Reproducibility of bunches of electrons accelerated in laser driven wakefields

Master's Thesis

by
Philip Siederer

LUND
UNIVERSITY

Supervisors:
Martin Hansson
Anders Persson

Department of Physics, Division of Atomic Physics
Faculty of Engineering, Lund University
SE-221 00 Lund, Sweden

Abstract

Laser wakefield acceleration (LWFA) utilises a high intensity laser pulse to excite a plasma density wave with associated large directed electric fields. These large electric fields can be used to accelerate electrons, making LWFA a good candidate for a compact electron beam source.

In this dissertation, the impact of small fluctuations in laser pulse energy on the generated electron bunch charge is investigated. The electrons are injected into the wave using the ionisation-induced injection scheme. To investigate the laser pulse energy fluctuations, a non-destructive measurement scheme is implemented into the beam line of the multi-terawatt laser at the Lund Laser Centre. To allow easier online measurements of the laser pulses, changes to the software used to control actuators, diagnostics and the laser shutter are made to allow image analysis to be performed on images captured during the experiment and the results of which are returned to users in real time.

It is found that by using a CMOS camera, imaging the leakage through a dielectric mirror, accurate online energy measurements can be performed to investigate the energy pulse fluctuations. Furthermore, it is found that there is a weak correlation between laser pulse fluctuations and electron bunch charge. However, the fluctuations in laser pulse energy can not be the only source of fluctuations in the electron bunch charge.

Acknowledgements

This work could not have been completed without the help from a great number of people. I would especially like to acknowledge and thank:

My two supervisors Martin Hansson and Anders Persson, for all their help and guidance, both of them have always made me feel welcome and have always take time out of their busy schedule to answer my questions. Claes-Göran Wahlström, for inviting me to join the high intensity group and for working around the time zone difference when I was in another country. The rest of the high intensity group, Jonas Björklund Svensson, Malay Dalui, Henrik Ekerfelt, Isabel Gallardo Gonzalez, Olle Lundh, Lovisa Senje, Jack Strand Berg and Kristoffer Svensson for contributing to a good work environment and answering my questions. Samuel Bengtsson, for sharing an office with me and putting up with my music selection.

My parents, for always supporting me.

Contents

1 Introduction

1.1 Background

On June 21, 2016, the new synchrotron facility MAX IV Laboratory was inaugurated in Lund. At the MAX IV Laboratory, electron bunches are accelerated by electric fields through radio-frequencies in a linear accelerator (linac) to be stored in either of two storage rings at the final energies of 1.5 GeV or 3 GeV, respectively. The length of the linear accelerator is 250 m and the need for this huge structure is due to that the maximum electric field that can be sustained in a radio-frequency cavity is approximately 100 MV/m. Electric fields beyond this limit will start to ionise the walls of the cavity and cause the accelerating field to break down. Because of this, if one wants to produce higher energies of electrons one can not simply increase the electric fields but instead one needs to build larger facilities where electrons get accelerated over a greater distance. For instance, the longest linear accelerator to date is located at the SLAC National Accelerator Laboratory in California and has a length of 3.2 km which is used to accelerate electrons up to 50 GeV.

These large projects are expensive (MAX IV costs an estimated 6 billion SEK) and alternative schemes of acceleration are being researched. A future alternative to accelerating radio-frequency cavity, may be to accelerate electrons in a plasma using a high intensity laser pulse in a scheme called laser wakefield acceleration (LWFA). In a laser wakefield acceleration experiment, a short, high power laser pulse is focused on a gas which becomes ionised by the high electric fields of the pulse, creating a plasma. The pulse will propagate through the plasma creating waves of electron density perturbation in its wake. Large electric fields are associated with these waves in which electrons can be accelerated. The electric fields in the plasma can reach strengths tens of thousands of times larger than the maximum fields in radio-frequency cavities and can as such accelerate electrons over much shorter distances to the same energies.

Laser wakefield acceleration was proposed in 1979 [1], but it took until 2004 before quasi-monoenergetic electron bunches were produced experimentally by three separate groups [2]–[4]. Since then, much research has been dedicated to LWFA and in recent experiments [5], it has been experimentally shown that electron bunches with energies on the scale of GeV can be generated.

LWFA shows much promise but there are still some major challenges that need to be resolved. A good electron source should fulfil the following criteria:

- It should be able to generate mono-energetic or close to mono-energetic electron bunches.
- It should be controllable in kinetic energy and charge.
- It should be reproducible.

While LWFA accelerates electrons to high enough energies, the energy and charge of the resulting bunches fluctuate significantly from pulse to pulse. This means one can not say with certainty at what energy the output electrons will be and how many electrons will be accelerated. Without solving these issues

LWFA will not see application as a source for high energy electrons. This is why this thesis focuses on what parameters are important to control to be able to manipulate the output of electrons.

1.2 Goals

In a recent paper from Lund University [6] it was found that by using an injection scheme called ionisation-induced injection (see chapter 2.9) one could achieve a linear regime where the amount of charge generated was linearly dependant on the electron density of the plasma. The aim of this thesis is to investigate if there is a similar regime dependent instead on the small shifts in laser pulse energy that occur naturally in the laser system from shot-to-shot. If this is the case, it may be advisory to implement upgrades to reduce the shot-to-shot differences in pulse energies. To investigate this phenomena it is required to find a nondestructive way to measure the intensity or energy of each laser pulse individually. This project also aims at implementing an easy to use solution to deliver immediate feedback on laser performance to the user to aid in the day-to-day activities at the Lund Laser Centre.

2 Laser wakefield acceleration

In this chapter the physics behind laser wakefield acceleration is introduced. The section starts with a general description of plasma and laser physics. Using these concepts, laser wakefield acceleration is then discussed. The chapter concludes with discussing different means of trapping electrons in plasma waves.

2.1 Laser pulses

For experiments in laser wakefield acceleration one needs to reach enormous intensities to ionise plasmas. Intensities on the scale of exawatt per cm^2 are common. To utilise this kind of intensity on one square centimetre for just 1 minute would already require more energy than the total energy production in Sweden for a year [7]. This is clearly unfeasible. The solution to this energy problem is twofold. Firstly, the laser light gets focused as tight as possible to maximise the intensity. Secondly, and most importantly, the laser is pulsed. The pulse is compressed to only a few femtoseconds, but during this time the intensities become as high as needed.

A femtosecond laser pulse can be described as an electromagnetic wave with a slowly varying carrying function as seen in figure 1. The propagation of light is governed by Maxwell's equation's

$$\begin{cases} \nabla \cdot \mathbf{E} = \dfrac{1}{\epsilon_0}\rho \\ \nabla \cdot \mathbf{B} = 0 \\ \nabla \times \mathbf{E} = -\dfrac{\partial \mathbf{B}}{\partial t} \\ \nabla \times \mathbf{B} = \mu_0(\mathbf{j} + \epsilon_0 \dfrac{\partial \mathbf{E}}{\partial t}) \end{cases} \tag{2.1}$$

where $\mathbf{E}$ is the electric field, $\mathbf{B}$ is the magnetic field, ϵ_0 is the permittivity in vacuum, ρ is the charge density, $\mathbf{j}$ is the current density and μ_0 is the permeability in vacuum.

A common solution for Maxwell's equation is an oscillating plane electromagnetic wave. Mathematically a linearly polarised, travelling wave can be described as

$$\begin{cases} \mathbf{E} = E_0 e^{i(\omega t - \mathbf{k}\cdot\mathbf{r})}\hat{\mathbf{e}}_{\perp,1} \\ \mathbf{B} = B_0 e^{i(\omega t - \mathbf{k}\cdot\mathbf{r})}\hat{\mathbf{e}}_{\perp,2} \end{cases} \tag{2.2}$$

where ω is the angular frequency and $\mathbf{k}$ the wave vector. The unit vectors, $\hat{\mathbf{e}}_{\perp,1}$ and $\hat{\mathbf{e}}_{\perp,2}$, are orthogonal to each other and to the wave vector (proven in section 2.4). Furthermore, the intensity of such an electromagnetic wave is linked to the electric field by

$$I = \frac{cn\epsilon_0}{2}|E_0|^2 \tag{2.3}$$

where n is the refractive index of the material, E is the peak of the electric field and I the intensity. A common convention in LWFA experiments is to use the

amplitude of the normalised vector potential instead of intensity. The vector potential, a, can be related to the intensity through

$$a = \sqrt{\frac{e^2}{2\pi^2 \epsilon_0 m_e^2 c^5} \lambda^2 I} \approx 0.85 \lambda[\mu\text{m}] \sqrt{I[10^{18}\text{W/cm}^2]} \tag{2.4}$$

This convention makes it easy to see when the electrons motions start to become relativistic which happens as a approaches 1.

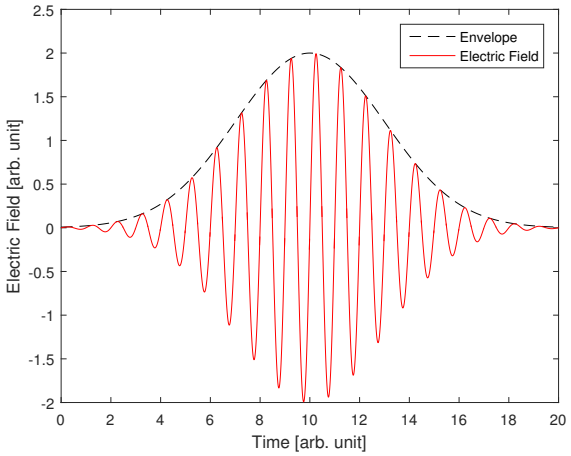

Figure 1: Illustration of an oscillating electric field with a Gaussian pulse shape.

2.2 Plasmas

Often forgotten, plasma is one of the four states of matter (the others being solids, liquids and gases). Although rather uncommon on earth, plasma is the most common state of matter in the universe [8]. A plasma is similar to a gas but is comprised of unbound charged particles whose total charge surmount to zero. This causes a few key differences, one of which is that while most interactions in gases are binary, i.e. two particles colliding, plasmas experience a collective behaviour as the ions and electrons are affected by the Coulomb force of all other particles. Plasmas can be created by high temperatures or, as in the case of laser wakefield experiments, by large electric fields.

2.3 Ionisation of atoms with electric fields

An atom can be simply modelled as electrons trapped in a potential well caused by the positive charge of the nucleus. The potential for the electron will be governed by the Coulomb potential which is

$$V = -\frac{Z_{eff} q^2}{4\pi\epsilon_0 |\mathbf{x}|}. \tag{2.5}$$

Here $\mathbf{x}$ is the distance from the nucleus, Z_{eff} is the net positive charge experienced by the electron and V is the potential. The electron has a rest energy as seen in the left image in figure 2. If the atom is subjected to an electric field, it will change the potential of the system. If a uniform electric field with the magnitude E_0 in the x-direction is applied over the atom the potential will become a barrier following the equation

$$V = -\frac{q^2}{4\pi\epsilon_0|\mathbf{x}|} - E_0 q x. \tag{2.6}$$

If the electric field is strong enough, the potential barrier will be pushed below the binding energy for the electron (right figure 2). When this happens the electron will escape and the atom will be ionised. The position of the peak of the barrier, x_b, can be found by calculating the stationary point of equation 2.6

$$x_b = \sqrt{\frac{Zq}{4\pi\epsilon_0 E_0}}. \tag{2.7}$$

Inserting x_b into equation 2.6, the value of the peak of the barrier can be found,

$$V_b = -\sqrt{\frac{Zq^3 E_0}{4\pi\epsilon_0}}. \tag{2.8}$$

To ionise an atom the barrier needs to fall below the energy of the bound electron, $-E_{ion}$. The minimum electric field needed is thus

$$E_{min} = \frac{E_{ion}^2 \pi\epsilon_0}{Zq^3}. \tag{2.9}$$

If the electric field is due to an oscillating electromagnetic field, the electric field can be linked to the intensity through equation 2.3, giving the minimal intensity required for ionisation

$$I_{ion} = \frac{c\epsilon_0^3 \pi^2 E_{ion}^4}{2Z^2 q^6} \tag{2.10}$$

For instance to ionise a hydrogen atom where the electron is bound by 13.6 eV requires an $1.4 \cdot 10^{14}$ W/cm^2 or a normalised vector potential of 0.008. For comparison the Lund multi-terawatt laser reaches intensities of the order of 10^{18} W/cm^2 when focused on a gas target.

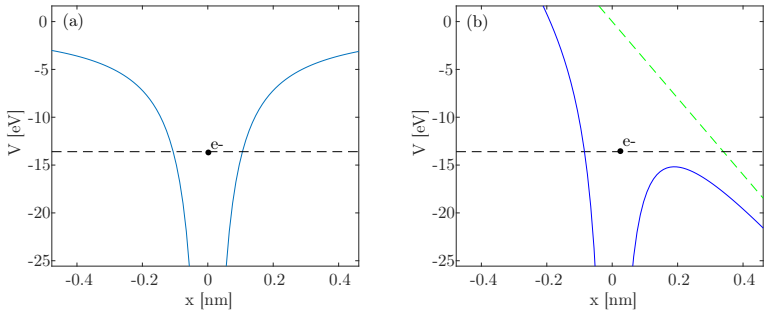

Figure 2: (a) An electron trapped in a potential well caused by a single proton. In figure (b) an external electric field is added (green dotted line) which causes the potential barrier to fall below the electron binding energy thus allowing it to escape.

2.4 Propagation of light in a plasma

To derive the fundamental properties of light travelling through a plasma it is assumed that the light entering the plasma is an infinite plane electromagnetic wave, i.e. can be described as

$$\begin{aligned} \mathbf{E} &= \mathbf{E}_0 e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)} \\ \mathbf{B} &= \mathbf{B}_0 e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)} \end{aligned} \tag{2.11}$$

Inserting equation 2.11 into the last two of Maxwell equations (equation 2.1) leads to

$$\begin{aligned} \mathbf{k}\times\mathbf{E} &= \omega\mathbf{B} \\ i\mathbf{k}\times\mathbf{B} &= \mu_0\mathbf{j} - i\omega\mu_0\epsilon_0\mathbf{E}. \end{aligned} \tag{2.12}$$

At low intensities one can derive an expression for the current density as [9]

$$\mathbf{j} = \frac{in_e q^2}{\omega m_e}\mathbf{E} \tag{2.13}$$

Inserting 2.13 into 2.12 leads to

$$\begin{aligned} \mathbf{k}\times\mathbf{E} &= \omega\mathbf{B} \\ \mathbf{k}\times\mathbf{B} &= (\frac{\mu_0 n_e q^2}{\omega m_e} - \omega\mu_0\epsilon_0)\mathbf{E} \end{aligned} \tag{2.14}$$

The first equation in 2.14 shows that $\mathbf{B}$ is orthogonal to $\mathbf{E}$ and $\mathbf{k}$. The second equation in 2.14 shows that $\mathbf{E}$ is orthogonal to $\mathbf{B}$ and $\mathbf{k}$. Thus $\mathbf{k}$, $\mathbf{E}$ and $\mathbf{B}$ are

all orthogonal as seen in figure 3. Using $\mathbf{B} = \frac{1}{\omega}(\mathbf{k} \times \mathbf{E})$ from the upper equation 2.14, the lower left part of equation 2.14 becomes

$$\mathbf{k} \times \mathbf{B} = \frac{1}{\omega}(\mathbf{k} \times (\mathbf{k} \times \mathbf{E}) = \frac{1}{\omega}\Big(\mathbf{k} \cdot (\mathbf{k} \cdot \mathbf{E}) - \mathbf{E} \cdot (\mathbf{k} \cdot \mathbf{k})\Big) = -\frac{1}{\omega}\mathbf{k}^2\mathbf{E} \tag{2.15}$$

Combining with the right side of the equation and using that $c = 1/\sqrt{\epsilon_0 \mu_0}$ one arrives at

$$\omega^2 = k^2c^2 + \frac{n_e q^2}{m_e \epsilon_0} \tag{2.16}$$

The square root of the last term is defined as the *plasma frequency*, ω_p

$$\omega_p \equiv \sqrt{\frac{n_e q^2}{m_e \epsilon_0}} \tag{2.17}$$

which rewrites 2.16 to

$$\omega^2 = k^2c^2 + \omega_p^2. \tag{2.18}$$

The plasma frequency is a fundamental property of a plasma. It is dependant only on the electron density of the plasma. If the plasma frequency, ω_p, is higher than the frequency of the light, ω, then the wave vector, k, must be imaginary. An imaginary wave vector means that no light can propagate in the medium. This means that for a given plasma one needs a frequency of light higher than the plasma frequency for the light to propagate in the plasma! In LWFA experiments the density of plasma is often varied while the frequency of the laser light is fixed. Because of this it is useful to define the *critical density*

$$n_c = \frac{m_e \epsilon_0}{q^2}\omega_{pe}^2. \tag{2.19}$$

If the electron density, n_e, is above the critical density then electromagnetic waves can not propagate and the plasma is called overdense. If n_e is smaller than n_c the waves can propagate in the plasma and the plasma is called underdense.

Using equation 2.16 the phase velocity and group velocity of a laser pulse can now be calculated

$$\begin{aligned} v_p &= \frac{\omega}{k} = \sqrt{c^2 + \frac{\omega_{pe}^2}{k^2}} \\ v_g &= \frac{d\omega}{dk} = \frac{c}{\sqrt{1 + \frac{\omega_{pe}^2}{k^2c^2}}}. \end{aligned} \tag{2.20}$$

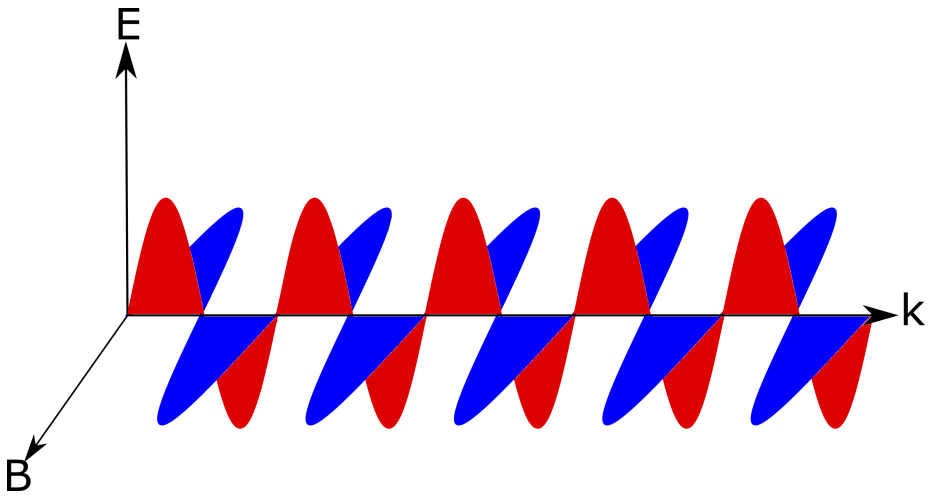

Figure 3: Illustration of an a single frequency oscillating electromagnetic wave travelling along the k-axis. Red corresponds to the electric field and blue to the magnetic field. Note that the wave vector **k**, electric field **E** and the magnetic field **B** are all orthogonal to each other.

2.5 Ponderomotive force

The ponderomotive force is a key concept in laser wakefield acceleration experiments. Due to this effect electrons are pushed away from the high intensity regions of the laser pulse, which in turn leads to the plasma waves that can be used to accelerate electrons. The following is a derivation of the ponderomotive force.

When electrons are subjected to an electromagnetic field they are affected by the Lorentz force

$$\mathbf{F} = m\frac{\partial \mathbf{v}}{\partial t} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}). \tag{2.21}$$

Assume that an electron is subjected to a monochromatic, linearly polarised, oscillating electric field. Furthermore, the electric field varies in magnitude radially symmetric from the optical axis. The electric field can be written as

$$\mathbf{E} = \mathbf{E}_s(\mathbf{r})\cos(\omega t) \tag{2.22}$$

with $\mathbf{E}_s$ being the magnitude dependant on the distance $\mathbf{r}$ from the optical axis. To calculate the Lorentz force only the first order terms are considered to begin with. The subscript $_1$ is used for the first order terms and they are evaluated at the initial position r_0. The $\mathbf{v} \times \mathbf{B}$ term is not of the first order and is thus assumed to be 0 in the first order case. This leads equation 2.21 to become

$$m\frac{\partial \mathbf{v}_1}{\partial t} = -q\mathbf{E}_s(\mathbf{r}_0)\cos(\omega t). \tag{2.23}$$

Integrating twice leads to

$$\begin{aligned} \mathbf{v}_1 &= -\frac{q}{m\omega}\mathbf{E}_s \sin(\omega t) \\ \mathbf{r}_1 &= \frac{q}{m\omega^2}\mathbf{E}_s \cos(\omega t) \end{aligned} \tag{2.24}$$

To calculate the second order a Taylor expansion is used to evaluate the electric field at $\mathbf{r}_0$

$$\mathbf{E}(\mathbf{r}) = \mathbf{E}(\mathbf{r}_0) + (\mathbf{r}_1 \cdot \nabla)\mathbf{E})|_{\mathbf{r}=\mathbf{r}_0} .... \tag{2.25}$$

Further it is assumed that $\mathbf{v}_1$ is a good approximation of the speed. Using Faraday's law of induction,

$$\nabla \times \mathbf{E} = -\frac{\delta \mathbf{B}}{\delta t} \tag{2.26}$$

leads to the magnetic field:

$$\mathbf{B}_1 = \frac{1}{\omega}\nabla \times \mathbf{E}_s|_{r=r_0} \sin(\omega t). \tag{2.27}$$

Inserting $\mathbf{B}_1$ and $\mathbf{v}_1$ into the Lorentz force equation 2.21 and averaging the force over one oscillation period the average force is calculated to

$$\mathbf{F}_{pond} = -\frac{q^2}{2\omega^2 m}\nabla\langle\mathbf{E}\rangle^2. \tag{2.28}$$

where $\langle\cdot\rangle$ signifies the average over one oscillation. This is called the ponderomotive force and has a few very interesting properties. The direction of the force is the same for both negatively and positively charged particles. The force is also inversely proportional to the mass of the particle which means it will affect electrons much more than protons due to that they are around 2000 times lighter than protons. The ponderomotive force pushes particles along the negative gradient of the square of the electric field i.e. particles get pushed away from high intensity regions. The ponderomotive force can, in a simplified model, be explained as follows.

When an electron is subjected to an oscillating electric field it starts to wiggle with the field. If the electric field is homogeneous it oscillates around a single point but if the field has a gradient the electron starts to drift along the negative gradient of the electric field while oscillating. As an example take a femtosecond pulse with a Gaussian profile. The ponderomotive force will then make the electrons wiggle away from the centre of the pulse but also give the electrons a longitudinal velocity due to the envelope of the pulse.

In LWFA experiments, the electric field is powerful enough to cause the electrons to oscillate at relativistic velocities. Due to this one needs to use the relativistic formulation of the ponderomotive force when describing these process accurately. Although the relativistic form differs from the non-relativistic form [10] the general effects are still the same.

2.6 Plasma waves

When a femtosecond laser pulse enters an underdense plasma the ponderomotive force described above will push away charged particles from high intensity regions. Due to the much larger mass of protons they can be assumed to be stationary for the duration of the interaction. The separation of electrons and protons will cause an electric force between the particles. When the laser pulse has passed the electric force will cause the electrons to rush back towards the protons. The electrons will overshoot and start oscillating around the stationary protons. The collective oscillation of electrons can be seen as a plasma density wave. This plasma wave has zero group velocity but the wave has a phase velocity equal to the group velocity of the laser light. An easy way to imagine this effect is to imagine each electron-proton pair as a spring. When the laser pulse passes it loads all the springs and lets them go when it leaves. Each spring will move independently with no interference from the other springs meaning that the group velocity is zero. However, the spring will form a wave due to them being released at different times. The wavefront will move with the same speed as the laser pulse passed them. This analogy gives an intuitive picture of how the plasma wave is formed. It is worth noting that this is a very simplified picture. In practice, plasma waves both have longitudinal and transverse components and the oscillating electrons interact with each other.

2.7 Linear plasma waves

It can be shown that the response from a plasma to a laser pulse with low amplitude ($a \ll 1$) is a harmonic oscillation [11]. After the laser pulse has passed, the density perturbation can be written as

$$\delta n = \Delta n \cos(k_p z - \omega_p t). \tag{2.29}$$

Using Poisson's equation

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} = -q\frac{\delta n}{\epsilon_0} \tag{2.30}$$

one can integrate this equation and arrive at

$$\mathbf{E} = -\frac{q\Delta n}{\epsilon_0 k_p} \sin(k_p z - \omega_p t)\mathbf{e}_z. \tag{2.31}$$

Here it can be seen that the electric field is phase shifted with respect to the density perturbation by $-\frac{\pi}{2}$. Also one can see that the higher the electron perturbation is, the higher the associated electric field is.

2.8 Nonlinear plasma waves and self-trapping

For femtosecond laser pulses of higher intensity the plasma response is not linear. It can be shown that the one dimensional response is [12]

$$\frac{\partial \phi^2}{\partial^2 \xi} = \frac{k_p^2}{2}\left(\frac{\delta n}{n_0} - 1\right) \tag{2.32}$$

where ϕ is the electrostatic potential normalised to $m_e c^2/e$, k_p is the plasma wave vector and $\xi = z - v_g t$. After solving equation 2.32 for the electrostatic potential, the associated plasma density wave can be calculated through

$$\frac{\delta n}{n_0} = \frac{k_p^2}{2}\left(\frac{1+a^2}{(1+\phi)^2} - 1\right). \tag{2.33}$$

The numerical solutions for Gaussian pulses can be seen in figure 4. As the intensity is increased, the electric field becomes sawtooth shaped while the electrons get bunched up close to the sharp edge of the electric field.

The plasma waves can not grow infinitely large with increasing laser intensity. At some point the plasma waves will lose their coherence and start to deform. The wave structure is however not immediately destroyed and the majority of electrons may still contribute to the plasma density wave. The electrons breaking from the wave may fall into the potential well behind the wave and become trapped. If the wave can still maintain its structure, the potential well can travel with the wave and accelerate the electrons trapped inside it. This process is called self-injection or self-trapping. The main advantage with self-trapping is that it can be accomplished using a relatively simple setup and due to this, this was the first scheme used to create quasi-monoenergetic electron beams from LWFA [2]–[4].

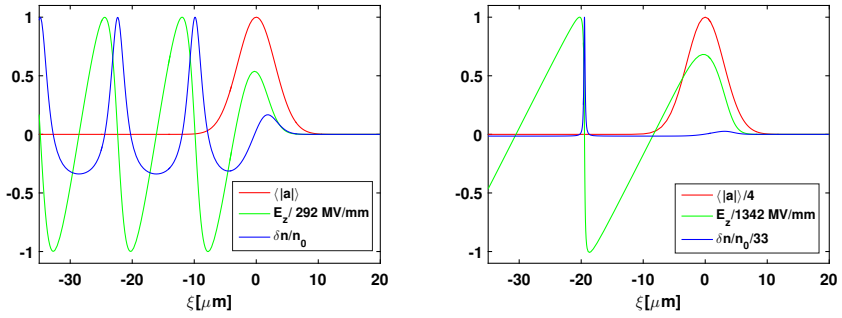

(a) (b)

Figure 4: Numerical solutions in one dimension of the electron density oscillation and the corresponding electric fields excited by a laser pulse with a Gaussian profile. The initial electron density, n_0, is $8 \cdot 10^{18}$ cm$^-3$. In (a) the non-linear effects are relatively small but by increasing the amplitude of the pulse (b) the effects become more pronounced causing the electric field to become sawtooth shaped.

To investigate how electrons behave when being accelerated, it is common to transform the movement of the electron into the reference frame of the plasma wave. In figure 5, one can see orbits of electrons in longitudinal phase space, where $\xi = z - v_g t$. The orbits track the kinetic energy and position of the electrons after it gets injected into the plasma wave. If one were to inject electrons close to the longitudinal dotted line of the figure 5, even electrons with low momentum will become trapped and accelerated with the wave (these orbits are the closed blue orbits in figure 5). In self-trapping schemes, most electrons will be injected injected at the points where the electric field changes rapidly. This is due to that the electron density is the highest in these regions (see figure 4). In figure 5, these electrons will have too low energy to become trapped (lower green orbits) and the plasma wave will simply overtake them. However, if the intensity of the laser pulse is increased, eventually even these electrons will become trapped.

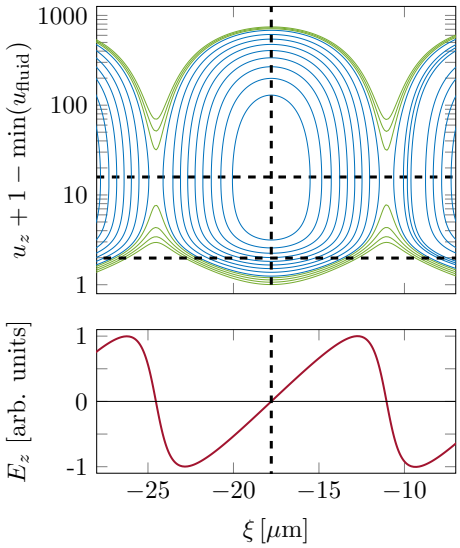

Figure 5: Longitudinal phase space orbits for electrons in a nonlinear plasma wave (above) with the associate electric field of the plasma wave (below). The lower green curves are orbits where electrons have low kinetic energy before the pulse arrives, these electrons do not become trapped. The blue curve shows orbits of electrons trapped within the potential wave. The green curves at the top are theoretical orbits for electrons with too high energy to become trapped. The lowest dashed curve shows where the zero longitudinal momentum is located. The upper horizontal dashed line shows the longitudinal momentum corresponding to an electron moving at the speed of the phase velocity of the plasma wave. (Figure courtesy of Martin Hansson)

2.9 Ionisation-induced trapping

The process of self-trapping relies on the breaking of plasma waves to get electrons injected. This means that the process of acceleration and injection are intrinsically linked together and if one wanted to change the injection, one would need to change the shape of the plasma wave. Other injection schemes aim at separating the process of injection from the process of acceleration, enabling parameters such as the number of accelerated electrons, the time and location of injection of electrons to be manipulated without changing the plasma wave.

This could lead to smaller shot-to-shot variations in number of accelerated electrons and give better control of the process in its entirety. One such controlled injection scheme is ionisation-induced trapping. To understand the idea behind ionisation-induced trapping one can look at figure 5. Looking at the middle of a plasma wave, even electrons at low kinetic energies will become trapped. Sadly, the electron density is very low at these points and as such very few electrons get injected. But if one could control the injection so that electrons get injected in the middle of the plasma wave one could accelerate electrons at a lot lower laser intensities. This is the idea of ionisation-induced trapping.

In an ionisation-induced trapping scheme, a gas mixture is used instead of just the one type of gas commonly used in self-trapping. When using a single low-Z gas, often helium or hydrogen, the gas will be fully ionised long before the peak of the laser pulse arrives. By adding another gas where the ionisation energy is matched so that a bound electron only escapes at the peak of the pulse, the electron will become trapped in the already formed plasma wave. Looking at figure 6, one can see that if an electron gets injected the peak of the laser pulse it follows an orbit where it gets accelerated and eventually overtakes the laser pulse. This is in stark contrast to an electron that gets released into the plasma at the beginning of the laser pulse which does not become trapped in the plasma wave.

Not only does this scheme have the benefit of achieving trapping at lower laser energies but also it is more robust as this scheme only relies on the laser pulse to inject electrons and not the electric fields associated with the plasma waves. In this work a mixture of hydrogen and nitrogen is used. Hydrogen is fully ionised already at an intensity corresponding to a normalised vector potential, a_0, of 0.008. The final two ionisation stages for nitrogen occurs at $a_0 = 2.2$ and $a_0 = 2.8$, an intensity only reached at the peak of the laser pulse. The final two electrons of the nitrogen atom will thus be released into the plasma close to the peak of the wave.

Worth noting is that during the experiments, the maximum reached intensity in vacuum corresponds to $a_0 \approx 1.3$, which is too low to ionise the final stages of nitrogen. However due to self-compression and self-focusing of the laser pulse in the plasma, sufficiently high intensity is reached for ionisation.

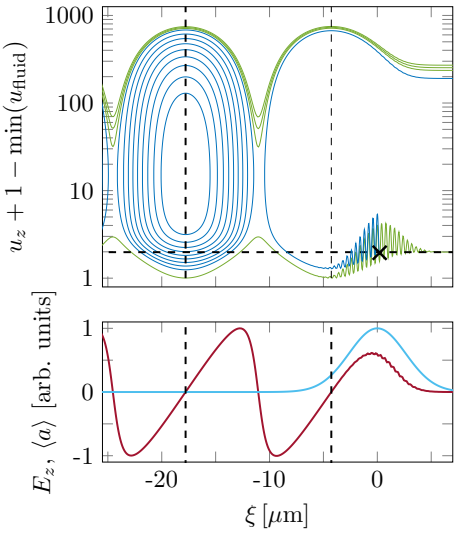

Figure 6: Longitudinal phase space orbits for electrons (above) with corresponding electric field (red, below), and the electric field due to the laser pulse (light blue, below). The green orbit corresponds to an electron being at rest before the laser pulse arrives and does not get trapped. An electron released into the plasma at the peak of the laser pulse (blue) gets trapped and eventually overtakes the laser pulse. Note that the trapped blue orbits in the second period after the laser pulse are purely theoretical as no orbits will be injected there in an ionisation-induced trapping scheme. (Figure courtesy of Martin Hansson)

3 Experiment control software

In this chapter the experiment control software used with the Lund multi-terawatt laser is described. An introduction to the system is given with an overview of the different parts and functions. The changes and additions made to the system to achieve the goals of this thesis are later on discussed.

3.1 The ShotManager system

In 2013, a system called "ShotManager" was designed and implemented to help control and get a better overview of the experiments in the high intensity laser laboratory at Lund Laser Centre. The system was written in C# and controls diagnostics and actuators of the experiments as well as the laser shutter. The program allows users to have an overview of how diagnostics are performing and allows users to set actuators before letting each laser pulse in to the experiment chamber. A schematic overview of the program is illustrated in figure 7. The system is comprised of four main parts: diagnostics units, actuator units, the controller and the data server. The system uses TCP/IP to communicate between the different units which allows actuators and diagnostics to be run on machines separate from the controller. Following is a brief summary of three of the components:

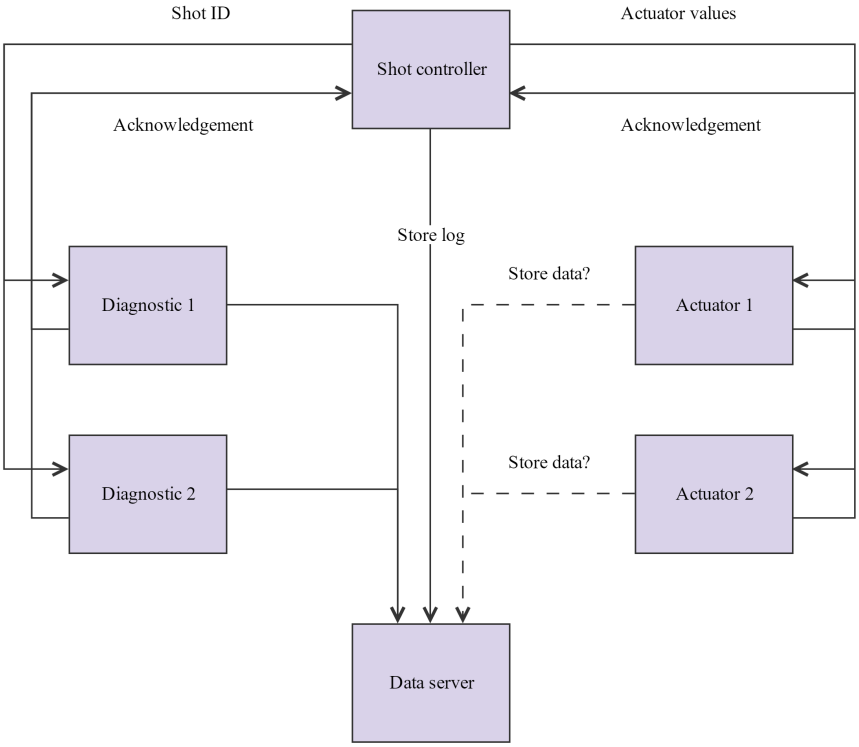

Figure 7: Overview and general workflow of the ShotManager system. One can see the communication between the shot controller, actuators, diagnostics and the data server. The dotted lines indicate optional data transmission as actuator do not necessarily need to store their measurements to the data server.

Controller

The controller is the central part of the system. It holds the graphical interface that gives the users the ability to overview and control the experiment. It also communicates with all diagnostic and actuators units connected to the experiment as well as controls the laser shutter which allows pulses into the experimental chamber. The controller has the ability to define sequences of shots where actuators can automatically be set to different values between shots. The main functions of the controller are:

- Set an error margin for actuators wherein their value must be before a shot is fired.
- Automatically name the shot with an identifier, shotID, and distribute it to all diagnostic units connected. The shotID also contains a path to the data server where the units should save their data.
- Set new values for actuators in the experiment.

- Await acknowledgement that diagnostic and actuator units are ready for the next shot.
- Send a shot to the experiment.
- Write data to a log file where data, such as shot time and actuator values, can be stored.

Diagnostics

The diagnostic units are units which acquire data from the experiment for a specific shot. In contrast to actuators they can not control any part of the experiment. Diagnostic units do not communicate their measurements back to the controller but only communicate to the controller if they are ready for the next shot. Typically, diagnostic units are cameras that save an image file to a predefined destination. The main functions of a diagnostic unit is:

- Receive a shot ID from the controller and acknowledge that the ID has been received.
- Take measurements of the shot and save to the correct location.
- Report when the diagnostic unit is ready for next shot.

Actuators

Actuators are units that control parts of the experiment. A user can set values from the controller for the actuator to follow e.g. pressure in the vacuum chamber or position of a target. Actuators communicate their value continuously with the controller. The main functions of the actuator are:

- Set their value to a value distributed from the controller.
- Continuously communicate current value back to the controller.

3.2 Improved system design

As previously discussed, one can connect two types of units to the controller. One group is actuators and the other is diagnostics. Actuators are devices that steer some kind of process in the experiment while diagnostics are units that record information about the shot and saves it. Diagnostics only report a Boolean value which indicate whether they succeeded in taking their measurement or not after a shot while actuators continuously report their value back to the controller. In this project a third type of device is wanted to measure the intensity of a laser pulse and send the value back to the controller. This device should report a measurement back to the controller after a shot is fired, as is the case with the diagnostic units, but contrary to the diagnostic unit it should not just be a Boolean but be able to be any type of value (integer, string, float etc.).

A possible solution would be to create a third type of unit but this seemed unnecessarily complicated since the new device shares so much with the diagnostic units. A better solution was to modify the diagnostic to report a value instead of just indicating if it is done and change the rest of the system structure to handle this change.

3.2.1 The new diagnostics unit

An important issue to keep in mind when upgrading the ShotManager system is the compatibility of the system. When the diagnostic unit is updated the old diagnostics still need to work with the new system. Therefore a good place to start is with updating the application CameraDiagnostic which is the program that runs on the cameras to let them communicate with the controller. The camera can handle two types of requests from the controller:

- Ping: Returns a value to confirm the unit is connected.
- NextShotID: Receive a path on where the next shot should be saved and makes the diagnostic wait for incoming shot.

The new units needed to be able to return a measured value to the controller which is why a third call was added named "MeasureValue". The call "MeasureValue" returns the value of the measurement and the unit of the measurement separated by a delimiter (in this case '\t'). For the old cameras one can set it to always report "N/A" as it does not return a measurement value. A common response from a energy measurement would be "0.6 \t Joules".

While CameraDiagnostic is the program that runs on the computer that diagnostic unit is connected to, the controller uses the program DiagnosticConnection to communicate with the diagnostic units. In figure 8a an overview of the process when a shotID is distributed is illustrated. In this scheme the controller, after sending a shotID, starts to check the directory where the diagnostic shall save its file. When a new file is registered in the directory the controller concludes that an image has been created by the diagnostic and registers a successful measurement. This scheme relies on the fact that all diagnostic units create a file. However, this was no longer the case with the new diagnostic unit. In figure 8b the new scheme is illustrated. This scheme no longer relies on the creation of files but instead relies on the controller asking the unit if it has completed a measurement. When a unit receives a shotID it sets a new Boolean variable named "shotReceived" to false. When the diagnostic has completed its measurement it sets the value to true. The controller can ask for the "shotReceived" variable by sending a new call named "isShotReceived". If the value returned is false, indicating the diagnostic value has not completed the measurement, the controller waits a set amount of time and asks again. If the value returned is true, indicating that the diagnostic has completed the measurement, the controller registers a successful measurement. With this new scheme, the software becomes more modular where users can define their own diagnostic units and set their own conditions for when the diagnostic has completed its action.

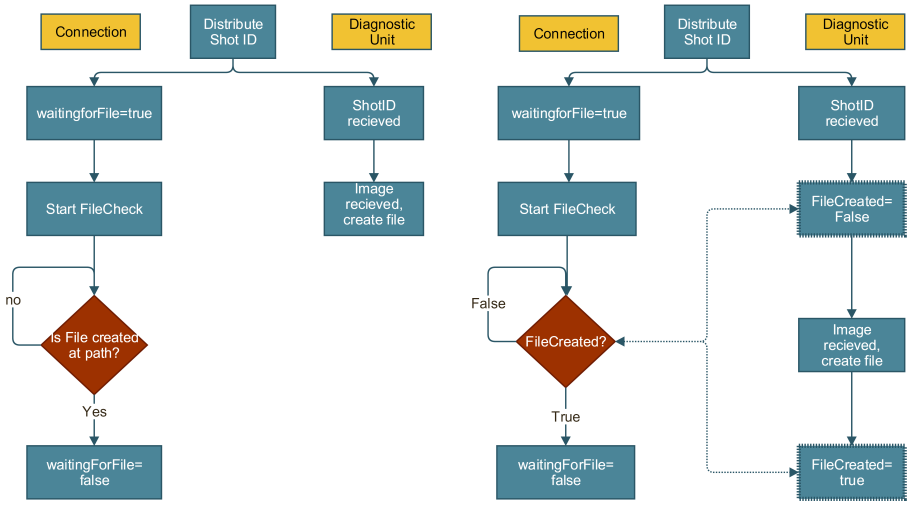

(a) The old version of CameraDiagnostic (b) The new version of CameraDiagnostic

Figure 8: Schematic over the workflow of CameraDiagnostic and the designed new workflow. While the old version relies on the connection associated with the shot controller to check if a new file has been created, the new workflow allows the diagnostic unit itself to communicate if it has measured a value back to the controller.

3.2.2 The ImageDiagnostic

A diagnostic unit, as discussed earlier, is defined as a unit which does some sort of observation after a shot is fired, the most obvious example being a camera. Another example would be a unit which calculates relevant data from an image such as the full-width-half-max or the intensity. This calls for a type of diagnostic that can do image processing. This diagnostic does not necessarily have to be the one that acquired the image but can be a completely separate unit. To fill this need the diagnostic program ImageDiagnostic was created. To use the ImageDiagnostic the user selects a directory to monitor for new image files (taken by a separate unit). When a shotID is received the diagnostic starts looking for a new image in the specified folder. When a file is created the program calls a MATLAB-script selected by the user and the output from the script is then saved as the value of ImageDiagnostic and can be read by the controller. This gives the user a highly modular solution where any type of file can have measurements performed on it as long as the user-defined MATLAB-script can handle it. Furthermore different scripts can be run as separate diagnostics units to make different calculations on the same image. A benefit of the ImageDiagnostic is that the user does not need to have any previous knowledge of how the ShotManager system functions or any TCP/IP knowledge to program new image analysis units.

A design issue that was run into when designing ImageDiagnostic and the new CameraDiagnostic was the seemingly arbitrary problem to get the current time from when to start checking for new files in the directory. Initially, a command

native to C# was called to get the current time but this resulted in problems when running on multiple computer systems. The file is often created on a data server that may have a different internal time compared to the machine running the ImageDiagnostic program and new files could often be missed. A workaround was found where when the ImageDiagnostic receives a shotID it creates a small temporary file in the directory where the image will appear. The program looks at the temporary file's creation time and starts to look for files created after the temporary file. Since creation time is set by the host of the directory both files will use the same time frame, thus solving the problem. The schematic over the program can be seen in figure 9.

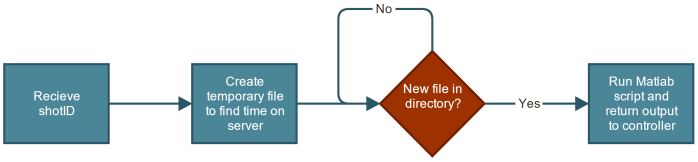

Figure 9: Schematic over the workflow of the ImageDiagnostic program from receiving a shotID to returning a value to the controller

3.2.3 A new workflow for shots

In figure 10 the original workflow of the controller for when a shot is fired is illustrated. After a shot is fired the controller waits for all diagnostic units and actuator units to become ready for the next shot. For diagnostics this entails that the workflow described in section 3.2.1 and seen in figure 8b should be completed. For actuators this means that they should fall within the error margin set by the user (table value in figure 10). If all units are not ready within a user defined time frame then the shot sequence is aborted.

The changes to the diagnostic units must be implemented into the main work flow of the controller. The values registered from the diagnostics must be shown to the user in the controller graphical interface and also be saved to a log file. For actuators, writing their data to the log file is performed before the shot is taken. This is done before the shot since the process of taking the shot will sometimes change the values of the actuators. For example, supplying the target gas changes the pressure in the vacuum chamber so if one measures the pressure after the shot, the value will not be the same as the experimental conditions. The readout of the diagnostic units, however, needed to be implemented after the shot was fired.

A difficulty noted was that the main workflow can not check which diagnostic units that are done as this is hidden from the main program of the controller. Instead it can ask if all diagnostics have completed their measurement. This causes some problems since if a diagnostic fails to complete its measurement the other working diagnostics should still return a value. Printing all values can be very hazardous, since if the diagnostic value has not updated since the last shot one gets the value from the wrong shot. This is something that could potentially give wrong data to the user and would also be very difficult to

detect. The solution implemented was to overwrite all values of diagnostic with a timeout message before a shotID is distributed. When diagnostics receive a shot they overwrite the timeout message with a measured value. When all diagnostics have received their shot the log table is updated with the new values. However, if one of the diagnostics does not become ready within the time limit, causing a timeout, the next shot is aborted and all values from the diagnostics are still reported. The diagnostic units which did not receive a new value keep their timeout message so when read it will say "Timeout" in the log file. By overwriting the measured values between each shot there is no chance for a diagnostic to report the wrong shot to the main system.

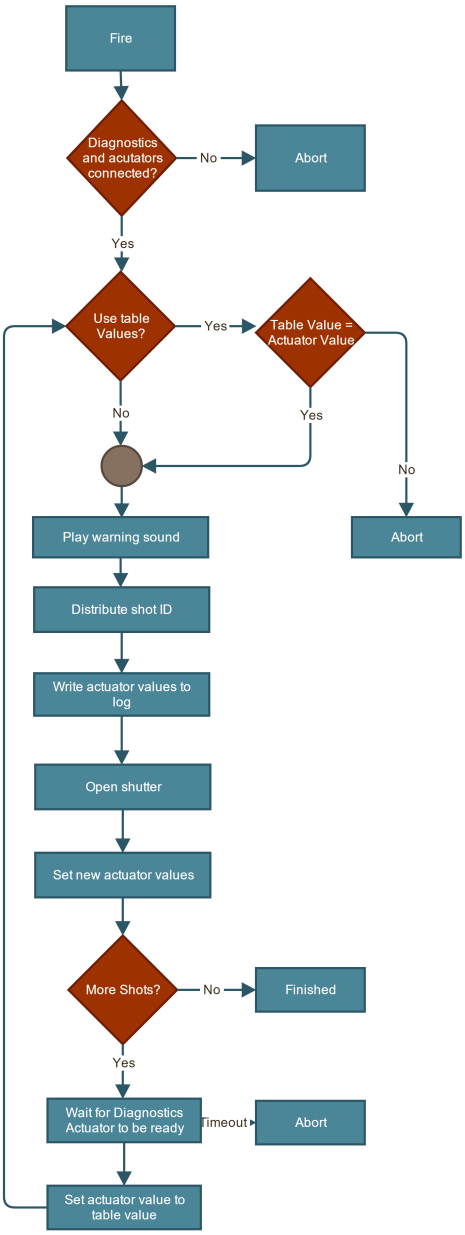

Figure 10: Workflow of the controller program from when a user inputs a fire command until the process is complete. The process can include multiple shots as well as an automatic change to actuator values before each shot.

4 Studies of ionisation-induced trapping

In this chapter experiments performed of laser wakefield acceleration within the scope of the thesis are discussed. The chapter starts with an overview of the setup used to perform LWFA experiments. The fundamentals of digital cameras and how to measure energy accurately is introduced. Finally, results are shown from the LWFA experiments, concluding the chapter.

4.1 Experimental setup

The experiment was conducted at the Lund High-Power Laser facility using a multi-terawatt laser. A schematic overview of the laser system is illustrated in figure 11. The laser system is based on an oscillator using Ti:Al_2O_3 crystal pumped with 532 nm light. Using a Kerr mode locking scheme, pulse trains of 80 MHz is produced with a central wavelength of 800 nm, bandwidth of 50 nm and an energy of approximately 5 nJ. Using a Pockels cell pulse picker 10 pulses per second are directed into a pre-amplifier where the energy of the pulses are increased from nJ to μJ. The pulses then enter a grating stretcher where they are stretched to 9 ps per nm bandwidth which corresponds to a pulse duration of around 450 ps if the bandwidth is assumed to still be 50 nm. The stretching is done to prevent ionisation of air, protect optics and to avoid second order effects when the power is increased.

The pulses are then further led into a regenerative amplifier utilising a Pockels cell to couple the pulses into the optical cavity. The pulses make 12 round trips in the optical cavity to be amplified to a few mJ. The pulses are then expanded to 8 mm in diameter and led into a 5 pass amplifier to amplify the pulses to approximately 400 mJ. The amplifying medium is a Ti:Al_2O_3 crystal with the dimensions 15 mm in length and 18 mm in diameter and is pumped by two frequency doubled Nd:YAG lasers. After this stage the beam is split into two arms. 250 mJ is led to other experiments while the remaining 150 mJ is led to the final amplification table. The beam is expanded to 16 mm in diameter before entering a 4 pass amplifier in a butterfly setup. The crystal is pumped by five Nd:YAG lasers creating a total pump energy of 7 J per pulse. The pulses are amplified to an energy between 0.6-1.4 J. Finally, the pulses are led to a grating compressor which compresses the pulses to approximately 40 fs. However, the grating compressor also lowers the energy of the pulse. This is due to that the compressor contains two diffraction gratings, each with a diffraction efficiency of approximately 0.9. The beam reflects off the two gratings 4 times which gives the theoretical best value of transmission of 65 %. Including the additional optical components, transmitting the laser pulses to the experiments, results in a total transmission measured to 50 % from after the final amplification stage to the target. After the compressor, the pulses must be kept in vacuum to prevent non-linear interaction with air which would destroy the beam quality.

The laser pulses were led into a vacuum chamber kept at pressure below 10^{-3} mbar where they were focused onto a pulsed gas jet supplied by a conical nozzle with its orifice approximately 1 mm away from the optical axis (see section 4.1). The gas used was a mixture of 99 % H_2 and 1 % N_2 to achieve ionisation-

induced trapping. The electrons accelerated which were propagating along the optical axis passed a dipole magnet where they were dispersed depending on their energy before finally hitting a scintillating screen (Kodak Lanex Regular). The backside of the scintillating screen was imaged onto and recorded by a CCD camera.

Two Firefly CMOS cameras were implemented in the beamline for online measurements of the pulse energy. One camera was setup to image the pulse leakage through a dielectric mirror after the final amplification stage but before the pulse enters the final compression. The other camera was set to image the farfield after final compression close to the interaction point. Both cameras had appropriate neutral density filters added to them to attenuate the pulse thus avoiding saturation and laser-induced damage to the camera. The cameras were then added as diagnostic units to the ShotManager system. A script was also connected through the ImageDiagnostic to give intensity measurements back to the user. A Gentec energy meter was able to block the beam before final compression to measure the energy of the pulses. The Gentec was used when changing the amplification power during the last amplification stage.§

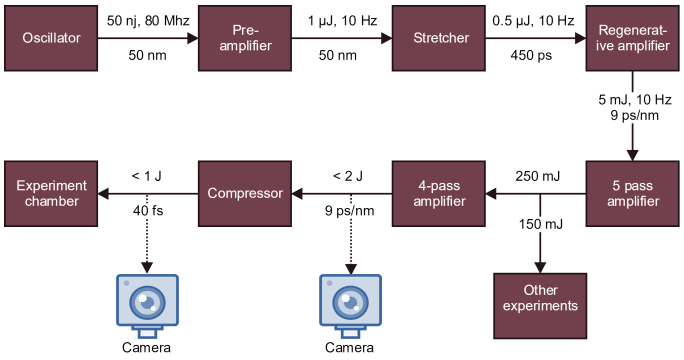

Figure 11: Overview of the multi-terawatt laser at Lund Laser Centre

Laser pulses were registered while scanning the backing pressure of the gas nozzle and intensity of the pulses. The pressure was varied between 1 and 12 bar and the energy of the laser changed by increments of 50 mJ with 4 shots being taken at each setting. Furthermore, 50 shots were taken at 6 bar and 0.7 J to see how the charge of the electron pulses depended on the laser fluctuations.

Calibration of the gas jet

An important parameter to control in LWFA is the density of the plasma as it affects, amongst other things, the phase velocity of the plasma wave and the amount of trapped electrons. The gas targets, which gets ionised to a plasma by the laser pulse, were supplied by a conical nozzle. Earlier experiments have calibrated the gas jet according to figure 12. The linear fit was used to convert the back pressure of the nozzle to the molecular density of the gas. During the

LWFA experiments, the gas is assumed to be fully ionised. By neglecting the multiple ionisation stages of nitrogen, the electron density is assumed to be half of the measured molecular density.

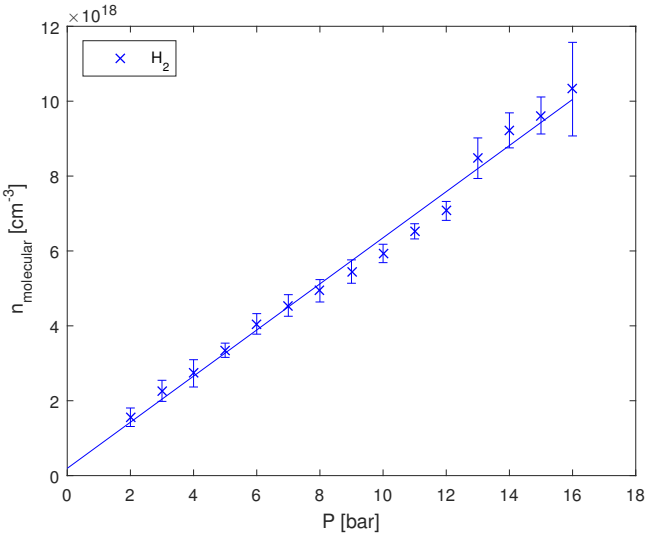

Figure 12: Calibration curve used to correlate the backing pressure of the gas nozzle to molecular density of the jet.

Calibration of the scintillating screen

To detect the accelerated electrons, a scintillating screen (Kodak Lanex Regular) was used. The screen was covered by a thin sheet of aluminium to stop the laser pulse from hitting the scintillator. Before the scintillator a dipole magnet was put that disperses the electrons dependent on their energy before hitting the scintillator screen. The light from the scintillator is captured by a camera. The scintillator sends out 6.95 $\cdot 10^9$ photons/pC/sr [13]. The optical length, L, between the camera and the scintillator screen was 64.5 cm. The lens of the camera had an f-number of 1.4 and a focal length of 50 mm. Using this the aperture of the lens is

$$A_{lens} = \pi(\frac{50 \text{ mm}}{1.4 * 2})^2 \tag{4.1}$$

and the collection angle

$$\theta_{collect} = \frac{A_{lens}}{L} = 0.0024 \text{ sr.} \tag{4.2}$$

Furthermore, the conversion factor between the impacting photons and image pixel counts for the camera was found by shining a HeNe-laser of a known power at the camera. By knowing the photon flux and the exposure time of the camera

the quantum efficiency was found to be 1.2. Converting the pixel value, I, to a charge, Q, hitting the Lanex screen follows the equation:

$$Q = \frac{I}{6.95 \cdot 10^9 \cdot 0.0024 \cdot 1.2 \cdot 7.09 \cdot 10^{OD}} \tag{4.3}$$

where OD is the optical density of an OD-filter if used.

4.2 Cameras as energy measurement devices

To measure the energy of the laser pulses, cameras are used for this experiment. There are multiple benefits with using cameras instead of specialised energy meters. While energy meters are focused on one task, cameras can give the user more information about the pulse such as information regarding the beam size, shape etc. Because of this, many cameras are often already implemented into a beamline and can then also be used as energy measurement devices. The following is a brief explanation of how laser energy is measured with a camera and how to get optimal measurements.

4.2.1 Image sensors

Inside a digital camera there is an image sensor chip which can be regarded as the heart of the camera. The lenses focus incoming light onto the chip which then detects incoming photons and converts it to a digital signal. The digital signal can then be analysed to reconstruct the image projected on the chip. Today generally two different types of image sensors are used; the Complementary Metal-Oxide-Semiconductor (CMOS) and the Charged-Coupled Device (CCD). In this thesis mostly CMOS cameras are used which is why this will be the focus of this section.

A CMOS chip is comprised of many separate pixels, each with a size of a few micrometres in general. Each pixel is comprised of a photo diode and if a photon hits the depletion region of the photo diode an electron-hole pair can form. After a set measurement time, each pixel contains the circuitry logic necessary to convert the stored charge to a voltage which is then subsequently converted to a digital signal and saved [14]. The voltage gives information of how many photons have hit the pixel during the measurement time. The higher voltage, the more photons hit that pixel. This means that one can extract information about the intensity of light shone on the chip.

4.2.2 Measurements and noise

When an image is formed, each pixel will get a numerical value corresponding to the charge accumulated. Using this value, one can then calculate the number of photons that hit the pixel. The sum of the value of all pixels will give a value corresponding the number of photons hitting the chip. To calibrate the camera so that the fluence (energy received per unit area over a set exposure time) values are no longer arbitrary one usually shines a known light source onto the

chip. By knowing the power of the light source and the exposure time of the camera, one can calculate the energy that hit the chip. Thus, one can convert the pixel values into energy.

Ideally there will be a linear dependency between number of incoming photons and the summarised value from the pixels. There are however, limitations of how accurate a measurement can be from a camera. An important factor is the signal-to-noise ratio. Noise is commonly a random signal added to the read out signal. The noise signal can be caused for a multitude of different reasons. For instance, thermal noise is due to electrons being excited over the band gap due to heating, thus giving a false read out. Noise is often randomised and can, because of this, often be reduced by applying some type of filter. One such filter is a median filter which sets the value of each pixel to the median of its neighbouring pixels. This causes the noise to decrease but sacrifices the resolution of the camera to achieve this goal. A median filter is effective when dealing with salt-and-pepper-noise, i.e. noise that causes random pixels to register very high values (white) or very low values (black).

Cameras may utilise gain to increase the brightness of an image. Gain increases the value of all pixels by a set amount and in doing so, darker parts become more visible but the noise will also increase. When calibrating the camera for energy measurement, the gain must be noted and accounted for. If gain is changed after calibration, all reported intensity values will be faulty which is why it's very important to be aware of the gain of the camera.

Cameras may also utilise gamma correction. The goal of gamma correction is to mimic the way which the human eye registers light which it does non-linearly. Gamma correction reallocates bits by decreasing the amount of information stored for darker shades, which the eye is worse at recognising, and allocates more bits to brighter hues [15]. While this is very useful for a human observer to make images look more natural, intensity measurements needs the unmodified values to accurately measure the energy of the laser pulses. Thus gamma correction is not desirable when using cameras as energy measurement devices and should be deactivated.

Another type of interference to the signal is the background signal. Typically a background signal comes from other external light sources than the one that is the target of investigation. To minimise the background, one can simply attempt to block as much of the external signal as possible e.g. measuring in a dark room. The source of the background signal is sometimes internal and can arise when signals are converted from analogue to digital. To remove the background signal, one can take advantage of the fact that the background is often not randomised but a steady signal and attempt to reduce it. One such simple scheme is to take a picture with no light source shining on the chip, a so called background picture, and then subtract that image from all measured values.

A more advanced scheme used, is for software to regenerate a background signal in the region of interest and then subtract it from the image. In the scheme later used during measurements for LWFA experiments, the user chooses 3 points in the image outside the region of interest. It then interpolates a background signal within the region of interest, which is then subtracted from the image. This

scheme has the advantage of being tailored to each image and is thus able to handle changes in the background better than a static image subtraction. However it takes more computational power and also assumes that the background has no higher orders of variation i.e. that the background is rapidly changing throughout the region of interest.

Another important thing to keep in mind when using cameras is to not allow them to overflow. Each pixel has a maximum number of photons it can store before it starts to overflow. If a pixel reaches its maximum value it can no longer register more photons and because of this additional energy will be ignored leading to an incorrect measurement.

4.3 Testing of cameras for energy measurement

4.3.1 Linearity test of cameras

As mentioned in section 4.2.2 one can use cameras to measure the intensity of a laser pulse. For calibration purposes it is often assumed that the relationship between laser intensity and pixel value is linear which may not be the case. For the experiment two different cameras were tested to check their suitability as energy measurement devices. One is a CMOS Firefly FFMV-03M2M-CS [16] and the other a CCD Dragonfly2 [17], both manufactured by Point Grey. A turn-key diode laser from Thorlabs with an optical fibre fitted to its output was used and the beam was focused on the camera being tested. The setup can be seen in figure 13. Images were captured at different power settings of the laser. For each power setting, a power meter was put into the beam line and the optical power was recorded. The pixel values were then correlated to the optical power measured by the power meter. A background image was taken without the laser beam and was subtracted from all images. A first order polynomial was then fitted to the acquired data points using the least-squares method to the measured data set. The results can be seen in figure 14. The largest deviation from the fitted curve to the measured value is just 1.2 % away from the measured value for the firefly and 0.3 % for the Dragonfly concluding that both cameras are linear enough to be used for energy measurements.

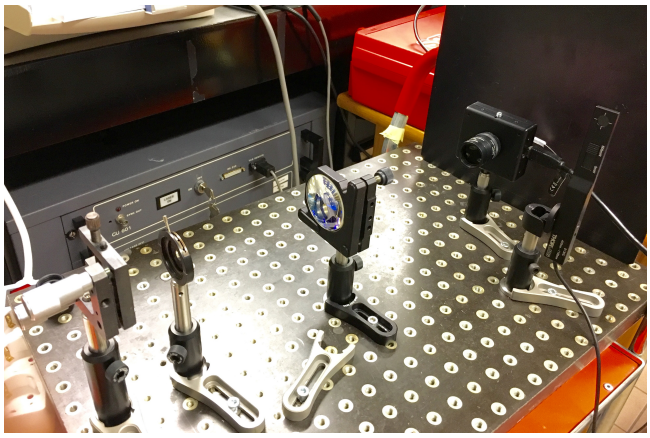

Figure 13: Image over the setup used to test the linearity of the cameras. To the right the Dragonfly camera can be seen as well as a power meter that can be flipped into the beam path of the diode laser.

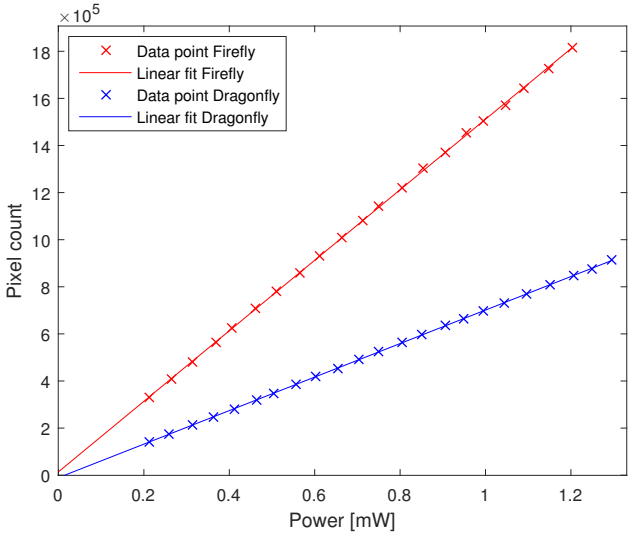

Figure 14: Linearity plot featuring the Firefly and Dragonfly cameras. Both cameras show good linearity and by extrapolating the data the pixel count is very close to 0 at 0 power indicating a good subtraction of background noise.

4.3.2 Measuring laser fluctuations

To further test the suitability of the cameras as intensity measurement devices, one of them was set up after the final amplification stage of the multi-terawatt laser (see section 4.1). The near-field energy distribution of the laser pulses were imaged onto the camera using the leakage through a dielectric mirror. An energy meter, Gentec QE95LP-S-MB-QE, was put into the main beam line, measuring the laser pulse energy. Only the Firefly camera was tested as it was concluded

that this camera fulfils the required specifications for the experiment. The images from the Firefly camera were subjected to a software background subtraction method which entailed regenerating a background signal as discussed in section 4.2.2. A median filter was also implemented, taking the median over sets of 5x5 pixels. The Gentec energy meter has a repeatability of more than 99.5 % [18]. The energy of 26 laser pulses were measured using both methods and the results are displayed in figure 15. The two methods were in excellent agreement and the average deviation in measured value was 3 mJ. It was therefore concluded that the camera was suitable to measure the fluctuations in laser pulse energy from shot to shot in laser wakefield experiments.

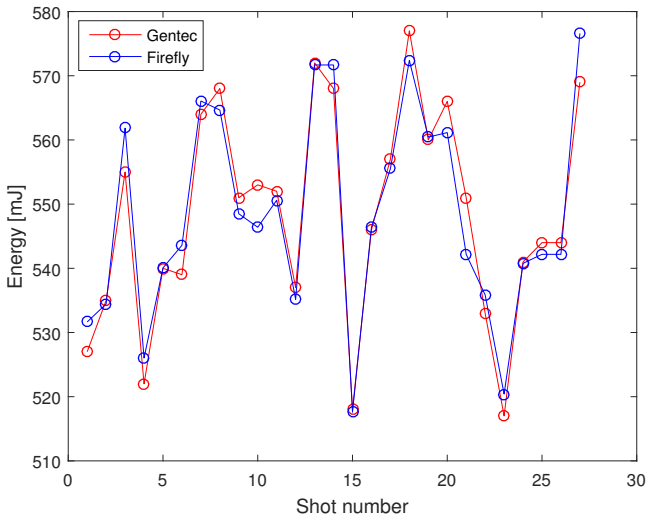

Figure 15: Comparison in measured energy by the Gentec energy meter and the Firefly CMOS camera. A good correlation can be seen between the two measurement methods, indicating that both are able to give accurate measurements of the laser pulses.

4.4 Results

4.4.1 Software

The software performed well and could easily run at 1 shot every 10 seconds which was the requirement as the time for the gas to evacuate the vacuum chamber is approximately 6 seconds. The most time consuming step in the process was the image capture process by the cameras. The start-up time for the scripts was approximately 0.1 seconds with multiple scripts being able to be run at the same time.

4.4.2 Camera

In figure , 2 images are shown captured by the two Firefly cameras. In figure 17, one can see the energy measured by the Gentec energy meter compared to the values from the two cameras. Since the Gentec energy meter blocks the pulses, it could not be used on each pulse for measurement. Instead, the Gentec was used to set the correct amplification for the laser pulses and then removed before the experiment was run. Due to this, the measured value on the x-axis in figure 17 is not specific for each pulse but rather a mean of all the pulses on that amplification setting.

The camera setup before compression shows good linearity, however the camera setup to image the farfield after compression is not functioning as well as anticipated. Looking at the measurement at 1 J and comparing it to the measurement at 1.1 J in figure 17b the values are the same. This makes it doubtful if the camera can be used for measurement as if the camera can not record a difference of 0.1 J with statistical accuracy, it will not be able to record the much smaller differences of the fluctuations of the laser at a fixed energy.

Because of the uncertainty of measurement by the second camera, all future energy measurements described in this thesis are done using the camera before the final compression stage and the linear fitting curve in figure 17a is used for calibrating the pixel values.

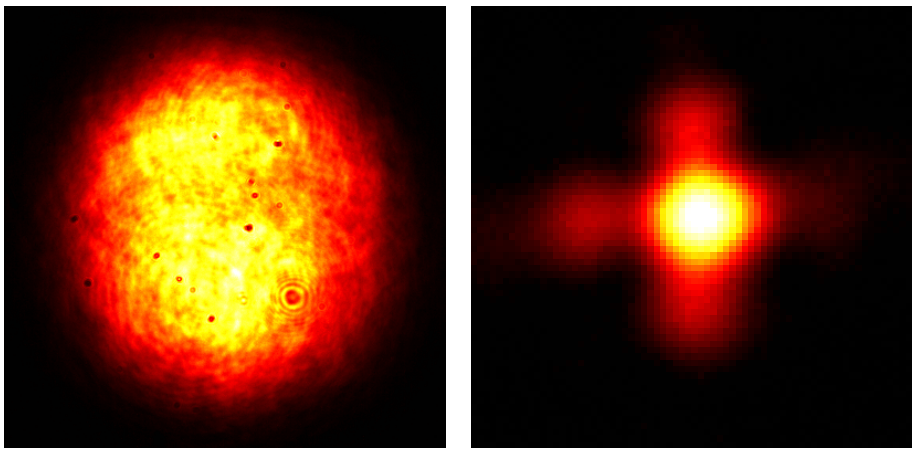

(a) Before compression (b) After compression

Figure 16: Images captured by the two firefly cameras during experiments.

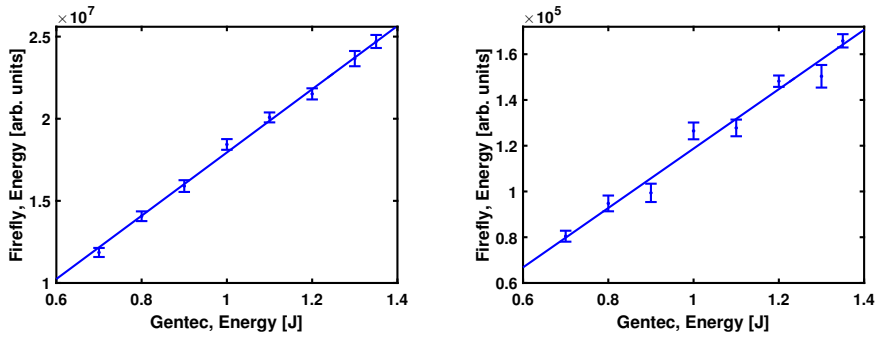

(a) Before compression (b) After compression

Figure 17: Correlation plot between energy measured by the Gentec energy meter (x-axis) and the two Firefly cameras (y-axis) together with a first order polynomial fitted to the data set. The measurement before the final compression (a) are satisfactory but the measurements after compression (b) are not accurate enough to measure pulse energy with the required precision.

4.4.3 Pressure dependence

Figure 18 shows the total charge generated as a function of the density of electrons in the plasma. The figure shows clearly that there is a specific density over which electrons are being trapped and accelerated within the laser wakefield. For a laser energy of 0.6 J this begins at an electron density of $2.8 \cdot 10^{18}$ cm^{-3}. This is further illustrated in figure 19 which shows two images of the lanex screen for densities below and above the threshold. Worth noting is that the density $2.8 \cdot 10^{18}$cm^{-3} seems to be right on the threshold as some shots did not produce any electrons. For a lower laser energy the electron density needs to be increased to generate electron beams. For example, the density required to produce electron beams with a laser pulse energy of 0.35 J is $5.3 \cdot 10^{18}$cm^{-3}.

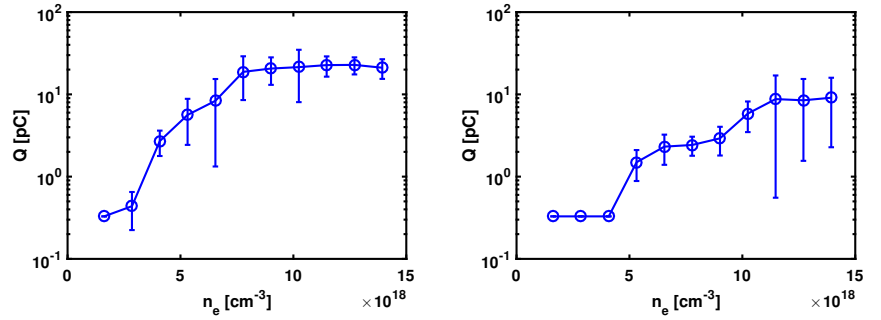

(a) Laser energy at 0.65 J per pulse (b) Laser energy at 0.35 J per pulse

Figure 18: Generated charge as a function of plasma density. Note that the threshold for generated charge is lower for a higher laser energy.

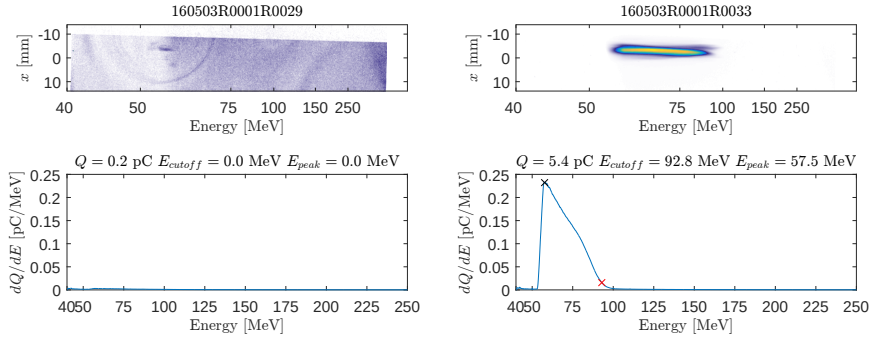

(a) $I = 0.65$ J, $n_e = 1.5 \cdot 10^{18}\text{cm}^{-3}$ (b) $I = 0.65$ J,$n_e = 2.8 \cdot 10^{18}\text{cm}^{-3}$

Figure 19: Image of the scintillating screen (top) and the corresponding charge in the beam (bottom). In figure (a) no electrons are trapped and accelerated. In figure (b) electrons are accelerated due to the plasma density being increased compared to figure (a).

4.4.4 Energy dependence

In figure 20 the total charge of accelerated electrons is shown for a sequence of shots in which the electron density of the plasma is kept constant while the energy of the laser pulses were varied. The aim was to study the dependence of the electron beam charge on laser pulse energy fluctuations. In the figure, the standard deviation is shown for the charge (y-axis) and laser energy (x-axis). The threshold when electrons start being trapped and accelerated is at 0.35 J.

A large fluctuation in laser intensity does not correlate to a large spread of generated charge as both shots at 0.4 J and 0.68 J show smaller deviation of laser energy but no smaller deviation in generated charge. A higher laser energy does not ensure generation of electrons as some shots taken at 0.55 J generate no charge and at this point the laser energy is far beyond the threshold for production of electrons.

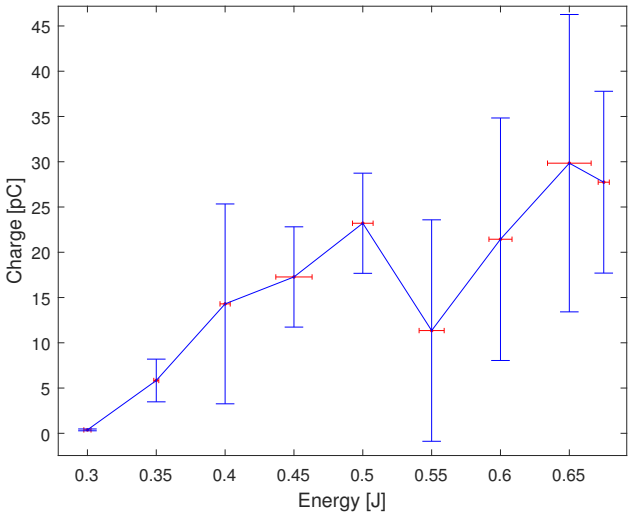

Figure 20: Plot over electron bunch charge as a function of laser pulse energy. Four shots were taken at each setting of the laser energy. The standard deviation (red) was calculated using energy measurements from the images of the laser beam using the Firefly camera before final compression of the laser beam. Electrons start being trapped an accelerated at 0.35 J.

4.4.5 Shot at a fixed setting

To further investigate the correlation between the fluctuations in laser pulse energy and generated electron charge, 50 shots were taken at an electron density of $7.8{\cdot}10^{18}\mathrm{cm}^{-3}$ and a laser energy centered around 0.7 J. The results can be seen in figure 21. The fluctuations in accelerated charge in this sequence are still large, however, a slight energy dependence may be seen when using a linear fit.

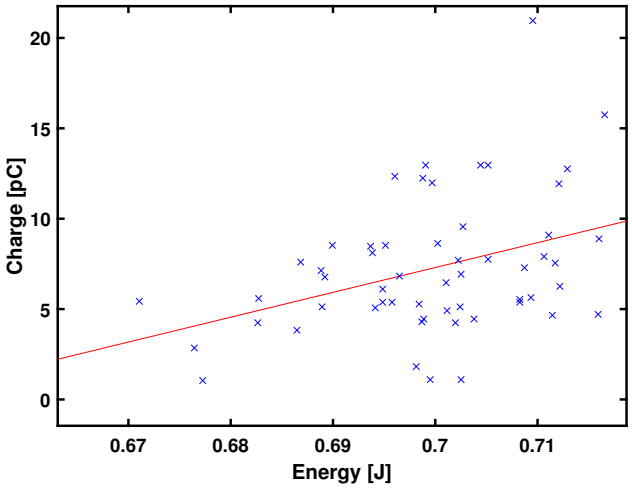

Figure 21: Generated charge as a function of laser pulse energy, measured using the Firefly camera before final compression of the laser pulse. The plasma is kept at a fixed density of $7.8{\cdot}10^{18}\mathrm{cm}^{-3}$. The red curve shows a first degree polynomial fit to the data set.

5 Discussion and Conclusions

5.1 Energy Dependence

The fluctuations in bunch charge were large as seen by the large error bars in figure 20. This makes it difficult to deduct any energy dependence with just four shots at each setting. The original goal was to see a dependence from individual shots at each setting but the fluctuations in generated charge were so large it is hard to deduct any dependence. When using only 4 shots at each setting, small fluctuations in laser pulse energy does not correlate to smaller charge fluctuations.

When both pressure of the gas jet and the laser gain is kept constant for many shots, one is able to see a very small correlation between bunch charge and laser energy as seen in figure 21 after 50 shots. This is however far from conclusive and many more shots need to be taken to be able to conclude the energy dependence. The fluctuations still vary greatly even when the pulse energy is the same. So while the laser energy may have an affect on smaller scales there seems to be other factors that have a greater impact on the amount of accelerated electrons. One such variable may be the fluctations in density of the gas jet. If one looks at the calibration curve of the gas jet in figure 12, the error bars are quite large. This may point at that there are fluctuations in the density of gas from jet pulse to jet pulse. More evidence pointing towards this is the fact that even at higher laser energies, well above the threshold for injection, there are still shots for which no electrons are accelerated. The electron density greatly affects the amount of charge in the electron bunches and may therefor be the reason for the large fluctuations. In the article from Lund University mentioned in section 1.2, the same laser system was used and much smaller fluctuations in generated charge was found. They did not, however, use a gas jet but a gas cell. In a gas cell a known amount of gas is contained inside a vial. The higher stability of a gas cell may answer the question of why the fluctuations were so much larger in this thesis compared to their work performed using the same laser system. When looking at other experiments that use a gas jet as a source for the plasma [19] the generated charge fluctuates on par with the results found in this thesis, further confirming that the stability of the gas jet is the limiting factor in this set up.

5.2 ShotManager System

Adding the new type of diagnostics to the ShotManager system was a success. In this thesis, only cameras were connected as diagnostic units but it is possible to connect other types of devices to the ShotManager including other type of energy measurement devices such as photo diodes and energy meters. The program allows for any type of diagnostic to be connected as long as it follows the communication protocol as described in section 3.2.1.

With the addition of the ImageDiagnostic program users can easily implement image analysis scripts that can perform calculations on acquired images and report the value back to the user while the experiment is running. By using the

ImageDiagnostic, new scripts can be implemented easily into the ShotManager system without the user having any knowledge of C#, TCP-protocol or the back end of the ShotManager system thus enabling users to easily implement tools that give immediate feedback on properties of the laser pulse. As currently implemented, only MATLAB scripts are able to be run on images. This is not optimal as this is limiting users to one specific software and requiring MATLAB to be installed on the machine running the ImageDiagnostic program. In the future compatibility to other programing languages can be added such as Python and C which would eliminate the need to have costly MATLAB licenses on machines as well as making the program more diverse.

Another feature that could be implemented in the future is for users to be able to do two dimensional scans. As it stands, only one parameter can be changed automatically between shots but with some changes to the ShotManager it would be possible to automate multiple actuators thus decreasing the amount of time spent by users manually changing values.

The biggest time sink for the day to day usage of the ShotManager program is the start-up time. Before using the system one needs to start up the main program on one computer and then startup many different applications for actuators and diagnostics on different computers. Efforts have been made to reduce the time needed by enabling the program to save its connections and ports upon closure. However some sort of automated startup sequence would be very useful and valuable to users. There are a few hurdles preventing this to be implemented. The main issue is that due to multiple computers being used it is very difficult to create a one-click startup process. However, some sort of standard configuration might be able to be added to the ShotManager system so users do not have to add all actuator units and diagnostic units manually to the controller when using a freshly installed version.

5.3 Cameras as a diagnostic tool

Utilising cameras as energy meters was met with mixed results. Under optimal conditions the chip is clearly able to measure energies precisely as seen during the testing with the HeNe laser. Furthermore, when using the camera to image the leakage through a dielectric mirror before final compression, where the pulse length is around 9 ps, the measurements are accurate enough to be able to observe the fluctuations in laser intensity. Looking over a large interval as done in figure 17 one can observe that the measured values fluctuate slightly from the linear graph. Because of this, using cameras as the only energy meter in a setup is not advisable. However, using cameras to see small fluctuations during the experiment works well.

To get the most accurate information about the laser pulse entering the experiment one wants to have the camera as close as possible to the experiment as optics following the measurement may alter the pulse. The optimal solution would be to have the camera imaging the laser pulse from within the vacuum chamber. However, it has been noted previously that keeping the Firefly cameras in vacuum for an extended period of time has caused the noise of the chip to increase.

In this experiment, the camera closest to the experiment used the leakage through the last dielectric mirror before the focusing mirror. however the results from this camera was unsatisfactory. Since a difference of 0.1 J did not make a noticeable difference to the readout, much smaller fluctuations as the ones wanted to be seen can not reliably be registered. One theory to explain that the results are bad could be that it is not the camera measuring incorrectly but the laser that is unstable. This theory can however be disproved, since the other camera measures a correlation between the laser pulse energy and the number of accelerated electrons. This concludes that this is a measurement issue.

One probable theory as to why the camera after compression functioned worse may be the spotsize. The full width half max is about 200 pixels on the camera before compression and 30 pixels on the camera after compression. While 30 pixels should be more than enough to theoretically measure the energy (90 pixels in total each with a 7-bit range), the background subtraction method may perform worse on a smaller spotsize. To further investigate, the spotsize should be increased to the same size as the spotsize before compression. The cameras should also change places to investigate if this is a hardware issue.

It was further found that the background subtraction method had a big impact on the measurement results. During the initial stages, a background subtraction method was used. An image with a blocked beam was taken and then subtracted from all images taken after. It was found that using software to create a background image gave better results. The main drawback of this method was the processing time. The script took approximately 3 seconds to run which in the case of this experiment was not a problem but might be for applications that need faster results. Then again, the goal with the feedback might not be to have as exact measurements as possible as this can be done after but more a general idea of where different parameters are at that moment in time.

5.4 Conclusion

The updates to the ShotManager system provided a good and modular approach for end-users to easily implement diagnostic units with the ability to give instant feedback during day to day activities. Furthermore, using cameras as energy measurement devices worked well when used in conjunction with specialised energy meters to measure small changes in laser intensities on a shot to shot basis.

The findings show that when using a gas jet the laser energy fluctuations may have a small effect on the generated charge from the laser wakefield acceleration but it is not a major contributing factor. The limiting factor seems to be the gas nozzle which produces the gas jets. To further investigate, if there exists a dependence on the laser fluctuation, another type of gas target should be used with higher stability. One such gas target would be a gas cell which would be a good choice as previous experiments have shown that a gas cell has a higher stability than a gas jet.

This thesis can finally conclude that, if upgrades were to be made to the multi-terawatt laser system at Lund University, upgrades to the stability of the laser should not be a primary concern if a gas jet is used as the gas target.

References

[1] T. Tajima and J. M. Dawson, "Laser electron accelerator", *Phys. Rev. Lett.*, vol. 43, pp. 267–270, 4 Jul. 1979. DOI: 10.1103/PhysRevLett.43.267.

[2] 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*, vol. 431, no. 7008, pp. 541–544, Sep. 2004. DOI: 10.1038/nature02963.

[3] C. G. R. Geddes, C. Toth, J. van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans, "High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding", *Nature*, vol. 431, no. 7008, pp. 538–541, Sep. 2004. DOI: 10.1038/nature02900.

[4] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick, "Monoenergetic beams of relativistic electrons from intense laser–plasma interactions", *Nature*, vol. 431, no. 7008, pp. 535–538, Sep. 2004. DOI: 10.1038/nature02939.

[5] C. E. Clayton, J. E. Ralph, F. Albert, R. A. Fonseca, S. H. Glenzer, C. Joshi, W. Lu, K. A. Marsh, S. F. Martins, W. B. Mori, A. Pak, F. S. Tsung, B. B. Pollock, J. S. Ross, L. O. Silva, and D. H. Froula, "Self-guided laser wakefield acceleration beyond 1 gev using ionization-induced injection", *Phys. Rev. Lett.*, vol. 105, p. 105 003, 10 Sep. 2010. DOI: 10.1103/PhysRevLett.105.105003.

[6] M. Hansson, T. L. Audet, H. Ekerfelt, B. Aurand, I. Gallardo González, F. Desforges, X. Davoine, A. Maitrallain, S. Reymond, P. Monot, A. Persson, S. Dobosz Dufrénoy, C.-G. Wahlström, B. Cros, and O. Lundh, "Localization of ionization-induced trapping in a laser wakefield accelerator using a density down-ramp", *Plasma Phys. Control. Fusion*, vol. 58, no. 5, p. 055 009, Mar. 2016. DOI: 10.1088/0741-3335/58/5/055009.

[7] *2015 var ett ar med stor elproduktion och rekordstor export av el*, Feb. 2016. [Online]. Available: http://www.energimyndigheten.se/nyhetsarkiv/2016/2015-var-ett-ar-med-stor-elproduktion-och-rekordstor-export-av-el/.

[8] F. F. Chen, *Introduction to Plasma Physics and Controlled Fusion*, ser. Introduction to Plasma Physics and Controlled Fusion v. 1. Springer, 1984, ISBN: 9780306413322.

[9] W. L. Kruer, *The Physics of Laser Plasma Interactions*, ser. Frontiers in physics. Westview Press, 2003, ISBN: 9780813340838.

[10] D. A. Jaroszynski, R. A. Bingham, and R. A. Cairns, *Laser-Plasma Interactions*, ser. Scottish Graduate Series. CRC Press, 2009, ISBN: 9781584887799.

[11] L. M. Gorbunov and V. I. Kirsanov, "The excitation of plasma waves by an electromagnetic wave packet", *Zhurnal Eksperimentalnoi i Teoreticheskoi Fiziki*, vol. 93, pp. 509–518, Aug. 1987.

[12] P. Sprangle, E. Esarey, and A. Ting, "Nonlinear interaction of intense laser pulses in plasmas", *Phys. Rev. A*, vol. 41, pp. 4463–4469, 8 Apr. 1990. DOI: 10.1103/PhysRevA.41.4463.

[13] A. Buck, K. Zeil, A. Popp, K. Schmid, A. Jochmann, S. D. Kraft, B. Hidding, T. Kudyakov, C. M. S. Sears, L. Veisz, S. Karsch, J. Pawelke, R. Sauerbrey, T. Cowan, F. Krausz, and U. Schramm, "Absolute charge calibration of scintillating screens for relativistic electron detection", *Rev. Sci. Instrum.*, vol. 81, no. 3, p. 033 301, 2010. DOI: 10.1063/1.3310275.

[14] G. C. Holst and T. S. Lomheim, *CMOS/CCD Sensors and Camera Systems*, ser. SPIE PM. JCD Publishing, 2007, ISBN: 9780819467300.

[15] E. Reinhard, W. Heidrich, P. Debevec, S. Pattanaik, G. Ward, and K. Myszkowski, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, ser. Morgan Kaufmann series in computer graphics. Elsevier Science, 2010, ISBN: 9780080957111.

[16] *Firefly data-sheet*, visited 2016-07-18, Point Grey. [Online]. Available: https://www.ptgrey.com/support/downloads/10141.

[17] *Dragonfly2 data-sheet*, visited 2016-07-18, Point Grey. [Online]. Available: https://www.ptgrey.com/support/downloads/10140.

[18] *Qe95 data-sheet*, visited 2016-06-16, Gentec-eo. [Online]. Available: https://gentec-eo.com/Content/downloads/specifications-sheet/QE95_2016_V1.0.pdf.

[19] S. P. D. Mangles, G. Genoud, M. S. Bloom, M. Burza, Z. Najmudin, A. Persson, K. Svensson, A. G. R. Thomas, and C.-G. Wahlström, "Self-injection threshold in self-guided laser wakefield accelerators", *Phys. Rev. ST Accel. Beams*, vol. 15, p. 011 302, 1 Jan. 2012. DOI: 10.1103/PhysRevSTAB.15.011302.