

Chapter 5

Continuum Generation

We have seen in the previous chapter that at powers high enough for self-focusing, interesting pulse propagation dynamics can result. Experiments above the critical power for self-focusing can result in the splitting of pulses into two or more sub-pulses. This pulse splitting arises in part because of the broad spectral bandwidths that are created by self-phase modulation and then separated temporally by group velocity dispersion. But, what happens at even higher input powers? Do the pulses continue to split into even more sub-pulses? And how broad does the spectrum get? It has been observed that at powers above the critical power for self-focusing, propagation in solids, liquids, and gases can produce an extremely broadened spectrum. The range of wavelengths produced often spans the entire visible spectrum and extends into both the IR and the UV. This phenomenon has been called spectral super-broadening or continuum generation. The continuum spectrum produced has been used in optical parametric amplification [120, 47, 121, 122], time-resolved broadband absorption and excitation spectroscopy [123, 124], and the generation of short pulses [125]. The exact process responsible for the generation of this continuum, however, is currently unknown.

Spectral super-broadening was first observed by Alfano and Shapiro in 1970 [126]. Their experiment involved propagating 530 nm picosecond pulses in glass with an input power of $\sim 1 \text{ GW/cm}^2$. At the output, they observed a spectrum that spanned from 400 to 700 nm. They also observed the formation of self-trapped filaments within

the glass. The first evidence of self-trapped filaments was reported in 1964 when extremely long, thin damage spots were observed in glass and other materials [127]. For self-trapped filaments to form, a process must be involved that arrests the self-focusing effect and allows the pulse to propagate without spreading. Self-trapped filaments have been explained in a variety of ways. Some of these explanations include a self-produced dielectric waveguide [128], a self-guiding structure similar to a fiber consisting of a Kerr cladding and a weakly ionized core [129], and as a moving focus [130]. One or more self-trapped filaments are often present with continuum generation. It has been asserted that regardless of the medium in which the continuum is generated, two requirements for continuum generation are self-trapped filaments and dielectric breakdown [131]. Although the validity of the first remains a question, the second has been proved incorrect. Evidence that dielectric breakdown is not necessary for continuum generation is found in the properties of the transmitted continuum beam. The spatial properties are only slightly modified from that of the incident beam [132], and the emitted continuum is collimated, modulated, and polarized [133]. Dielectric breakdown should result in emission that is neither collimated nor polarized. The necessity of self-trapped filaments has been brought into question by an experiment involving continuum generation in a thin ($500 \mu\text{m}$) jet of ethylene glycol [43]. The length of the sample is small enough that self-trapped filaments may not form. The role of self-focusing is also minimized in this experiment. It is commonly believed, however, that self-focusing plays an important role in the continuum generation process. That self-focusing is important is evidenced by the fact that continuum generation only occurs with input powers above the critical power for self-focusing [134].

The white light continuum appears as a central beam of white light surrounded by rings of color. These rings do not extend continuously from the central white light, but occur as distinct rings of red, green, and blue light, often separated from the central white light beam — and from each other — by rings of darkness. The higher

frequency light appears in the outermost ring, and the lower frequency light forms the ring nearest the white light continuum. This arrangement of frequency components is exactly opposite that which would be predicted by diffraction. An example of white light continuum generated in sapphire is shown in Fig. 5.1. The central white light continuum is also observed to be both spatially [133, 41] and spectrally modulated [131, 132, 134, 41]. Super-broadened spectra have been observed in a variety of materials including glasses such as fused silica and BK-7 [133, 135]; crystals such as calcite, quartz, sodium chloride and potassium bromide [133, 136]; semiconductors such as gallium arsenide and zinc selenide [136]; gases such as neon, argon, krypton, xenon, nitrogen and air [137, 132, 134, 138, 129]; and liquids such as water, methanol, carbon tetrachloride, carbon diselenide, and benzene [131, 135]. Both picosecond and femtosecond pulses have been used to generate continua, with incident pulses spectrally centered from around 248 nm [134] to around $9.3 \mu\text{m}$ [136]. Continuum spectra under all conditions share certain similarities including the central white light continuum surrounded by rings of color, an asymmetric pulse spectrum with the blue spectral wing extending further than the red, a dependence on input power rather than intensity, and a polarization that is in the same direction as that of the input pulse. These striking similarities indicate that the mechanism responsible for continuum generation is very general, having little to do with the spectroscopic features of the material involved. The width of the spectrum generated is, however, dependent on the medium in which it is generated [124], but it does not scale with the size of the nonlinear index of refraction as might be expected. In fact, in materials with high n_2 , and therefore stronger self-focusing and self-phase modulation, the spectrum is actually narrower than for materials with low n_2 [135].

Several approaches have been taken to explain the observed spectra. On first observing continuum generation in glass, Alfano and Shapiro attributed the broad spectrum to a stimulated, nondegenerate four wave mixing process [126]. This four wave




Figure 5.1: Far-field image of continuum generated in sapphire.

mixing paradigm was championed in many following papers by Penzkofer *et al.* [139, 140, 141, 142]. Penzkofer's work asserts that there are several processes that can be involved in continuum generation. The primary process is the amplification of quantum noise by a stimulated, parametric four-photon interaction. In certain materials, and under certain conditions, stimulated Raman scattering and a second parametric process driven by the pump laser and stimulated Raman light come into play. These generated frequencies can result in a cascaded frequency upconversion effect, generating a broad spectrum. The asymmetry of the spectrum is explained as absorption of the red and infrared components of the pulse [141]. Experiments by this same group indicate that the on-axis spectrum consists of a spectrally broadened central peak, with low intensity wings extending into the UV and IR. The four-wave mixing theory predicts broad wings, but it does not account for the significant broadening near the input wavelength. To explain this effect, self-phase modulation is invoked as a special case of four-wave mixing in which the amplified wavelengths are already present in the pulse and do not need to arise from quantum noise. Although this four-wave mixing theory does predict broad spectral wings, it predicts a peaked spectrum corresponding to Raman resonances in the material. The experimentally observed spectra are often modulated, but the peaks do not always occur at the predicted frequencies. In addition, this theory fails to account for self-focusing in the medium, and it fails to address the ring structure observed surrounding the central white light continuum. Although it predicts a very broadened spectrum for special cases where the resonant structure of the material is conducive to strong Raman scattering, the theory does not appear to be as universal as experiments indicate continuum generation is.

Several attempts have been made to model continuum generation by more rigorously solving the nonlinear wave equation for pulse propagation. The wave equation, accounting for self-phase modulation, four-wave mixing, and self-steepening and neglecting dispersion, predicts an asymmetric spectral broadening. The predicted broad-

ening, however, is less than the observed broadening [143]. Another work, in which the same basic equation is solved using a different method, is able to predict broader pulses [144]. In both cases, however, dispersion is assumed to be zero. This assumption is valid when comparing to continua generated in gases, but it is invalid in experiments involving most liquids and solids. That dispersion plays an important role in propagation has already been demonstrated in Chapter 4 of this thesis where it is shown that dispersion leads to temporal pulse broadening and arrests the collapse of a self-focusing pulse. In addition, self-phase modulation implies increased broadening with a higher nonlinear index of refraction in contrast to experimental evidence.

It has been observed that self-trapped filaments are nearly always present during continuum generation. These self-trapped filaments can only form if there is a process occurring in the medium that stops the collapse of the pulse due to self-focusing. The approach can then be taken that the processes involved in halting the collapse and allowing the formation of self-trapped filaments are most likely also responsible, at least in part, for continuum generation [114]. It was proposed that plasma formation, and the associated appearance of free electrons, was sufficient to stop the collapse and simultaneously enhance SPM such that a broad spectrum results [145]. Although this predicts a broad blue-shifted spectrum, the associated loss of energy is much higher than that observed [137]. Another twist on the influence of ionization was proposed [129] that describes the self-trapped filament as a leaky anti-guiding structure with a Kerr type cladding and a slightly ionized core. The nonlinear index is high at the edge of the trapped pulse but falls off toward the center due to the influence of free electrons. This antiguide cannot support stable propagation modes because of the lack of total internal reflection and thus loses some energy through refraction. The theory proposes that quasi-stable modes exist and that refraction losses are responsible for the off-axis ring structure surrounding the white-light continuum in the far-field. Spectral broadening of the guided light is attributed to SPM.

The most promising theory to date describes continuum generation in condensed media and is again based on the generation of free electrons [135, 146]. In this theory, Brodeur and Chin assert that it is not plasma formation that generates the free electrons, but a multiphoton excitation (MPE) of electrons from the valence band to the conduction band. This MPE is a bandgap dependent process, and experiments in several liquids and solids have shown a dependence of the self-focal diameter on the bandgap of the material [135, 146]. The basic premise of this theory is that as the pulse undergoes self-focusing, different temporal slices of the pulse reach a maximum intensity where MPE occurs. MPE generates essentially free electrons, which serve to decrease the index of refraction and arrest the focusing. The sudden change in index of refraction results in a sudden drop in the phase. A change in phase with respect to time gives the instantaneous frequency. In this way, MPE generates a broad blue-shifted spectral wing. This theory is based on a moving focus model of pulse self-focusing [130, 147, 114], which is also applied to the continuum problem by Strickland and Corkum [41]. According to the moving focus model, the pulse is divided into discrete temporal slices that undergo self-focusing independently. Because each slice has a different power, various slices focus to different positions along the propagation axis. These slices in effect form temporally narrow transient sub-pulses. According to Strickland and Corkum, these sub-pulses generate continua via an on-axis phase retardation. The observed continuum spectrum is the summation of the broadened spectra produced by each of these transient sub-pulses. The spatial and spectral modulation characteristic of the observed continuum can be explained as far-field interference resulting from the multiple broadbandwidth sources. The more modern theory relies on MPE instead of an on-axis phase shift and thereby inherently includes a dependence on the band gap of the medium.

5.1 On the Verge of Continuum Experiments

In an effort to better understand the continuum generation process, and in light of the excellent agreement between the (3+1)-D theory and our FROG measurements above the critical power for self-focusing presented in Chapter 4, our group pushed FROG measurements into the region just above the continuum formation threshold [119]. During these measurements, blue light was clearly visible by eye. To ensure that the phase matching bandwidth of the second harmonic generating crystal in the FROG device was adequate, the 100 μm BBO crystal was replaced by a 50 μm BBO crystal. As in the pulse splitting measurements, the laser beam was spatially filtered and focused to a FWHM spot size of 70 μm at the entrance face of a 3.0 cm long fused silica sample. SHG FROG measurements of the on-axis far-field were performed before and after propagation through the sample. The input pulses were nearly transform-limited with an intensity duration of ~ 90 fs FWHM.

Figures 5.2 (a) and (b) show the measured on-axis intensity after propagation for input peak powers of 6.8 and 7.4 MW, respectively. From the pulse splitting experiments performed in Chapter 4 we might expect the pulse to split into many more sub-pulses with these higher input powers. As the peak input power is increased, however, the sub-pulses begin to coalesce into a single broad pulse [119]. Calculations using Eq. 4.14 without the nonparaxiality term are shown in Figs. 5.2 (c) and (d). The input powers used in the calculation match those of the experiment to within 15%. In both the experiment and the calculations, as input power is increased, the central peak broadens and merges with the leading and trailing peaks.

Calculations at lower powers have shown that the splitting of the pulse into two sub-pulses arises first in the far-field and progresses to the near-field with increasing input power. Multiply-split pulses are observed in the far-field when the near-field is fully split into two sub-pulses. At higher input powers, however, multiply split pulses

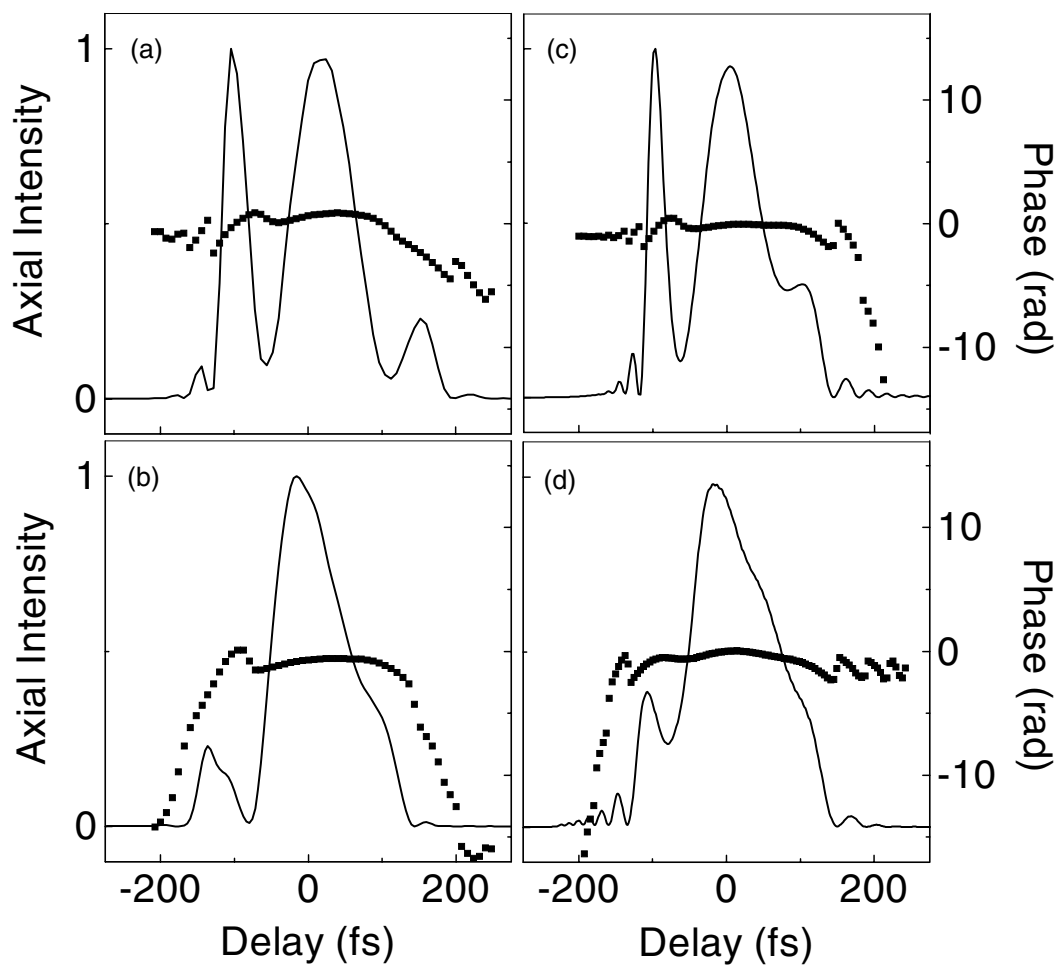


Figure 5.2: Measured axial intensities and phases for input powers of (a) 6.8 MW and (b) 7.4 MW. The corresponding calculated intensities and phases are shown in (c) and (d). (From Ref. [119])

do not appear in the near-field. Although the near-field is modulated, distinct multiple splitting arises only from the build-up of phase in the far-field as described previously and does not arise from the same physical processes responsible for the initial pulse splitting.

Measured and calculated spectra corresponding to the temporal intensity profiles displayed in Figs. 5.2 (b) and (d) are shown in Figs. 5.3 (a) and (b), respectively. The spectra are plotted on a log scale to emphasize the long tail extending to both sides of the central peak. By contrast, a bandwidth-limited, 100 fs Gaussian pulse centered at 800 nm would have wings with an intensity three orders of magnitude lower than those shown here just 50 nm on either side of the peak. These tails are thus evidence of weak continuum generation. The short wavelength tail of the calculated spectrum falls off faster than that of the measured spectrum. As indicated by the theory of Chin and Brodeur [135, 146], inclusion of multiphoton excitation in the theory may account for this difference.

5.2 Full Continuum Experiments

To improve and test the theory, good experimental data with which to compare is essential. Careful studies of continuum generation are difficult because of the broad bandwidths involved. With spectra spanning several hundred nanometers, an accurate intensity autocorrelation is difficult, if not impossible, because of the phase-matching bandwidth required. Since FROG is an autocorrelation-based device, increasing the input power above that used for Fig. 5.2 (b) quickly results in spectrally broad pulses for which the validity of SHG FROG comes into question. In addition, continuum pulses diverge quickly after exiting the medium and result in very low on-axis intensities. These intensities can be sufficiently low that a SHG FROG signal is difficult to obtain. Near-field SHG FROG measurements could be performed by imaging the output face of the sample onto the entrance aperture of the FROG apparatus. Such a measurement




Figure 5.3: (a) The measured axial spectrum for an input power of 7.4 MW, and (b) the corresponding calculated spectrum. (From Ref. [119])

would suffer the same phase-matching bandwidth limitations, however, and near-field SHG FROG measurements have not been tested in simpler propagation regimes. For these reasons, temporal measurements of the field after propagation are not practical. Therefore, careful measurements of the spectrum and other empirical parameters are used for comparisons with theory at this time.

All data presented to this point have concentrated on the on-axis field before and after propagation. In an attempt to gain insight into the continuum generation process, we add a spatial dimension to the spectral measurements. Qualitative spatial-spectral, far-field measurements have been available for years in the form of verbal and pictorial descriptions of a white light beam surrounded by rings of color. The only previous quantitative measurement, by Strickland *et al.* [41], recorded the spatial spectrum of the far-field using a spectrometer and CCD. This chapter reports the first careful measurements of the near-field spatial spectrum of continuum pulses (Section 5.2.4). Detailed investigations of power loss and continuum behavior as a function of input power and timelapse are also reported in Sections 5.2.2 and 5.2.3.

For these measurements, spatially filtered pulses are focused at the input of the fused silica to a spot size of $33 \mu\text{m}$ FWHM or a $26 \mu\text{m}$ radius at $1/e^2$ of intensity with a 10 cm focal length lens. The fused silica sample length is reduced to 0.318 cm. These changes are made to better facilitate comparisons with theory. Self-focusing plays an important role in continuum generation, reducing pulse radii during their passage through the sample by an order of magnitude to approximately a few microns. These dramatic changes in beam size require different stages of the propagation to be calculated on different size grids. To minimize the associated difficulties, we focus to a smaller initial spot size to reduce the amount of spatial beam compression that occurs during propagation. A shorter sample can thus be used since the propagation-limited beam size is reached sooner.

5.2.1 Qualitative Observations

The generated continuum exits the sample and propagates 10 cm to a white card where it is observed by eye. Observations are also made after the beam travels 35 cm to a white card. These observations differ from those at 10 cm only in that the beam is much larger and therefore weaker by eye. Subtle changes in the continuum behavior are more easily observed on the card placed 10 cm after the sample, so those observations are reported here.

Continuum spectra are generated with input peak powers ranging from 2.7 to 12.6 MW. Self-trapped filaments are not visible by eye in the sample; the sample is short enough, however, that they may be difficult to observe under these conditions. At the lowest powers, the beam on the card looks like a small red spot. As input power is increased, the red spot grows larger until white light begins to appear in the center, surrounded very closely by rings of blue and green. I should point out that in these experiments the continuum never goes through a stage where the central red spot is surrounded by a blue ring as was seen in the earlier experiments. This is most likely a result of the smaller focusing geometry used here. It is probable that in selecting input powers to sample, we took too large a step and missed the transition from a red spot to red surrounded by blue before the central white light spot appears. At an input power of 4.5 MW the green and blue rings are larger, and the white center takes on a yellow-orange hue. Following this, green frequencies begin adding into the central beam and then it appears white again. At an input power of 5.8 MW, a red outer ring appears separated from the white/blue/green structure by a large black ring. As the input power is increased further, the red and blue rings grow fatter. At an input power of 7.4 MW, a thin yellow ring forms around the outer edge of the red ring. This yellow ring is tenuous, and seems to appear and disappear many times per second. The white center again turns a yellow-orange color at a peak input power of 9.9 MW. And at 10.8

MW rings of black separate each ring of color surrounding the again white center. At this point, we see a strong white light central beam surrounded by a ring structure that consists, going outward from the center, of a thin black ring, a thicker blue ring, a second thin black ring, a thin green ring, a thick black ring, and a thick red ring with a thin yellow lining at the outermost edge. Increasing the power to 12.6 MW results in a similar far-field structure with the addition of a second outer red ring. Again, all color rings are separated by areas of darkness. The appearance of this ring structure is reminiscent of some kind of interference, but the mechanism that is actually channeling light off axis is unknown. Also, the appearance of two distinct outer red rings begs the question of whether one ring may be the result of the continuum generation process while the other is actually fundamental light being funneled off axis, or if both are created via the same mechanism.

This data is presented as if one always observes a distinct set of concentric rings in the far-field. In reality, though, one occasionally observes a pattern that resembles a series of overlapping concentric ring sets, each with a different spatial center. By slightly moving the sample, however, the far-field ring structure can be returned to the desired concentric ring pattern. It is also interesting to note that the power throughput is different in these two cases. Power measured after the sample is on the order of 10% higher when a single ring structure is observed. These patterns have been reported once before and were attributed to fluctuations due to plasma buildup [131]. I believe, however, that these differing patterns can be explained by considering the formation of self-trapped filaments. If a single ring structure results from continuum generation along a single trapped filament, then these more complicated patterns probably arise from the formation of multiple filaments. Continuum is generated along each filament and the resulting concentric ring patterns overlap in the far-field. For all of the data presented in this chapter, care was taken to ensure that only a single ring structure was present during the measurement.

It is interesting to note that although self-focusing and continuum generation are power dependent effects, the power figures presented in this chapter seem to indicate that continuum is occurring at the same powers at which we previously reported pulse splitting. One possible explanation for this is that the experimental set-up for these experiments is different from the pulse splitting experiments in that a shorter sample and a smaller focal spot are used. The pulse enters the sample with a much higher initial intensity than in the pulse splitting experiments. In those experiments, the pulse is spectrally broadened and GVD acts to spread the pulse in time, initiating pulse splitting. By the time the pulse reaches the same peak intensity used in the continuum experiments, significant spectral broadening has already occurred. Therefore, we should expect the propagation dynamics to be somewhat different in the two cases.

5.2.2 Power Throughput as a Function of Intensity

The power measurements presented here and in the next section were performed using a large area photodiode detector based power meter. The continuum light generated during propagation diverges rapidly after exiting the sample. Therefore, the power meter was placed as close as possible to the output face of the sample to collect all of the emitted light, including both the central white light and the colored rings.

Figure 5.4 shows the percentage of power lost during propagation as a function of input power. This percentage was determined by measuring the power with and without the sample present in the beam path. The two lowest power points in Fig. 5.4 correspond to situations where no central white light beam is observed. As the white light continuum appears, the losses jump to $\sim 15\%$. As power is increased, the overall trend is for the losses to increase as well up to an input power of 8.0 MW. There are three points, however, that do not follow this trend. At input powers of 4.5 and 10 MW, denoted by circles in Fig. 5.4, the losses are much lower than expected. Interestingly, it is at these powers that a yellow-orange color is observed in the central white-light




Figure 5.4: The percentage of power lost during propagation as a function of input power. The circles represent points at which the central far-field beam is yellow-orange instead of white.

region of the beam. The very high loss observed at 6.8 MW does not seem to correlate to any of the observed continuum behaviors.

Overall, it is evident that the processes involved in continuum generation involve a loss of overall power during propagation. This lends credence to the idea that multi-photon excitation or an absorption process may be involved. In fact, calculations based on the (3+1)-D theory presented earlier and including four-photon absorption predict accumulated losses of 12-14% for an input power of 12 MW [148].

5.2.3 Power Throughput as a Function of Time

Power measurements are also performed as a function of time. For these measurements, a post holder is attached to the optical table behind the sample, and a collar is placed on the post of the power meter head to ensure that the power meter could be removed and replaced to the same position. First the power without the sample present is measured. Next the beam is blocked, the sample is added into the beam path, and the power readings are recorded as soon as the beam is unblocked. The power meter is then removed from the beam path while the beam continues to propagate through the sample. The power meter is replaced and power measurements are made at intervals of several seconds to several minutes for approximately one half hour. The detector is sensitive to changes in temperature, therefore the power meter is removed between measurements to avoid unnecessary heating. After each set of measurements, the power is again measured without the sample present to check that the laser power has not drifted during the measurement.

When the laser first traverses the sample, the power loss discussed in the previous section is evident. At low input powers, after the beam is incident on the sample for several minutes, the output power begins to rise. After sufficient time, the power at the output is very nearly the same as the input power. At higher input powers, the output begins to rise after a shorter time. For an input peak power of 13 MW the output power

matches the input power after only 15 seconds. During the time the output power is rising, the far-field continuum structure also changes. Initially, a typical ring structure including the thin yellow lining around the outermost red ring is observed as described in Section 5.2.1. As the output power increases with time, the outer yellow lining disappears from around the red ring and simultaneously, the central white light region takes on a yellow hue. The initial power loss followed by a return of the output power to the value of the input is likely a result of a saturable absorption in the material. That the output power rises sooner with higher input powers is consistent with a saturable absorption.

After longer periods of time, the output power again begins to drop. Again, the amount of time that passes before this event occurs decreases with increasing input power. With an input power of 3.7 MW, just above the power where a strong central white light spot is observed, the output power drops by roughly 40% after 10 minutes. After the power drops, the far-field has lost its white light characteristics and appears as a red spot that is roughly the same diameter as the white light spot that preceded it. When the sample is moved slightly such that the beam traverses a different spatial path through it, the original white light continuum beam returns. These observations suggest that even at the lowest powers for continuum generation damage is occurring in the fused silica. Damage by short IR pulses in fused silica has been reported previously and is associated with a rapid buildup of electrons in the conduction band [149]. This damage thus provides further evidence that free electrons may be implicated in the continuum generation process. A different study on the formation of color centers in glasses asserts that continuum generation is actually responsible for the damage to the glass [150]. The study finds that color centers are formed by one- or two-photon absorption of the short wavelength continuum light generated by the propagation. This same study, however, finds no evidence of the formation of color centers in fused silica, although they may have formed and not been observed.

5.2.4 Near-Field Spatial Spectrum

Measurements of the spatial spectrum are performed by imaging the exit face of the sample onto the entrance slits of a 0.25 meter imaging spectrometer with an $f = 10$ cm lens and a magnification of 16. A radial slice of the beam is selected by the entrance slit, dispersed, and imaged onto a 16 bit CCD camera. A mask in front of the slit allows the central, high intensity portion of the beam to be blocked when measuring the radial extremes so that more than four orders of magnitude of intensity variation can be acquired. A long pass filter is employed when measuring the high frequency side of the spectrum to ensure that only the first order diffracted light from the grating is collected. Neutral density filters are placed in the beam path when necessary to prevent camera saturation. The entire spectrum cannot be captured in a single measurement because the grating is more dispersive than the CCD element is large. Therefore, a series of 10 to 14 spatial-spectral measurements corresponding to different wavelength regimes are necessary to obtain the full spatial spectrum at each input power. A corresponding background spectrum is subtracted from each individually measured spectrum to account for the differing exposure times and filtering used in each measurement. Each spatial spectrum is normalized such that its spectral edge matches the edge of the one obtained for the wavelength range adjoining it and the full spatial spectrum is pieced together. When measuring spectra over such a broad range of wavelengths, it is also important to consider the efficiencies of both the spectrometer grating and the CCD. The grating efficiency is known from 300 to 1000 nm. While the camera efficiency is only calibrated from 625 nm to 975 nm, the general shape of the efficiency curve is known over a larger range. Since continuum spectra usually fall off quickly on the long wavelength side, the limitations for the red components are not severe. The short wavelength spectrum, however, often extends well below 625 nm. Fortunately, the grating efficiency of the 300 groove/mm grating used here is increasing in the region of 300 to

roughly 600 nm where it peaks. Over this same range, the camera efficiency is falling off such that combined efficiency is relatively flat over this range. For this data a flat efficiency with a value of the known efficiency at 625 nm is assumed over the range 400 nm to 625 nm.

Figure 5.5 shows the near-field spatial spectrum for an input peak power of 12 MW. Only one measurement of the radial position at each wavelength range is shown in this plot. During the measurement, the continuum character changed such that the data from the low intensity wing was no longer related to the higher intensity spatial data shown here. This is probably a result of the sample being damaged with time as discussed in the previous section. Rings of color would be represented in this plot as a particular wavelength showing up primarily at two distances, equidistant from the central peak. Since rings are observed to form with lower frequencies at the outermost edge, these rings should appear as a v-shaped peak or set of peaks that opens toward the distance axis in Fig. 5.5. It is evident from this data that the ring structure exists only in the far-field. Spatial modulation of the near-field occurs near the on-axis portion of the beam. Although the far-field continuum spectrum appears to have strong radial symmetry, the near-field spatial modulation does not reflect this symmetry. An interesting feature is the peak around 580 nm. This peak is roughly six orders of magnitude less intense than the peak near 800 nm. This peak is better illustrated by plotting the on-axis spectrum as shown in Fig. 5.6. This spectrum also exhibits a strong spectral asymmetry in the generation of continuum. The short wavelength side of the spectrum extends much further than the long wavelength side.

Lineouts of a spatial spectrum obtained with a lower input peak power of 3.6 MW are shown in Figs. 5.7 and 5.8. Figure 5.9 shows the on-axis spectrum in this case. As in the higher power case, the on-axis spectrum extends further to short wavelengths than do the spectra at other radii. In this low power case, however, the peak at 580 nm is nonexistent.




Figure 5.5: Spatial spectrum for an input power of 12 MW. Radial distance is plotted in mm and wavelength is plotted in nm.




Figure 5.6: On axis spectrum corresponding to Fig. 5.5.




Figure 5.7: Lineouts of spectral intensity from a spatial-spectral measurement of continuum generated with an input power of 3.6 MW. The plot at 1.37 mm is the on-axis spectrum.




Figure 5.8: Lineouts of spectral intensity from a spatial-spectral measurement of continuum generated with an input power of 3.6 MW. The plot at 1.37 mm is the on-axis spectrum.




Figure 5.9: On-axis spectrum for an input power of 3.6 MW

5.2.5 Discussion

Preliminary on-axis calculations based on the (3+1)-D equation used to model pulse splitting with the inclusion of four-photon absorption are capable of reproducing some general features of the continuum [148]. For instance, the theory predicts asymmetric spectral broadening just as observed experimentally and predicts an extensively broadened spectrum with spectral wings extending from 400 nm to 1.4 microns. This theory also predicts a peak around 580 nm in the high input power case. The calculated peak, however, is only three orders of magnitude less than the peak intensity, while the measured peak is nearly six orders of magnitude less intense than the peak. As discussed previously, the inclusion of four photon absorption also predicts a loss of power during propagation that is comparable to the power loss we observe.

The theory of Chin and Brodeur also generally reproduces experimental observations. It predicts an asymmetric, very broad spectrum and accounts for the far-field ring structure as interference from continua generated from multiple broad-bandwidth sources. The theory is also based on multiphoton excitation which the power data presented here supports. Comparisons with experiment reveal, however, that although it predicts sufficient broadening, it also predicts a blue wing that is an order of magnitude more intense than that observed by experiment. In addition, the theory relies on a moving focus model of self-focusing that does not present a realistic representation of the self-trapped filaments. In fact, an experiment was performed in 1969 that demonstrates that filaments are indeed trapped and not a result of a moving focus [151]. In this experiment, self-trapped filaments were sent through an aperture such that if filaments actually existed, they would pass through unaffected. But if they were actually a result of a moving focus, they would be destroyed by diffraction. The filaments persisted after propagation through the aperture. The reverse experiment was also performed. The self-focused pulse was propagated through a glass film such that a filament would be de-

stroyed by diffraction, but the larger focusing beams would pass through and continue to focus. In this case, filaments were selectively destroyed. For this reason, a moving focus model is probably not the best representation to be used. Also, the moving focus model predicts the formation of short temporal sub-structures during propagation. One might argue that perhaps short temporal sub-structures are formed as a result of pulse splitting. The (3+1)-D NLSE used to model our pulse splitting experiments, however, never yields distinct multiple splittings in the near-field. Therefore, it is unlikely that the continuum is actually generated by short-duration sub-pulses. In addition, this theory relies on one-dimensional calculations even though self-focusing is involved. Since the transverse dimensions of the beam are being changed dramatically by the propagation, a three-dimensional calculation is clearly more appropriate.

The question remains: what are the necessary components of a reasonable theory for continuum generation? Because continuum generation occurs only above the critical power for self-focusing, the theory should account for changes to the spatial dimension of the beam during propagation. Based on the power measurements presented here, clearly a loss mechanism such as multiphoton absorption or excitation is important. The data here also emphasize that a damage mechanism should be considered. The work of Chin and Brodeur clearly indicates that multiphoton excitation has a probable role in the propagation. Their work also points out the dependence on bandgap that can be accounted for by MPE. Counter to the work of Chin and Brodeur, a successful theory should include, or at least not explicitly exclude, the formation of self-trapped filaments. The formation of self-trapped filaments also offers an explanation of the ring structure associated with continuum in similar fashion to the moving focus model. If continuum is generated all along the filament, the light generated at different points in the medium will interfere to give a ring structure. The ring structure may also be explained by including a saturable process in the theory. Ring formation as a result of saturation has been predicted by Marburger [114]. Theory should com-

pare not only to the on-axis spectrum, but to the entire spatial spectrum. Descriptions of the far-field spatial spectrum are available both in the literature and in this chapter. This chapter also provides the first measurement of the near-field spatial spectrum. In addition, a successful theory should include the dispersion of the material. Continuum spectra produced in gases have shown a modulation that was initially associated with all continuum spectra. Continuum spectra in glass that do not exhibit this modulation have typically been explained as resulting from multiple shot experiments where the spectral modulation would be averaged out. I believe, however, that continuum spectra in glass do not show a strong modulation because the various spectral components are dispersed by GVD.

Although the processes involved in continuum generation cannot fully be explained at this time, some of the main issues involved have been elucidated. The observations presented here should provide a useful comparison for developing theories. Understanding continuum generation remains an interesting and important unsolved problem.