain would be appropriate.

Most of the above-mentioned aspects have already received substantial attention from the conventional accelerator community, as their effect on the operation of FELs driven by conventional accelerators is often comparable. This highlights an overriding theme concerning the proposed operation of an FEL by a laser-wakefield accelerator: increased and enhanced engagement of the LWFA community with the conventional accelerator community is paramount to the success of such a project. Fortunately, the interactions between the two communities have intensified in recent years, and have been overwhelmingly constructive — with many researchers either transferring between the two fields, or working at the (increasingly blurred) interface between them.

As laser-wakefield accelerators mature, one field within conventional accelerator physics that is becoming especially relevant is bunch diagnostics. In particular, the development of reliable and consistent methods for measuring key bunch parameters such as energy, charge and emittance to high precision is crucial. While in the early days the laser-acceleration of quasi-monoenergetic electron bunches to any energy was a breakthrough in its own right, these bunches could only ever be relevant in real-world applications if we are able to characterize them sufficiently well — which would in turn be a step towards consistently reproducing the desired characteristics.

Nevertheless, one diagnostic that does not have a ready-made equivalent in conventional accelerator physics is a diagnostic for bunch lengths of just a few femtoseconds, such as those produced by laser-wakefield accelerators. In the present thesis we demonstrated the ability to reconstruct such short longitudinal profiles based on broadband measurements of coherent transition radiation. This appears to be a particularly fruitful field of work that deserves further pursuit. One aspect of this work could involve improvements and better characterization of the bunch profile reconstruction algorithm, for instance in order to better establish the reliability with which the algorithm can determine the "correct" profile from a given spectrum.

Based on our experience, it should be possible to develop a production-quality "black box" CTR diagnostic setup, which can be seamlessly implemented into a LWFA experiment in order to provide real-time longitudinal bunch profile information, with minimum or no disruption to the bunch itself. As we saw in Chapter 5, a relatively simple initial analysis of the CTR spectra can uncover a wealth of information about both the bunch and the acceleration process, and further studies are only likely to reveal more.

Apart from obtaining longitudinal profile information from the CTR spectrum, the near-field transverse profile of CTR emission at visible wavelengths appears to contain information about the transverse profile of the bunch itself, as shown in Section 4.4.1. While our

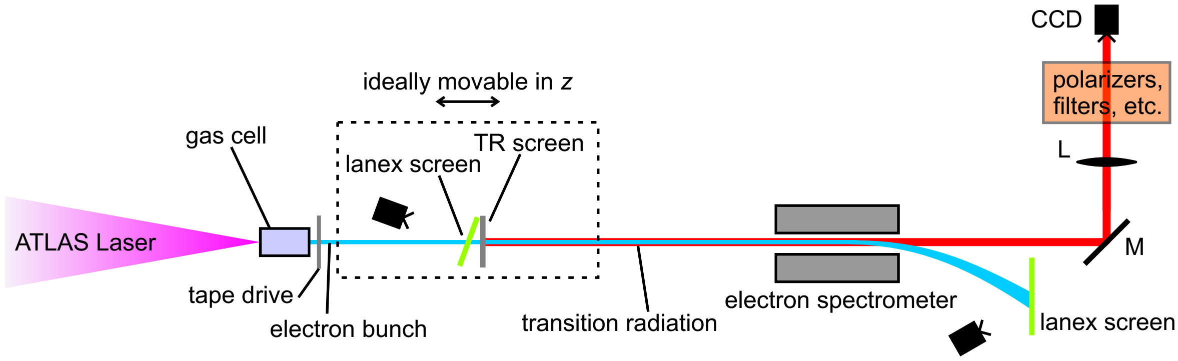

Figure 7.1: Schematic of a proposed experiment to evaluate the viability of a CTR-based diagnostic for the transverse bunch profile. The laser beam, entering from the left, accelerates an electron bunch in the gas cell. The bunch (light blue path) passes through a steel tape and through a lanex screen (green) and a transition radiation screen (gray), mounted back-to-back. The transition radiation profile is imaged in the near-field and processed by the algorithm outlined in Sec. 4.4.1, in order to obtain a transverse profile that is compared to the one recorded on the lanex screen. If necessary, the transition radiation may be filtered by a polarizer or a spectral filter.

results do not provide conclusive evidence on the reliability of such a transverse profile diagnostic, a follow-up experiment will aim to study this in the near future. The basic setup for this experiment is outlined in Figure 7.1, where the key difference with respect to previous experiments is that a lanex screen will be mounted back-to-back with the TR screen, in order to establish the degree of correlation between the bunch profile deduced from the transition radiation diagnostic and the actual transverse profile.

Finally, as others have pointed out [218, 219], if bunches of the lengths measured here are used to drive free-electron lasers at resonant wavelengths that are a significant fraction of the bunch length, the longitudinal bunch shape itself may lead to coherent spontaneous emission, which could even enhance the FEL start-up process. Further studies of effects in this context are warranted, and could be carried out with a new class of non-SVEA non-undulator-period-averaged codes that have emerged recently, such as AURORA, although 3-D simulations, e.g. using the code described in Ref. [221], may be more accurate. On a related note, coherent spontaneous emission may in some cases lead to unexpectedly high output levels, which may be mistaken for actual FEL gain. In order to eliminate any doubts, it would be advisable to include multiple diagnostics that could identify free-electron lasing, such as CTR-based diagnostics to measure the associated microbunching at the resonant wavelength [225–228].

With this, we conclude the present discussion of the realization of an FEL driven by laser-accelerated electrons. Intensifying work in this field by a number of research groups has the clear potential to turn such a system into reality in the coming years. This would bring us to yet another watershed moment in the explosive development of laser-wakefield accelerators: machines that would allow us even greater freedom to explore the extreme scales of our world.

Appendix A

Optimization of FEL parameters

In order to conveniently optimize the FEL working point, the author created a MATLAB-based graphical user interface for parameter optimization. The scalings employed by it are derived from the formulae presented in Sections 2.3.2, 2.3.3, 2.3.4 and 3.2.1. Notably, the 3-D Xie scaling and the bunch length scaling of Sec. 3.2.1 are implemented. The undulator is assumed to be a planar hybrid NdFeB undulator as described in Sec. 2.3.4.

A screenshot of the interface is pictured in Figure A.1. The user first selects a target operation wavelength. They can then select values for two parameters out of the undulator period, undulator gap and the electron energy, and allow the code to choose the third in a way that the target output wavelength is achieved. The undulator parameter a_u is automatically calculated for the given combination of λ_u and undulator gap. The β-functions of the electron beam in the undulator are either based on natural undulator focusing in y (see Sec. 2.3.3) and regular quadrupole focusing in x (so that the mean transverse size over the undulator length is minimized), or on a regular focusing lattice.

The selected normalized emittance of the electron bunch is used to obtain the geometrical emittance and the corresponding mean transverse bunch size. The charge and bunch length are used to obtain the peak current for a Gaussian bunch. These parameters, together with the electron energy spread, are fed into the scaling laws of Sec. 3.2.1 in order to evaluate the 1-D, 3-D, and effective gain lengths, and the saturation length. The values of the Xie scaling parameters η_d, η_ε, η_γ, as well as the bunch length scaling parameter, η_z, are displayed in order to give the user an idea about the factors impacting gain at the selected working point.

Contour plots of the variation of saturation length with undulator period, undulator gap, electron energy, energy spread, bunch charge, bunch length, and normalized emittance allow the user to visualize the dependence of FEL performance on these parameters. This allows for the selection of suitable values of the undulator parameters and the target electron energy, and for an evaluation of the dependence of FEL performance on bunch quality. The user is able to click on any point of interest on the contour plots in order to automatically retrieve the parameters corresponding to that point.

When a suitable working point has been selected, the code can automatically generate

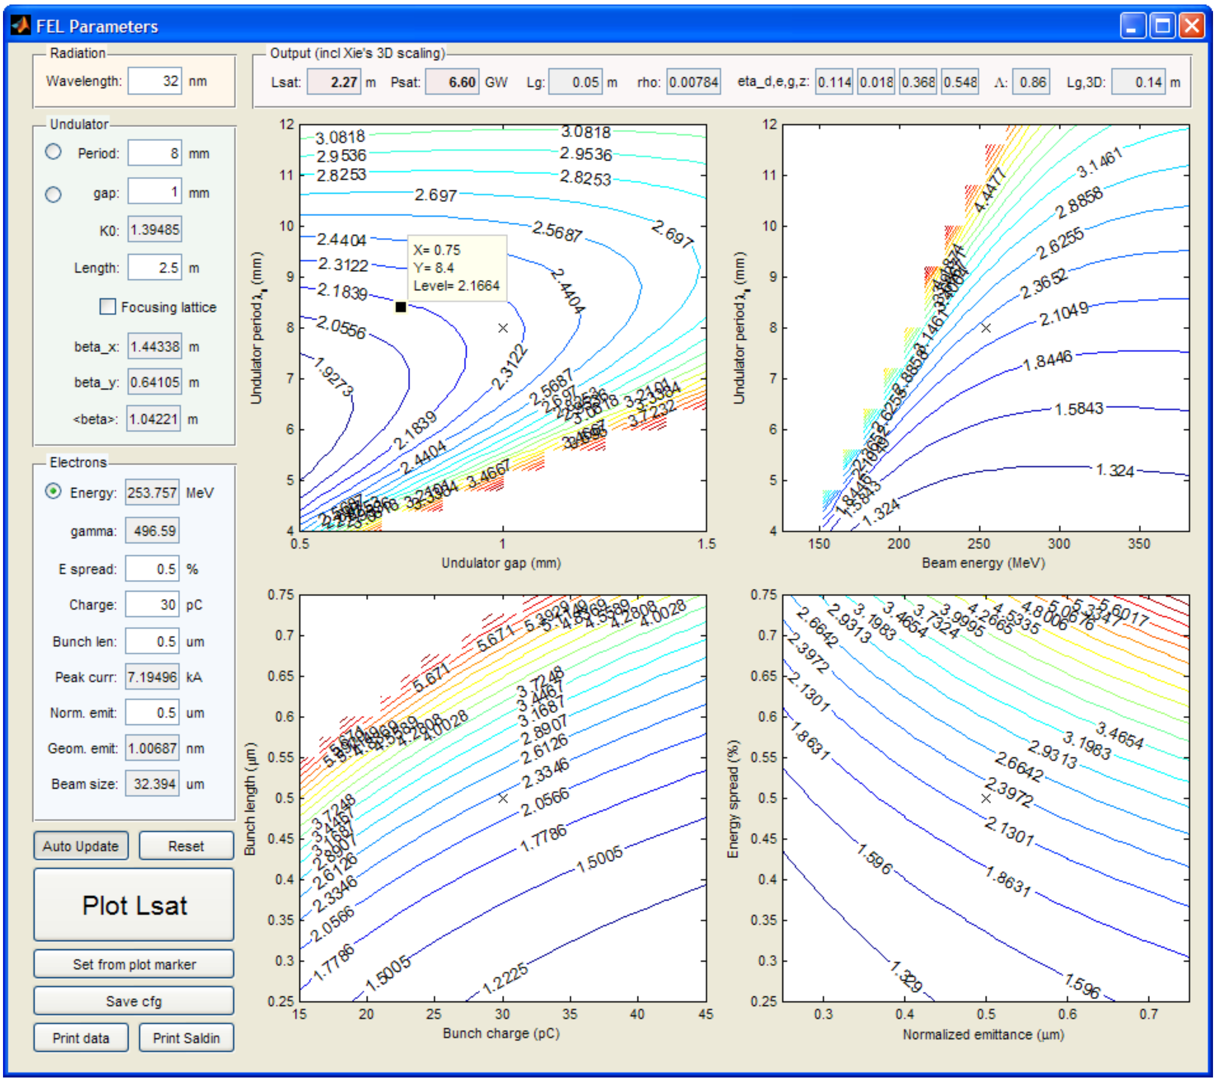

Radiation
Wavelength: 32 nm

Output (incl Xie's 3D scaling)
Lsat: 2.27 m
Psat: 6.60 GW
Lg: 0.05 m
rho: 0.00784
eta_d,e,g,z: 0.114 0.018 0.368 0.548
Λ: 0.86
Lg,3D: 0.14 m

Undulator
○ Period: 8 mm
○ gap: 1 mm
K0: 1.39485
Length: 2.5 m
☐ Focusing lattice
beta_x: 1.44338 m
beta_y: 0.64105 m
<beta>: 1.04221 m

Electrons
● Energy: 253.757 MeV
gamma: 496.59
E spread: 0.5 %
Charge: 30 pC
Bunch len: 0.5 um
Peak curr: 7.19496 kA
Norm. emit: 0.5 um
Geom. emit: 1.00687 nm
Beam size: 32.394 um

Figure A.1: Screenshot of FEL parameter optimization interface.

input files for steady-state and time-dependent GENESIS simulations based on the parameters, as well as input files for simulating the transverse bunch dynamics in MAD. Batch files for running the resource-intensive GENESIS simulations on a cluster are also automatically generated.

Appendix B

Conventional 1-D FEL Code Under the SVEA

A time-dependent 1-D FEL code can be used to quickly estimate FEL performance under a proposed set of parameters. Here we describe such a code, which was used at several points during the work described in this thesis. The equations underpinning it are equivalent to those used in the code PERSEO, described in Ref. [208], and it has a comparable feature set. Indeed, the present code has been benchmarked against PERSEO in the steady-state as well as the time-dependent regime. It is also equivalent to the code AURORA described in Sec. 6.3 — in the case where the longitudinal profiles of the radiation envelope and electron bunch vary slowly in time, as quantified in Sec. 6.2. In order to contrast the features of the present code with those of AURORA, throughout this thesis it is referred to as the "conventional SVEA code".

The code employs a simulation window in θ that spans L resonant wavelengths $\lambda_0 \equiv \lambda_{\mathrm{R}}$, with one "bin" per wavelength. For steady-state simulations, $L = 1$. The radiation field is represented by the envelope of the normalized magnetic vector potential, $a_0(z,\theta) = eA_0(z,\theta)/mc$, where the complex value of a_0 is recorded for each bin. In terms of this envelope, the full field is $A_x = A_0 e^{-i\omega_0 t}$, where $\omega_0 = 2\pi c/\lambda_0$. The corresponding electric field is $E_x = -\partial A_x/\partial t$. A fixed number, P, macroparticles are used to represent the electrons within each bin. Since the primary purpose of this code was to carry out SVEA simulations for comparison with equivalent AURORA simulations, the initial degree and phase of microbunching for each bin were obtained from the binned initial macroparticle distribution in AURORA. The latter, as described in Sec. 6.3, is generated using Poissonian statistics on the macroparticle weights.

After this initialization procedure, the macroparticle energies and positions, and radiation amplitude and phase, are propagated according to the following equations:

$$\frac{\mathrm{d}\theta_i}{\mathrm{d}z} = k_{\mathrm{u}} - \frac{k_0}{2\gamma_i^2}(1 + a_{\mathrm{u}}^2); \tag{B.1}$$

$$\frac{\mathrm{d}\gamma_i}{\mathrm{d}z} = -k_0 \frac{a_\mathrm{u}[JJ]}{\sqrt{2}\gamma} \Re\left[a_0 e^{i\theta_i}\right] \times G; \tag{B.2}$$

$$\frac{\mathrm{d}a_0}{\mathrm{d}z} = \frac{e^2}{\sqrt{2} m \epsilon_0 c^2} \frac{a_\mathrm{u}[JJ]}{k_0} \left\langle \frac{n_e}{\Sigma_r} \frac{e^{-i\theta_i}}{\gamma_i} \right\rangle_{\lambda_0} \times G. \tag{B.3}$$

Here subscripts i represent the macroparticle indices within the bin. In Eq. (B.3) a_0 and n_e are the values for the current bin, and the average in the angled brackets is over all macroparticles in the bin. G is a gain scaling factor that is nominally unity, but can be reduced in order to account for effects that may degrade gain, such as those described by Xie [148] and discussed in Sec. 3.2.1 here. All other parameters are as introduced in Sections 2.3.2 and 6.3, and $k_0 = 2\pi/\lambda_0$.

The above equations are integrated over one undulator period at a time, using a Runge-Kutta 4th order formula [229]. After each of these steps of size $\Delta z = \lambda_\mathrm{u}$ through the undulator, the radiation field values are shifted by one bin in the $+\theta$ direction in order to account for the slippage of the radiation over the electrons, as described in Sec. 2.3.2. A periodic grid is used, whereby radiation slipping out of the front of the window is re-inserted at the back. Finally, a graphical user interface allows monitoring of the output in real time.

Appendix C

FEL electron energy modulation due to a broadband seed

In a free-electron laser, a radiation seed co-propagating with the electrons modulates their energy. Through the dependence of the evolution of electrons' longitudinal positions on their energy (Eq. (6.24)), this may lead to microbunching. Therefore, it is instructive to consider the modulation of electron energy caused by an arbitrary (broadband) seed, and how this seed may be reduced to a form that is amenable for use in a conventional SVEA code, yet has the same modulating effect.

From the relativistic Lorentz equation, the change in electron energy is

$$\frac{\mathrm{d}\gamma}{\mathrm{d}t} = \frac{e}{mc}\mathbf{E}\cdot\boldsymbol{\beta}, \tag{C.1}$$

where $\mathbf{E}$ is the electric field and $\boldsymbol{\beta}$ is the normalized electron velocity. We assume a seed field that is linearly polarized along x and a planar undulator in which the transverse normalized electron velocity is $\beta_x = (\sqrt{2}a_\mathrm{u}/\gamma)\cos(k_\mathrm{u}z)$, where the parameters are as in Section 2.3.2. For this simplified treatment we neglect higher-order terms in the longitudinal motion of the electrons down the undulator, thus setting $z = \beta_0 ct$, where $\beta_0 = 1-(1+a_\mathrm{u}^2)/(2\gamma_\mathrm{R}^2)$ is the mean normalized velocity along z. In terms of the angular variables $\zeta = k_\mathrm{u}z$ and $\varphi = k_\mathrm{R}(z-ct)$, such that $\theta = \zeta+\varphi$, the energy modulation becomes

$$\frac{\mathrm{d}\gamma}{\mathrm{d}\zeta} = \frac{\sqrt{2}ea_\mathrm{u}}{m\beta_0c^2k_\mathrm{u}\gamma}E_x(\varphi)\cos(\zeta). \tag{C.2}$$

We consider an initially cold beam of electrons at the resonant energy γ_R. While evaluating the modulation in energy, i.e. change in γ, we assume $\gamma \equiv \gamma_\mathrm{R}$ on the r.h.s. of equation (C.2), and correspondingly assume the longitudinal locations of the electrons in θ remain constant. The change in γ for an electron at ponderomotive location θ, $\Delta\gamma(\theta) = \gamma(\theta)-\gamma_\mathrm{R}$, induced by the field $E_x(\varphi)$ is therefore

$$\Delta\gamma(\theta) = \eta\int \mathrm{d}\zeta\, E_x(\theta-\zeta)\cos(\zeta), \tag{C.3}$$

where the constants have been collected in $\eta = \sqrt{2}ea_\mathrm{u}/(m\beta_0c^2k_\mathrm{u}\gamma_\mathrm{R})$.

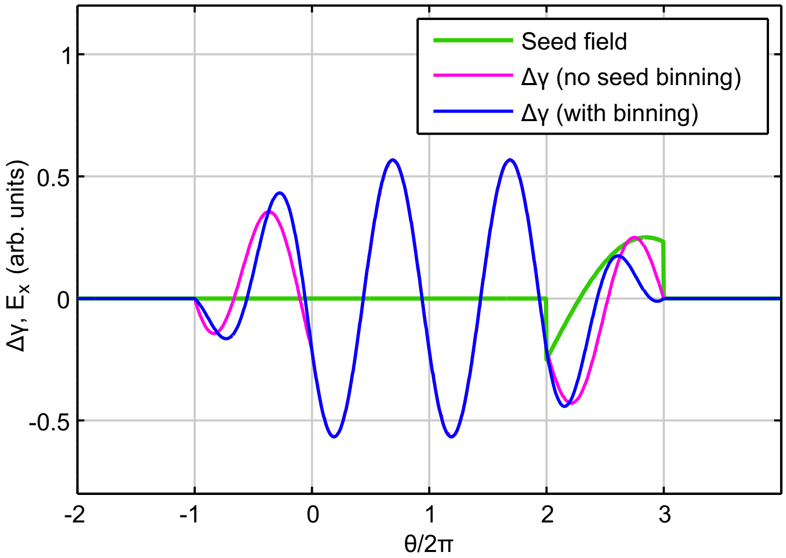

Figure C.1: Energy modulation due to $E_x(\varphi) = 0.5[\sin(5\varphi/17) - 0.5]$ on the interval $\varphi \in [0, 2\pi]$ integrated from $\zeta = -2\pi$ to $\zeta = 4\pi$. As expected, the full and binned seeds lead to identical modulation over $\theta \in [0, 4\pi]$, however not in regions of θ with which only part of the seed has interacted.

Let us assume that $E_x(\varphi)$ is zero everywhere except in the interval $\varphi \in [0, 2\pi]$ — in a conventional SVEA code this corresponds to a single radiation bin of width λ_{R}. In order to find the energy modulation at a fixed location θ due to this seed we therefore need to integrate over $\zeta \in [\theta - 2\pi, \theta]$. Expressing Eq. (C.3) as an integral over φ, we obtain

$$\begin{aligned}\Delta\gamma(\theta) &= \eta \int_0^{2\pi} \mathrm{d}\varphi E_x(\varphi) \cos(\theta - \varphi) \\ &= \eta\, \Re\left[\int_0^{2\pi} \mathrm{d}\varphi E_x(\varphi) e^{i(\theta-\varphi)}\right] \qquad \text{(C.4)} \\ &= \eta\, \Re\left[e^{i\theta} E_1\right], \qquad \text{(C.5)}\end{aligned}$$

where $E_1 = \int_0^{2\pi} \mathrm{d}\varphi E_x(\varphi) e^{-i\varphi}$ is the complex amplitude of the fundamental frequency component of the full seed E_x. Taking the real part of the entire integral rather than just the complex exponential in Eq. (C.4) is valid since E_x is a real-valued function. Equation (C.5) is precisely the method by which energy modulation is calculated in conventional FEL codes, where the radiation is resolved on a gird of spacing $\geq \lambda_{\mathrm{R}}$. Thus, it appears that if the full seed field over a resonant wavelength is converted to a single complex amplitude, E_1, it should lead to an equivalent energy modulation and therefore bunching in a conventional SVEA code.

Finally, a limiting feature of this binning approach is the fact that it gives the correct energy modulation only when the *entire* field, over $\varphi \in [0, 2\pi]$, has interacted with an electron at certain θ. The modulation differs for electrons that either started off within the seed field (and have only interacted with its tail), or for electrons over which the field is currently passing (which have only interacted with its head). This situation is illustrated in Figure C.1, where the binned field has induced the correct modulation on electrons in the ponderomotive range $\theta \in [0, 4\pi]$, but not on electrons in the ranges $\theta \in [-2\pi, 0]$ or $\theta \in [4\pi, 6\pi]$.

Bibliography

[1] T. C. Kriss and V. M. Kriss, “History of the operating microscope: From magnifying glass to microneurosurgery”, *Neurosurgery* **42**, 899–907 (1998). ⇒ doi:10.1097/00006123-199804000-00116

[2] D. B. Williams and C. B. Carter, *Transmission electron microscopy : a textbook for materials science : I : basics.*, Plenum, New York (1996). ⇒ online

[3] J. D. Watson and F. H. C. Crick, “Molecular structure of nucleic acids: A structure for deoxyribose nucleic acid”, *Nature (London)* **171**, 737–738 (1953). ⇒ doi:10.1038/171737a0

[4] J. P. Glusker and K. N. Trueblood, *Crystal Structure Analysis: A Primer*, IUCr Texts on Crystallography. Oxford University Press, third ed. (2010). ⇒ online

[5] W. I. F. David, K. Shankland, L. B. McCusker, and C. Baerlocher (eds.), *Structure Determination from Powder Diffraction Data*, IUCr Monographs on Crystallography. Oxford University Press (2002). ⇒ online

[6] J. Lipfert and S. Doniach, “Small-Angle X-Ray scattering from RNA, proteins, and protein complexes”, *Annual Review of Biophysics and Biomolecular Structure* **36**, 307–327 (2007). ⇒ doi:10.1146/annurev.biophys.36.040306.132655

[7] “European XFEL — research — possibilities”. ⇒ online

[8] C. Pellegrini and J. Stöhr, “X-ray free-electron lasers—principles, properties and applications”, *Nucl. Instrum. Meth. Phys. Res. A* **500**, 33 (2003). ⇒ doi:10.1016/s0168-9002(03)00739-3

[9] R. Neutze, R. Wouts, D. van der Spoel, E. Weckert, and J. Hajdu, “Potential for biomolecular imaging with femtosecond x-ray pulses”, *Nature (London)* **406**, 752–757 (2000). ⇒ doi:10.1038/35021099

[10] H. N. Chapman, *et al.*, “Femtosecond diffractive imaging with a soft-X-ray free-electron laser”, *Nat. Phys.* **2**, 839–843 (2006). ⇒ doi:10.1038/nphys461

[11] S. Doniach, “Studies of the Structure of Matter with Photons from an X-ray Free-Electron Laser”, *J. Synchrotron Rad.* **3**, 260–267 (1996). ⇒ doi:10.1107/S0909049596009739

[12] T. Pfeifer, C. Spielmann, and G. Gerber, “Femtosecond x-ray science”, *Rep. Prog. Phys.* **69**, 443–505 (2006). ⇒ doi:10.1088/0034-4885/69/2/R04

[13] J. Arthur *et al.*, “Linac Coherent Light Source (LCLS) conceptual design report”, *Tech. rep.*, SLAC, CA, U.S.A. (2002). ⇒ `online`

[14] P. Emma, *et al.*, “First lasing and operation of an ångstrom-wavelength free-electron laser”, *Nat. Photon.* **4**, 641–647 (2010). ⇒ `doi:10.1038/nphoton.2010.176`

[15] H. N. Chapman, *et al.*, “Femtosecond x-ray protein nanocrystallography”, *Nature (London)* **470**, 73–77 (2011). ⇒ `doi:10.1038/nature09750`

[16] M. M. Seibert, *et al.*, “Single mimivirus particles intercepted and imaged with an x-ray laser”, *Nature (London)* **470**, 78–81 (2011). ⇒ `doi:10.1038/nature09748`

[17] T. Tanaka and T. Shintake, “SPring-8 compact SASE source (SCSS) conceptual design report”, *Tech. rep.*, SPring-8, Hyogo, Japan (2005). ⇒ `online`

[18] T. Shintake, *et al.*, “Stable operation of a self-amplified spontaneous-emission free-electron laser in the extremely ultraviolet region”, *Phys. Rev. ST Accel. Beams* **12**, 070701 (2009). ⇒ `doi:10.1103/PhysRevSTAB.12.070701`

[19] M. Altarelli *et al.*, “The european X-Ray Free-Electron laser technical design report”, *Tech. rep.*, DESY, Hamburg, Germany (2007). ⇒ `online`

[20] W. Ackermann *et al.*, “Operation of a free-electron laser from the extreme ultraviolet to the water window”, *Nat. Photon.* **1**, 336–342 (2007). ⇒ `doi:10.1038/nphoton.2007.76`

[21] T. Tajima and J. M. Dawson, “Laser electron accelerator”, *Phys. Rev. Lett.* **43**, 267–270 (1979). ⇒ `doi:10.1103/PhysRevLett.43.267`

[22] T. Katsouleas, “Accelerator physics: Electrons hang ten on laser wake”, *Nature (London)* **431**, 515–516 (2004). ⇒ `doi:10.1038/431515a`

[23] P. Chen, J. M. Dawson, R. W. Huff, and T. Katsouleas, “Acceleration of electrons by the interaction of a bunched electron beam with a plasma”, *Phys. Rev. Lett.* **54**, 693–696 (1985). ⇒ `doi:10.1103/PhysRevLett.54.693`

[24] C. Joshi, W. B. Mori, T. Katsouleas, J. M. Dawson, J. M. Kindel, and D. W. Forslund, “Ultrahigh gradient particle acceleration by intense laser-driven plasma density waves”, *Nature (London)* **311**, 525–529 (1984). ⇒ `doi:10.1038/311525a0`

[25] C. E. Clayton, K. A. Marsh, A. Dyson, M. Everett, A. Lal, W. P. Leemans, R. Williams, and C. Joshi, “Ultrahigh-gradient acceleration of injected electrons by laser-excited relativistic electron plasma waves”, *Phys. Rev. Lett.* **70**, 37–40 (1993). ⇒ `doi:10.1103/PhysRevLett.70.37`

[26] C. Joshi, T. Tajima, J. M. Dawson, H. A. Baldis, and N. A. Ebrahim, “Forward raman instability and electron acceleration”, *Phys. Rev. Lett.* **47**, 1285–1288 (1981). ⇒ `doi:10.1103/PhysRevLett.47.1285`

[27] K. Nakajima, *et al.*, “Observation of ultrahigh gradient electron acceleration by a Self-Modulated intense short laser pulse”, *Phys. Rev. Lett.* **74**, 4428–4431 (1995). ⇒ `doi:10.1103/PhysRevLett.74.4428`

[28] C. A. Coverdale, C. B. Darrow, C. D. Decker, W. B. Mori, K. C. Tzeng, K. A. Marsh, C. E. Clayton, and C. Joshi, "Propagation of intense subpicosecond laser pulses through underdense plasmas", *Phys. Rev. Lett.* **74**, 4659–4662 (1995). ⇒ `doi:10.1103/PhysRevLett.74.4659`

[29] C. Joshi, "The development of laser- and beam-driven plasma accelerators as an experimental field", *Phys. Plasmas* **14**, 055501 (2007). ⇒ `doi:10.1063/1.2721965`

[30] D. Strickland and G. Mourou, "Compression of amplified chirped optical pulses", *Opt. Commun.* **56**, 219–221 (1985). ⇒ `doi:10.1016/0030-4018(85)90120-8`

[31] G. Mourou and D. Umstadter, "Development and applications of compact high-intensity lasers", *Phys. Fluids B* **4**, 2315–2325 (1992). ⇒ `doi:10.1063/1.860202`

[32] V. Yanovsky, *et al.*, "Ultra-high intensity- 300-TW laser at 0.1 hz repetition rate", *Opt. Express* **16**, 2109–2114 (2008). ⇒ `doi:10.1364/OE.16.002109`

[33] S. P. D. Mangles, *et al.*, "Monoenergetic beams of relativistic electrons from intense laser-plasma interactions", *Nature (London)* **431**, 535–538 (2004). ⇒ `doi:10.1038/nature02939`

[34] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J. P. Rousseau, F. Burgy, and V. Malka, "A laser-plasma accelerator producing monoenergetic electron beams", *Nature (London)* **431**, 541–544 (2004). ⇒ `doi:10.1038/nature02963`

[35] C. G. R. Geddes, C. Tóth, J. van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans, "High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding", *Nature (London)* **431**, 538–541 (2004). ⇒ `doi:10.1038/nature02900`

[36] W. P. Leemans, B. Nagler, A. J. Gonsalves, C. Tóth, K. Nakamura, C. G. R. Geddes, E. Esarey, C. B. Schroeder, and S. M. Hooker, "GeV electron beams from a centimetre-scale accelerator", *Nat. Phys.* **2**, 696–699 (2006). ⇒ `doi:10.1038/nphys418`

[37] H. P. Schlenvoigt, *et al.*, "A compact synchrotron radiation source driven by a laser-plasma wakefield accelerator", *Nat. Phys.* **4**, 130–133 (2008). ⇒ `doi:10.1038/nphys811`

[38] M. Fuchs, *et al.*, "Laser-driven soft-X-ray undulator source", *Nat. Phys.* **5**, 826–829 (2009). ⇒ `doi:10.1038/nphys1404`

[39] F. J. Grüner, *et al.*, "Design considerations for table-top, laser-based VUV and X-ray free electron lasers", *App. Phys. B* **86**, 431–435 (2007). ⇒ `doi:10.1007/s00340-006-2565-7`

[40] K. Nakajima, "Compact x-ray sources: Towards a table-top free-electron laser", *Nat. Phys.* **4**, 92–93 (2008). ⇒ `doi:10.1038/nphys846`

[41] T. P. A. Ibbotson, *et al.*, "Laser-wakefield acceleration of electron beams in a low density plasma channel", *Phys. Rev. ST Accel. Beams* **13**, 031301 (2010). ⇒ `doi:10.1103/PhysRevSTAB.13.031301`

[42] C. Thomas, G. Rehm, R. Bartolini, N. Delerue, and S. Bajlekov, "Single shot emittance measurement from beam size measurement in a drift section", In *Proceedings of the 1st International Particle Accelerator Conference*, p. 1167. Kyoto, Japan (2010). ⇒ `online`

[43] M. Pittman, S. Ferré, J. P. Rousseau, L. Notebaert, J. P. Chambaret, and G. Chériaux, "Design and characterization of a near-diffraction-limited femtosecond 100-TW 10-Hz high-intensity laser system", *App. Phys. B* **74**, 529–535 (2002). ⇒ `doi:10.1007/s003400200838`

[44] C. Tóth, J. Faure, J. van Tilborg, C. G. R. Geddes, C. B. Schroeder, E. Esarey, and W. P. Leemans, "Tuning of laser pulse shapes in grating-based compressors for optimal electron acceleration in plasmas", *Opt. Lett.* **28**, 1823–1825 (2003). ⇒ `doi:10.1364/OL.28.001823`

[45] W. P. Leemans, *et al.*, "The BErkeley lab laser accelerator (BELLA): A 10 GeV laser plasma accelerator", *AIP Conf. Proc.* **1299**, 3–11 (2010). ⇒ `doi:10.1063/1.3520352`

[46] K. A. Nugent, "X-ray imaging: Caught in a spin", *Nat. Phys.* **5**, 17–18 (2009). ⇒ `doi:10.1038/nphys1171`

[47] E. Esarey, P. Sprangle, J. Krall, and A. Ting, "Overview of plasma-based accelerator concepts", *IEEE Trans. Plasma Sci.* **24**, 252–288 (1996). ⇒ `doi:10.1109/27.509991`

[48] E. Esarey, C. B. Schroeder, and W. P. Leemans, "Physics of laser-driven plasma-based electron accelerators", *Rev. Mod. Phys.* **81**, 1229–1285 (2009). ⇒ `doi:10.1103/RevModPhys.81.1229`

[49] T. P. Rowlands-Rees, *Laser Wakefield Acceleration in the Hydrogen-Filled Capillary Discharge Waveguide*, Ph.D. thesis, University of Oxford (2009). ⇒ `online`

[50] P. Sprangle, E. Esarey, A. Ting, and G. Joyce, "Laser wakefield acceleration and relativistic optical guiding", *Appl. Phys. Lett.* **53**, 2146–2148 (1988). ⇒ `doi:10.1063/1.100300`

[51] J. M. Dawson, "Nonlinear electron oscillations in a cold plasma", *Phys. Rev.* **113**, 383–387 (1959). ⇒ `doi:10.1103/PhysRev.113.383`

[52] P. Mora and T. M. Antonsen, "Electron cavitation and acceleration in the wake of an ultraintense, self-focused laser pulse", *Phys. Rev. E* **53**, R2068–R2071 (1996). ⇒ `doi:10.1103/PhysRevE.53.R2068`

[53] A. Pukhov and J. Meyer-ter Vehn, "Laser wake field acceleration: the highly non-linear broken-wave regime", *App. Phys. B* **74**, 355–361 (2002). ⇒ `doi:10.1007/s003400200795`

[54] S. Gordienko and A. Pukhov, "Scalings for ultrarelativistic laser plasmas and quasimonoenergetic electrons", *Phys. Plasmas* **12**, 043109 (2005). ⇒ `doi:10.1063/1.1884126`

[55] W. Lu, C. Huang, M. Zhou, W. B. Mori, and T. Katsouleas, "Nonlinear theory for relativistic plasma wakefields in the blowout regime", *Phys. Rev. Lett.* **96**, 165002 (2006). ⇒ `doi:10.1103/PhysRevLett.96.165002`

[56] W. Lu, M. Tzoufras, C. Joshi, F. S. Tsung, W. B. Mori, J. Vieira, R. A. Fonseca, and L. O. Silva, “Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime”, *Phys. Rev. ST Accel. Beams* **10**, 061301 (2007). ⇒ doi:10.1103/PhysRevSTAB.10.061301

[57] 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). ⇒ doi:10.1103/PhysRevE.65.056505

[58] S. P. D. Mangles, *et al.*, “Laser-Wakefield acceleration of monoenergetic electron beams in the first Plasma-Wave period”, *Phys. Rev. Lett.* **96**, 215001 (2006). ⇒ doi:10.1103/PhysRevLett.96.215001

[59] S. Kneip, *et al.*, “Bright spatially coherent synchrotron x-rays from a table-top source”, *Nat. Phys.* **6**, 980–983 (2010). ⇒ doi:10.1038/nphys1789

[60] F. S. Tsung, R. Narang, W. B. Mori, C. Joshi, R. A. Fonseca, and L. O. Silva, “Near-GeV-energy Laser-Wakefield acceleration of Self-Injected electrons in a Centimeter-Scale plasma channel”, *Phys. Rev. Lett.* **93**, 185002 (2004). ⇒ doi:10.1103/PhysRevLett.93.185002

[61] S. P. D. Mangles, A. G. R. Thomas, O. Lundh, F. Lindau, M. C. Kaluza, A. Persson, C. G. Wahlström, K. Krushelnick, and Z. Najmudin, “On the stability of laser wakefield electron accelerators in the monoenergetic regime”, *Phys. Plasmas* **14**, 056702 (2007). ⇒ doi:10.1063/1.2436481

[62] S. P. D. Mangles, A. G. R. Thomas, C. Bellei, A. E. Dangor, C. Kamperidis, S. Kneip, S. R. Nagel, L. Willingale, and Z. Najmudin, “Self-Guided wakefield experiments driven by Petawatt-Class ultrashort laser pulses”, *IEEE Trans. Plasma Sci.* **36**, 1715–1721 (2008). ⇒ doi:10.1109/TPS.2008.2000961

[63] D. H. Froula, *et al.*, “Measurements of the critical power for Self-Injection of electrons in a laser wakefield accelerator”, *Phys. Rev. Lett.* **103**, 215006 (2009). ⇒ doi:10.1103/PhysRevLett.103.215006

[64] D. Umstadter, J. K. Kim, and E. Dodd, “Laser injection of ultrashort electron pulses into wakefield plasma waves”, *Phys. Rev. Lett.* **76**, 2073–2076 (1996). ⇒ doi:10.1103/PhysRevLett.76.2073

[65] 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). ⇒ doi:10.1103/PhysRevLett.79.2682

[66] J. Faure, C. Rechatin, A. Norlin, A. Lifschitz, Y. Glinec, and V. Malka, “Controlled injection and acceleration of electrons in plasma wakefields by colliding laser pulses”, *Nature (London)* **444**, 737–739 (2006). ⇒ doi:10.1038/nature05393

[67] C. Rechatin, *et al.*, “Controlling the Phase-Space volume of injected electrons in a Laser-Plasma accelerator”, *Phys. Rev. Lett.* **102**, 164801 (2009). ⇒ doi:10.1103/PhysRevLett.102.164801

[68] C. Rechatin, X. Davoine, A. Lifschitz, A. B. Ismail, J. Lim, E. Lefebvre, J. Faure, and V. Malka, “Observation of beam loading in a Laser-Plasma accelerator”, *Phys. Rev. Lett.* **103**, 194804 (2009). ⇒ `doi:10.1103/PhysRevLett.103.194804`

[69] V. Malka, J. Faure, C. Rechatin, A. B. Ismail, J. K. Lim, X. Davoine, and E. Lefebvre, “Laser-driven accelerators by colliding pulses injection: A review of simulation and experimental results”, *Phys. Plasmas* **16**, 056703 (2009). ⇒ `doi:10.1063/1.3079486`

[70] H. Kotaki, *et al.*, “Electron optical injection with Head-On and countercrossing colliding laser pulses”, *Phys. Rev. Lett.* **103**, 194803 (2009). ⇒ `doi:10.1103/PhysRevLett.103.194803`

[71] 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). ⇒ `doi:10.1103/PhysRevE.58.R5257`

[72] C. G. R. Geddes, K. Nakamura, G. R. Plateau, Cs, C. E. Michel, E. Esarey, C. B. Schroeder, J. R. Cary, and W. P. Leemans, “Plasma-Density-gradient injection of low Absolute-Momentum-spread electron bunches”, *Phys. Rev. Lett.* **100**, 215004 (2008). ⇒ `doi:10.1103/PhysRevLett.100.215004`

[73] J. Faure, C. Rechatin, O. Lundh, L. Ammoura, and V. Malka, “Injection and acceleration of quasimonoenergetic relativistic electron beams using density gradients at the edges of a plasma channel”, *Phys. Plasmas* **17**, 083107 (2010). ⇒ `doi:10.1063/1.3469581`

[74] A. J. Gonsalves, *et al.*, “Tunable laser plasma accelerator based on longitudinal density tailoring”, *Nat. Phys.* (2011).

[75] K. Schmid, A. Buck, C. M. S. Sears, J. M. Mikhailova, R. Tautz, D. Herrmann, M. Geissler, F. Krausz, and L. Veisz, “Density-transition based electron injector for laser driven wakefield accelerators”, *Phys. Rev. ST Accel. Beams* **13**, 091301 (2010). ⇒ `doi:10.1103/PhysRevSTAB.13.091301`

[76] E. Oz, *et al.*, “Ionization-Induced electron trapping in ultrarelativistic plasma wakes”, *Phys. Rev. Lett.* **98**, 084801 (2007). ⇒ `doi:10.1103/PhysRevLett.98.084801`

[77] C. McGuffey, *et al.*, “Ionization induced trapping in a laser wakefield accelerator”, *Phys. Rev. Lett.* **104**, 025004 (2010). ⇒ `doi:10.1103/PhysRevLett.104.025004`

[78] A. Pak, K. A. Marsh, S. F. Martins, W. Lu, W. B. Mori, and C. Joshi, “Injection and trapping of Tunnel-Ionized electrons into Laser-Produced wakes”, *Phys. Rev. Lett.* **104**, 025003 (2010). ⇒ `doi:10.1103/PhysRevLett.104.025003`

[79] C. E. Clayton, *et al.*, “Self-Guided laser wakefield acceleration beyond 1 GeV using Ionization-Induced injection”, *Phys. Rev. Lett.* **105**, 105003 (2010). ⇒ `doi:10.1103/PhysRevLett.105.105003`

[80] J. S. Liu, *et al.*, “All-Optical cascaded laser wakefield accelerator using Ionization-Induced injection”, *Phys. Rev. Lett.* **107**, 035001 (2011). ⇒ `doi:10.1103/PhysRevLett.107.035001`

[81] B. B. Pollock, *et al.*, "Demonstration of a narrow energy spread, approx. 0.5 GeV electron beam from a Two-Stage laser wakefield accelerator", *Phys. Rev. Lett.* **107**, 045001 (2011). ⇒ `doi:10.1103/PhysRevLett.107.045001`

[82] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva, "Beam loading in the nonlinear regime of Plasma-Based acceleration", *Phys. Rev. Lett.* **101**, 145002 (2008). ⇒ `doi:10.1103/PhysRevLett.101.145002`

[83] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva, "Beam loading by electrons in nonlinear plasma wakes", *Phys. Plasmas* **16**, 056705 (2009). ⇒ `doi:10.1063/1.3118628`

[84] C. Rechatin, *et al.*, "Characterization of the beam loading effects in a laser plasma accelerator", *N. J. Phys.* **12**, 045023 (2010). ⇒ `doi:10.1088/1367-2630/12/4/045023`

[85] E. Esarey, P. Sprangle, J. Krall, and A. Ting, "Self-focusing and guiding of short laser pulses in ionizing gases and plasmas", *IEEE J. Quant. Elec.* **33**, 1879–1914 (1997). ⇒ `doi:10.1109/3.641305`

[86] P. Sprangle, E. Esarey, and A. Ting, "Nonlinear theory of intense laser-plasma interactions", *Phys. Rev. Lett.* **64**, 2011–2014 (1990). ⇒ `doi:10.1103/PhysRevLett.64.2011`

[87] P. Sprangle, E. Esarey, and A. Ting, "Nonlinear interaction of intense laser pulses in plasmas", *Phys. Rev. A* **41**, 4463–4469 (1990). ⇒ `doi:10.1103/PhysRevA.41.4463`

[88] T. P. A. Ibbotson, *et al.*, "Investigation of the role of plasma channels as waveguides for laser-wakefield accelerators", *N. J. Phys.* **12**, 045008 (2010). ⇒ `doi:10.1088/1367-2630/12/4/045008`

[89] C. G. Durfee and H. M. Milchberg, "Light pipe for high intensity laser pulses", *Phys. Rev. Lett.* **71**, 2409–2412 (1993). ⇒ `doi:10.1103/PhysRevLett.71.2409`

[90] P. Volfbeyn, E. Esarey, and W. P. Leemans, "Guiding of laser pulses in plasma channels created by the ignitor-heater technique", *Phys. Plasmas* **6**, 2269–2277 (1999). ⇒ `doi:10.1063/1.873503`

[91] S. M. Hooker, D. J. Spence, and R. A. Smith, "Guiding of high-intensity picosecond laser pulses in a discharge-ablated capillary waveguide", *J. Opt. Soc. Am. B* **17**, 90–98 (2000). ⇒ `doi:10.1364/JOSAB.17.000090`

[92] A. Butler, D. J. Spence, and S. M. Hooker, "Guiding of High-Intensity laser pulses with a Hydrogen-Filled capillary discharge waveguide", *Phys. Rev. Lett.* **89**, 185003 (2002). ⇒ `doi:10.1103/PhysRevLett.89.185003`

[93] T. P. Rowlands-Rees, *et al.*, "Laser-Driven acceleration of electrons in a partially ionized plasma channel", *Phys. Rev. Lett.* **100**, 105005 (2008). ⇒ `doi:10.1103/PhysRevLett.100.105005`

[94] C. D. Decker, W. B. Mori, K. C. Tzeng, and T. Katsouleas, "The evolution of ultra-intense, short-pulse lasers in underdense plasmas", *Phys. Plasmas* **3**, 2047–2056 (1996). ⇒ `doi:10.1063/1.872001`

[95] E. Brunetti, *et al.*, “Low emittance, high brilliance relativistic electron beams from a Laser-Plasma accelerator”, *Phys. Rev. Lett.* **105**, 215007 (2010). ⇒ `doi:10.1103/PhysRevLett.105.215007`

[96] C. M. S. Sears, A. Buck, K. Schmid, J. Mikhailova, F. Krausz, and L. Veisz, “Emittance and divergence of laser wakefield accelerated electrons”, *Phys. Rev. ST Accel. Beams* **13**, 092803 (2010). ⇒ `doi:10.1103/PhysRevSTAB.13.092803`

[97] O. Lundh, *et al.*, “Few femtosecond, few kiloampere electron bunch produced by a laser-plasma accelerator”, *Nat. Phys.* **7**, 219–222 (2011). ⇒ `doi:10.1038/nphys1872`

[98] A. Buck, M. Nicolai, K. Schmid, C. M. S. Sears, A. Savert, J. M. Mikhailova, F. Krausz, M. C. Kaluza, and L. Veisz, “Real-time observation of laser-driven electron acceleration”, *Nat. Phys.* **7**, 543–548 (2011). ⇒ `doi:10.1038/nphys1942`

[99] J. Osterhoff, *et al.*, “Generation of stable, Low-Divergence electron beams by Laser-Wakefield acceleration in a Steady-State-flow gas cell”, *Phys. Rev. Lett.* **101**, 085002 (2008). ⇒ `doi:10.1103/PhysRevLett.101.085002`

[100] R. Weingartner, *et al.*, “Imaging laser-wakefield-accelerated electrons using miniature magnetic quadrupole lenses”, *Phys. Rev. ST Accel. Beams* **14**, 052801 (2011). ⇒ `doi:10.1103/PhysRevSTAB.14.052801`

[101] J. G. Gallacher, *et al.*, “A method of determining narrow energy spread electron beams from a laser plasma wakefield accelerator using undulator radiation”, *Phys. Plasmas* **16**, 093102 (2009). ⇒ `doi:10.1063/1.3216549`

[102] K. Nakamura, W. Wan, N. Ybarrolaza, D. Syversrud, J. Wallig, and W. P. Leemans, “Broadband single-shot electron spectrometer for GeV-class laser-plasma-based accelerators”, *Rev. Sci. Instrum.* **79**, 053301 (2008). ⇒ `doi:10.1063/1.2929672`

[103] K. Nakamura, A. J. Gonsalves, C. Lin, A. Smith, D. Rodgers, R. Donahue, W. Byrne, and W. P. Leemans, “Electron beam charge diagnostics for laser plasma accelerators”, *Phys. Rev. ST Accel. Beams* **14**, 062801 (2011). ⇒ `doi:10.1103/PhysRevSTAB.14.062801`

[104] A. Buck, *et al.*, “Absolute charge calibration of scintillating screens for relativistic electron detection”, *Rev. Sci. Instrum.* **81**, 033301 (2010). ⇒ `doi:10.1063/1.3310275`

[105] C. B. Schroeder, E. Esarey, C. Benedetti, Cs, C. G. R. Geddes, and W. P. Leemans, “Design considerations for plasma accelerators driven by lasers or particle beams”, *AIP Conf. Proc.* **1299**, 60–66 (2010). ⇒ `doi:10.1063/1.3520397`

[106] V. Malka, J. Faure, Y. Glinec, A. Pukhov, and J. P. Rousseau, “Monoenergetic electron beam optimization in the bubble regime”, *Phys. Plasmas* **12**, 056702 (2005). ⇒ `doi:10.1063/1.1869498`

[107] J. Faure, Y. Glinec, G. Gallot, and V. Malka, “Ultrashort laser pulses and ultrashort electron bunches generated in relativistic laser-plasma interaction”, *Phys. Plasmas* **13**, 056706 (2006). ⇒ `doi:10.1063/1.2180727`

[108] B. Hidding, *et al.*, “Generation of quasimonoenergetic electron bunches with 80-fs laser pulses”, *Phys. Rev. Lett.* **96**, 105004 (2006). ⇒ `doi:10.1103/PhysRevLett.96.105004`

[109] K. Nakamura, B. Nagler, C. Tóth, C. G. R. Geddes, C. B. Schroeder, E. Esarey, W. P. Leemans, A. J. Gonsalves, and S. M. Hooker, "GeV electron beams from a centimeter-scale channel guided laser wakefield accelerator", *Phys. Plasmas* **14**, 056708 (2007). ⇒ `doi:10.1063/1.2718524`

[110] A. D. Debus, *et al.*, "Electron bunch length measurements from Laser-Accelerated electrons using Single-Shot THz Time-Domain interferometry", *Phys. Rev. Lett.* **104**, 084802 (2010). ⇒ `doi:10.1103/PhysRevLett.104.084802`

[111] A. G. R. Thomas, *et al.*, "Effect of Laser-Focusing conditions on propagation and monoenergetic electron production in Laser-Wakefield accelerators", *Phys. Rev. Lett.* **98**, 095004 (2007). ⇒ `doi:10.1103/PhysRevLett.98.095004`

[112] S. Karsch, *et al.*, "GeV-scale electron acceleration in a gas-filled capillary discharge waveguide", *N. J. Phys.* **9**, 415 (2007). ⇒ `doi:10.1088/1367-2630/9/11/415`

[113] A. G. R. Thomas, *et al.*, "Monoenergetic electronic beam production using dual collinear laser pulses", *Phys. Rev. Lett.* **100**, 255002 (2008). ⇒ `doi:10.1103/PhysRevLett.100.255002`

[114] S. Kneip, *et al.*, "Near-GeV acceleration of electrons by a nonlinear plasma wave driven by a Self-Guided laser pulse", *Phys. Rev. Lett.* **103**, 035002 (2009). ⇒ `doi:10.1103/PhysRevLett.103.035002`

[115] K. Schmid, *et al.*, "Few-Cycle Laser-Driven electron acceleration", *Phys. Rev. Lett.* **102**, 124801 (2009). ⇒ `doi:10.1103/PhysRevLett.102.124801`

[116] T. Matsuoka, *et al.*, "Self-guided laser wakefield acceleration using ablated plasma targets", *Plasma Phys. Controlled Fusion* **51**, 095003 (2009). ⇒ `doi:10.1088/0741-3335/51/9/095003`

[117] A. Popp, *et al.*, "All-Optical steering of Laser-Wakefield-accelerated electron beams", *Phys. Rev. Lett.* **105**, 215001 (2010). ⇒ `doi:10.1103/PhysRevLett.105.215001`

[118] C. M. Huntington, *et al.*, "Current filamentation instability in laser wakefield accelerators", *Phys. Rev. Lett.* **106**, 105001 (2011). ⇒ `doi:10.1103/PhysRevLett.106.105001`

[119] H. Goldstein, *Classical Mechanics*, Addison-Wesley Pub. Co., second ed. (1980). ⇒ `online`

[120] H. Wiedemann, *Particle Accelerator Physics*, Springer-Verlag, Berlin, third ed. (2007). ⇒ `online`

[121] S. Humphries, *Charged particle beams*, Wiley, New York, first printing ed. (1990). ⇒ `online`

[122] K. Floettmann, "Some basic features of the beam emittance", *Phys. Rev. ST Accel. Beams* **6**, 034202 (2003). ⇒ `doi:10.1103/PhysRevSTAB.6.034202`

[123] F. Ciocci, G. Dattoli, A. Torre, and A. Renieri, *Insertion devices for synchrotron radiation and free electron laser*, World Scientific series on synchrotron radiation techniques and applications, vol. 6. World Scientific, Singapore (2000). ⇒ `online`

[124] A. Hofmann, *The Physics of Synchrotron Radiation*, Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press (2004). ⇒ `online`

[125] T. Eichner, *et al.*, "Miniature magnetic devices for laser-based, table-top free-electron lasers", *Phys. Rev. ST Accel. Beams* **10**, 082401 (2007). ⇒ `doi:10.1103/PhysRevSTAB.10.082401`

[126] S. Becker, *et al.*, "Characterization and tuning of ultrahigh gradient permanent magnet quadrupoles", *Phys. Rev. ST Accel. Beams* **12**, 102801 (2009). ⇒ `doi:10.1103/PhysRevSTAB.12.102801`

[127] G. Fubiani, J. Qiang, E. Esarey, W. P. Leemans, and G. Dugan, "Space charge modeling of dense electron beams with large energy spreads", *Phys. Rev. ST Accel. Beams* **9**, 064402 (2006). ⇒ `doi:10.1103/PhysRevSTAB.9.064402`

[128] G. Geloni, E. Saldin, E. Schneidmiller, and M. Yurkov, "Longitudinal wake field for an electron beam accelerated through an ultrahigh field gradient", *Nucl. Instrum. Meth. Phys. Res. A* **578**, 34–46 (2007). ⇒ `doi:10.1016/j.nima.2007.05.120`

[129] F. J. Grüner, C. B. Schroeder, A. R. Maier, S. Becker, and J. M. Mikhailova, "Space-charge effects in ultrahigh current electron bunches generated by laser-plasma accelerators", *Phys. Rev. ST Accel. Beams* **12**, 020701 (2009). ⇒ `doi:10.1103/PhysRevSTAB.12.020701`

[130] S. Becker, P. Hilz, F. Grüner, and D. Habs, "Retardation effects in space charge calculations of high current electron beams", *Phys. Rev. ST Accel. Beams* **12**, 101301 (2009). ⇒ `doi:10.1103/PhysRevSTAB.12.101301`

[131] C. W. Mansell, "Studies of the effect of coulomb forces in an electron bunch", *Tech. rep.*, University of Oxford (2008). ⇒ `online`

[132] S. Bajlekov, R. Bartolini, N. Delerue, G. Doucas, S. M. Hooker, K. Peach, D. Urner, and J. S. Wark, "OPALS: The oxford plasma accelerator light source project", In *Proceedings of the 30th International Free Electron Laser Conference*, p. 163. Gyeongju, Korea (2008). ⇒ `online`

[133] S. Reiche, *Numerical Studies for a Single Pass High Gain Free-Electron Laser*, Ph.D. thesis, Universität Hamburg, Hamburg (1999). ⇒ `online`

[134] D. Attwood, *Soft X-Rays and Extreme Ultraviolet Radiation: Principles and Applications*, Cambridge University Press (2007). ⇒ `online`

[135] K. J. Kim, "Characteristics of synchrotron radiation", *AIP Conf. Proc.* **184**, 565–632 (1989). ⇒ `doi:10.1063/1.38046`

[136] T. Tanaka and H. Kitamura, "SPECTRA: a synchrotron radiation calculation code", *J. Synchrotron Rad.* **8**, 1221–1228 (2001). ⇒ `doi:10.1107/S090904950101425X`

[137] Z. Huang and K. J. Kim, "Review of x-ray free-electron laser theory", *Phys. Rev. ST Accel. Beams* **10**, 034801 (2007). ⇒ `doi:10.1103/PhysRevSTAB.10.034801`

[138] C. Pellegrini, “Design considerations for a SASE X-ray FEL”, *Nucl. Instrum. Meth. Phys. Res. A* **475**, 1–12 (2001). ⇒ doi:10.1016/S0168-9002(01)01527-3

[139] P. Sprangle and A. T. Drobot, “The linear and Self-Consistent nonlinear theory of the electron cyclotron maser instability”, *IEEE Trans. Microwave Theory* **25**, 528–544 (1977). ⇒ doi:10.1109/TMTT.1977.1129151

[140] H. Haus, “Noise in free-electron laser amplifier”, *IEEE J. Quant. Elec.* **17**, 1427–1435 (1981). ⇒ doi:10.1109/JQE.1981.1071289

[141] G. Dattoli, A. Marino, A. Renieri, and F. Romanelli, “Progress in the hamiltonian picture of the free-electron laser”, *IEEE J. Quant. Elec.* **17**, 1371–1387 (1981). ⇒ doi:10.1109/JQE.1981.1071268

[142] R. Bonifacio, C. Pellegrini, and L. M. Narducci, “Collective instabilities and high-gain regime in a free electron laser”, *Opt. Commun.* **50**, 373–378 (1984). ⇒ doi:10.1016/0030-4018(84)90105-6

[143] W. M. Fawley, “Algorithm for loading shot noise microbunching in multidimensional, free-electron laser simulation codes”, *Phys. Rev. ST Accel. Beams* **5**, 070701 (2002). ⇒ doi:10.1103/PhysRevSTAB.5.070701

[144] D. Ratner, *et al.*, “FEL gain length and taper measurements at LCLS”, In *Proceedings of the 31st International Free Electron Laser Conference*, p. 221. Liverpool, U.K. (2009). ⇒ online

[145] L. Giannessi, *et al.*, “Self-amplified spontaneous emission free-electron laser with an energy-chirped electron beam and undulator tapering”, *Phys. Rev. Lett.* **106**, 144801 (2011). ⇒ doi:10.1103/PhysRevLett.106.144801

[146] R. Bonifacio and F. Casagrande, “The superradiant regime of a free electron laser”, *Nucl. Instrum. Meth. Phys. Res. A* **239**, 36–42 (1985). ⇒ doi:10.1016/0168-9002(85)90695-3

[147] E. Saldin, E. Schneidmiller, and M. Yurkov, “Coherence properties of the radiation from x-ray free electron laser”, *Opt. Commun.* **281**, 1179 (2008). ⇒ doi:10.1016/j.optcom.2007.10.044

[148] M. Xie, “Exact and variational solutions of 3D eigenmodes in high gain FELs”, *Nucl. Instrum. Meth. Phys. Res. A* **445**, 59–66 (2000). ⇒ doi:10.1016/S0168-9002(00)00114-5

[149] L. Giannessi, “Harmonic generation and linewidth narrowing in seeded FELs”, In *Proceedings of the 26th International Free Electron Laser Conference*, p. 37. Sincrotrone Trieste, Frascati, Italy (2004). ⇒ online

[150] G. Lambert, *et al.*, “Injection of harmonics generated in gas in a free-electron laser providing intense and coherent extreme-ultraviolet light”, *Nat. Phys.* **4**, 296–300 (2008). ⇒ doi:10.1038/nphys889

[151] T. Hara, *et al.*, “First observation of the 61.5 nm seeded FEL at the SCSS test accelerator”, In *Proceedings of the 32nd International Free Electron Laser Conference*. Malmö, Sweden (2010). ⇒ online

[152] P. B. Corkum and F. Krausz, “Attosecond science”, *Nat. Phys.* **3**, 381–387 (2007). ⇒ `doi:10.1038/nphys620`

[153] M. Lewenstein, P. Balcou, M. Ivanov, A. L’Huillier, and P. B. Corkum, “Theory of high-harmonic generation by low-frequency laser fields”, *Phys. Rev. A* **49**, 2117–2132 (1994). ⇒ `doi:10.1103/PhysRevA.49.2117`

[154] L. H. Yu, *et al.*, “High-Gain Harmonic-Generation Free-Electron laser”, *Science* **289**, 932–934 (2000). ⇒ `doi:10.1126/science.289.5481.932`

[155] R. Bonifacio, L. De Salvo Souza, P. Pierini, and E. T. Scharlemann, “Generation of XUV light by resonant frequency tripling in a two-wiggler FEL amplifier”, *Nucl. Instrum. Meth. Phys. Res. A* **296**, 787–790 (1990). ⇒ `doi:10.1016/0168-9002(90)91307-W`

[156] C. S. Hwang, J. C. Jan, C. S. Chang, S. D. Chen, C. H. Chang, and T. M. Uen, “Development trends for insertion devices of future synchrotron light sources”, *Phys. Rev. ST Accel. Beams* **14**, 044801 (2011). ⇒ `doi:10.1103/PhysRevSTAB.14.044801`

[157] P. Elleaume, J. Chavanne, and B. Faatz, “Design considerations for a 1Å SASE undulator”, *Nucl. Instrum. Meth. Phys. Res. A* **455**, 503–523 (2000). ⇒ `doi:10.1016/S0168-9002(00)00544-1`

[158] K. Halbach, “Permanent magnet undulators”, *Le Journal de Physique Colloques* **44**, 211–216 (1983). ⇒ `doi:10.1051/jphyscol:1983120`

[159] Y. Ding, *et al.*, “Measurements and simulations of ultralow emittance and ultrashort electron beams in the linac coherent light source”, *Phys. Rev. Lett.* **102**, 254801 (2009). ⇒ `doi:10.1103/PhysRevLett.102.254801`

[160] J. B. Rosenzweig, *et al.*, “Generation of ultra-short, high brightness electron beams for single-spike SASE FEL operation”, *Nucl. Instrum. Meth. Phys. Res. A* **593**, 39–44 (2008). ⇒ `doi:10.1016/j.nima.2008.04.083`

[161] S. Reiche, P. Musumeci, C. Pellegrini, and J. Rosenzweig, “Development of ultra-short pulse, single coherent spike for SASE x-ray FELs”, *Nucl. Instrum. Meth. Phys. Res. A* **593**, 45–48 (2008). ⇒ `doi:10.1016/j.nima.2008.04.061`

[162] R. Bonifacio, C. Maroli, and N. Piovella, “Slippage and superradiance in the high-gain FEL: Linear theory”, *Opt. Commun.* **68**, 369–374 (1988). ⇒ `doi:10.1016/0030-4018(88)90234-9`

[163] R. Bonifacio, B. W. J. McNeil, and P. Pierini, “Superradiance in the high-gain free-electron laser”, *Phys. Rev. A* **40**, 4467–4475 (1989). ⇒ `doi:10.1103/PhysRevA.40.4467`

[164] G. Dattoli, L. Giannessi, P. L. Ottaviani, and C. Ronsivalle, “Semi-analytical model of self-amplified spontaneous-emission free-electron lasers, including diffraction and pulse-propagation effects”, *J. App. Phys.* **95**, 3206–3210 (2004). ⇒ `doi:10.1063/1.1645979`

[165] C. B. Schroeder, W. M. Fawley, E. Esarey, and W. P. Leemans, “Design of an XUV FEL driven by the laser-plasma accelerator at the LBNL LOASIS facility”, In *Proceedings of the 28th International Free Electron Laser Conference*. BESSY, Berlin, Germany (2006). ⇒ `online`

[166] C. B. Schroeder, W. M. Fawley, F. Grüner, M. Bakeman, K. Nakamura, K. E. Robinson, C. Tóth, E. Esarey, and W. P. Leemans, "Free-electron laser driven by the LBNL laser-plasma accelerator", *AIP Conf. Proc.* **1086**, 637–642 (2009). ⇒ `doi:10.1063/1.3080982`

[167] S. Reiche, "GENESIS 1.3: a fully 3D time-dependent FEL simulation code", *Nucl. Instrum. Meth. Phys. Res. A* **429**, 243–248 (1999). ⇒ `doi:10.1016/S0168-9002(99)00114-X`

[168] E. L. Saldin, E. A. Schneidmiller, and M. V. Yurkov, "Statistical properties of radiation from VUV and x-ray free electron laser", *Opt. Commun.* **148**, 383–403 (1998). ⇒ `doi:10.1016/S0030-4018(97)00670-6`

[169] R. Bonifacio, L. De Salvo, P. Pierini, N. Piovella, and C. Pellegrini, "Spectrum, temporal structure, and fluctuations in a high-gain free-electron laser starting from noise", *Phys. Rev. Lett.* **73**, 70–73 (1994). ⇒ `doi:10.1103/PhysRevLett.73.70`

[170] X. He, *et al.*, "Spatial and spectral properties of the high-order harmonic emission in argon for seeding applications", *Phys. Rev. A* **79**, 063829 (2009). ⇒ `doi:10.1103/PhysRevA.79.063829`

[171] I. M. Frank and V. L. Ginzburg, "Radiation of a uniform moving electron due to its transition from one medium into another", *J. Phys. (USSR)* **9**, 353 (1945).

[172] P. A. Cherenkov, "Visible emission of clean liquids by action of gamma radiation", *Doklady Akademii Nauk SSSR* **2**, 451 (1934). ⇒ `online`

[173] V. L. Ginzburg and V. N. TSytovich, *Transition radiation and transition scattering*, A. Hilger, Bristol, U.K. (1990). ⇒ `online`

[174] J. D. Jackson, *Classical electrodynamics*, John Wiley & Sons, New York, third ed. (1999). ⇒ `online`

[175] S. Casalbuoni, B. Schmidt, P. Schmüser, and B. Steffen, "Far-infrared transition and diffraction radiation — Part I: Production, diffraction effects and optical propagation", *Tech. rep.*, DESY (Germany) (2005). ⇒ `online`

[176] H. Loos, *et al.*, "Observation of coherent optical transition radiation in the LCLS linac", In *Proceedings of the 30th International Free Electron Laser Conference*, p. 485. Gyeongju, Korea (2008). ⇒ `online`

[177] M. Castellano and V. A. Verzilov, "Spatial resolution in optical transition radiation beam diagnostics", *Phys. Rev. ST Accel. Beams* **1**, 062801 (1998). ⇒ `doi:10.1103/PhysRevSTAB.1.062801`

[178] Y. Glinec, J. Faure, A. Norlin, A. Pukhov, and V. Malka, "Observation of fine structures in Laser-Driven electron beams using coherent transition radiation", *Phys. Rev. Lett.* **98**, 194801 (2007). ⇒ `doi:10.1103/PhysRevLett.98.194801`

[179] C. Lin, *et al.*, "Observation of long range coherent OTR from LPA electron beams", In *Proceedings of SPIE Optics + Optoelectronics: Laser Acceleration of Electrons, Protons, and Ions* (K. W. D. Ledingham, W. P. Leemans, E. Esarey, S. M. Hooker, K. Spohr, and P. McKenna, eds.). SPIE, SPIE (2011). ⇒ `online`

[180] W. P. Leemans, *et al.*, “Observation of terahertz emission from a Laser-Plasma accelerated electron bunch crossing a Plasma-Vacuum boundary”, *Phys. Rev. Lett.* **91**, 074802 (2003). ⇒ `doi:10.1103/PhysRevLett.91.074802`

[181] J. van Tilborg, *et al.*, “Temporal characterization of femtosecond Laser-Plasma-accelerated electron bunches using terahertz radiation”, *Phys. Rev. Lett.* **96**, 014801 (2006). ⇒ `doi:10.1103/PhysRevLett.96.014801`

[182] B. Schmidt, C. Behrens, H. Delsim-Hashemi, P. Schmüser, and S. Wesch, “Longitudinal structure of electron bunches at the micrometer scale from spectroscopy of coherent transition radiation”, In *Proceedings of the 11th European Particle Accelerator Conference*, p. 130. Genoa, Italy (2008). ⇒ `online`

[183] K. Németh, B. Shen, Y. Li, H. Shang, R. Crowell, K. C. Harkay, and J. R. Cary, “Laser-Driven coherent betatron oscillation in a Laser-Wakefield cavity”, *Phys. Rev. Lett.* **100**, 095002 (2008). ⇒ `doi:10.1103/PhysRevLett.100.095002`

[184] H. Delsim-Hashemi, *Infrared Single Shot Diagnostics for the Longitudinal Profile of the Electron Bunches at FLASH*, Ph.D. thesis, DESY (2008). ⇒ `online`

[185] M. A. Green and M. J. Keevers, “Optical properties of intrinsic silicon at 300 K”, *Progress in Photovoltaics: Research and Applications* **3**, 189–192 (1995). ⇒ `doi:10.1002/pip.4670030303`

[186] D. Sayre, “Some implications of a theorem due to Shannon”, *Acta Crystallographica* **5**, 843 (1952). ⇒ `doi:10.1107/S0365110X52002276`

[187] R. W. Gerchberg and W. O. Saxton, “A practical algorithm for the determination of the phase from image and diffraction plane pictures”, *Optik (Jena)* **35**, 237 (1972).

[188] J. R. Fienup, “Phase retrieval algorithms: a comparison”, *Appl. Optics* **21**, 2758–2769 (1982). ⇒ `doi:10.1364/AO.21.002758`

[189] V. Elser, “Phase retrieval by iterated projections”, *J. Opt. Soc. Am. A* **20**, 40–55 (2003). ⇒ `doi:10.1364/JOSAA.20.000040`

[190] S. Marchesini, *et al.*, “Massively parallel x-ray holography”, *Nat. Photon.* **2**, 560–563 (2008). ⇒ `doi:10.1038/nphoton.2008.154`

[191] V. Elser, I. Rankenburg, and P. Thibault, “Searching with iterated maps”, *Proc. Nat. Acad. Sci.* **104**, 418–423 (2007). ⇒ `doi:10.1073/pnas.0606359104`

[192] S. Marchesini, H. He, H. N. Chapman, S. P. H. Riege, A. Noy, M. R. Howells, U. Weierstall, and J. C. H. Spence, “X-ray image reconstruction from a diffraction pattern alone”, *Phys. Rev. B* **68**, 140101 (2003). ⇒ `doi:10.1103/PhysRevB.68.140101`

[193] H. N. Chapman, *et al.*, “High-resolution ab initio three-dimensional x-ray diffraction microscopy”, *J. Opt. Soc. Am. A* **23**, 1179–1200 (2006). ⇒ `doi:10.1364/JOSAA.23.001179`

[194] 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). ⇒ `doi:10.1103/PhysRevE.69.016501`

[195] A. Popp, *Dynamics of electron acceleration in laser-driven wakefields: Acceleration limits and asymmetric plasma waves (working title)*, Ph.D. thesis, Ludwig-Maximilians-Universität München (2011).

[196] I. Blumenfeld, *et al.*, "Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator", *Nature (London)* **445**, 741–744 (2007). ⇒ `doi:10.1038/nature05538`

[197] B. Hidding, T. Königstein, J. Osterholz, S. Karsch, O. Willi, and G. Pretzler, "Monoenergetic energy doubling in a hybrid Laser-Plasma wakefield accelerator", *Phys. Rev. Lett.* **104**, 195002 (2010). ⇒ `doi:10.1103/PhysRevLett.104.195002`

[198] K. H. Pae, I. W. Choi, and J. Lee, "Self-mode-transition from laser wakefield accelerator to plasma wakefield accelerator of laser-driven plasma-based electron acceleration", *Phys. Plasmas* **17**, 123104 (2010). ⇒ `doi:10.1063/1.3522757`

[199] A. G. R. Thomas and K. Krushelnick, "Betatron x-ray generation from electrons accelerated in a plasma cavity in the presence of laser fields", *Phys. Plasmas* **16**, 103103 (2009). ⇒ `doi:10.1063/1.3237089`

[200] K. T. Phuoc, S. Corde, R. Shah, F. Albert, R. Fitour, J. P. Rousseau, F. Burgy, B. Mercier, and A. Rousse, "Imaging electron trajectories in a Laser-Wakefield cavity using betatron X-Ray radiation", *Phys. Rev. Lett.* **97**, 225002 (2006). ⇒ `doi:10.1103/PhysRevLett.97.225002`

[201] H. Kotaki, Y. Hayashi, K. Kawase, M. Mori, M. Kando, T. Homma, J. K. Koga, H. Daido, and S. V. Bulanov, "Manipulation and electron-oscillation-measurement of laser accelerated electron beams", *Plasma Phys. Controlled Fusion* **53**, 014009 (2011). ⇒ `doi:10.1088/0741-3335/53/1/014009`

[202] S. P. D. Mangles, *et al.*, "Electron acceleration in cavitated channels formed by a petawatt laser in Low-Density plasma", *Phys. Rev. Lett.* **94**, 245001 (2005). ⇒ `doi:10.1103/PhysRevLett.94.245001`

[203] S. I. Bajlekov, W. M. Fawley, C. B. Schroeder, R. Bartolini, and S. M. Hooker, "Simulation of free-electron lasers seeded with broadband radiation", *Phys. Rev. ST Accel. Beams* **14**, 060711 (2011). ⇒ `doi:10.1103/PhysRevSTAB.14.060711`

[204] S. Krinsky and L. H. Yu, "Output power in guided modes for amplified spontaneous emission in a single-pass free-electron laser", *Phys. Rev. A* **35**, 3406–3423 (1987). ⇒ `doi:10.1103/PhysRevA.35.3406`

[205] S. Krinsky and Z. Huang, "Frequency chirped self-amplified spontaneous-emission free-electron lasers", *Phys. Rev. ST Accel. Beams* **6**, 050702 (2003). ⇒ `doi:10.1103/PhysRevSTAB.6.050702`

[206] S. Krinsky, "Transient analysis of free-electron lasers with discrete radiators", *Phys. Rev. E* **59**, 1171–1183 (1999). ⇒ `doi:10.1103/PhysRevE.59.1171`

[207] G. Sansone, *et al.*, "Isolated Single-Cycle attosecond pulses", *Science* **314**, 443–446 (2006). ⇒ `doi:10.1126/science.1132838`

[208] L. Giannessi, "Overview of Perseo, a system for simulating FEL dynamics in Mathcad", In *Proceedings of the 28th International Free Electron Laser Conference*, p. 91. BESSY, Berlin, Germany (2006). ⇒ `online`

[209] L. Giannessi, M. Quattromini, P. Musumeci, M. Nisoli, G. Sansone, S. Stagira, and S. de Silvestri, "Analysis of the process of amplification in a single pass FEL of high order harmonics generated in a gas jet", In *Proceedings of the 28th International Free Electron Laser Conference*, p. 248. BESSY, Berlin, Germany (2006). ⇒ `online`

[210] B. W. J. McNeil, D. Dunning, N. R. Thompson, and B. Sheehy, "The use of HHG at 4GLS", In *Proceedings of the 28th International Free Electron Laser Conference*, p. 234. BESSY, Berlin, Germany (2006). ⇒ `online`

[211] J. Wu, P. R. Bolton, J. B. Murphy, and X. Zhong, "Free electron laser seeded by ir laser driven high-order harmonic generation", *Appl. Phys. Lett.* **90**, 021109 (2007). ⇒ `doi:10.1063/1.2431455`

[212] C. Maroli, "One-dimensional free-electron laser equations without slowly varying envelope approximation", *Opt. Commun.* **208**, 155–161 (2002). ⇒ `doi:10.1016/S0030-4018(02)01569-9`

[213] C. K. Birdsall and A. B. Langdon, *Plasma Physics via Computer Simulation*, Series in Plasma Physics. Taylor & Francis, London (2004). ⇒ `online`

[214] N. Piovella, "High gain free electron laser amplifiers starting from coherent and incoherent spontaneous emission", *Phys. Plasmas* **6**, 3358–3368 (1999). ⇒ `doi:10.1063/1.873575`

[215] B. W. J. McNeil, G. R. M. Robb, and D. A. Jaroszynski, "Self-amplification of coherent spontaneous emission in the free electron laser", *Opt. Commun.* **165**, 65–70 (1999). ⇒ `doi:10.1016/S0030-4018(99)00222-9`

[216] B. W. J. McNeil and G. R. M. Robb, "Self-amplified coherent spontaneous emission in the planar wiggler free-electron laser", *Phys. Rev. E* **65**, 046503 (2002). ⇒ `doi:10.1103/PhysRevE.65.046503`

[217] B. W. J. McNeil, M. W. Poole, and G. R. M. Robb, "Unified model of electron beam shot noise and coherent spontaneous emission in the helical wiggler free electron laser", *Phys. Rev. ST Accel. Beams* **6**, 070701 (2003). ⇒ `doi:10.1103/PhysRevSTAB.6.070701`

[218] Z. Huang and K. J. Kim, "Effects of bunch density gradient in high-gain free-electron lasers", *Nucl. Instrum. Meth. Phys. Res. A* **445**, 105–109 (2000). ⇒ `doi:10.1016/S0168-9002(00)00122-4`

[219] C. Maroli, V. Petrillo, and M. Ferrario, "One-dimensional free-electron laser equations without the slowly varying envelope approximation", *Phys. Rev. ST Accel. Beams* **14**, 070703 (2011). ⇒ `doi:10.1103/PhysRevSTAB.14.070703`

[220] S. I. Bajlekov, R. Bartolini, and S. M. Hooker, "Simulating sub-wavelength temporal effects in a seeded FEL driven by laser-accelerated electrons", In *Proceedings of the 31st International Free Electron Laser Conference*, p. 119. Liverpool, U.K. (2009). ⇒ `online`

[221] L. T. Campbell, R. Martin, and B. W. J. McNeil, "A fully 3D unaveraged, Non-Localised electron, Parallelized-Computational model of the FEL", In *Proceedings of the 31st International Free Electron Laser Conference*, p. 115. Liverpool, U.K. (2009). ⇒ `online`

[222] C. Penman and B. W. J. McNeil, "Simulation of input electron noise in the free-electron laser", *Opt. Commun.* **90**, 82–84 (1992). ⇒ `doi:10.1016/0030-4018(92)90333-M`

[223] J. Chen, S.-I. Chu, and J. Liu, "Time-frequency analysis of molecular high-harmonic generation spectrum by means of wavelet transform and wigner distribution techniques", *J. Phys. B.* **39**, 4747 (2006). ⇒ `doi:10.1088/0953-4075/39/22/017`

[224] G. Geloni, E. Saldin, E. Schneidmiller, and M. Yurkov, "Longitudinal impedance and wake from XFEL undulators. impact on current-enhanced SASE schemes", *Nucl. Instrum. Meth. Phys. Res. A* **583**, 228–247 (2007). ⇒ `doi:10.1016/j.nima.2007.09.019`

[225] A. Tremaine, J. B. Rosenzweig, S. Anderson, P. Frigola, M. Hogan, A. Murokh, C. Pellegrini, D. C. Nguyen, and R. L. Sheffield, "Observation of Self-Amplified Spontaneous-Emission-induced Electron-Beam microbunching using coherent transition radiation", *Phys. Rev. Lett.* **81**, 5816 (1998). ⇒ `doi:10.1103/PhysRevLett.81.5816`

[226] A. Tremaine, *et al.*, "Experimental characterization of nonlinear harmonic radiation from a visible Self-Amplified spontaneous emission Free-Electron laser at saturation", *Phys. Rev. Lett.* **88**, 204801 (2002). ⇒ `doi:10.1103/PhysRevLett.88.204801`

[227] A. Lumpkin, B. Yang, W. Berg, Y. Chae, N. Sereno, R. Dejus, C. Benson, and E. Moog, "Utilization of CTR to measure the evolution of electron-beam microbunching in a self-amplified spontaneous emission (SASE) free-electron laser (FEL)", *Nucl. Instrum. Meth. Phys. Res. A* **475**, 462–469 (2001). ⇒ `doi:10.1016/S0168-9002(01)01688-6`

[228] A. H. Lumpkin, *et al.*, "Comprehensive z-dependent measurements of electron-beam microbunching using COTR in a saturated SASE FEL", *Nucl. Instrum. Meth. Phys. Res. A* **483**, 394–401 (2002). ⇒ `doi:doi:10.1016/S0168-9002(02)00350-9`

[229] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, third ed. (2007). ⇒ `online`