Impact of Plasma Dynamics On Femtosecond Filamentation

Thesis submitted to the Faculty of Graduate and Postdoctoral Studies in partial fulfillment of the requirements for the degree:

Master of Science in Physics

OTTAWA-CARLETON INSTITUTE OF PHYSICS

DEPARTMENT OF PHYSICS

FACULTY OF SCIENCE

Author:

Rhys Emms

Supervisor:

Dr. Thomas Brabec

May 25, 2016

UNIVERSITY OF OTTAWA

Abstract

In this thesis we ran a series of 2D simulations of femtosecond laser pulses filamenting in air using the FDTD method, a saturable Lorentz oscillator model of air [1], and two separate models of plasma: a Drude model where the plasma density is static in space, and a particle-in-cell model where plasma is free to migrate throughout the simulation space. By comparing matched pairs of simulations, which varied in pulse size, duration, and intensity, we can gauge the impact plasma dynamics has upon the evolution of a filamenting laser pulse. From these tests we determine that, while there are some visible differences between dynamic and static plasmas, plasma dynamics do not significantly alter the evolution of the pulse.

Acknowledgments

I would like to thank my supervisor, Dr. Brabec, for his insights and guidance through my studies. I would also like to thank Dr. Charles Varin for his invaluable knowledge into the arcane workings of massively parallel FDTD-PIC codes and his tolerance for my unending questions about programming problems. Furthermore, without Dr. Varin's development of the saturable Lorentz oscillator model for air, this project would have taken many months more of work to complete. I would like to thank Dr. Christopher Brady and Dr. Keith Bennett of the University of Warwick for writing and maintaining the EPOCH program which has tormented me so for these past few years. Lastly, I must thank the maintenance and support staff of Calcul Québec's Briarée computer for their assistance in getting my code up and running on their machine.

Contents

List of Figures

List of Tables

1 Introduction

1.1 A Brief History of Laser Filamentation

Following the invention of the laser in 1960 a new age of optical studies began. The laser enabled the exploration of a myriad of known phenomena which had previously been untestable. Additionally, the laser brought to light a slew of entirely new and enigmatic problems. One such discovery occurred in 1964 when Hercher [13] was performing a study on the failure modes of glasses used in ruby lasers. When observing the damage left by laser pulses fired into glass he noticed that increasing the power of the laser pulse did not simply increase the size of the damaged area. Instead there appeared to be some threshold power after which the laser pulses produced long, thin, filaments of plasma within the glass. As the input pulses were not initially intense enough to generate a plasma then there must have been some mechanism at work within the glass which was intensifying the pulse to the point that ionization could take place. These filaments soon attracted some attention and a number of studies were done to explain their origins however, progress was slow. Laser filamentation is a highly nonlinear optical phenomenon, this makes an analytical approach to the problem difficult. While a numerical approach is certainly more straightforward the computers of the day were simply not powerful enough to handle the scope of the simulations required. This hindered research into filaments for many years.

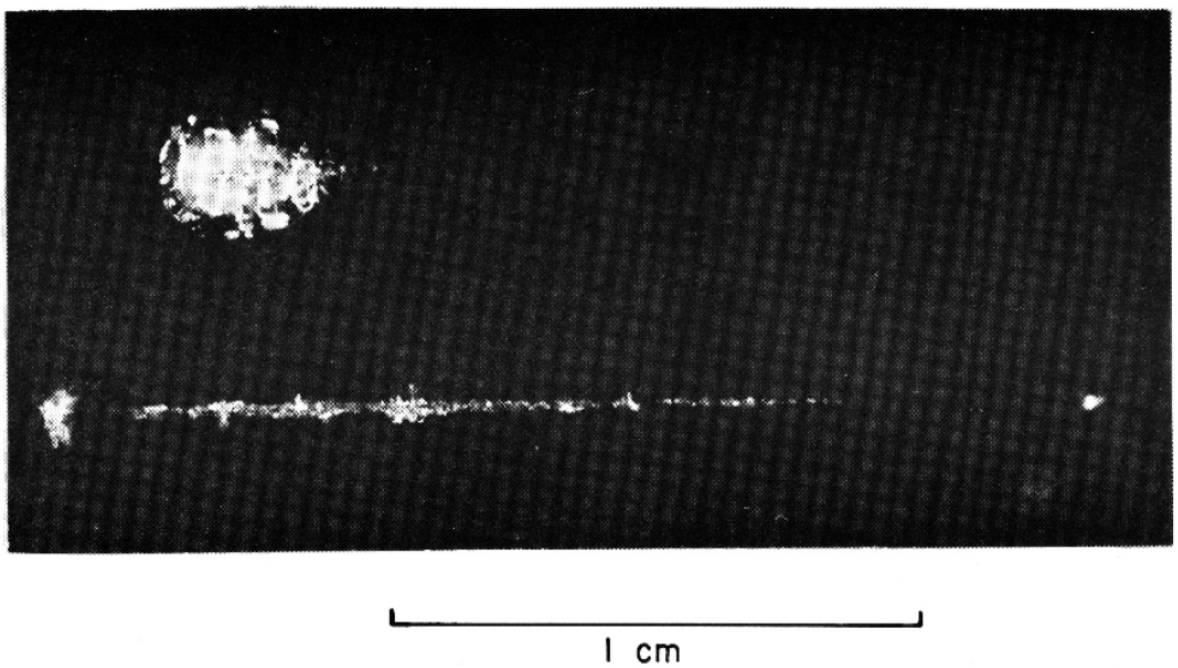

Figure 1.1: Photographs of the differing damage tracks left by a low and high power laser pulse in glass. *Source:* [13]

Studies finally began to make progress in the early 1990's with the advent of practical and powerful femtosecond lasers as well as the rapid increase in computer power. In the 1960's the lasers used in filamentation

experiments were limited to nanosecond pulses; the plasma generated by such pulses exhibits highly chaotic motions which are difficult to model. By reducing the duration of a laser pulse to the order of hundreds of femtoseconds the plasma generated within the pulse is restricted to motions largely in lock-step with the laser's electric field. Additionally, the mechanism of plasma generation changes from being primarily collisional to field ionization as the laser-plasma interaction time is far shorter than the timescale over which collisional ionization occurs. This simplified the modeling process dramatically and allowed the first major numerical studies of filamentation to be done. [21]

1.2 Previous Work

Filamentation in air has been the subject of much research as there are many applications for such filaments. For example one could perform laser spectroscopy on distant objects by controlling the position at which the plasma filament forms. Alternatively the filament can be used to generate a white light laser which can then be used to perform highly accurate LIDAR scans for imaging or chemical analysis through backscattering. A more dramatic example is the manipulation of weather; by making a filament which stretches from the ground to a cloud the plasma of the filament forms a conducting channel through which lightning can be discharged, or rain may be summoned by seeding cloud formation. [29]

Since air is such an interesting medium in which to experiment there have been many theoretical studies of filamentation in air. The review papers published by Couairon and Bergé [9, 8] list a variety of different beam propagation equations for modeling filamentation. Most of these models are variations on either the Nonlinear Envelope Equation (NEE) developed by Brabec and Krausz [19] or the nonlinear Schrödinger equation. The continuing trend among these models is the incorporation of ever more terms in their propagation equations: field effects, plasma effects, effects of a nonlinear medium. Ultimately there are inherent limitations to these models as they can not incorporate all effects while remaining tractable for simulation work.

Most commonly in these models the plasma is assumed to be frozen; the relatively heavy ions are assumed to be static and the motion of the electrons is assumed to follow strictly collinearly the applied laser field. True plasma dynamics effects, such as ponderomotive force effects and plasma evolution, driven by the plasma charge fields, are neglected; this results in a free electron density profile that only grows in time with the ionization of particles in the local area. This assumption has been made so far without tests in order to keep the problem computationally manageable. Furthermore, as the duration of the pulse increases, the validity of this assumption becomes questionable.

1.3 Motivation and Goals

The motivation behind this project came from the simplistic treatment of the plasma used in many papers. At some point treating the electrons in a filament as free electrons quivering in the applied field must break down and the impact of their dynamics within the plasma will begin to affect the filament. The goal of this project was to determine under which conditions the internal dynamics of the plasma have a noticeable impact on the evolution of a filamenting laser pulse and to determine what that impact is.

1.4 Project Outline

Firstly we needed two separate models for plasma: a dynamic model where the plasma particles were free to move throughout the simulation space and a static model where they were chained to the place of their birth. Then we needed a way to propagate a laser field with as high a fidelity as possible; necessary so that no plasma effects are missing due to an omitted field effect. Lastly we required a model for the linear and nonlinear response of the medium in which the laser is propagating. For this we developed a model of the induced polarization in a nonlinear medium [1] which, when coupled to the laser propagation routine, replicates the effects of propagation through a medium. By combining these components into a singular program we are able to simulate an ionizing laser pulse traveling through a nonlinear medium and thus reproduce a filament.

To test for a difference between plasma models a series of simulations were run in matched pairs: one using the static model and another using the dynamic model. In our simulations we shot laser pulses of various size, duration, intensity, and wavelength through air and tracked them for propagation times ranging between 20-40 picoseconds. By comparing the simulations after propagation we can determine whether the different plasma models have caused the outcomes of two otherwise identical simulations to diverge. It was our expectation that for very short pulses, durations on the order of 10 fs, there would be no appreciable differences between the two models. We expected a difference could be forced by altering the, duration, and waist size of the pulses which would elicit behaviours from the dynamic plasma model which are not supported by the static plasma model; what we found instead was that, while altering these parameters does elicit different behaviours from the two models of plasma, the two models of plasma are practically equivalent.

2 Theory of Filamentation

Early on in filamentation research two different models were put forth to explain the observed plasma filaments: the moving focus model [24], and the self-trapping model. [22] The moving focus model posited that the filament observed was being produced by the driving laser pulse being self-focused down into a streak rather than a point, a product of the significantly different powers successive time slices possess, which then generated the plasma filament. The second model, self-trapping , argued that the filament was the product of the dynamic balance between self-focusing and plasma defocusing. Despite its early success in matching nanosecond filament experiments the moving focus model fails to explain many features of a filament and has lost favor and the self-trapping model has become the dominant interpretation of filamentation. To understand this model we first need to understand the nonlinear effects it relies upon. In this section we describe some of the features of filamentation as well as the nonlinear processes which shape the evolution of a filament.

2.1 Optical Kerr Effect

When a strong EM wave passes through certain materials the polarization induced may not scale linearly with the applied field. The Kerr effect due to bound electrons, which occurs in most every transparent material, introduces a term to the electric susceptibility which scales with the intensity of the electric field.

2.1.1 Origin of the Kerr Nonlinearity

To first order the potential well binding an electron to a particle can be described as a quadratic curve. As the distance from the particle grows however this approximation breaks down and additional polynomial terms, such as in 2.1, must be added to maintain the validity of the approximation. Electrons in such a potential experience a linear restoring force when their dislocation from the center is small but begin to feel a nonlinear force as their dislocation grows. When subjected to a strong external oscillating field the electrons trapped by the potential are driven into anharmonic motion; this then causes the induced polarization of the particle to depend nonlinearly on the applied field.

$$U(r) \simeq C_1 r^2 + C_2 r^3 + C_3 r^4 + ... \tag{2.1}$$

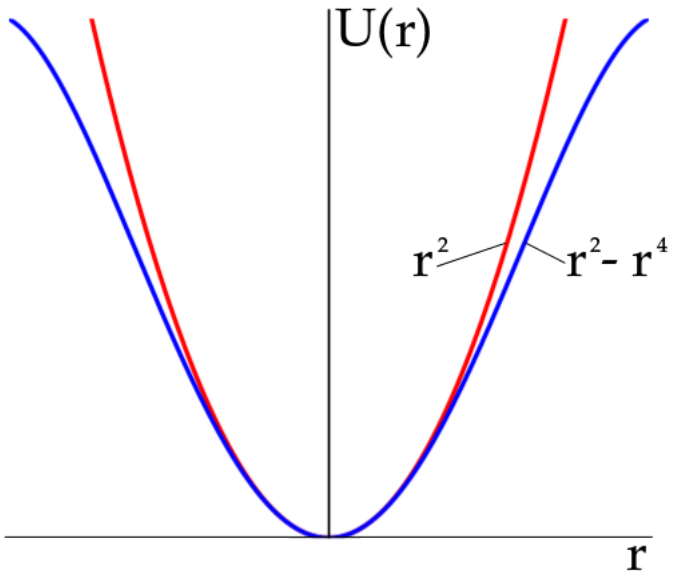

Figure 2.1: In this toy model of a Kerr medium's potential well it is clear that just the quadratic term is sufficient for short ranges but quickly fails at larger distances as the quartic term comes to dominate.

In a Kerr medium the potential well is centrosymmetric, that is to say $U(r) = U(-r)$, and as such all odd powers in the polynomial expansion of the potential must be zero. What this means is that the lowest order nonlinear term in the induced polarization, as written in equation 2.2, depends cubicly on the applied electrical field.

$$\mathbf{P}(\omega) \simeq \varepsilon_o \left(\chi^{(1)} + \frac{3}{4}\chi^{(3)} |\mathbf{E}|^2 \right) \mathbf{E} \tag{2.2}$$

The polarization can be rewritten in its general form as $\mathbf{P}(\omega) = \varepsilon_o \chi(\omega, E)\mathbf{E}$ by bundling the nonlinear dependency into the electric susceptibility such that $\chi = \chi^{(1)} + \frac{3}{4}\chi^{(3)}|E|^2 + ...$ In most materials terms

higher than $\chi^{(3)}$ are quite small and so we can discard them; this lets us rewrite χ into the tidy form of equation 2.3.

$$\chi = \chi^{(1)} + \frac{3}{2cn_o\varepsilon_o}\chi^{(3)}I \tag{2.3}$$

Where n_o is the linear index of refraction, and $I = \frac{cn_o\varepsilon_o}{2}|E|^2$ is the intensity of the applied electric field. Generally $\chi^{(3)}$ is a 3x3x3 tensor but in a Kerr medium it reduces down to a scalar dependant on the frequency of the applied field.

The Kerr nonlinearity produces a variety of different effects: self-phase modulation, self steepening, and the Kerr modulational instability to name a few. [28] While these certainly do impact the propagation of a filamenting laser pulse it should suffice to cover just two other effects which have the most significant impact on a pulse's development.

2.1.2 Third Harmonic Generation

As a laser pulse propagates through a Kerr medium a small portion of the pulse is converted into the third harmonic of the laser's frequency. This arises from the cubic term in the polarization which allows it to support odd numbered harmonics. Typically the harmonics are quite weak compared to the fundamental frequency, this can be seen in figure 2.2. In the figure a small but significant portion of the original signal has been converted into the third harmonic and a much smaller amount changed to the 5^{th}, 7^{th}, etc.

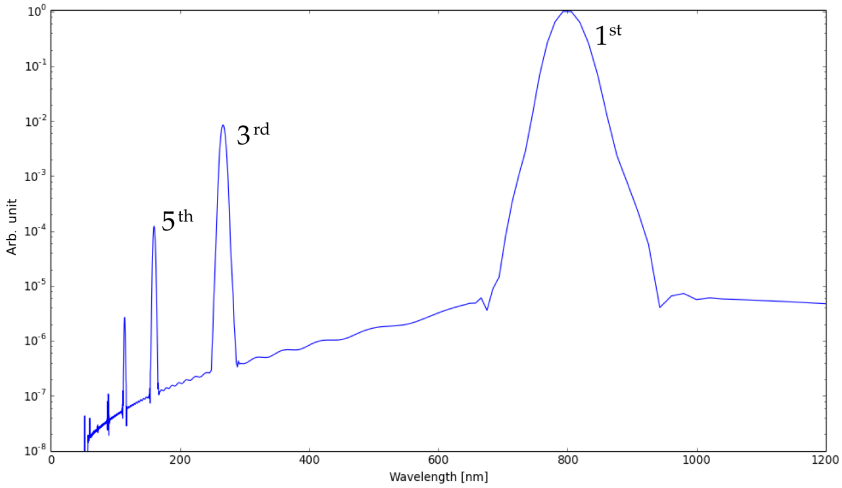

Figure 2.2: Spectrum of a pulse propagated through air for 1 ps normalized about the fundamental frequency. Visible here are the peaks at the fundamental 800 nm as well as the 3^{rd}, 5^{th}, and even 7^{th}. Higher modes appear to be present but are poorly resolved and extremely weak.

In our simulations we must have a spatial resolution fine enough to resolve the various wavelengths present in the spectrum. Despite the presence of higher harmonics we need only resolve down to the third harmonic as the higher modes are very much weaker and do not significantly impact the outcome of our simulations.

2.1.3 Self-focusing

Using equation 2.3 the refractive index of a Kerr medium can be rewritten to show that it too depends on the intensity of the applied field.

$$\begin{aligned} n(\omega, I) &= \sqrt{1+\chi(\omega, I)} \simeq 1 + \frac{\chi(\omega, I)}{2} \\ &= 1 + \frac{\chi^{(1)}(\omega)}{2} + \frac{3\chi^{(3)}(\omega)}{4cn_o\varepsilon_o} I \\ &= n_o + n_2 I \end{aligned} \tag{2.4}$$

Where n_2 is the nonlinear index of refraction and n_o the linear.

Consider now a laser pulse whose radial amplitude profile is gaussian, $E(r) = E_o e^{-\frac{r^2}{w_o^2}}$. As the intensity of such a pulse varies radially then so too will the refractive index. In a positive n_2 material the intense core of the pulse will experience a stronger index of refraction and so it lags behind the comparatively weaker wings. The difference in speeds between core and wings causes a concavity to develop in the wavefront of the pulse which in turn produces a focusing effect. This self-focusing effect causes a focused pulse to focus more strongly and reduces the divergence of a defocusing pulse. As the power of the pulse increases these nonlinear effects become increasingly more pronounced.

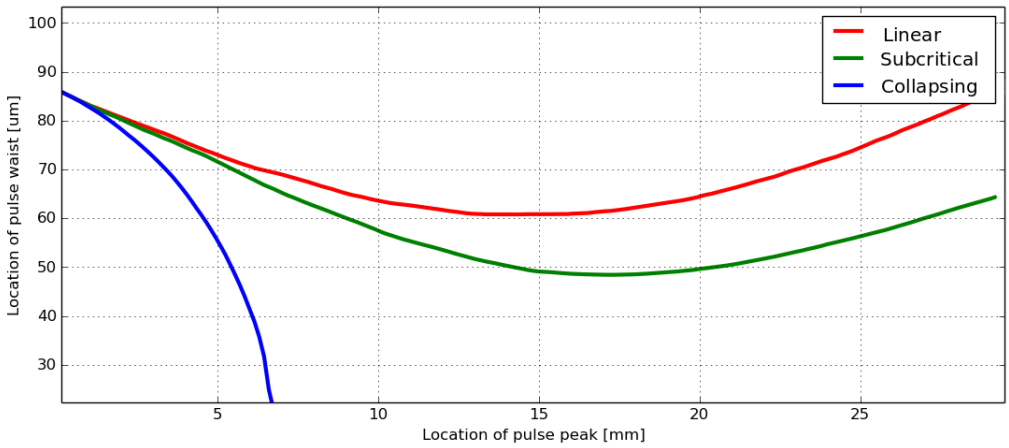

Figure 2.3: When a focused laser pulse enters air it is affected in three possible ways: a weak pulse (red) experiences only the linear index and so focuses and diverges as normal, a stronger pulse (green) experiences some self-focusing and so is focused to a smaller size then diverges more slowly, an extremely strong pulse (blue) will collapse upon itself as the self-focusing out-competes the natural divergence of the pulse.

For very powerful pulses it is possible that the self-focusing effect overcomes the natural divergence of the pulse [9]; the critical power for this behaviour is given by $P_{cr} = \alpha \frac{\lambda^2}{4\pi n_o n_2}$, where α is a parameter governed by the pulse's spatial profile. As a super-critical pulse is continuously focused to a smaller size its intensity increases to conserve energy which further strengthens the focusing effect. This feedback loop causes the pulse to rapidly collapse upon itself and drives the intensity to incredible levels. The collapse of the pulse can not continue indefinitely though; the dramatic increase in intensity will at some point trigger photo-ionization

of the Kerr medium. With the formation of plasma the collapse of the beam is halted as the plasma works to defocus the pulse.

2.2 Plasma Generation, Defocusing, and Dynamics

When the intensity of a laser beam is high enough its electric field will tear electrons away from their parent particles, ionizing the medium and generating a partial plasma. The interactions between the laser and its plasma result in some highly nonlinear behaviours which can significantly alter the shape and behaviour of the laser pulse. In this section we discuss the origin of the plasma and some effects it has on the laser which produced it.

2.2.1 Field Ionization

It has long been known that materials may be ionized when exposed to photons of a sufficiently high energy. This is because when an atom absorbs such a photon one of its electrons must jump to a higher energy state. If the energy absorbed by the electron is greater than the binding energy of the atom then the electron is ejected and the atom ionized. Single photon ionization is however, not the only means by which light can ionize a particle.

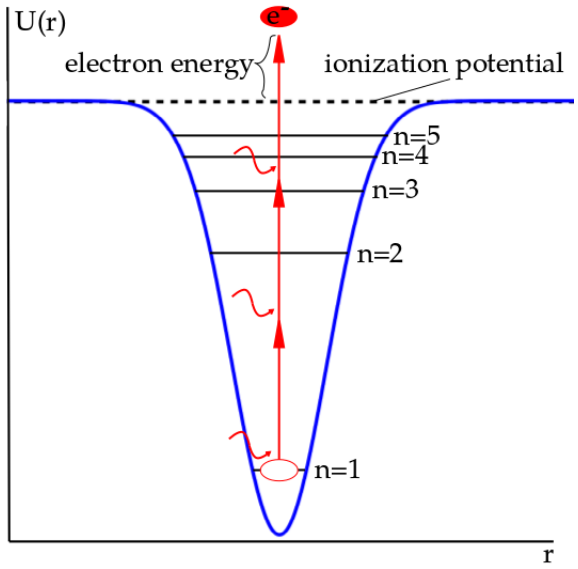

Figure 2.4: An atom which simultaneously absorbs a sufficient number of photons will be ionized. The electron released by the atom will carry away with it the difference between the ionization potential and absorbed photon energy.

Should an optical field be sufficiently intense, such as that found in a strong laser pulse, it is possible for an atom to experience multi-photon ionization (MPI). In MPI the energy carried by an individual photon is insufficient to photo-ionize a particle but, due to the high intensity, it becomes probable that an atom will simultaneously absorb several photons and thus an electron will gain sufficient energy to escape the

potential which binds it. In this case the ejected electron possesses some post ionization kinetic energy $E_k \geq K\hbar\omega - U_I$, where U_I is the ionization potential and K the minimum number of photons with frequency ω needed to ionize. In equation 2.5 we give the MPI ionization rate, where I is the field intensity and σ_K is a value characteristic of the ionizing medium. This equation is only valid for relatively weak fields as it is unbounded in infinitely strong fields. At higher field strengths a second mechanism of field ionization takes over.

$$w_{MPI}(I) = \sigma_K I^K \tag{2.5}$$

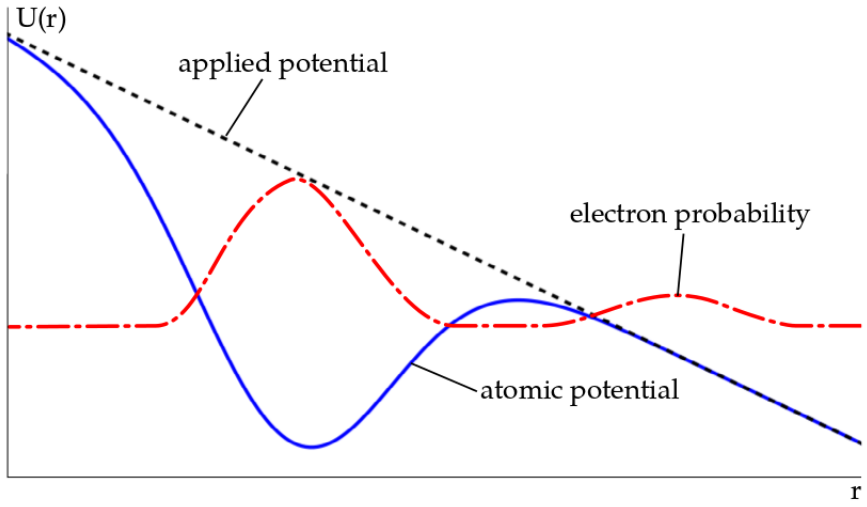

Figure 2.5: The suppression of an atomic potential by a strong external field increases the probability of a trapped electron being found outside the atom's potential well. In a strong enough field this tunneling probability is a major source of ionization.

$$w_{ADK}(E) = \frac{U_I^{2n^*}}{n^*\Gamma\left(n^*+l^*+1\right)\Gamma\left(n^*-l^*\right)}\left(\frac{3E}{\pi\left(2U_I\right)^{3/2}}\right)^{1/2} \frac{(2l+1)\left(l+|m|\right)!}{2^{|m|}\left(|m|!\right)\left(l-|m|\right)!}\left(\frac{\left(2U_I\right)^{3/2}}{E}\right)^{2n^*-|m|-1} exp\left(-\frac{2\left(2U_I\right)^{3/2}}{3E}\right) \tag{2.6}$$

As the amplitude of the laser's electric field increases it will come to rival the field which binds electrons to their parent particle. When this happens the potential barrier surrounding the particle will be suppressed and the probability of an electron tunneling through it becomes non-negligible. In equation 2.6 we give the tunnel ionization rate formulated by Ammosov, Delone, and Krainov (ADK). [5] It should be noted that the ADK rate gives the rate of tunnel ionization and as such, this drops to zero for low fields strengths as tunneling becomes improbable. At high fields ADK underestimates ionization as Barrier Suppression Ionization (BSI) allows bound electrons to escape their potential wells directly.

2.2.2 Plasma Defocusing

Consider now a linearly polarized laser pulse propagating through a plasma, $E = E_o cos(\omega t)$. As the pulse propagates its electric field induces in the plasma particles an oscillation which is collinear with the field such that the electron velocity is $v = \frac{q}{m\omega} E_o cos(\omega t)$. The oscillation of the ions is usually so small that they are practically motionless; their hefty masses make them unresponsive to optical frequency oscillations. The electrons, being at least three orders of magnitude lighter, respond strongly to the laser field and in doing so generate a current density which is out of phase with the driving field, $J = \frac{q_e^2 n_e}{m_e \omega} E_o sin(\omega t)$. This current then generates a polarization, $P = -\frac{q_e^2 n_e}{m_e \omega^2} E_o cos(\omega t) = \varepsilon_o \chi(\omega) E$, in which the electric susceptibility is negative. In a complete plasma this results in the index of refraction being a complex function but in a partially ionized medium it causes the local index of refraction to be reduced.

$$\chi_p(\omega) = -\frac{n_e q_e^2}{\varepsilon_o m_e \omega^2} \tag{2.7}$$

where n_e is the electron population density and q_e and m_e are the electron charge and mass respectively.

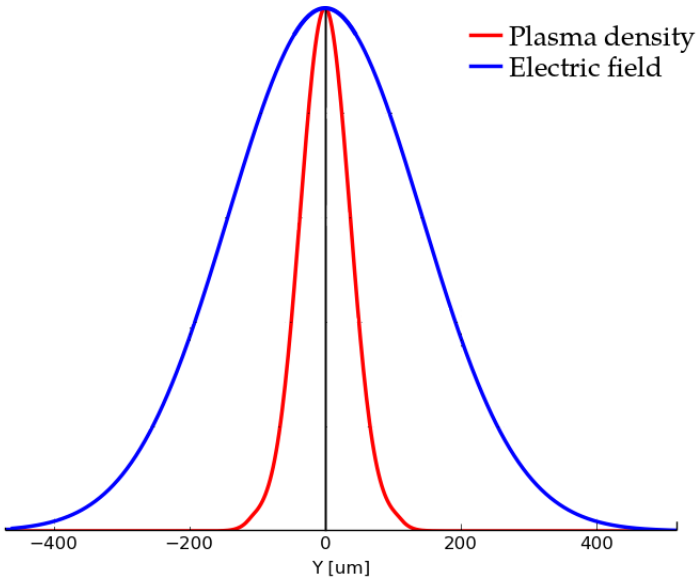

Figure 2.6: Due to the highly nonlinear scaling of the ionization rate the plasma generated by a gaussian pulse has a density profile much narrower than electric field.

Consider now an ionizing gaussian pulse as in figure 2.6. The plasma generated at the core of the pulse will be much denser than that generated nearer the wings. Due to the dependance of χ_p on plasma density this means the index of refraction at the core of the pulse is substantially lower. This means the core of the beam is traveling somewhat faster than the wings of the pulse and so the core pulls ahead to form a convex, defocusing, wave front.

The defocusing effect of a plasma is particularly strong. For instance, in air, an 800 nm pulse with an intensity of the order 10^{13}W/cm^2 can generate a plasma with a density approximately 4% that of air which causes

a change in the refractive index $\Delta n_p \sim$ -10^{-4} whereas the change caused by the Kerr effect $\Delta n_{nl} \sim 10^{-5}$. Clearly the plasma density required to counteract the collapse of a pulse does not need to be particularly high.

2.2.3 Ponderomotive Force

For a linearly polarized laser pulse the oscillations induced in a plasma work to reduce the electric susceptibility, as stated in the previous section. In general however, the forces acting on the plasma particles are more complicated. The Lorentz force, equation 2.8, controls the motions of charge particles:

$$\mathbf{F}_L = q\left(\mathbf{E} + \mathbf{v} \times \mathbf{B}\right) \tag{2.8}$$

where q and $\mathbf{v}$ are the charge and velocity of the particle and $\mathbf{E}$ and $\mathbf{B}$ are the electric and magnetic field at the particle's position.

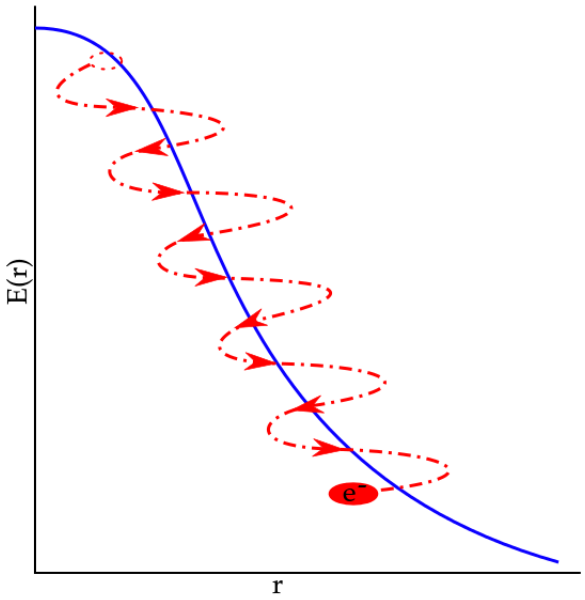

Figure 2.7: An electron, subject to the ponderomotive force, is excluded from regions of high field strength and forced into low field regions.

In a continuous, uniform, oscillating electric field an electron is a simple harmonic oscillator but in a spatially varying electric field, such as in figure 2.7, additional forces are at play. Over the course of the electric field's cycle an electron may be shuttled from a region of high field to a low field and back; as a result of this the electron feels an unbalanced different restoring force and thus it does not return to its original position. In an EM field there is an additional drift caused by the $\mathbf{v} \times \mathbf{B}$ force which curves the electron's trajectory in the direction of the field's propagation. Due to the phase mismatch between $\mathbf{v}$ and $\mathbf{B}$ an electron feels a force pushing along the laser's propagation vector. After some work the Lorentz force can be used to show that these two effects result in an effective force which excludes charged particles from regions of high field intensity, this is the ponderomotive force.

$$\mathbf{F}_p = -\frac{q_e^2}{4m_e\omega^2}\nabla\left(|E|^2\right) \tag{2.9}$$

The ponderomotive force scales with the gradient of the intensity of the laser field. This implies that plasma generated by an extremely small and intense laser pulse would feel a powerful force expelling it from the center of the pulse. In extreme circumstances this effect can be exploited to produce laser-wake accelerators but in our simulations the impact is likely far more modest. It may be the case that the ponderomotive force causes the plasma density profile to become broader, less sharply peaked, as electrons are forced from the core. This may weaken the defocusing effect of the plasma.

2.3 Self-trapping Model

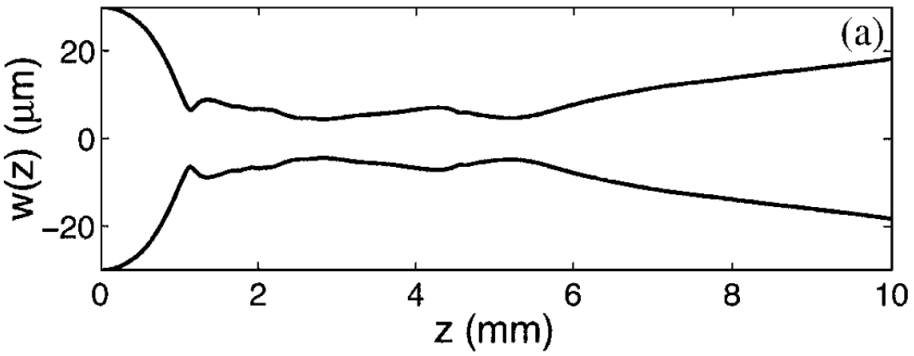

Figure 2.8: A pulse with a super-critical power self focuses down until a plasma begins to form. Once plasma generation has begun the pulse will gradually cycle between defocusing and focusing, producing a string of plasma filaments. *Source:* [23]

Understanding now the opposing natures of plasma and the Kerr effect the self-trapping model of filamentation is rather intuitive. As a laser pulse traveling in a Kerr medium collapses upon itself its intensity rises. Once the intensity rises above the ionization threshold a plasma abruptly forms which halts and then reverses the collapse of the pulse. The pulse can then propagate for an extended distance as the Kerr effect and plasma defocusing compete with each other, leaving behind a tail of luminescent plasma. As the pulse continues its propagation it gradually loses intensity due to energy losses from ionization. Eventually the pulse's intensity falls below the ionization threshold and ceases to generate a plasma, terminating the filament. If the pulse retains enough power it may then undergo another round of self-focusing which triggers the formation of a secondary filament; this cycle of focusing-defocusing can be repeated a number of times, producing a train of filaments.

2.3.1 Properties of a Filament

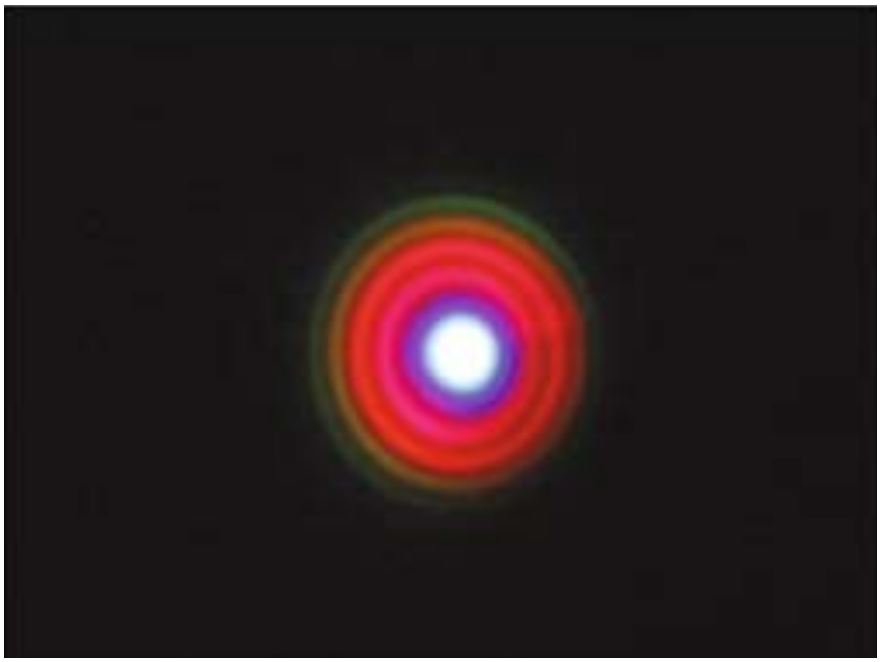

Figure 2.9: Cross sectional image of a filamenting 800 nm, 45 fs, 5 mJ, laser pulse after propagating through 25 m of air. *Source:* [26]

While a pulse filaments, its shape and spectral content undergo massive alterations. Figure 2.9 demonstrates some of these features quite clearly. At the center of the filament we notice a bright white spot; during the course of its propagation the core of the largely monochromatic laser pulse has been converted into a white light laser through super continuum generation. Super continuum generation is a catch-all term for any nonlinear effect which causes substantial spectral broadening. In filaments this it is due to self-phase modulation from the Kerr nonlinearity as well as the electron oscillations in the plasma.

The second notable change in the frequency content of the pulse occurs in the colored rings which surround the core. As a pulse filaments it casts off these colored rings which, when viewing their spatio-temporal distribution, appear as concentric conical rings; near the core the spectral content of these rings is dominated by the frequency of the laser pulse, rings further out become tinged with increasingly higher frequencies.

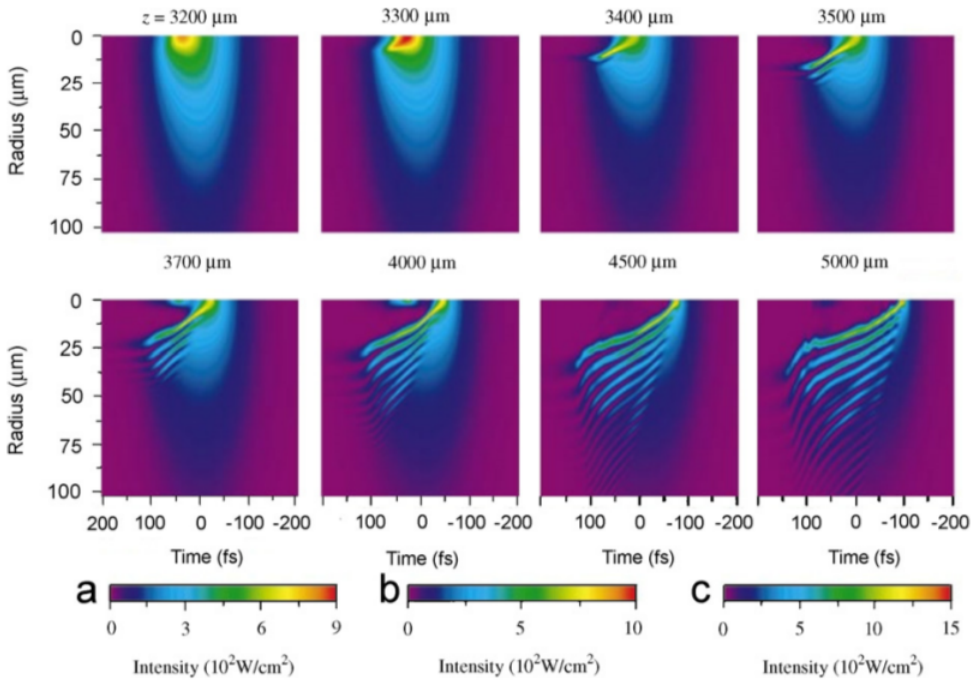

Figure 2.10: Spatio-temporal evolution of a filament in fused silica. *Source:* [25]

When viewing the spatio-temporal distribution of a filamenting laser pulse we can see how the various features of a filament develop. As the pulse propagates the plasma generated at the center of the pulse allows the central core of the pulse to compress into an optical shockfront much shorter than the original pulse. Meanwhile the flanks of the pulse develop into long, intense, tails. Were we to extrapolate this structure into 3D it would look something like a hollowed out bullet. As the pulse propagates, successive layers begin to peel off from the tails; these are the rings in the conical emission. Not clearly show in figure 2.10 is how a refocusing event, as described in section 2.3, occurs. The head of the central shockfront continually casts off layers which feed into the the center of the pulse where it is redistributed into layers of concentric tubular structures. As the pulse propagates these tubes grow in strength and size, slowly they collapse upon themselves through self-focusing and, as they fuse, form a secondary shockfront within the first. This secondary front siphons energy from the leader and eventually becomes the dominant one. This process will repeat itself until the pulse power falls below the critical power for self-focusing.

3 Methods and Models

For this project we chose to use air as the medium in which to run our simulations. This choice was made for a number of reasons. On the computational side, modeling the plasma dynamics of air is simpler as it is much less dense than other filament forming materials like glass or water, thus particle collisions can be safely neglected. Furthermore the linear and nonlinear responses of air have been well documented, as has its ionization characteristics. On the practical side, filaments in air are of particular interest as they can be exploited for a number of interesting applications. In this section we describe the methods and models

which were employed to model the propagation of a laser pulse through air. Specifically we describe how each model we use operates as well as how they are implemented and validated.

3.1 Beam Propagation

As mentioned in section 1.4 our intent was to include all possible field effects in our simulations. There are several methods by which this can be accomplished. The method we chose, the Finite Difference Time Domain (FDTD) method, does this by directly solving Maxwell's equations. FDTD is a very popular method as it scales with simulation size in an approximately linear fashion, somewhat better than other methods like the NlogN scaling of the Pseudo Spectral Time Domain (PSTD) method.

3.1.1 Finite Difference Time Domain Method

In 1966 Yee showed that when the differential forms of Maxwell's curl equations are discretized in both time and space they may be rearranged into a pair of fully explicit field propagation equations with second order accuracy. [11]

$$\begin{aligned} \nabla \times \mathbf{E} &= -\frac{\partial \mathbf{B}}{\partial t} \\ \nabla \times \mathbf{B} &= \mu_o \left(\mathbf{J} + \varepsilon_o \frac{\partial \mathbf{E}}{\partial t} \right) \end{aligned}$$

$$E_x^{n+1/2,i,j,k} = \frac{\Delta t}{\varepsilon_o} \left(\frac{1}{\mu_o} \left(\frac{B_z^{n,i,j,k+1/2} - B_z^{n,i,j,k-1/2}}{\Delta y} - \frac{B_y^{n,i,j+1/2,k} - B_y^{n,i,j-1/2,k}}{\Delta z} \right) - J_x^{n,i,j,k} \right) - E_x^{n-1/2,i,j,k} \tag{3.1}$$

$$B_x^{n+1,i,j,k+1/2} = \Delta t \left(\frac{E_z^{n+1/2,i,j,k+1} - E_z^{n+1/2,i,j,k-1}}{\Delta y} - \frac{E_y^{n+1/2,i,j+1,k} - E_y^{n+1/2,i,j-1,k}}{\Delta z} \right) - B_x^{n-1,i,j,k+1/2} \tag{3.2}$$

Where n is the time index and, i, j, k are spatial indices in x, y, and z, respectively. Δt, Δx, Δy, and Δz are the step sizes for time, x, y, and z respectively.

By spatially staggering the electric and magnetic fields, as shown in figure 3.1, the spatial derivatives in equations 3.1 and 3.2 become centered finite differences with second order accuracy. A similar trick can be used to achieve second order accuracy in time by leap-frog integrating the field equations, first **E** then **B**.

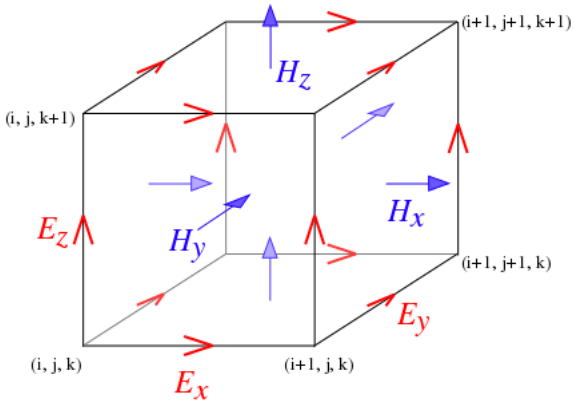

Figure 3.1: In a 3D Yee lattice the electric and magnetic field components are spatially staggered to increase FDTD accuracy without increasing spatial resolution. *Source: http://commons.wikimedia.org/*

The stability of the FDTD method is governed by the Courant condition which demands that the grid points of the simulation space be causally linked to each other. When making a step through time the grid point being updated calls on the grid points which surround it and as such these points can not be further away that light can travel during that time step. This puts a cap on the time step $\Delta t \leq \frac{\Delta x}{c}$. In higher dimensions we must take into account propagation in diagonal directions and thus the time step must be made smaller such that $\Delta t \leq \frac{\Delta x}{\sqrt{n}c}$, where n is the dimensionality of the simulation. In theory a simulation can be run at the maximum time step but in practice it us usually a better idea to run with a somewhat smaller step as instabilities may creep in through peculiarities of the system or other operations being carried out within the simulation. In EPOCH, the FDTD program we use, the time step is by default 95% the maximum possible.

3.1.2 EPOCH - the Extendable PIC Open Collaboration

EPOCH is an FDTD-PIC plasma simulator produced chiefly by Chris Brady of the University of Warwick but has had much work done on it by many others. The program was originally written for simulating such violent processes as inertial confinement fusion but has since been expanded upon to allow it to simulate a broad array of physical phenomena. We chose EPOCH as the basis for our simulations as it boasts three major features:

- The entire program is already optimized for massively parallel computation. This is a very attractive feature as our simulations require massive computing power and writing a fully parallel code is a great deal of work.
- EPOCH is written in a modular fashion which makes it easy to modify and implement new features. Since our work required new models to be included this meant we needed only to write new modules to be integrated into the overall program.
- EPOCH has already built into it ionization and particle-in-cell routines, both of which are very difficult to implement and parallelize. We expand upon these features later in section 3.5.

3.1.3 Pulse Tracking

Simulating a filamenting laser pulse means we must follow it over a distance which is much larger than the size of the pulse. Simulating the entirety of the space through which the pulse passes would be hugely inefficient. It is far more efficient to instead only simulate the space immediately surrounding the pulse. In most numerical studies this is accomplished by shifting the reference frame from the lab to the retarded time frame where $t' = t - \frac{|\mathbf{r}-\mathbf{r}'|}{v_g}$, where v_g is the group velocity of the pulse in the medium. This is not something which can be easily implemented into EPOCH so instead we use a numerical trick which mimics the retarded time frame by shifting the boundaries of the simulation space. By knowing the average group velocity of the pulse in the medium we can calculate how many grid points the pulse will traverse during a time step of the simulation, calculated by $n_{shift} = \left\langle \frac{v_g \Delta t}{\Delta x} \right\rangle$. Once the pulse has crossed at least one grid point all field values, electric, magnetic, current densities, polarizations, are shifted back by the number of grid points crossed and the grid points at the leading edge of the simulation are zeroed. Should the pulse travel a non-integer number of grid points then the non-integer component is allowed to accumulate over successive time steps until it is greater than one at which time the fields are shifted an additional grid point. This 'moving window' changes none of the physics and effectively keeps the pulse stationary in the simulation space.

3.2 Polarizable Nonlinear Air Background

The bare FDTD routine solves only for the propagation of a field therough a vacuum; to account for the effects of a medium a secondary model must be coupled to the FDTD. A common method uses an auxillary oscillator model for the induced polarization. From such a model you can find the polarization current density, $\mathbf{J}_p = \frac{d\mathbf{P}}{dt}$, which is then fed into the global current density of equation 3.1. Doing this is functionally equivalent to solving Maxwell's equations in a medium but neatly avoids the implicit differentiation scheme that entails. To get the correct polarization current we needed a model which could accurately replicate both the linear and nonlinear polarization of a Kerr medium in an explicitly integrable way. One such model, as given by Jackson et al. [12] is an anharmonic oscillator for air. Early trials using this model proved it to be unsuitable for our purposes as it is totally unstable at the field intensities used in our simulations. In order to add the Kerr nonlinearity to EPOCH we required a new model which was stable even at the intensities found inside filaments.

3.2.1 Saturable Lorentz Oscillator

The saturable Lorentz oscillator model as described by C. Varin et al. in [1] was the solution to our polarization problems. It replicates the saturability of the nonlinear index in high fields which prevents the model from destabilizing but, unfortunately, its nonlinearity is negative which results in self-defocusing. To get this model to produce self-focusing we rewrite the model such that it reads:

$$\frac{d^2\mathbf{P}}{dt^2} = \omega_o^2 \varepsilon_o \chi^{(1)} \left(2 - \frac{1}{1 + |\mathbf{E}|^2 \chi^{(3)}/\chi^{(1)}} \right) \mathbf{E} - \omega_o^2 \mathbf{P} \tag{3.3}$$

Where ω_o is the resonance frequency of the air, $\chi^{(1)}$ and $\chi^{(3)}$ are, respectively, the linear and nonlinear electrical susceptibilities.

With this correction made the model is ready to be implemented into EPOCH.

3.2.2 Implementation of Nonlinear Background

Equation 3.3 can be discretized and rearranged to give:

$$\begin{aligned} \mathbf{J}^{n+1/2,i,j,k} &= \mathbf{J}^{n-1/2,i,j,k} + \Delta t \left(\omega_o^2 \varepsilon_o \chi^{(1)} \left(2 - \frac{1}{1 + |\mathbf{E}^{n,i,j,k}|^2 \chi^{(3)}/\chi^{(1)}} \right) \mathbf{E}^{n,i,j,k} - \omega_o^2 \mathbf{P}^{n,i,j,k} \right) \\ \mathbf{P}^{n,i,j,k} &= \mathbf{P}^{n-1,i,j,k} - \Delta t \mathbf{J}^{n+1/2,i,j,k} \end{aligned} \tag{3.4}$$

Where n is the time index and, i, j, k are spatial indices in x, y, and z, respectively.

To include the index of refraction of air we assign each grid point in the simulation space a polarization and polarization current. As in FDTD, the nonlinear background is leap-frog integrated to give it second order accuracy. At the end of the update the polarization current is added to the global current density shown in equation 3.1.

To configure the nonlinear background to approximate air we set the nonlinear electric susceptibility at $\chi^{(3)}$= 1.305x10^{-25} m^2/W, somewhat higher than the value in air for an 800 nm laser for which $\chi^{(3)}$= 1.062x10^{-25} m^2/W. It is a reasonable simplification to make $\chi^{(3)}$ constant as the central frequency of our simulations does not substantially change, as can be seen in figure 2.2 the central frequency remains dominant. Getting the correct linear response turned out to be slightly trickier as the Sellmeier equation which describes it, equation 3.5, has two resonance bands. Were our driving frequency near one of the resonance frequencies, ω_1 and ω_2, we would have to use two separate oscillators to model the linear response properly. Fortunately the frequencies we are interested in are far from resonance and we can thus approximate the linear response with one oscillator by simply scaling $\chi_1^{(1)} = \frac{\Omega_1^2}{\omega_1^2} a$, where the scaling factor a = 1.2287.

$$\chi^{(1)}(\omega) = \frac{\Omega_1^2}{\omega_1^2 - \omega^2} + \frac{\Omega_2^2}{\omega_2^2 - \omega^2} \tag{3.5}$$

Ω_1 = 6.411x10^{14} rad/s	Ω_2 = 1.092x10^{14} rad/s
ω_1 = 2.906x10^{16} rad/s	ω_2 = 1.427x10^{16} rad/s

Table 3.1: Constants defining the two resonance bands of the Sellmeier equation for air. *Source:* [2]

3.2.3 Validation of Implementation

To test our implementation of the nonlinear background we ran a simulation in two dimensions where only FDTD and saturable air model were active. For the test a collimated pulse, gaussian in both space and time, was propagated from its waist through 2.7 cm of air. The critical power for a gaussian pulse in air is ∼3.2 GW, our pulse was given a peak power ∼9.8 times the critical power so as to ensure its prompt collapse. Throughout propagation measurements of the pulse's waist were made.

Wavelength	Initial waist	Intensity	Propagation time
800 nm	200 μm	5x10^{13} W/cm^2	90 ps
Polarization	Propagation vector	Axial resolution	Transverse resolution
$\hat{y}$	$\hat{\mathbf{k}} = \hat{x}$	3800 cells	900 cells

Table 3.2: Simulation parameters of nonlinear background test run.

The data collected from the EPOCH simulation was then compared to a simple paraxial wave equation model:

$$\frac{\partial}{\partial x}U = \frac{i}{2k_o}\nabla_\perp^2 U + i\frac{\omega_o}{c}n_2 |U|^2 U \tag{3.6}$$

Where $U = U_o e^{i\phi}$ is the complex amplitude of the laser which contains both the field amplitude, U_o, and phase information, ϕ. To simulate a pulse using this model we drop down to two dimensions (x,y) and discretize it:

$$U^{j+1,k} = U^{j,k} + \Delta x\left(\frac{i}{2k_o}\frac{U^{j,k+1} - 2U^{j,k} + U^{j,k-1}}{\Delta y^2} + i\frac{\omega_o}{c}n_2 \left|U^{j,k}\right|^2 U^{j,k}\right) \tag{3.7}$$

Where j and k are spacial indicies for x and y, respectively. In this form solving for the evolution of a laser pulse's envelope is as simple as integrating over the direction transverse to its propagation. For the paraxial simulation we initialize the complex envelope such that it corresponds to a transverse slice taken at the geometric center of a gaussian pulse identical to that of the EPOCH simulation: the envelope is a gaussian with a waist is 200 μm and initial phase of 0 (envelope is entirely real).

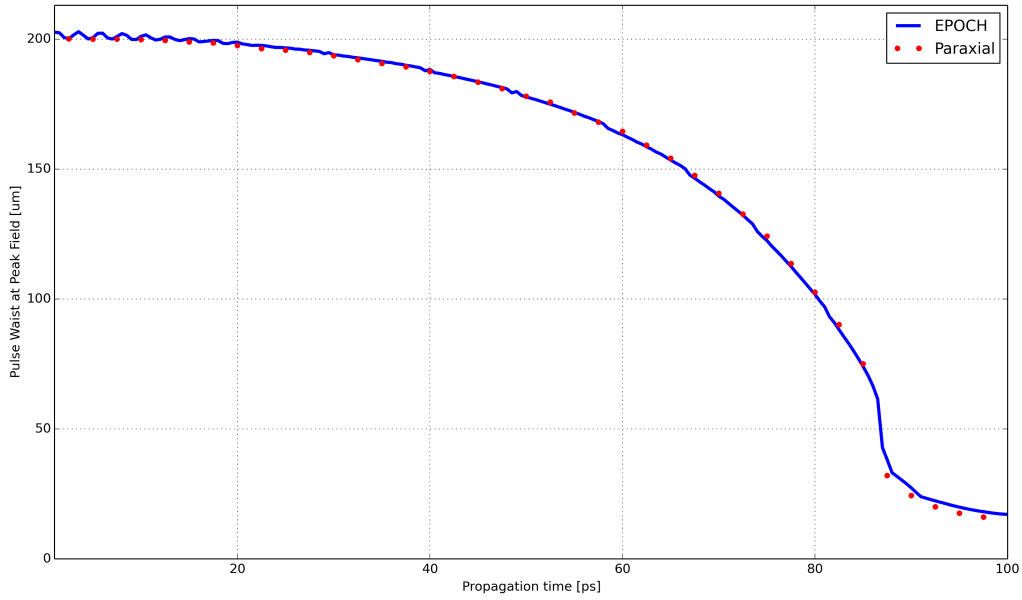

Figure 3.2: A comparison between the paraxial wave model of equation 3.6 and the saturable oscillator model implemented in EPOCH.

When comparing the output of the two models we can see that there is a small disagreement between the

paraxial model and EPOCH. This is likely due to the fact that the location in which the pulse waist was measured in the two simulations was necessarily different. In the paraxial model measuring the waist of the pulse is then as easy as finding the position at which the envelope amplitude is equal to $|U_p|/e$, where $|U_p|$ is the peak field amplitude, which remains at the geometric center of the pulse. As EPOCH does not explicitly output the pulse envelope we first recover it using the method outlined in appendix A. We then find the location of the peak field along the central axis and measure the pulse waist at this point. As the location of the peak field may differ from the geometric center of the pulse this may cause the slight difference between the two simulations. Overall the two models are in very good agreement which confirms the correct implementation of nonlinear polarizable background in EPOCH.

3.3 Calculating Ionization Rates

In his 2013 doctoral thesis, Alistair Lawrence-Douglas [7] implemented into EPOCH an ionization routine. While this routine is very well implemented, and has been rigorously tested, there was one glaring flaw in this implementation which we needed to correct for our simulations. For the field strengths we were interested in the ionization rate was calculated using one of two rate equations; if the electric field strength was relatively low then the MPI rate, equation 2.5, was used and if it was high enough the ADK tunneling ionization rate, equation 2.6, was used. The problem with this is that the two rate equations are valid over a relatively narrow range of field strengths and disagree strongly with each other for intermediate fields. To remedy this, Lawrence-Douglas applied a crude numerical patch. For fields where the MPI rate was greater than the ADK rate the ionization rate was held constant, as shown in figure 3.3. Unfortunately the varying intensity found in a filamenting pulse can straddle this awkward patch. To keep the ionization rate physical we had to implement an ionization rate equation which remained valid in both the MPI and tunneling regimes.

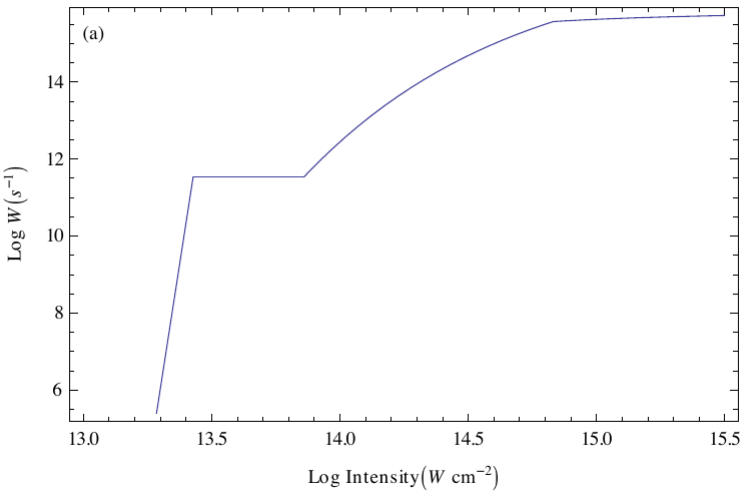

Figure 3.3: Ionization rates in vanilla EPOCH are held constant for field values where the MPI rate is higher than the tunneling rate. *Source:* [7]

3.3.1 Perelemov, Popov, Terentyev Ionization Rate

In 1965 Keldysh published a paper [3] in which he derived a rate equation which accurately predicted the ionization rate for hydrogen in both the MPI and tunneling ionization regimes. To fix the ionization rates in EPOCH we implemented a rate equation based on this model which was derived by Perelemov et al. in 1967. [4] The PPT rate equation was chosen over the original Keldysh formulation as it accounts for electron-ion interactions, making it somewhat more accurate. In its entirety the PPT rate equation is rather a formidable calculation but, by assuming the quantum numbers l and m are zero, it can be reduced down to the the following:

$$\begin{aligned} w_{PPT} =& \frac{4\sqrt{2}}{\pi\hbar}|C_{n^*l^*}|^2 U_I \frac{\gamma^2}{1+\gamma^2}\left(\frac{2F_o}{E\sqrt{1+\gamma^2}}\right)^{(2n-3/2)} \\ & \exp\left(\frac{-E_o}{E\gamma}\left[\operatorname{asinh}(\gamma)\left(1+\frac{1}{1+2\gamma^2}\right)-\frac{\sqrt{1+\gamma^2}}{2\gamma}\right]\right) \\ & \sum_{k\geq\nu_o}^{\infty}\exp\left(-2\left(\operatorname{asinh}(\gamma)-\frac{\gamma}{\sqrt{1+\gamma^2}}\right)(k-\nu)\right)\Phi\left(\sqrt{\frac{2\gamma}{\sqrt{1+\gamma^2}}(k-\nu)}\right) \end{aligned} \tag{3.8}$$

Where U_I and U_H are the ionization potentials of the target and hydrogen atoms, respectively. Below are the definitions of the other variables of equation 3.8.

$\gamma = \frac{\omega\sqrt{2m_e U_I}}{q_e E}$ is the Keldysh adiabacity parameter. It is a comparison between electron tunneling frequency and laser frequency. A $\gamma < 1$ means the time it takes an electron to tunnel ionize is comparable to the duration of an optical cycle and thus it is probable that the electron tunneled through the potential barrier instead of multi-photon ionizing.

$n = Z\sqrt{\frac{U_H}{U_I}}$ is the quantum number and Z the residual ion charge.

$\nu = \frac{U_I}{\hbar\omega}\left(1+\frac{1}{2\gamma^2}\right)$ is the fractional number of photons required to ionize the target atom, thus the minimum number of photons required for multi-photon ionization is $\nu_o = \langle\nu\rangle + 1$ where $\langle\nu\rangle$ is the integer component of ν.

$E_o = \frac{q_e^5 m_e^2}{\hbar^4(4\pi\varepsilon_o)^3}\left(\frac{U_I}{U_H}\right)^{3/2}$ is the characteristic electric field binding the electrons to the atomic core.

$|C_{n^*,l^*}|^2 = \frac{2^{2n^*}}{n^*\Gamma(n^*+l^*+1)\Gamma(n^*-l^*)}$ is a dimensionless constant and a function of the effective quantum numbers $n^* = \sqrt{\frac{U_H}{U_I}}$, $l^* = -n - \sqrt{\frac{U_H}{U_I}}$

The final component is the function $\Phi\left(\frac{2\gamma}{1+\gamma^2}(k-\nu)\right)$ where Φ is the Dawson integral, $\Phi(x) = e^{-x^2}\int_0^x e^{y^2}dy$.

The PPT rate, as stated above, does a very good job of predicting the ionization rates of hydrogen-like atoms but consistently underestimates the rate of ionization of molecules. This is unfortunate for us as the medium we wish to be ionizing, air, is largely composed of the diatomic molecules N_2 and O_2. All is not lost however, as others have already encountered this difficulty and found various fixes can be applied to boost the ionization rates to match reality. For instance Talebpour et al. [14] found that a modification could be made to Z to boost up the ionization rate. The method we chose instead appends a prefactor to the rate equation. This prefactor, found by Mishima et al. [6], boosts the ionization rate substantially to match experimental observations.

$$W_{Mol} = \frac{16\left(2\gamma + 3\right)}{3\left(1+\gamma^2\right)^2} \tag{3.9}$$

3.3.2 Implementation

In EPOCH we calculate the ionization rate using equation 3.8. We set the quantum numbers l and m to zero as accounting for them would only alter the rate of plasma formation while we are more interested in maintaining the physicality of the ionization rate through the transition between MPI and tunnel ionization.

The first hurdle we had to pass in order to calculate the PPT ionization rate was the solution to the Dawson integral, $\Phi\left(x\right)$. Solving the Dawson integral to an acceptable degree of accuracy turned out to be surprisingly challenging. We tried numerous means of solution such as expansion approximations, built in FORTRAN functions, direct numerical integration, and all failed in some respect. Ultimately we found the ideal solution was a 5^{th} order real approximation. This expansion remains accurate down to the sixth decimal place across all ranges of x.

$$\Phi\left(x\right) \simeq x\frac{1 + x\left(a_1 + \left(b_1 + x\left(c_1 + x\left(d_1 + e_1 x\right)\right)\right)\right)}{1 + x\left(a_2 + \left(b_2 + x\left(c_2 + x\left(d_2 + x\left(e_2 + 2x0.000178997\right)\right)\right)\right)\right)} \tag{3.10}$$

	a_i	b_i	c_i	d_i	e_i
1	0.1049934947	0.0424060604	0.0072644182	0.0005064034	0.0001789971
2	0.7715471019	0.2909738639	0.0694555761	0.0140005442	0.0008327945

Table 3.3: Constants used in 5th order real approximation of Dawson's integral.

With the Dawson integral handled we are able to solve for the PPT rate but doing so requires computing the infinite sum:

$$S = \sum_{k \geq \nu_o}^{\infty} \exp\left(-2\left(\operatorname{asinh}(\gamma) - \frac{\gamma}{\sqrt{1+\gamma^2}}\right)\left(k - \nu\right)\right) \Phi\left(\sqrt{\frac{2\gamma}{\sqrt{1+\gamma^2}}\left(k-\nu\right)}\right) \tag{3.11}$$

If we had an infinite amount of time then solving for S would be no obstacle but, since our time is finite, we instead elected to incrementally solve for S until it converges on its true value. In EPOCH we cut off the summation of S when the value found in the latest iteration is 10^{-10} percent different from the previous iteration. For large values of γ this only takes a few dozen iterations but for small values the sum may not converge on a value for many thousands, or perhaps millions, of iterations. In such an instance we terminate the sum on its millionth iteration as it is usually well enough converged by this point.

Since calculation of S can be so cumbersome we precalculate a table of possible S values at the beginning of a simulation. In our simulations U_I and ω are invariant which makes S a function of γ. With a table of S values we can calculate the sum using a linear interpolation between table entries instead of calculating it on the fly. This is tremendously quicker than finding S directly and drops calculation times down to a usable level without compromising the accuracy of the PPT rate.

$$S\left(\gamma\right) \simeq \frac{S^{k+1} - S^k}{\Delta\gamma}\left(\gamma - n\Delta\gamma - \gamma_{min}\right) + S^k \tag{3.12}$$

Where k is the table index found by taking the integer component of $k = \langle(\gamma - \gamma_{min})/\Delta\gamma\rangle$ where γ_{min} is the smallest value of γ for which S is calculated.

In the current implementation a table of S values is loaded into memory at runtime for every energy level of every ionizing species specified in the input file. This table is comprised of 2.6 million entries for γ values within the range $\gamma_{min} \leq \gamma \leq 500$ where $\gamma_{min} = \gamma(10E_o)$, the adiabaticity parameter for a field ten times the strength of the input pulse amplitude. The table is terminated at $\gamma = 500$ as ionization at such a high γ is highly improbable. In the likely event that $\gamma_{min} < 1$ the table will be broken into two sections with one million entries for $\gamma < 1$ and the remainder for $1 \leq \gamma \leq 500$. This splitting is done because the sum begins to grow and oscillate rapidly when $\gamma < 1$, necessitating a great many data points for accurate interpolation.

3.3.3 Validation

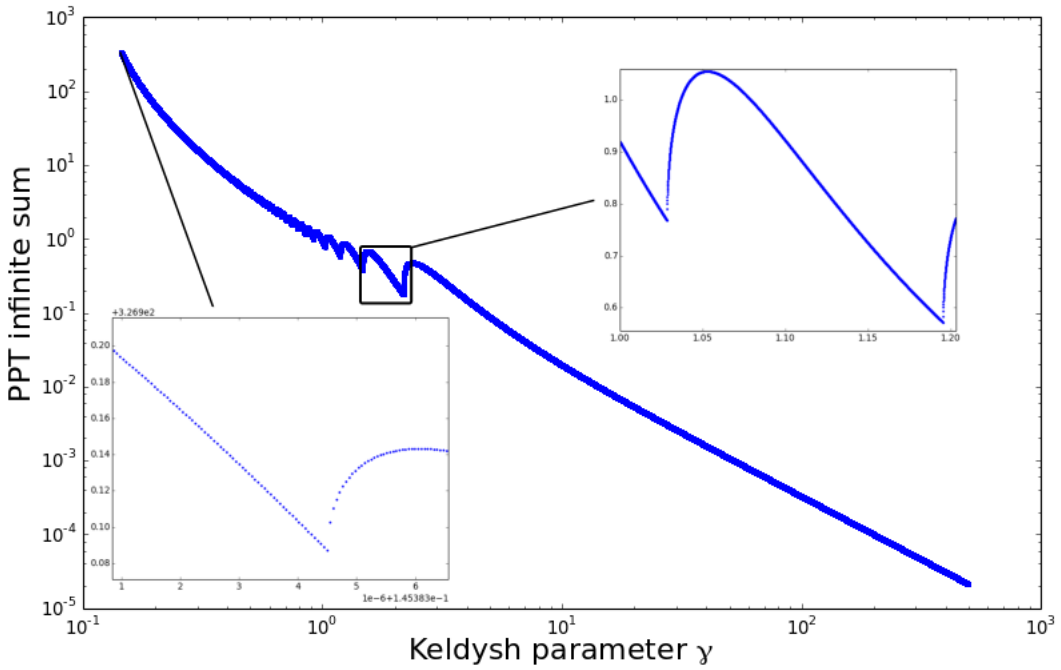

Figure 3.4: Entries within an S table produced for an 800 nm pulse with an intensity of 5x10^{13} W/cm^2 ionizing air. From the zoomed insets individual entries can be discerned.

From figure 3.4 it is clear that the sampling rate is more than sufficient to completely resolve all the features of the sum and as such the interpolations performed between points are most certainly accurate. As the S table has been shown to be reliable we can be confident in the rates computed by our ionization routine.

3.4 Static Plasma Modeling

The static plasma background was intended to mimic the simplified model of plasma used in filamentation studies up until now. In such a model the plasma is approximated to first order with the Drude model and the plasma density is calculated using the ionization rate. To incorporate such a model into EPOCH we implemented a plasma background in a manner similar to the nonlinear air background.

3.4.1 Drude Plasma Background

The Drude model was first proposed in 1900 as a model for the motion of free electrons in a conducting medium but it also serves as a first order approximation of a plasma. In a Drude-like plasma the positive ions are assumed stationary as their motions are many orders of magnitude slower than the responsive and fast moving electrons. Thus the electric susceptibility of the model is simply as shown in equation 2.7.

With an expression for the electric susceptibility we can immediately find an expression for the induced polarization in the plasma as a function of frequency and plasma density. By performing a derivation very similar to what was done for the air background in [1] we arrive at a second order ODE describing the evolution of the plasma polarization in time.

$$\mathbf{P}_p(\omega) = \varepsilon_o \chi_p \mathbf{E}(\omega) \tag{3.13}$$

$$\omega^2 \mathbf{P}_p(\omega) = \frac{n_e q_e^2}{m_e} \mathbf{E}(\omega)$$

$$-\frac{d^2 \mathbf{P}_p(t)}{dt^2} = \frac{n_e q_e^2}{m_e} \mathbf{E}(t) \tag{3.14}$$

The final component needed to complete the description of the plasma background is a means for the plasma density to evolve. We do this by modeling the evolution of the plasma density as a Poisson distribution as described in [10]:

$$n_e(\mathbf{r}, t) = n_o \left(1 - e^{\int_0^t w(E(t))dt}\right) \tag{3.15}$$

Where $w(E(t))$ is the ionization rate and n_o is the number density of the material being ionized.

3.4.2 Implementation of Static Plasma Background

By discretizing 3.14 we arrive at a pair of equations which describe the plasma polarization and plasma current.

$$\begin{aligned} \mathbf{J}_p^{n+1/2,i,j,k} &= \mathbf{J}_p^{n-1/2,i,j,k} + \frac{n_e q_e^2}{m_e} \mathbf{E}^{n,i,j,k} \Delta t \\ \mathbf{P}_p^{n,i,j,k} &= \mathbf{P}_p^{n-1,i,j,k} - \mathbf{J}_p^{n+1/2,i,j,k} \Delta t \end{aligned} \tag{3.16}$$

Where n is the time index and i, j, k are spacial indicies.

To track the population density at each grid point we need to discretize equation 3.15 which is is simply written as:

$$n_e^{n,i,j,k}(\mathbf{r}) = n_o \left(1 - e^{-\sum_0^{n,i,j,k} w(E^{n,i,j,k})\Delta t}\right) \tag{3.17}$$

As the time integral is in an exponential it is not helpful to store the plasma population directly; far easier instead to store $\sum_0^n w(E^{n,i,j,k})\Delta t$ as the new plasma density can be computed with:

$$n_e^{n+1,i,j,k} = n_o \left(1 - e^{-\left(\sum_0^n w(E^{n,i,j,k}) + w\left(E^{n+1,i,j,k}\right)\right)\Delta t}\right) \tag{3.18}$$

The process of updating the plasma background during each time step is quite simple. First the PPT ionization rate is calculated for each grid point whose $\gamma < 500$. This ionization rate is then added to the integrated ionization rate and the plasma population density is then found. With the density in hand the plasma current and polarization are then leap-frog integrated and the plasma current is fed into 3.1.

To account for energy lost by the pulse to the generation of the new plasma we add a correction current to the global current density:

$$\mathbf{J}_c^{n,i,j,k} = \frac{U_I \mathbf{E}^{n,i,j,k}}{\Delta t \left|E^{n,i,j,k}\right|^2} \left(n_e^{n+1,i,j,k} - n_e^{n,i,j,k}\right) \tag{3.19}$$

By making this correction to the current the electrical field is weakened by an amount of energy equivalent to the amount needed to liberate the electrons freed in that time step.

3.4.3 Validation

To test the plasma background we show that the defocusing it produces in a pulse is as we expect. According to Lawrence-Douglas' thesis [7] the divergence of beam in plasma can be approximated by equation 3.20. As the density of the plasma background is calculated using equation 3.17 the divergence angle of the pulse should scale linearly with the density of the ionizing medium.

$$\theta_I \simeq \frac{1}{\sigma_s} \int \frac{q_e^2 n_e\left(z\right)}{m_e \varepsilon_o \omega^2} dz \tag{3.20}$$

Where σ_s is the spot size of the beam.

To test this we ran six simulations where the density of the ionizing medium was varied. In each simulation a 25 fs long pulse with a wavelength of 800 nm, an intensity of 1x10^{14} W/cm^2, and a waist of 50 μm was propagated for 200 fs. At the end of its propagation the fluence of the pulse was calculated and the waist of this fluence is measured. With the fluence waist we can then calculate the divergence angle as $\theta_I = arctan\left(\frac{w_f}{ct_p}\right)$, where t_p is the propagation time and w_f is the waist of the fluence.

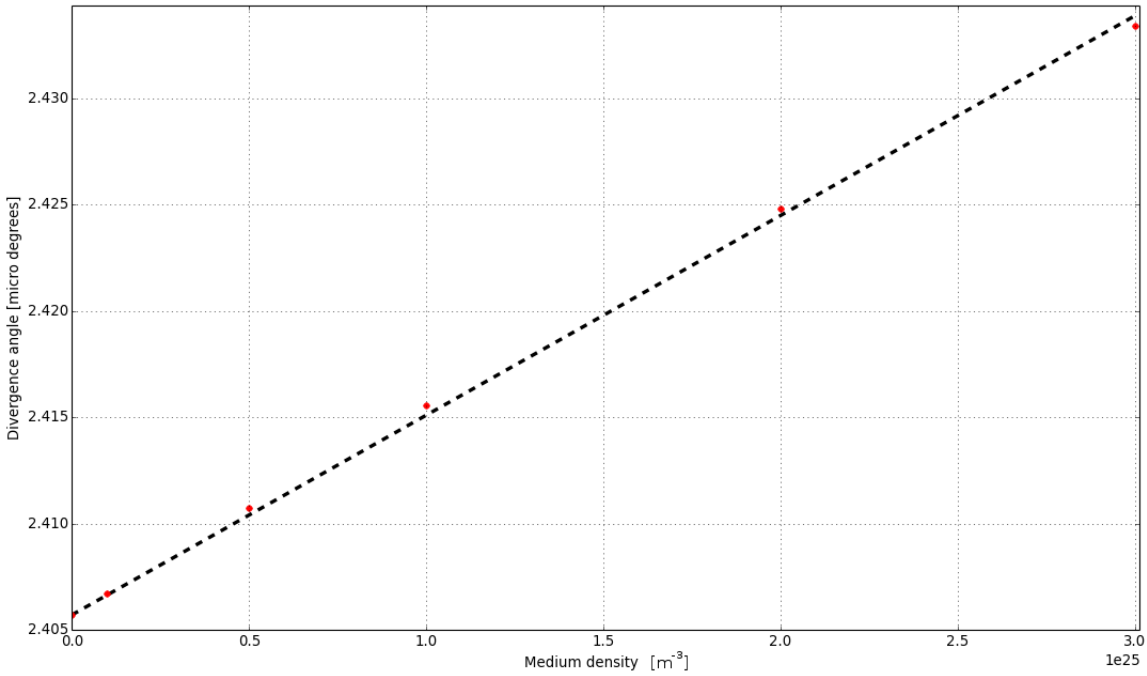

Figure 3.5: Divergence of a 25 fs, 800 nm, pulse as a function of ionizing medium density. On this plot we have the measured values in red and a fit line in black.

From the above figure we can see that the divergence angle of the plasma does indeed vary linearly with the density and as such we can be confident that the plasma background is working as intended.

3.5 Dynamic Plasma Modeling

The dynamic plasma background, in contrast to the static, is intended to be as complete as possible. The most complete way to model a plasma would be to simulate every individual particle and solve for their individual motions using Newton's laws and the Lorentz force; this would be equivalent with solving Liouville's equation for the classical plasma. Such a thing might be feasible for small simulation spaces or a very rarefied medium but our simulations are neither. For such plasmas a similar method may be used but a critical assumption must be made. When the particle population of a simulation is very high, super-particles can be made out of statistically significant groups of particles as their individual motions are very similar to each other. The motion of these super particles is then computed as though they were single particles with the combined mass and charge of their constituents. This is the Particle-In-Cell (PIC) method. PIC is a test particle solver for the Vlasov equation, which is essentially the Liouville equation without collisions.

3.5.1 Plasma Generation in the PIC Method

The ionization routine implemented into EPOCH is rather complicated as the collective nature of super-particles allows them to exist in partially ionized states. Rather than give a complete breakdown of the internal operations of the routine it will suffice to describe the general operation of the routine and the

changes which were made to it. In the dynamic plasma model we initialize the simulation by populating it with static, ionizable, neutral, super-particles with specified ionization potentials. During each time step the ionization routine runs through the list of these particles to determine whether any have ionized. A simplified flow chart of the process is shown in figure 3.6.

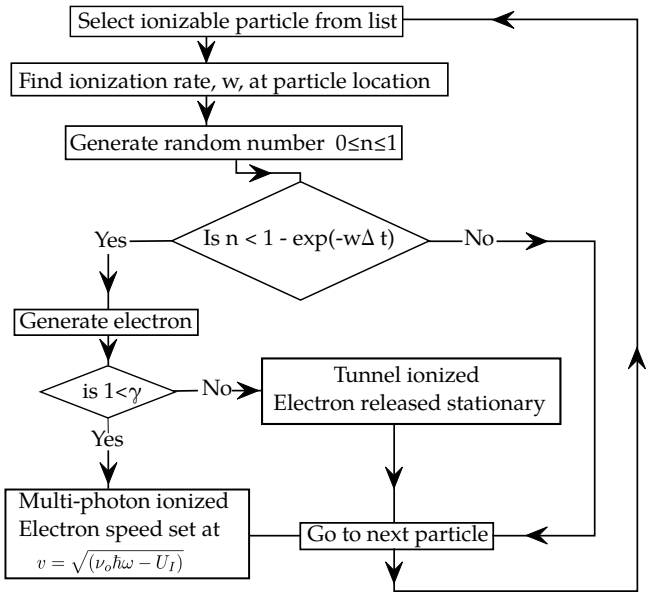

Figure 3.6: Simplified PPT ionization routine implemented in EPOCH.

In the event that a neutral particle does ionize it is converted into an ion and electron. The electrons are then given an additional boost in velocity in the direction of the electric field if they were freed by MPI. For our simulations the ions were assumed to be stationary for the same reason as in the static model of plasma; they are far too massive to move appreciably in the femtoseconds they can interact with the filamenting laser pulse. The final step of the ionization routine is to account for the energy lost by the laser pulse to plasma generation. In the dynamic background this is done by feeding a current density correction into the global current density as described in [7]:

$$\mathbf{J}_c^{n,i,j,k} = \frac{N\epsilon_t \mathbf{E}^{n,i,j,k}}{\Delta t \left|E^{n,i,j,k}\right|^2} \tag{3.21}$$

Where N is the particle weight at a grid point, ϵ_t is the energy used to ionize the super-particle. As the super particle represents many individuals this is some multiple of the ionization energy.

Once particles have been ionized the newly freed electrons are handed over to the PIC routine to solve for their motion through the simulation space.

3.5.2 Particle-in-Cell Method

In PIC, Newton's laws and the Lorentz force are combined and discretized to produce an explicitly integrable set of equations for the velocities and positions of the particles being simulated.

$$\begin{aligned} \frac{d\mathbf{v}}{dt} &= \frac{q\left(\mathbf{E}+\mathbf{v}\times\mathbf{B}\right)}{m} \\ \mathbf{v}^{n+1/2} &= \Delta t\frac{q\left(\mathbf{E}^{n}+\mathbf{v}^{n-1/2}\times\mathbf{B}^{n-1/2}\right)}{m}+\mathbf{v}^{n-1/2} \\ \mathbf{r}^{n+1} &= \Delta t\mathbf{v}^{n+1/2}+\mathbf{r}^{n} \end{aligned} \tag{3.22}$$

When solving for the motions of electrons in an FDTD simulation some effort is required to find the fields at the locations of the electrons. As the electrons are not bound to the spatial grid the value of the fields at their locations is not known; they must be found by performing an average of the nearest grid points weighted by the electrons' proximity to these grid points. Weighting the average is a somewhat tricky operation as it can be difficult to determine how many grid points to use and the weight they should carry. Suffice it to say that what is done in EPOCH is to apply a weighting curve which helps to smooth out the discontinuities in the electric field felt by an electron as it moves between grid points. Once the fields are known the equations are leap-frogged through time, achieving second order accuracy.

The final step of the routine is to account for the currents produced by the motion of the electrons. To do this the current density is found at every grid point by using the relation $\mathbf{J} = n_e q\mathbf{v}$ where n_e is the electron density at the grid point. Since the electron density isn't known at the grid points it is calculated in the same manner as the fields at the electrons' positions but in reverse. Once the average electron density is known the current can be added to the global current density of equation 3.1 and thus the laser pulse is influenced by the dynamic plasma.

It should be mentioned that while the particles in PIC can freely interact through coulomb forces they can not collide with each other. Although EPOCH supports particle collisions we do not enable the feature as the time-scale over which collisions are probable is much larger than the laser-plasma interaction time. Thus including collisions only serves to increase the computational load of the dynamic plasma.

3.5.3 Validation

Because the ionization of particles is randomly determined there must be a sufficient number of ionizable particles per cell to prevent unphysical gaps in the plasma density. The number of particles per cell we use in our simulations was chosen based on the stimulations of plasma defocusing which were run while validating the original ionization routine in [7]. In these tests 32 particles per cell (p/c) were used and the spatial resolution in the direction of propagation was set to be approximately 16 cells per wavelength (c/λ). In our simulations our spatial resolution must be considerably higher with a minimum of $\sim$37 c/λ; the higher spatial resolution lets us reduce the particle density down to 10 p/c without significantly impacting the smoothness of the generated plasma density.

To show that the combination of low particle count and high spatial resolution is valid we ran a pair of simulations of an ionizing laser pulse traversing a simulation space. In both simulations a 25 fs laser pulse with a waist of 200 μm and a wavelength of 800 nm was propagated for 1 ps through ionizable particles with

an ionizing potential of 12.63 eV, an odd ionization potential used for reasons explained in section 4.1. In one simulation the resolution in the direction of propagation is 20 c/λ with 30 p/c and the other with 37 c/λ and 10 p/c

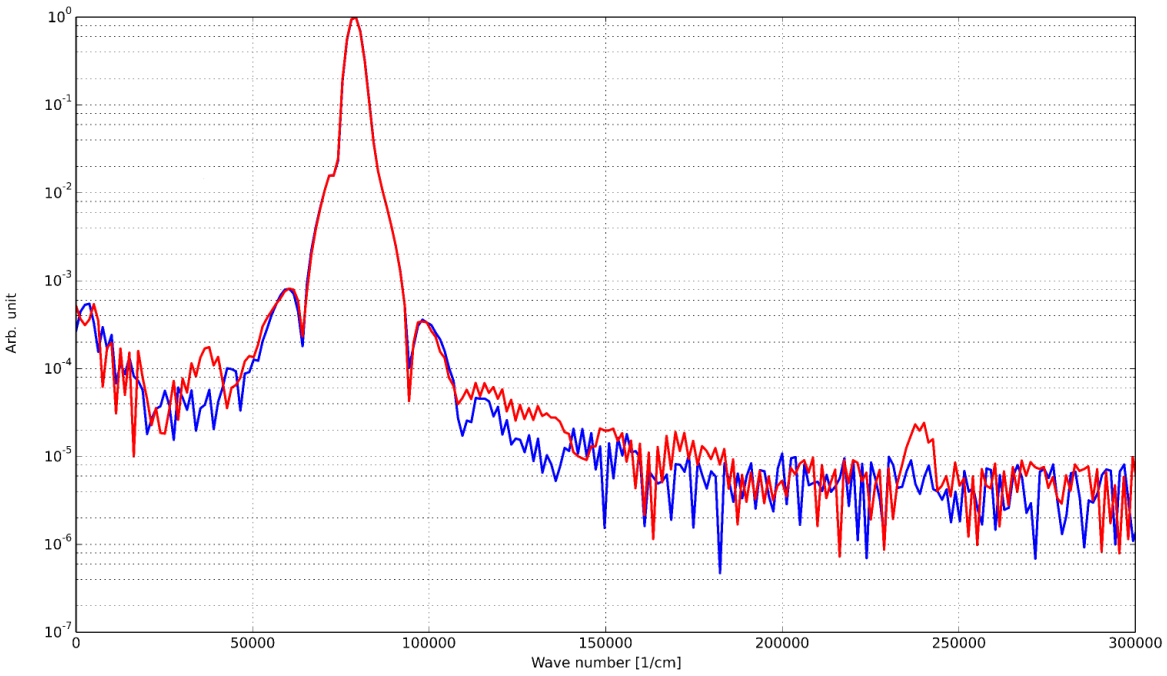

Figure 3.7: Axial spectrum of free particle density, normalized to the peak spectral component. In red is the 10 p/c with 37 c/λ. In blue is the 30 p/c with 20 c/λ.

When observing the on-axis spectrum of the two simulations we can see that the central frequency of both simulations is practically identical. The most telling sign that the particle densities/spatial resolution trade off is equivalent is shown in the noise generated by the free electrons in the simulations. As the noise level is practically equivalent in both simulations we know the plasma densities must be equivalently smooth. One point where the two curves noticeably disagree appears in the high resolution run near the third harmonic of the fundamental; this peak is due to harmonics being generated in the plasma background. Since the two simulations are otherwise in good agreement with each other we can safely conclude that our choice of 10 p/c is an acceptable particle count for our simulations.

4 Simulations and Analysis

4.1 Simulation Setup

4.1.1 Simplifications Made to Simulations

In order to make our simulations more tractable we needed to first make some simplifications to their parameters. The greatest simplification we make is to reduce the dimensionality of our simulations from three to two as a full 3D simulation is, computationally speaking, extremely challenging. Dropping the third dimension does have some consequences on the filamentation process; as there is only one dimension which is perpendicular to a pulse's propagation vector the self-focusing/defocusing effects are weakened. Consider again the paraxial model of equation 3.6, as there is only one transverse dimension and the laser pulse is radially symmetric the $\nabla_{\perp}^{2}$ term is effectively reduced by half, this then reduces the strength of the self-focusing effect. For a similar reason the plasma defocusing is also weakened.

The second major simplification was to model the ionizing components of air as one, singly ionizable, species. The reason for this is that the primary source of ions in filaments in air is singly ionized oxygen. [27] Furthermore, since the ions in our simulations are held fixed, the masses of different ions is irrelevant when considering plasma dynamics. As the plasma dynamics are unaffected and the plasma density is not significantly impacted we model the ionizable medium of both plasma models as being comprised solely of "oxygen" with a number density of $5\text{x}10^{24}$ m^{-3}, which is the number density for the oxygen content of atmospheric air. It should be noted that the ionizing medium is listed in quotes as "oxygen", this is because we have artificially raised the ionization potential from 12.063 eV to 12.63 eV; this was done to slightly reduce the generated plasma density so that it is in closer agreement with that measured in experiments.

The final major simplification we made has to do with the initial configuration of our laser pulses. Looking back at figure 2.3 we can see that a filament forming pulse may spend many picoseconds being self-focused down to the point at which its intensity surpasses the ionization threshold. During this time no plasma is formed and as such it holds no interest for us. Additionally, since simulating even one picosecond is not a trivial exercise, simulating the initial self-focusing is a significant waste of computational resources. In our simulations we choose to initialize our pulses at their waist with powers $\sim$10 P_{cr} and waist sizes comparable to those of filaments in air.[9] By doing this we skip over the initial self-focusing of the pulse, as the initial intensities are already sufficient to ionize the medium the pulse immediately begins to filament. The effect of this is that our pulses are not affected by the spacial/spectral reshaping which occurs during the self-focusing stage of filamentation; thus our results will not be directly comparable to most experimental and numerical studies.

4.1.2 Data Analysis

During the filamentation process the amplitude envelope, which defines the distribution of the laser pulse in space, undergoes a radical transformation. As previously mentioned, the pulse envelope is not an explicitly known quantity and as such we must recover it from our simulations. By recovering the amplitude envelope we can then image the intensity distribution of the pulse in space.

The deformation of the pulse envelope is, of course, not the complete picture. As a pulse filaments the frequency content of different parts of the pulse can vary substantially. For this reason we plot the spectral

distribution of the pulse. The Fourier transform is calculated in the direction of propagation and then plotted against the transverse spatial dimension.

The final measurement we make is plotting the density distribution of the free electrons generated by the pulse. For the static plasma model this is as simple as plotting the result of equation 3.15. Plotting the density of the dynamic plasma is more complicated as the random fluctuations in the density give it a spotty appearance. To help smooth out these fluctuations, and to make the plot more readable, we box-car average across five transverse cells and eleven axial cells. This does not completely erase the fluctuations but it does reduce their extremity.

By combining these measurements we can gain an understanding of how the pulse has been redistributed through space, how its spectral content has been altered, and where the plasma is being generated. When comparing such plots for simulations using the static and dynamic models of plasma we can search for features which have developed differently, or not developed at all, as a result of using the different models. In this way we can determine the impact plasma dynamics has on the development of the pulse.

4.2 Testing Pulse Duration

It was our expectation that the two plasma models would produce nearly identical results for very short pulses but begin to diverge from each other as pulse duration was increased. One reason for this thought was that the plasma generated by longer pulses is denser by virtue of the increased time over which ionization can occur; plasma forces should be stronger for a denser plasma. The second reason was that the effects of plasma dynamics become more pronounced over longer laser-plasma interaction times.

To test this hypothesis we ran three pairs of simulations; each pair consisted of simulations that were identical in all respects save that one used a static plasma and the other a dynamic plasma. The majority of the parameters defining all six simulations were the same, only the duration and propagation time of the pulse were varied. In table 4.1 the parameters common to all simulations are listed, in table 4.2 are the parameters common only to the simulation pairs.

Wavelength	Intensity	Pulse type	Waist	Pulse polarization
800 nm	$5\text{x}10^{13}$ W/cm^2	Gaussian	200 μm	$\hat{y}$
Propagation vector	Transverse resolution	Transverse size	"Oxygen" density	"Oxygen" ionization potential
$\hat{k} = \hat{x}$	900 cells	2 mm	$5\text{x}10^{24}$ m^{-3}	12.63 eV

Table 4.1: Simulation parameters common to all duration test simulations.

Duration	Axial size	Axial resolution	Propagation time
30 fs	60 μm	3400	20 ps
300 fs	600μm	28005	30 ps
500 fs	1200 μm	56010	30 ps

Table 4.2: Simulation parameters shared only by matched simulation pairs.

4.2.1 Results of the 30 fs Pulse

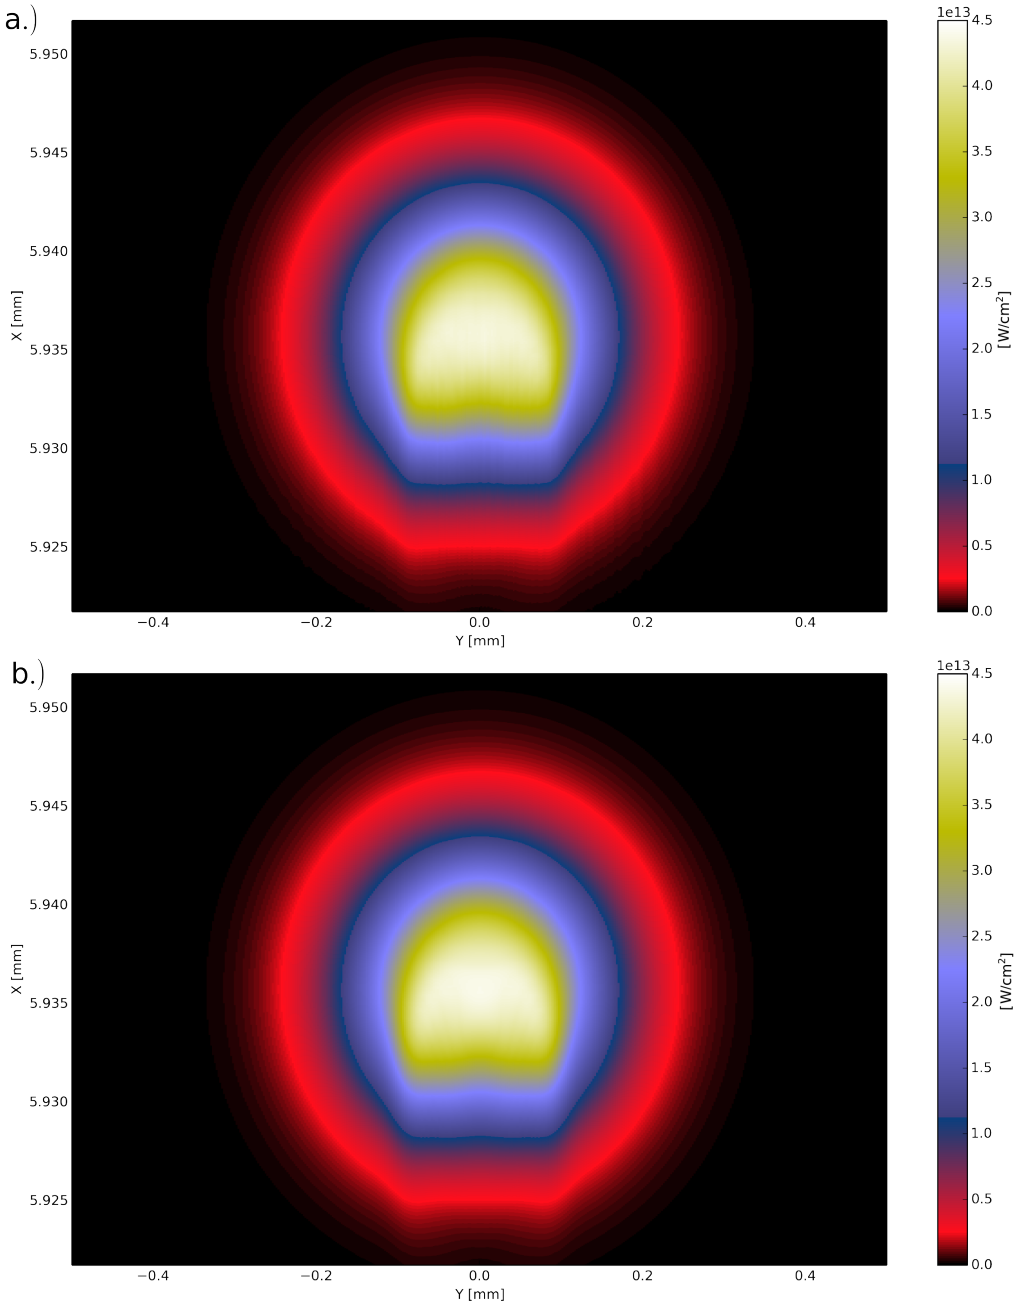

Figure 4.1: Spatial intensity distribution of a 30 fs pulses in air after 20 ps of propagation. a.) Dynamic plasma model b.) Static plasma model.

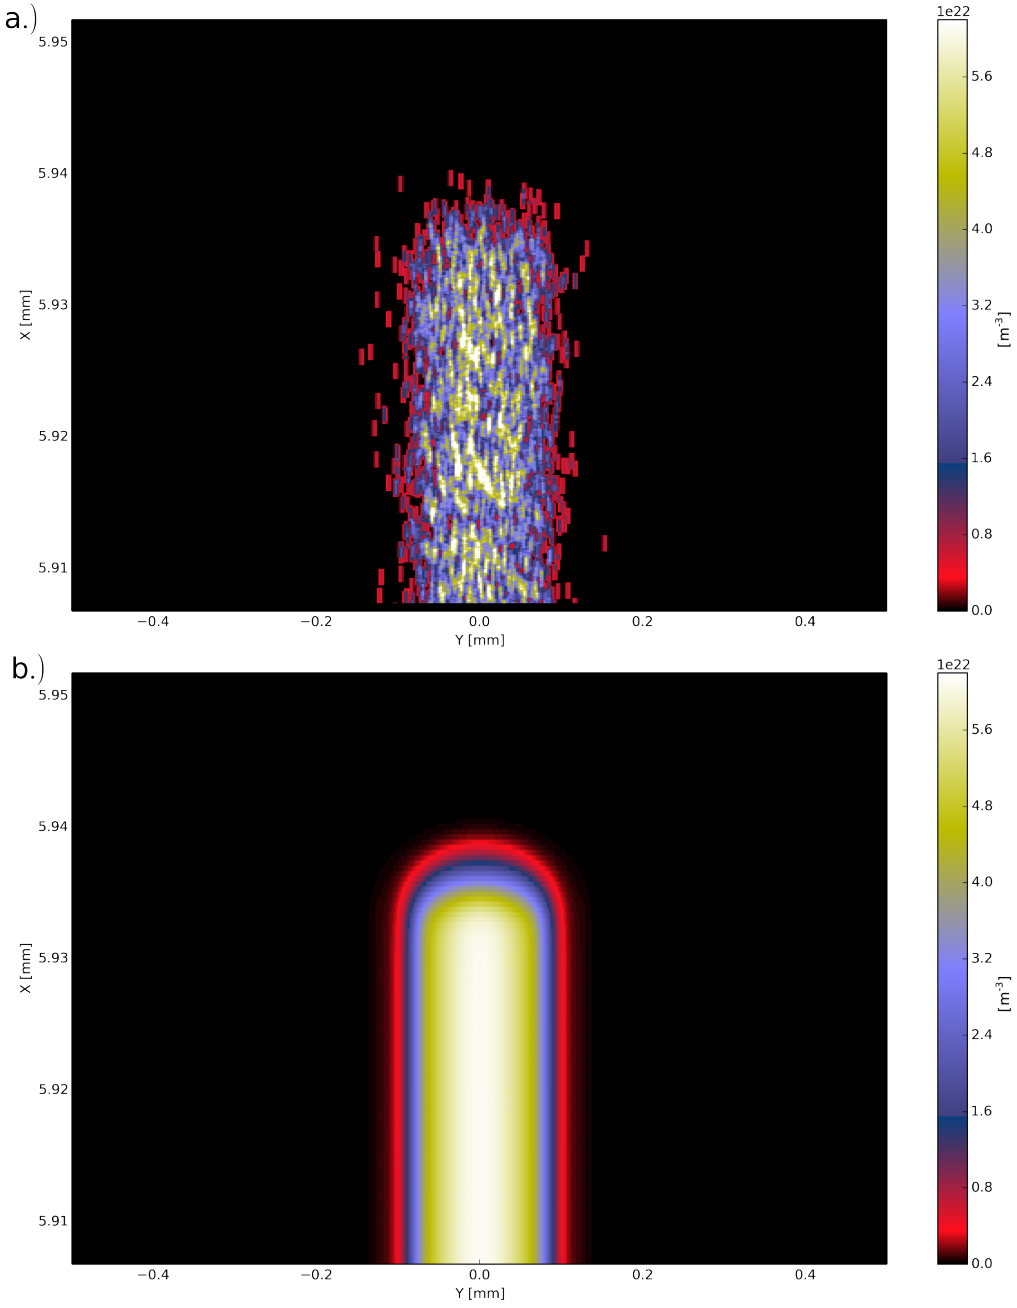

Figure 4.2: Plasma density distributions for the 30 fs pulse simulations after 20 ps of propagation in air. a.) Dynamic plasma model b.) Static plasma model.

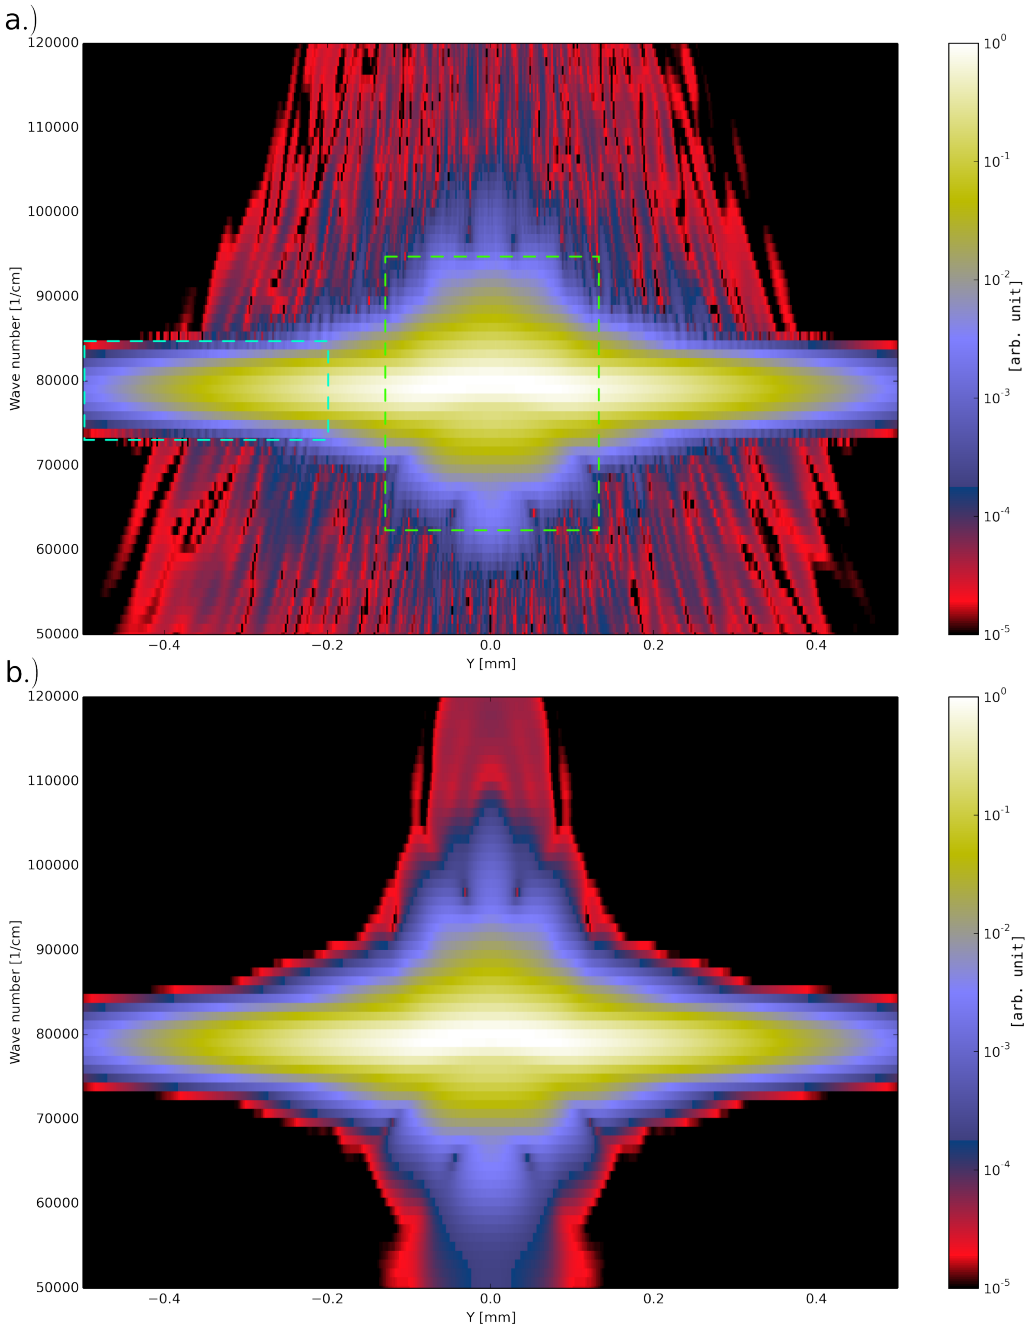

Figure 4.3: Spectral distribution as a function of the transverse spatial coordinate of the 30 fs pulse simulations after 20 ps of propagation in air. Frequency amplitude is normalized to the highest frequency component. The color bar is a logarithmic scale. Highlighted in these plots is: the original frequency content of the pulse in cyan, and the much increased spectral content at the center of the pulse in green. a.) Dynamic plasma model b.) Static plasma model.

From observing the figure 4.1 we can see that in both simulations the pulse has entered the initial stages of filamentation. As we expected there is scarcely any visible difference between the intensities of the two plots. This similarity is most reassuring as, when looking at the plasma density plots of figure 4.2, we can see that

the characteristics of of the plasma densities (smooth and continuous for the static plasma, fluctuating and blocky for the dynamic) are quite different.

Looking now to figure 4.3 we can see that there appears to be something quite interesting happening within the pulse. Stretching across the center of both plots is a broad beam of the driving frequency, approximately 80000 cm^{-1}. Looking at the left and right ends of this beam (cyan) gives us a good estimate of how broad the pulse's spectrum was initially, looking towards the center (green) we can see that the pulse has undergone substantial broadening of its spectrum; this is the onset of super-continuum generation. At the center of the pulse we can see that wavelengths between 1000 - 550 nm are present with a strength which is a significant fraction of the strength of the central wavelength. Although the spectra of the two plasma models appear to be different at first glance, a closer inspection shows them to be practically the same. One obvious difference between the static and dynamic simulations is the presence of random spots throughout the dynamic simulation spectrum. These are the result of the noise that the free electrons in the PIC routine generate as they move through the simulation space.

4.2.2 Results of the 300 fs Pulse

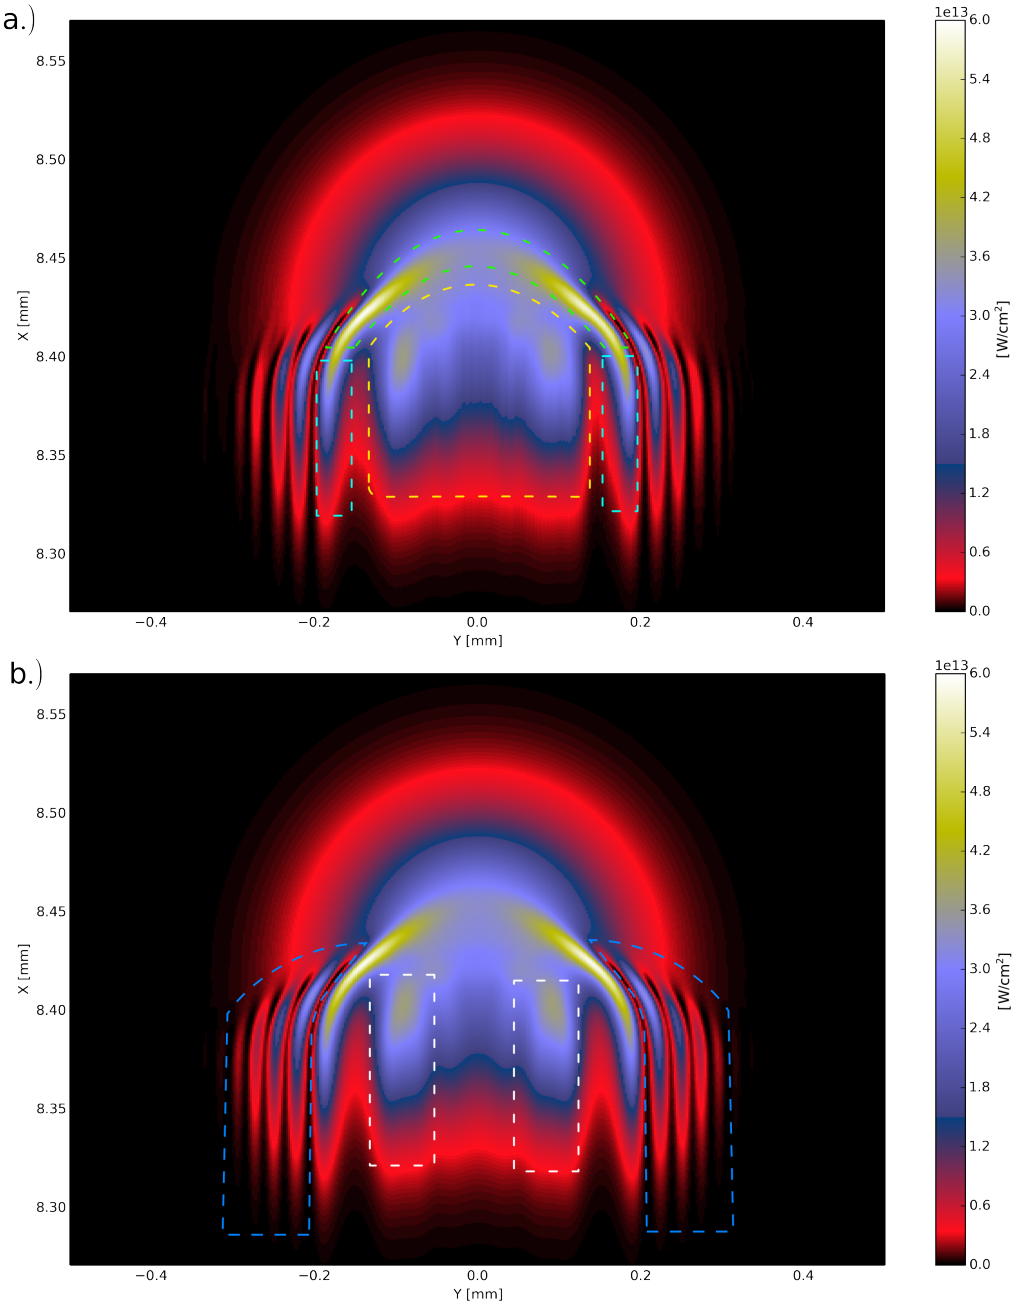

Figure 4.4: Spatial intensity distribution of a 300 fs pulses in air after 30 ps of propagation. Highlighted in these plots are: the shockfront in green, the shockfront tails in cyan, the internal structure in yellow, the outer bands of the internal structure in blue, and lastly the ringed conical emissions in white. a.) Dynamic plasma model. b.) Static plasma model.

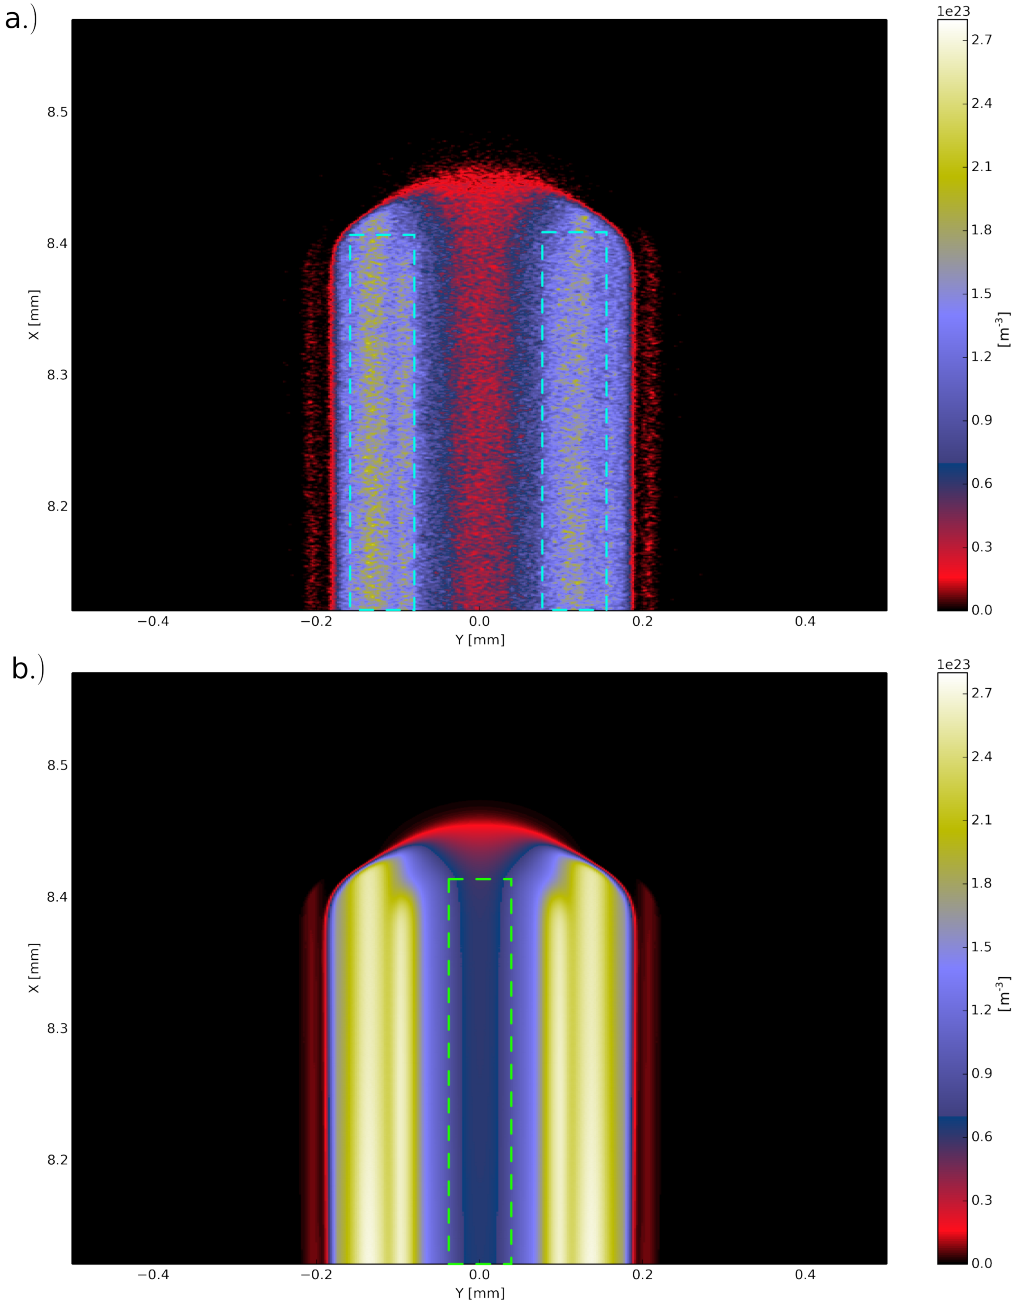

Figure 4.5: Plasma density distributions for the 300 fs pulse simulations after 30 ps of propagation in air. Highlighted in these plots are: the central region of reduced plasma density in green, and the asymmetrical sections of the dynamic plasma density in cyan. a.) Dynamic plasma model b.) Static plasma model.

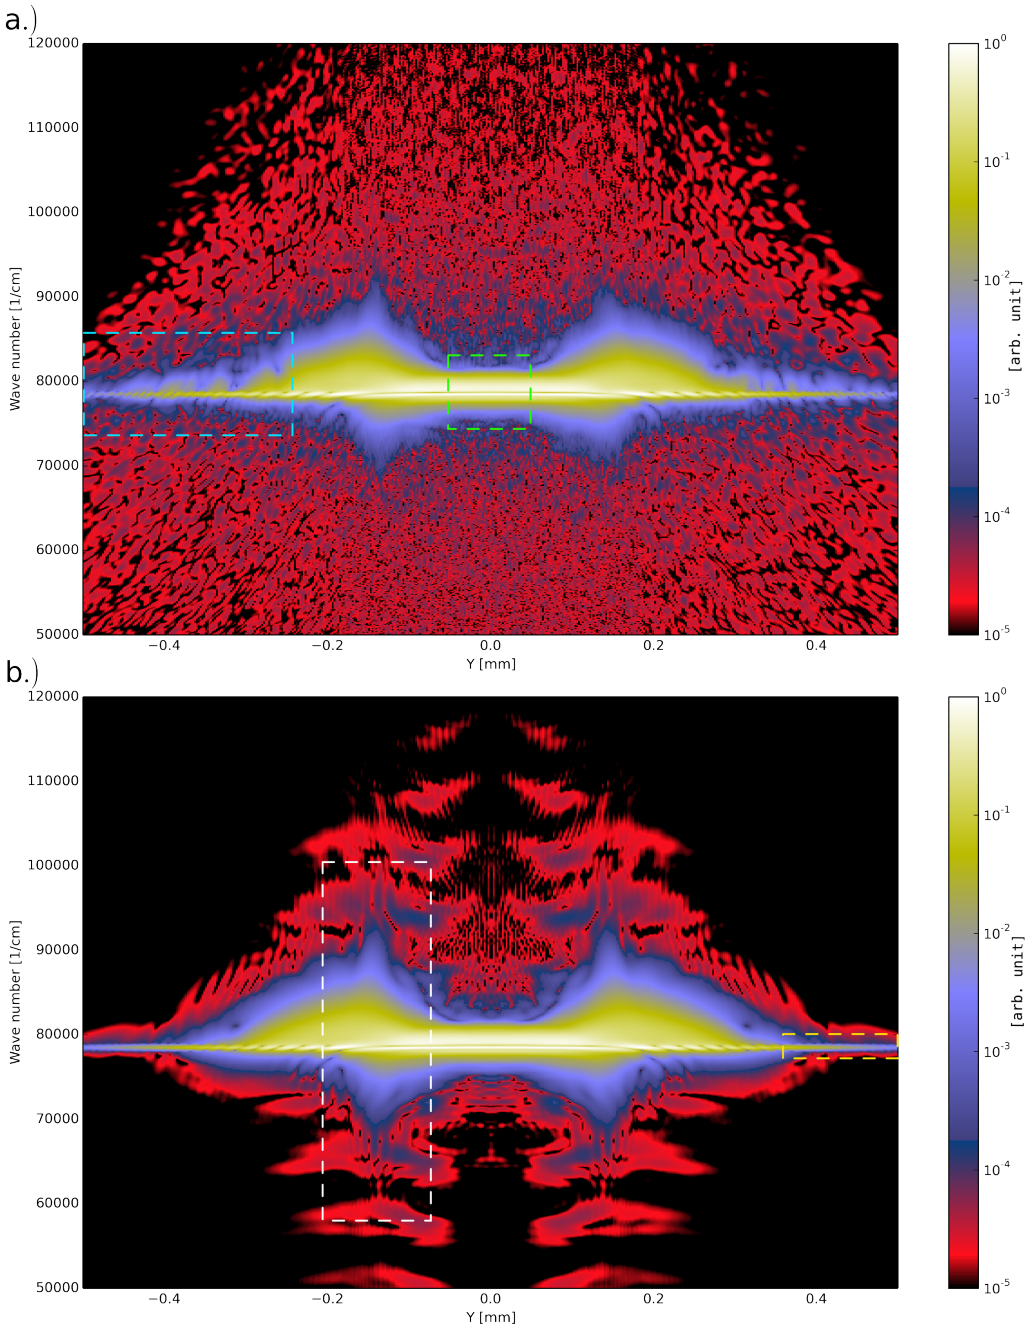

Figure 4.6: Spectral distribution of the 300 fs pulse simulations after 30 ps of propagation in air. Frequency amplitude is normalized to the highest frequency component. The color bar is a logarithmic scale. Highlighted in these plots is: the weakly broadened core of the pulse in green, the extended high frequency content of the dynamic plasma in cyan, a spectral horn in white, and the initial frequency content in yellow. a.) Dynamic plasma model b.) Static plasma model.

In figure 4.4 we can see that the 300 fs pulse has developed significantly differently than the 30 fs pulse. Firstly we can clearly see the shockfront (green) and its tails (cyan). Flanking the shockfront tails are numerous parallel bands (blue); these bands are the rings of the conical emissions given off by the filament.

Within the center of the pulse, between the tails of the shockfront, we see the internal structure (yellow) of the filament and within it the two parallel bands (yellow) whose fusion prefaces a refocusing event.

It is interesting to note that the rings of the static plasma pulse are more clearly visible than those of the dynamic plasma pulse. In fact you can clearly make out three rings in the static plasma but only two in the dynamic plasma. This would imply that the mechanism behind their generation is somehow being inhibited in the dynamic plasma model. When comparing the internal structure of the filament we can see that it is more clearly defined in the static plasma pulse than in its dynamic plasma twin; this suggests that the refocusing of the pulse is being hindered.

When looking at figure 4.5 we first notice that there is a dearth of plasma within the center of the filament (green). We can also see that the density of the dynamic plasma is asymmetrical about its axis (cyan). This is expected as the ionization of a particle is randomly determined and thus there are random fluctuations in the plasma density. At a small scale these fluctuations produce a noisiness in electric field, on a larger scale these fluctuations can cause the pulse to develop asymmetrically. This is not an unphysical result as such asymmetries are seen in experiment such as is seen figure 2.9 where you can see that the conical emission is weaker on one side of the filament.

Looking now at figure 4.6 we can see that the core of the pulse (green) has not exploded in frequency content like the 30 fs pulse simulations did. This is because the increase in pulse duration significantly reduces the initial bandwidth (yellow) of the pulse and thus the processes which generate the super-continuum are hindered. Most noticeable in these spectra are the horns (white) protruding into both the red and blue ends of the spectrum. The tips of these horns, near Y = $\pm$ 150 μm, appear to roughly align with the region of the envelope plots wherein the large internal bands overlap with the intense region of the leading shockfront. The most likely explanation for this is that the majority of these other frequencies are being generated in the shockfront and are then being enhanced further by the internal structure of the pulse.

Again we can see that the free electron noise in the dynamic plasma simulation masks weak features visible in the static plasma. Noise aside, there appear to be some small differences between the two spectra. It seems that the horns of the dynamic plasma pulse do not reach as deeply into the red or blue but the region of enhanced spectral content (blue) does stretch across a greater distance. As this region extends past the visible rings seen in the intensity distributions, and because there is no plasma within this region, this must be the colored fringes of the conical emission.

4.2.3 Results of the 500 fs Pulse

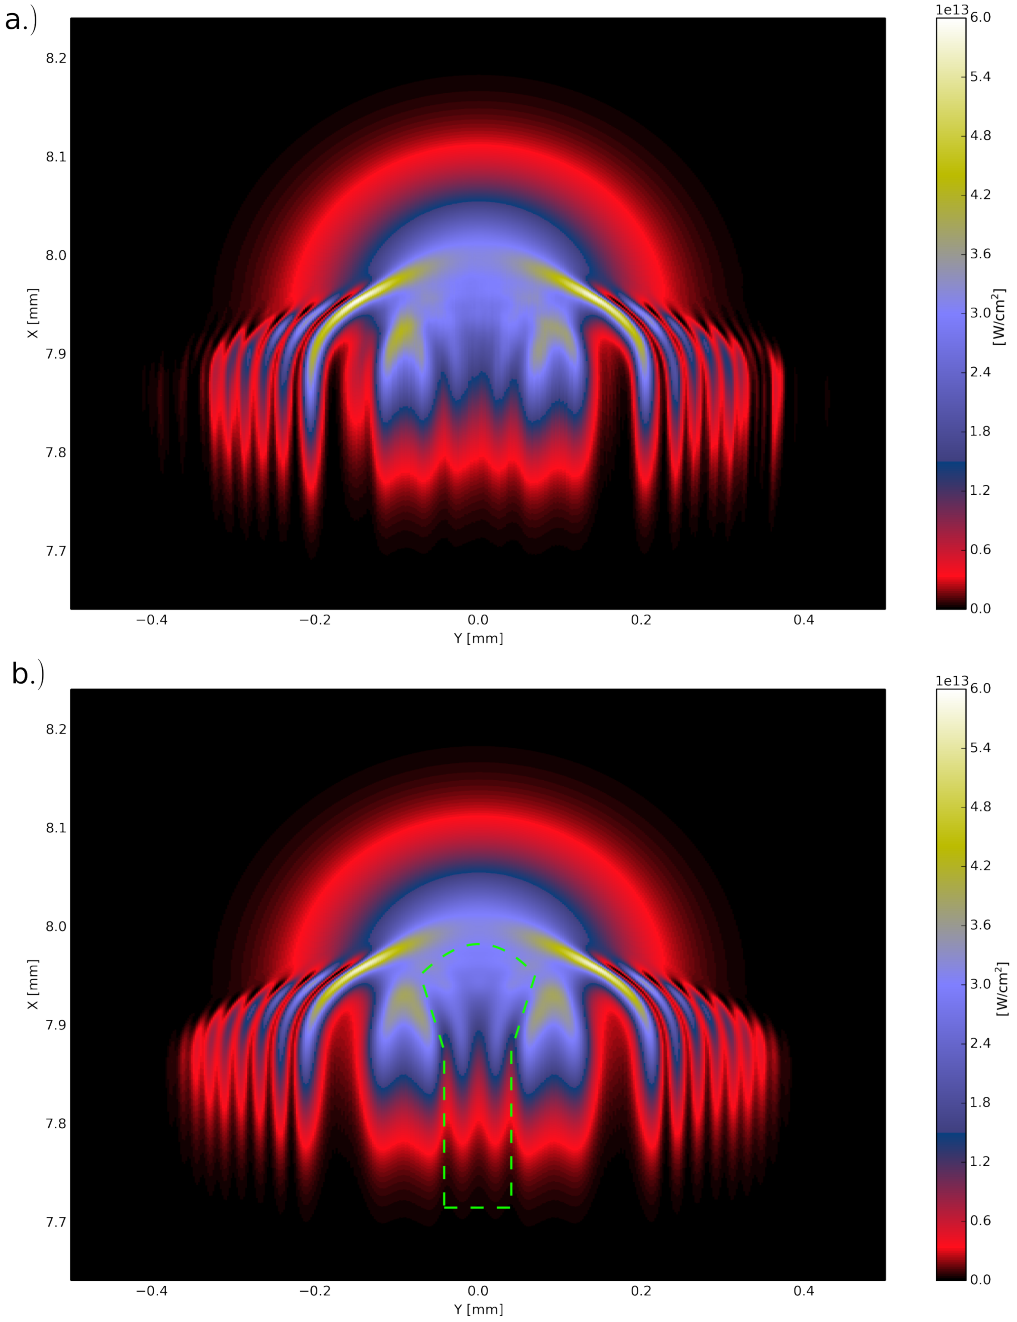

Figure 4.7: Spatial intensity distribution of a 500 fs pulses in air after 30 ps of propagation. Highlighted in green are the secondary bands of the internal structure, indicative of imminent refocusing. a.) Dynamic plasma model b.) Static plasma model.

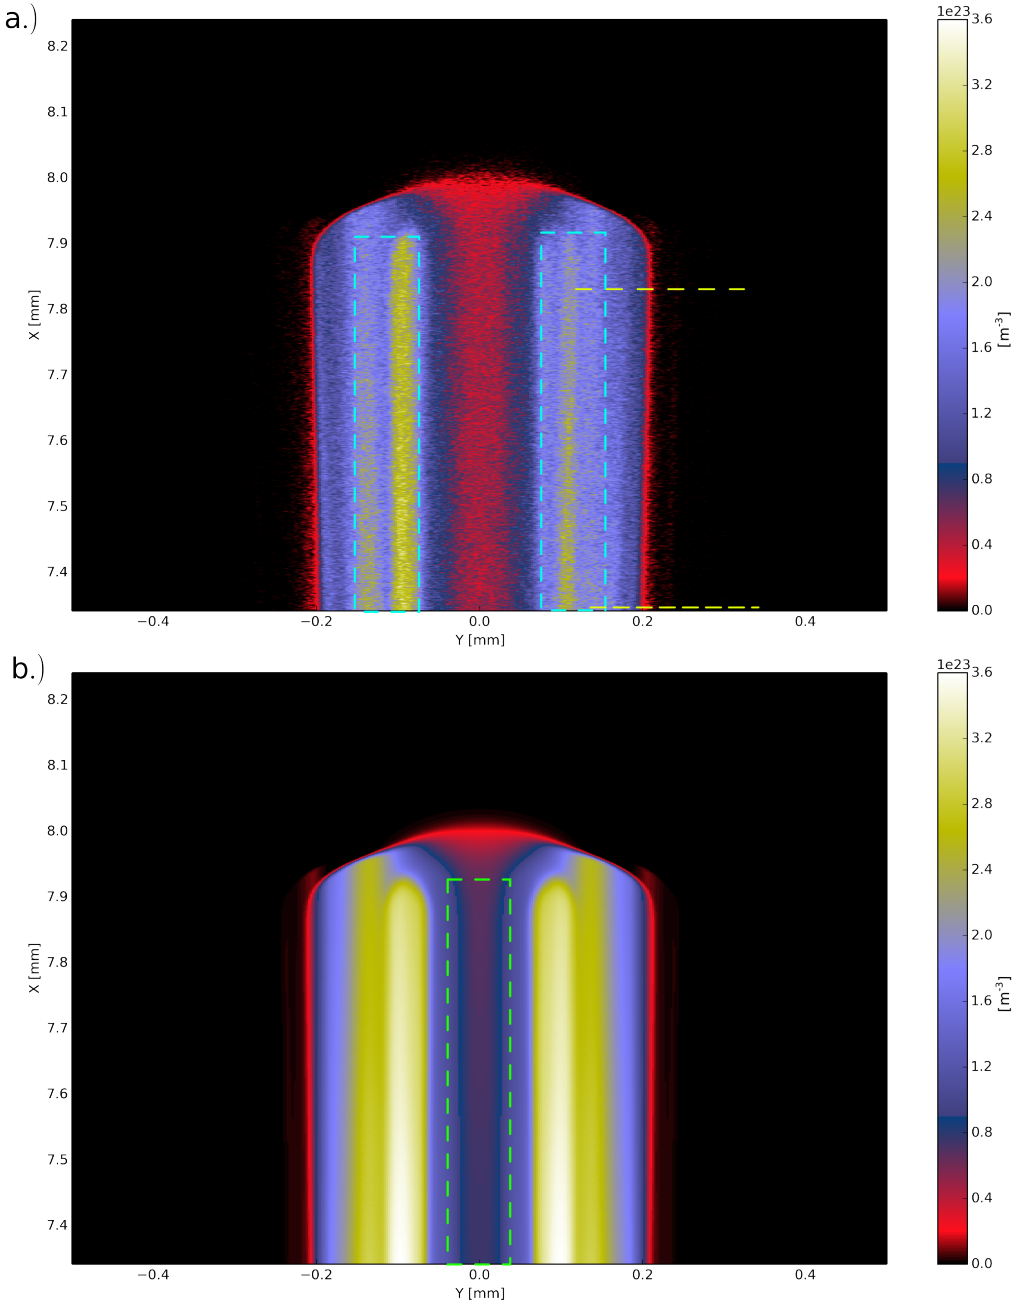

Figure 4.8: Plasma density distributions for the 500 fs pulse simulations after 30 ps of propagation in air. Highlighted in these plots are: the central region of reduced plasma density in greem, the asymmetrical sections of the dynamic plasma density in cyan, and the approximate locations of the slices taken for figure 4.10 in yellow. a.) Dynamic plasma model b.) Static plasma model.

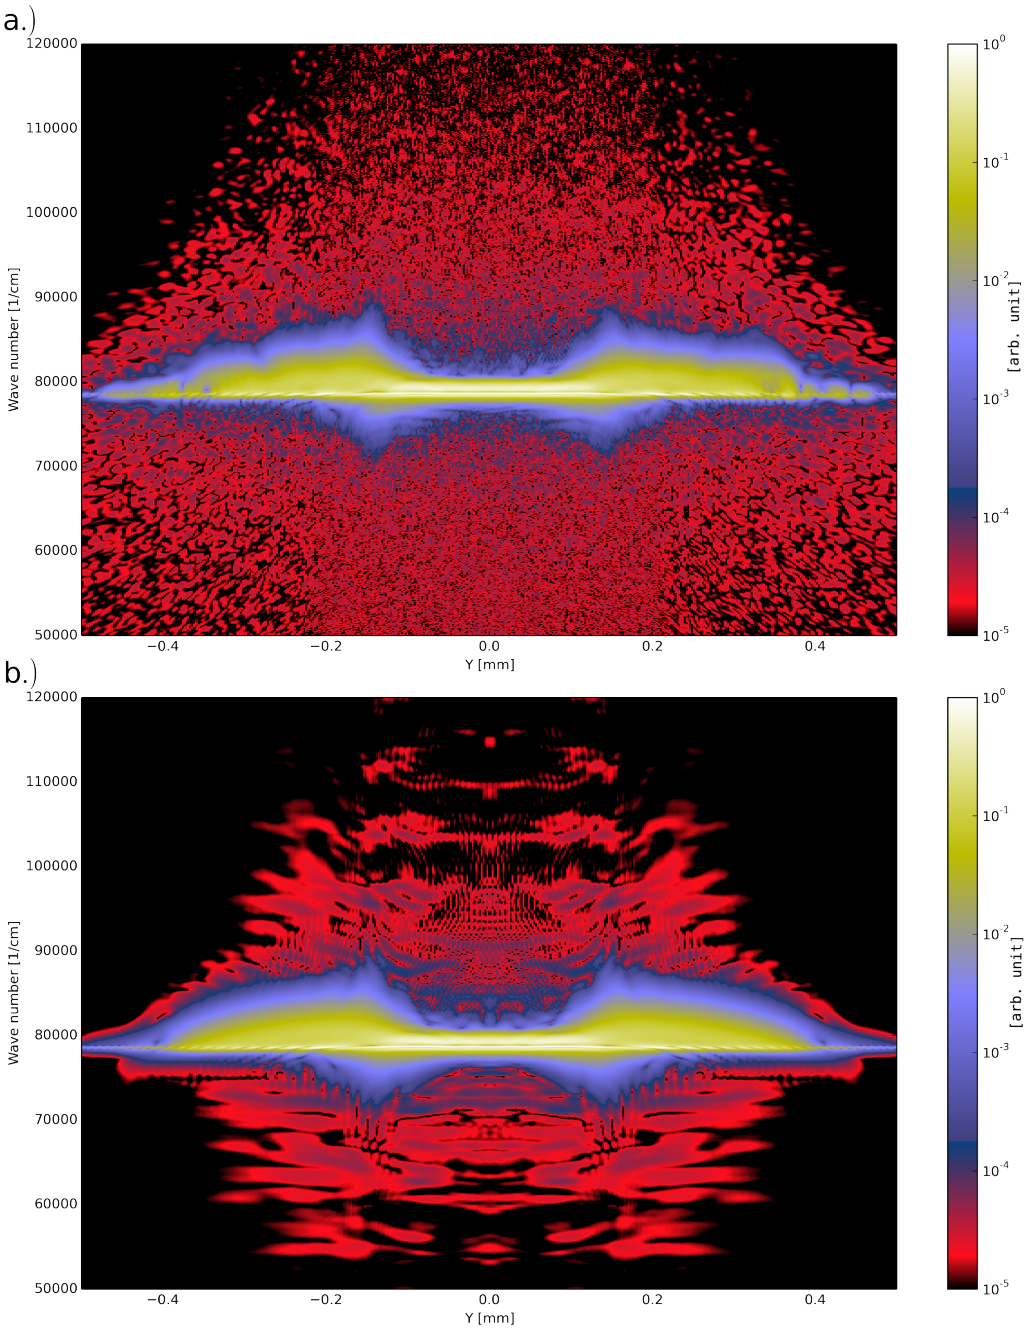

Figure 4.9: Spectral distribution of the 500 fs pulse simulations after 30 ps of propagation in air. Frequency amplitude is normalized to the highest frequency component. The color bar is a logarithmic scale. a.) Dynamic plasma model b.) Static plasma model.

The intensity distributions of figure 4.7 bear a strong resemblance, unsurprisingly, to the 300 fs simulation. There are three noticeable differences between the two: there are at least three more rings visible in the conical emissions, the shockfront is flatter, its tails are spaced farther apart, and the internal structure of the pulse is more complex; note the presence of secondary bands which have formed a funnel-like structure

(green). The reason for these changes must be due to the increased plasma generation and interaction times. The stronger plasma response forces the tails of the pulse further from the core and causes them to shed rings at a higher rate. Meanwhile the internal structure has reached a more evolved state as its growth has been fueled by the expansion of the shockfront.

Again we can see that the dynamic plasma simulation seems less evolved than its static plasma twin. Again it is missing a ring present in the static simulation. We can also see that the tails of the pulse are both narrower and more intense in the dynamic plasma; this would imply that the forces working to disperse the tails are more effective in the static plasma model. Lastly we notice that the innermost bands of the internal structure at the core of the pulse are nearer to each other in the static plasma simulation than they are in dynamic plasma while the outer bands are somewhat smaller in the dynamic model than they are in the static model. As the internal bands are closer to touching in the static plasma simulation it would seem that it is closer to a refocusing event than is the dynamic plasma is.

In figure 4.8, as in the 300 fs simulation, there is an asymmetry in the dynamic plasma density about the central axis (cyan). Again we can see that there is a central channel in which the plasma density is relatively low (green). As this channel is significantly narrower than it was in the 300 fs pulse we can conclude that the 500 fs pulse is nearer to a refocusing event; this channel will vanish when the internal structure of the filament collapses in on itself and forms a secondary shockfront.

Aside from the asymmetry in the dynamic plasma there does not appear to be any gross differences between the two density plots, there is, however, a small one. Looking at the outer edges of the dynamic plasma we notice that the width of the transverse plasma density profile is slightly tapered in the wake of the pulse. It is unlikely this narrowing is a random fluctuation in density as it appears on both sides of the pulse. As this feature is absent in the static plasma it must be a consequence of the dynamic plasma. As the plasma in the wakte of the pulse is being pushed up a charge gradient we can immediately rule out static plasma field effects. The most probable cause for the tapering is the ponderomotive force as the edges of the plasma correspond with the location of the shockfront tails. The ponderomotive force exerted by the tails of the pulse push the plasma generated by the shockfront into the center of the filament and thus the tails of the dynamic plasma pulse are not being subjected to the same defocusing plasma effects their static plasma twins feel. This seems to indicate that the ponderomotive force is active in these simulations.

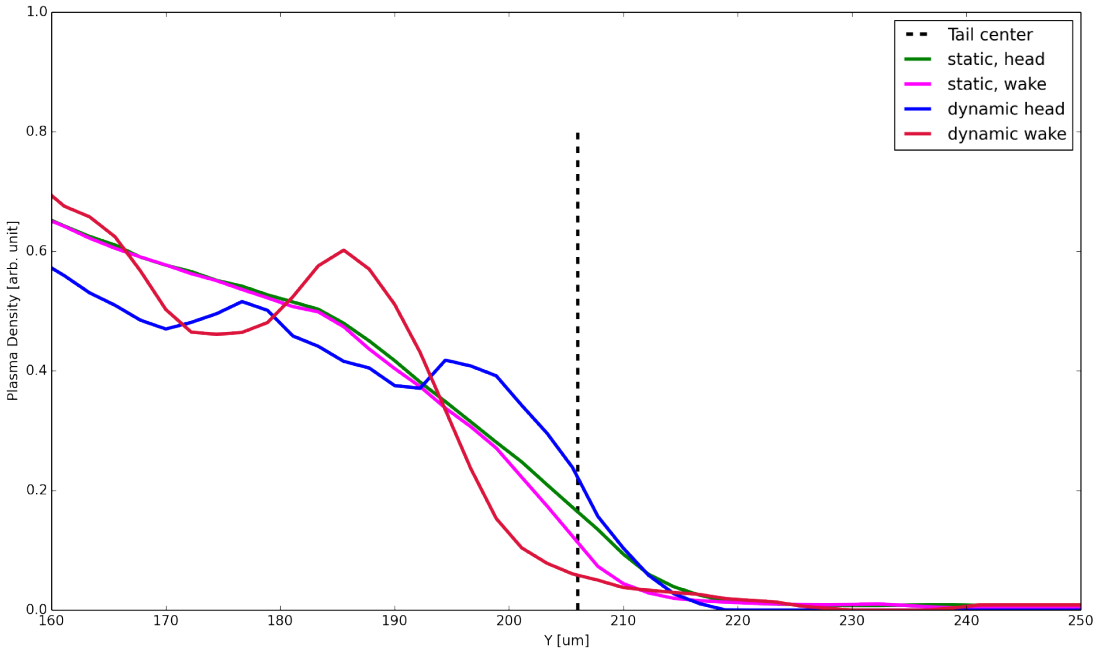

Figure 4.10: Shown here are transverse slices of the dynamic plasma density of figure 4.8 taken at x $\simeq$ 7.8 mm (blue) and 7.3 mm (red). In the inset figure the transverse slices of the static plasma density taken from x $\simeq$ 7.8 mm (green) and 7.3 mm (magenta). While the static plasma density does not change much over the 0.5 mm between slices the dynamic plasma does.

In figure 4.9 we see that the cores of the 500 fs pulses contains even fewer frequencies than does the 300 fs pulses. This is not surprising as the initial frequency content of the pulses is narrower due to the increased duration. Again we see a rapid increase in spectral content around where the intensity of the tail begins to pick up, Y = $\pm$ 100 μm, but this time we see that the frequency content is less than that found in the 300 fs pulses, this too is a result of the smaller initial pulse bandwidth.

Comparing the spectral distributions of the two plasma models we see that the spectral content of the conical emission is somewhat less colorful in the dynamic plasma. Curiously it seems that the dynamic plasma spectrum sports asymmetric lumps in the spectra of its conical emissions. This asymmetry is the result of the random nature of the plasma density causing the conical emissions to be somewhat irregular.

4.2.4 Analysis of Duration Tests

From these tests it would seem that there are some small differences between the static plasma and dynamic plasma. In both the 500 fs and 300 fs simulations the conical emissions of the static plasma simulation were larger and more intense. Furthermore the internal structures of the 500 fs and 300 fs simulations appeared to be more developed in the static simulations, though this is less apparent than the absence of rings. Both the conical emission and refocusing are powered by the leading shockfront which implies that the difference in the two plasma models must originate from differing behaviour in this region. As both the field gradient and field intensity are extremely high within the shockfront it is not unreasonable to think that, perhaps, the ponderomotive force is strong enough to push out a small number of of electrons. As even a small change

in plasma density can result in a significant change in the index of refraction it may be the case that the ponderomotive force is slightly reducing the plasma density in the shockfront and its tails and thus the index of refraction is higher and the conical emissions are weaker. Although it is not a particularly strong effect, the slight narrowing of the plasma density in the 500 fs dynamic plasma simulation suggests that this may be the case. This would also explain why the internal structure of the pulse is less strongly affected; as the center of the shockfront is less sharply defined, and the plasma is much less dense, the ponderomotive force is far weaker and thus its evolution would only be slightly altered.

To determine whether it truly is the ponderomotive force which lies at the root of these differences we needed to run some simulations in which the effect will be magnified. Looking at equation 2.9 we see that it scales with the gradient of the field intensity so, by shrinking the pulse size while keeping the same intensity, we should be able to increase this gradient and thus increase the magnitude of the ponderomotive force.

4.3 Testing Pulse Size

4.3.1 Narrowed Pulse

This pair of simulations was identical to the previous 300 fs pulse simulations except the waist of the pulses was cut down by half. We expected that doing this would increase the field gradient in the transverse direction and thereby increase the ponderomotive force. The parameters of this simulation are given in table 4.3.

Wavelength	Intensity	Polarization	Propagation vector
800 nm	5x10^{13} W/cm^2	$\hat{y}$	$\hat{k} = \hat{x}$
Pulse Type	Waist	Duration	Propagation time
Gaussian	100 μm	300 fs	30 ps
Axial size	Transverse size	Axial resolution	Transverse resolution
600 μm	1 mm	28005 cells	900 cells
"Oxygen" density		"Oxygen" ionization potential	
5x10^{24} m^{-3}		12.63 eV	

Table 4.3: Parameters of the narrowed pulse simulations.

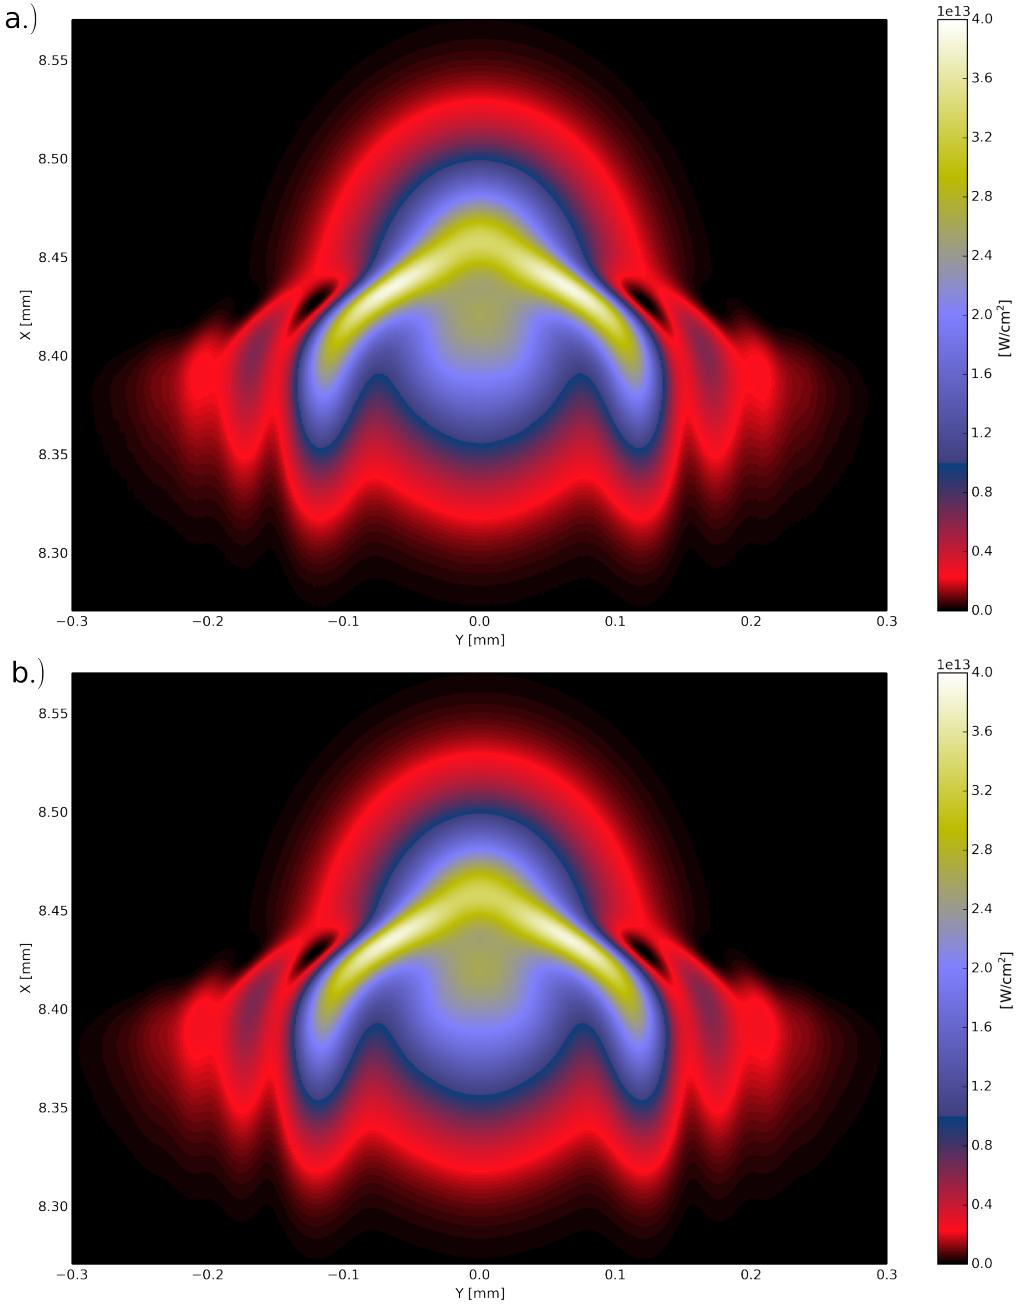

Figure 4.11: Spatial intensity distribution of the 100 μm wide 300 fs pulse simulations after 30 ps of propagation in air. a.) Dynamic plasma model b.) Static plasma model.

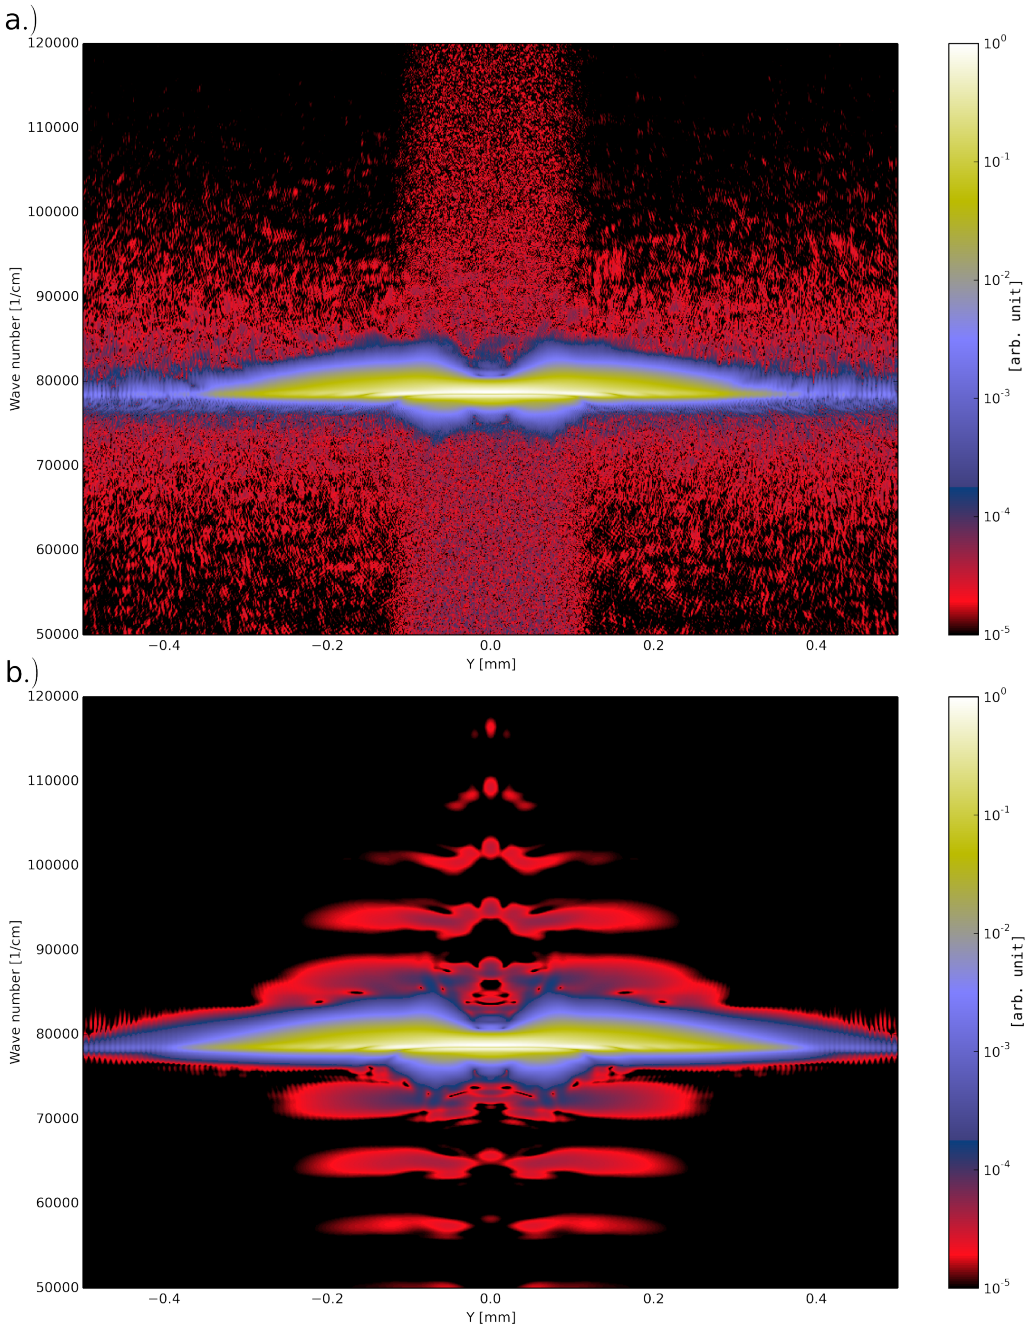

Figure 4.12: Spectral distribution of the 100 μm wide 300 fs pulse simulations after 30 ps of propagation in air. Frequency amplitude is normalized to the highest frequency component. The color bar is a logarithmic scale. a.) Dynamic plasma model b.) Static plasma model.

From the intensity distribution plots of figure 4.11 we can see that the simulations are once again largely indistinguishable. A well formed shockfront has developed and grown a pair of tails in both simulations. Also present are some weak ringed conical emissions. It is clearly evident however, that these structures are less sharply defined and far less intense than those present in the previous 300 fs pulse simulations.

Within the spectral distributions of figure 4.12 we can see that the horn-like structure we had previously

associated with the shockfront and the internal structure is present but its tips are dull and its frequency content is quite limited. As the shockfront is quite broad and the internal structure has yet to form we can conclude that the horns are primarily the product of the shockfront.

Just about the only noticeable differences in the intensity distribution is that the ring of the static plasma simulation is slightly larger and more intense than the ring of the dynamic plasma simulation. The spectral distribution plots also fail to show us any obvious differences. There is, of course, the presence of free electron noise in the dynamic plasma simulation but, aside from this, the two spectra are nearly identical. Clearly size alone was not sufficient to produce an obvious difference.

4.3.2 Intensified Narrow Pulse

For our next test we attempted to provoke an obviously different plasma response by increasing the intensity of the previous simulation. In the previous simulation, by shrinking the pulse size by half without altering the intensity, the peak power of the pulse dropped from $\sim$9.8 P_{cr} to $\sim$2.5 P_{cr} which weakened the Kerr effect substantially. Without the Kerr effect to tighten up the tails and shockfront we were not getting as steep a field gradient as we needed. To correct for this the intensity of the pulse was raised from 5x10^{13} W/cm^2 to 2x10^{14} W/cm^2, raising the power of the pulse back up to 9.8 P_{cr}. By doing this however we would completely saturate the plasma as the "oxygen" medium would completely ionize. This was undesirable, as we found out in early tests, as it produces unphysical behaviours. To prevent the plasma from saturating out we increase the density and ionization potential of the ionizing medium. This way the high intensities of the pulse generate a plasma which is only a few percent the density of the ionizing medium but still denser than the plasmas generated in previous simulations. The parameters of this simulation are given in table 4.4.

Wavelength	Intensity	Polarization	Propagation vector
800 nm	2x10^{14} W/cm^2	$\hat{y}$	$\hat{k} = \hat{x}$
Pulse Type	Waist	Duration	Propagation time
Gaussian	100 μm	300 fs	23 ps
Axial size	Transverse size	Axial resolution	Transverse resolution
600 μm	1.4 mm	28005 cells	1260 cells
"Oxygen" density		"Oxygen" ionization potential	
1x10^{25} m^{-3}		18.576 eV	

Table 4.4: Parameters of narrowed and intensified pulse simulations.

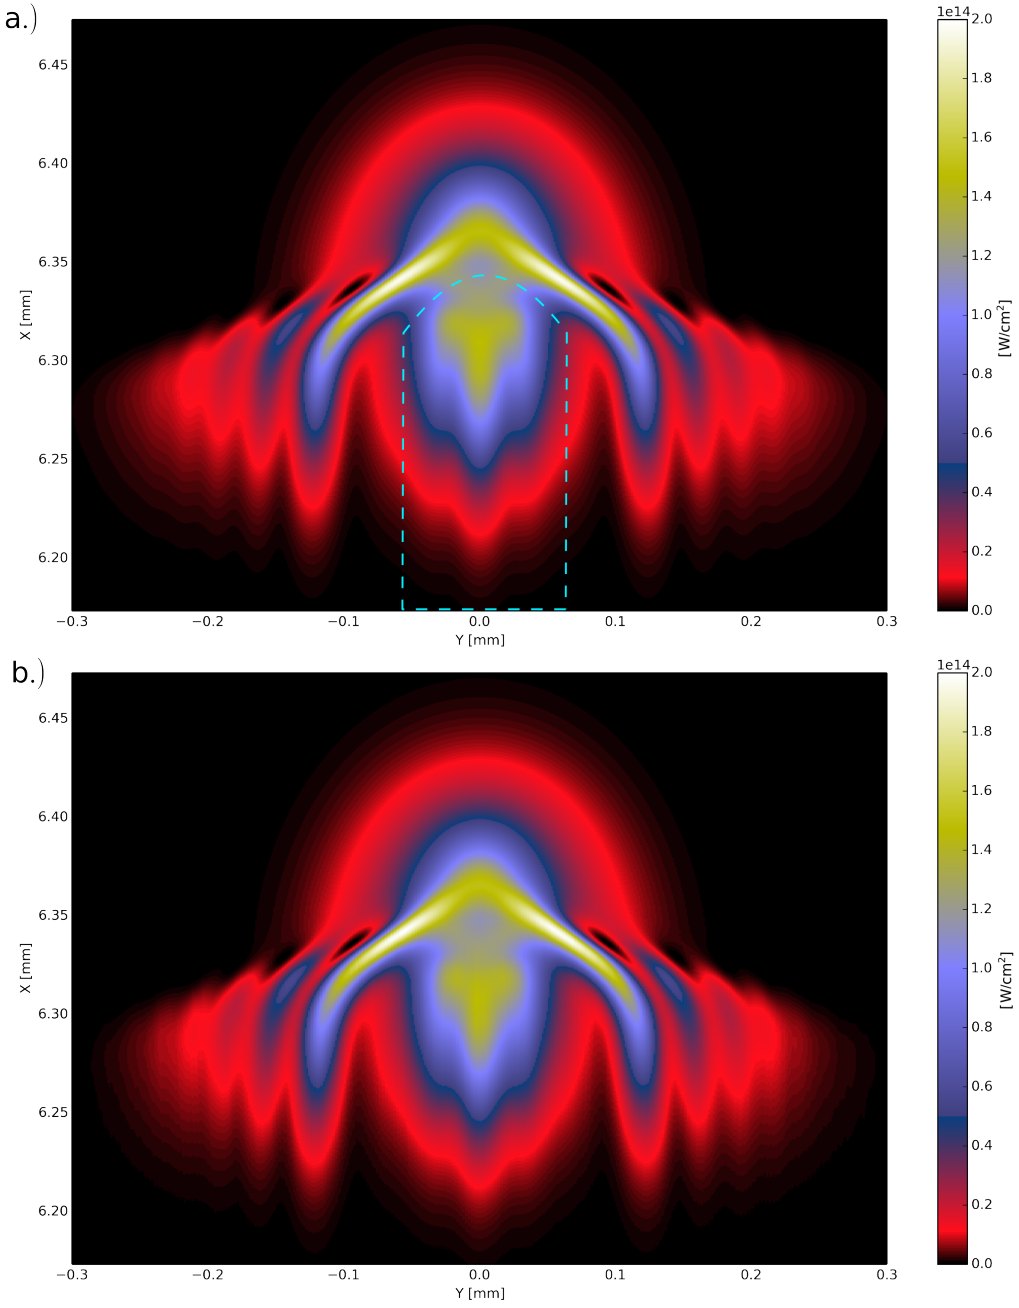

Figure 4.13: Spatial intensity distribution of the intense 100 μm wide 300 fs pulses after 23 ps of propagation in air. Highlighted in cyan is the collapsing internal structure, indicating a refocusing event has begun. a.) Dynamic plasma model b.) Static plasma model.

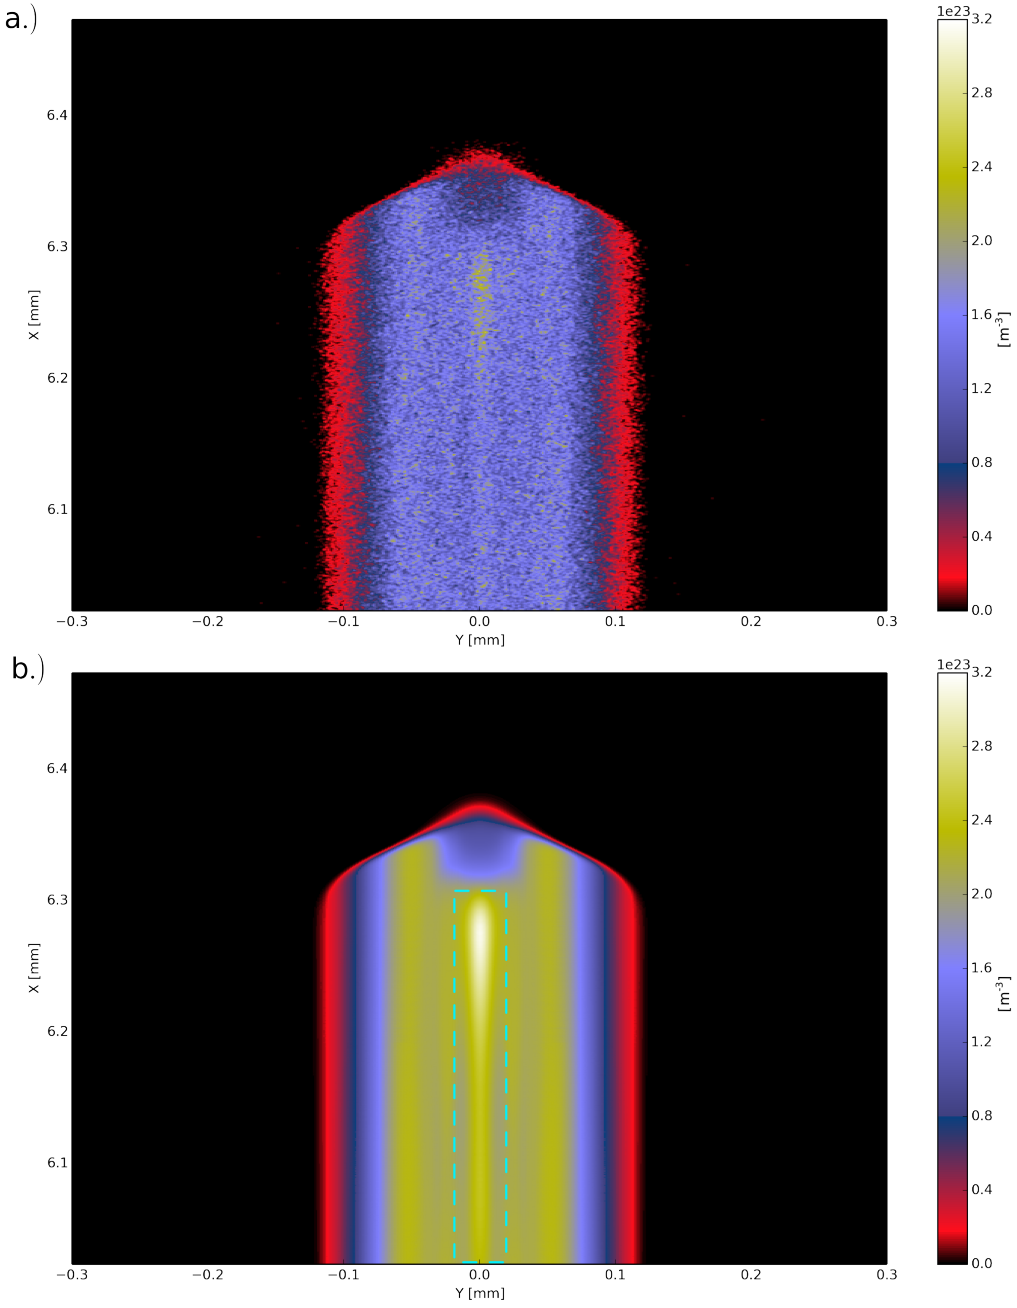

Figure 4.14: Spectral distribution of the intense 100 μm wide 300 fs pulse simulations after 23 ps of propagation in air. Highlighted in cyan is the plasma peak resulting from the refocusing of the filament. a.) Dynamic plasma model b.) Static plasma model.

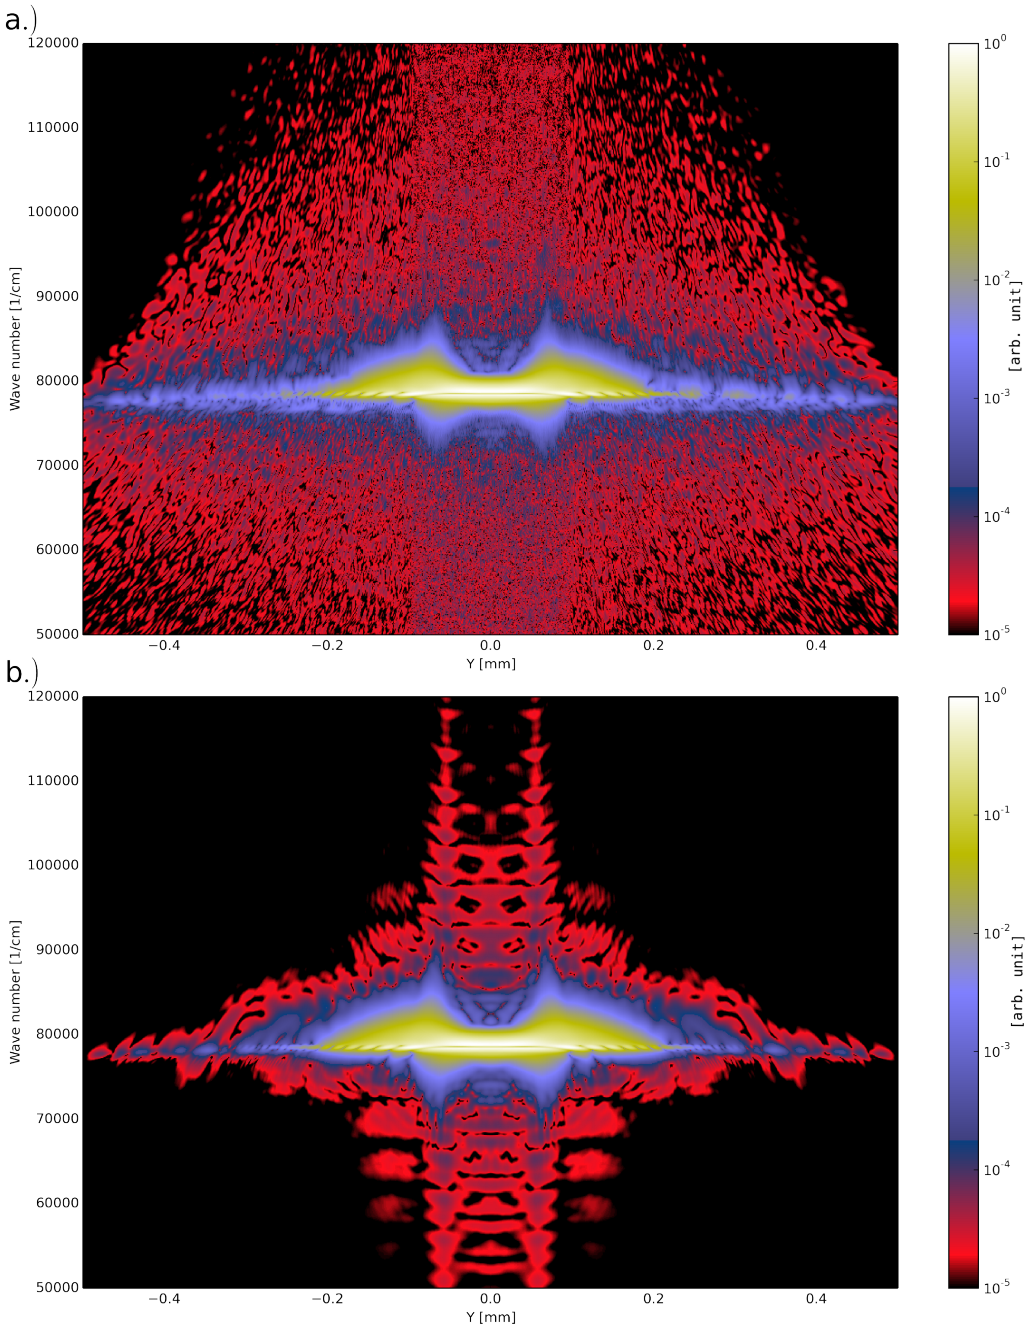

Figure 4.15: Spectral distribution of the intense 100 μm wide 300 fs pulse simulations after 23 ps of propagation in air. Frequency amplitude is normalized to the highest frequency component. The color bar is a logarithmic scale. a.) Dynamic plasma model b.) Static plasma model.

From figure 4.13 we can see that increasing the intensity of the pulse has indeed reproduced the features we were looking for. In addition to the ringed conical emissions we can see that the internal structure of the pulse (cyan) has collapsed in on itself, triggering a refocusing event. It's important for us to note that the pulse has only been propagated for 23 picoseconds but has already evolved to a point well beyond where any of the others have. Unfortunately, the two plots are virtually indistinguishable from each other. Were

the ponderomotive force at the root of the differences observed in the duration tests then we would have expected to see at least some difference between the two simulations. Instead the differences are even less apparent than they were for the 200 μm wide, 300 fs long test.

Looking at figure 4.14 we can see that the two are very similar. Aside from the fuzziness inherent in the dynamic plasma density we do not see any noticeable differences. In contrast to the density plots of the 200 μm wide 300 fs and 500 fs pulses we can see that the very center of the pulse is where the plasma is the most dense; this is due to the beginning of the formation of a secondary shockfront (cyan). Conspicuously missing from the dynamic density plot is the narrowing of the transverse plasma density profile we had seen in the 500 fs simulation. Apparently the tails of the shockfront are either insufficiently intense or long enough to induce that effect.

Finally, within figure 4.15 we can see that the center of the pulse contains a greater diversity of frequencies than did the 200 μm wide 300 fs pulse simulations. This tells us that the higher intensity of the shockfront has allowed for it to compress down to a shorter duration. We also see that the horns of the spectrum are quite broad, compared to previous runs; this is likely due to the compact nature of the filament as a whole. Unsurprisingly we see that the spectra of the two models are, like the intensity distributions, very similar.

4.3.3 Analysis of Size Tests

Reducing the size of the pulse and increasing its intensity should have increased the magnitude of the ponderomotive force substantially. The fact that we saw even less of a difference in the size tests than we did in the duration tests indicates that our changes did not massively increase the ponderomotive force. It may be that we did not allow the narrowed, intensified, simulations to run for a long enough time and thus did not observe any cumulative effects plasma dynamics may have had upon them, alternatively the pulses may have needed to be longer.

5 Conclusion

5.1 Summary of Work

Over the course of this project we simulated numerous laser pulses propagating through air. All of these simulations exhibited behaviours and features which are characteristic of filamentation. In the 30 fs pulse simulations we saw super-continuum generation from an optical shockfront. In the other simulations we clearly saw the formation of ringed conical emissions whose spectral content matched expectations. Finally we saw the beginnings of a refocusing event in the 100 μm, 2x10^{14} W/cm^2, simulation. Having seen these things we are confident that our simulations do indeed contain filaments.

In our simulations we tested two models of a plasma, the static and dynamic models, so as to determine whether plasma dynamics played a significant role in the evolution of a filament. From varying the duration of a pulse up to 500 fs we saw that there were some visible differences; the conical emissions were consistently weaker and less developed and the tails of the shockfront were narrower in the dynamic plasma simulations. The origin of these differences was unclear but the narrowing of the plasma density profile in the wake of the 500 fs pulse simulations suggests that the ponderomotive force is at work. However, when we ran simulations

intended to increase the ponderomotive force, the simulations turned out to be even more similar to each other. Based on the outcomes of the tests of duration we can conclude that plasma dynamics do have some impact on pulses with durations several hundreds of femtoseconds long but that effect is quite small. As the tests of the ponderomotive force failed to show a significant difference between the two models we must conclude that the differences seen in the duration tests were not due entirely to the ponderomotive force.

5.2 Possible Future Studies

A reasonable simulation that could next be run would be one where we increase the wavelength of the filamenting laser pulse. As the ponderomotive force is dependant on the inverse square of the laser frequency, doubling the wavelength should quadruple the ponderomotive force experienced by the free electrons. By running the intense 300 fs, 100 μm, pulse simulation again but with a wavelength of 1600 nm we would hopefully see some effects of the ponderomotive force. If we did not then we could be certain that the ponderomotive force is not a significant player for femtosecond laser filaments.

In our simulations we propagated our pulses over distances less than a centimeter. Given that filaments in air can stretch over *meters* it is not unreasonable to think that we have not run our simulations for long enough. It is possible that the effects of plasma dynamics are so weak that their impact is not visible in the early stages we simulated and only over long distances is the cumulative effect visible. If this is the case then it would be desirable to run at least one simulation for a much longer time, possibly over numerous refocusing events. By running such a simulation we could be certain that the duration of the propagation is not the cause of our failure to see any significant difference between the static and dynamic plasma models. Such a simulation would require very long computation times, unfortunately.

By losing the third dimension both the plasma defocusing and Kerr self-focusing are weakened. It is possible that this has negatively impacted our simulations as it may have reduced the already weak plasma effects below a noticeable level. It would be beneficial to run a fully 3D simulation of a filamenting pulse. In order to do this some simplifications would need to be made in order to make the simulation possible. The simulation space must be made smaller to reduce the huge number of spatial grid points and as such the maximum duration of a pulse would probably be limited to around 100 fs. The number of particles per cell in the dynamic plasma would have to be reduced to the absolute bare minimum as the PIC component of the code is by far the heaviest.

References

[1] C. Varin , G. Bart, R. Emms, T Brabec, "Saturable Lorentz model for fully explicit three-dimensional modeling of nonlinear optics", Opt. Exp. 23 (2014)

[2] P. E. Ciddor, "Refractive index of air: new equations for the visible and near infrared," Appl. Opt. 35, 1566–1573 (1996).

[3] L. Keldysh, "Ionization in the field of a strong electromagnetic wave", JETP 47, (1964)

[4] A. Perelomov, V. Popov, M. Terent'ev, "Ionization of atoms in an alternating electric field", JETP 50, 1393-1409 (1966)

[5] M. Ammosov, N. Delone, V. Krainov, "Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field", JETP 91, (1986)

[6] K. Mishima, M. Hayashi, J. Yi, S. Lin, H. Selzle, E. Schlag, "Generalization of Keldysh's theory", Phys Review A, 66, (2002)

[7] A. Lawrence-Douglas, "Ionisation effects for laser-plasma interactions by particle-in-cell code", U. Warwick, (2013)

[8] L. Bergé, S. Skupin, R. Nuter, J. Kasparian, J.-P. Wolf, "Ultrashort filaments of light in weakly ionized, optically transparent media", Rep. Prog. Phys., 70, (2007)

[9] A. Couairon , A. Mysyrowicz, "Femtosecond filamentation in transparent media", Phys. Reports, 441, (2007)

[10] A. Scrinzi, M. Geissler, T. Brabec, "Ionization above the coulomb barrier", Physical Review Letters, 83, (1999)

[11] K. Yee, "Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media", IEEE Transactions on Antennas and Propagation, 14 (1966)

[12] J. Jackson, D. Gordon, M. Helle, J. Pnano, "Fully explicit nonlinear optics model in a particle-in-cell framework," J. Comput. Phys.s, 250, (2013)

[13] M. Hercher, "Laser-induced damage in transparent media", J. Opt. Soc. Am. 54, (1964)

[14] A. Talebpour, J. Yang, S. Chin, "Semi-empirical model for the rate of tunnel ionization of N_2 and O_2 molecule in an intense Ti:sapphire laser pulse", Opt. Com., 163, (1999)

[15] S. Chin, "Femtosecond Laser Filamentation", New York, Springer, (2010)

[16] Progress in Ultrafast Intense Laser Science, Berlin, Springer, (2006)

[17] S. Tzortzakis, B. Lamouroux, A. Chiron, S. Moustaizis, D. Anglos, M. Franco, B. Prade, A. Mysyrowicz, "Femtosecond and picosecond ultraviolet laser filaments in air: experiments and simulations", Opt. Com., 197, (2001)

[18] S. Chekalin, A. Dokukina, A. Dormidonov, V. Kompanets, E. Smetanina, V. Kandidov, "Light bullets from a femtosecond filament", J. Phys. B: At. Mol. Opt. Phys., 48, (2015)

[19] T. Brabec, F. Krausz, "Nonlinear optical pulse propagation in the single-cycle regime", Phys. Rev. Lett., 78, (1997)

[20] R. Kupfer, B. Barmashenko, I. Bar, "Computational modeling of laser-plasma interactions: Pulse self-modulation and energy transfer between intersecting laser pulses", Phys. Rev. E, 88, (2013)

[21] A. Braun, G. Korn, X. Liu, D. Du, J. Squier, G. Mourou, "Self-channeling of high-peak-power femtosecond laser pulses in air", Opt. Lett. 20, (1995)

[22] R. Chiao, E. Garmire, C. Townes, "Self-trapping of optical beams", Phys. Rev. Lett., 13 (1964)

[23] A. Couairon , L. Bergé, "Self-guided propagation of ultrashort IR laser pulses in fused silica", Phys. Rev. Lett., 87, (2001)

[24] V. Lugovoi, A. Prokhorov, "A possible explanation of the small-scale self-focusing filaments", JETP Lett. 7, (1968)

[25] K. Ishikawa, H. Kumagai, K. Midorikawa, "High-power regime of femtosecond-laser pulse propagation in silica: Multiple-cone formation", Phys. Rev. E, 66 (2002)

[26] S. Chin, "The physics and the challenge of the propagation of powerful femtosecond laser pulses in optical media", La Phys. au Can., 66 (2004)

[27] C. Ma and W. Lin, "The life cycle of infrared ultra-short high intensity laser pulses in air", Eur. Phys. J. D, 69 (2015)

[28] R. W. Boyd, "Nonlinear Optics 3rd Ed.", San Diego, Academic Press (2010)

[29] H. Sun, J. Liu, C. Wang, J. Ju, Z. Wang, W. Wang, X. Ge, C. Li, S. L. Chin, R. Li, Z. Xu , "Laser filamentation induced air-flow motion in a diffusion cloud chamber", Opt. Exp. 21 (2013)

A Pulse Envelope Recovery

To recover the amplitude envelope of a 1D laser pulse we follow an extremely simple routine, as shown in the figure below.

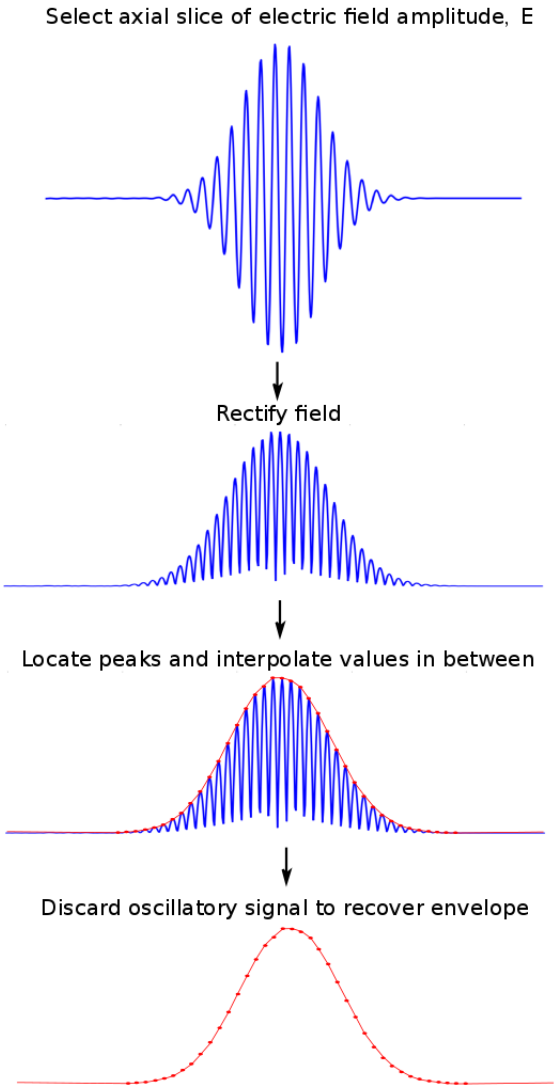

This procedure can be expanded to work in 2D; for a laser pulse propagating along the X-axis we can perform the procedure for every grid point along the Y-axis. In this way the envelope is recovered in strips. This method is crude but highly effective and accurate. Seen in the image below is an example of this process. In the figure we show first the oscillating electric field a 10 fs 800 nm pulse and below it the envelope recovered from it.