

1

Introduction to Strong Field Physics

1.1 Definition of Strong Field Physics

Strong field physics (or “high field physics” in some of the literature) refers to phenomena that occur during the interaction of intense electromagnetic waves with matter of various forms. While it is possible to create “strong” static electric or magnetic fields in the laboratory, by far the highest field strengths that can be produced in the lab occur in the electromagnetic fields of a focused high-power laser pulse. For example, while the strongest DC electric fields produced in the lab rarely exceed one MV per cm (limited by the ability of materials to withstand plasma breakdown at high fields), a petawatt-class high-intensity laser can produce oscillatory electric fields with values of over one TV per cm, nearly six orders of magnitude greater than any laboratory-produced DC field. The highest DC magnetic fields produced by laboratory magnets are around one megagauss, whereas a focused petawatt-class laser can produce an oscillatory magnetic field with peak values of many gigagauss.

The interaction of such high-intensity focused electromagnetic radiation with matter can lead to exotic physics. While strong field interactions have been accessed with microwave radiation (Gallagher 1992), traditionally, strong field physics has been studied with intense optical and near-infrared (IR) pulses generated by high-intensity lasers. These interactions occur in a regime in which the electric field of the optical wave dominates the motion and dynamics of electrons subject to these fields. They are characterized by interactions that are often highly nonlinear. At the highest intensities that are accessible, the motion of electrons can become relativistic during each optical cycle, and the magnetic field of the light pulse becomes important in affecting the motion of electrons in the field.

There is no generally accepted definition for when an electromagnetic field is high enough amplitude to enter the “strong field” regime, and, to a certain extent, the threshold for strong field physics will depend on the particular situation. However, to guide the reader it is interesting to ask at what focused intensity might we expect to encounter strong field phenomena. There are two ways to look at this question.

From the standpoint of the interaction of an intense laser field with a free atom, it is fair to say that the strong field regime is entered when the light field is intense enough that perturbation theory breaks down in the quantum mechanical description of the interaction

with the bound electrons. This breakdown occurs when the light intensity is high enough that the peak electric field of the wave $E_0 = \sqrt{8\pi I/c}$ (where I is the light intensity and c is the speed of light) approaches the atomic unit of electric field $E_a = e/a_B^2 = 5.1 \times 10^9$ V/cm, the field felt by an electron in a hydrogen atom (where a_B is the Bohr radius and e is the charge of the electron). Light acquires this electric field at an intensity of 3.5×10^{16} W/cm². Because perturbation theory relies on the convergence of a perturbation series, in practice, nonperturbative effects become manifest at fields that are a few percent of this value. Consequently, strong field effects in focused laser interactions with atoms become evident at intensities above about 10^{14} W/cm².

Alternatively, from the standpoint of laser light interacting with the free electrons in a plasma, we might argue that the strong field regime will be entered when the laser can drive oscillations of the free electrons with an energy that is comparable to or exceeds the thermal energy of the electrons in that plasma. When this occurs, the laser field dominates the bulk motion of the plasma electrons. Since plasmas begin to form at electron temperatures comparable to atomic ionization potentials, many of the plasmas encountered in the lab have temperatures of ~ 10 eV up to a few keV. Since the quiver energy of an electron in a near-infrared field acquires a value of a few eV at intensities just under 10^{14} W/cm², we can fairly say that the strong field regime of laser plasma interactions is entered at focused intensity above 10^{14} W/cm².

With these two alternative views of laser–matter interaction, we will take as a starting point for this book that strong field physics is accessed at laser intensities above 10^{14} W/cm² (or in a few cases just below this). These days, such intensities are considered quite modest and can be produced with rather compact, tabletop lasers. The upper end of our realm of study is limited only by the experimental ability to create and focus very high-peak-power lasers. At the time of this writing, the highest intensities produced have been in the vicinity of a few times 10^{22} W/cm² and there are lasers which will soon produce intensity one order of magnitude higher.

Describing strong field phenomena over an intensity window spanning nine orders of magnitude is daunting. However, many of the theoretical techniques for describing these interactions are applicable over a wide range of intensity. Only when the motion of electrons becomes relativistic in the laser field (an effect which occurs at intensity around 10^{18} W/cm² in most near-IR fields) do the theoretical descriptions require amendment. Such intensities are also characterized by high magnetic fields and optical forces. For example, in a pulse with intensity of 10^{18} W/cm², an intensity fairly easily accessed by modern tabletop ultrafast lasers, the peak electric field is 3×10^{10} V/cm and the optical magnetic field is 100 MG. The light pressure, I/c , is ~ 0.3 Gbar. At the time of the writing of this book, the highest-peak-power lasers can reach 100 TV/cm fields and > 10 Tbar light pressures.

1.2 Historical Overview

Theoretical considerations of how intense light interacts with matter are not particularly new and predate the invention of the laser. High-intensity light excitation of a multi-photon process was considered as early as 1931 when Maria Goeppert-Mayer discussed

two-photon absorption in her PhD thesis (Goeppert-Mayer 1931). It was not until after the demonstration of the laser, however, that a true multiphoton process was observed in the laboratory when, in 1961, Peter Franken working at the University of Michigan observed second harmonic conversion of a ruby laser pulse in a quartz crystal (Franken *et al.* 1961). It was the rapid development in laser technology in the 1960s and early 1970s that led to a rapid increase in peak power enabling true strong field physics studies. Q-switching (McClung and Hellwarth 1962) and mode-locking (Demaria *et al.* 1966) were technological advancements that permitted the construction of lasers with peak power over 1 GW and focusable intensity entering the strong field regime.

How a strong field interacts with a free electron in the relativistic regime was a topic studied by a number of authors in the early 1960s (often in an astrophysical context), and a reasonable date for the dawn of the field of strong field physics can be marked by the 1964 theoretical publication by Brown and Kibble on the relativistic dynamics of a free electron in a laser field (Brown and Kibble 1964). However, the theoretical study of strong field physics really began in earnest with a classic paper by Russian physicist L. V. Keldysh in the following year (Keldysh 1965). In this paper, the rate of ionization of an atom or ion in a strong laser field was first derived with a nonperturbative quantum theory. At the time, this work was largely ignored in the West. Keldysh's theoretical work amazingly predicted phenomena such as tunnel ionization and high-order above-threshold ionization, effects which became the focus of strong field physics experiments a decade later and have been at the center of strong field research for many years.

The first real experimental observation of a nonperturbative strong field physics effect occurred in the groundbreaking experiment of Agostini *et al.* in 1979 at the Saclay lab in France (Agostini *et al.* 1979). This work was one of the pioneering discoveries that led to the 2023 Nobel Prize in Physics. Their experiment observed, for the first time, truly nonperturbative multiphoton effects in laser–atom interactions by examining photoelectron production from intense six-photon ionization of Xe atoms at intensity up to 4×10^{13} W/cm². They found that electrons were ejected during ionization with energy higher than that expected from absorption of the minimum number of photons needed for ionization. At the highest intensities studied in that experiment, the ejected electrons were emitted in a number of energy peaks separated in energy by one-photon quanta with almost equal electron yield over the first four or five peaks, an effect that was coined “above threshold ionization” (ATI) shortly after its observation (Fabre *et al.* 1982). Despite Keldysh's prediction of this very effect 15 years earlier, this experiment was greeted with surprise by the atomic physics community at the time, as it was expected that these higher-order peaks would be emitted with exponentially decreasing amplitude as predicted by lowest-order perturbation theory. This observation of nonperturbative effects sparked a long campaign of experiments and theoretical work on strong laser field ionization of atoms and ions that continues to this day.

The early theoretical work of Keldysh was subsequently elaborated on by two authors in the 1970s and 1980s, F. Faisal and H. Reiss (see in Faisal's book of 1987; and Reiss 1980), leading to so-called Keldysh–Faisal–Reiss or KFR theories, which represent the primary basis for analytic strong field theory in atoms to the present. Experimental work in strong field physics exploded at about the same time, propelled at first by the

development of mode-locking, permitting production of picosecond and, shortly thereafter, femtosecond laser pulses which could then be amplified with broadband dye laser amplifiers (Shank and Ippen 1974). However, the field really skyrocketed with the revolutionary development of chirped pulse amplification (CPA) in 1985 by Gerard Mourou and Donna Strickland and at the University of Rochester, a discovery which led to the award of the 2018 Nobel Prize in Physics to Mourou and Strickland (Strickland and Mourou 1985). This laser technology advance permitted construction of tabletop-scale lasers with powers well in excess of 1 terawatt, and such systems quickly proliferated around the world, particularly in the US, France and Britain. It was soon realized that this technology might lead to lasers with power above 1 petawatt (Maine *et al.* 1987).

Work on atomic ionization made possible by the development of CPA and inspired by a number of Russian theoretical works that followed Keldysh's initial paper (Perelomov *et al.* 1966) showed that strong field ionization could often be described by tunneling theories (Augst *et al.* 1991). This intellectual leap has inspired much of the modern understanding of strong field ionization and a number of other strong field phenomena in atoms, ions and molecules.

Those early experiments in strong field multiphoton ionization were followed by the first observation of nonperturbative nonlinear optical phenomena in high-order harmonic generation at the University of Illinois, Chicago in 1987 (McPherson *et al.* 1987). This group found that interactions of an intense laser pulse with a gas of atoms led to emission of a range of high harmonics of the laser frequency. The initial observation of high harmonics was striking in that a span of harmonics were emitted which extended to high orders with almost constant intensity, completely at odds with lowest order perturbation theory. This early observation of high-harmonic generation in Chicago was followed up by experiments from the Saclay group in France which observed a number of curious trends in the character of the harmonic spectra and confirmed the formation of a "plateau" in the harmonic spectrum over a large number of harmonic orders (L'Huillier and Balcou 1993). This line of work by Anne L'Huillier led to her 2023 Nobel Prize in Physics.

In fact, not long after the first high harmonic experiments in the mid 1990s very high nonlinear orders, >100 , were reported (Chang *et al.* 1997), resulting in the production of coherent light well into the soft X-ray region. These experiments demonstrated, in a dramatic manner, that at the intensities now available, quantum mechanical multiphoton processes with hundreds or even thousands of photons were possible. High-harmonic generation (or just "HHG") continues to be studied actively. One of the most mysterious aspects of these high harmonic studies was explained nearly simultaneously in 1993 by Ken Kulander at Lawrence Livermore National Laboratory (Kulander *et al.* 1993) and Paul Corkum at The National Research Council Canada (Corkum 1993). Both surmised that the extent of the so-called harmonic plateau could be explained by a relatively simple quasi-classical model of the laser-driven electrons in an ionizing atom. This "simple-man's" semiclassical treatment is now the basis for much of our understanding of strong field ionization, above-threshold ionization, and high-harmonic generation.

For a time, a push to produce the shortest wavelengths possible by high-harmonic generation drove the research field. However, research in high-harmonic generation was

reenergized by a number of remarkable proposals in the mid and late 1990s that suggested that the broad, coherent spectrum of the high harmonic comb could lead to generation of pulses with duration under 1 fs, in the attosecond regime. The first attosecond pulse was demonstrated in the lab (with pulse duration of 650 as) by Ferenc Krausz *et al.* in Vienna (Hentschel *et al.* 2001). This work led to Krausz's share of the 2023 Nobel Prize in Physics. This demonstration has essentially spawned an entire subfield of ultrafast research devoted to generating and using ever-shortening XUV pulses produced by controlled HHG (Krausz and Corkum 2002). Pulses under 100 as are now routinely generated and characterized in labs around the world.

Another important development in strong field atomic physics arose in 1992 with the experimental observation at Lawrence Livermore National Laboratory of nonsequential double ionization of He atoms in femtosecond 800 nm pulses at intensity around 10^{15} W/cm² (Fittinghoff *et al.* 1992). (In fact, the experimental signature of this effect was apparent in ionization experiments performed in France as early as the mid 1980s (L'Huillier *et al.* 1982).) This observation challenged the long-held assumption that intense lasers interacted almost solely with one electron (the most loosely bound electron) at a time, resulting in sequential, uncorrelated stripping of the electrons from the atom as the light intensity increased. Details of this effect were illuminated in a classic, careful experiment performed by Walker and DiMauro shortly after the report of Fittinghoff (Walker *et al.* 1994). Though innumerable theoretical and computational studies have been performed to understand this multielectron physics, it turns out that the simple quasi-classical model of Kulander and Corkum describes this effect for the most part. The recollision of a tunnel-ionized electron driven back into the parent ion by the oscillating laser field can liberate a second electron. This phenomenon is now well understood, and there are a host of experimental results in the literature describing the various nuances of this effect.

The vast majority of effort in strong field studies in the atomic regime has focused on ionization and high-harmonic generation from single atoms. However, a number of fascinating effects of strong field interactions with small molecules have been investigated since the 1980s. Early studies examined the fate of molecular bonds in modest laser intensity, exploring the so-called bond softening that results from the interaction of the laser field with the bonding electrons. A particularly puzzling observation made in the late 1980s in studies of the explosion of diatomic molecules multiply ionized by an intense laser was the topic of much discussion in the strong field community for a few years in the 1990s. Experiments on the Coulomb explosion of diatomic molecules were initially inexplicable: the energies of the ions ejected from the Coulomb explosion almost always seemed to occur at an ion separation somewhat larger than the equilibrium separation of the molecular bond, an effect which appeared to be largely independent of laser pulse duration (Schmidt *et al.* 1994). This mystery was solved later in the 1990s when it was realized that molecules preferentially ionize when the nuclear separation increases to a so-called "critical separation" (Chelkowski and Bandrauk 1995; Posthumus *et al.* 1995).

To a certain extent, strong field ionization of molecules continues to confound investigators to some degree. Molecules almost always tend to show ionization rates that are

substantially lower than atoms with similar ionization potentials (Talebpour *et al.* 1996). This effect seems to be partially understood in the context of modern tunneling and strong field approximation theories that account for the extended, nonspherical nature of the wavefunction of bonding and antibonding molecular orbitals. However, to a large extent, the specifics of this effect are not yet understood and anomalous behavior in molecular ionization continues to be observed. It would seem that multielectron physics is much more important in molecular strong field interactions than it is in atoms, and this has challenged modern computational simulations of the interactions.

For a time starting in 1993 with inexplicable results out of the Rhodes group at the University of Illinois Chicago (McPherson *et al.* 1993), there was a period of interest in the study of strong field interactions of laser pulses with clusters of atoms. The first anomalous experimental signature was copious X-rays emitted when gas jets that carried these large clusters were irradiated at intensity $>10^{16}$ W/cm². It was soon realized by Ditmire and coworkers at Lawrence Livermore National Laboratory (hereafter denoted as LLNL) that, in fact, these small clusters actually formed small “nanoplasmas” under intense irradiation that could efficiently absorb the incoming laser light and lead to bright X-ray emission and fast ejection of highly charged ions (Ditmire *et al.* 1995). For the most part, intense near-IR laser pulse interactions with clusters are now largely understood in the context of this nanoplasma model, and, at this time, most research on clusters has shifted to study of these clusters in intense XUV and X-ray pulses. A number of exotic phenomena have been observed in these strong field cluster interactions including the production of D₂ nuclear fusion neutrons in a gas of deuterium clusters (Ditmire *et al.* 1999).

While the study of strong field physics had its origins in the study of atomic ionization, it was also realized early on that strong field interactions with plasmas would manifest unique effects, not only through the ionization of atoms and ions in the plasma but in the collective motion of the plasma electrons driven by the strong forces of an intense laser pulse. Strong field studies in plasmas paralleled the atomic and molecular physics studies of the 1980s, 1990s and 2000s.

One of the first pioneering proposals for exploiting strong field interactions in underdense plasmas came with the classic paper of Tajima and Dawson in 1979 in which they proposed accelerating electrons in the wake of a plasma wave set up by the passage of an intense laser pulse through the plasma (Tajima and Dawson 1979). Using the ponderomotive forces of the intense light field, they surmised that very high accelerating gradients could be created in the traveling wave behind a laser pulse which could accelerate electrons to very high (GeV) energies over distances of only a few centimeters. This proposal sparked a vigorous research effort into this so-called “wakefield acceleration.” At the time of this writing nonlinear plasma waves produced by femtosecond pulses at intensity $>10^{19}$ W/cm² propagating through ionized gases have accelerated electrons to energies of ~ 10 GeV over lengths of only a couple of centimeters (Gonsalves *et al.* 2019). The acceleration in highly nonlinear waves, now referred to as bubble acceleration, has been exploited in experiments around the world.

Another flurry of activity in this field was initiated around 1990 when it was proposed by Burnett and Corkum that strong field ionization of atoms in a gaseous target might be

able to set up the conditions for gain in the soft X-ray region (Burnett and Corkum 1989). Though a robust, high-gain, field-ionized recombination X-ray laser was never really realized experimentally, this line of research has led to demonstration of a number of compact femtosecond laser-driven XUV and soft X-ray lasers based on other schemes. About the time that the community was exploring intense ultrafast production of plasmas in gas targets, a remarkable and completely unexpected observation was made at the University of Michigan when intense femtosecond pulses were weakly focused in air (Braun *et al.* 1995). They observed an extremely long plasma filament produced by the laser, extending for many meters, down the hall from the laser lab. This effect was, at first, unexplained but is now understood to result from an interplay between self-focusing of the intense pulses in the air and refraction of the production of the plasma, yielding a “moving” focus for different slices of the pulse. Many subtle aspects of this visually arresting effect have been elucidated in experiments over the past 25 years.

Strong field interactions in underdense plasma have led to the observation of other nonlinear phenomena, particularly when laser intensities have entered the relativistic regime. These phenomena include relativistic self-focusing, self-phase modulation and nonlinear forward-directed Raman scattering in gaseous plasma targets. Furthermore, exotic effects such as betatron X-ray emission from oscillating electrons in nonlinear plasma wakes have been observed in these experiments.

Finally, studies of high-intensity laser interactions with solid targets go back to the earliest days of the laser. After “giant” laser pulses were produced by Q-switching in the early 1960s, studies of the explosion of plasmas created by irradiation of a solid target were published. Experiments in the strong field regime began in earnest with the availability of joule-class, subpicosecond lasers from the invention of CPA in the mid 1980s. Most early studies were essentially phenomenological experimental studies of the X-rays produced by these solid target plasmas. However, studies of high-intensity laser interactions with solid targets saw an enormous increase in activity in the early 1990s with the proposal of the so-called fast ignition concept. This idea came soon after a remarkable initial study by a group at Stanford (Kmetec *et al.* 1992), followed by experiments at Berkeley and LLNL in the early 1990s that showed that very efficient generation of MeV “hot” electrons and MeV photons accompanied irradiation of solids at intensity approaching 10^{18} W/cm². Study of the collisionless absorption mechanisms that lead to these multi-MeV hot electrons has been the topic of many high-intensity laser-plasma studies for 20 years since the observation of efficient hot electron production in the 1990s.

The fast ignition concept inspired by this research was forwarded by Max Tabak and coworkers at LLNL (Tabak *et al.* 1994). Tabak’s idea suggested that an intense picosecond laser could produce a high-energy (many joules) burst of multi-MeV electrons which could be injected into the compressed fuel of an inertial confinement fusion (ICF) implosion. These hot electrons could serve to heat the compressed fuel and ignite it, triggering a fusion ignition burn.

The fast ignition proposal energized the high-intensity laser-plasma community and essentially led to 20 years of extensive research on the production and propagation of hot electrons in solid targets irradiated at relativistic intensity. At the time of the writing of this

book, it is generally believed that the hot electron generation efficiency and behavior of the high-peak-current hot electron bunches in the compressed plasma of an ICF experiment are not favorable for achieving ignition with any reasonable amount of short pulse laser energy (say <200 kJ). Nonetheless, the worldwide research on high-intensity laser production of hot electrons has led to a comprehensive understanding of how intense laser pulses couple to electrons in an overdense plasma and how such electrons propagate. This has resulted in many studies of exotic phenomena such as kilotesla magnetic field generation and positron production in the laser plasma (Cowan *et al.* 1999).

Fast ignition and hot electron research also propelled the development of high-peak-power laser technology as it was clear early on that fast ignition would demand picosecond pulsed lasers with peak power over 1 PW. The late 1990s saw the demonstration of the first petawatt laser at LLNL on the NOVA ICF laser, a project motivated in large part by the fast ignition idea (Perry *et al.* 1999). This first demonstration of a petawatt of peak power has led to a proliferation of PW peak power lasers around the world, and an explosion of strong field research at relativistic intensities. After the large aperture grating development at LLNL which enabled that first PW laser on NOVA, a handful of large Nd:glass-based PW lasers emerged at national laboratories such as the Rutherford Appleton Lab in the UK and the ILE in Osaka, Japan. The past 25 years have seen numerous additional PW lasers at scales from 40 J to 500 J see completion; currently dozens of laser labs around the world operate lasers with power ~ 1 PW. In fact, a number of lasers at peak power near or in excess of 10 PW have now been constructed.

A particularly active research area which spun out of the research on fast ignition is in proton and other ion acceleration. The PW laser on NOVA at LLNL made a surprising observation of multiple tens of MeV proton ejection from thin metal targets at intensity above 10^{20} W/cm² in some of that laser's early experiments (Snively *et al.* 2000). Fast ion ejection from pulsed laser irradiation of solid target plasmas was, in fact, a well-known phenomenon, with observations dating back to the earliest laser plasma experiments in the early 1960s. However, the LLNL PW results were remarkable in the high energy of the observed protons (>10 MeV) and the efficiency with which these protons were produced (with ~ 10 percent of the total laser energy emerging in the pulse of fast protons from the back of the target). After a short period of controversy about the mechanism for this surprising hot proton ejection, it is now well established that these protons arise from the hot electrons produced in a solid target interaction and the sheath fields these electrons produce as they attempt to exit out the back side of a thin foil. This so-called target normal sheath acceleration (TNSA) is now well understood and well characterized.

The broad proton spectra that are produced by TNSA has prompted, since 2000, a vigorous research effort into alternate ion acceleration mechanisms which might yield higher proton energies than TNSA and which could produce nearly monoenergetic bursts of MeV protons or heavier ions. This research has been much of the impetus to understand solid density plasma interactions at highly relativistic intensity where radiation pressure becomes important. Currently research into ion acceleration by PW-class lasers is among the most active areas of strong field plasma physics research worldwide. The fast electrons produced in PW-class laser interactions with solids led to the triggering of various nuclear reactions

in the targets (Cowan *et al.* 2000). Such laser-induced nuclear reactions have been one of the many byproducts of these interactions which have presented the promise of laser-driven nuclear applications (like deployment of compact neutron sources or radioactive transmutation). Applications like this drive much of modern research in the field. Laser-driven ion acceleration has led to something of a new renaissance in strong field laser-plasma studies and a rebirth in interest in fast ignition. It is now thought that the production of intense proton bursts could be an alternate, viable way to fast ignition (Roth *et al.* 2001). At this writing this research avenue is again gaining momentum because of the recent demonstration of fusion ignition at the LLNL NIF ICF facility and the promise of laser-driven fusion energy (Abu-Shawareb *et al.* 2002).

The past two decades the past two decades have seen strong field research advance in many ways. Particularly exciting progress continues in the manipulation of attosecond extreme ultraviolet pulses via high-harmonic generation. Relativistic intensities are now regularly generated by many labs and laser-plasma interactions in this regime are being better understood all the time. Dramatic progress has been made in laser-driven plasma acceleration of electrons since 2010. In fact, relativistic effects, where the change in the mass of the electron which comes from its acceleration by the laser field to velocity approaching the speed of light within one optical cycle are now regularly seen in laser plasma experiments at intensity above 10^{20} W/cm². This electron mass increase leads to exotic plasma physics such as relativistic plasma transparency where a normally optically opaque dense plasma becomes transparent to the laser light because of the electron mass shift, or relativistic self-focusing where the optical properties of the plasma are changed by this mass change, focusing the laser light. In fact, QED effects might start playing a role in laser plasma interactions at the extreme intensities now attainable. At the time of writing this book, the experimentally obtainable intensity frontier is approaching 10^{23} W/cm². Intensities exceeding 10^{24} W/cm² should be reached within the next decade.

1.3 Outline of This Book

This book is intended to introduce many of the fundamental concepts underlying modern strong field physics research. These concepts span descriptions of intense light interactions with single electrons, individual atoms, ensembles of atoms in molecules and clusters, and many charged particles in plasmas. This book does not represent a comprehensive review of modern strong field physics research and is not a survey of recent results in the field. No attempt is made to discuss specific experimental results that confirm the phenomena presented (though citations to such work are often given). Instead, the basic phenomena underlying the more complex effects observed in strong field physics will be discussed, and the basic equations needed to describe these high field effects will be derived. If a more detailed review of the various aspects of strong field physics is desired, there have been a number of excellent review articles published in recent years and a number of focused textbooks. A listing of some of these complementary books can be found in the bibliography.

The book begins with a comprehensive review of the technology employed to access strong field physics. Then, the text, in a series of chapters, marches through discussions of high-intensity laser pulse interactions with systems of increasing complexity. This begins with an examination of strong field interactions with free electrons, in which the considerations are essentially classical but demand relativistic dynamics descriptions. This is followed by discussions on interactions with atoms and then small molecules, topics which mandate a dive into quantum descriptions of the interactions, though I attempt to utilize quasi-classical models when I can. Building on these quantum strong field models, the book then turns to a key aspect of strong field nonlinear optics, generation of high-order harmonics of the laser field in gases of atoms and molecules. After this, the book essentially moves to many-body systems in which understanding relies heavily on concepts in plasma physics. This set of chapters begins with a discussion of high-intensity interactions with microscopic clusters of atoms. The chapters then include a consideration of strong field interactions with macroscopic scale plasmas, first with underdense plasmas which are low enough density to allow light propagation within them and then concluding with a chapter on interactions with solid density plasmas which are overdense and reflect the light wave. This ordering of chapters is designed so that models of the more complex systems can be built from the physics explored in earlier chapters on smaller systems. I attempt to write each chapter such that it can essentially stand alone as a comprehensive description of that aspect of strong field phenomena; however, the utilization of models from earlier chapters as building blocks for models describing the more complex systems does mean that the reader will benefit from a sequential study of the material in the various chapters.

1.4 Units, Variables and Mathematical Notation

For the entirety of the book, I use CGS units, unless otherwise stated. When mixed units are employed, those units will be listed after the variable in brackets (e.g. intensity, I [W/cm²]).

All scalar variables will be denoted by symbols in italic text. Vectors (traditional 3-vectors) will be given as bold-faced text symbols (e.g. momentum, $\mathbf{p}$). When a vector is considered, if the same symbol is employed in nonboldface, italicized text, this will be implied to mean the magnitude of that vector (e.g. $|\mathbf{p}_0| = p_0$). When 4-vectors are denoted, they will be given by italicized nonboldface text symbols with Greek symbol subscript/superscript counters (e.g. the four-momentum is p_μ with contravariant counterpart p^μ).

I will attempt to hold to as many of the widely used conventions in naming variables as possible. Because of the frequent use of a particular symbol in different senses in the literature (e.g. γ can denote the relativistic factor in special relativity, the Keldysh parameter, or the ratio of specific heats in thermodynamics), we will make liberal use of subscripts in naming many variables. To aid the reader with the bewildering array of variable symbols, an extensive (though incomplete) list of variable symbols is provided in the end-of-book Appendix. Because of the extensive employment of variables describing the laser field

throughout all chapters in the book, to minimize clutter, the symbols used to describe the frequency, wavelength and wavenumber of the incident laser will always be written without a subscript (i.e. the laser's wavelength, wavenumber, frequency and angular frequency will always be written λ , k , ν , ω). The mass of the electron will always be denoted by an unsubscripted m ; the masses of other particles will be written with a descriptive subscript (e.g. mass of the proton: m_p). Generally speaking, when any frequency is denoted with an ω_{whatever} it is meant to denote an “angular” frequency which is $2\pi \times$ a cycle frequency, which will always be denoted as ν_{whatever} .

Inevitably I will make approximations in our discussions. I will employ the $=$ symbol to mean exactly equal. When $\cong$ is employed it will mean that the two sides are approximately equal and the approximation made is meant to be rather accurate, which means the reader should interpret the sides of the equation as nearly quantitatively equivalent (e.g. $a_0^2 + 1 \cong a_0^2$ for large a_0). When the symbol $\approx$ is used, the implication is that the two sides of the equation are only roughly equal, perhaps to no better than an order of magnitude. Finally, when $\sim$ is used to relate two mathematical expressions, it means that the left-hand quantity simply scales as the right-hand side, with prefactors dropped for illustrative purposes (e.g. $\text{energy} \sim v^2$). Hence the two sides will not necessarily be numerically equal to the same order of magnitude or even have the same units.