

Chapter 1

Introduction: Historical Background

The field of high-intensity laser interaction with matter, although barely two decades old, is already bursting with enough exotic phenomena to keep researchers busy for years to come. Since the invention of ‘chirped pulse amplification’ in 1985, progress in short pulse laser technology has been unrelenting, see Fig. 1.1. Pulse durations have come down from a picosecond to less than 5 femtoseconds (10^{-15} s); whereas focused intensities have skyrocketed six orders of magnitude. At present, several laboratories around the world are now promising intensities in excess of 10^{21} Wcm $^{-2}$.

In view of the impending escalation suggested by this diagram, a shift into the uncharted territory of highly relativistic laser particle physics, it is perhaps appropriate to take a brief look at the various disciplines from which this new field has grown. These contributory fields are numerous and diverse: they include laser physics, atomic physics, plasma physics, and lately even astrophysics and elementary particle physics. Many of the theoretical models described in later chapters can be traced back to one of these more classical areas. This does not mean that researchers in this business have been able to get their theories or ideas ‘off the shelf’ — on the contrary: the extreme conditions under which light and matter are forced to coexist during such interactions have posed a continual challenge to both theoreticians and experimentalists alike.

The key to understanding the underlying physics in these interactions is to realize that ordinary matter — whether solid, liquid or gas — will be rapidly ionized when subjected to high intensity irradiation. The electrons released are then immediately caught in the laser field, and oscillate with a characteristic energy (the right hand column in Fig. 1.1) which then dictates the subsequent interaction physics.

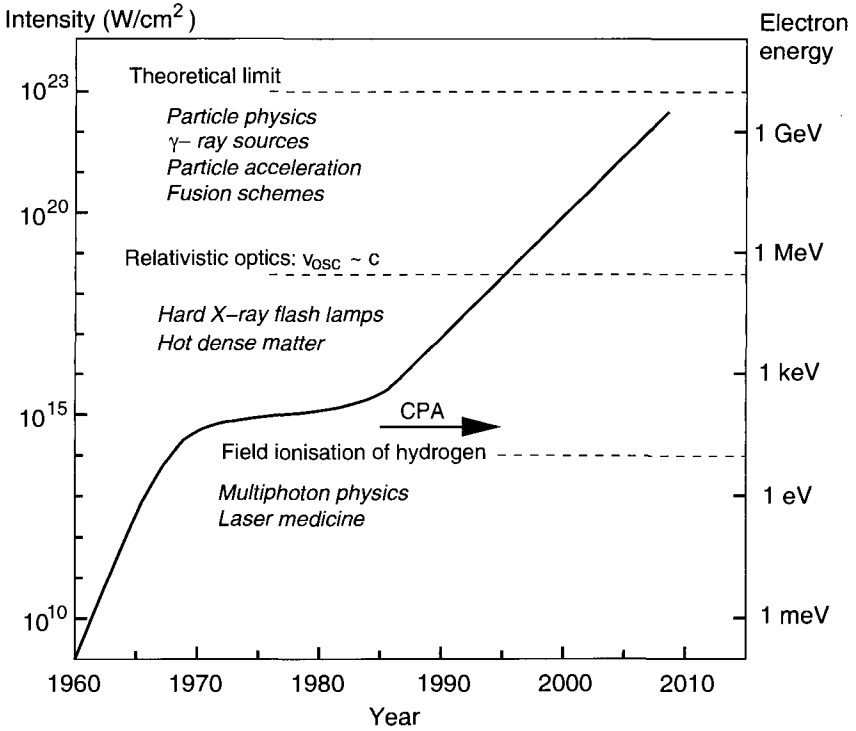


Fig. 1.1 Progress in peak intensity since the invention of the laser in 1960.

1.1 Multiphoton Physics

One of the most familiar examples of light inflicting a change in material properties is the photoelectric effect, predicted by Einstein a century ago (Einstein, 1905) and validated experimentally by Millikan a decade later (Millikan, 1916). This process, the ejection of an electron from an atom by a single photon, occurs when the photon energy $\hbar\omega$ matches the height of the atomic potential barrier, I_p , which the electron experiences in the vicinity of the ion, i.e. $\hbar\omega = I_p$. Even for the outer shells of most atoms, this energy still runs into many electron-volts, equivalent to photon wavelengths well into the ultraviolet (UV) range. For inner shells ($I_p \sim \text{keV}$), one needs hard x-rays to induce photo-ionisation.

With standard lasers (operating wavelengths $0.25 \mu\text{m} - 13.4 \mu\text{m}$), one cannot observe the photoelectric effect on normal material because $\hbar\omega \ll$

I_p . As lasers grew more powerful in the 1960s and 70s, however, (Fig. 1.1) it became possible to consider *multiphoton* ionisation, expressed by the condition

$$n\hbar\omega = I_p.$$

Thus, instead of one very energetic photon, an electron absorbs n photons of moderate energy (for example, laser photons with $\hbar\omega \approx \text{eV}$) and is subsequently ejected. Some of the pioneering work on this was done at CEA-Saclay, France by the group of Mainfray and Manus (1991), including the discovery of *above-threshold ionisation* by Agostini *et al.* (1979). Another ground-breaking event during this period, prompted by the development of chirped pulse amplification, was the prediction of *high-harmonic generation* via recombination of multiphoton-ionized electrons by Shore and Knight (1987). This phenomenon was confirmed experimentally soon afterwards (McPherson *et al.*, 1987; Ferray *et al.*, 1988) by a intensive series of experimental investigations in a number of labs worldwide. Today it is possible to generate hundreds of harmonics down to nanometre wavelengths, a feat which has heralded a new era of highly compact, coherent XUV light sources.

1.2 Single-Electron Interaction with Intense Electromagnetic Fields

This subject actually goes much further back than most newcomers to the femtosecond business would probably guess; predating the invention of the laser at the beginning of the 1960s by some margin. One of the first analyses of the behavior of free electrons in the presence of intense radiation was made by Volkov (1935), who introduced the concept of a *dressed state* to describe the enhanced inertia experienced by an electron when oscillating in an electromagnetic field. Later, motivated by the first experiments with synchrotrons, Schwinger made a detailed analysis of the radiated power emitted by accelerated electrons (Schwinger, 1949), pointing out that radiation is preferentially emitted in the direction of motion at high energies.

These early ideas were refined and explored with more urgency when the invention of the laser brought the prospect of an experimental means to study relativistic photon-electron physics in the laboratory (Brown and Kibble, 1964; Vachaspati, 1962; Eberly and Sleeper, 1968; Sarachik and Schappert, 1970). These authors defined the figure of merit for ‘laser-

electron' interaction by a dimensionless parameter q (or q^2), given by:

$$q = \frac{eE_L}{m\omega c}, \quad (1.1)$$

where e , m and c are the electronic charge, electron mass and speed of light respectively; E_L is the laser electric field strength and ω the light frequency. Needless to say, these theoretical works all lamented the impossibility of achieving truly relativistic conditions ($q > 1$) with the optical lasers available at the time, but speculated that they might be might one day be reached with 'future' technology. Forty years later, this wishful thinking has become reality, and in Chapter 3 we will see how these visionary models have inspired new lines of experimental investigation.

Independently of the laser's arrival, astrophysicists were beginning to suggest mechanisms for cosmic ray generation in the vicinity of pulsars via the interaction of intense electromagnetic (EM) radiation with free electrons (Ostriker and Gunn, 1969; Gunn and Ostriker, 1971). The numbers involved here are of course vastly different from laser-plasma interactions: pulsar radiation has a frequency between 0.3 and 30 Hz, and magnetic field strengths near the star surface in the region of 10^{12} Gauss; causing gargantuan oscillation amplitudes of particles in the surrounding plasma. The interaction physics is very similar however, and is just one of many instances of *scalable laboratory astrophysics*, where one can emulate the conditions in astrophysical objects using high-power lasers.

1.3 Nonlinear Wave Propagation

The fact that plasmas can support large-amplitude, nonlinear waves has been known for almost as long as plasma physics established itself as a mainstream branch of science. Early seminal works by Akhiezer and Polovin (1956) and Dawson (1959) set the scene for numerous studies on the behavior of both large-amplitude Langmuir (electrostatic) waves and the propagation of high-intensity electromagnetic radiation in plasmas (Davidson, 1972). One attraction of this then rather obscure field was the tantalizing possibility of producing long-lived solitons in plasmas (Zakharov, 1972; Decoster, 1978; Shukla *et al.*, 1986).

The publication of a method for laser-acceleration of electrons in underdense plasmas by Tajima and Dawson (1979) sparked off a fresh wave of interest in wave propagation. This enthusiasm also encompasses members of the accelerator community, who are actively on the look-out for

viable alternatives to conventional linear accelerators and synchrotrons as these devices approach their physical and economic limits. At the time of writing, the world record for laser-plasma acceleration of electrons lies at around 350 MeV, not breathtaking compared to the 50 GeV per electron/positron beam at SLAC, until one realizes the former was achieved with a few millimetres of plasma rather than several kilometres of super-cooled accelerator structure. Although there is some way to go before laser-plasma acceleration can compete with existing facilities, scaling up to the Petawatt powers now available will soon make the tabletop GeV electron accelerator a reality.

1.4 Metal Optics

The story of laser interactions with solids also has a nineteenth century prologue. The simple observation that polished metals behave as almost perfect reflectors, whereas other materials either absorb or transmit light, could not be satisfactorily explained until Drude set out his ‘Electron Theory’ (Drude, 1900). Although solid state physics has advanced beyond recognition since then, the so-called *Drude model* of electron conduction still retains its appeal and usefulness in describing the main features of metal optics (see, for example: Ashcroft and Mermin, 1976, Chapter 1).

Drude’s original idea — just three years after J. J. Thomson’s discovery of the electron — was to suppose that the atoms in a metal somehow share a limited number of ‘valence’ electrons, forming a conduction band. These can wander as far as they like from their parent atoms, carrying current and heat through the material in the process. For an element with mass density ρ and atomic weight A , the free electron density is given by $n_e = N_A Z^* \rho / A$, where N_A is Avogadro’s constant and Z^* is the number of valence electrons per atom.

The conductivity of a metal will depend on the rate at which these free electrons are slowed due to collisions with the ions. Mathematically this can be expressed by the relation: $\sigma_e = n_e e^2 \tau / m_e$, where τ is the collision or relaxation time. Even today, precise theoretical treatments to determine the relaxation time in solids remain a challenge, since the latter depends on details of the crystal structure, electronic configuration and so on. For this reason, nearly all the available data on metallic conductivity has been obtained through experimental measurements using Ohm’s law: $\mathbf{j} = \sigma_e \mathbf{E}$. This technique — applying a DC voltage and finding the

resulting current — can be recast in a form suitable for *optical* diagnostics via the AC conductivity:

$$\sigma(\omega) = \frac{\sigma_e}{1 - i\omega\tau}.$$

When combined with Maxwell's equations, this leads to a complex dielectric constant:

$$\epsilon = 1 - \frac{\omega_p^2}{\omega(\omega + i\nu)},$$

where

$$\omega_p^2 = 4\pi n_e e^2 / m_e$$

is the *plasma frequency* of the valence electrons, and $\nu \equiv \tau^{-1}$ is their relaxation rate or collision frequency. This expression basically predicts that a metal will only become transparent for radiation with wavelengths $\lambda < \lambda_p = 2\pi c / \omega_p$, which turns out to be in the UV (200–400 nm) range. Standard laser light (0.5–10 μm) will be reflected or absorbed, depending on the collision time.

Now, anyone with some basic plasma physics can see that the Drude model is almost identical to the standard theory of collisional laser absorption in plasmas, except that in the latter case, τ *can* be calculated with some accuracy. Suffice to say that metal optics provides an excellent starting point for studying reflectivity (and other transport properties) of short pulse laser-produced plasmas (Godwin, 1972). In Chapter 5 we will see how this leads to an almost seamless transition between the solid and plasma states of matter.

1.5 Long Pulse Laser-Plasma Interactions

A common misconception about femtosecond laser-plasma research is that it has little or nothing to do with the long pulse laser-plasma interaction issues relevant to inertial confinement fusion (ICF). It is certainly true that one cannot do ICF with a table-top laser (at least not yet, anyway!). The scientific (and often political) demarkation which prevails today masks a broad thematic overlap, and there are plenty of researchers who happily jump back and forth between the two fields on a daily basis. Whatever one's personal view of ICF as a future energy source, and its unavoidable relationship with nuclear weapons programs, it has to be said that the

advances in the short pulse field would have been nowhere near as rapid without the considerable prior scientific and technical knowledge of laser science and interaction physics generated by the ICF programs over the last 30 years. Luckily, training in ICF physics is not necessary to work on short pulse interactions, but it is nonetheless useful to know where to find the original literature, on, say, hydrodynamics, parametric instabilities, energetic particle generation, and so on. A brief outline of ICF would therefore appear to be in order, even if to mainly draw contrasts between the new femtosecond phenomena and this *long pulse regime* in later chapters.

Laser fusion became official in 1972 (having previously been under military classification) with the publication of a classic but over-optimistic paper in *Nature* by Nuckolls *et al.*, (1972). In this work, the authors describe how a small micrometer-sized pellet filled with deuterium and tritium fuel can be compressed to enormous densities by irradiating it with laser beams focused symmetrically onto its surface (Brueckner and Jorna, 1974). By converting the laser energy into thermal plasma energy, the pellet shell material ablates radially outwards, thus pushing the fuel inwards via a rocket-like reaction. By virtue of the spherical symmetry, the fuel implodes, eventually reaching densities ρ of several hundred gcm^{-3} and temperatures T of 10 keV (10^7 degrees Kelvin), thus meeting the requirements for thermonuclear confinement encapsulated by the product:

$$nT\tau > 10^{15} \text{ keV s cm}^{-3}. \quad (1.2)$$

This condition — known as the *Lawson criterion* — basically states that the fuel must burn up and release its energy before the capsule blows apart, leading to a requirement for the *areal fuel density* ρR , where R is the final capsule radius.

The standard *hot-spot* scenario (Lindl, 1995) currently being pursued by the major laboratories in the USA, France and Japan, requires a central $\rho R = 0.3$, which, after allowing for all the inefficiencies and target design considerations, translates into a laser energy requirement of around 1 MJ to achieve gain. We will return to these constraints later in section 7.5, where some speculative new ideas to reduce this driver requirement with the help of an additional short pulse laser are described.

These more recent schemes aside, the whole process of target ablation, implosion and ignition is essentially determined by hydrodynamics: any plasma physics involved is almost invariably destructive, putting additional constraints on the laser and target design. These *coronal* processes (so-called by analogy with stellar objects) are illustrated in Fig. 1.2: as we




Fig. 1.2 Laser-plasma interactions in the corona of an imploding micro-balloon.

shall see in later chapters, nearly all of them turn up again in femtosecond interactions. Parametric instabilities such as Raman and Brillouin scattering (Kruer, 1988) are generally bad for the implosion because they generate fast electrons. Because of their long range, these electrons preheat the target core, making the compression less efficient. Resonance absorption (which we will meet in Sec. 5.5.1) is also undesirable for this reason: in fact, up to 50% of the laser energy can be wasted on superheated electrons in this way. The classic signature of such ‘anomalous absorption’ processes is a bi-Maxwellian electron distribution, see Fig. 5.17. In the case of resonance absorption, the hot electron component has a temperature T_H given by (Forsslund *et al.*, 1975a; Estabrook *et al.*, 1975):

$$T_H \simeq 14 (I\lambda^2)^{1/3} \text{ keV}, \quad (1.3)$$

where I is the laser intensity in units of 10^{16} Wcm $^{-2}$ and λ the laser wavelength in microns. The dependence on the product $I\lambda^2$ (originating in the quiver motion of the electrons in the laser field) means that long wavelength lasers (e.g.: CO $_2$) tend to produce hot electrons much more readily (than, say KrF) for a given intensity. This was a major factor in the strategic switch to shorter wavelengths ($\lambda < 0.5\mu\text{m}$) in ICF programs at the beginning of the 1980s. A major advantage of *indirectly* (X-ray) driven schemes is that hot electron generation is almost completely avoided (Lindl, 1995).

We have already seen that the rules of the game describing ultrafast, ultra-high-intensity (UHI) interactions have to be rewritten. For example, the short timescales make it necessary to discard steady-state or adiabatic models in favour of fully dynamic non-equilibrium formulations. On the other hand, some of the computational tools developed over 30 years ago to investigate hot electron generation in nanosecond laser-plasma context are used today to model femtosecond interactions at intensities of 10^{21} Wcm $^{-2}$, not least in the *fast ignitor* context (Sec. 7.5). In fact many of the codes used for UHI interactions can be traced back to ICF applications: atomic physics packages, fluid codes, particle-in-cell or Vlasov codes. These have usually been adapted to meet contemporary demands: for instance by using time-dependent ionization models to cope with the breakdown of local thermal equilibrium (LTE); by allowing for relativistic electron dynamics; or by introducing wave propagation into hydrodynamic models to handle the laser pressure and absorption more accurately. Details of these models are deferred to Chapter 6, but note that there are many areas — including many optical and diagnostic techniques — where the connection with ICF can turn out to be very fruitful.

1.6 Femtosecond Lasers

Hopefully, it will be clear by now that the intention of this book is not to supply assembly instructions for a Terawatt laser system, but rather to examine the physics which can be explored with the help of such a device. The material contained in the chapters which follow is thus very much aimed at those researchers either physically or mentally gathered around the target chamber (see Fig. 5.15), whose knowledge of the laser operation consists primarily of the four magic numbers: wavelength, energy, pulse duration and focal spot size. That said, the only way to appreciate where

these numbers are conjured up from is to lift the lid off the box and take a peek at the optical wizardry inside.




Fig. 1.3 Front end of the Jena Ti:Sapphire femtosecond laser.

If we do this, we will find something resembling a huge BRIO-BuilderTM set with an awful lot of glass in it, see Fig. 1.3. Make no mistake, though: these devices represent the state-of-the-art in optical engineering.¹ We can make more sense of this optical jungle with the help of the schematic diagram in Fig. 1.4, where we notice that the ‘laser’ actually comprises several autonomous units, namely: an *oscillator*, producing low-energy short pulses at regular intervals; a *stretcher*, converting the femtosecond pulse into a 50–200 ps *chirped* pulse; one or more amplifier stages to increase the pulse energy by a factor of 10^7 – 10^9 ; and finally a *compressor*, performing the exact optical inverse of the stretcher to deliver an amplified pulse of the same duration as the oscillator.

The key technique in this business is *chirped-pulse amplification* (CPA), which since its first application to lasers by Gerard Mourou and co-workers in 1985 (Strickland and Mourou, 1985; Maine *et al.*, 1988), is now exploited by nearly all existing high-power, femtosecond systems. At first sight, the procedure seems unwarranted: why bother with all the stretching and compressing when one could amplify the short pulse from the oscillator straight away? The answer is simply that most of the optical components will

¹A tip for theoreticians being shown around a laser lab for the first time: resist the temptation to fiddle with the pieces on the table — it can take several days to put right again.




Fig. 1.4 Schematic diagram of the Jena TW laser.

be overheated and damaged if the fluence exceeds a level of around $0.16 \text{ Jcm}^{-2} \tau_{ps}^{1/2}$, where τ_{ps} is the pulse duration in picoseconds. For maximum efficiency, however, the fluence should be near the saturation level of the amplifying medium, which is 1 Jcm^{-2} for Ti:sapphire. Therefore, the pulse length throughout the amplifier chain has to be at least 40 ps (typical values are 100–200 ps) so that the components (e.g. mirror coatings) have time to cool by thermal conduction.

The stretcher-compressor combination separates the pulse generation and amplification stages, permitting standard techniques to be used for the amplifier chain, and furthermore leaving room for more advanced phase-compensation devices at either end. Before we consider these refinements, let us return to the ‘front end’, where the pulse is generated. Femtosecond laser sources are *mode-locked*: the output pulse is made up of a superposition of many electromagnetic waves (or laser modes), and is *transform limited*, so that $\tau_p \sim 1/\Delta\nu$, where $\Delta\nu$ is the bandwidth. Clearly a large bandwidth is essential to generate a short pulse. Consider, for example, a 10 fs Gaussian pulse, for which $\Delta\nu\tau = 0.44$. This gives $\Delta\nu = 4.4 \times 10^{13} \text{ Hz}$, which for a central wavelength of 800 nm, translates to:

$$\Delta\lambda = \Delta\nu \frac{\lambda^2}{c} = 94 \text{ nm}.$$

This is an absolute requirement of the lasing medium in the oscillator, as well as the subsequent amplification stages. Most materials have bandwidths far below this, but titanium-doped sapphire crystals have a remarkable gain curve ranging from 700 nm to 1100 nm, with $\Delta\lambda = 230 \text{ nm}$; the largest bandwidth of any known material. This property, together with its high saturation fluence of nearly 1 Jcm^{-2} and high damage threshold, make Ti:sapphire the ideal choice for most femtosecond laser systems in operation today.

How do we create such a pulse in the first place? In particular, a means of converting a continuous-wave (cw) pump laser beam into a train of short pulses is needed. This is done by exploiting another fortuitous property of the Ti:sapphire crystal: namely, that it also acts as a nonlinear focusing lens. This so-called nonlinear Kerr nonlinearity (Shen, 1984) occurs when the intensity inside the crystal exceeds 10^{11} Wcm^{-2} , so that by blocking the unfocused cw mode with a slit, one ends up with a single, short pulse bouncing to-and-fro in the oscillator. As it stands, this arrangement is unstable because dispersion will cause the pulse to spread out longitudinally (red wavelengths are faster than blue ones). To compensate, a set of prisms

is inserted between the crystal and back mirror, which imposes an equal but opposite dispersion on the pulse (blue faster than red). The combination of cavity plus dispersion-correction constitutes a highly stable and reliable femtosecond laser, which these days can even be purchased ‘off-the-shelf’. The true art of femtosecond laser physics comes when we try to amplify the pulse, as we shall see shortly.

The fs pulse from the oscillator typically contains just a few nJ of energy; yet we wish to increase its power to the TW level and beyond. As already mentioned, to avoid damaging optical elements, the pulse must first be stretched by a factor of 10^3 – 10^4 or so. This is usually done with a pair of diffraction gratings, which impose a positive chirp on the pulse, so that longer wavelengths emerge before shorter wavelengths. Early CPA systems (Maine *et al.*, 1988) actually used enormous km-length optical fibres to broaden the bandwidth via self-phase modulation. However, the stretch-factors obtained by this method are limited because there is no practical means of correcting the nonlinear, high-order phase distortion introduced by the fibre, ultimately restricting the final pulse length to values of just under 1 ps.

The stretched pulse can now be amplified — a process which is usually split into two or more stages, depending on the final beam energy required. The Jena system comprises three different modules: a *regenerative* preamplifier and two multipass *power* amplifiers. Most of the gain of the system ($\sim 10^7$) is obtained by the regenerative amplifier, which works in a very similar fashion to the laser resonator itself. The difference is that it seeded by the chirped pulse, which subsequently makes up to 20 round trips through a low gain medium, before being switched out by a combined Pockels cell/polarizer combination. Since the regenerative amplifier eventually becomes saturated at a few mJ, additional cavity-free techniques are then used to amplify up to and beyond the 1 J level. The multipass configuration in Fig. 1.4 is typical of modern tabletop-TW systems, producing gains of 10^2 – 10^3 .

The amplified, long chirped pulse is then recompressed using a grating pair (or quadruplet in the Jena system shown), ideally reducing the pulse length back down to a value slightly above the one originally emitted by the oscillator. Some pulse lengthening is inevitable due to a combination of nonlinear dispersion effects and gain-narrowing. The latter arises because the amplifier medium preferentially enhances central wavelengths over peripheral ones, leading to a reduction in bandwidth and hence lengthening of pulse duration.

Due to the high-quality beam profile which can be produced by Ti:Sapphire amplifier systems, the pulses can be focused to an almost transform-limited spot containing most of the laser energy, see Fig. 1.5. This is the form of the laser pulse we want for interaction physics: the maximum possible photon density for a given wall-plug energy.




Fig. 1.5 Sharp end of the Jena Ti:sapphire femtosecond laser system: a focal spot of $3 \mu\text{m}$ diameter containing more than 50% of the pulse energy. The peak intensity reached here is $4 \times 10^{19} \text{ Wcm}^{-2}$.

Naturally there are variations of this generic short pulse system depending on the intended application. What I have just described is a multi-TW system typical for a university laboratory. For a high-end user facility, offering powers above 100 TW, one still has to go to one of the dedicated laser laboratories listed in Table 1.1. On the other hand, there is also a growing number of smaller, *high-repetition-rate* (kHz) systems for high throughput applications such as x-ray sources. More recently, state-of-the-art *few-cycle* (< 10 fs) lasers have opened up a new field of attosecond physics, a theme which crops up again in Chapter 7.

As seen in the table, there are currently three or four Petawatt Nd:glass lasers coming on line over the next two years (VULCAN has already passed the PW mark). These will deliver focused intensities of about 10^{21} Wcm^{-2} — perhaps a bit higher after some tuning and by increasing the energy — but only up to a point. The damage threshold on the compression gratings or focusing optics is around 0.3 Jcm^{-2} , so a Petawatt facility with 400 J pulse energy needs gratings with least 40 cm in diameter. To be on the safe side, the VULCAN laser at RAL uses 1 m-diameter gratings, right at the

Table 1.1 Multi-Terawatt laser systems and laboratories worldwide.

Name	Laboratory	Country	Type	λ (nm)	Energy (J)	Pulse length (fs)	Power (TW)	Focal spot (μm)	Intensity (Wcm^{-2})
Petawatt ^a	LLNL	USA	Nd:glass	1053	700	500	1300	–	$> 10^{20}$
VULCAN ^b	RAL	UK	Nd:glass	1053	423	410	1030	10	1.06×10^{21}
PW laser ^c	ILE	Japan	Nd:glass	1054	420	470	1000	30	10^{20}
PHELIX ^d	GSI	Germany	Nd:glass	1064	500	500	1000	–	–
LULI 100TW	LULI	France	Ti:Sa	800	30	300	100	–	–
APR 100 TW	APR	Japan	Ti:Sa	800	2	20	100	11	2×10^{19}
HERCULES	FOCUS	USA	Ti:Sa	800	1.2	27	45	(1)	(8×10^{21})
ALFA 2	FOCUS	USA	Ti:Sa	800	4.5	30	150	(1)	(10^{22})
Salle Jaune	LOA	France	Ti:Sa	800	0.8	25	35	–	10^{19}
Lund TW	LLC	Sweden	Ti:Sa	800	1.0	30	30	10	$> 10^{19}$
MBI Ti:Sa	MBI	Germany	Ti:Sa	800	0.7	35	20	–	$> 10^{19}$
Jena TW	IOQ	Germany	Ti:Sa	800	1.0	80	12	3	5×10^{19}
ASTRA	RAL	UK	Ti:Sa	800	0.5	40	12	–	10^{19}
USP	LLNL	USA	Ti:Sa	800	1 (10)	100 (30)	10 (100)	–	5×10^{19}
UHI 10	CEA	France	Ti:Sa	800	0.7	65	10	–	5×10^{19}

^a 1996–1999

^b Petawatt performance achieved on October 5, 2004.

^c Projected upgrade of PWM — PetaWatt Module.

^d Commissioned for end 2005.

limit of present manufacturing capability.

So just as with the old (non-CPA) pulsed technology, the only way to increase the power without destroying the laser components is to reduce the pulse length further: a point where the smaller Ti:sapphire systems have a distinct advantage thanks to their broader bandwidth. Fortunately, an extension of CPA — optical parametric chirped pulse amplification, or OPCPA — has been proposed by Ross *et al.*, (1997) designed to achieve exactly that. This technique takes a conventional laser pump beam at any wavelength (Nd:glass, Ti:sapphire or KrF) and amplifies a large bandwidth of the chirped pulse, resulting in 10–20× shorter pulses after recompression than with the usual CPA scheme. Thus it is possible that the PW-class lasers may be upgradable to the 10 PW level via OPCPA without too much additional cost or reconstruction.