applied optics

Visualizing spatiotemporal pulse propagation: first-order spatiotemporal couplings in laser pulses

Michelle Rhodes,[1,2,*] Zhe Guang,[1] Jerrold Pease,[1] and Rick Trebino[1]

[1]Georgia Institute of Technology, 837 State Street, Atlanta, Georgia 30332, USA
[2]Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore, California 94550, USA
*Corresponding author: mrhodes3@gatech.edu

Received 19 December 2016; revised 6 March 2017; accepted 6 March 2017; posted 8 March 2017 (Doc. ID 283099); published 6 April 2017

Even though a general theory of first-order spatiotemporal couplings exists in the literature, it is often difficult to visualize how these distortions affect laser pulses. In particular, it is difficult to show the spatiotemporal phase of pulses in a meaningful way. Here, we propose a general solution to plotting the electric fields of pulses in three-dimensional space that intuitively shows the effects of spatiotemporal phases. The temporal phase information is color-coded using spectrograms and color response functions, and the beam is propagated to show the spatial phase evolution. Using this plotting technique, we generate two- and three-dimensional images and movies that show the effects of spatiotemporal couplings.

OCIS codes: (320.0320) Ultrafast optics; (320.5550) Pulses.

https://doi.org/10.1364/AO.56.003024

1. INTRODUCTION

The general theory of first-order spatiotemporal distortions [1] provides a complete and general framework for understanding first-order beam distortions/couplings. It described the complete and sometimes unintuitive relationships between well-known spatiotemporal couplings like spatial chirp, angular dispersion, and pulse-front tilt. It also drew attention to the angular counterpart to pulse-front tilt, known as "time versus angle" or the "ultrafast lighthouse effect," as it was first named in that work. Perhaps most importantly, it identified and described four previously unconsidered *imaginary* coupling terms, such as "wave-front rotation," that necessarily accompany the above-mentioned real coupling terms. These latter terms can be quite unintuitive.

Unfortunately, despite various attempts to describe and plot these new distortions, their effects are not always particularly obvious. The authors of Ref. [1] attempted to show phase couplings by plotting phase fronts in space versus time and frequency. While these plots are certainly correct, they do not impart much understanding about the appearance of such a pulse or how it evolves as it propagates.

The theory of spatiotemporal distortions makes it clear that the phase is just as important as the amplitude in describing pulses. Therefore, plotting the spatiotemporal phase is necessary to fully understand coupled fields. However, the largest barrier to fully displaying electric fields in time and space is a lack of intuitive methods for plotting their phase. Indeed, the impact of phase is often difficult to appreciate even when it is plotted in a straightforward way.

As a result, we here report visualization methods that rely on indirect, but more intuitive, displays of the phase in order to plot pulses. Specifically, we display the temporal phase by color-coding the distribution of spectral energy over time and space, and we display the spatial phase by propagating the electric field and showing how the beam evolves in space and time. Showing the impact of spatial and temporal phase on the pulse rather than the numerical value of the phase is the key to intuitively displaying the electric field. This approach is quite general and can be applied to arbitrary pulses, provided that the electric field is known at the input plane.

We believe that our work on pulse visualization provides an unprecedented opportunity to study and understand spatiotemporal couplings in ultrashort pulses, particularly for the elusive phase couplings. Some aspects of these visualization methods have been briefly described previously, as we have applied them to visualizing full spatiotemporal measurements of interesting pulses [2,3]. Here, we will describe in detail our most recent approach to plotting ultrashort pulses before applying these methods to displaying distorted pulses.

2. VISUALIZING PULSES

The goal of our pulse visualization technique is to be able to view a pulse as it would appear naturally (if one could freeze time and study it with a magnifying glass)—as a three-dimensional glob of

light, with all of its spatial, temporal, and spectral structures readily apparent. If the electric field $E(x, y, t)$ is known in time and space at some plane, $z = z_0$, it is straightforward to plot three-dimensional surfaces of constant intensity in x, y, and t to display the spatiotemporal amplitude. In addition, we color-code the spectral energy distribution to show the impact of the temporal phase, and we propagate the pulse to show the impact of the spatial phase. From the propagated field, we can make a movie of the pulse evolving or simply show several z planes in a still image. Figure 1 shows a fairly complex example pulse with interesting temporal and spatial phases. This pulse has a cubic spectral phase, resulting in the middle (green) frequencies arriving first, followed by beating between concurrent large (blue) and small (red) frequencies. As a result, the leading (right) edge of the pulse is green, which fades first to white (indicating the presence of the whole bandwidth) and then to purple. Surfaces of constant intensity show that the satellites are much weaker than the main pulse. This pulse is also diverging and experiences a small amount of angular dispersion. The same pulse plotted after 12 mm of free-space propagation is larger in space, and the red and blue colors are beginning to spatially separate. The angular dispersion also causes the pulse at $z = 0$ to be slightly tilted. Using color to indicate the temporal phase and propagation to indicate spatial phase allows a large amount of information to be condensed into a single, informative plot. The visualization code used to generate the plots in this paper can be found online [4,5].

A. Plotting the Temporal Phase

It is fairly common for the temporal phase of pulses to be represented in color using its first derivative, the instantaneous frequency. However, because the instantaneous frequency only represents the average frequency at a given time, it cannot indicate the bandwidth or spectral energy distribution of the pulse. Interestingly, the pulse in Fig. 1 has a constant instantaneous frequency. This ambiguity is particularly troubling for plotting spatiotemporal distortions, where the local bandwidth and degree of color separation are very interesting. If there is no distinction between having all frequencies present and only having the center (average) frequency present, it will be very difficult to interpret the plotted electric field.

To represent the temporal phase of pulses more meaningfully, we will rely on a more indirect representation of the phase. The spectrogram (or sonogram), also known as the short-time Fourier transform, has been used for many years, specifically for its strength in showing the distribution of the signal energy in time and frequency [6]. The first proponent of the spectrogram referred to it as the "physical spectrum" because it intuitively displays the time-frequency characteristics of various signals [7]. A spectrogram is defined as

$$Sp(x, y, \tau, \omega) = \left| \int_{-\infty}^{+\infty} E(x, y, t) g(t - \tau) \exp(-i\omega t) \mathrm{d}t \right|^2. \quad (1)$$

As we will discuss in detail in the next section, a spectrogram depends on the gate function $g(t)$ used to construct it and is not unique for a given signal $E(t)$. Because the spectrogram describes the distribution of spectral energy over time, it is ideal for this application.

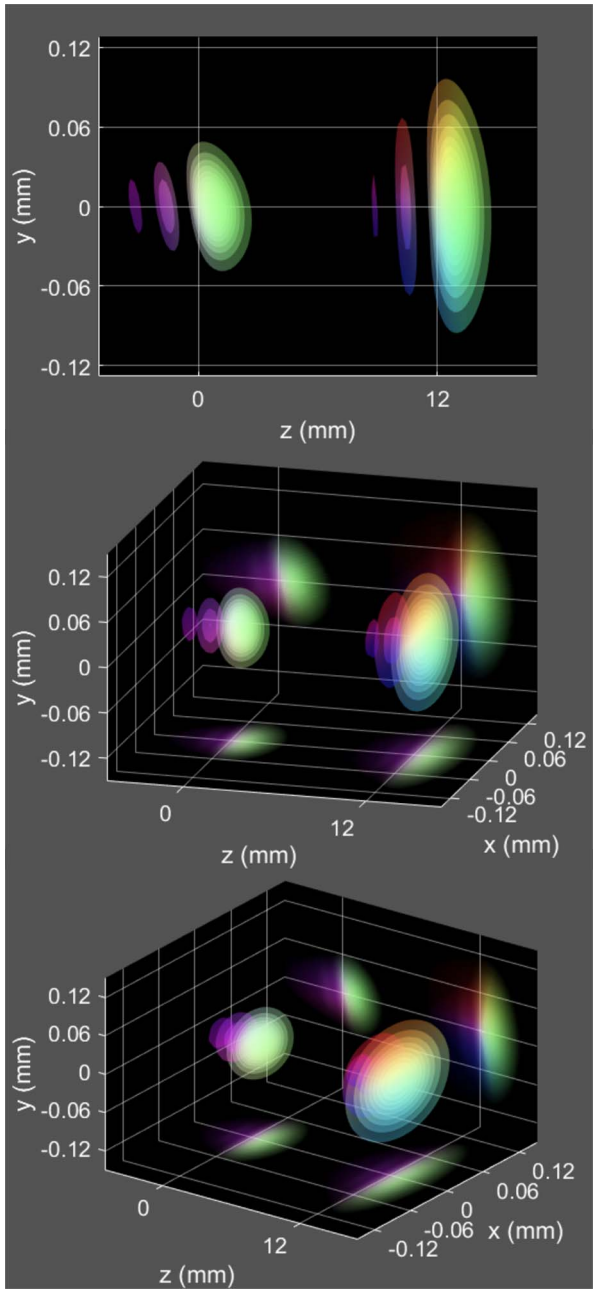

Fig. 1. Diverging pulse with cubic spectral phase and angular dispersion. The pulse is shown as it would appear if it could be photographed at two different z-positions while propagating to the right. Surfaces of constant intensity are plotted every 10% of the peak intensity at each z-distance. The cubic spectral phase causes the green light to arrive first, followed by red and blue light. After propagating ~12 mm, angular dispersion causes the red and blue light to spatially separate, as indicated by a red tint near the top of the pulse and a blue tint near the bottom. (top) Side view of the pulse, highlighting the temporal structure. (middle) Angled view of the pulse, showing three-dimensional capability (see Visualization 1). The electric field is projected onto the bottom and back walls to show the x and y profiles. (bottom) Different angled view of the same pulse, emphasizing the projections. Note that the apparent shape of the three-dimensional pulse changes when viewed from a different angle.

To color-code the spectrogram for display purposes, we define red, green, and blue response functions $R(\omega)$, $G(\omega)$, and $B(\omega)$, centered at the lower, middle, and higher frequencies, respectively, in the pulse spectrum. We then use overlap integrals to compute how much red, green, and blue energy is present at each time and spatial position in the pulse:

$$R(x,y,\tau) = \int_{-\infty}^{+\infty} Sp(x,y,\tau,\omega)R(\omega)\mathrm{d}\omega, \quad (2)$$

$$G(x,y,\tau) = \int_{-\infty}^{+\infty} Sp(x,y,\tau,\omega)G(\omega)\mathrm{d}\omega, \quad (3)$$

$$B(x,y,\tau) = \int_{-\infty}^{+\infty} Sp(x,y,\tau,\omega)B(\omega)\mathrm{d}\omega. \quad (4)$$

These functions are scaled afterwards such that each color function has the same total energy, normalizing the colors to the pulse spectrum. As a result, the pulse will appear white when all the frequencies in its spectrum are present. A highly saturated color indicates a narrow local spectrum. These color functions are ideal for showing how the spectral characteristics of the pulse's spatial profile change over time.

Without further adjustments, the brightness of the color also indicates the pulse intensity. Intuitively, times and positions that contain no pulse energy would be shown as black. While this is useful for plotting two-dimensional cross sections [2], three-dimensional plotting does not benefit from having low-intensity areas appear nearly black. Because surfaces of constant intensity already show the intensity distribution, the color functions only need to show the local spectrum. The final step in determining the plotted color is to boost the brightness of each voxel while maintaining the ratio of red, green, and blue in that voxel. This step allows us to plot three-dimensional surfaces at low intensities without making the pulse appear overly dark.

1. Choosing Gates and Color Responses

The degree of color saturation in the plotted pulse is determined by two factors. The choice of the spectrogram gate pulse plays a very large role. Shorter gate pulses generate spectrograms with better temporal resolution but poorer spectral resolution. The spectrogram signal at each delay is shorter in time and consequently more broadband. Given the same spectral response functions, a spectrogram created with a shorter gate will result in lower color saturation and a whiter-looking pulse. Conversely, choosing a longer gate pulse—implying poorer temporal resolution in the spectrogram—results in better spectral resolution and hence higher color saturation.

Generally speaking, the most informative results come from achieving an optimal balance between the temporal and spectral resolutions of the spectrogram. Often, the best compromise for ultrashort pulses is to choose a gate pulse close to the same temporal duration as the pulse itself [8]. However, in some cases, it may be interesting to emphasize either good spectral resolution or good temporal resolution. When the signal has variations in instantaneous frequency that must be resolved, the optimal gate length scales inversely with the rate of change of the instantaneous frequency [9]. In other words, the more complicated and quickly varying the signal is in time, the shorter the gate pulse must be to display these variations.

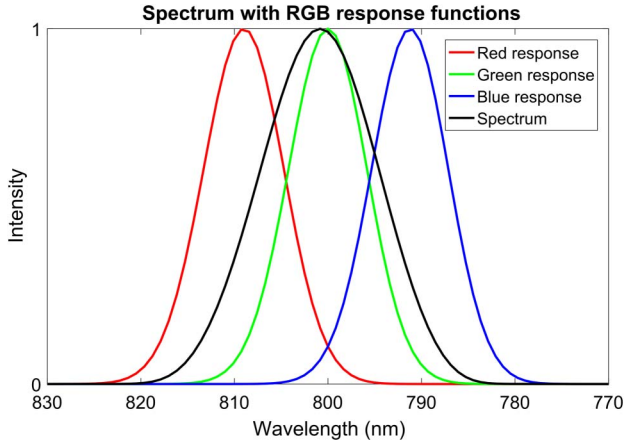

Fig. 2. Example choice of RGB response functions with respect to the pulse spectrum (black). This shows the target global spectrum for the coupled pulses plotted in this paper and the corresponding response functions used. Note that RGB function widths and their separations determine the appropriate overlapping with the pulse spectrum (and with each other).

The shortest gate pulse (a temporal δ-function) provides the best temporal resolution but always gives an infinitely wide spectrum, resulting in a plot that shows no spectral energy variations. The arrival time of spatial structures can be precisely shown, but the pulse will appear white at all times. The longest gate pulse (a continuous-wave signal) gives a perfect frequency resolution, but all the arrival time information is lost. The spatial distribution of the colors will be very well defined, but it will not evolve in time. Both of these cases result in the loss of information in the plots. It is always important to understand that the choice of gate function has a strong impact on displaying frequency variations. Color-response functions cannot recover frequency resolution that has been lost in making a spectrogram with too short a temporal gate function.

While the width and placement of spectral response functions certainly play a role in the color saturation of the plots, the choice of these functions is somewhat less critical due to normalization. There should be some overlap between the functions chosen (see Fig. 2) so that there are no areas of zero or near-zero response within the spectrum of the pulse. Otherwise, there may be unnatural-looking divisions of color in the plots or regions that look monochromatic despite actual presence of many different frequencies. If there is a large overlap between the response functions, then the color differentiation will be weak, and the pulse will simply appear white. Similar to choosing a gate pulse, there are trade-offs and compromises in choosing spectral response functions depending on the plotting purposes.

B. Plotting the Spatial Phase

In order to display the temporal phase, we explored its relationship to the local spectrum of the pulse. Similarly, to plot the spatial phase, we explore the Fourier conjugate domain. The spatial frequencies k_x and k_y are equivalent to the off-axis propagation vectors of the pulse. This strongly suggests that the best way to understand the spatial phase is to show how

the pulse changes as it propagates. This requires computing the free-space propagation of the electric field.

We use the angular-spectrum description of diffraction [10,11], which has been shown to be equivalent to the first Rayleigh–Sommerfeld solution [12]. After computing the two-dimensional spatial Fourier transform of the original electric field, the field at any z plane is given simply by

$$E(k_x,k_y,\omega,z)=E(k_x,k_y,\omega,0)\exp\left(iz\sqrt{\frac{\omega^2}{c^2}-k_x^2-k_y^2}\right), \quad (5)$$

where k_x and k_y are the spatial frequencies or off-axis k-vectors in the x and y dimensions. Using this model, propagation is described as a simple transfer function. $E(x, y, t, z)$ is easily determined for plotting purposes by three inverse Fourier transforms. We plot the field as a function of the local time around a given value of z, so only one set of inverse transforms is required to plot the field at each z plane.

Of course, when using angular spectrum propagation, care must be taken to avoid aliasing the transfer function at large values of k and z. In order to prevent aliasing from introducing errors into the propagated field, we simply zero out the transfer function for values of k_x and k_y above the Nyquist sampling limit $k_{\max}$ (as suggested by Ref. [13]):

$$k_{\max}=\left(\lambda\sqrt{(2\Delta kz)^2+1}\right)^{-1}. \quad (6)$$

We also take the additional step of checking how many discrete points lie between $-k_{\max}$ and $k_{\max}$. If too few points are in the non-zero region, the spatial grid is expanded to increase the spatial frequency sampling. These steps are sufficient to avoid distortions in the localized fields and short distances of interest to us.

The propagated electric field can be plotted at several values of z to show the effects of the spatial phase on the pulse evolution. While plotting a few planes in z is often sufficient to communicate the relevant phase information, we can also create animations and movies showing the evolution of the pulse shape through many planes. These movies may be necessary to understand very complex beams.

1. Beam Shape and Time Scaling

There is a final, subtle issue in plotting these fields that deserves some discussion. Ultrashort pulses are short in time and therefore are longitudinally small along their propagation direction. For example, a 500 fs pulse is only 150 μm long along its propagation direction. Because we often would like to show the evolution of a pulse over tens of millimeters or more, the longitudinal length of such a pulse must be increased so that its temporal structure will be visible in the plot. In many cases, the implicit local time axis can be scaled by any amount that conveys the desired information. However, the topic of this publication inspires an additional constraint. Because we are interested in displaying tilted pulses and other spatiotemporal distortions, we choose to scale time such that the angle of any tilted pulses will be faithfully represented. In other words, we require

$$\tan\ \alpha=\frac{c\Delta t_{\text{real}}}{\Delta x_{\text{real}}}=\frac{\Delta z_{\text{plotted}}}{\Delta x_{\text{plotted}}}, \quad (7)$$

where α is the tilt angle of the pulse. If we define scales s_x and s_z for the x- and z-axes (in μm/pixel or similar units), then we can relate the real and plotted quantities:

$$\Delta x_{\text{real}}=s_x\Delta x_{\text{plotted}}, \quad (8)$$

$$\gamma c\Delta t_{\text{real}}=s_z\Delta z_{\text{plotted}}. \quad (9)$$

Here, γ is the dimensionless scale factor for time that we desire. Plugging it into Eq. (7) and solving for γ, we find

$$\gamma=\frac{s_z}{s_x}. \quad (10)$$

Because the z-axis usually has a much larger range of values than the x-axis, γ is often fairly large.

We use the same axis scales and ranges for all plots in this paper, and the temporal scaling is therefore consistent as well. While this type of scaling may not be possible or desirable for all pulse length ranges and propagation distances, in this case, we believe that it is a useful and informative choice.

3. REVIEW OF SPATIOTEMPORAL COUPLINGS

Among all spatiotemporal couplings, the most commonly observed are the first-order spatiotemporal couplings [14–18]. These couplings can be described by a set of Gaussian equations in four Fourier-transform-conjugate domains comprised of pairs of dimensions [1,19]. In these equations, the coupling coefficients are complex, with the real and imaginary parts having different effects on the pulse. The real coefficients are called amplitude coupling terms because they affect the energy distribution, and the imaginary parts are called phase coupling terms because they affect the phase fronts of the beam and the relative phases of the frequency components.

Starting from the space and time dimensions (in the spatiotemporal domain), the expression for a first-order coupled field is

$$E(x,t)\propto\exp\{Q_{xx}x^2+2Q_{xt}xt-Q_{tt}t^2\}. \quad (11)$$

For simplicity, we keep our discussion of couplings in the x and t dimensions, so there is no need to explicitly write the y-dependence. For this expression to describe a Gaussian pulse, we require the real part of Q_{xx} to be negative and the real part of Q_{tt} to be positive. Otherwise, the field will diverge in space or time. The real part of Q_{xt} is referred to as the pulse-front tilt (PFT), and it can cause the energy on one side of the beam to arrive earlier than on the other side. The imaginary part is called the wave-front rotation, which describes how the direction of the phase fronts changes with time in the pulse.

We can express the same electric field in terms of different variables by taking the Fourier transform of one or both dimensions. Because Gaussians Fourier-transform to Gaussians, we will be able to write similar expressions in other domains. Taking the Fourier transform of Eq. (11) with respect to time leads us to the space-frequency (or spatiospectral) domain, where the field is

$$E(x,\omega)\propto\exp\{R_{xx}x^2+2R_{x\omega}x\omega-R_{\omega\omega}\omega^2\}. \quad (12)$$

Here, the real part of the coupling term $R_{x\omega}$ is called the spatial chirp, which separates the frequencies in the pulse spatially. The imaginary part causes different frequencies in the pulse to have

differently tilted phase fronts, an effect called wave-front tilt dispersion.

If we instead take one Fourier transform with respect to space instead of time, we obtain an expression in terms of time and spatial frequency. The spatial frequency k (also known as the off-axis propagation vector k_x) is related to the spatial shape of the pulse and to the phase fronts which determine the propagation direction. The field in the $k-t$ domain is then given by

$$E(k,t) \propto \exp\{P_{kk}k^2 + 2P_{kt}kt - P_{tt}t^2\}. \quad (13)$$

The real part of P_{kt} is called time versus angle, or the ultrafast lighthouse effect. It causes the early parts of the pulse to propagate in a different direction from the later parts. The imaginary part of P_{kt}, referred to as the angular temporal chirp, is somewhat less intuitive. It means that different propagation directions have different relative phases as time evolves.

The fourth domain is the frequency and spatial frequency domain. In this domain, the electric field is

$$E(k,\omega) \propto \exp\{S_{kk}k^2 + 2S_{k\omega}k\omega - S_{\omega\omega}\omega^2\}. \quad (14)$$

The amplitude coupling between the frequency and spatial frequency is angular dispersion, a well-known effect produced by prisms and other optics. Different frequencies propagate in different directions. The phase coupling term, called the angular spectral chirp, is less well known. This term causes different propagation directions to gain different relative phases for different frequencies.

Since all of these domains are inter-related by Fourier transforms, a coupling in any of these domains means that there are related couplings in other domains. Using simulations and pulse visualization, we will explore how each of these real and imaginary couplings affect the appearance of the pulse.

4. SIMULATION DETAILS

All simulated pulses are based on the same Gaussian pulse in time and space. This pulse has a full width at half-maximum (FWHM) temporal length of 64 fs and an FWHM beam width of 94 μm [see Fig 3(a)]. Its spectrum is centered at 800 nm. A gate pulse with a temporal FWHM length of 67 fs is used to generate the spectrograms used in coloring the plots. To show how the couplings interact with spectral chirp, 1000 fs^2 of chirp is added in some cases. To show how the couplings interact with spatial phase, a radius of curvature of 15 mm (focusing) is added in some cases. Figure 3(b) shows both chirp and spatial phase added to the standard pulse. The intensity

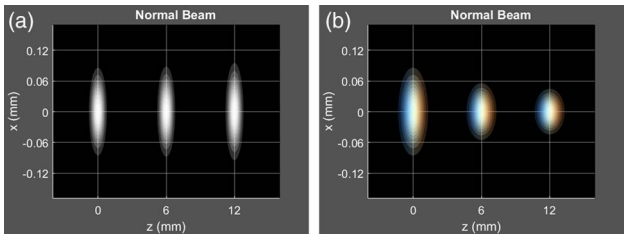

Fig. 3. (a) Collimated, unchirped beam with no distortions. (b) Focusing, chirped beam with no spatiotemporal distortions. The pulse is longer in time as a result of the chirp.

profile and phase parameters are translated into each domain, and Eqs. (11)–(14) are used to define the coupled electric fields.

We introduce spatiotemporal couplings that are normalized to the pulse parameters. As in Ref. [1], we define the amplitude couplings in the form

$$\mathrm{Re}\{R_{x\omega}\} = \rho\sqrt{-\mathrm{Re}\{R_{xx}\}\mathrm{Re}\{R_{\omega\omega}\}}, \quad (15)$$

where ρ is a constant that determines how strong the coupling is. For valid electric fields, ρ is between −1 and 1. We use $\rho = 0.9$ for the real couplings plotted in this paper.

Curiously, there are no such rules for the strength of imaginary couplings. Attempting to derive restrictions on imaginary couplings based on their relations to real couplings shows that any value for an imaginary coupling will produce real couplings with values of ρ between −1 and 1 (see Appendix A). The relationships between couplings in adjacent domains (separated by one Fourier transform) follow the general form [1]

$$Q_{xt} = \frac{i}{2}\frac{R_{x\omega}}{R_{\omega\omega}}. \quad (16)$$

Because of this, we choose the imaginary couplings to have the form

$$\mathrm{Im}\{R_{x\omega}\} = 2(0.9)\sqrt{-\mathrm{Re}\{R_{xx}\}\mathrm{Re}\{R_{\omega\omega}\}}. \quad (17)$$

This choice produces imaginary couplings with similar effects to real couplings in many cases.

As discussed in Ref. [1], introducing real couplings causes global pulse characteristics to differ from local ones. For example, introducing an $R_{x\omega}$ (spatial chirp) term causes the global bandwidth to be larger than the local bandwidth and the global beam width to be larger than the local beam width. The relationship between the local bandwidth $\Delta\omega_L$ and the global bandwidth $\Delta\omega_G$ is

$$\Delta\omega_L = \Delta\omega_G\sqrt{1-\rho^2}. \quad (18)$$

Interestingly, it is the local bandwidth that is directly related to the spectral profile in the pulse expression:

$$\Delta\omega_L = 1/\left(2\sqrt{\mathrm{Re}\{R_{\omega\omega}\}}\right). \quad (19)$$

This means that when introducing spatial chirp into a pulse equation, the overall bandwidth of the pulse will become larger if the local bandwidth (the $R_{\omega\omega}$ term) remains the same. While this makes sense from a mathematical perspective, it is at odds with our physical capabilities in manipulating ultrashort pulses. Because only nonlinear effects can introduce new frequencies into a pulse, we choose to hold the global bandwidth constant when introducing couplings. This means reducing the local bandwidth when adding couplings in the frequency domain, instead of expanding the global bandwidth. Consequently, $R_{\omega\omega}$ and $S_{\omega\omega}$ must change when real couplings are added to the pulse. When considering real couplings in the time domain, the same physically motivated constraints produce opposite effects. Holding the bandwidth of a pulse constant means that the local pulse duration should remain the same, because introducing shorter temporal features increases the bandwidth. This means that Q_{tt} and P_{tt} remain the same when adding real couplings, and the global pulse duration increases. For

further visual consistency, similar adjustments are applied in the x domain to keep the global beam width constant.

In order to make the effects of the spatiotemporal distortions as clear as possible, we consider only one transverse coordinate, x or k_x. The y or k_y dependence is kept as a simple Gaussian. We choose to plot a two-dimensional side view of the pulses to most clearly show the relationship between the spatial and temporal structures. While these pulses can easily be shown in three dimensions using our visualization method, the side view is more helpful for seeing the effects we are attempting to study.

We begin our discussion of spatiotemporal couplings with the x–ω domain, where the simulation results are the most intuitive. From there, we will explore similarities and differences between the spatiospectral domain couplings and couplings in other domains.

5. $R_{x\omega}$: SPATIAL CHIRP AND WAVE-FRONT TILT DISPERSION

Figure 4(a) shows a spatially chirped pulse. Comparing this pulse to Fig. 3(a), we see that spatially separating colors to this degree significantly increases the temporal pulse length. A second interesting difference (that we will see consistently for all coupled beams) is that the spatially chirped pulse does not stay collimated like the undistorted beam in Fig. 3(a) does. As the pulse propagates, the local spot size of each color expands, resulting in less color separation and a shortening of the local pulse duration. Group-delay arguments predict the emergence of pulse-front tilt as the pulse propagates.

Adding spectral chirp [see Fig. 4(b)] has a predictable effect on the pulse, causing the red side of the beam to arrive earlier and the blue side to arrive later. Interestingly, the frequencies appear to be separated along a diagonal and not along strict horizontal lines, as one might expect from Ref. [20]. This effect can be understood by realizing that each horizontal slice has a range of frequencies present, and the chirp causes the lower frequencies in that range to arrive first.

The effect of wave-front tilt dispersion on the beam is very different from spatial chirp [see Fig. 5(a)]. A phase coupling between x and ω implies that the tilt of the phase fronts is different for each frequency. Because light propagates in a direction perpendicular to its phase fronts, the varying tilt for each frequency means that different frequencies propagate in different directions. This is very similar to the well-known effect of angular dispersion. Accordingly, it comes as no surprise that wave-front tilt dispersion appears very similar to what we would expect for angular dispersion, including introducing negative

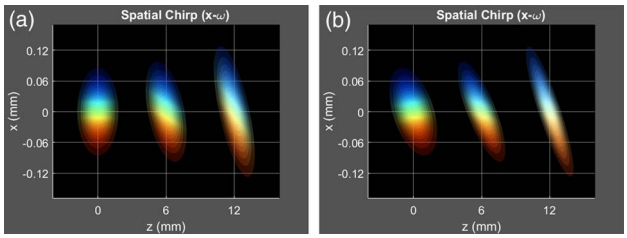

Fig. 4. Spatial chirp. (a) Spatial chirp in a collimated pulse with no spectral chirp (see Visualization 2). (b) Spatial chirp in a collimated pulse with spectral chirp.

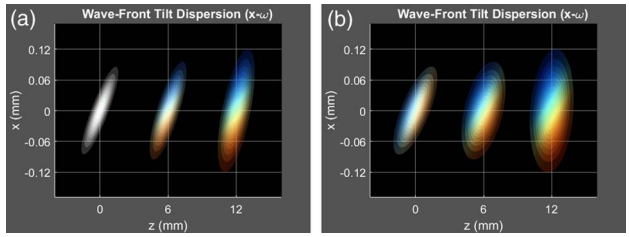

Fig. 5. Wave-front tilt dispersion. (a) Wave-front tilt dispersion in a collimated pulse with no spectral chirp. (b) Wave-front tilt dispersion in a collimated, chirped pulse.

group-delay dispersion in the pulse. In fact, we will see that for collimated beams with no chirp, the eight real and imaginary distortions can be grouped into two categories with nearly identical effects. One category contains pulse-front tilt, wave-front tilt dispersion, angular dispersion, and angular temporal chirp. The second category contains spatial chirp, wave-front rotation, time versus angle, and angular spectral chirp. These relationships may be inferred from the tables and calculations in Ref. [1], but they become abundantly clear when plotting the electric fields.

Adding spectral chirp to the pulse results in an identical temporal chirp across the whole beam [see Fig. 5(b)]. The applied chirp of 1000 fs^2 very nearly cancels out the negative group-delay dispersion generated by the coupling, and the dispersed colors consequently have a very similar arrival time after 12 mm of propagation.

6. Q_{xt}: PULSE-FRONT TILT AND WAVE-FRONT ROTATION

Since the x–t domain is related to the x–ω domain by a single temporal Fourier transform, there will be many similarities between couplings in the two domains. The mathematical relationship between the couplings is

$$Q_{xt} = \frac{i}{2}\frac{R_{x\omega}}{R_{\omega\omega}}, \qquad R_{x\omega} = -\frac{i}{2}\frac{Q_{xt}}{Q_{tt}}. \tag{20}$$

For pulses with no spectral or temporal chirp, $R_{\omega\omega}$ and Q_{tt} are real, and a real coupling in one domain corresponds to an imaginary coupling in the other. However, we see that introducing spectral or temporal chirp will break the symmetry between these two domains.

As expected, a beam with pulse-front tilt only [see Fig. 6(a)] is extremely similar to the plot in Fig. 5(a) of wave-front tilt dispersion. The most noticeable difference is that the pulse shape appears to be flipped about the x-axis, indicating that positive pulse-front tilt corresponds to negative wave-front tilt dispersion. Introducing chirp has a particularly dramatic effect that differentiates this pulse significantly from the pulse with wave-front tilt dispersion and spectral chirp. The temporal chirp causes the pulse to be quite long in time. The highly saturated colors in the $z = 0$ pulse indicate that the bandwidth in each location is small.

Despite introducing the same amount of chirp as in Figs. 4(b) and 5(b), the pulse in Fig. 6(b) has much greater temporal separation between colors. This can be explained by considering

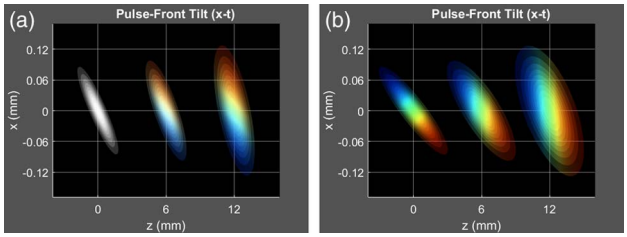

Fig. 6. Pulse-front tilt. (a) Pulse-front tilt in a collimated pulse with no temporal chirp (see Visualization 3). (b) Pulse-front tilt in a collimated pulse with temporal chirp.

local and global times. Looking at the algebraic expression for this pulse,

$$E(x,t) \propto \exp\left[-\frac{x^2}{4\Delta x_L^2}+\frac{2\rho x t}{4\Delta x_L \Delta t_L}-\frac{t^2}{4\Delta t_L^2}+i\beta t^2\right], \quad (21)$$

we see that we can rewrite the amplitude coupling term as an x-dependent modification of the local arrival time of the temporal pulse envelope:

$$E(x,t) \propto \exp\left[-\frac{(1-\rho^2)x^2}{4\Delta x_L^2}-\frac{1}{4\Delta t_L^2}\left(t-\rho x\frac{\Delta t_L}{\Delta x_L}\right)^2+i\beta t^2\right]. \quad (22)$$

However, the chirp term in this expression does not depend on x at all. As a result, the chirp uses a global pulse time, not a local one. Because the global pulse time has a larger range than the local pulse time, the temporal phase can have a large range of values. In fact, adding pulse-front tilt and temporal chirp to the field expression in this way actually increases the pulse bandwidth. The electric field can be modified so that the chirp also uses an x-dependent local time, resulting in a pulse that is nearly identical to Fig. 5(b). However, it is important to note that using a local chirp introduces an imaginary component in the Q_{xt} coupling term. The modified pulse would have both pulse-front tilt and wave-front rotation.

Another interesting aspect of Fig. 6(b) is the similarity between the $z=0$ pulse in this plot and the $z=0$ pulse in Fig. 4(b) of the spatial chirp with the spectral chirp. The pulse with spatial and spectral chirps appears to have some pulse-front tilt. Akturk *et al.* [1] derived an expression for the pulse-front tilt in terms of the spatiospectral coupling terms (see Appendix B):

$$\frac{\mathrm{Re}\{Q_{xt}\}}{\mathrm{Re}\{Q_{tt}\}} = -2\,\mathrm{Im}\{R_{x\omega}\}+2\,\mathrm{Im}\{R_{\omega\omega}\}\frac{\mathrm{Re}\{R_{x\omega}\}}{\mathrm{Re}\{R_{\omega\omega}\}}. \quad (23)$$

Here, the pulse-front tilt is normalized by the local pulse duration, and the spatial chirp is normalized by the local bandwidth. Accordingly, one would say that pulse-front tilt is caused by wave-front tilt dispersion and also by the combination of the spectral chirp and spatial chirp. For positively chirped pulses, the imaginary part of $R_{\omega\omega}$ is positive, and so pulse-front tilt and spatial chirp have the same sign, while pulse-front tilt and wave-front tilt dispersion have opposite signs. The unchirped pulse with pulse-front tilt [Fig. 6(a)] has all of its tilt generated by wave-front tilt dispersion. The opposite signs are evident in the opposite tilt angles. The pulse with spatial and spectral chirp [Fig. 4(b)] has no wave-front tilt dispersion, and all of its tilt therefore comes from the second term in Eq. (23). The pulse in Fig. 6(b) experiences pulse-front tilt generated by both effects. The lack of wave-front tilt dispersion in the pulse with spatial and spectral chirp [Fig. 4(b)] explains why this pulse looks very different from the pulse with pulse-front tilt and temporal chirp [Fig. 6(b)] after the $z=0$ plane. The divergent evolution of these pulses also underscores the need to plot the pulse evolution as part of understanding the electric field.

The unusual appearance of the temporally chirped pulse with pulse-front tilt has interesting implications for experiments involving pulse-front tilt. One common method of introducing pulse-front tilt is to interact with a diffraction grating at a non-zero incidence angle [21–23]. A chirped pulse diffracting off a grating would retain the same temporal profile for all transverse positions. The resulting beam should therefore look like the wave-front tilt dispersion in Fig. 5(b), not the pulse-front tilt in Fig. 6(b). The plots therefore imply that a diffraction grating primarily introduces wave-front tilt dispersion, and a chirped pulse interacting with a grating therefore experiences wave-front rotation as well as pulse-front tilt. These plots allow us to cultivate a much better understanding of how to describe real-world couplings using equations.

Another experimental insight that can be gained from Fig. 6 relates to the pulse tilt angle. The pulse-front tilt angle is the natural way to quantify the pulse-front tilt in many situations. One would expect the tilt angle to be essentially equivalent to the strength of the coupling between x and t. However, both pulses in Fig. 6 have the same coupling strength, and the chirped pulse is tilted at a larger angle than the unchirped pulse. The tilt angle therefore depends on more than just the coupling strength, and it would be useful to have an expression relating the pulse-front tilt angle to the parameters used in the Gaussian beam equations. If we define the pulse-front tilt angle α as

$$\tan\,\alpha = \frac{\Delta z}{\Delta x}, \quad (24)$$

then it can be shown that the tilt angle depends on the Gaussian pulse parameters as

$$\tan\,\alpha = c\frac{\mathrm{Re}\{Q_{xt}\}}{\mathrm{Re}\{Q_{tt}\}} = \frac{c\rho\Delta t_L}{\Delta x_L} = \frac{c\rho\Delta t_G}{\Delta x_G}. \quad (25)$$

It is interesting to note that the tilt angle depends on both the coupling strength ρ and the ratio of the pulse duration to the beam size. Mathematically, as the local pulse duration Δt_L increases due to chirp, the pulse-front tilt angle also increases. In contrast, in the experimental grating geometry often used to create pulse-front tilt, the tilt angle of the pulse bouncing off of the grating depends only on the angle of incidence onto the grating. Physically, the tilt angle remains constant regardless of the pulse length or chirp. This means that mathematically describing pulse-front tilt generated by a grating requires changing the coupling strength as the pulse length changes, even though the experimental apparatus is fixed. Such details become much more apparent when the theoretical and experimental spatiotemporal fields can be directly plotted.

The effect of wave-front rotation on a collimated, unchirped pulse is shown in Fig. 7(a). As anticipated, the pulse is nearly

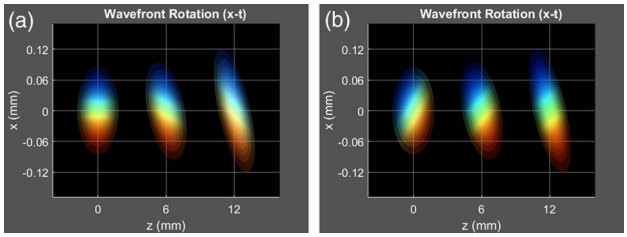

Fig. 7. Wave-front rotation (a) Wave-front rotation in a collimated beam with no temporal chirp. (b) Wave-front rotation in a collimated beam with temporal chirp.

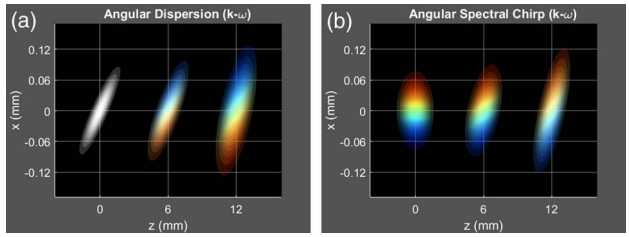

Fig. 8. k–ω domain couplings: angular dispersion and angular spectral chirp. (a) Angular dispersion in a collimated pulse with no spectral chirp. (b) Angular spectral chirp in a collimated pulse with no spectral chirp.

identical to a pulse with spatial chirp. Unlike the spatially chirped pulse, however, this pulse does not gain pulse-front tilt when temporal chirp is added. Wave-front rotation is a phase coupling, and it therefore cannot change the spatial or temporal intensity profile of the pulse, even in the presence of temporal chirp.

The addition of temporal chirp means that the pulse in Fig. 7(b) has both spatial chirp and wave-front tilt dispersion. The distribution of colors in this plot becomes much more intuitive when considering the sum of the color profiles for spectrally chirped pulses with spatial chirp and wave-front tilt dispersion. Similar to the pulse with spectral chirp and wave-front tilt dispersion, this pulse has a stronger temporal separation of colors than the pulse with spectral chirp and spatial chirp. As the beam expands due to propagation, the spatial chirp becomes more pronounced.

Plots of pulse-front tilt and wave-front rotation show that introducing temporal phase into the field expressions creates pulses that have a mixture of spatial chirp and wave-front tilt dispersion. We will see that this rule of thumb can be applied generally when adding phase in the dimension not shared by the two domains.

7. $S_{k\omega}$: ANGULAR DISPERSION AND ANGULAR SPECTRAL CHIRP

The $\boldsymbol{k}$–ω domain is related to the x–ω domain by a single spatial Fourier transform, and their couplings are therefore related by

$$S_{k\omega} = \frac{i}{2}\frac{R_{x\omega}}{R_{xx}}, \qquad R_{x\omega} = -\frac{i}{2}\frac{S_{k\omega}}{S_{kk}}. \tag{26}$$

When the pulse is collimated, R_{xx} and S_{kk} are real. In this case, the angular dispersion is proportional to the wave-front tilt dispersion, and the angular spectral chirp is proportional to the spatial chirp (see Fig. 8). In these cases, the effects of angular dispersion and angular spectral chirp are identical to effects we have already considered. Positive angular spectral chirp corresponds to negative spatial chirp, but the two effects are otherwise the same.

The previous section explored symmetry breaking between the x–ω and x–t domains due to chirp. The symmetry between x–ω and k–ω is broken by spatial phase, including focusing or defocusing of the beam. For example, wave-front tilt dispersion and angular dispersion have the same effect on collimated pulses even in the presence of spectral chirp. If there is no spatial phase, a pulse with only angular dispersion (and no angular spectral chirp) has only wave-front tilt dispersion and no spatial chirp. However, adding spatial phase to the pulse expressions creates very different electric fields (see Fig. 9). While the pulse with wave-front tilt dispersion focuses somewhat as expected, the pulse with angular dispersion becomes very strongly dispersed in the transverse dimension. Adding spatial phase to a pulse with angular dispersion creates a beam that has both wave-front tilt dispersion and spatial chirp.

Similarly to our discussion of pulse-front tilt, we can derive an expression for the angular dispersion in terms of the x–ω couplings:

$$\frac{\mathrm{Re}\{S_{k\omega}\}}{\mathrm{Re}\{S_{kk}\}} = -2\mathrm{Im}\{R_{x\omega}\} + 2\mathrm{Im}\{R_{xx}\}\frac{\mathrm{Re}\{R_{x\omega}\}}{\mathrm{Re}\{R_{xx}\}}. \tag{27}$$

Here, the angular dispersion is normalized by the beam divergence, and the spatial chirp is normalized by the beam size. Angular dispersion can therefore be caused by wave-front tilt dispersion and by the combination of wave-front curvature and spatial chirp. The real parts of S_{kk} and R_{xx} are always negative, and the wave-front curvature is negative for focusing beams, so the angular dispersion and wave-front tilt dispersion have the same sign and spatial chirp has an opposite sign. Figure 5(a) [or equivalently, Fig. 8(a)] shows the wave-front tilt dispersion contribution to the angular dispersion. Figure 10(a) shows the other contribution: a focusing pulse with spatial chirp. The opposite signs of angular dispersion and spatial chirp in Eq. (27) are evident in the opposite color gradients in Figs. 9(b) and 10(a). Because wave-front tilt dispersion and spatial chirp belong to the same domain, the pulse with wave-front tilt dispersion has zero spatial chirp and vice versa. The focusing

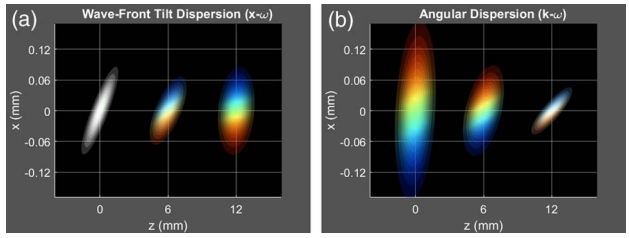

Fig. 9. Differences between wave-front tilt dispersion and angular dispersion. (a) Wave-front tilt dispersion in a focusing pulse with no chirp. (b) Angular dispersion in a focusing pulse with no chirp (see Visualization 4).

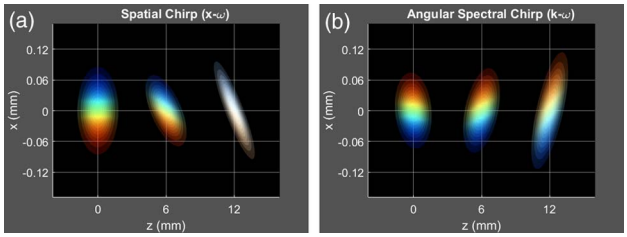

Fig. 10. Differences between spatial chirp and angular spectral chirp. (a) Spatial chirp in a focusing pulse with no spectral chirp. (b) Angular spectral chirp in a focusing pulse with no spectral chirp.

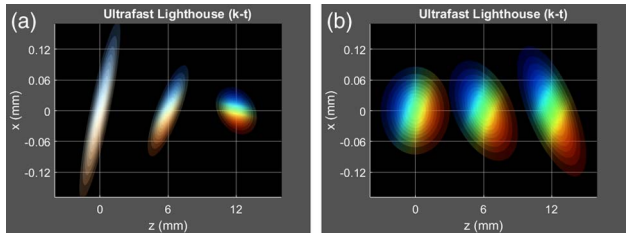

Fig. 11. Pulses with time versus angle, also known as the ultrafast lighthouse effect. (a) Time versus angle in a focusing beam with no temporal chirp (see Visualization 5). (b) Time versus angle in a collimated beam with temporal chirp.

pulse with angular dispersion [Fig. 9(b)] includes contributions to the total angular dispersion from both effects.

The large beam width of the focusing pulse with angular dispersion highlights a parallel with our consideration of pulse-front tilt. The combination of temporal phase and pulse-front tilt resulted in a pulse that was long in time and had a broader bandwidth [Fig. 6(b)]. In this case, adding spatial phase in combination with angular dispersion results in a pulse with a large convergence angle and a larger beam size. Although we attempted to keep the beam size consistent between all pulses considered by maintaining the same local beam divergence, we were not successful for pulses with spatial phase. This highlights the complicated relationships between beam parameters in the presence of spatiotemporal couplings and spatial and temporal phases.

Figure 10 shows the equivalent symmetry breaking for spatial chirp and angular spectral chirp. Without spatial phase, these two terms are equivalent (with a relative minus sign) and have identical effects on a pulse. Adding spatial phase differentiates them, although to a less dramatic degree than the previous case. The pulse with spatial chirp focuses more cleanly, while the pulse with angular spectral chirp still has significant color dispersion after 12 mm of propagation. This makes sense, because focusing wave-front tilt dispersion creates significant color separation [see Fig. 9(a)], and a focusing pulse with angular spectral chirp has both a spatial chirp and wave-front tilt dispersion.

We should explicitly note that spatial phase breaks the symmetry between the x domains and the k domains. This means that a focusing beam with pulse-front tilt looks just like a focusing beam with wave-front tilt dispersion (besides the relative minus sign), because these are x-domain couplings with similar effects. A focusing beam with wave-front rotation is the same as a focusing beam with spatial chirp for the same reasons. Likewise, the spectral phase differentiates the ω domains and the t domains, meaning that angular dispersion and angular spectral chirp respond to spectral phase in the same way as the wave-front tilt dispersion and spatial chirp, respectively. The reader is encouraged to refer to similar plots when a specific combination of coupling and phase has been omitted for brevity.

8. P_{kt}: TIME VERSUS ANGLE AND ANGULAR TEMPORAL CHIRP

The k–t domain is two Fourier transforms away from the x–ω domain, and so the relationship between the couplings in these domains is complicated:

$$P_{kt} = \frac{1}{4}\frac{R_{x\omega}}{R_{xx}R_{\omega\omega} + R_{x\omega}^2}. \tag{28}$$

However, we can leverage our understanding of other domains to predict the behavior of the k–t couplings.

For collimated beams with no chirp, time versus angle is identical to spatial chirp, angular spectral chirp (with a sign flip), and wave-front rotation. When focusing, time versus angle is best expressed as a combination of wave-front rotation and pulse-front tilt:

$$\frac{\mathrm{Re}\{P_{kt}\}}{\mathrm{Re}\{P_{kk}\}} = -2\,\mathrm{Im}\{Q_{xt}\} + 2\,\mathrm{Im}\{Q_{xx}\}\frac{\mathrm{Re}\{Q_{xt}\}}{\mathrm{Re}\{Q_{xx}\}}. \tag{29}$$

Because the real parts of P_{kk} and Q_{xx} are negative, and the wave-front curvature $\mathrm{Im}\{Q_{xx}\}$ is negative for focusing beams, time versus angle has the same sign as wave-front rotation and the opposite sign from pulse-front tilt. As with angular dispersion, the beam width becomes very large when spatial phase is added [see Fig. 11(a)]. The focusing term is large enough that the resulting pulse more closely resembles focusing pulse-front tilt [see Fig. 9(a) of focusing wave-front tilt dispersion] than wave-front rotation [Fig. 7(a)].

When adding spectral phase to a pulse with time versus angle [see Fig. 11(b)], the pulse behaves similarly to wave-front rotation [see Fig. 7(b)]. Alternatively, the field can be viewed as the combination of angular spectral chirp and angular dispersion:

$$\frac{\mathrm{Re}\{P_{kt}\}}{\mathrm{Re}\{P_{tt}\}} = -2\,\mathrm{Im}\{S_{k\omega}\} + 2\,\mathrm{Im}\{S_{\omega\omega}\}\frac{\mathrm{Re}\{S_{k\omega}\}}{\mathrm{Re}\{S_{\omega\omega}\}}. \tag{30}$$

For positively chirped pulses, the imaginary part of $S_{\omega\omega}$ is positive, meaning that time versus angle has the same sign as angular dispersion and the opposite sign from the angular spectral chirp. The second term, angular dispersion with spectral phase, is identical to wave-front tilt dispersion with spectral phase [Fig. 5(b)]. Both angular dispersion and angular spectral chirp appear to have a sizeable contribution to the pulse shape.

Figure 12 shows pulses with angular temporal chirp. Interestingly, the focusing pulse [Fig 12(a)] is not very different at all from a collimated beam [see Fig. 8(a) of angular dispersion]. Much like angular spectral chirp, adding spatial phase to this pulse has very little effect. It is not clear why the spatial phase has such a dramatic effect on the pulse shape for the amplitude couplings [see Fig. 8(b) of angular dispersion

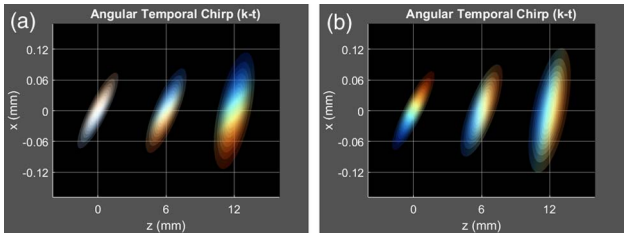

Fig. 12. Pulses with angular temporal chirp. (a) Angular temporal chirp in a focusing pulse with no temporal chirp. (b) Angular temporal chirp in a collimated pulse with temporal chirp.

and Fig. 11(a) of time versus angle] while having such a small impact on the phase couplings.

The combination of angular temporal chirp with temporal chirp is shown in Fig. 12(b). As we should expect, this pulse is very similar to the chirped pulse with pulse-front tilt [Fig. 6(b)], although the angle and colors are flipped spatially.

9. CONCLUSIONS

In general, it becomes clear that first-order phase couplings can be understood in terms of how they relate to amplitude couplings. Once one understands when two closely related couplings behave similarly and when they differ, it becomes simple to predict how a given phase coupling will affect an ultrashort pulse. In addition, visualizing a coupling as the sum of couplings in other domains is useful and informative.

We conclude that the described method for displaying the intensity and phase of ultrashort pulses should be extremely helpful in displaying and understanding both amplitude and phase variations. The ability to display the electric field of any pulse is an important step forward in our ability to understand and visualize distorted beams. This plotting technique is very general and can be applied to any localized electric fields of interest.

APPENDIX A: PROOF THAT IMAGINARY COUPLING TERMS CAN TAKE ANY VALUE

Assume that a pulse in the x–t domain has wave-front rotation and no pulse-front tilt ($\mathrm{Re}\{Q_{xt}\} = 0$). We will also assume that the pulse is collimated and not chirped ($\mathrm{Im}\{Q_{xx}\} = \mathrm{Im}\{Q_{tt}\} = 0$) (for simplicity only; the proof still holds for pulses with spatial and/or spectral phase). The same pulse considered in the x–ω domain has spatial chirp. The x–ω coefficients can be expressed in terms of the x–t coefficients [1]:

$$R_{xx} = Q_{xx} + \frac{Q_{xt}^2}{Q_{tt}} = \mathrm{Re}\{Q_{xx}\} + \frac{-\mathrm{Im}\{Q_{xt}\}^2}{\mathrm{Re}\{Q_{tt}\}}, \tag{A1}$$

$$R_{x\omega} = -\frac{i}{2}\frac{Q_{xt}}{Q_{tt}} = \frac{\mathrm{Im}\{Q_{xt}\}}{2\mathrm{Re}\{Q_{tt}\}}, \tag{A2}$$

$$R_{\omega\omega} = \frac{1}{4Q_{tt}} = \frac{1}{4\mathrm{Re}\{Q_{tt}\}}. \tag{A3}$$

The coupling strength ρ of the spatial chirp must be between −1 and 1 for valid electric fields:

$$\rho^2 = \frac{\mathrm{Re}\{R_{x\omega}\}^2}{-\mathrm{Re}\{R_{xx}\}\mathrm{Re}\{R_{\omega\omega}\}} < 1. \tag{A4}$$

By substituting Eqs. (A1)–(A3) into Eq. (A4), we obtain

$$\rho^2 = \frac{\frac{1}{4}\frac{\mathrm{Im}\{Q_{xt}\}^2}{\mathrm{Re}\{Q_{tt}\}^2}}{-\left(\mathrm{Re}\{Q_{xx}\} - \frac{\mathrm{Im}\{Q_{xt}\}^2}{\mathrm{Re}\{Q_{tt}\}}\right)\frac{1}{4\mathrm{Re}\{Q_{tt}\}}} = \frac{\frac{\mathrm{Im}\{Q_{xt}\}^2}{\mathrm{Re}\{Q_{tt}\}}}{-\mathrm{Re}\{Q_{xx}\} + \frac{\mathrm{Im}\{Q_{xt}\}^2}{\mathrm{Re}\{Q_{tt}\}}}. \tag{A5}$$

Keeping in mind that the real part of Q_{xx} is always negative (and the real part of Q_{tt} is always positive), it becomes clear that the expression in Eq. (A5) will always be less than one, regardless of the value of $\mathrm{Im}\{Q_{xt}\}$.

APPENDIX B: DERIVATION OF EQ. (23)

Using Eq. (A2), we can write the x – t coupling term as

$$Q_{xt} = 2i \qquad R_{x\omega}Q_{tt}. \tag{B1}$$

The pulse-front tilt is the real part of the coupling term, given by

$$\mathrm{Re}\{Q_{xt}\} = -2\,\mathrm{Im}\{R_{x\omega}\}\mathrm{Re}\{Q_{tt}\} - 2\,\mathrm{Re}\{R_{x\omega}\}\mathrm{Im}\{Q_{tt}\}. \tag{B2}$$

Using Eq. (A3), we find that

$$\frac{\mathrm{Im}\{R_{\omega\omega}\}}{\mathrm{Re}\{R_{\omega\omega}\}} = -\frac{\mathrm{Im}\{Q_{tt}\}}{\mathrm{Re}\{Q_{tt}\}}. \tag{B3}$$

Using this relationship, we then write

$$\frac{\mathrm{Re}\{Q_{xt}\}}{\mathrm{Re}\{Q_{tt}\}} = -2\,\mathrm{Im}\{R_{x\omega}\} + \mathrm{Re}\{R_{x\omega}\}\frac{\mathrm{Im}\{R_{\omega\omega}\}}{\mathrm{Re}\{R_{\omega\omega}\}}. \tag{B4}$$

Funding. National Science Foundation (NSF) (1307817); Georgia Research Alliance (GRA)

Acknowledgment. This work was performed in part under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

REFERENCES

1. S. Akturk, X. Gu, P. Gabolde, and R. Trebino, "The general theory of first-order spatio-temporal distortions of Gaussian pulses and beams," Opt. Express **13**, 8642–8661 (2005).
2. Z. Guang, M. Rhodes, M. Davis, and R. Trebino, "Complete characterization of a spatiotemporally complex pulse by an improved single-frame pulse-measurement technique," J. Opt. Soc. Am. B **31**, 2736–2743 (2014).
3. Z. Guang, M. Rhodes, and R. Trebino, "Measurement of the ultrafast lighthouse effect using a complete spatiotemporal pulse-characterization technique," J. Opt. Soc. Am. B **33**, 1955–1962 (2016).
4. M. Rhodes, "3D laser pulse visualization," https://doi.org/10.6084/m9.figshare.4719193.v1.
5. Trebino Group Website, http://frog.gatech.edu/code.html.
6. W. Koenig, H. Dunn, and L. Lacy, "The sound spectrograph," J. Acoust. Soc. Am. **18**, 19–49 (1946).
7. W. D. Mark, "Spectral analysis of the convolution and filtering of non-stationary stochastic processes," J. Sound Vib. **11**, 19–63 (1970).
8. R. Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Springer, 2012).
9. L. Cohen, *Time–Frequency Analysis* (Prentice-Hall, 1995).
10. J. Ratcliffe, "Some aspects of diffraction theory and their application to the ionosphere," Rep. Prog. Phys. **19**, 188–267 (1956).

11. J. W. Goodman, *Introduction to Fourier Optics* (Roberts & Company, 2005).
12. G. C. Sherman, "Application of the convolution theorem to Rayleigh's integral formulas," J. Opt. Soc. Am **57**, 546–547 (1967).
13. K. Matsushima and T. Shimobaba, "Band-limited angular spectrum method for numerical simulation of free-space propagation in far and near fields," Opt. Express **17**, 19662–19673 (2009).
14. H. Vincenti and F. Quéré, "Attosecond lighthouses: How to use spatiotemporally coupled light fields to generate isolated attosecond pulses," Phys. Rev. Lett. **108**, 113904 (2012).
15. J. A. Wheeler, A. Borot, S. Monchocé, H. Vincenti, A. Ricci, A. Malvache, R. Lopez-Martens, and F. Quéré, "Attosecond lighthouses from plasma mirrors," Nat. Photonics **6**, 829–833 (2012).
16. H. Vincenti, A. Borot, T. Hammond, K. T. Kim, J. Wheeler, C. Zhang, T. Ruchon, T. Auguste, J. Hergott, and D. Villeneuve, "Applications of ultrafast wavefront rotation in highly nonlinear optics," J. Phys. B **47**, 124004 (2014).
17. A. Zaukevičius, V. Jukna, R. Antipenkov, V. Martinėnaitė, A. Varanavičius, A. P. Piskarskas, and G. Valiulis, "Manifestation of spatial chirp in femtosecond noncollinear optical parametric chirped-pulse amplifier," J. Opt. Soc. Am. B **28**, 2902–2908 (2011).
18. G. Zhu, J. Van Howe, M. Durst, W. Zipfel, and C. Xu, "Simultaneous spatial and temporal focusing of femtosecond pulses," Opt. Express **13**, 2153–2159 (2005).
19. S. Akturk, X. Gu, P. Bowlan, and R. Trebino, "Spatio-temporal couplings in ultrashort laser pulses," J. Opt. **12**, 093001 (2010).
20. S. Akturk, X. Gu, E. Zeek, and R. Trebino, "Pulse-front tilt caused by spatial and temporal chirp," Opt. Express **12**, 4399–4410 (2004).
21. T. C. Wong, M. Rhodes, and R. Trebino, "Single-shot measurement of the complete temporal intensity and phase of supercontinuum," Optica **1**, 119–124 (2014).
22. T. C. Wong and R. Trebino, "Single-frame measurement of complex laser pulses tens of picoseconds long using pulse-front tilt in cross-correlation frequency-resolved optical gating," J. Opt. Soc. Am. B **30**, 2781–2786 (2013).
23. R. Wyatt and E. E. Marinero, "Versatile single-shot background-free pulse duration measurement technique, for pulses of subnanosecond to picosecond duration," Appl. Phys. **25**, 297–301 (1981).