Journal of Applied Physics TUTORIAL scitation.org/journal/jap

Highly reliable measurement of ultrashort laser pulses

Cite as: J. Appl. Phys. **128**, 171103 (2020); doi: 10.1063/5.0022552
Submitted: 26 July 2020 · Accepted: 10 October 2020 ·
Published Online: 5 November 2020

R. Trebino,[1] R. Jafari,[1] S. A. Akturk,[1] P. Bowlan,[2] Z. Guang,[1] P. Zhu,[1,3] E. Escoto,[4] and G. Steinmeyer[4,a)]

AFFILIATIONS

[1]School of Physics, Georgia Institute of Technology, 837 State St. NW, Atlanta, Georgia 30332, USA
[2]Physical Chemistry and Applied Spectroscopy, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA
[3]Shanghai Institute of Optics and Fine Mechanics, Chinese Academy of Sciences, Shanghai 201800, China
[4]Max Born Institute for Nonlinear Optics and Short Pulse Spectroscopy, Max-Born-Straße 2a, 12489 Berlin, Germany

[a)]**Author to whom correspondence should be addressed:** steinmey@mbi-berlin.de

ABSTRACT

The past 30 years have seen spectacular progress in the development of techniques for measuring the complete temporal field, and even the complete *spatiotemporal* field, of ultrashort laser pulses. The challenge has been to measure a pulse without the use of a shorter event or an independent known reference pulse, neither of which is typically available. We begin with autocorrelation, the first such "self-referenced" pulse-measurement method ever proposed, which measures only a rough pulse length, and we describe its limitations. One such limitation is the presence of a somewhat unintuitive "coherent artifact," which occurs for complicated pulses and also when averaging over a train of pulses whose shapes vary from pulse to pulse. We then describe the most important modern techniques capable of measuring the complete temporal intensity and phase of even complicated ultrashort pulses, as well as their ability (or inability) to measure such unstable pulse trains. A pulse reliably measured with such a device can then be used as a reference pulse in conjunction with another technique, such as spectral interferometry or holography, to measure pulses otherwise unmeasurable by a self-referenced technique. Examples include techniques for measuring low-intensity pulse(s) and for measuring the complete spatiotemporal intensity and phase of arbitrary pulse(s). This Tutorial is limited to well-established, proven methods, but other methods whose description proves instructive will be discussed.

I. INTRODUCTION

Ultrashort light pulses are the shortest events ever created. Pulses as short as tens of attoseconds (10^{-18} s) have been generated, and it is now routine to generate pulses less than 100 fs long. To use them effectively and to determine how to make them shorter, less structured, and more stable from pulse to pulse, it is important to be able to measure them.

This task seems particularly difficult because, to measure an event in time, it seems that one would need a shorter event with which to time it. For example, resolving the action of a bubble popping requires a strobe light with a flash shorter than the time it takes for the bubble to pop. Then, to measure the strobe light intensity vs time requires a detector with an even shorter response time. So one might reasonably conclude that precisely measuring the *shortest* event is *impossible*.

We will return to this dilemma later, but we must first ask precisely what it is that we are trying to measure about a pulse, and that is its electric field as a function of time and space, which can potentially be a complicated function of all four spatiotemporal coordinates. We will temporarily ignore the field's spatial dependence, assuming that the temporal quantities do not depend on the transverse position, and concentrate on the pulse's dependence on time only. We can then write the pulse electric field as

$$\mathcal{E}(t) = \frac{1}{2}\sqrt{I(t)}\exp\{i\,[\omega_0 t - \phi(t)]\} + c.c., \tag{1}$$

where t is the time in the reference frame of the pulse, ω_0 is a carrier angular frequency on the order of 10^{15} rad/s for visible and near-IR light, and $I(t)$ and $\phi(t)$ are the temporal intensity (often simply called the "intensity") and temporal phase (often simply called the "phase") of the pulse. We have ignored the proportionality constant and its corresponding units and so use these terms in their generic mathematical senses because we are only concerned with pulse *shapes* and not absolute magnitudes.

As usual, "*c.c.*" means *complex conjugate* and is required to yield a real pulse field. However, we will use the equivalent *analytic signal* representation and ignore the complex conjugate term, yielding a complex pulse field and simplifying the mathematics significantly. The *c.c.* can be re-added anytime. As a result, and because the center frequency, ω_0, is easy to measure using a spectrometer, the quantity generally desired in a measurement is the pulse *complex amplitude*,

$$E(t) \equiv \sqrt{I(t)}\exp[-i\phi(t)]. \tag{2}$$

The pulse field in the frequency domain, $\tilde{E}(\omega)$, is the Fourier transform of the time domain field and is usually separated into its spectral intensity $S(\omega)$ (often simply called the "spectrum") and spectral phase $\varphi(\omega)$,

$$\tilde{E}(\omega) = \sqrt{S(\omega)}\exp[-i\varphi(\omega)]. \tag{3}$$

Note that the temporal phase, ϕ, and spectral phase, φ, are both called "phi," but they are different (ϕ is a function of time and φ is a function of frequency), so we have used different Greek letter forms to distinguish them. Also, as is customary in optics, when we mean the spectral intensity or the spectral phase, we will always use the adjective "spectral." Otherwise, the temporal quantities are indicated. As with their temporal analogs, we are only concerned with their shapes and not their absolute magnitudes, so we ignore any proportionality constants. See Fig. 1 for a schematic drawing of a simple, common pulse and these quantities for it.

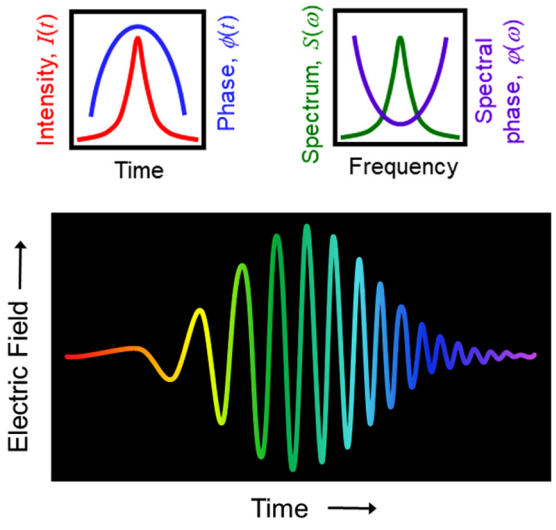

FIG. 1. Upper plots: Intensity and phase vs time (upper left) and the spectrum and spectral phase vs frequency (upper right). The curves plotted here are for a positively "chirped" pulse, whose red colors precede its blue ones. In the lower plot, the electric field of a positively chirped white pulse is shown using the color of its instantaneous frequency ($\omega_{inst}(t) = \omega_0 - d\phi/dt$) at the particular time. Reproduced with permission from www.frog.gatech.edu.

Before we begin discussing the methods for measuring pulses, we must define what we mean by terms like "short pulse," "very short pulse," and "extremely short pulse." Simply put, they range roughly from picoseconds to a few femtoseconds (even attoseconds) in length, respectively. Nonetheless, techniques and the devices based on them will have varying applicability, depending on the pulse wavelength and other parameters, so we will necessarily be a bit vague about pulse lengths, except in specific cases. Finally, we will refrain from any attempt to be quantitative about the meaning of the term "ultrashort," as it has evolved over the years from meaning slightly sub-nanosecond to, currently, at most hundreds of femtoseconds.

Also, a few words on the philosophy of pulse-measurement devices are in order. First, a pulse-measurement device should be able to completely measure pulses such as the simple one above but also much more complicated pulses as well. The reason for this requirement is that, if a measurement technique *can* only measure simple pulses, it *will* only measure simple pulses. In other words, when presented with a complicated pulse, a technique that can only measure simple pulses would likely yield an incorrect simple pulse, rather than the correct complicated pulse. This incorrect simple result is unacceptable. Second, the measurement device should itself be simple, with a minimal number of components and knobs, so that non-experts in the field of pulse measurement should be able to use it. Also, tweaking a knob of a measurement device often changes the resulting measurement—a highly undesirable situation—so it is also important that the device does not yield results that are sensitive to its alignment. The best way to accomplish this is for the device to have minimal knobs in the first place. Third, the technique should have a minimum of "ambiguities"—incorrect pulses that have the same measured trace as the correct pulse. Although we will see that ambiguities are unavoidable, some are "trivial" ambiguities, that is, can easily be calculated and/or removed and which we can usually live with. Unfortunately, in many techniques, there are "nontrivial" ambiguities, which can be neither calculated nor removed and so are unacceptable. When a device can only measure simple pulses, it is usually because the device has uncountably many nontrivial ambiguities when confronted with a more complicated pulse. In addition, the device should also yield some sort of feedback as to whether the measurement is correct or not and whether the assumptions of the measurement are satisfied. Such feedback is not usually available in optical devices, but we will see that it is important for pulse-measurement devices to have this capability, and, fortunately, it turns out to be available, so we now demand it.

Another issue arises because a device usually must average over many, possibly very differently shaped, pulses, that is, pulses with different intensities and phases. But, by definition and design, devices can only provide one result. This is, of course, an unsolvable problem; no one answer can be correct, so this problem provides an excellent litmus test for pulse-measurement devices. In this case, how well does the device's "measured pulse" approximate an actual "typical pulse"? Does the device at least yield the correct pulse length? And, critically, does the device provide an indication of the instability?

Note that we absolutely do *not* desire that the technique measure the *average* intensity and phase vs time or frequency.

Measuring the average spectral phase is particularly unacceptable because the average spectral phase will yield a pulse significantly simpler and shorter than is actually present when instability occurs. This is because a flat or linear spectral phase yields the shortest possible pulse for a given spectrum.[1,2] So averaging the spectral phase over many different complicated spectral phases (each corresponding to a different long and complicated pulse) yields a flat spectral phase, corresponding to a much shorter pulse—clearly a highly inaccurate result. Getting the pulse length at least approximately correct is clearly a necessary minimum standard in the presence of instability.

Worse, there is unfortunately no such thing as a "pulse-shape stability meter," so determining the presence or absence of instability necessarily also falls to the pulse-measurement technique. As a result, indicating pulse-shape instability will prove to be one of the most important tasks of a pulse-measurement device, one that most methods lack.

Indeed, while many pulse-measurement techniques have been proposed, very few meet even half of these criteria. This article will focus mainly on devices and general techniques that have these important properties. It will also include discussions of some that do not in part because, for the applications these techniques address, we cannot do better at this time and so the reader can see how things can go wrong.

We will also initially spend a good bit of time on older, less sophisticated methods, not because they are particularly useful (they lack most of the desired properties) but because they are the raw materials from which more sophisticated methods that do have the required properties are built. Also, understanding them and their limitations helps us to see why the more sophisticated methods work.

We will also distinguish between "self-referenced" techniques and those that require a well-characterized reference pulse to make a measurement. A self-referenced technique is necessary when measuring a pulse directly from a laser, when no previously measured time-synchronized reference pulse is available.

As mentioned previously, in principle, it would seem that self-referenced pulse measurement is impossible, as a shorter event seems required to measure an event in time, and the only available event—the pulse itself—is only as short as itself, not shorter. However, this argument was shown to be a mere myth in 1991,[3] and it is, in fact, not only possible but actually quite easy for a pulse to completely measure itself in time.[2] Indeed, self-referenced techniques with all the above required properties exist and will be discussed at length here.

Once a pulse has been measured using only itself (and, of course, some optical components), that pulse can then be used as a reference pulse for other, much more difficult pulse-measurement problems, such as the measurement of a very weak pulse or a pulse's complete *spatiotemporal* intensity and phase.

Our standards for techniques that use a reference pulse for such difficult tasks will be much lower with regard to pulse-shape stability measurement. This is mainly because, in general, no self-referenced methods exist that can make such measurements. Also, once established by a self-referenced technique as having a stable shape, a pulse train can be relatively certain to remain stable when simply propagating around an optical table. So, the stability measurement requirement can be relaxed in this case. Of course, if a technique that requires a reference pulse can also confirm pulse-shape stability in multi-shot measurements, all the better. Indeed, for measuring very complicated pulses generated from a stable pulse train using, say, an optical fiber in which high-order nonlinear-optical processes occur, pulse-shape stability cannot be counted on, and a technique that indicates the presence of instability is critical, even though it uses a reference pulse.

Over the decades, far too many pulse-measurement techniques have been proposed to all be discussed here (this article is a Tutorial, rather than a review), so only those that have been established to meet the above set of standards commonly expected from such devices will be discussed. Some promising new techniques have been proposed and demonstrated but have not yet been shown to meet all the above standards.

We will begin with a discussion of self-referenced techniques because, without a reference pulse, techniques that require one are useless.

II. MEASURING THE SPECTRUM

Of the above four quantities that define a pulse's electric field, it has only been generally possible to measure the pulse spectrum, $S(\omega)$. Spectrometers perform this task well and are readily available. The most common spectrometer involves diffracting a collimated beam off a diffraction grating and focusing the diffracted beam onto a camera. But interferometers also work (see Fig. 2).

Fourier transform spectrometers operate in the time domain and measure the transmitted integrated intensity from a Michelson interferometer vs delay τ, which is often called the light's *field*

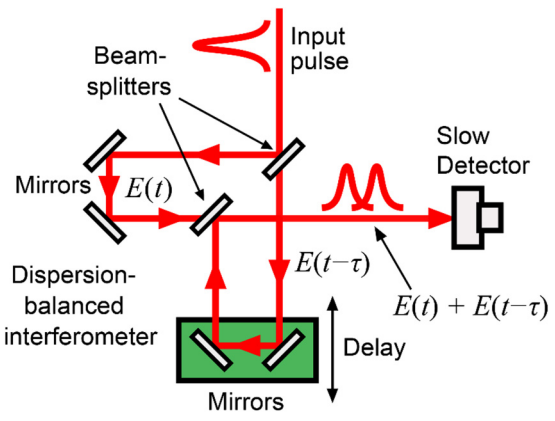

FIG. 2. Experimental layout for a Fourier transform spectrometer. Note that this Michelson-like interferometer design is "dispersion balanced" (has the same amount of glass, and hence dispersion, in each path), which is required for high-quality fringes in the presence of broadband light. It also has one-third the propagation length through glass as that of a standard Michelson interferometer, which is important when a pulse is extremely short and can be distorted by propagation through glass. Reproduced with permission from www.frog.gatech.edu.

autocorrelation or its *interferogram*,

$$\Gamma^{(2)}(\tau) = \int_{-\infty}^{\infty} E(t)\, E^*(t-\tau)\, dt, \tag{4}$$

neglecting constant terms. The interferogram's Fourier transform (with respect to delay τ) is simply the spectrum, a result known as the *Autocorrelation Theorem* (which is a special case of the more commonly known Wiener–Khinchin theorem describing the convolution of a function with its complex conjugate),

$$|\tilde{E}(\omega)|^2 = \mathscr{F}\left\{\int_{-\infty}^{\infty} E(t)\, E^*(t-\tau)\, dt\right\}, \tag{5}$$

where $\mathscr{F}$ indicates the Fourier transform with respect to τ, and the tilde ($\sim$) over the field indicates the Fourier transform.

Thus, all spectrometers, whether diffraction grating or Fourier transform devices, yield the spectrum, and only the spectrum.

A. The spectrum and one-dimensional phase retrieval

It is actually more interesting than it may seem to ask what information is, in fact, available from the spectrum. Obviously, if we have only the spectrum, what we lack is precisely the spectral phase. But what if we know that the pulse intensity vs time is definitely zero outside a finite range of times? Or at least asymptotes quickly to zero as $t \to \pm\infty$? Knowledge of finite support of the field in time is a great deal of additional information, which could, in principle, yield the pulse's spectral phase when combined with the known spectrum.

This class of problems is called the *one-dimensional phase-retrieval problem* for the obvious reasons that the problem is one dimensional and we know the spectral magnitude and are trying to retrieve the spectral phase using this additional information. It turns out that, even with additional information, the one-dimensional phase-retrieval problem is nearly always highly *ill-posed*[4] in the sense that there are many, often infinitely many, ambiguities—pulses that correspond to a given spectrum and that satisfy any additional constraints such as those mentioned above. The one-dimensional phase-retrieval problem is unsolvable in almost all cases of practical interest, even when additional information is included.

Specifically, even with this additional information, there are obvious, or "trivial," ambiguities.[4] Clearly, if the complex amplitude $E(t)$ has a given spectrum, then adding a phase shift, yielding $E(t)\exp(i\phi_0)$, does not change the spectrum. The same is true for a translation, $E(t-t_0)$, and the same is true for the complex-conjugated mirror image, $E^*(-t)$, which corresponds to a time reversal. As mentioned earlier, however, most researchers can live with these trivial ambiguities, hence the name. They are known, simple, and usually can be removed if additional information is available. But are there other, more difficult or even impossible-to-remove, "nontrivial" ambiguities?

Unfortunately, the answer is nearly always yes. In two classic papers written in 1956 and 1957, Akutowicz showed that knowledge of the spectrum in conjunction with the additional knowledge that $E(t)$ is of finite duration—often called *finite support*—is still insufficient to uniquely determine $E(t)$.[5,6] Indeed, he showed that infinitely many pulse fields usually satisfy these constraints. For example, a Gaussian spectrum can have any amount of linear chirp, and so can correspond to an intensity vs time that is also Gaussian, but with any pulse width. Of course, a Gaussian spectrum can have almost any higher-order phase distortion, as well.

III. THE INTENSITY AUTOCORRELATION

The *intensity autocorrelation*, $A^{(2)}(\tau)$, was the first pulse-measurement technique introduced in the 1960s, and it is based on using the pulse to measure itself.[7–15] It involves splitting the pulse into two, variably delaying one with respect to the other, and spatially overlapping the two pulses, usually in some instantaneously responding nonlinear-optical medium, such as a second-harmonic-generation (SHG) crystal (see Fig. 3). An SHG crystal produces light at twice the frequency of the input light with a field given by (ignoring constants),

$$E_{sig}^{SHG}(t,\tau) = E(t)\, E(t-\tau), \tag{6}$$

where τ is the delay between the two pulses. This field has an intensity (the squared magnitude of the electric field) proportional to the product of the intensities of the two input pulses,

$$I_{sig}^{SHG}(t,\tau) = I(t)\, I(t-\tau). \tag{7}$$

Because detectors (even streak cameras) are typically too slow to time-resolve $I_{sig}^{SHG}(t,\tau)$, this measurement necessarily yields an

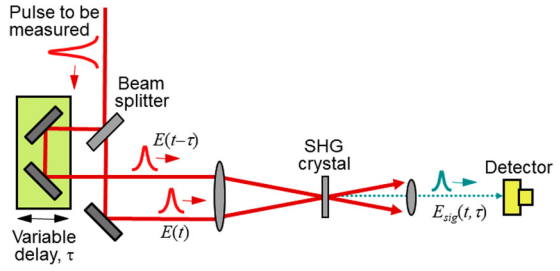

FIG. 3. Experimental layout for an intensity autocorrelator using second-harmonic-generation (SHG). A pulse is split into two replicas, one is variably delayed with respect to the other, and the two replicas are overlapped in an SHG crystal. The SHG pulse energy is measured vs delay, yielding the autocorrelation trace. Other effects, such as two-photon absorption or fluorescence caused by this process can also yield the autocorrelation, using similar beam geometries.[11,13,14,16] In this figure and others, a lens is shown focusing the beam into the crystal, but, for extremely short pulses, when propagation through glass unacceptably distorts the pulse, a curved mirror can be used in its place. Also, because intensity autocorrelation is not interferometric, compensation for dispersion from the beam splitter is not needed. Reproduced with permission from www.frog.gatech.edu and Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic Publishers, Boston, 2002). Copyright 2002 Springer.

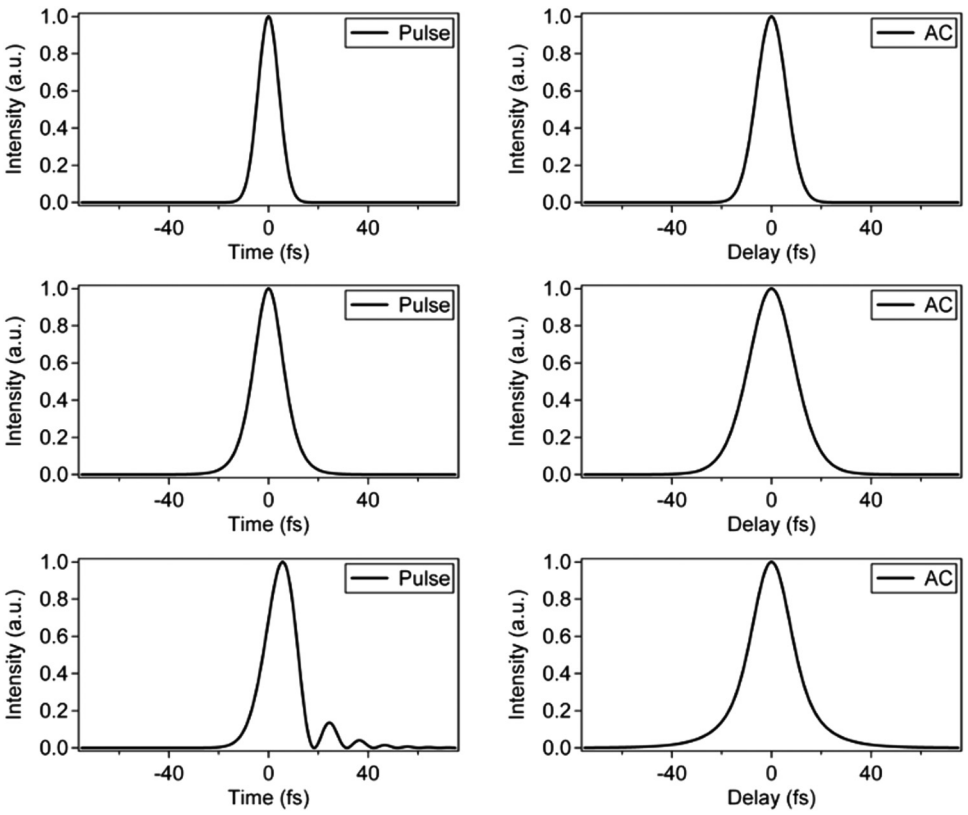

FIG. 4. Examples of theoretical pulse intensities and their intensity autocorrelations (AC). Left: Intensities vs time. Right: The intensity autocorrelation corresponding to the pulse intensity to its left. Top row: A 10 fs Gaussian intensity. Middle row: A 7 fs sech^2 intensity. Bottom row: A pulse whose intensity results from cubic spectral phase. Both scales are in arbitrary units. Note that, because the pulse is measuring itself, the autocorrelation loses details of the pulse, and, as a result, all these pulses have similar autocorrelations. Reproduced with permission from Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic Publishers, Boston, 2002). Copyright 2002 Springer.

integral over time,

$$A^{(2)}(\tau) = \int_{-\infty}^{+\infty} I(t)\, I(t-\tau)\, dt. \tag{8}$$

Equation (8) is the definition of the intensity autocorrelation, or, for short, simply the autocorrelation. This equation is different from the *field* autocorrelation [Eq. (4)], which provides only the information contained in the spectrum.

It is clear that an (intensity) autocorrelation yields some measure of the pulse length because no second-harmonic intensity will result if the pulses do not overlap in time; thus, a relative delay of one pulse length will typically reduce the SHG intensity by about a factor of two.

Figure 4 shows some simple pulses and their intensity autocorrelations. Because the intensity autocorrelation only attempts to provide a measure of the pulse intensity vs time $I(t)$ and makes no attempt to measure the phase, only intensities are shown.

A. The autocorrelation and one-dimensional phase retrieval

We can learn more about the autocorrelation by applying the Autocorrelation Theorem to Eq. (8), yielding

$$\tilde{A}^{(2)}(\omega) = |\tilde{I}(\omega)|^2, \tag{9}$$

where $\tilde{I}(\omega)$ is the Fourier transform of the intensity vs time [note that $\tilde{I}(\omega)$ is not the spectrum, $S(\omega)$]. In words, the Fourier transform of the autocorrelation is the mag-squared Fourier transform of the intensity. In other words, if we know the autocorrelation of an intensity, we know the magnitude, but not the phase of the Fourier transform of the quantity we wish to find, $I(t)$.

If this statement sounds familiar, it should. *It is another one-dimensional phase-retrieval problem!*

Thus, autocorrelation also suffers from trivial and nontrivial ambiguities. Figure 5 gives examples of different pulse intensities that have the same autocorrelation.

The approach taken by users of autocorrelation has been to assume a simple pulse shape, such as a Gaussian or a hyperbolic secant squared, and divide the width of their autocorrelation trace by the theoretical factor calculated for that pulse shape to obtain a possible pulse length. Although highly simplified models of ultrafast lasers yielded such theoretical pulse lengths, this assumption regarding the pulse shape has little to no theoretical justification in practice, but no better approach was available until the early 1990s, and that approach did not involve simple autocorrelation.

B. Autocorrelations of complicated pulses and noisy pulse trains

Nowhere does the lack of power of the autocorrelation to reveal structure in a pulse reveal itself more than in the

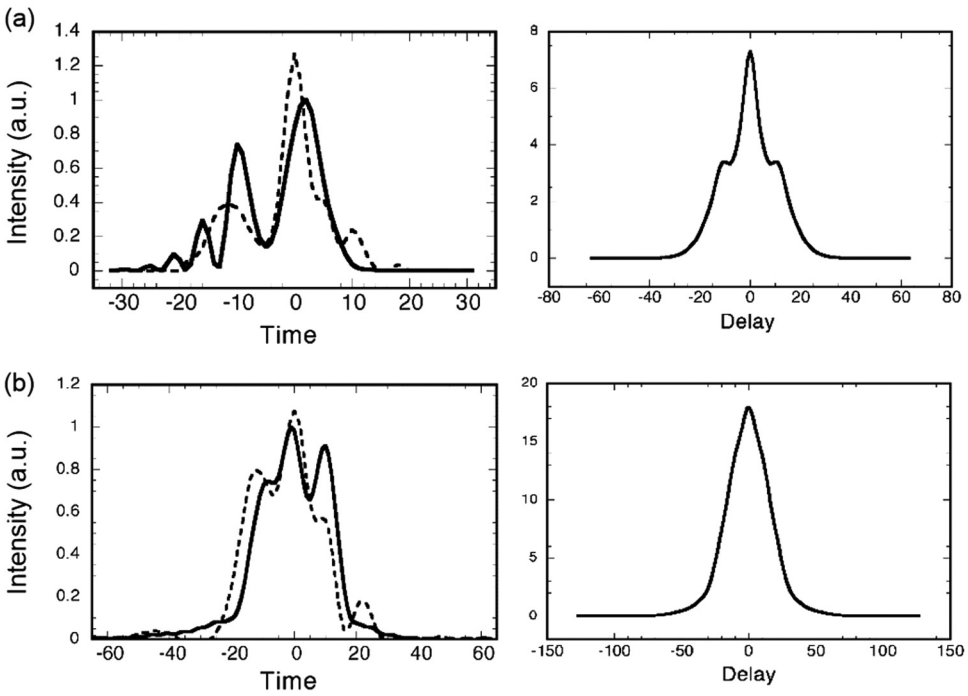

FIG. 5. (a) Left: Two pulse intensities that yield numerically equivalent autocorrelations. Right: Their intensity autocorrelation. (b) Left: Two additional pulse intensities that yield numerically equivalent autocorrelations. Right: Their intensity autocorrelation. In this case, despite their structured shapes, the intensity autocorrelations of these pulses never vary by more than the thickness of the above curve from a Gaussian. Reproduced with permission from Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic Publishers, Boston, 2002). Copyright 2002 Springer.

measurement of complicated pulses, where, unfortunately, there happens to be a great deal of structure to be revealed. In fact, for complicated pulses, it can be shown that, as the intensity increases in complexity (i.e., more structure), the autocorrelation actually becomes *simpler* and approaches a simple shape of a narrow spike on a pedestal, *independent of the intensity structure* (see Fig. 6).[17]

The narrow central spike, called the *coherence spike* or *coherent artifact* of approximate width τ_c (the coherence time), sits on top of a broad *pedestal* or *wings* of the approximate pulse length τ_p. In this case, when only a single complicated pulse (or a train of identical complicated pulses) is to be measured, we call this spike the *single-pulse* coherent artifact.

Even relatively simple pulses can yield autocorrelations comprising a coherence spike on a pedestal if the measurement averages over a noisy train of them, in which their shape varies.[18] Consider, for example, double pulses. Figure 7 shows some double pulses and their autocorrelations, which have three bumps.

Now, when a laser double pulses, it typically does so quite randomly. An over-pumped laser, for example, will often emit a train of double pulses with different, random separations for each double-pulse in the train. Since a typical ultrafast laser emits pulses at a very high repetition rate (100 MHz), and most autocorrelators are multi-shot devices anyway (measuring the SHG energy for only one delay at a time), the autocorrelator will necessarily average over the autocorrelations of many such pulses.

This situation will also produce a trace that contains two components, a narrow central coherence spike sitting on top of a broad pedestal, whose height will typically be much less than the value of 1/2 we saw in the last case. Clearly, the coherence spike is a rough measure of the individual pulses within the double-pulse, and the pedestal indicates the distribution of double-pulse separations. Again, while it would be tempting to try to derive the pulse length from the coherence spike—especially now that the pedestal seems so weak in comparison—the true pulse length is related not to the coherence spike but to the pedestal.

Now consider a related problem: averaging over a train of pulses with randomly varying complicated intensity pulse shapes. Because each pulse has the same simple autocorrelation, this situation clearly also yields precisely the same autocorrelation trace as shown in Fig. 6.

Unfortunately, in the 1960s, when autocorrelation was first introduced, a few researchers, desiring to claim a shorter (more "exciting") pulse, neglected the background and confused the coherent artifact for the pulse length. The correct interpretation of such traces was provided by Fisher and Fleck.[19] This confusion has been considered an embarrassing mistake, but it was not at all obvious at that time that such traces would arise in a technique that integrated over an always-positive quantity (the intensity) and yielded an always-positive measured quantity. Coherence is generally considered to only be a property of waves, which go negative and yield fringes that can cancel out in phase-sensitive measurements. Indeed, Fisher and Fleck never used the term "coherent." But the term has stuck, and it can be justified because complicated intensities vs time can be thought of as having (always-positive) oscillatory components.

To see just how careful one must be when using autocorrelation, consider the case of partial mode-locking, a common problem in ultrashort-pulse lasers, which can lead to highly unstable and complicated pulse shapes. Here, we consider only slightly complicated random pulse shapes consisting of a stable 12 fs flat-phase Gaussian component plus a longer random component. Autocorrelations are shown in Fig. 8 for a stable 12 fs pulse train and also for two such

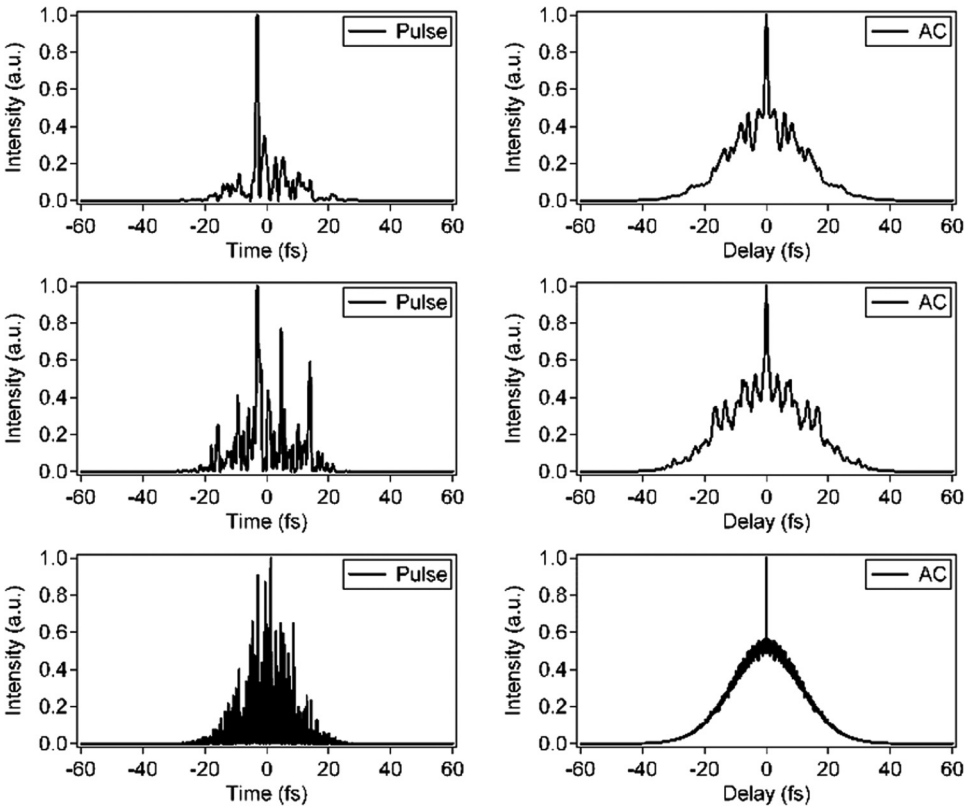

FIG. 6. Complicated intensities with Gaussian slowly varying envelopes with increasing amounts of intensity structure (left) and their autocorrelations (right). As the pulse increases in complexity (from top to bottom), the autocorrelation approaches the simple coherence-spike-on-a-pedestal shape, independent of the pulse intensity structure. Note that, as the structure increases in complexity, the coherence spike narrows along with the structure, while the pedestal continues to reveal the approximate width of the envelope of the intensity. Reproduced with permission from Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic Publishers, Boston, 2002). Copyright 2002 Springer.

random pulse trains. Note that the autocorrelation for these unstable pulse trains approaches the same shape as for complicated individual pulses and variable-separation double pulses (bottom row). Note also that when the coherence spike is only a factor of two or so shorter than the base (middle row), it blends in with the base and the two components cannot be distinguished from each other, yielding an autocorrelation trace that looks like that of a simple stable pulse train and yields a pulse length that is considerably shorter than the actual pulse. Thus, even for only slightly random pulse shapes and long before autocorrelation trace approaches the classic spike-on-a-pedestal shape, autocorrelation can lead to significant errors in the pulse length—by a factor of 2 or more.

Because this case involves multiple pulses, we call it the multi-pulse coherent artifact. In autocorrelation, the single- and multi-shot coherent artifacts are essentially indistinguishable, but they will be distinguishable in techniques that will be discussed later.

So, to summarize, even when the pedestal is weak, do not make the mistake of misidentifying the coherence spike as an indication of the pulse length! And also recognize that even an innocent-looking autocorrelation trace could in fact be the sum of a base and coherent artifact and yield pulses that are significantly shorter and simpler than in fact are present.

C. The autocorrelation and spectrum—In combination

If the autocorrelation by itself does not determine the intensity and the spectrum by itself does not determine the field, why not just use *both* measures in combination and see what the two quantities *together* yield?

Unfortunately, for ultrashort laser pulses, we do not have the spectrum and the *intensity*. We have the spectrum and the *autocorrelation*. Of course, as we have seen, the autocorrelation does not uniquely determine the intensity. So, this process can at best yield only *a* possible pulse field, not *the* pulse field. Indeed, for complicated pulses, because the autocorrelation contains so little information, this procedure is doomed to fail.

Even for simple pulses, no analytical work has been performed on this topic (it is mathematically very difficult). However, Chung and Weiner[20] have performed numerical computations and found numerous nontrivial ambiguities, in addition to the obvious direction-of-time ambiguity. Figure 9 shows one of the many ambiguities they found.

There are several variations on this theme, involving different types of autocorrelations, for example. While no one has taken the time to evaluate them as Chung and Weiner have for the above scheme, they may work in some simple cases. But it is doubtful that they perform much better in general, because an autocorrelation always becomes simpler as the pulse becomes more complicated, with the result that information is clearly lost. So, all clearly have many nontrivial ambiguities. As a result, attempts to save this approach have largely been abandoned, with a good reason.

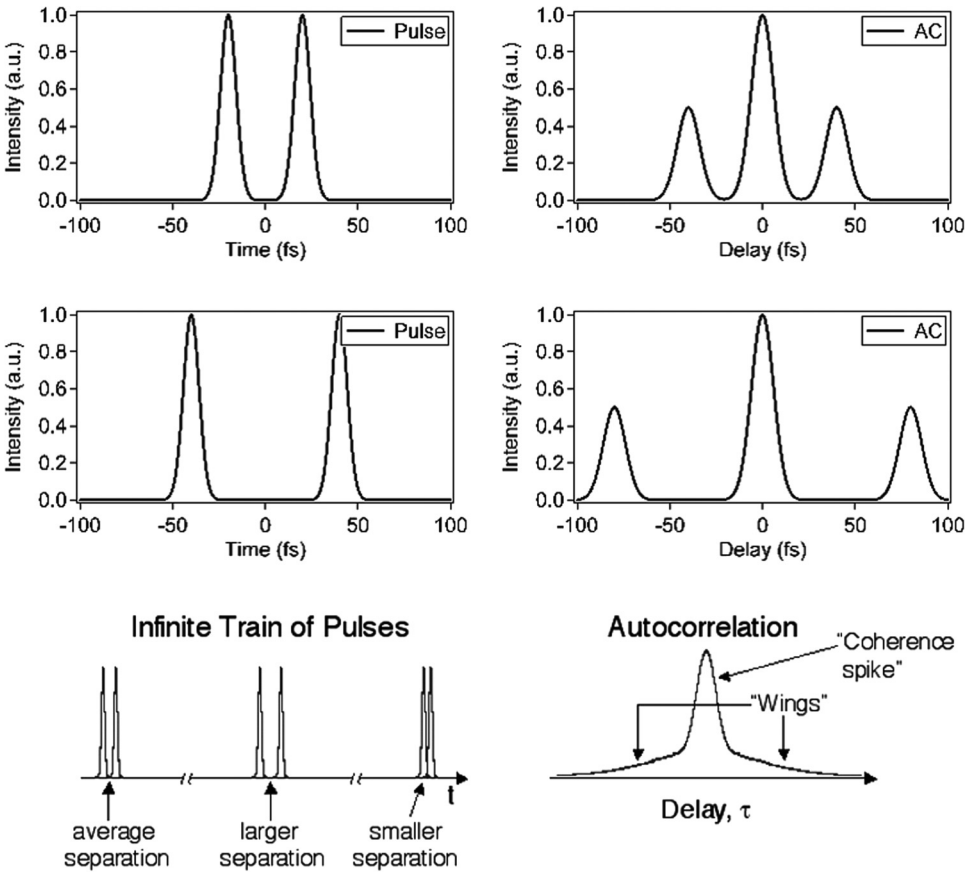

FIG. 7. Examples of theoretical double-pulse intensities and their intensity autocorrelations. Left: Intensities vs time. Right: The intensity autocorrelation corresponding to the intensity to its left. Top row: Two pulses (10 fs Gaussians) separated by four pulse lengths. Second row: The same two pulses separated by eight pulse lengths. Third row: A train of double pulses with varying separation. A multishot autocorrelation measurement (third row, right) averages over many double pulses. Note that the structure has washed out in the autocorrelation due to the averaging over many double pulses in the train, each with its own separation. The pulse length is better estimated by the width of the pedestal than by the width of the coherence spike. Reproduced with permission from Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic Publishers, Boston, 2002). Copyright 2002 Springer.

D. Interferometric autocorrelation

Another version of autocorrelation is called the *interferometric autocorrelation*, *phase-sensitive autocorrelation,* or *fringe-resolved autocorrelation* (FRAC). It was introduced by Jean-Claude Diels in 1983[21–27] and quickly became popular. Interferometric autocorrelation involves measuring the second-harmonic energy vs delay from an SHG crystal placed at the output of a Michelson interferometer (see Fig. 10). In other words, FRAC involves performing an autocorrelation measurement using collinear beams, so that the second-harmonic light created by the interaction of the two different beams combines coherently with that created by each individual beam. As a result, interference occurs due to the coherent addition of the three beams, and interference fringes occur vs delay, as well as a background out to ±∞ due to the SHG from the individual beams. This nonzero background contrasts with the zero background in intensity autocorrelation, which is also often referred to as the *background-free autocorrelation* when interferometric autocorrelation is also being discussed.

The expression for the second-order interferometric autocorrelation trace is

$$I_{FRAC}(\tau) = \int_{-\infty}^{+\infty} \left| \left[E(t) + E(t-\tau)\right]^2 \right|^2 dt \qquad (10)$$

$$= \int_{-\infty}^{+\infty} \left| E(t)^2 + 2\,E(t)\,E(t-\tau) + E(t-\tau)^2 \right|^2 dt. \qquad (11)$$

Note that, if the $E(t)^2$ and $E(t-\tau)^2$ terms were removed from the above expression, only the cross term, $2E(t)E(t-\tau)$ would remain, yielding the usual expression for background-free autocorrelation. These new terms, integrals of $E(t)^2$ and $E(t-\tau)^2$, are due to SHG of each individual pulse, and their interference, both with each other and with the cross term, yields the additional information in the interferometric autocorrelation that is not present in intensity autocorrelation. Multiplying out all the terms yields four terms: a constant, the intensity autocorrelation, a term related to the interferogram, and the interferogram of the second harmonic. Unfortunately, the mathematics involved is complicated, and interferometric autocorrelation does not yield the pulse or even its intensity and must also be curve-fit to a guessed field. Example interferometric autocorrelation traces are shown in Fig. 11.

Chung and Weiner also shed some light on the issue of how well the interferometric autocorrelation determines pulses by calculating interferometric autocorrelation traces for the pairs of pulses that yielded ambiguities in the spectrum/autocorrelation approach. They found that the resulting traces of the pairs of pulses had very similar interferometric autocorrelation traces (see Fig. 12).

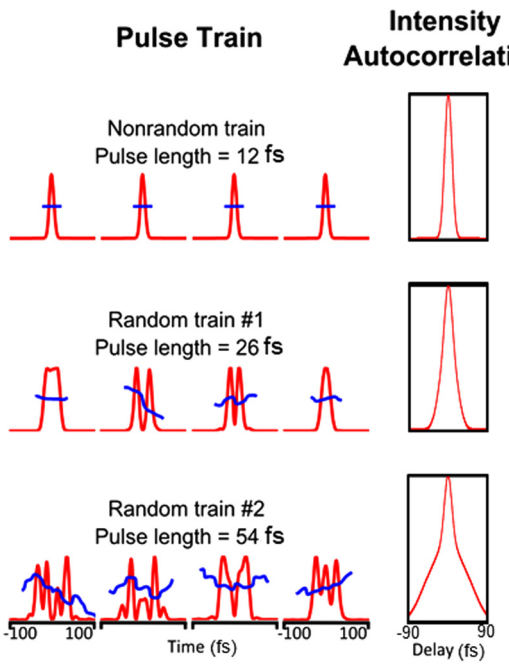

FIG. 8. Examples of theoretical random complicated intensities (red) and phases (blue, but the phase is irrelevant for autocorrelation) and their intensity autocorrelations. Top: Stable simple intensities vs time and their autocorrelation, which correctly yields the pulse length. Middle: The sum of the stable pulse train of the top row and a random component and its corresponding intensity autocorrelation. Bottom: The sum of the stable pulse train of the top row and a highly random, more complicated component and its corresponding intensity autocorrelation. Note that the structure has washed out in the middle and bottom autocorrelations due to the averaging over many double pulses in the train. Also, note that the middle row autocorrelation's coherent artifact is long enough that it plus the background look like an autocorrelation of a simple stable pulse and would easily be confused for one. Therefore, this autocorrelation trace yields a factor of almost two shorter pulses than is actually present. Reproduced with permission from www.frog.gatech.edu.

In short, interferometric autocorrelation yields essentially the same information as the autocorrelation and the spectrum, but it does so in one measurement and so is considerably more convenient and convincing.

Also, Diels and co-workers showed that, once a field has been fit to an interferometric autocorrelation trace, the direction of time could be determined by including a second interferometric autocorrelation measurement—actually a fringe-resolved *cross* correlation—in which some glass is placed in one of the interferometer arms. This breaks the symmetry and yields an asymmetrical trace. Then, assuming that the dispersion of the glass is known, Diels and co-workers showed that the two traces could be used to completely determine the pulse field in some cases.[25] Again, however, no study has ever been published on this algorithm's performance, and this approach is rarely used due to its complexity.

E. Third-order autocorrelations

The inadequacies of autocorrelation and its interferometric cousin were not lost on those who used them. As a result, many

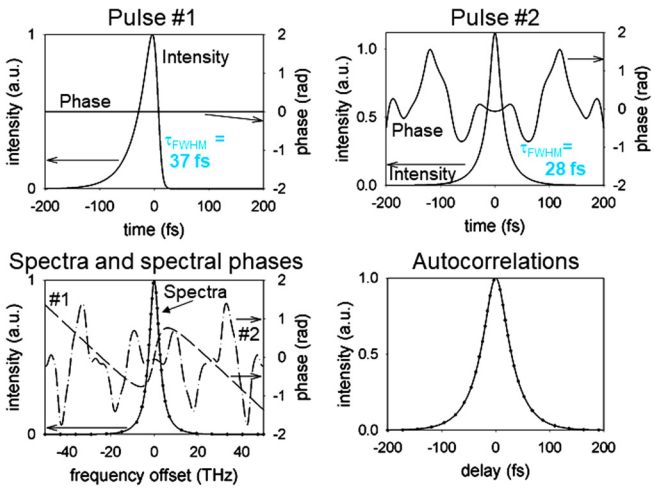

FIG. 9. Two pulses (top row) with different intensities and phases, which yield numerically identical autocorrelations (bottom right) and spectra (bottom left). The spectral phase of both pulses is shown as the dashed curves at bottom left. Reproduced with permission from Reproduced with permission from Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic Publishers, Boston, 2002). Copyright 2002 Springer.

variations exist, and an important one is the *third-order intensity autocorrelation,* or just the *third-order autocorrelation.*[28–32] Third-order autocorrelations have been generated using various nonlinear-optical effects, such as third-harmonic-generation (THG) and fluorescence involving three input photons.[30]

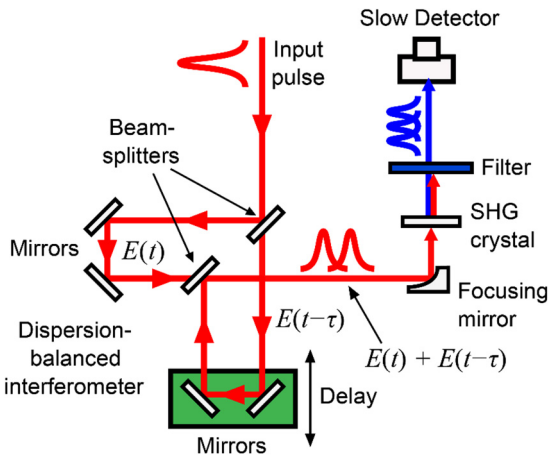

FIG. 10. Experimental layout for interferometric autocorrelation. The correct setup uses a dispersion-balanced Michelson interferometer, as shown here. A curved focusing mirror is shown, rather than a lens, to emphasize its application to extremely short pulses, for which propagation through glass must be minimized. Reproduced with permission from www.frog.gatech.edu.

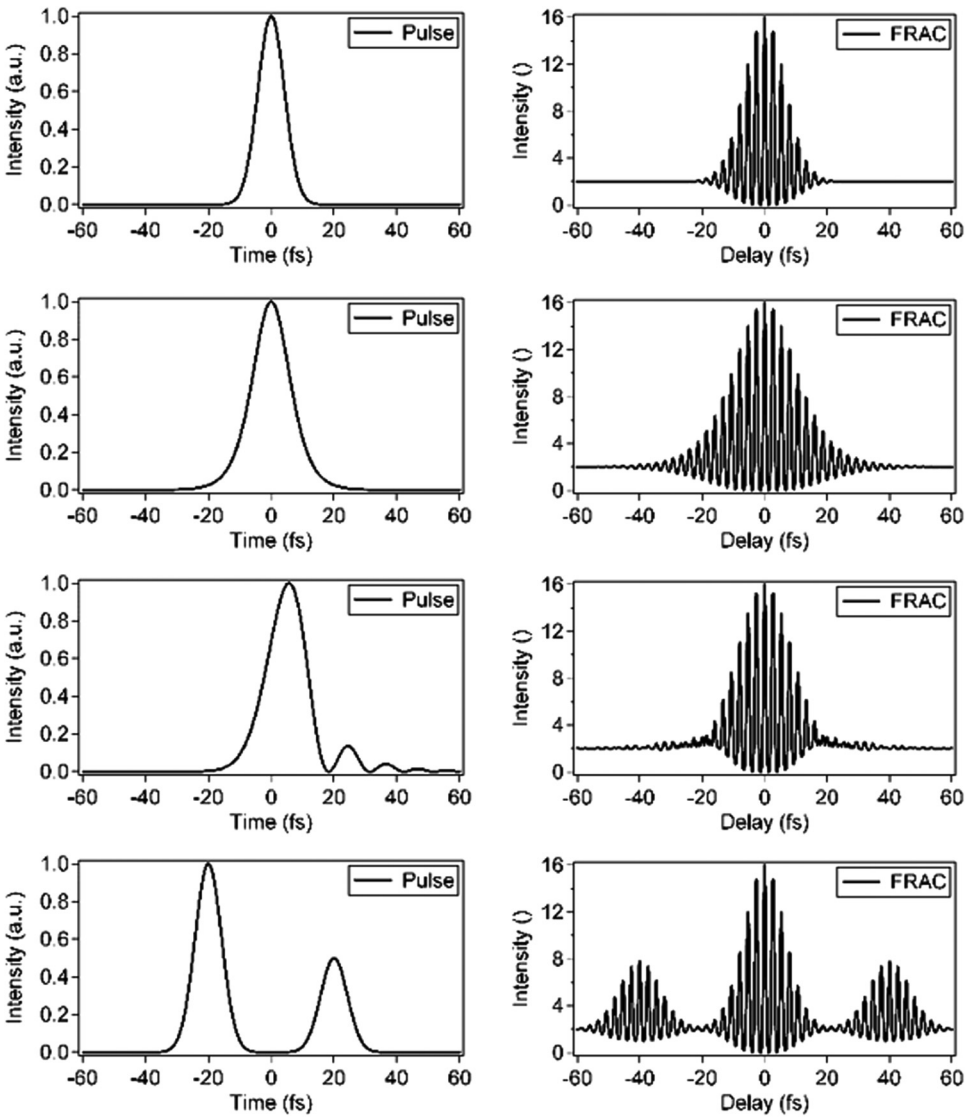

FIG. 11. Pulses and their interferometric autocorrelation or FRAC traces. Top row: A 10 fs Gaussian intensity. Second row: A 7 fs sech^2 intensity. Third row: A pulse whose intensity results from the third-order spectral phase. Note that the satellite pulses due to third-order spectral phase, which were invisible in the intensity autocorrelation, actually can be seen by looking carefully in the wings of the interferometric autocorrelation trace. Fourth row: A double pulse. Reproduced with permission from Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic Publishers, Boston, 2002). Copyright 2002 Springer.

Because $I(t)$ and $I(t-\tau)$ appear in the third-order autocorrelation *asymmetrically* (one is squared and the other is not), a third-order autocorrelation is symmetrical only if the intensity that produces it is. The asymmetry is not overwhelming, but it can be sufficient for determining whether a satellite pulse is a pre-pulse or a post-pulse.

Third-order autocorrelation also provides at best a rough estimate of the pulse length, and the third-order autocorrelation of a complicated pulse or a train of random pulses is similar to the second-order autocorrelation of such a pulse or train: a coherence spike on top of a broad pedestal.[17]

Third-order (and other higher-order) nonlinearities are also weaker than second-order ones and hence require more pulse energy. As a result, they do not work well for unamplified pulses from typical ultrafast laser oscillators, but they are useful for amplified pulses and UV pulses (for which SHG cannot be performed and where third-order nonlinearities are stronger).

Indeed, third-order autocorrelation is critical in ultrahigh-intensity settings, where a pre-pulse can have enough intensity to damage the sample before the main pulse arrives.[33] For this application, one of the pulses is frequency-doubled, and it and the original pulse overlap in a THG crystal, where the third harmonic is generated and then detected. As a result, any scattered light occurs only at the input frequency and the second and fourth harmonics, so no scattered light at the third harmonic exists to mask the desired signal pulse. Extremely high dynamic range measurements have been accomplished in this manner, and this device is the only known method for such applications. Indeed, in such applications, it is rarely necessary to know more about such satellite pulses than their arrival time and relative intensity, so third-order autocorrelation is ideal for this purpose.

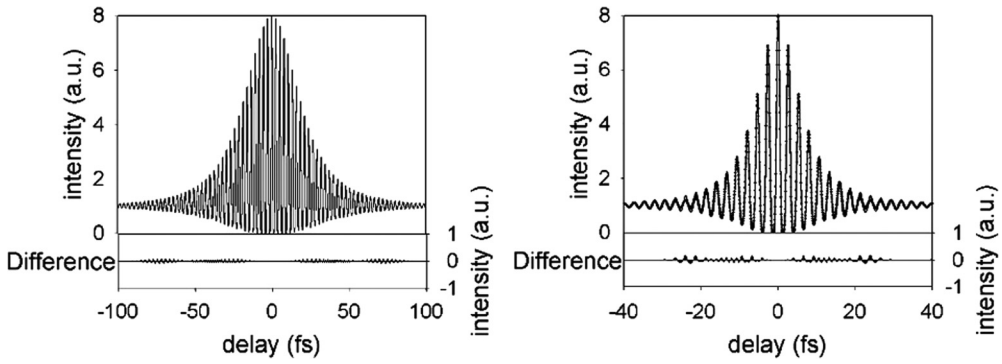

FIG. 12. Left: interferometric autocorrelation traces of the pair of pulses from Fig. 9. The difference between the two interferometric autocorrelation traces is plotted below. Right: interferometric autocorrelation traces of the same pulses but shortened by a factor of 5. Note that, in both cases, the two interferometric autocorrelation traces are very similar. Note also that the interferometric autocorrelation traces are even more difficult to distinguish as the pulse lengths decrease. Reproduced with permission from Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic Publishers, Boston, 2002). Copyright 2002 Springer.

F. Cross correlation

Occasionally, a shorter event is available to measure a pulse. In this case, a cross correlation can be performed (see Fig. 13). The cross correlation, $C^{(2)}(\tau)$, is given by

$$C^{(2)}(\tau) = \int_{-\infty}^{+\infty} I(t)\, I_g(t-\tau)\, dt, \tag{12}$$

where $I(t)$ is the unknown intensity and $I_g(t)$ is the gate pulse (shorter event) intensity vs time.

Although cross correlation yields no information about the phase, the cross correlation precisely yields the intensity when a much shorter gate pulse is available. This property can easily be seen by substituting $\delta(t)$ for $I_g(t)$, which yields $I(t)$ precisely. In fact, it is not even necessary to know the gate pulse—just that it is much shorter. The problem is that a much shorter pulse is not usually available.

G. Autocorrelation conclusions

Despite their drawbacks, ambiguities, and generally unknown information content, the autocorrelation and the spectrum remained the standard measures of ultrashort pulses for over 20 years, largely for lack of better methods. They allowed rough estimates for pulse lengths and time-bandwidth products (TBPs), and they helped researchers to make unprecedented progress in the development of sources of ever-shorter light pulses. However, the drawbacks of autocorrelation and its relatives began to severely limit their progress in the late 1980s, when the inability to measure the pulse spectral phase became the limiting factor in the generation of shorter pulses.

Fortunately, complete-intensity-and-phase measurement techniques became available in 1991. As a result, autocorrelation and its relatives, which are prone to badly under-estimating pulse lengths, are no longer appropriate measures of pulses and are rapidly becoming obsolete, except for checking for satellite pulses in ultrahigh-intensity applications. They have been discussed here mainly for historical reasons and because they are the building blocks from which more modern techniques are constructed.

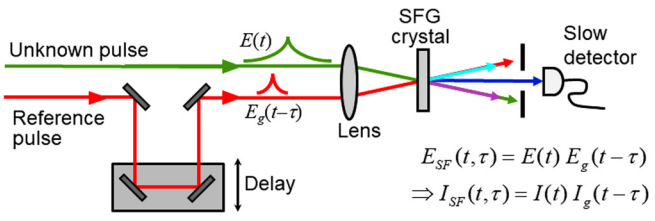

FIG. 13. A cross-correlator. A shorter reference pulse can gate a longer one in a sum-frequency generation (SFG) crystal and yield the intensity vs time of the unknown pulse. Reproduced with permission from www.frog.gatech.edu.

IV. THE TIME–FREQUENCY DOMAIN

With the failure of the combination of the spectrum and autocorrelation to determine the pulse, the next step in any effort to develop a self-referenced pulse-measurement technique is obvious: the spectrum *of* the autocorrelation.[34–36] It is remarkable that this simple approach was not considered in detail until 1991.[37] It is likely because this approach seemed unlikely to succeed. Another likely reason was that its mathematics appears complicated. It involves a hybrid domain: the *time–frequency domain*.[1,38] Indeed, the key mathematical ideas were not discovered until the mid-1980s. That this simple innovation solves the problem so completely and easily certainly is counterintuitive, as we shall see.[2]

Measurements in the time–frequency domain involve both temporal *and* spectral resolution simultaneously. This intermediate domain has received much attention in acoustics and applied mathematics research but, at the time, had received only scant use in optics and then only for qualitative measurements. One notable exception, however, was the idea of spectro-chronography,[39] which, although not widely used, was the principle and work that partially inspired the work that follows.

FIG. 14. The musical score is a plot of an acoustic waveform's frequency (vertically) vs time (horizontally), with information on the bottom regarding the intensity. Here, the wave increases in frequency with time. It also begins at low intensity (pianissimo), increases to a high intensity (fortissimo), and then decreases again. Musicians call this waveform a "scale," but ultrafast laser scientists refer to it (roughly) as a "linearly chirped pulse." Reproduced with permission from www.frog.gatech.edu.

In reality, everyone is actually quite familiar with the time–frequency domain. A well-known example of it is the *musical score*, which is a plot of a sound wave's short-time spectrum vs time. Specifically, this visualization involves breaking the sound wave up into short pieces and plotting each piece's spectrum (vertically) as a function of time (horizontally). So, the musical score is a function of both time and frequency (see Fig. 14).

The mathematically rigorous version of the musical score is the spectrogram, $\Sigma_g(\omega,\tau)$,[40]

$$\Sigma_g(\omega,\tau) \equiv \left| \int_{-\infty}^{\infty} E(t)\, g(t-\tau)\, \exp(-i\omega t)\, dt \right|^2, \tag{13}$$

where $g(t-\tau)$ is a variable delay gate function, and the subscript on Σ indicates that the spectrogram uses the gate function $g(t)$. Figure 15 is a graphical depiction of the spectrogram, showing a linearly chirped Gaussian pulse and a rectangular gate function, which gates out a piece of the pulse. For the case shown in Fig. 15, the rectangular function gates a relatively weak,

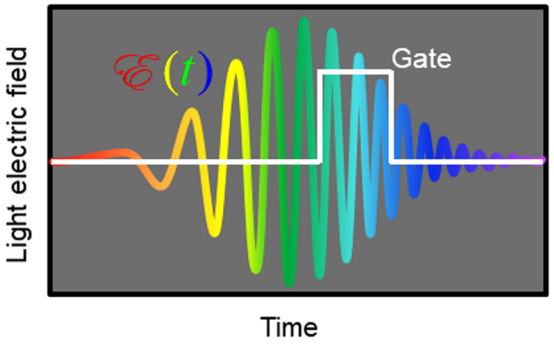

FIG. 15. Graphical depiction of the spectrogram. A gate function gates out a piece of the waveform (here a linearly chirped Gaussian pulse), and the spectrum of that piece is measured or computed. The gate is then scanned through the waveform and the process is repeated for all values of the gate position (i.e., delay). Reproduced with permission from www.frog.gatech.edu.

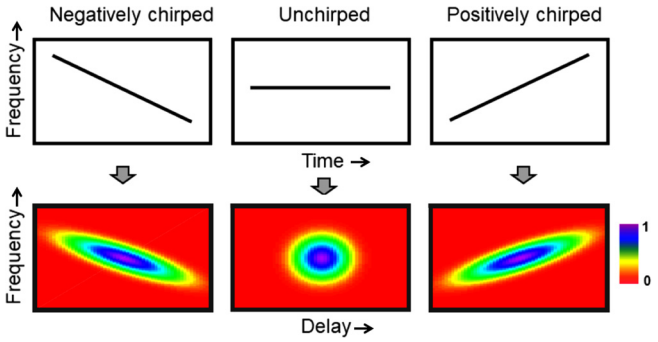

FIG. 16. Spectrograms (bottom row) for linearly chirped Gaussian pulses (top row), all with the same spectrum, but different amounts of chirp and using a Gaussian gate pulse. The spectrogram, like the musical score, reflects the pulse instantaneous frequency vs time. It also yields the pulse intensity vs time: notice that the shortest pulse (center) has the narrowest spectrogram in time. Looking at the spectrogram sideways yields the frequency domain quantities: the group delay vs frequency, with its height yields the spectrum. Reproduced with permission from www.frog.gatech.edu.

high-frequency region in the trailing part of the pulse, etc. The spectrogram is the set of spectra of all gated chunks of $E(t)$ as the delay, τ, is varied.

The spectrogram is a highly intuitive display of a waveform. Some examples of it are shown in Fig. 16, where it can be seen that the spectrogram intuitively displays the pulse instantaneous frequency vs time, and the pulse intensity vs time is also evident in the spectrogram. Importantly, knowledge of the spectrogram of $E(t)$ is sufficient to essentially completely determine $E(t)$[40,41] (except for a few trivial ambiguities, such as the absolute phase, which are typically of little interest in optics problems).

A. Frequency-resolved optical gating (FROG)

As in autocorrelation, it will be necessary to use the pulse to measure itself, in other words, to gate the pulse with itself, and to make a spectrogram of the pulse, it is necessary to spectrally resolve the resulting gated piece of the pulse. The frequency-resolved optical gating (FROG) technique, introduced in 1991,[3] measures such an "autospectrogram" of the pulse and was the first method (and, arguably, so far the only method) to completely solve the pulse-measurement problem, having all the required characteristics mentioned earlier.[37,42–49]

In its simplest form, FROG is any autocorrelation-type measurement in which the autocorrelation signal beam is spectrally resolved.[37,44,45] In other words, instead of measuring the autocorrelation signal energy vs delay, which yields an autocorrelation, FROG involves measuring the signal *spectrum* vs delay.

As an example, consider the usual SHG autocorrelation geometry. The autocorrelator's signal field is $E_{sig}(t,\tau) = E(t)E(t-\tau)$. Spectrally resolving $E_{sig}(t,\tau)$ yields the Fourier transform of the signal field with respect to time, and detectors, of course, measure

the squared magnitude, so the SHG FROG signal trace is given by

$$I_{FROG}(\omega, \tau) = \left| \int_{-\infty}^{\infty} E_{sig}(t, \tau) \exp(-i\omega t)\, dt \right|^2. \quad (14)$$

Note that the SHG FROG trace is a spectrogram in which the pulse field gates the pulse itself.

So why can $E(t)$ be obtained from its FROG trace? Letting the Fourier transform of the signal field with respect to delay be $E_{sig}(t,\Omega)$, where Ω is the conjugate frequency variable to delay, the FROG trace is the squared magnitude of the two-dimensional Fourier transform of $E_{sig}(t,\Omega)$. Once this quantity is found, it easily yields the pulse field. Now, because the FROG trace is the squared magnitude of this function of two variables, this problem is called, quite reasonably, the *two-dimensional phase-retrieval problem*.[2,4]

B. FROG and the two-dimensional phase-retrieval problem

Earlier, we discussed the *one-dimensional* phase-retrieval problem and saw that it was ill-posed and so a poor choice for the mathematics behind a pulse-measurement technique. Quite unintuitively, however, the two-dimensional phase-retrieval problem has been shown to have an essentially *unique solution* and is a *solved* problem.[4] This interesting and useful fact follows from the fact that the Fundamental Theorem of Algebra, which holds for polynomials of one variable, fails for polynomials of two variables.

The two-dimensional phase-retrieval problem, when finite support is the case, has only "trivial" ambiguities. In addition, there is an extremely small probability that another, nontrivial, ambiguous solution may exist, but no such ambiguity has ever been found in FROG.

In FROG, we actually do not have finite support because no function can be finite in extent in both time and frequency. However, FROG has another, much stronger constraint. We know that, for SHG, $E_{sig}(t,\tau) = E(t)E(t-\tau)$. This mathematical relation is a very strong constraint—the *mathematical form* that the nonlinear-optical signal field can have. Hence, we refer to this constraint as the mathematical-form constraint or nonlinear-optical constraint. Other versions of FROG, which use other nonlinear-optical processes, have slightly different but analogous constraints.

This additional information turns out to be sufficient, and thus, the problem is solved.[37] Indeed, it is solved in a particularly robust manner, with many other advantageous features, such as feedback regarding the validity of the data.[47,50,51] Thus, FROG is able to measure an event using the event itself and does not require a shorter one. In fact, using a shorter event is actually *undesirable*: substituting a delta-function for the gate pulse in the expression for the spectrogram [Eq. (13)] yields only the pulse intensity vs time and not the phase. Except for trivial ambiguities, FROG was recently rigorously proven to yield unique solutions.[52]

The mathematical-form constraint removes the direction-of-time, or inversion, ambiguity in all but one FROG variation (SHG). Also, some FROG versions have single-parameter ambiguities in the relative phases of well-separated pulses in time. Finally, no known technique is able to measure the relative phases of well-separated modes in frequency.

1. The FROG algorithm

There are many different FROG pulse-retrieval algorithms, all based on phase-retrieval approaches. They start with an initial guess for the field $E(t)$, usually random noise (see Fig. 17); unlike common minimization schemes, it is not necessary to start with a good guess (although we will see shortly that a good initial guess is a much better approach and hence yields a much more reliable result). A signal field $E_{sig}(t,\tau)$ is then generated using the relevant expression for the FROG trace for the beam geometry being used. This field is then Fourier transformed with respect to t in order to generate the signal field $E_{sig}(\omega,\tau)$ in the frequency domain. The measured FROG trace $I_{FROG}(\omega,\tau)$ is then used to generate an improved signal field $E'_{sig}(\omega,\tau)$ by realizing that the squared magnitude of $E_{sig}(\omega,\tau)$ should be equal to $I_{FROG}(\omega,\tau)$, so this step involves simply replacing the magnitude of $E_{sig}(\omega,\tau)$ with the square root of the measured trace to generate $E'_{sig}(\omega,\tau)$. $E'_{sig}(\omega,\tau)$ is then transformed back into the time domain by applying an inverse Fourier transform (IFT). In the last step of the cycle, the modified signal field $E'_{sig}(t,\tau)$ is used to generate a new guess for $E(t)$, and the process is repeated. Ideally, each iteration of the algorithm generates a better signal field, which eventually approaches the correct complex electric field.

The algorithm that has made FROG a technique that can quickly measure virtually every imaginable pulse is the *generalized projections* (GP) algorithm,[2,49] which, in the absence of noise, generally converges to the correct solution, even for very complicated pulses, with an accuracy only limited by the host-computer system's numerical precision (or experimental noise). It is also very versatile: the algorithm can be modified for any nonlinear-optical interaction, even some slow ones,[53] used to measure a pulse.

The essence of the generalized projections technique is graphically displayed in Fig. 18. Consider it a Venn diagram in which the entire figure represents the set of all complex functions of two variables, i.e., potential signal fields, $E_{sig}(t,\tau)$. The set of signal fields

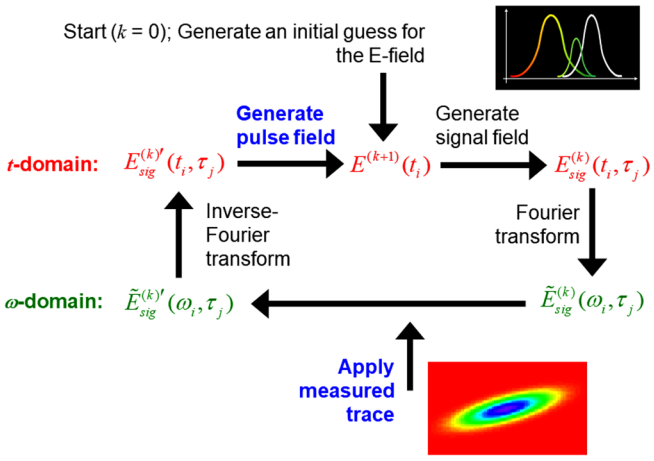

FIG. 17. Schematic of a generic FROG algorithm. k indicates the kth iteration. Reproduced with permission from www.frog.gatech.edu.

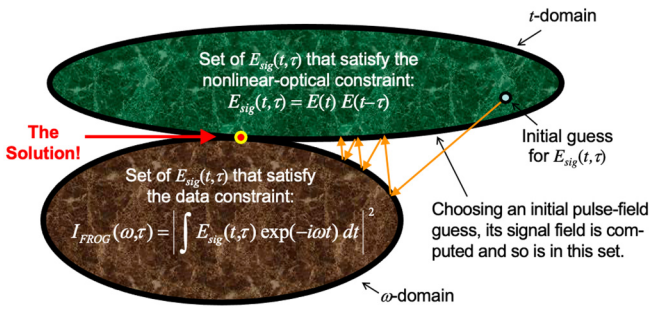

FIG. 18. Generalized projections applied to SHG FROG. The two equations are considered as constraints on the function, $E_{sig}(t,\tau)$, which, when found, yield $E(t)$, the pulse field. Moving to the closest point in one constraint set and then the other yields convergence to the solution. Although the mathematical-form constraint for SHG FROG is shown, other FROG geometries can be treated analogously. Reproduced with permission from www.frog.gatech.edu.

satisfying the data constraint are indicated by the lower elliptical region, while those satisfying the mathematical-form constraint are indicated by the upper elliptical region. The signal-pulse field satisfying both constraints corresponds to the intersection of the two elliptical regions and is the solution, uniquely yielding the pulse field, $E(t)$.

The solution is found by making function *projections*, which have simple geometrical analogs. We begin with an initial guess at an arbitrary point in signal field space. In the first iteration, we make a projection onto one of the constraint sets, which consists of moving to the point in that set (in function space) closest to the initial guess. From this point, we then project onto the other set, moving to the point in that set closest to the first iteration. This process is continued until the solution is reached.

When the two constraint sets are convex, i.e., all line segments connecting two points in each constraint set lie entirely within the set, convergence is guaranteed.

Unfortunately, the constraint sets in FROG are not convex. In this case, the projection is not necessarily unique, and the computed projection is called a *generalized projection*. The technique is then called generalized projections (GP), and convergence cannot be guaranteed. Although it thus is conceivable that the algorithm may stagnate (that is, not converge), this approach is in practice fairly robust in FROG problems. In a recent study of the GP algorithm, in the presence of noise, convergence occurred on the first guess for simple pulses about 75% of the time. The use of additional guesses yields convergence most of the time. For very complicated pulses (TBP ~ 100) in the presence of noise, convergence occurs about half of the time. Since pulses are rarely that complicated, and no other self-referenced technique has even been proposed that can measure such complicated pulses unless a reference pulse is available, this performance has generally been sufficient in most cases.

When a FROG algorithm fails, one can rerun the algorithm multiple times, but it would be far preferable to have a pulse-retrieval algorithm that converges every time. As a result, such an algorithmic approach has recently been developed. It is called the Retrieved-Amplitude N-grid Algorithm (RANA) approach.[54–56] It uses a much better initial guess—the *actual spectrum,* which, it turns out, can actually be extracted directly from the measured FROG trace. Obtaining this initial guess is possible because the frequency marginal (the integral of the trace with respect to delay) can be shown to be the autoconvolution of the spectrum for SHG FROG, which can be inverted to yield the spectrum using the simple fact that the Fourier transform of a finite-width function (e.g., the spectrum) is continuous and has continuous derivatives. Versions of FROG that use other optical nonlinearities also directly yield the spectrum but use slightly different approaches.

The RANA approach also uses smaller, coarser grids for initial iterations and so operates very quickly in such early iterations as a result. Finally, because it operates so quickly, it can use about a dozen initial guesses, all with the correct spectrum. The RANA approach can use any FROG algorithm but so far has been demonstrated using the GP algorithm. It achieves 100% convergence for even extremely complicated pulses with TBPs of 100 on tests of over 25 000 pulses for several FROG beam geometries and in the presence of significant noise in the measured traces.[54–56] It is freely available for download at www.frog.gatech.edu.

C. FROG beam geometries

Because FROG is a spectrally resolved autocorrelation, every nonlinear-optical process that can be used to make an autocorrelator can also be used to construct a FROG.[2] Here, we consider only the most common FROG geometries. SHG FROG achieves the best signal-to-noise ratios because it is the strongest (lowest order) nonlinearity and its signal beam is a different color, so scattered light is easily filtered. As a result, SHG is the most commonly used variation. Its apparatus and those of other common FROG geometries are shown in Figs. 19–21.

Figures 22(a) and 22(b) show theoretical FROG traces for the most common beam geometries, including polarization-gate (PG) FROG, self-diffraction (SD) FROG, and THG FROG. Pulses shown are all Gaussian intensity pulses and include the Fourier

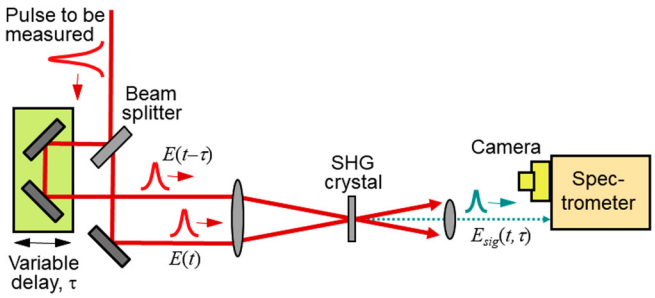

FIG. 19. Experimental apparatus for SHG FROG. The apparatus for third-harmonic-generation (THG) FROG is similar, but the third harmonic is generated instead of the second and in a slightly different direction. Reproduced with permission from www.frog.gatech.edu.

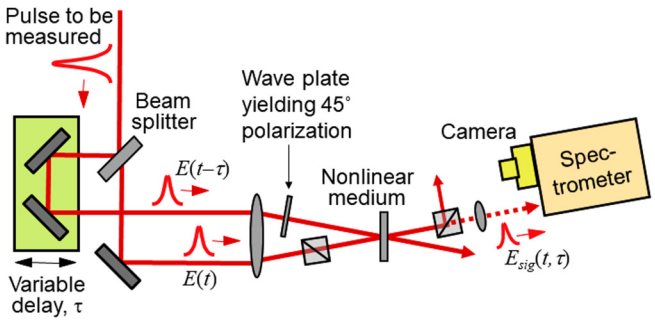

FIG. 20. Polarization-gate (PG) FROG apparatus. A pulse is split into two: one (the "gate" pulse) has its polarization rotated by 45° and is variably delayed, and the other (the "probe" pulse) passes through crossed polarizers. Then, the two pulses are overlapped in a piece of glass. The 45°-polarized gate pulse induces birefringence in the glass, which slightly rotates the polarization of the probe pulse causing it to leak through the polarizers if the pulses overlap in time. The leakage pulse spectrum is measured with respect to delay, producing the third-order autocorrelation trace. Reproduced with permission from www.frog.gatech.edu.

transform-limited (flat-phase) pulse, a pulse with negative chirp, a pulse with positive chirp, and a pulse with self-phase modulation, which do not change the pulse's intensity but distorts its spectrum. Note the symmetrical SHG FROG traces, which yield a trivial ambiguity in the direction of time. This ambiguity can be removed by inserting a piece of (dispersive) glass and making a second measurement; only one direction of time is consistent with both measurements.

D. Properties of FROG

It can accurately be said that, while intensity autocorrelation yields a blurry black and white picture of the pulse, FROG yields a high-resolution full-color image of it. Indeed, the pulse intensity and phase may be estimated simply by looking at the experimental

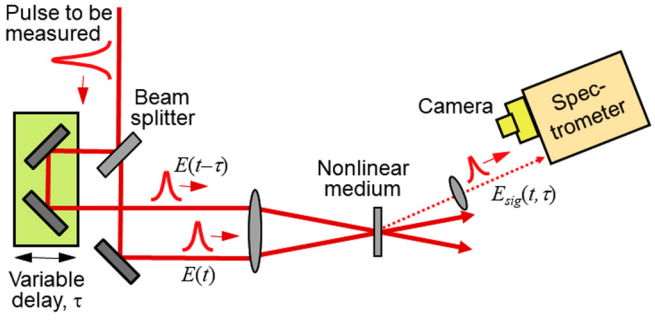

FIG. 21. Self-diffraction FROG. The pulse is split into two, delayed, and recombined in a third-order nonlinear medium, as in the previous figure, but here, the pulses induce a grating in the glass, which diffracts one of the pulses into a new direction off to the side, at $2\mathbf{k}_1 - \mathbf{k}_2$. Reproduced with permission from www.frog.gatech.edu.

FROG trace, or the iterative RANA approach may be used to reliably retrieve the precise intensity and phase vs time or frequency.

FROG has many useful features. It is very accurate; no significant approximations are made regarding the pulse. All that must be assumed in FROG is a nearly instantaneously responding medium (SHG crystals work very well in second order, as does glass in third order), and even that assumption has been shown to be unnecessary, as the medium response can be included in the pulse-retrieval algorithm.[53] Similarly, any known systematic error in the measurement may also be modeled in the algorithm,[50,51] although this modeling is not generally necessary, except for extremely short pulses (<10 fs) or for exotic wavelengths. Systematic error can often be removed by simple preprocessing of the measured trace.[51] Also, FROG completely determines the pulse with essentially infinite temporal resolution.[44,50] Interestingly, this extremely high temporal resolution can be obtained by using delay increments that are as large as, or even larger than, the time scale of the structure. The reason for flexibility in temporal delay steps is that the short-time information is also obtained from large frequency-offset measurements.

Another useful and important feature that is unique to FROG is the presence of feedback regarding the validity of the measurement data. The FROG trace is a time–frequency plot, that is, an $N \times N$ array of points, which is then used to determine N intensity points and N phase points, that is, $2N$ points. There is thus significant over-determination of the pulse intensity and phase—there are many more degrees of freedom in the trace than in the pulse. As a result, a measured trace that has been contaminated by systematic error is unlikely to correspond to an actual pulse. Thus, convergence of the FROG algorithm to a pulse whose trace agrees well with the measured trace virtually assures that the measured trace is free of systematic error, as well as additional effects that will be discussed shortly.

In practice, FROG has been shown to work very well in the IR,[57,58] visible,[59] UV,[60–62] and x-ray[63,64] regions of the spectrum. It has been used to measure pulse lengths from tens of attoseconds[64] to several ns.[65] It has measured pulses from aJ[66,67] to J in energy. It can measure simple near-transform-limited pulses to extremely complicated pulses with time-bandwidth products in excess of 1000.[68] It can use nearly any fast nonlinear-optical process that might be available. If an autocorrelator or a cross-correlator can be constructed to measure a given pulse, then making a FROG is straightforward since measuring the spectrum of it is usually easy.

1. Single-shot FROG

For high-repetition-rate ultrashort-pulse lasers, there often is not much variation from pulse to pulse. As a result, the delay in an autocorrelator or FROG may often be scanned in time with confidence that the pulse has not changed during the scan, but not always.

Indeed, some amplified laser systems have considerably lower repetition rates, and significant pulse-to-pulse variations are expected. In this latter case, a single-shot method is required. But how to accomplish this is not immediately obvious because the delay must somehow be scanned during a single pulse.

Fortunately, a single-shot autocorrelation or FROG trace can be obtained by mapping the delay onto the transverse position at

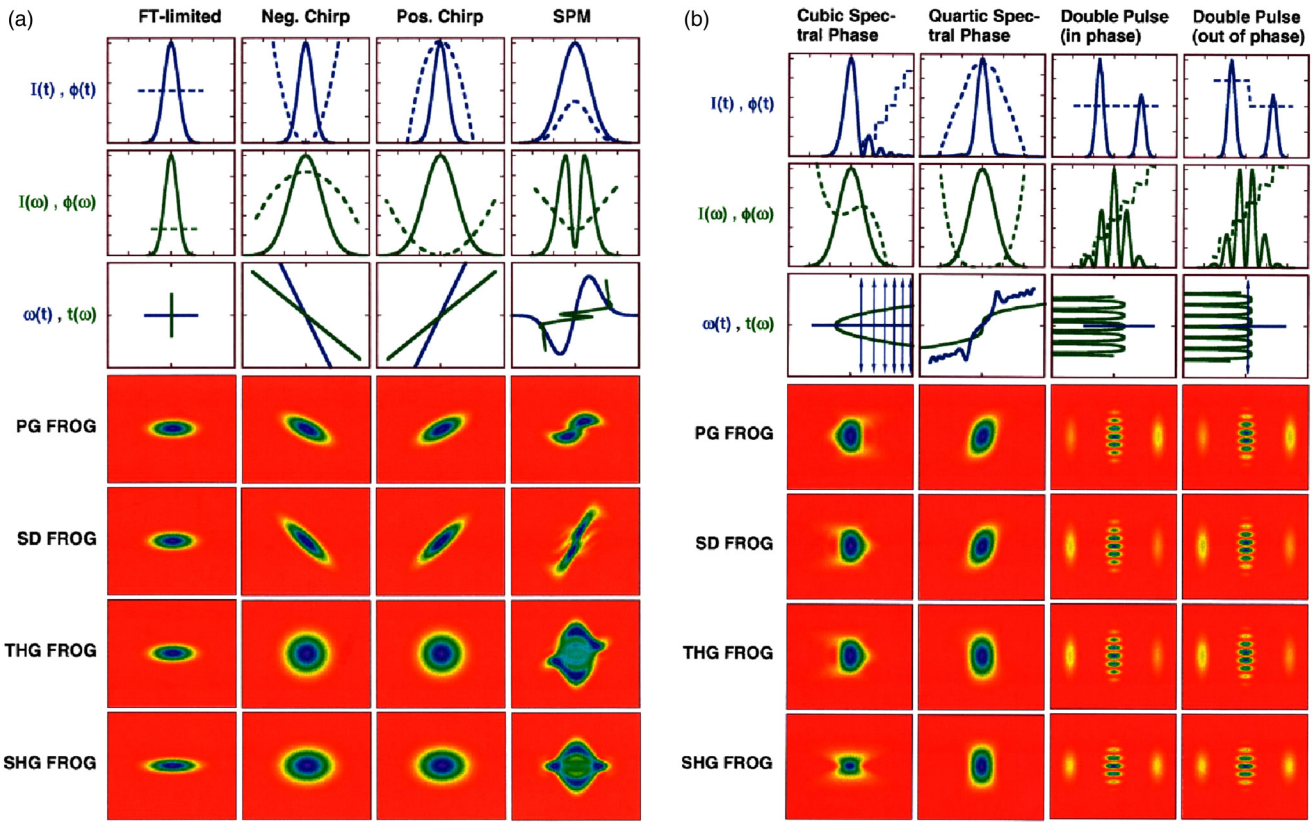

FIG. 22. (a) Comparison of traces for common ultrashort-pulse distortions for the most common FROG beam geometries. (b) Comparison of traces for common ultrashort-pulse distortions for the most common FROG beam geometries for additional pulses. Reproduced with permission from Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic Publishers, Boston, 2002). Copyright 2002 Springer.

the nonlinear medium and spatially resolving the autocorrelation signal by imaging the nonlinear medium onto a camera—an idea first implemented in 1970.[69] This involves crossing the two beams in the nonlinear-optical crystal at a large angle, so that, on the left, one pulse precedes the other, and, on the right, the other precedes the one (see Fig. 23).[15,70,71] In this manner, the delay ranges from a negative value on one side of the crystal to a positive value on the other. Usually, we focus with a cylindrical lens or mirror, so the beams are line shaped at the crystal, and the range of delays is greater. Spectral resolution occurs later in the perpendicular direction.

2. Near-single-cycle pulse measurement

The shortest pulse for a given wavelength is only a single cycle long (~2.7 fs for an 800 nm wavelength pulse). Measuring such incredibly short—and incredibly broadband—pulses can be very challenging. Propagation through essentially all glass must be avoided to avoid distortions to the pulse due to dispersion in the glass over the pulse's broad spectrum. Also, achieving SHG for all pulse wavelengths requires a very thin (~10 μm thick) crystal, which is challenging. Fortunately, such crystals are now common. Because the pulse can propagate through any amount of glass after the crystal, such thin crystals can be placed on substrates of any thickness for added durability. Finally, the SHG efficiency can vary significantly across the entire spectrum of the pulse, although the effect of this variation can easily be corrected for.[51]

Geometrical distortions can also potentially plague the measurement, causing the trace to broaden unacceptably. For example, as shown in Fig. 23, if the delay were to be scanned in the usual multi-shot manner using a delay stage, a range of delays would be sampled on any given shot, yielding a somewhat broadened autocorrelation or FROG trace, and the minimum amount of resulting pulse lengthening has been shown to be on the order of a fraction of one cycle.[51] The presence of this "transverse geometrical smearing" means that all of the beam geometries that involve crossed beams and scanning the delay will fail to accurately measure a near-single-cycle pulse. As a result, such multi-shot geometries should *not* be used for near-single-cycle measurements. We will discuss simple alternative FROG geometries that avoid this problem shortly, but many authors, unaware of this fact, have

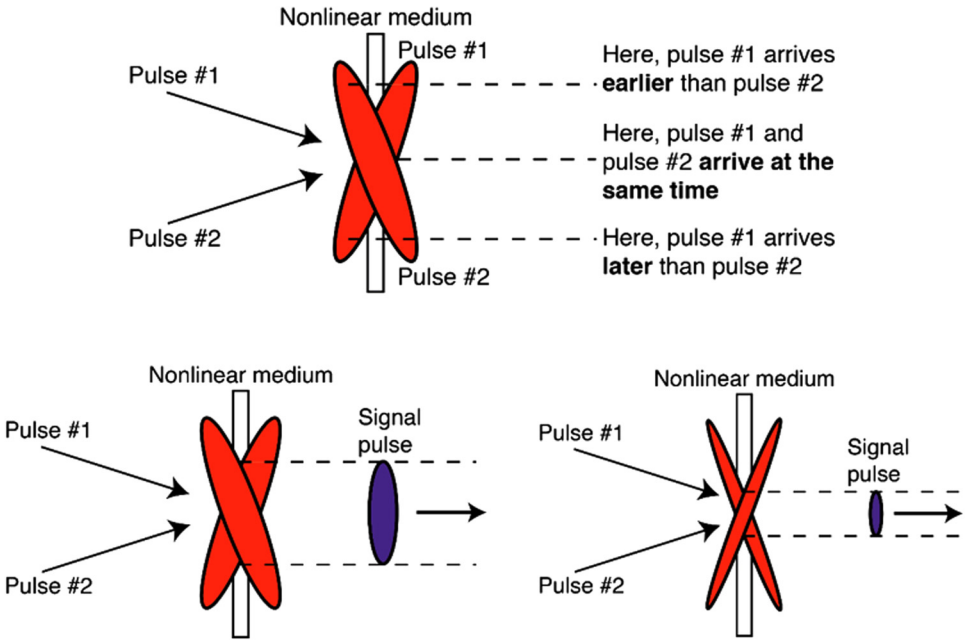

FIG. 23. Crossing beams at a large angle causes the delay to vary across the nonlinear medium, thus mapping delay onto transverse position. The length of the pulse is indicated by the width of the nonlinear-optical signal beam. Reproduced with permission from www.frog.gatech.edu.

erroneously concluded that FROG cannot measure such short pulses, so the reader should be cautioned not to draw such conclusions.

In addition, "longitudinal" geometrical smearing (in which the delay can vary along the nonlinear-optical signal beam direction) can occur in some crossed-beam geometries (except for SHG). This effect can also broaden traces in the delay direction (for more information on this effect, see Ref. 2).

One way to avoid all geometrical smearing effects in measurements of near-single-cycle pulses was pioneered by Steinmeyer and co-workers[72] and is to use collinear beams, that is, a spectrally resolved interferometric autocorrelator, in what is called Interferometric FROG (IFROG) (see Figs. 24 and 25). The collinear beams completely avoid all geometrical distortions. IFROG involves propagation through only a thin beam splitter and the nonlinear-optical crystal and so is ideal for near-single-cycle pulses. The interference fringes present in this case arguably increase its sensitivity to spectral-phase distortions, which are more likely in such short pulses. Multiple traces can be extracted from a single IFROG trace by Fourier filtering, each of which can be simultaneously used to retrieve the pulse using optimization techniques.[73,74] This additional redundancy further enhances the robustness of the pulse retrieval. Moreover, the interference fringes also contain the exact delay positions with interferometric precision, allowing the correction of timing jitter, which can be quite relevant for very short pulses.

Alternatively, single-shot FROG geometries work very well for near-single-cycle pulse measurements because the delay vs position is measured and used for the delay range, so no transverse geometrical smearing can occur. Indeed, single-shot crossed-beam SHG FROG is also free of longitudinal geometrical smearing and so is also ideal for performing such measurements, especially if an actual single-shot measurement is required due to a very low repetition rate or the presence of pulse train instability (something IFROG cannot be used for). The best such approach was pioneered by Akturk and co-workers using the beam geometry shown in Fig. 26, which has no glass in the beam until the SHG crystal (after which glass does not affect the result).[75] A measured pulse using this

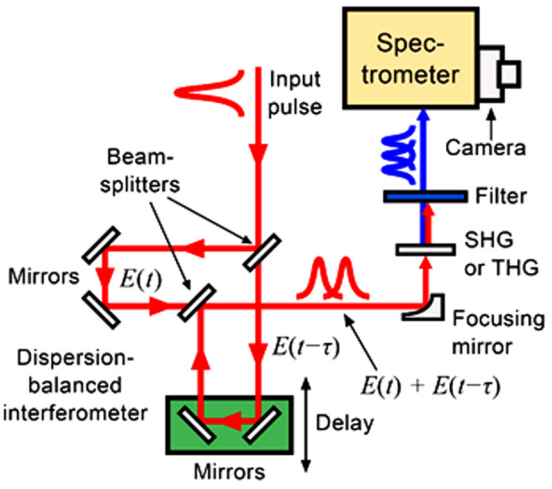

FIG. 24. Schematic for interferometric FROG (IFROG) for measuring near-single-cycle pulses. Note that a curved focusing mirror replaces the usual lens to avoid propagation through glass. Also, the beam splitters are thin, so the dispersion introduced by their glass is also minimal and also designed to be identical for each beamlet yielding a "dispersion-balanced" interferometer—important for yielding correct, symmetrical fringes. Reproduced with permission from www.frog.gatech.edu.

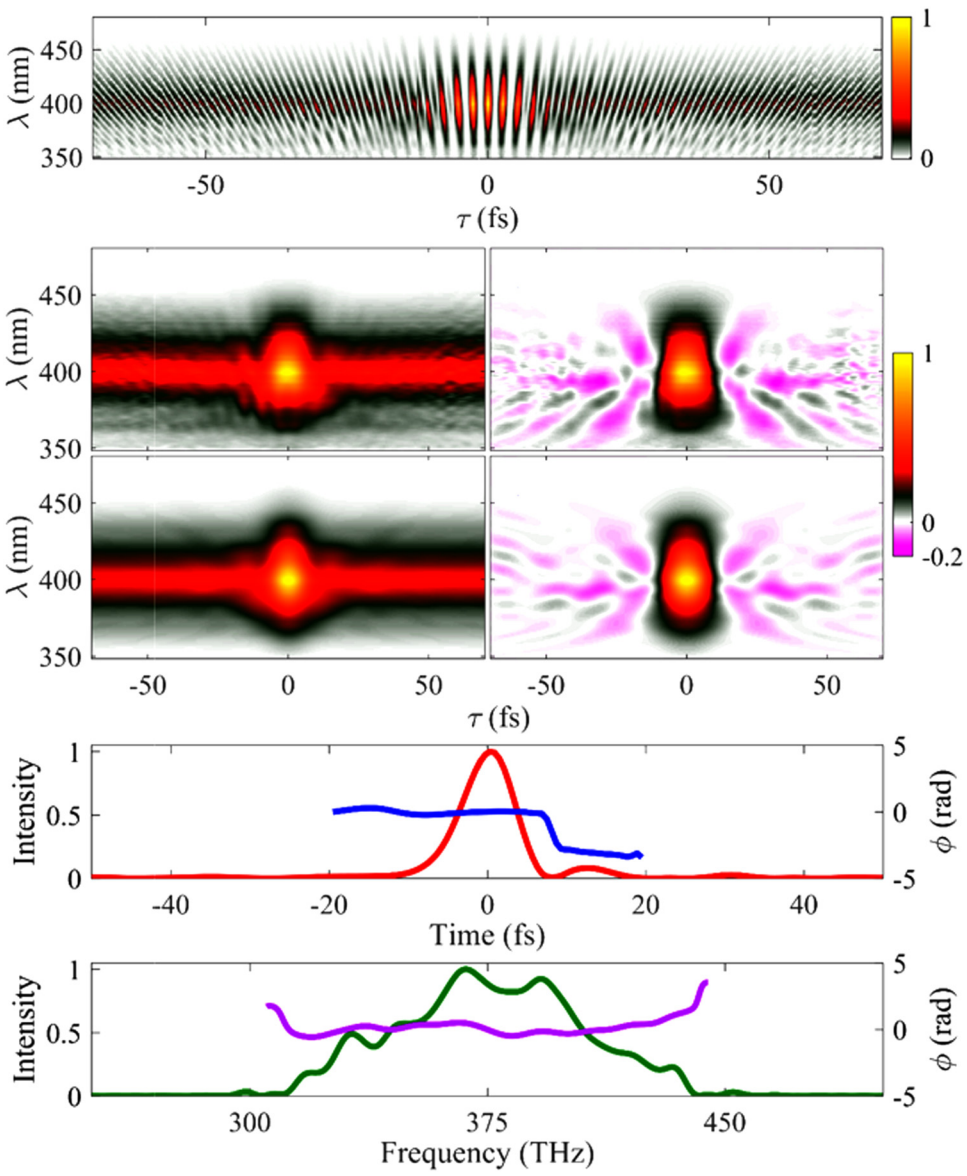

FIG. 25. Near-single-cycle pulse measured using IFROG. Top row: measured trace. Note the fringes. Middle row: Measured and retrieved Fourier-filtered trace (with fringes removed). Bottom rows: retrieved pulse in the time and frequency domains. Figure provided by G. Steinmeyer and reproduced with permission from Hyyti *et al.*, Rev. Sci. Instrum. **88**, 103102 (2017). Copyright 2017 AIP Publishing LLC.

geometry is shown in Fig. 27. This approach is less sensitive than IFROG, however, and it requires a spatially smooth beam.

E. FROG and the coherent artifact

We have seen that, when presented with a complicated pulse, variable-separation or variable-relative-phase double pulses, or trains of pulses with varying complicated shapes, autocorrelation yields an inconvenient and misleading coherent artifact. Interestingly, FROG and its variations not only measure much more information but also deal very well with pulse-shape instability and the resulting coherent artifact.

First, as the pulse to be measured becomes very complicated, the FROG trace also becomes very complicated, not simpler, as in autocorrelation. This feature is good: the pulse information is not lost in its measurement, and FROG's performance (with the RANA approach) has been studied for complicated pulses with TBPs as large as 100, even in the presence of noise, and the algorithm does very well.[76] In other words, FROG does not suffer from a single-pulse coherent artifact.

For variable-separation or variable-relative-phase double pulses, FROG yields traces that clearly indicate the presence of a double pulse, but frequency fringes in the center lobe that indicate the pulses' relative phase (which would be present for stable double pulses; see Fig. 28) tend to wash out.[77] However, because such a

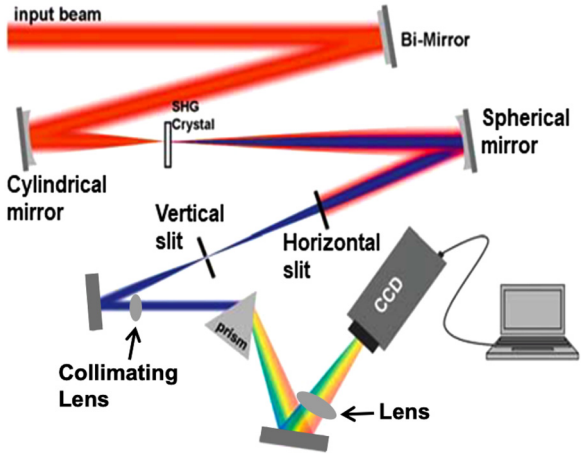

FIG. 26. Single-shot, single-cycle SHG FROG setup of Akturk *et al.* The "bi-mirror" consists of two mirrors next to each other, each at a slightly different angle, which spit the beam into two beams, each propagating at a slightly different angle, yielding crossing beams. Note that a curved focusing mirror replaces the usual lens to avoid propagation through glass. The pulse propagates through no glass at all before 5 μm-thick BBO crystal (and so is not distorted at all), essential for such extremely short pulses. This single-shot beam geometry can be used to measure a single pulse or to average over many pulses. Reproduced with permission from www.frog.gatech.edu, and Akturk *et al.*, J. Opt. Soc. Am. B **25**(6), A63 (2008). Copyright 2008 The Optical Society.

fringeless trace does not correspond to a possible FROG trace, the FROG algorithm still correctly yields a double pulse. While it often yields an incorrect relative height for the two pulses, the relative height mismatch between sub-pulses is easily corrected using additional information in the trace.[77]

Finally, when a FROG measurement must average over a train of randomly shaped complicated pulses, it also yields a trace with a spike in it that cannot correspond to a trace of a single pulse—a multi-shot coherent artifact. Fortunately, significant discrepancies between the measured and retrieved traces result and make it clear that the pulse train is unstable (see Fig. 29).

While it was well-known that discrepancies between measured and retrieved FROG traces could indicate pulse-shape instability,[68] the above simulations were not performed until a just few years ago,[77–81] when they resolved a mystery that had remained unsolved for two decades. FROG simulations using multiple initial guesses, and, more recently, the RANA approach, always showed very reliable algorithm convergence;[49,76] even in the presence of considerable noise, experimental FROG measurements often showed discrepancies between measured and retrieved traces, even after trying multiple initial guesses. For many years, such discrepancies were routinely ascribed to algorithm non-convergence by those performing the measurements. However, it is now clear that such discrepancies are due to pulse-shape instability. These discrepancies can even be quantified and separated from the measured trace, and can be shown to be directly caused by pulse-to-pulse instability,[82] and such pulse-shape variations can occur not only from pulse to

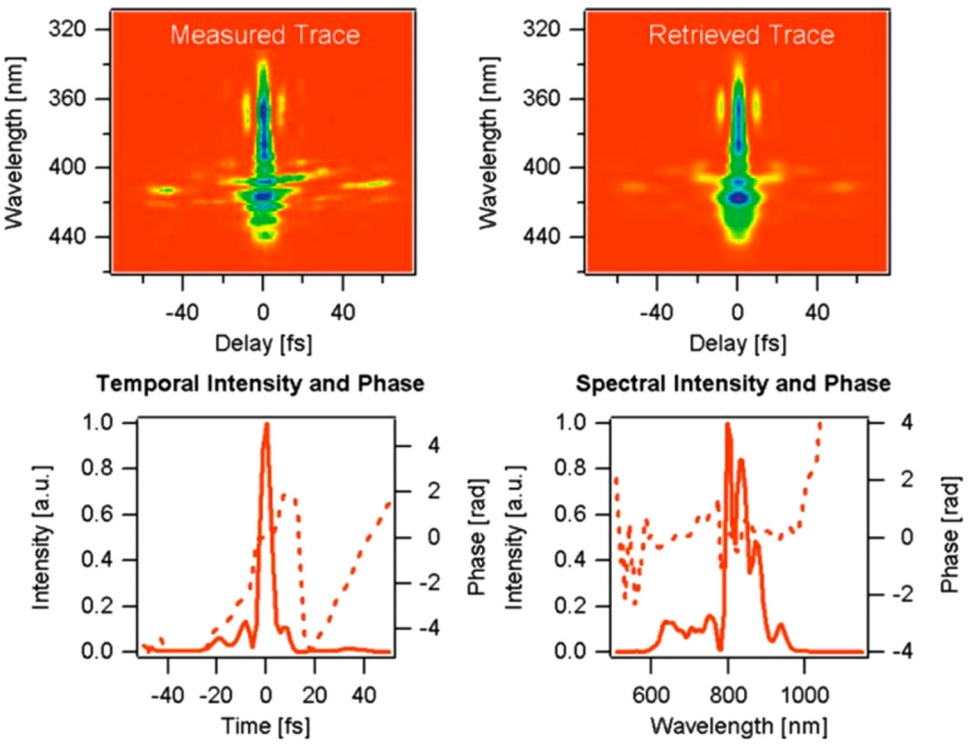

FIG. 27. Near-single-cycle SHG FROG measurements of Akturk *et al.*, of a 4.9 fs pulse. Measured and retrieved traces (top) and the measured intensity and phase vs time (lower left) and frequency (lower right). Reproduced with permission from Akturk *et al.*, J. Opt. Soc. Am. B **25**(6), A63 (2008). Copyright 2008 The Optical Society.

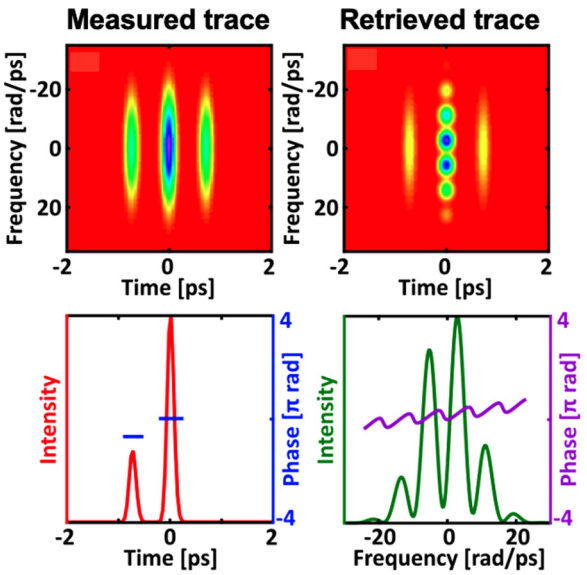

FIG. 28. Simulations of SHG FROG measurements of double pulses with random relative phases and separations. Note that the central lobe's fringes, which ordinarily yield the pulse relative phase, have washed out due to the random relative phase. So, the relative phase is no longer an important parameter. But the retrieved pulse relative heights are wrong. However, the retrieved sidelobes are too weak, and the heights from the measured trace can be used to correct this difference. Reproduced with permission from www.frog.gatech.edu.

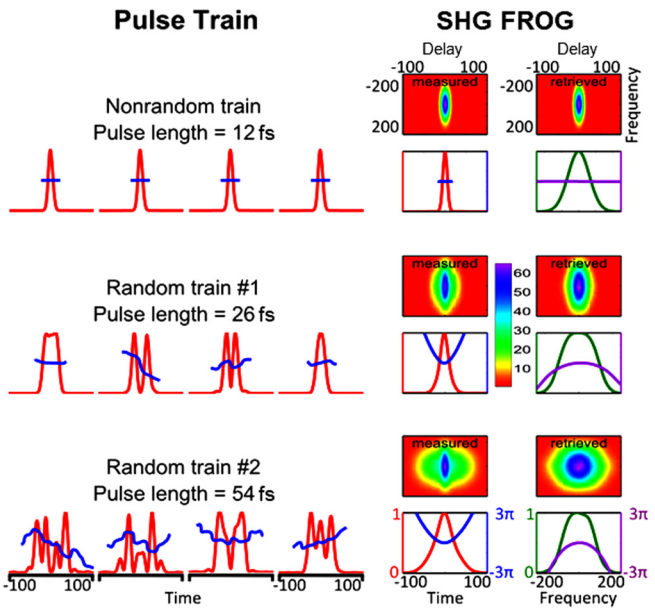

FIG. 29. SHG FROG traces and measurements of nonrandom (top row) and random (middle and bottom rows) trains of pulses. Note that FROG yields major discrepancies between the measured and retrieved traces when the pulse train is random. This discrepancy is an excellent indication of pulse-shape instability. Other versions of FROG also indicate instability in the same manner with a retrieved pulse closer to a typical pulse in the train. Reproduced with permission from www.frog.gatech.edu, and Ratner *et al.*, Opt. Lett. **37**, 2874 (2012). Copyright 2012 The Optical Society.

pulse but also from place to place in the beam, and so can occur even in single-shot measurements, although this latter effect requires more study.

In other words, like autocorrelation, FROG has a multi-pulse coherent artifact, but FROG's 2D traces with massive data redundancy and newly developed extremely reliable pulse-retrieval algorithm (RANA) allow FROG to provide unambiguous identification of pulse-shape instability or its absence. Also, all FROG versions yield correct average pulse lengths and TBPs, with third-order FROG devices also yielding the approximate structure of a typical pulse.

F. Cross correlation FROG (XFROG)

When a well-characterized reference pulse is available and can be used to gate the unknown pulse in a FROG setup, the resulting measurement is referred to as cross correlation FROG, or XFROG.[83] The resulting trace is precisely a spectrogram. XFROG is the most reliable version of FROG and so is essentially always advised when a well-characterized (not necessarily shorter) reference pulse is available. In a study of the simple GP XFROG algorithm for pulses with TBPs as large as 100, it retrieved the correct pulses with a 100% reliability, even in the presence of significant noise in the trace.[76]

Indeed, where XFROG truly excels is for measuring extremely complicated pulses. It is the only known method for measuring ultrabroadband supercontinuum, whose TBP can be as high as 10 000. Shown below in Fig. 30 are XFROG measurements of a train of supercontinuum pulses averaged over 10^{11} laser shots. Notice that the measured and retrieved traces disagree in their fine structure—a clear indication of pulse-shape instability. This conclusion was confirmed by single-shot spectral measurements.[68] These measurements showed that the supercontinuum was arguably the most complicated and unstable laser pulse train ever generated (or measured).

XFROG is also useful for measuring pulses at difficult-to-detect frequencies, such as in the mid-IR (~2 to ~40 μm), where cameras are very expensive and typically still have only ~10^4 pixels. Using a near-IR reference pulse, the nonlinear interaction can be chosen to produce the XROG trace at visible frequencies so that a standard Si detector can be used. Such measurements have been done using sum- or difference-frequency generation between the mid-IR and a near-IR reference pulse in crystals.[84,85] Similarly, to use XFROG to measure mid-IR pulses with bandwidths spanning several octaves, four-wave mixing up-conversion, a $\chi^{(3)}$ process, can be performed in a gas such as air.[86] In these cases, measuring the spectrum vs delay of the new visible or UV pulse yields an XFROG trace from which the mid-IR pulse can be retrieved, provided that the reference near-IR pulse is known.

In addition, by use of the optical-parametric-amplification nonlinearity, gain can also be achieved along with the usual temporal gating. Gains of up to 10^6 are possible, allowing the

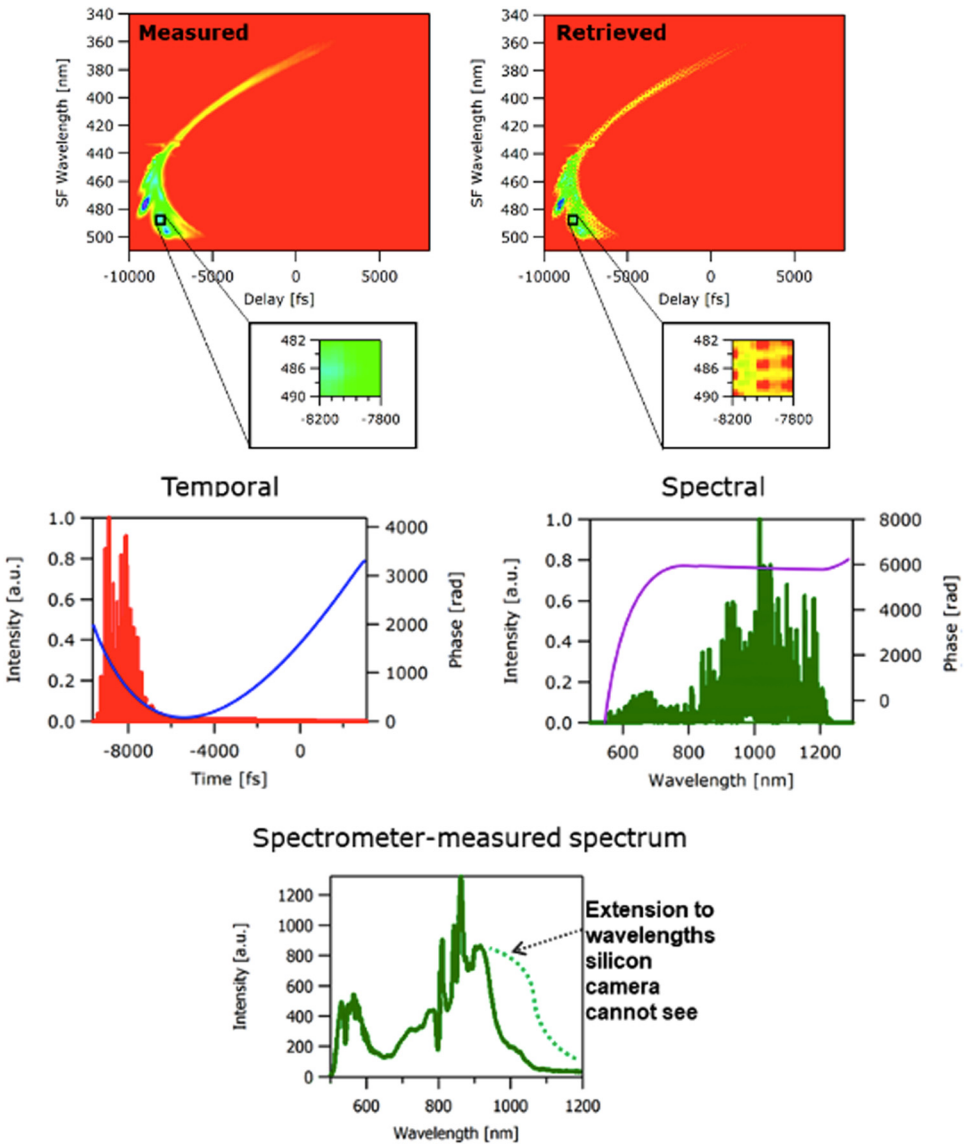

FIG. 30. FROG measurement of supercontinuum. Note the major discrepancies between the measured and retrieved traces because the pulse train is random. The FROG-measured spectrum is more accurate than the spectrometer-measured spectrum. Reproduced with permission from www.frog.gatech.edu.

measurement of extremely weak pulses. In this manner, continuum pulses with energies of aJ were measured.[67]

G. Very simple FROG: GRENOUILLE

Recall that device simplicity is also important, and, while FROG is not complicated, a significantly simpler design is possible (see Fig. 31).

This simpler device, like its other relatives in the FROG family of techniques, has a frivolous amphibian name: GRating-Eliminated No-nonsense Observation of Ultrafast Incident Laser Light E-fields (GRENOUILLE, which is the French word for "frog"). It and its operating principles are shown in more detail in Figs. 32–34.[2,87]

GRENOUILLE first involves replacing the beam splitter, delay line, *and* beam combining optics with a *single* simple element, a Fresnel biprism[87] (see Figs. 31 and 32), which accomplishes all these tasks by itself. Second, in seemingly blatant violation of the SHG phase-matching-bandwidth requirement, it also involves replacing the thin SHG crystal with a *thick* SHG crystal (see Figs. 31 and 33), which not only gives considerably more signal (signal strength scales as the approximate square of the thickness) but also simultaneously replaces the spectrometer.

How does GRENOUILLE work? Consider the Fresnel biprism first (see Fig. 32). It is a prism with an apex angle close to 180°, whose refraction crosses the two resulting beamlets at an angle—exactly what is required in conventional single-shot autocorrelator and FROG beam geometries, in which the relative beam delay is

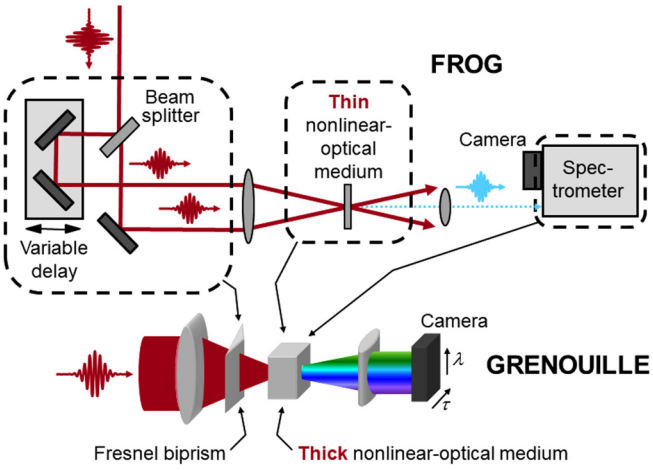

FIG. 31. SHG FROG device and its simpler version, GRENOUILLE. While SHG FROG is a fairly simple device, there are a few components of it that can be replaced with even simpler ones. Reproduced with permission from www.frog.gatech.edu.

mapped onto horizontal position at the crystal. However, unlike conventional single-shot geometries, beams that are split and crossed by a Fresnel biprism are *automatically aligned* in space and time, a significant simplification. Then, as in standard single-shot geometries, the crystal is imaged onto a camera, where the signal is detected vs position (i.e., delay) in, say, the horizontal direction.

FROG also involves spectrally resolving a pulse that has been time-gated by itself. GRENOUILLE combines both of these operations in a single *thick* SHG crystal.[87,88] As usual, the SHG crystal performs the self-gating process: the two pulses cross in the crystal with variable delay. However, in addition, the thick crystal has a relatively small phase-matching bandwidth, so the phase-matched

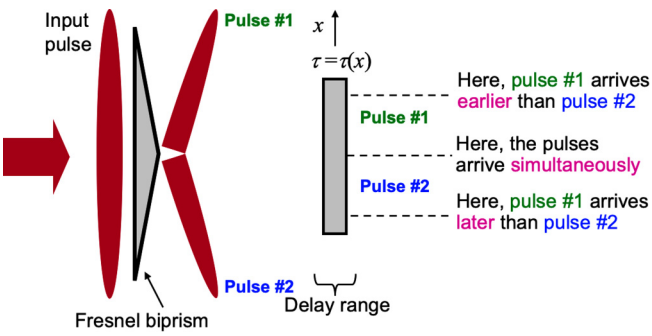

FIG. 32. The Fresnel biprism (a prism with a near-180° apex angle) refracts each half of the pulse by different amounts, crossing them at the crystal. This crossing maps delay onto transverse position and yields a single-shot autocorrelation, which, when spectrally resolved in the perpendicular direction and imaged onto a camera, yields a single-shot SHG FROG device. Reproduced with permission from www.frog.gatech.edu.

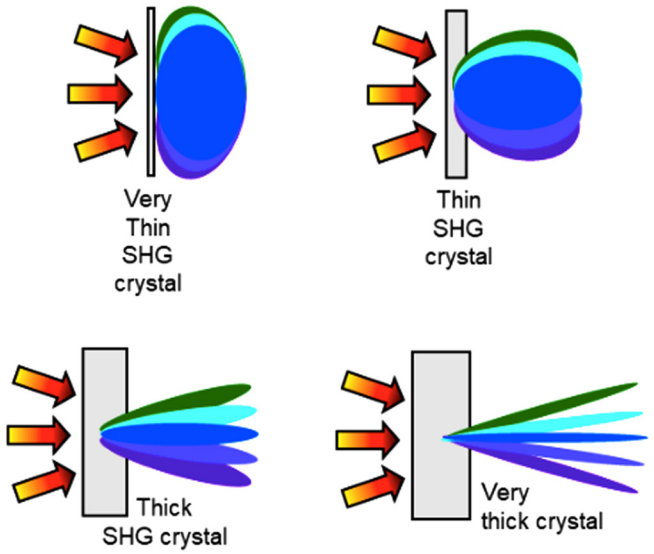

FIG. 33. SHG crystals of various thicknesses illuminated by converging broadband light and polar plots (in green through violet) of the generated colors vs crystal exit angle. Note that the very thin crystal (ordinarily required in pulse-measurement techniques) generates the second harmonic of all colors in the forward direction. The very thick crystal, on the other hand, generates a small range of wavelengths in each direction and, in fact, acts like a dispersive element used in a spectrometer. Note also that the thick crystal generates considerably more SH in the relevant directions. Reproduced with permission from www.frog.gatech.edu.

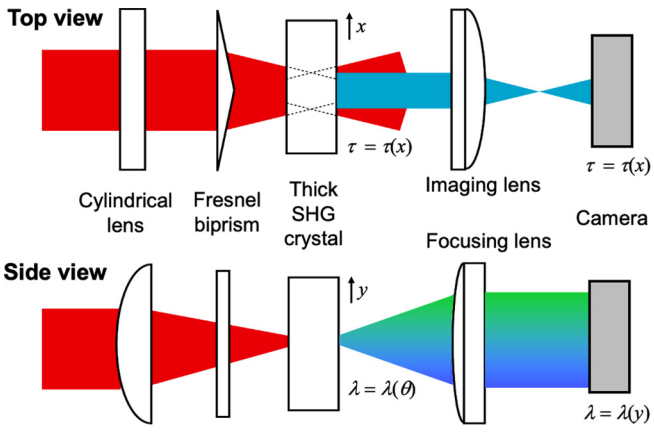

FIG. 34. Side and top views of the GRENOUILLE beam geometry. Typically, the anamorphic lens between the crystal and camera has a focal length of f in the vertical direction to map angle and wavelength out of the crystal onto vertical position at the camera and $f/2$ in the horizontal direction to image delay at the crystal onto horizontal position at the camera. Note that the beam becomes a vertical line just before the camera, a convenient place for a slit to filter out any extraneous beams, ensuring optimal signal-to-noise ratio. Reproduced with permission from www.frog.gatech.edu.

wavelength produced by it varies with angle (see Fig. 38). Thus, the thick crystal also acts as the dispersive element of a *spectrometer*. The ability of a thick nonlinear-optical medium to act as a low-resolution spectrometer was realized many years ago,[89,90] but pulses then were longer and more narrowband, so its rediscovery and use in pulse measurement had to wait until pulse bandwidths increased and pulse lengths significantly decreased.

Two additional cylindrical lenses complete the device (see Fig. 34). The first cylindrical lens must focus the beam into the thick crystal tightly enough to yield a range of crystal incidence (and hence exit) angles large enough to include the entire spectrum of the pulse. After the crystal, a cylindrical lens then maps the crystal exit angle onto position at the camera, with wavelength a near-linear function of (vertical) position.

GRENOUILLE has many appealing characteristics. It has very few elements and so is inexpensive and compact, with no alignment required. It naturally operates in single shot. Due to its thick crystal, it is more sensitive than other pulse-measurement devices. Furthermore, because GRENOUILLE produces (in real-time, directly on a camera) traces identical to those of SHG FROG, it yields the full pulse intensity and phase (except the direction of time and other trivial ambiguities). In addition, several feedback mechanisms on the measurement accuracy that are already present in the FROG technique also work with GRENOUILLE, allowing confirmation of—and confidence in—the measurement. Finally, GRENOUILLE is extremely simple to set up and align: it involves no beam-splitting, no beam-recombining, and no scanning of the delay, and so has *zero* sensitive alignment degrees of freedom. GRENOUILLE also measures first-order spatiotemporal distortions.[91,92]

GRENOUILLE can also measure relatively long (ps) pulses; indeed, it works even better for such pulses. Use of a pentagonal SHG crystal combines the Fresnel biprism and crystal into one component, yielding a device with only three optical elements. Using such a device, pulses up to 20 ps have been measured.[93]

Extremely short pulses, on the other hand, prove difficult for GRENOUILLE. Such pulses lengthen in the biprism and first lens, but simple theoretical backpropagation of the pulse through these elements removes this lengthening. Alternately, an all-reflective GRENOUILLE can be built, using a "Fresnel bi-mirror,"[94] allowing the measurement of pulses as short as ~15 fs. However, for shorter, and hence broader-band, pulses, the crystal group velocity dispersion begins to approach the crystal group velocity mismatch, and a GRENOUILLE cannot be designed for such a measurement.

H. Measuring two pulses simultaneously

Quite frequently, labs require two different-color pulses for an application, such as an excite-probe chemistry experiment, or a pulse at one wavelength is generated by a laser, and another pulse at a new, more useful wavelength is generated from it for the desired application. One could build two FROGs, one for each wavelength, but a single device would certainly be preferred, especially one that can be modified from an already existing excite-probe beam geometry.

This task can most easily be accomplished using a setup called Double Blind FROG. This name comes from the mathematical problem called Blind Deconvolution, in which two unknown

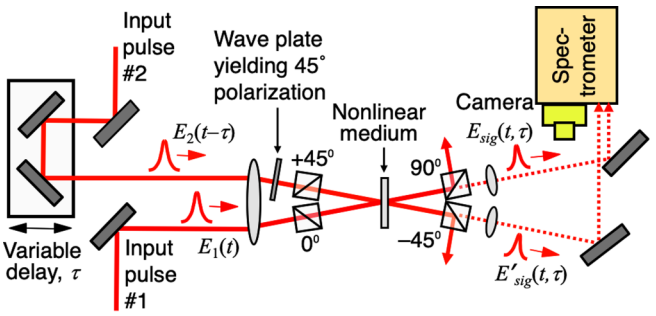

FIG. 35. Beam geometry for the measurement of two different-color pulses using Double Blind FROG. Note that each pulse gates the other, yielding two separate PG FROG traces. Reproduced with permission from www.frog.gatech.edu.

functions are determined from their convolution. For one-dimensional functions, this deconvolution is obviously not possible, but, unintuitively, it is often possible for two-dimensional problems.[95,96] Alas, despite the fact that FROG pulse retrieval is a two-dimensional problem, it is not possible in FROG. However, if *two* traces are generated, it not only becomes possible but also proves to be very reliable and the corresponding apparatus not too complicated—and easily constructed from an existing excite-probe beam geometry.

This technique is called Double Blind FROG.[97,98] It works by one pulse gating the other in a PG XFROG beam geometry, while, simultaneously, the other gates the one in the same geometry, but viewed from a 45° angle (see Fig. 35)

The pulse-retrieval algorithm uses the XFROG algorithm with one pulse treated as the known reference pulse on odd-numbered iterations and the other treated as the known reference pulse on even-numbered iterations. This technique and its accompanying algorithm work extremely well and can measure extremely short pulses due to the polarization-gating beam geometry's extremely large bandwidth. Double Blind FROG can also measure pulses at two completely different wavelengths.[99]

I. Other self-referenced methods

FROG was the first and remains the most powerful and popular pulse-measurement technique, so it has been described in detail here (and it is the subject of an entire book[2]). Since FROG's introduction, however, many additional methods have emerged. Many are simply variations on autocorrelation and so are not useful, except in unique circumstances. Most are interferometric, and so will be considered in Sec. V, specifically devoted to such approaches.

Many, however, are not known to be reliable as yet or can only measure very simple pulses. Others are very complicated and/or very expensive. Some, for example, require the use of a pulse shaper, which is a very expensive piece of equipment, and something whose performance requires an additional, separate pulse-measurement device to check.

Many others, as we will see shortly, measure *only* the coherent artifact. While lasers are more stable now than in the 1960s,

many are *not* stable, and simple misalignment, a rickety platform, or an overly powerful or unstable pump laser can turn an otherwise stable laser into an unstable one. Also, lasers at the edge of technology, such as those emitting near-single-cycle pulses, are often unstable and so absolutely require a check on their output pulse stability.

In addition, mode-locking instabilities have often been observed in, for example, dye lasers due to their short upper-state lifetimes. This problem was considered to be solved with the advent of much more stable solid-state lasers. However, semiconductor lasers or passive laser combs lack the stabilizing action of a long upper-state lifetime and have, therefore, often shown indications of a coherent artifact, and their autocorrelation measurements have often been misinterpreted.

A technique that only measures the coherent artifact cannot distinguish a stable train of short simple pulses from an unstable train of long complicated pulses. Not being able to distinguish the typically *most* desirable situation from the *least* desirable one is the worst-case scenario for any type of measurement device. Such methods are unacceptable and so should not even be considered. They will not be treated here, except to show how the coherent artifact measurement can (and often does) occur in pulse measurement and to illustrate this point.

Before we continue, it is important to remind the reader that simple and obvious standards for pulse-measurement techniques have been established.[80] In our earlier discussion of them, we mentioned that the method should indicate pulse-shape instability if it is present, but we can now be more specific: if there is even a remote possibility that pulse-shape instability might be present, it is extremely important that the pulse-measurement method *not* yield only the coherent artifact. Unfortunately, interferometric methods are likely to measure only the coherent artifact.

V. SPECTRAL INTERFEROMETRY

FROG and its variations work well for a wide variety of pulses, but they fail for very weak pulses due to the required nonlinear-optical process if a more powerful reference pulse is not available. Also, for very complicated pulses, FROG works, but the required data collection and pulse retrieval can be very slow. In these cases, methods using a FROG-measured reference pulse are useful, and XFROG is often used in this case. However, if a reference pulse train is measured to be stable, then a method that measures only the coherent artifact (as is the case for most interferometric methods) can be acceptable, as features such as fringe visibility can then confirm the unknown pulse's stability.

One such technique is spectral interferometry (SI)[100] (see Fig. 36), first introduced by Froehly and co-workers.[100] SI is, in principle, very simple, and it can measure a light pulse's intensity and phase for very complicated (i.e., large TBP) pulses, reaching TBP values of ~65 000, or very weak pulses with energies of zeptojoules or less.

SI requires a completely characterized reference pulse (in intensity and phase) whose spectrum contains that of the unknown pulse. Fortunately, weak and/or complicated pulses usually do not exist "in a vacuum." The processes that create them, whether a pulse shaper or propagation through a plasma or biological tissue,

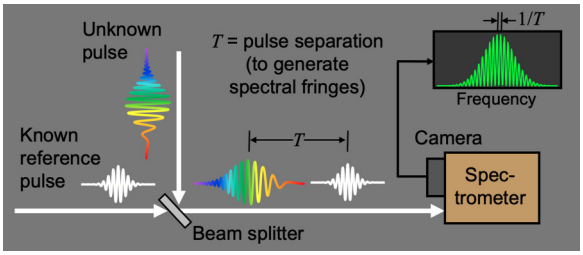

FIG. 36. Basic SI experimental setup. The reference and unknown beams are recombined collinearly and then sent into a spectrometer. Because the pulses expand in time inside the spectrometer, they interfere even in the presence of a delay, yielding spectral fringes. Reproduced with permission from www.frog.gatech.edu.

generally begin with a much stronger, simpler pulse with a spectrum that contains, and often is the same as or at least contains, that of the unknown pulse as the input. Thus, a more intense, simpler reference pulse, which can be measured using FROG, is often available to use to measure the more difficult pulse. The combination of FROG and SI is often called TADPOLE or Temporal Analysis by Dispersing a Pair Of Light E-fields.[101]

SI involves simply measuring the spectrum of the sum of the unknown pulse and a known pulse. The resulting spectral interferogram is given by

$$S_{SI}(\omega) = S_{ref}(\omega) + S_{unk}(\omega) + 2\sqrt{S_{ref}(\omega)}\sqrt{S_{unk}(\omega)}\cos[\varphi_{unk}(\omega) - \varphi_{ref}(\omega) + \omega T], \quad (15)$$

where $S_{ref}(\omega)$ and $S_{unk}(\omega)$ are the reference and unknown pulse spectra, $\varphi_{ref}(\omega)$ and $\varphi_{unk}(\omega)$ are their spectral phases, and T is the delay between the two pulses.

The spectral phase is encoded in the phase of the interference fringes and is easily retrieved.

The most common method for retrieving a pulse from its SI trace takes advantage of the delay between the reference and unknown beams. The spectral interferogram is Fourier transformed with respect to frequency to a "pseudo-time" domain ("pseudo" because it is only being used for data reduction, and the resulting quantity has no simple physical meaning). The delay is chosen so that, in the Fourier domain, the different components of the interferogram are well separated along the pseudo-time axis. As a result, the first two terms in Eq. (15), the "DC" terms, are centered at $t = 0$. The cosine separates into its positive and negative components, called the "AC" terms, each centered at $\pm T$. Only one of the two "AC" terms is retained and is then translated to $t = 0$ and inverse Fourier transformed. What results is the product of the reference and unknown fields,

$$\sqrt{S_{ref}(\omega)}\sqrt{S_{unk}(\omega)}\exp\{i[\varphi_{unk}(\omega) - \varphi_{ref}(\omega)]\}, \quad (16)$$

(or the complex conjugate) which, after dividing out by the reference field (or its complex conjugate), yields the unknown field

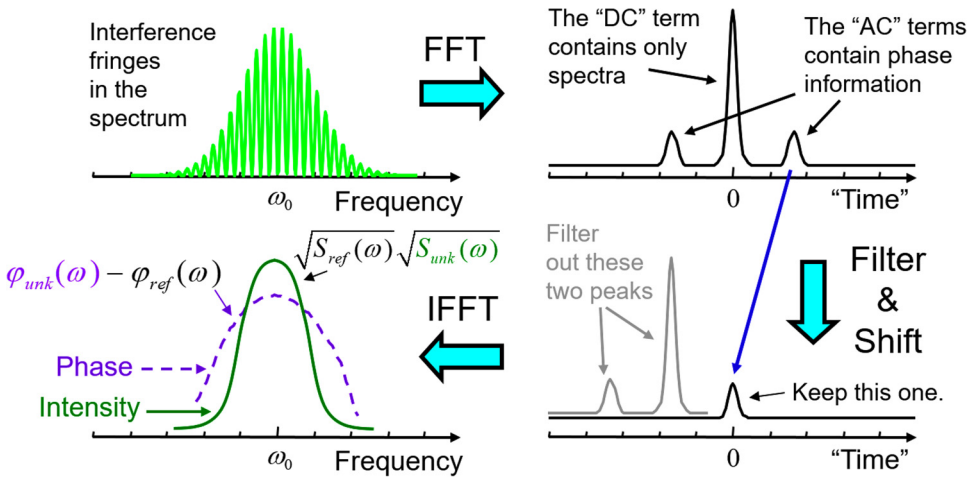

FIG. 37. FTSI concept. The spectral interferogram (top, left) is Fourier transformed from ω to a pseudo-time domain, t. A sufficiently large delay T is needed so that there are three well-separated nonzero regions (top right). The two sidebands are the phase-containing component of the data, so either one of these is retained (bottom, right), translated to $t = 0$, and inverse Fourier transformed, yielding the unknown pulse's spectrum and spectral phase (bottom, left). Reproduced with permission from www.frog.gatech.edu.

spectrum and spectral phase. When using this retrieval algorithm, the approach is often called Fourier transform spectral interferometry (FTSI) (Fig. 37).[101]

A. Advantages and disadvantages of spectral interferometry

SI uses an inherently single-shot geometry, and its retrieval algorithm is fast and direct. Also, SI is a "homodyne" method because what is measured is the squared magnitude of the sum of the reference and unknown (spectral) fields. This means that, because the key measured quantity is the cross term of this squared magnitude, which is a product of the two fields, a strong reference field can effectively amplify a weak unknown field. Choosing the reference pulse to be M times more intense than the unknown pulse produces peak-to-peak spectral fringes that are $4\sqrt{M}$ times as intense as the spectrum of the unknown pulse; for $M = 100$, the fringes are 40 times more powerful than the unknown pulse spectral intensity in the absence of the reference pulse. As a result, for a 100 MHz rep-rate pulse train, SI was able to measure a pulse with only 42 zJ of energy or 1/5 of a photon per laser pulse,[101] and pulses weaker by a factor of 1000 should be measurable using it, provided that the pulse train is stable. See Fig. 38 for these results. Notice that the spectral fringe contrast is easily visible even though the unknown pulse is much weaker than the reference pulse.

It should always be remembered that, in the presence of pulse-shape fluctuations, SI measures only the coherent artifact and so does not see fluctuations in a pulse train, whose fringes wash out. This is fine when there is perfect stability, as in the case above, when the coherent artifact is equal to the pulse itself (in which case one does not refer to the retrieved pulse as an "artifact"). As a result, it is very important to measure the reference pulse using a technique that establishes its stability, such as FROG. Then, a 100% SI fringe visibility confirms the stability of the unknown pulse as well (except for unequal beam intensities and/or misalignments that reduce the fringe visibility, and which can be accounted for). This statement can be said for all variations of SI and also for holographic pulse-measurement techniques to be discussed later. Fortunately, this requirement is easily accomplished when the reference pulse is strong and hence measurable with, for example, FROG.

Also, in this form, SI requires perfectly collinear beams, with perfectly matched spatial modes. This makes SI very difficult to align, despite its seemingly trivial-looking apparatus. To see this, consider that crossing two pulses with large beams at some noncritical angle has only one sensitive alignment degree of freedom (delay), whereas making them collinear has five (delay, two angles, and two spatial coordinates). Fortunately, there are simpler alternatives, and many involve crossed beams instead of collinear ones.

But first, we will consider a recent implementation of SI using optical fibers, known as STARFISH (Spatio-Temporal

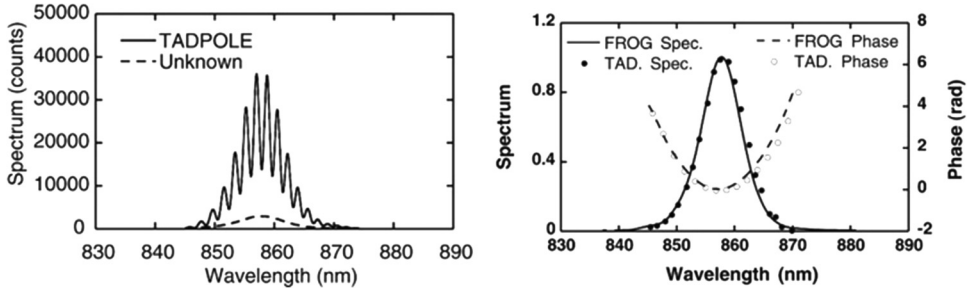

FIG. 38. Left: Spectral interferogram for a 42 zJ pulse, and the spectrum of the pulse (here labeled "Unknown"). Right: Retrieved spectrum and spectral phase. The fringe visibility here is not 100% due to unequal intensities of the reference and unknown pulses, not instability. Reproduced with permission from Fittinghoff *et al.*, Opt. Lett. **21**, 884 (1996). Copyright 1996 The Optical Society.

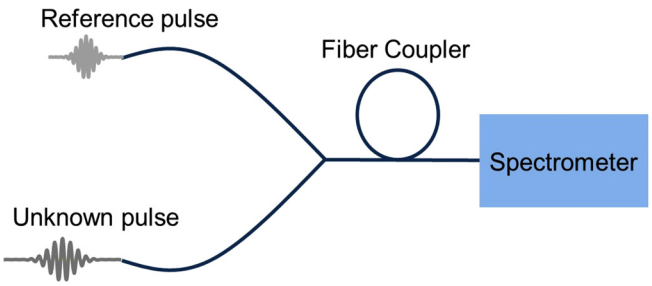

FIG. 39. STARFISH: This very convenient experimental setup for SI uses a fiber spectrometer, a fiber coupler and two fibers to introduce the reference and unknown pulses. Reproduced with permission from www.frog.gatech.edu.

Amplitude-and-phase Reconstruction by Fourier transform of Interference Spectra of High-complexity-beams), that solves most of these problems, provided that the beam is large and only a small fraction of it needs be considered[102] (see Fig. 39). The reference and unknown pulses are coupled into two, equal-length fibers whose outputs are combined with a fiber coupler and sent into the spectrometer. This setup ensures collinear beams and matched modes. If the fibers have equal lengths, both the reference and unknown pulses see identical dispersions, which cancel out and so can be ignored. In practice, it is not possible to make the lengths exactly the same, but the small residual difference can be measured and then subtracted from all measurements. The fibers are extremely helpful, since the beams can now move around without affecting the alignment of the spectral interferometer.

The main drawback of STARFISH is the required spectral fringes and the resulting loss of spectral resolution due to the need for spectral filtering in the pulse-retrieval algorithm.[103] Typically, at best, only about 1/5 of the spectral resolution is retained. This drawback severely limits STARFISH's (and SI's) ability to measure even moderately complicated pulses.

B. Crossed-beam spectral interferometry

To avoid this loss of spectral resolution, a variation of SI instead makes spectrally resolved *spatial* interference fringes. This approach involves a very convenient experimental setup and also solves SI's alignment problems. In addition, it yields spatial resolution in one dimension.[104–108] It is shown in Fig. 40.

The SI spectrometer's usual linear detector is replaced with a camera. Now the idea is to make the interference fringes along the spatial dimension instead by setting the reference-unknown pulse relative delay to ∼0 and crossing the two beams at a small angle. The crossing angle plays the same role as the delay in usual SI. After Fourier transforming, now along the *spatial dimension,* the "AC" and "DC" terms separate due to their different *spatial frequencies*. This step leaves the data along the frequency axis unchanged so that the unknown pulse is reconstructed with the spectrometer's full spectral resolution.

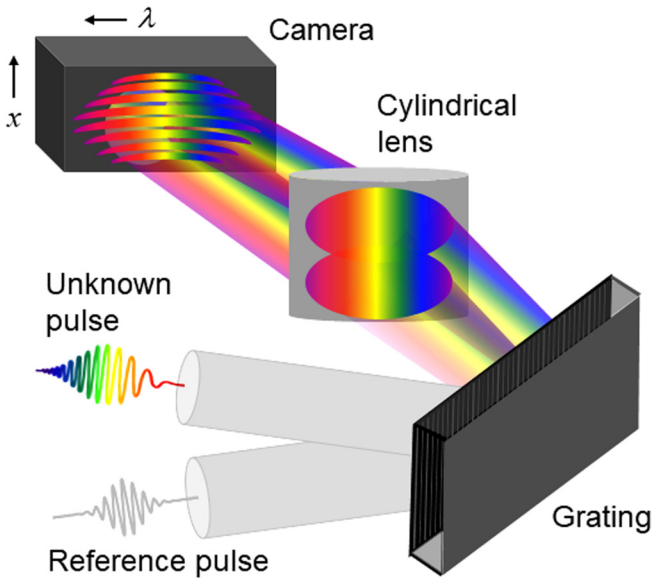

FIG. 40. Crossed-beam spatial interferometry. The reference and unknown beams cross at the camera of a spectrometer, or equivalently at the slit of an imaging spectrometer. Mirrors in one of the two beam paths are tilted to make the beams cross. The spatial interference of the two beams at each frequency within the pulse is recorded. From this measurement, $\tilde{E}_{unk}(\omega)$ can be reconstructed with the full resolution of the spectrometer. It also yields spatial resolution in one dimension (across the fringes). Reproduced with permission from www.frog.gatech.edu.

The equation for a spectrally resolved spatial interferogram is shown below,

$$S(\omega, x) = S_{ref}(\omega) + S_{unk}(\omega) + 2\sqrt{S_{ref}(\omega)}\sqrt{S_{unk}(\omega)}\cos[\varphi_{unk}(\omega) - \varphi_{ref}(\omega) + 2kx\sin\theta]. \quad (17)$$

In Eq. (17), θ is the crossing half angle, k is the wavenumber, and x is the vertical dimension in which the beams are crossing (perpendicular to the propagation direction). Figure 41 shows such interferograms for two different pulse shapes, assuming an identical spectrum for the reference and unknown pulses and flat spectral phase for the reference. What is especially convenient about these interferograms is that the curvature of the interference fringes is proportional to the spectral phase difference. So, qualitatively, the spectral phase of the pulse can be immediately seen by the eye. This feature has been used for real-time adjustments of a pulse compressor. Curve fitting has been applied to very accurately extract the unknown pulse from the spatially resolved spectral interferogram,[106,109] and it has the advantage that, if some spatial information is present, there is no resolution loss along this dimension.

Figure 42 shows the pulse-retrieval algorithm for crossed-beam spectral interferometry.

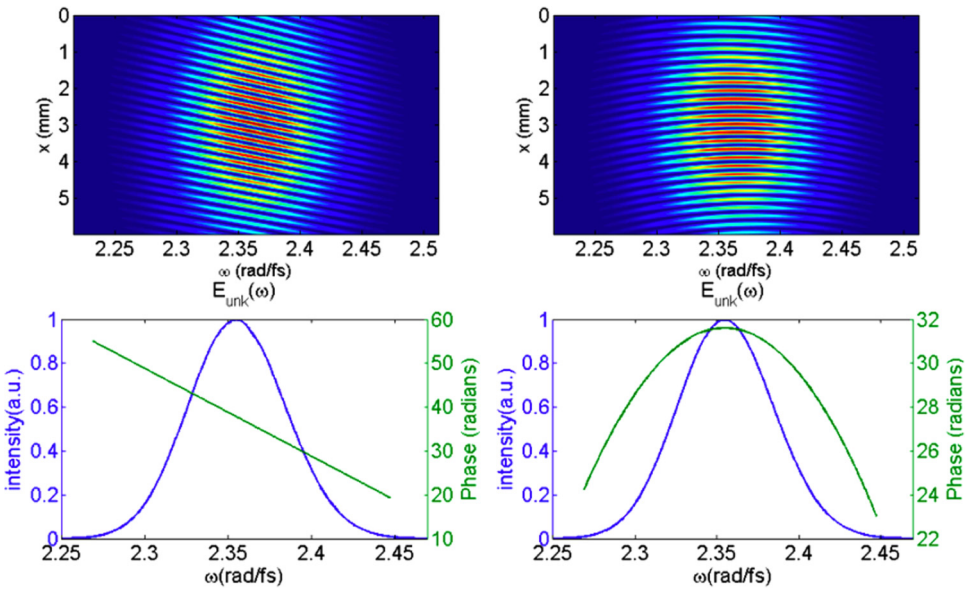

FIG. 41. Simulations of crossed-beam spatial interferograms for two different unknown pulse shapes, assuming a reference pulse with a flat spectral phase and the same spectrum. Notice that the shape of the fringes mirrors the spectral-phase difference. Reproduced with permission from Bowlan, "Measuring the spatiotemporal electric field of ultrashort pulses with high spatial and spectral resolution," Ph.D. thesis (Georgia Institute of Technology, 2009).[162]

C. Practical measurement of weak and complicated pulses: SEA TADPOLE

Spatially Encoded Arrangement for Temporally Dispersing a Pair of Light E-fields (SEA TADPOLE)[106,110] is another experimentally simple version of crossed-beam spectral interferometry that uses optical fibers to ensure perfect mode-matching and a lens to ensure collimated beams (see Fig. 43) entering a simple imaging spectrometer. Like other crossed-beam variations on SI, it avoids alignment problems and is able to use the entire spectrometer resolution.

In SEA TADPOLE, the reference and unknown beams are coupled into approximately equal-length fibers. At their output ends, the fibers are mounted with a small gap between them, and the diverging beams emerging from them are collimated with a single spherical lens. Because of the fibers' displacement from the optic axis, the beams are collimated, but one propagates downward, and the other upward. Eventually, the beams cross, resulting in spatial interference fringes. A camera is placed at the beam's crossing point. In the other dimension, a diffraction grating and a cylindrical lens map each wavelength to a different camera position (that is, act as a spectrometer). The Fourier filtering algorithm described earlier is used to reconstruct the pulse, but now the filtering occurs in the spatial domain, rather than the spectral domain.

SEA TADPOLE is very easy to use in practice, collinear beams are unnecessary, and a low-resolution home-made spectrometer is used. It has measured complicated pulses with TBPs as high as 100,[105] with the only limitation being the spectrometer. An example of a SEA TADPOLE measurement of the spectrum and phase added by a pulse shaper is shown in Fig. 44. Note the intuitive shape of the fringes, mirroring the actual spectral phase of the pulse.

The bottom left plot of the spectra reveals a very interesting feature of such measurements. The spectrum in green was that obtained from the SEA TADPOLE measurement. The spectrum in magenta was obtained by simply blocking the reference beam and measuring the unknown pulse spectrum using the same

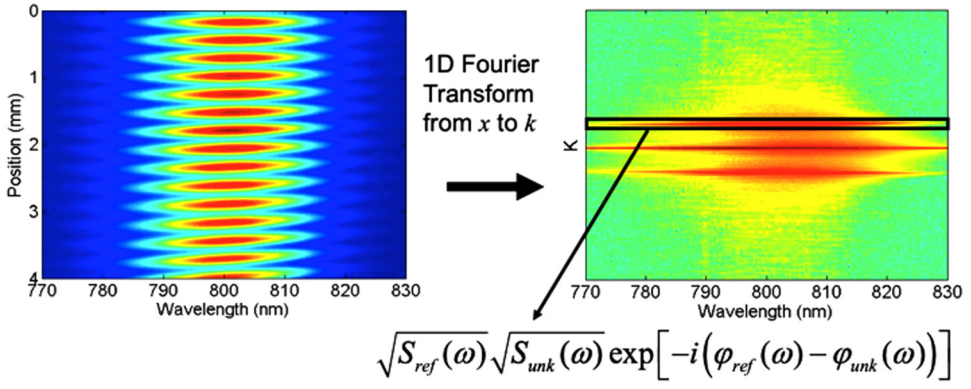

FIG. 42. Fourier filtering along the spatial axis. Take a 1D Fourier transform to k_x where the data separate into three bands. Use either of the sidebands, which contain the phase information, and inverse Fourier transform to obtain $\tilde{E}_{ref}(\omega)\tilde{E}_{unk}(\omega)^*$ without any loss of spectral resolution. Reproduced with permission from Bowlan *et al.*, Opt. Express **14**, 11892 (2006). Copyright 2006 The Optical Society.

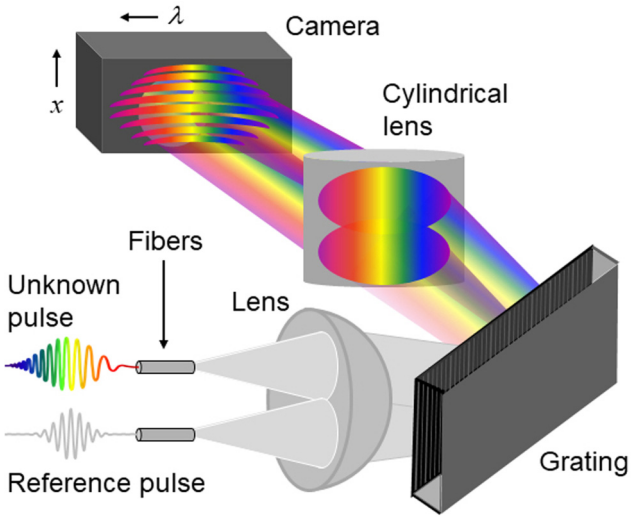

FIG. 43. SEA TADPOLE experimental setup. Reproduced with permission from www.frog.gatech.edu.

spectrometer, but now acting as a standard spectrometer, rather than a spectral interferometer. Note that, despite the use of the same spectrometer, the SEA TADPOLE-measured spectrum has considerably higher spectral resolution.

Such high spectral resolution is due to the fact that SEA TADPOLE measures the spectral *field*, not the spectral *intensity*, and convolutions of the spectrometer point-spread function with potentially negative functions, like the spectral field, do not broaden as much as positive-definite functions, like the spectral intensity.[105]

1. Measuring very complicated pulses in time: MUD TADPOLE

At each pixel on the detector of an SI spectrometer, the pulses stretch in time and have the duration of the inverse of the spectrometer spectral resolution. So, they temporally overlap for this amount of time. Because two pulses can only interfere when temporally overlapping, the task of measuring longer, more complicated pulses using any version of SI, including SEA TADPOLE, would, in principle, require a higher-resolution spectrometer, in order to further broaden the reference pulse. If the unknown pulse is longer than the inverse of the spectrometer resolution, then a

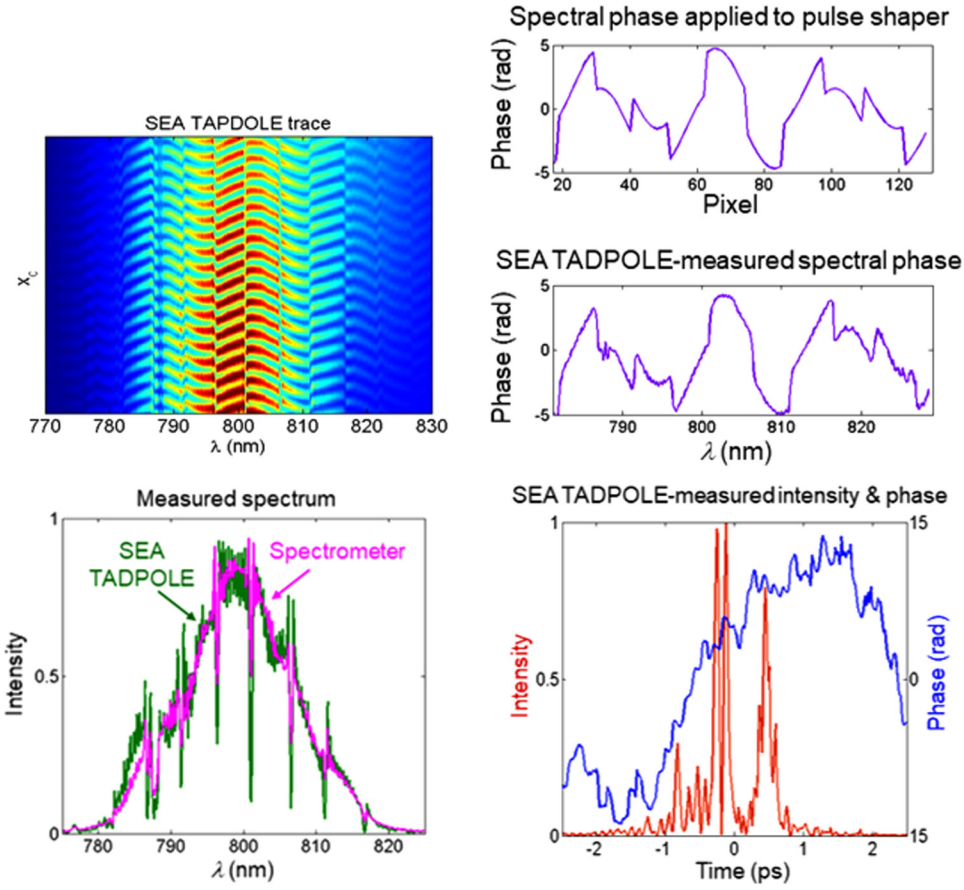

FIG. 44. SEA TADPOLE measurement of a shaped pulse with a TBP of 30. The measured interferogram is shown at the top left. The top, right plots are the extracted spectral phase introduced by the pulse shaper, compared to that applied, showing good agreement, and having the same shape as the interference fringes. The temporal intensity and phase are shown at the bottom, right. The spectrum retrieved from the SEA TADPOLE measurement is shown in green (lower left). Also shown in the power left plot, in magenta, is the spectrum measured directly with the SEA TADPOLE spectrometer, obtained simply by blocking the reference beam. Note that the SEA TADPOLE-measured spectrum is better resolved. Reproduced with permission from Bowlan *et al.*, J. Opt. Soc. Am. B **25**, A81 (2008). Copyright 2008 The Optical Society.

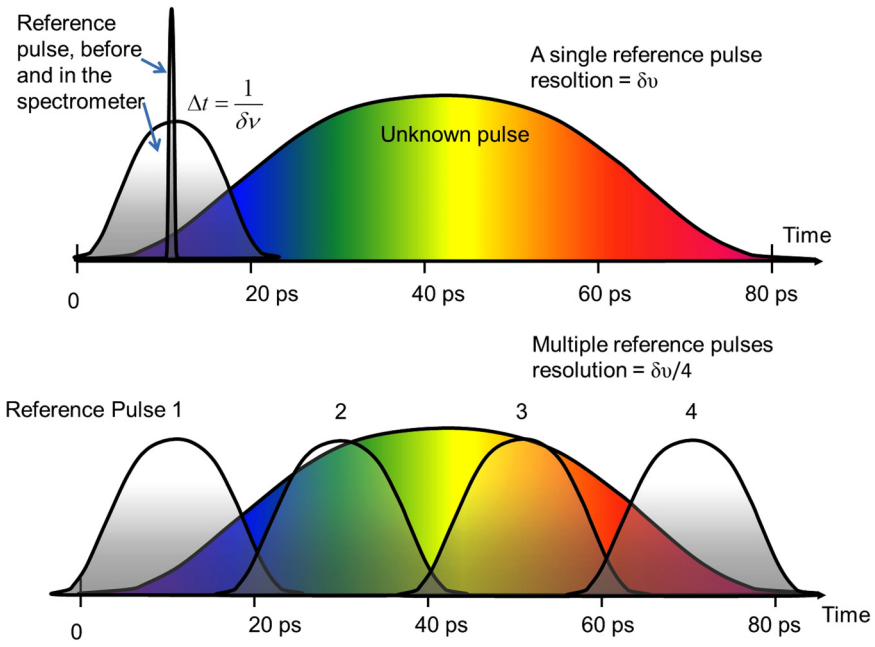

FIG. 45. In a spectrometer, two pulses can only interfere with each other for as long as they overlap, which is determined by the inverse of the spectral resolution. *N* reference pulses could then be used to extend the spectral resolution by *N* times, as depicted here. The color in the unknown pulse represents the instantaneous frequency, showing that it has a complicated spectral phase. Reproduced with permission from www.frog.gatech.edu.

single reference pulse will not be able to interfere with all of it. The key to overcoming this limitation, that is, to beating the spectrometer's resolution, is to use *multiple*, delayed reference pulses. The use of a reference pulse with multiple delays is the key idea of MUltiple Delay TADPOLE (MUD TADPOLE).[111–113] This idea is illustrated in Fig. 45.

This idea is straightforward to implement in a multi-shot scanning setup[111–113] and was accomplished using SEA TADPOLE for the individual measurements. Then, in the time domain, these temporal pieces are concatenated to yield the entire unknown pulse, which will now be resolved with a spectral resolution given by the inverse of the scanning range. Figure 46 illustrates this process.

MUD TADPOLE has succeeded in measuring a pulse with a TBP of 65 000.

The limiting factor in multi-shot MUD TADPOLE is the dynamic range of the camera. The camera only sees fringes from the temporal region of the unknown pulse that temporally overlaps with the reference pulse. All other temporal regions yield only relatively constant background. For long enough unknown pulses, the fringes will no longer be discernable.

The main disadvantage of this setup for MUD TADPOLE is that it is multi-shot and so is not applicable to pulse trains in which every pulse is different. Also, the measurable spectral width is still limited by the spectrometer, and, as with other SI-based approaches, the reference pulse spectrum must contain the unknown pulse spectrum.

2. Single-shot MUD TADPOLE

It is straightforward to extend the multiple-reference pulse idea to a single-shot measurement geometry. It turns out to be possible using a trick that is often employed for single-shot pulse measurement, tilting the pulse in order to achieve the required large range of delays. The variably delayed reference pulse overlaps with the unknown pulse at different heights.[114] Alternatively, different angles can correspond to different delays, yielding different fringe spacings for different temporal chunks of the pulse. In other words, the fringe spacing codes for the delay (see Fig. 47).

TBPs as large as 4500 have been measured on a single shot with this device,[114] a significant improvement over all other approaches. The main limitation is not yet the dynamic range of even an inexpensive eight-bit camera, but the pulse front tilt and hence the delay range producible by a diffraction grating. With other, more dispersive optics, or just larger diffraction gratings and cameras, single-shot MUD TADPOLE should be scalable to even more complicated pulses.

VI. SPECTRAL PHASE INTERFEROMETRY FOR DIRECT ELECTRIC-FIELD RECONSTRUCTION (SPIDER)

Frustrated by the slow convergence of the FROG algorithm on available computers in the 1990s, researchers sought a non-iterative, that is, "direct" self-referenced intensity-and-phase pulse-measurement technique. They looked to SI, which, as we have seen, has a direct pulse-retrieval algorithm. In its usual configuration, however, SI cannot be self-referencing because it measures only phase differences, and, if the unknown pulse interferes with itself, the spectral-phase difference necessarily vanishes and is at most only the deliberately introduced delay between the two replicas of the pulse. Either way, it yields no pulse phase information.

However, *nonlinear-optical* variations of SI were introduced that overcame this limitation, allowing the measurement of pulses without the need for a reference pulse. The first and most popular

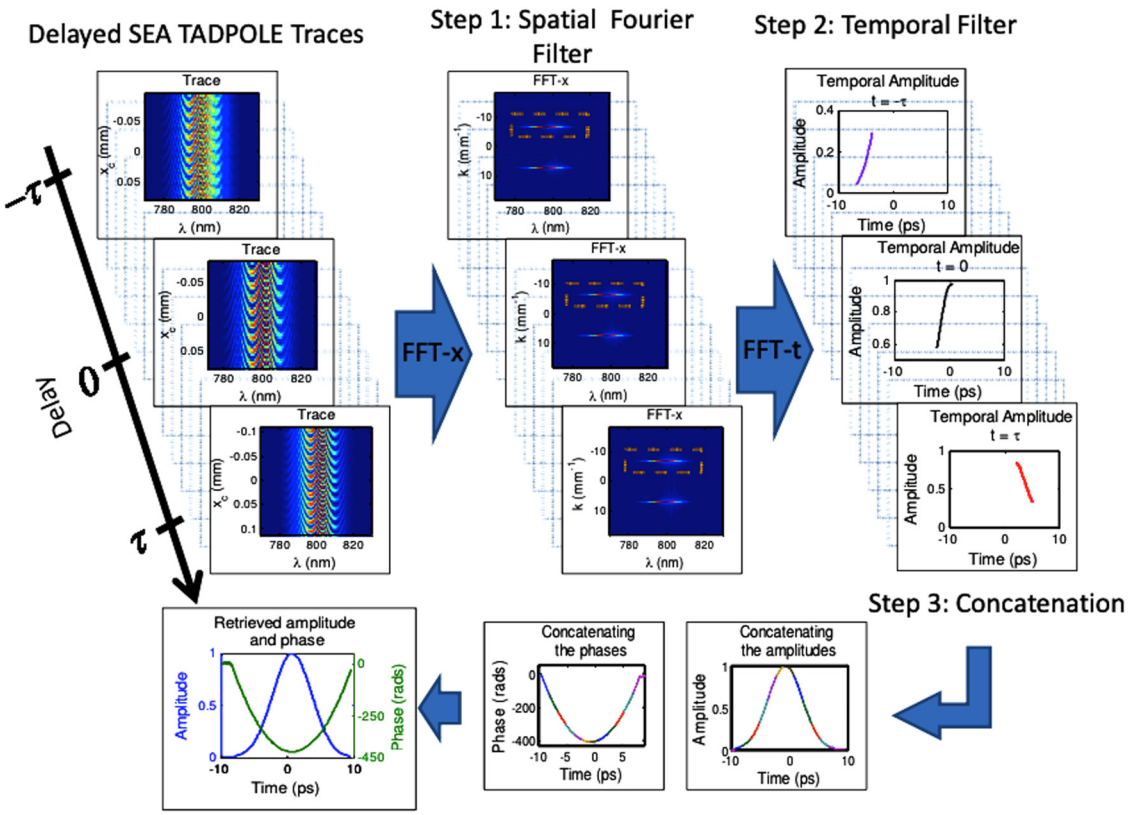

FIG. 46. Multi-shot MUD TADPOLE principle. The left column shows the SEA TADPOLE traces measured at a few different reference pulse delays. $E(t)$ is extracted from these just as in SEA TADPOLE (middle and right columns). The concatenation, done in the time domain, is illustrated at the bottom. The individual colors illustrate individual measurements. The final concatenated, reconstructed pulse is shown at the bottom, left. Reproduced with permission from Opt. Express **18**, 6583 (2010). Copyright 2010 The Optical Society.

has been Spectral Phase Interferometry for Direct Electric-field Reconstruction (SPIDER)[115] (see Figs. 48 and 49), and another is self-referenced spectral interferometry.[116]

SPIDER involves performing SI on the pulse and a *frequency-shifted* version of itself. This frequency shift $\delta\omega$ is called the spectral shear, and it can be induced in a self-referenced way by sum-frequency generation between two replicas of the pulse under test and a third strongly chirped replica of the pulse. With some simplification, the SPIDER trace is given by

$$S_{SPIDER}(\omega) = S(\omega) + S(\omega+\delta\omega) + 2\sqrt{S(\omega)}\sqrt{S(\omega+\delta\omega)}\cos[\varphi(\omega+\delta\omega) - \varphi(\omega) + \omega T], \tag{18}$$

where T is the delay between the two replicas. The SPIDER phase (the quantity inside the cosine) can be approximated by

$$\Delta\varphi_{SPIDER} \equiv \varphi(\omega+\delta\omega) - \varphi(\omega) + \omega T \approx \delta\omega\frac{d\varphi}{d\omega} + \omega T. \tag{19}$$

The derivative, $d\varphi/d\omega$, is the group delay, $\tau(\omega)$, of the pulse vs frequency. Consequently, any deviations of the fringe phase from linearity are due to variations in the pulse group delay.

SPIDER does not measure the spectrum and instead uses a separate, independent measurement of it. SPIDER's ambiguities are similar to those of FROG: the absolute phase, the pulse arrival time, and relative phases of multiple pulses and multiple modes. Like most FROG versions, it does not have an ambiguity in the direction of time.

The SPIDER phase $\Delta\varphi_{SPIDER}(\omega)$ between the two pulses can be analytically retrieved from the measured $S_{SPIDER}(\omega)$ interferogram by extracting the phase of its inverse Fourier transform (with a specified finite frequency range). This algorithm is known as the Takeda algorithm[117] for numerical phase demodulation. As in standard spectral interferometry, SPIDER relies on a Fourier filtering approach. Then, it isolates one of the modulation sidebands, e.g., the positive modulation sideband, from which one can simply extract $\Delta\varphi(\omega) = \omega T$ by applying the complex logarithm (or arg function) to this modulation term. An alternative method for reconstructing the complex-valued sideband is given by the Hilbert transform,[118] and a third method is the use of wavelets.[119] All

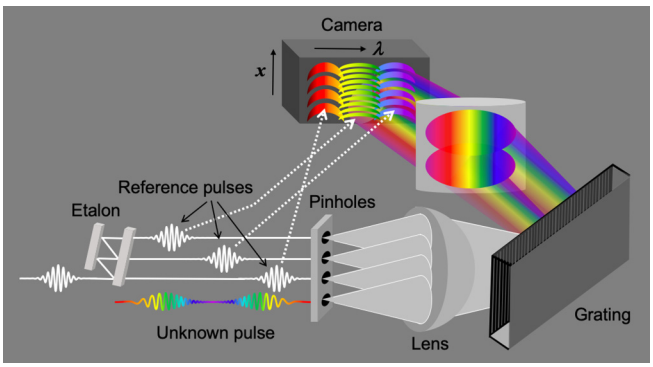

FIG. 47. Single-shot MUD TADPOLE. Multiple delays are generated using an etalon, and each delayed reference pulse interferes with a section of the unknown pulse at the camera. Note that, due to the spatial displacement introduced by etalon, each delayed reference pulse has different crossing angles with the unknown pulse, and hence yields a different fringe spacing. The unknown field is retrieved from multiple SEA TADPOLE traces, separable from each other on a single camera frame based on their fringe spacings, and concatenating the results obtained from each trace. Reproduced with permission from www.frog.gatech.edu.

approaches ultimately isolate the underlying spectral phase difference, removing the constant group delay corresponding to the interference of identical pulses, and any deviation from a linear relation vs ω indicates the SI of two pulses with different phases.

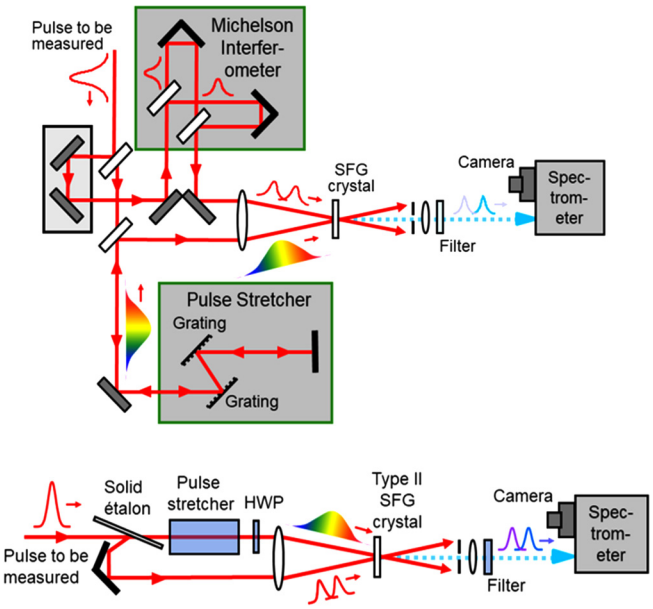

FIG. 48. SPIDER apparatus. Top: As originally proposed. Bottom: Experimentally simplified version in which the grating pulse stretcher has been replaced with a simple block of glass whose dispersion stretches the pulse. Also, a simple etalon replaces the Michelson interferometer. The latter design is used mainly for few-cycle pulses. Reproduced with permission from www.frog.gatech.edu.

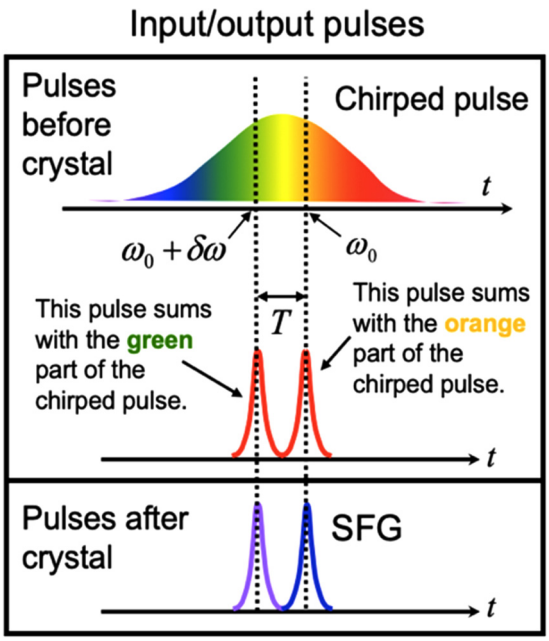

FIG. 49. SPIDER principle. The undistorted double pulse performs sum frequency generation with two different temporal regions and hence two different wavelengths of the chirped pulse. This procedure yields two pulses, one wavelength-shifted and delayed compared to the other. As a result, performing spectral interferometry with them yields a phase that is ideally the derivative of the spectral phase. Reproduced with permission from www.frog.gatech.edu.

While analytical pulse retrieval in SPIDER is very fast and may look very robust at first sight, there is nevertheless a hidden caveat: isolated application of Eq. (19) to a measured SPIDER trace typically does not lead to meaningful results without a calibration measurement. This calibration measurement is performed using the identical beam-splitting scheme for generating the two unchirped replicas but replacing the sum-frequency arrangement with second-harmonic-generation. This second measurement results in an interferogram $S'_{SPIDER}(\omega)$, which is processed as above to yield $\tau'(\omega)$, which then needs to be subtracted from the previously measured $\tau(\omega)$ to remove the constant group delay as well as any spectral calibration errors from the measurement.

Because SPIDER is usually used to measure very broadband pulses, higher-order corrections are typically required to directly obtain the spectral phase of a pulse from Eq. (18) without subtracting the reference phase. As these corrections affect the SPIDER measurement in the same way as the reference measurement, one can completely avoid a careful wavelength calibration of the spectrograph by the reference subtraction.

Ignoring any possible issues with the spectrograph calibration, the remaining source of error in a SPIDER measurement is the spectral shear, $\delta\omega$. To avoid the need for frequent recalibration of the setup, it is preferred to implement both the pulse separation T and the dispersion of the stretched pulse using solid-state components, e.g., employing an etalon and a glass block, respectively, as

shown in Fig. 48 (bottom). This arrangement provides two replica pulses and one chirped pulse with nearly identical pulse energies and has been shown to provide the optimum efficiency in the sum-frequency process.[120]

Other variants of SPIDER address some of the problems inherent in this method. The first improvement was zero-additional phase SPIDER (ZAP-SPIDER),[121] in which the role of the replicas and the chirped pulse is inverted, and two chirped pulses with a separation T interact with one short replica that can be generated by a surface reflection. Consequently, the latter does not experience any temporal stretching in beam splitter substrates or in an etalon, which makes ZAP-SPIDER a better approach for the ultraviolet range. If the SPIDER setup uses two chirped pulses, one can also replace the dispersive element by two narrowband interference filters, which filter out different spectral portions from the input pulse.[122] This replacement is one of the central ideas of two-dimensional spectral shearing interferometry (2DSI),[123] which additionally exploits the concept of spatial shearing interferometry, which is also utilized in the spatially encoded arrangement for SPIDER (SEA-SPIDER),[124] analogous to SEA TADPOLE. While traditional SPIDER only requires a one-dimensional detector array in the Fourier plane of the spectrograph, both 2DSI and SEA-SPIDER rely on two-dimensional interferograms, require area-scan rather than line-scan cameras, and an imaging spectrograph. The increased experimental complexity pays off in the averaging effect from the multiple parallel interferograms recorded in the individual lines of a CCD camera. Finally, another, less related, method is self-referenced SI,[116] which involves performing SI between the pulse field and its cube (which is usually shorter), which has been used on occasion to measure amplified pulses.

A SPIDER-measured pulse is shown in Fig. 50.

All multi-shot versions of SPIDER, however, have the same significant drawback: if there is any pulse-shape instability, *they measure only the coherent artifact.*[77–81] This problem occurs because SPIDER is inherently interferometric; random components in the pulse yield fringes that cancel out of the measurement, yielding only an approximately constant background, which is necessarily ignored in SPIDER pulse retrieval. Ignoring the background leaves only the stable component to yield SPIDER fringes and hence the retrieved pulse (see Fig. 51).

This drawback can alternatively be seen by realizing that SPIDER measures the average spectral phase of the pulses in a train. Long complicated pulses have complicated spectral phases. But if the pulses are all different, their average spectral phase tends to be a simple flat curve, which corresponds to the shortest pulse for a given spectrum. This average spectral phase is precisely the frequency domain description of the coherent artifact.

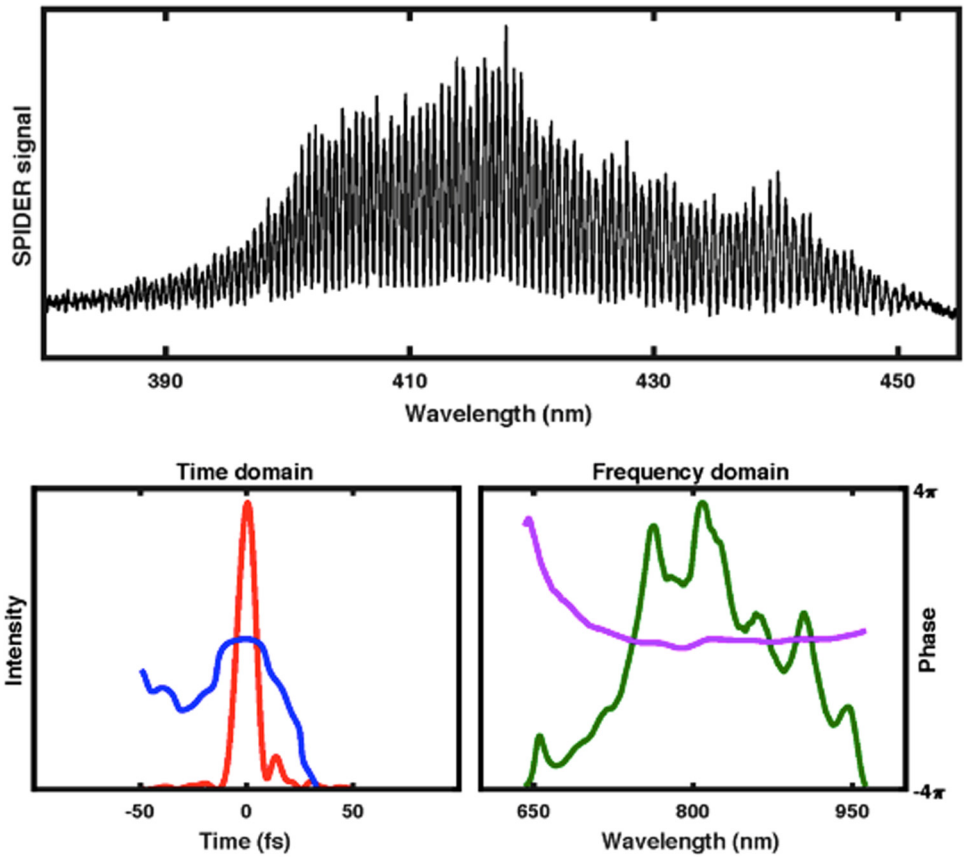

FIG. 50. SPIDER measurement. Top: Measured SPIDER trace. Bottom: Retrieved pulse. SPIDER trace provided by G. Steinmeyer and reproduced with permission from Hyyti *et al.*, Opt. Lett. **42**, 2185 (2017). Copyright 2017 The Optical Society.[163]

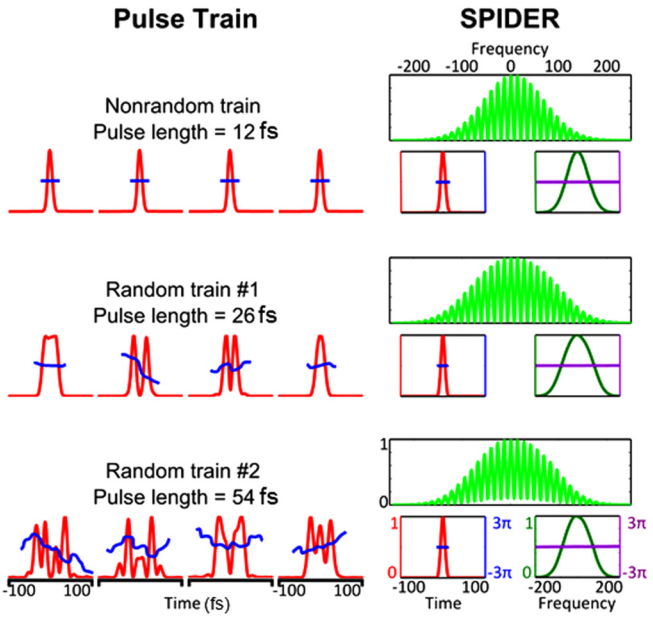

FIG. 51. Simulated SPIDER measurements of a stable and two unstable trains of pulses (the same trains used for analogous studies for autocorrelation and FROG earlier). Note that SPIDER only measures the coherent artifact (the stable component of the pulse). Therefore, SPIDER cannot distinguish a stable train of short simple pulses from an unstable train of long complicated pulses. In general, SPIDER yields only a lower bound on the pulse length. Reproduced with permission from Ratner *et al.*, Opt. Lett. **37**, 2874 (2012). Copyright 2012 The Optical Society.

The background could, in principle, yield clues as to the pulse train stability. Unfortunately, the background in SPIDER cannot be distinguished from identical background due to benign misalignment effects that are unavoidable in SPIDER devices. This ambiguity results in problems determining the stability and pulse length of pulses in a train of pulses. As shown in the simulations of Fig. 51, a mere 2% background corresponds to an under-estimate of the pulse length by a factor of 2.2, and 10% corresponds to an under-estimate of a factor of 4.5. As a result, like other techniques that only measure the coherent artifact, *SPIDER cannot distinguish a stable train of short simple pulses from an unstable train of long* complicated *pulses*.

Consequently, in general, SPIDER only yields a lower bound on the pulse length. Also, it does not see satellite pulses at all when their relative phase is random,[77] which is often the case when lasers double-pulse due to over-pumping—one of the main reasons for making pulse measurements. Unfortunately, readily available devices, such as spectrometers, also cannot distinguish between these diametrically opposed types of pulses. For example, spectral fringes due to unstable double pulses or variable complicated pulses also cancel out in measured spectra. In general, as mentioned earlier, the task of determining a pulse train's shape stability falls to the pulse-measurement technique.

Single-shot measurements could help, but variations from *place to place* in the beam could also give rise to similar effects (although the effect of these spatial variations of field on the measured spectrum by spectrometer has not been studied as yet), and SPIDER is rarely used for single-shot measurements.

There is a way to be certain that a SPIDER measurement is of a stable pulse train and is accurate, and that is if it has a 100% fringe visibility.[77–81] Unfortunately, SPIDER measurements rarely achieve this value, and the question always remains as to whether the background that is present is due to benign misalignment issues or instability. A mere 2% background due to instability can correspond to a pulse length error of a factor of two or more.

While SPIDER has been extensively tested and compared with FROG for well-known stable lasers, the above issue necessarily presents a difficult dilemma for researchers who use SPIDER on a new laser: because there is no feedback regarding the measurement quality and SPIDER measures only the coherent artifact, it is prone to yielding shorter pulses in the presence of otherwise undetected instability. It is true that standard ultrafast lasers are more stable than ever, but many are not, especially extremely short pulses at the edge of technology, for which SPIDER is often used, so users of SPIDER should be very cautious that measurements using it do not underestimate their pulse length. Also, it should be pointed out that the original motivation for SPIDER—that the FROG pulse-retrieval algorithm was slow—has entirely disintegrated due to the many-orders-of-magnitude increase in computing speed that has occurred over the past three decades.

On the other hand, SPIDER is very useful when the average spectral phase is the actual desired quantity. This is the case, for example, when tuning the dispersion of a pulse compressor.

VII. SPATIOTEMPORAL PULSE MEASUREMENT

The propagation of an ultrashort pulse is fundamentally *spatiotemporal*, meaning that its electric field cannot usually be separated into a product of purely temporal and spatial fields due to unavoidable spatiotemporal couplings.[125] Even the simple cases of diffraction of ultrashort pulses from a circular aperture[126] or the mere focusing of them[104,107,127–132] results in complicated spatiotemporal structures. As a result, measuring the pulse vs time, averaging over spatial coordinates (x and y), or measuring the pulse vs x and y, averaging over time, are not always useful. Complete-spatiotemporal-resolution measurements, on the other hand, enable direct visualization of these phenomena, so that they can be studied and understood in greater detail.

Complete spatiotemporal pulse-measurement techniques determine $\tilde{E}(x,y,z,\omega)$ or, equivalently, $E(x,y,z,t)$. Unfortunately, no method exists that, by itself, measures all this information.

Fortunately, it is not necessary to measure the z-dependence because it can be obtained from $\tilde{E}(x,y,\omega)$ using a diffraction integral. Unfortunately, we will see that, even this quantity cannot usually be measured in its entirety, and some information will be missing from it, in which case this additional information must be obtained.

Most multi-shot spatiotemporal-measurement methods are based on spectral interferometry (SI).[133,134]

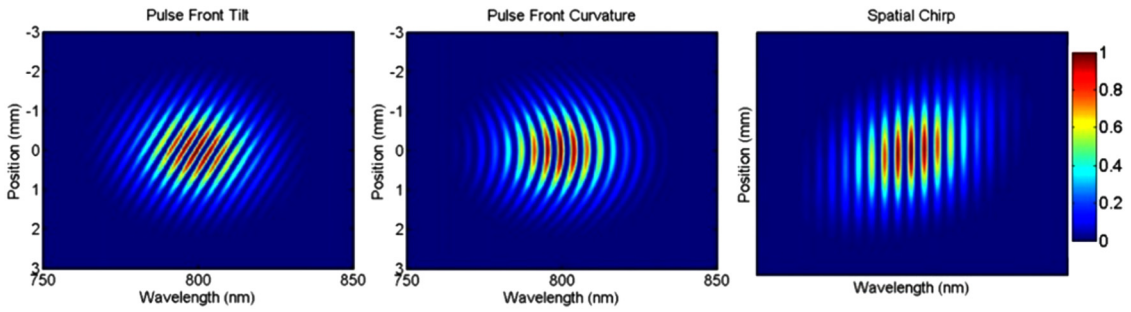

FIG. 52. Calculated spatially resolved spectral interferograms of pulses with three different spatiotemporal couplings. Left: pulse front tilt. Center: pulse front curvature. Right: spatial chirp. It is assumed that the transform-limited reference pulse is free of spatiotemporal couplings. Reproduced with permission from Akturk *et al.*, J. Opt. 12, 093001 (2010). Copyright 2010, IOP Publishing.

A. Spatially resolved spectral interferometry: One spatial dimension

It is straightforward to simply perform standard SI, using the other dimension of the camera to obtain spatial information of the beam in one spatial direction (see Fig. 52). Such an approach has been used to measure beams with spatiotemporal couplings, such as the pulse front distortions caused by lenses.[135,136] The drawback is that a pulse cannot be measured directly at a focus, where the beam is too small, so the beam had to be first recollimated, by propagating a second time through the lens. The advantage is that $\tilde{E}(y,\omega)$ can be retrieved from a single camera frame.

Spatially resolved spectral interferometry was also recently adapted for spatiotemporal measurements of a pulse with a huge pulse front tilt of ~89.9°. This extreme pulse front tilt was obtained using an etalon as the spectrometer's dispersive element, rather than a diffraction grating.[137] Analogously, spectrally resolved spatial interferometry can provide similar information.[138–140]

B. Fiber-based scanning spatiotemporal pulse measurement: Two and three spatial dimensions

SEA TADPOLE and STARFISH both naturally extend to spatiotemporal measurement in gathering spatial information by scanning a small fiber tip or probe through the light beam of interest. Using such methods, SI has been integrated into a scanning microscope to collect spatiotemporal information from a small sample.[109] Such methods have the same constraints as SI; for example, they require a well-characterized reference pulse that contains the spectrum of the unknown pulse, which now must also be spatially uniform. A significant benefit of using an optical fiber for introducing the unknown beam into a measurement device is that the fiber provides spatial resolution—as small as a few nanometers. The fiber collects only the light from a small spatial sample. The tip of the fiber can then be freely scanned through three-dimensional space, measuring $\tilde{E}(\omega)$ at each fiber position, yielding the spatiotemporal electric field of the unknown pulse $\tilde{E}_{unk}(x,y,z,\omega)$ with high spatial, temporal, and spectral resolutions.

In scanning SEA TADPOLE (see Fig. 53), a single-mode fiber in the unknown arm is mounted on an *x,y,z* translation stage, so that the spectrum and spectral phase of the unknown pulse are determined at multiple fiber positions in space.[141–143] 3D-scanning SEA TADPOLE was adapted to a measurement from a supercontinuum laser source (with sub-3 fs temporal resolution) for the generation and measurement of Airy-Bessel light bullets.[144]

Because of the high sensitivity of SEA TADPOLE, even very large beams can be measured, in which only a tiny fraction of the

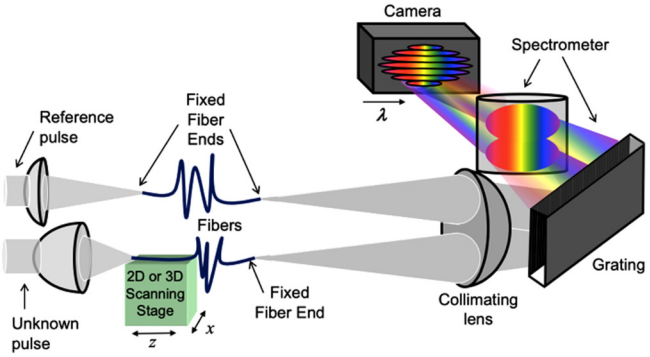

FIG. 53. Scanning SEA TADPOLE experimental setup. The laser beam is split into the reference and unknown arms in a Mach–Zehnder interferometer. The spatially uniform reference pulse goes through a delay line and is focused into one of the single-mode fibers (for nearly unlimited bandwidth the so-called endlessly single-mode photonic crystal fibers can be used). $E(t)$ of the reference pulse can be characterized by sending it with a flip mirror to a FROG device. In the unknown arm of the interferometer, the light is directed to the optical system being studied (the "experiment") and the resulting field is sampled at every *x*, *y*, and *z* point of interest by the other single-mode fiber. The outputs of the fibers are placed at the entrance of an imaging spectrometer with a few mm separation between them, so after the collimating lens, they cross at a small angle at the camera, which records the 2D spectrally resolved spatial interference pattern. The interference is measured for each fiber position and $\tilde{E}(\omega)$ is reconstructed from each interferogram yielding $\tilde{E}(x, y, z, \omega)$. Reproduced with permission from Bowlan *et al.*, Opt. Express **15**, 10219 (2007). Copyright 2007 The Optical Society.

pulse energy is coupled in the small fiber tip. Also, a SEA TADPOLE measurement at a single fiber position is proportional to the product of the reference and unknown fields, so the much more intense reference pulse homodynes and effectively amplifies the unknown pulse, making such a measurement even more sensitive.

The spatial resolution of scanning SEA TADPOLE can be improved to <1 μm using a near field scanning optical microscopy (NSOM) fiber probe, rather than a single-mode fiber.[145] NSOM fibers have been used in the past to measure the spatial intensity distribution of tightly focused continuous-wave lasers. Using an NSOM probe with an aperture diameter of 500 nm, the complete electric field of focused pulses with numerical apertures (NAs) as high as 0.44 and tiny features in their intensity of <1 μm have been measured with SEA TADPOLE.[104] NSOM probes have also been used in SEA TADPOLE to measure the refractive index and group velocity in a waveguide structure by collecting the local evanescent field into the fiber probe.[146] In addition, NSOM probes have been combined with Fourier transform spectral interferometry to characterize the local plasmons excited by ultrashort laser pulses in gold nanostructures.

Unfortunately, scanning in SEA TADPOLE cannot be performed stably enough to avoid a loss of the absolute *spatial* phase $\varphi_0(x,y,\omega_0)$. In some cases, for example, to measure spatiotemporal couplings or to see the pulse front, the spatial phase is not of interest. However, there are many benefits to having this additional spatial information. If it is known, the field $\tilde{E}(x,y,\omega)$ only needs to be measured at one value of z, and then it can be numerically propagated to any other plane using a diffraction integral.

Fortunately, it has been demonstrated that this information is still actually contained in the measured SEA TADPOLE data as long as $\tilde{E}(x,y,\omega)$ is measured for at least two values of z.[147] Using a simple *phase-diversity* approach, which involves Fresnel-transforming back and forth and replacing the intensity or spectrum by the known intensity or spectrum (often called a Gerchberg–Saxton-like phase-retrieval algorithm), the absolute spatial phase of a pulse can be very accurately and reliably numerically recovered. In addition, the complete spatiotemporal field at all locations is more accurately determined by this process.

C. Spatiotemporal measurement examples

The first motivation for scanning SEA TADPOLE was to measure the spatiotemporal field of focused pulses propagating using various lenses, including simple plano-convex lenses, aspheric lenses made of molded PMMA, achromatic doublets, and microscope objectives and lenses with numerical apertures as high as 0.44.[104,134] Figure 54 shows one such measurement exhibiting the so-called "fore-runner pulse" resulting from the combination of diffraction at the edge of the lens and chromatic aberration.[130–132] In this measurement, a 500 nm NSOM probe was used to resolve the sub-micrometer spatial features.

Another application of SEA TADPOLE has been to measure superluminal Bessel-X pulses, first demonstrated in Ref. 148. SEA TADPOLE yields highly accurate measurements of the group velocity. Although the superluminal speed of Bessel-X pulses had been measured before,[149] the full spatiotemporal field of a Bessel-X pulse had never been directly recorded with simultaneous high spatial and temporal resolutions.[141] Figure 55 shows the spatiotemporal field $E(x,z,t)$ from these measurements, side-by-side with the corresponding simulations. These results can be considered "snapshots in flight" of the Bessel-X pulse.

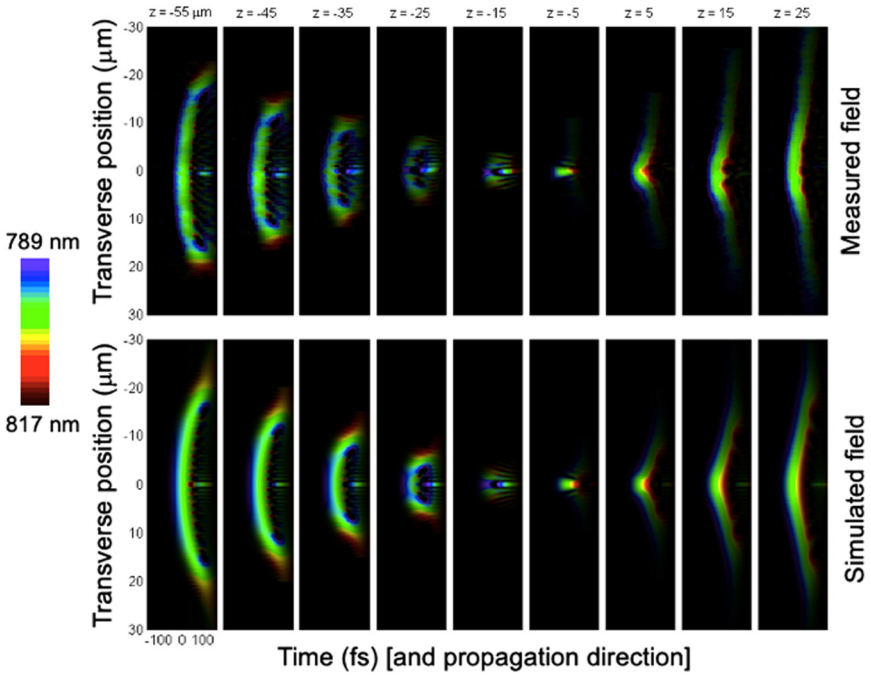

FIG. 54. Measured $E(x,z,t)$ for a pulse focused with a 0.44 NA aspheric lens. Each box shows $E(x,t)$ at a certain distance from the focus (z) written above the box. The color in the plot is the instantaneous frequency of the pulse as indicated by the color bar. Here the color also varies along the x-direction due to the severe chromatic aberrations that are present. The combination of overfilling the lens and chromatic aberrations result in the additional "fore-runner" pulse ahead of the main pulse, for $z < 0$. Reproduced with permission from Bowlan *et al.*, Opt. Express **16**, 13663 (2008). Copyright 2008 The Optical Society.

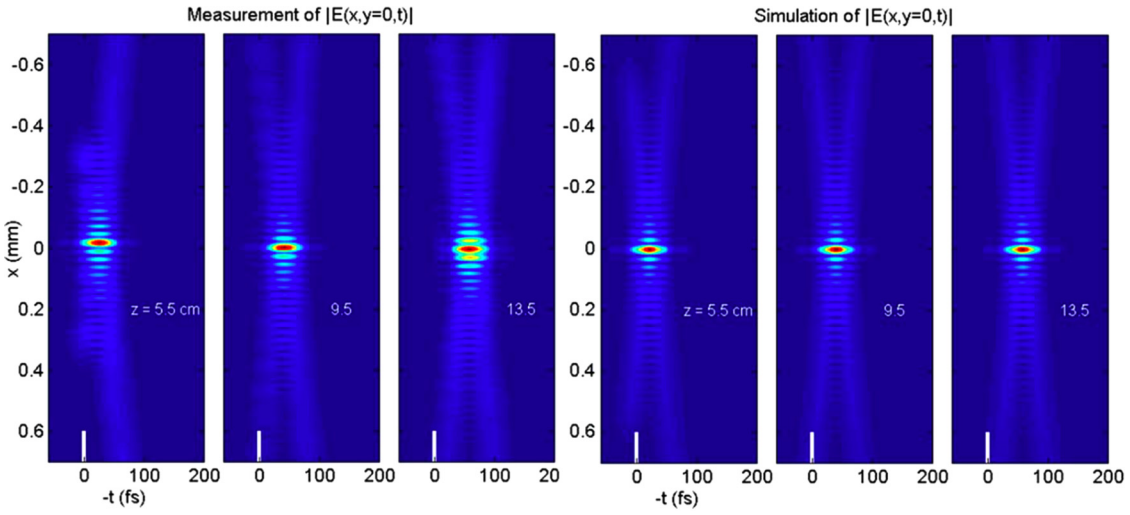

FIG. 55. Left: the measured spatiotemporal amplitude of the Bessel-X pulse at three different distances after the axicon. Right: the corresponding simulations. Amplitude is indicated by the color, and is normalized for each field to have a maximum of 1. The white bar on the time axis emphasizes the zero of time—a reference frame moving along with the reference pulse at the speed of light. Reproduced with permission from Bowlan *et al.*, Opt. Lett. **34**, 2276 (2009). Copyright 2009 The Optical Society.

Note that the central maximum of the pulse has a width of only ~20 μm and does not significantly diverge over 8 cm of propagation, a well-known and important characteristic of Bessel pulses. In comparison, a Gaussian beam of the same waist size would have expanded by 26 times. The superluminal velocity of the pulse is apparent in these plots because the reference pulse, determining the center of the time-windows, propagates at the speed of light (and is indicated by the small white bars on the snapshots and labeled by $t = 0$). The measured group velocity of this Bessel-X pulse is within 0.001% of the theoretically predicted value, 1.000 12c. Similarly, SEA TADPOLE has been used to study accelerating and decelerating Bessel pulses[150] and subluminal pulsed Bessel beams generated by diffractive axicons.[142,144]

Another interesting measurement made using SEA TADPOLE is the formation and evolution of Arago's spot in the shadow region behind an opaque circular disk that leads to a better understanding of diffraction phenomena (see Fig. 56).[151] The SEA TADPOLE technique has also been used for the generation and characterization of ultrabroadband Airy pulses that exhibit nondispersive behavior.[152,153]

D. Other spatiotemporal-measurement methods

Other methods have been developed for measuring the complete spatiotemporal intensity and phase on a multi-shot basis. One such technique, also with nanometer-scale spatial resolution, is called nanoFROG.[154] It involves making a FROG trace for the beam at a point in space using a nanometer-sized nonlinear crystal placed at that point. This technique yields the intensity and phase vs time for that point, with the usual trivial FROG ambiguities. While they comprise a larger set of ambiguities than the spatial phase ambiguity that arises in SEA TADPOLE, they should also be removable using the phase-diversity approach described earlier.[147]

E. Spatiotemporal measurement on a single shot: STRIPED FISH

Although the aforementioned complete spatiotemporal pulse-measurement techniques work well, and some even have sub-micrometer spatial resolution, they are all inherently multi-shot. But many applications—in particular, high-intensity extremely-low-rep-rate pulse measurements—require single-shot operation.

In attempting to take complete spatiotemporal measurement to three or even four dimensions, an obvious observation must be made. There are three spatial dimensions and only one temporal dimension, so it makes more sense to adapt a *spatial*-measurement technique to spatiotemporal measurement than to adapt a *temporal*-measurement technique. Arguably, the most common and powerful spatial intensity-and-phase technique is holography. Thus, it is reasonable to expect that spectrally resolved digital holography could be used to measure the complete spatiotemporal intensity and phase of ultrashort pulses, indeed, on a single shot.

The difference between the spatiotemporal measurement problem and standard holography is that holograms are usually generated using *monochromatic* beams. Consequently, they have no temporal dependence and so need only yield the spatial intensity and phase. Pulses, on the other hand, are *broadband*. Also, the relative phases of the various frequencies of the pulse must be measured in order to obtain the correct temporal dependence.

So the obvious approach is to record the three-dimensional unknown field information, $\widetilde{E_{unk}}(x,y,\omega)$, using multiple two-dimensional digital holograms, each at a different frequency. By using a digital camera with high pixel count and small pixel size, it is possible to illuminate different regions of the imaging sensor with digital holograms, each recorded at a different frequency.

In order to determine the relative phases of the many different-frequency holograms, the reference pulse intensity and

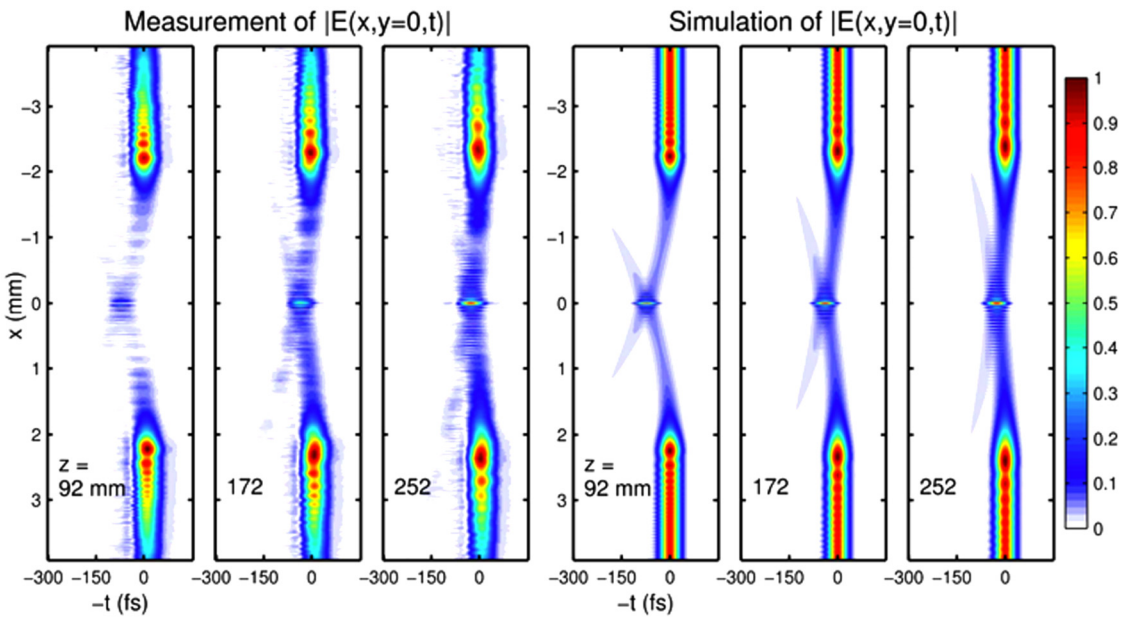

FIG. 56. Formation and evolution of the Arago spot behind an opaque disk 4 mm in diameter. The magnitude of the electric field E is shown at three different propagation distances z in pseudo-color code according to the color bar (white has been taken for the zero of the scale in order to better reveal areas of weak field). Reproduced with permission from Saari *et al.*, Opt. Express **18**, 11083 (2010). Copyright 2010 The Optical Society.

phase must be measured in time or frequency.[155] The reference pulse must also have relatively simple, known spatial properties and so must be spatially filtered to remove any possible spatiotemporal couplings.

Finally, as usual, the device should be simple and usable by those not skilled in pulse measurement. A device with these characteristics is STRIPED FISH (Spatially and Temporally Resolved Intensity and Phase Evaluation Device: Full Information from a Single Hologram).[155,156] At the time of this writing, STRIPED FISH is the only technique ever proposed to measure the complete spatiotemporal intensity and phase on a single shot.

A schematic of the STRIPED FISH apparatus is shown in Fig. 57. It comprises a very simple setup of only a coarse two-dimensional diffractive optical element (DOE), an interference bandpass filter (IBPF), imaging optics, and a camera. STRIPED FISH uses a previously spatially smoothed and temporally characterized reference pulse, accomplished at an earlier point using a spatial filter and a FROG measurement of the spatially filtered pulse. These operations can be performed on a replica of the unknown pulse, making the technique self-referenced. The pulse to be measured and the known reference pulse cross at a small vertical angle on the DOE, which simultaneously generates multiple pairs of beams at divergent different angles.

The DOE is also rotated slightly, so the horizontal propagation angle is different for each beam pair. Because the IBPF's transmission wavelength varies with horizontal incidence angle, it wavelength-filters each pair of beams to be quasi-monochromatic and with different center wavelengths. The beam pairs then overlap at the camera, generating an array of quasi-monochromatic holograms, each at a different wavelength. The spatiotemporal information of the unknown pulse is contained within multiple holograms, which are recorded simultaneously on the camera frame.

In the recorded STRIPED FISH trace, for each hologram on the camera, a Fourier filtering algorithm (similar to that of SEA

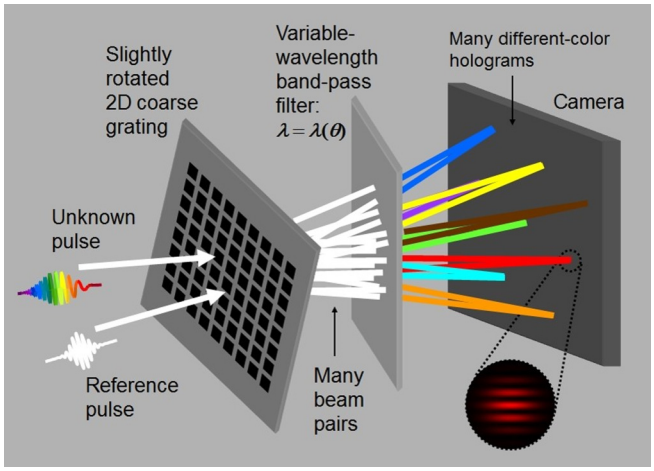

FIG. 57. Conceptual schematic of STRIPED FISH. Reproduced with permission from www.frog.gatech.edu, and Guang *et al.*, J. Opt. Soc. Am. B **31**, 2736 (2014). Copyright 2014 The Optical Society.

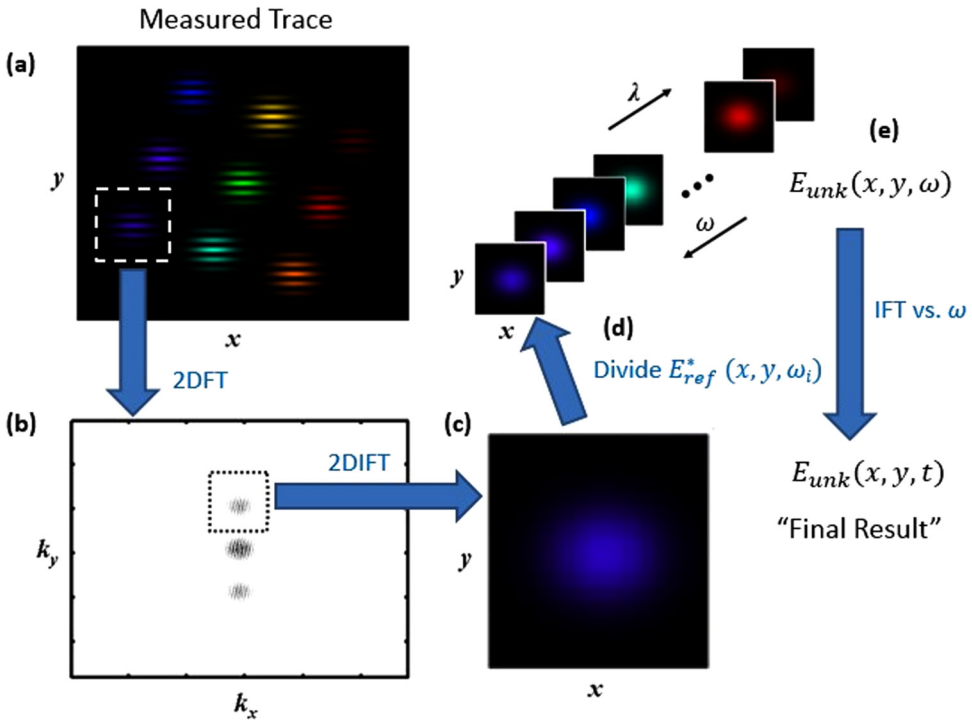

FIG. 58. Illustration of the STRIPED FISH retrieval algorithm. Amplitudes are plotted for complex quantities. (a) Conceptual schematic of multiple measured holograms of different frequencies are recorded on the camera. Axes are simply the two spatial coordinates of the camera frame. (b) A hologram of a certain frequency ω_i is selected. A two-dimensional Fourier transform (2DFT) is taken over spatial dimensions x and y. (c) The oscillating interference term is extracted and inversely processed by a 2DFT into the spatial domain, obtaining a product term $\tilde{E}_{unk}(x,y,\omega_i)\,\tilde{E}_{ref}^*(x,y,\omega_i)$. (d) Dividing the product term by the conjugated reference field $\tilde{E}_{ref}^*(x,y,\omega_i)$ yields the unknown spatial field at frequency ω_i, $\tilde{E}_{unk}(x,y,\omega_i)$. (e) Performing steps (b) through (d) for every hologram yields $\tilde{E}_{unk}(x,y,\omega)$; then $E_{unk}(x,y,t)$ is obtained by an inverse Fourier transform (IFT) into the time domain. Reproduced with permission from Gabolde and Trebino, J. Opt. Soc. Am. B **25**, A25 (2008). Copyright 2008 The Optical Society.

TADPOLE) is applied to obtain the unknown field $\tilde{E}_{unk}(x,y,\omega)$ at that frequency. Once the unknown fields at all frequencies are obtained, an inverse Fourier transform vs frequency is performed to convert the spatio-spectral field $\tilde{E}_{unk}(x,y,\omega)$ to the spatiotemporal pulse profile $E_{unk}(x,y,t)$. The pulse-retrieval algorithm is shown in Fig. 58.

Also, once $\tilde{E}_{unk}(x,y,\omega)$ is obtained, diffraction integrals can be used to propagate the field into different z-locations as well.

The STRIPED FISH trace, comprising multiple spectral digital holograms, is quite informative. After spatial filtering, the reference pulse contains no spatiotemporal coupling, and the spatio-spectral information of the unknown pulse is encoded in the STRIPED FISH trace.[157] Specifically, for a certain frequency, the unknown pulse's spatial structure is contained within each hologram: the spatial amplitude is represented by the intensity distribution and the spatial phase by the fringe shape across that hologram. Likewise, for each location, the unknown pulse's spectral information is reflected by multiple holograms: the spectral amplitude is represented by the intensity variations and the spectral phase is indicated by the fringe shifts among different holograms.

To illustrate these effects, simulated STRIPED FISH traces are shown in Fig. 59 for the cases of temporal double pulses and spatial double pulses. From the traces [Figs. 59(a) and 59(b)], we can clearly see the spectral intensity variations among different holograms. Also, from their fringe-shifting (e.g., the upper left hologram shows a shifted fringe pattern compared to the lower right one), we know that the spectral phase of the unknown pulse is varying with respect to different holograms or different frequencies.

Similarly, the spatial effects are demonstrated by considering a pair of spatial double pulses. Two spatial pulses are assumed to be propagating in the same direction, the left of which is half the amplitude (therefore quarter intensity) as the right one. To show the spatial phase variation, a π phase jump was introduced between the two component pulses. As shown in Figs. 59(c) and 59(d), the left pulse appears dimmer than the right one, indicating different spatial amplitudes. Also, in the middle of each hologram, we observe a fringe discontinuity due to the introduced spatial phase jump.

Thus, it is possible to quickly identify various significant spatial–spectral (or spatiotemporal) structures of the unknown pulse even before complete retrieval, by observing the recorded STRIPED FISH trace on camera.

STRIPED FISH has been used to measure spatiotemporally complicated ultrashort pulses. A measurement of simple spatially chirped pulse is shown in Fig. 60, where the measured pulse intensity

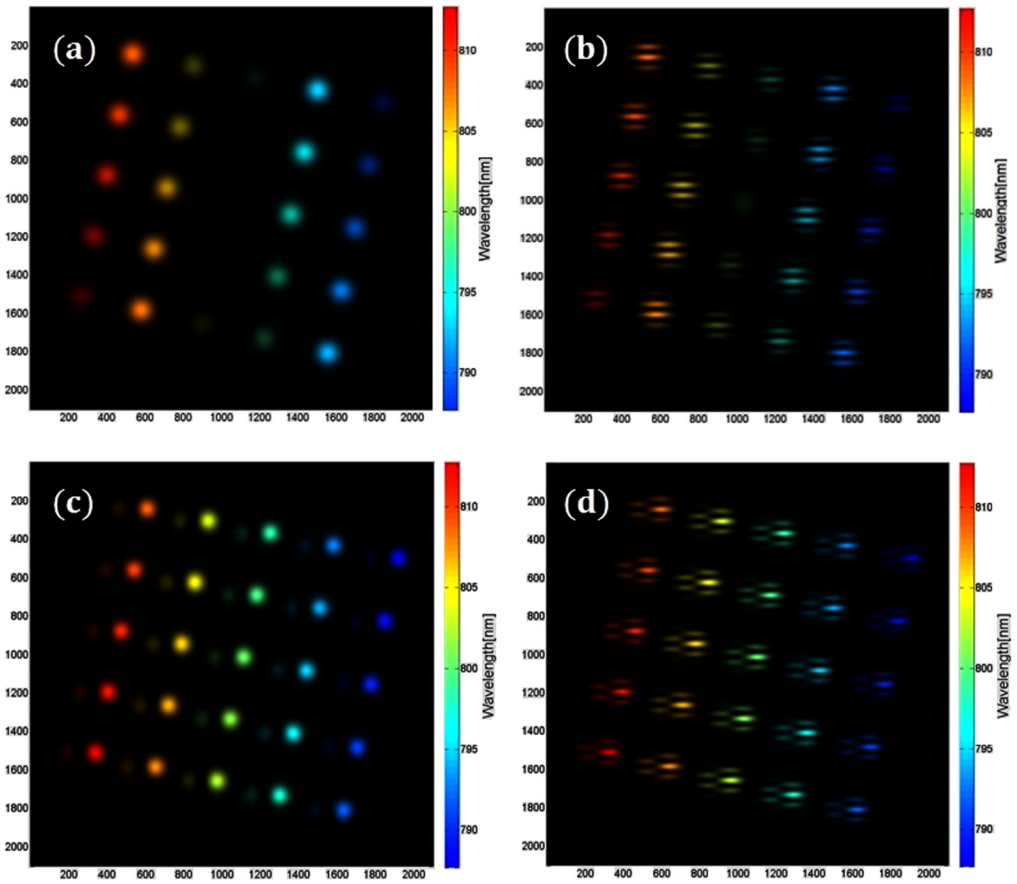

FIG. 59. STRIPED FISH traces for double pulses. (a) Pattern of the temporal double pulses, with equal intensities and π phase jump between the two component pulses in the absence of the reference pulse. (b) The STRIPED FISH holograms of the temporal double pulses. (c) Pattern of the spatial double pulses, with the left pulse of one fourth intensity as the right pulse in the absence of the reference pulse. A π phase jump was introduced between the two pulses. (d) The STRIPED FISH holograms of the spatial double pulses. Reproduced with permission from Guang *et al.*, Appl. Opt. **54**, 6640 (2015). Copyright 2015 The Optical Society.

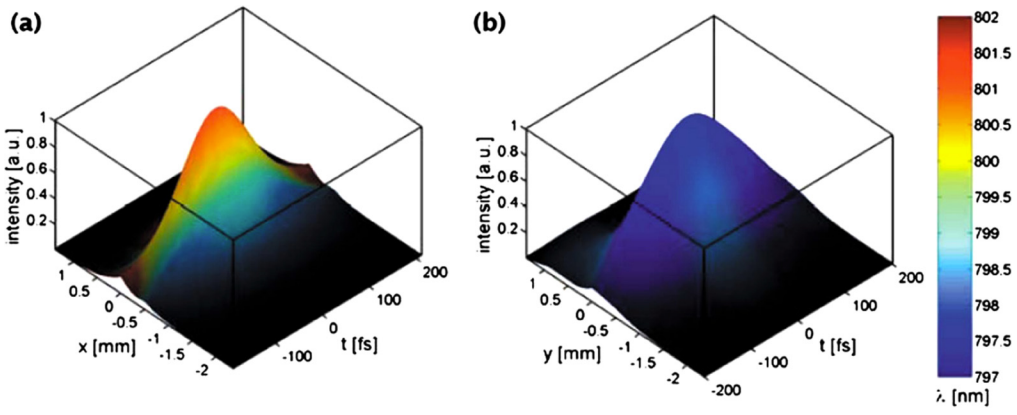

FIG. 60. STRIPED FISH measurement of a spatially chirped pulse. (a) Measured intensity along *x* and *t*. (b) Measured intensity along *y* and *t*. Color is the instantaneous frequency. Reproduced with permission from Gabolde and Trebino, J. Opt. Soc. Am. B **25**, A25 (2008). Copyright 2008 The Optical Society.

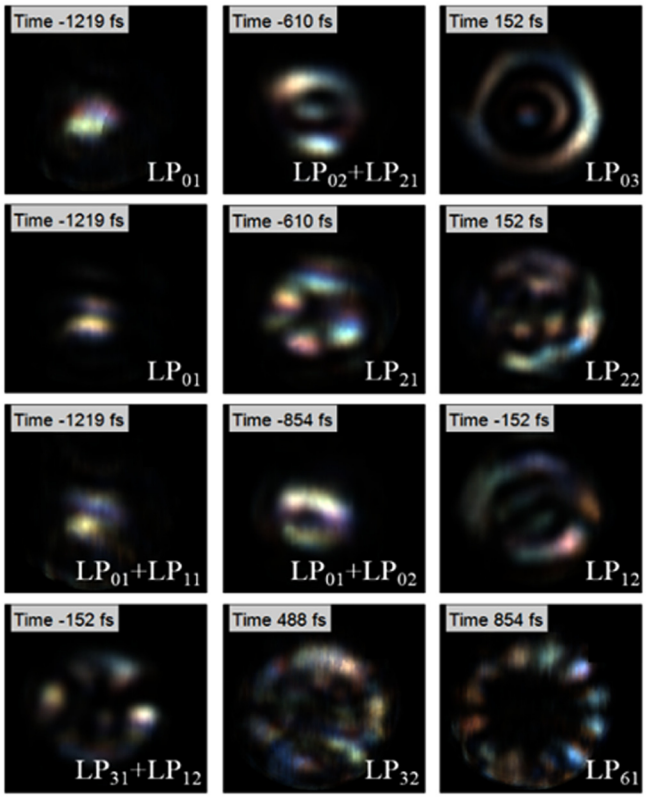

FIG. 61. Measured delay-scanned STRIPED FISH results of output pulses from multi-mode optical fibers with different fiber coupling situations. Each row contains three snapshots of the output pulses measured by delay-scanned STRIPED FISH for a different coupling arrangement. Top row: Centered coupling situation. Second row: Small offsets in both *x*- and *y*-direction situation. Third row: Small offset in only the *x*-direction. Fourth row: Large offset. All the snapshots show the spatial intensities (by brightness) and frequencies present (by color) at different times. Every snapshot has one or two dominant LP modes, which are labeled in the snapshots. Reproduced with permission from Zhu *et al.*, Opt. Express 25, 24015 (2017). Copyright 2017 The Optical Society.

is plotted vs *x*-*t* and *y*-*t*. The height and brightness denote the field intensity, and the color reflects the instantaneous wavelength. It is clear that the pulse shows a spatial frequency chirp along the *x*-direction whereas it has essentially no chirp along the *y*-direction.

A typical STRIPED FISH setup can generate ∼30 holograms at different frequencies. Each hologram typically is a ∼300 × ∼300 array of spatial pixels. As a result, the maximum measurable space-time-bandwidth product is on the order of 1 000 000.

If pulses more complicated in time must be measured, more holograms, but smaller in size, could be generated. However, given by the practical bandwidth of the bandpass filter, STRIPED FISH is limited in its temporal range. The longest measurable pulse is the reciprocal of this spectral width, so pulses as long as ∼10 ps can be measured using the narrowest-band available filters. This temporal range is usually more than adequate for measuring most ultrahigh-intensity pulses.

Alternatively, the temporal-range limitation can be overcome if a stable pulse train is to be measured, and scanning in one dimension (the delay) can be performed. In this case, the longest possible pulse length is the delay range, which can be ∼ns in length. One example of such cases is spatiotemporally complicated ultrashort pulses from multi-mode optical fibers. The pulses are stretched and modulated spatiotemporally by intermodal delay, modal dispersion, and material dispersion in the fibers. So, delay-scanned STRIPED FISH would be suitable for such measurement. Snapshots of movies of pulses with different fiber coupling situations from multi-mode optical fibers measured using this approach are shown in Fig. 61.

Such complicated pulses are best displayed using movies based on spectrograms due to the massive amount of information measured about the pulse. This plotting scheme avoids artifacts in plotting instantaneous frequency, e.g., when the pulse contains all of its spectrum, the pulse appears green (instantaneous frequency) instead of white (what it should be).[158–161]

VIII. CONCLUSIONS

Despite some confusion and missteps, the field of ultrashort-pulse measurement has progressed spectacularly in the past three decades. It has evolved from a position well behind that of pulse generation to now leading it in most areas. In 1990, only blurry black and white images were measurable and then only for simple pulses. Worse, such pulse measurements contained misleading artifacts. Today, it is possible to measure high-definition full-color spatiotemporal images of most pulses, including those as complicated in time and in space-time as have ever been intentionally generated, and pulse-shape (in)stability can also be determined easily and with confidence. Finally, pulse retrieval is now highly reliable, even for extremely complicated pulses and in the presence of significant noise. What remains is to take advantage of these powerful methods to better understand the pulses that arise in laboratory settings and the processes that they can be used to study.

ACKNOWLEDGMENTS

The Georgia Tech authors would like to acknowledge that this work was supported in part by the National Science Foundation (NSF) under Grant No. ECCS-1609808. We thank one of the anonymous reviewers for a number of highly helpful remarks that led to significant clarifications throughout the manuscript and, in particular, of Fig. 14. It should further be noted that Rick Trebino owns a small company that sells pulse-measurement devices.

DATA AVAILABILITY

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

REFERENCES

[1]L. Cohen, *Time-Frequency Analysis* (Prentiss-Hall, Englewood Cliffs, NJ, 1995).
[2]R. Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic Publishers, Boston, 2002).
[3]D. J. Kane and R. Trebino, paper presented at the Optical Society of America Annual Meeting, San Jose, CA, 1991.

[4]H. Stark, *Image Recovery: Theory and Application* (Academic Press, Orlando, FL, 1987).
[5]E. J. Akutowicz, Trans. Am. Math. Soc. **83**, 179–192 (1956).
[6]E. J. Akutowicz, Proc. Am. Math. Soc. **84**, 234–238 (1957).
[7]K. L. Sala, G. A. Kenney-Wallace, and G. E. Hall, IEEE J. Quantum Electron. **16**(9), 990–996 (1980).
[8]D. M. Rayner, P. A. Hackett, and C. Willis, Rev. Sci. Instrum. **53**(4), 537–538 (1982).
[9]E. S. Kintzer and C. Rempel, Appl. Phys. B **42**, 91–95 (1987).
[10]O. L. Bourne and A. J. Alcock, Rev. Sci. Instrum. **57**(12), 2979–2982 (1986).
[11]J. I. Dadap, G. B. Focht, D. H. Reitze, and M. C. Downer, Opt. Lett. **16**(7), 499–501 (1991).
[12]G. J. Dixon, Laser Focus World **33**(9), 99–102, 104–105 (1997), available at https://pascal-francis.inist.fr/vibad/index.php?action=getRecordDetail&idt=2143406.
[13]M. H. R. Hutchinson, I. A. McIntyre, G. N. Gibson, and C. K. Rhodes, Opt. Lett. **12**(2), 102–104 (1987).
[14]A. Tünnermann, H. Eichmann, R. Henking, K. Mossavi, and B. Wellegehausen, Opt. Lett. **16**(6), 402–404 (1991).
[15]R. Wyatt and E. E. Marinero, Appl. Phys. **25**, 297–301 (1981).
[16]J. A. Giordmaine, P. M. Rentzepis, S. L. Shapiro, and K. W. Wecht, Appl. Phys. Lett. **11**(7), 216–218 (1967).
[17]R. Trebino, E. K. Gustafson, and A. E. Siegman, J. Opt. Soc. Am. B **3**, 1295 (1986).
[18]A. Birmontas, R. Kupris, A. Piskarskas, V. Smil'gyavichyus, and A. Stabinis, Sov. J. Quant. Electron. **12**(6), 792–794 (1982).
[19]R. A. Fisher and J. J. A. Fleck, Appl. Phys. Lett. **15**(9), 287–290 (1969).
[20]J.-H. Chung and A. M. Weiner, IEEE J. Sel. Top. Quant. Electron. **7**(4), 656–666 (2001).
[21]J. C. Diels, J. J. Fontaine, and F. Simoni, in *Proceedings of the International Conference on Lasers* (STS Press, McLean, VA, 1983), pp. 348–355.
[22]J. C. Diels, J. J. Fontaine, I. C. McMichael, and F. Simoni, Appl. Opt. **24**(9), 1270–1282 (1985).
[23]J. C. Diels, Proc. SPIE **533**, 63–70 (1985).
[24]C. Yan and J. C. Diels, J. Opt. Soc. Am. B **8**(6), 1259–1263 (1991).
[25]J. C. M. Diels, J. J. Fontaine, N. Jamasbi, and M. Lai, paper presented at the Conference on Lasers & Electro-Optics, 1987.
[26]S. P. Le Blanc, G. Szabo, and R. Sauerbrey, Opt. Lett. **16**(19), 1508–1510 (1991).
[27]G. Szabó, Z. Bor, and A. Müller, Opt. Lett. **13**(9), 746–748 (1988).
[28]J. Etchepare, G. Grillon, and A. Orszag, IEEE J. Quantum Electron. **19**(5), 775–778 (1983).
[29]J. Janszky and G. Corradi, Opt. Commun. **60**(4), 251–256 (1986).
[30]N. Sarukura, M. Watanabe, A. Endoh, and S. Watanabe, Opt. Lett. **13**(11), 996–998 (1988).
[31]H. Schulz, H. Schuler, T. Engers, and D. von der Linde, IEEE J. Quantum Electron. **25**(12), 2580–2585 (1989).
[32]R. Fischer, J. Gauger, and J. Tilgner, AIP. Conf. Proc. **172**, 727–729 (1988).
[33]P. F. Curley, G. Darpentigny, G. Cheriaux, J. P. Chambaret, and A. Antonetti, Opt. Commun. **120**(1–2), 71–77 (1995).
[34]A. Watanabe, H. Saito, Y. Ishida, and T. Yajima, Opt. Commun. **63**(5), 320–324 (1987).
[35]A. Watanabe, S. Tanaka, H. Kobayashi, Y. Ishida, and T. Yajima, Rev. Sci. Instrum. **56**(12), 2259–2262 (1985).
[36]E. B. Treacy, J. Appl. Phys. **42**(10), 3848–3858 (1971).
[37]R. Trebino and D. J. Kane, J. Opt. Soc. Am. A **10**(5), 1101–1111 (1993).
[38]L. Cohen, Proc. IEEE **77**(7), 941–981 (1989).
[39]A. Freiberg and P. Saari, J. Quantum Electron. **19**(4), 622–630 (1983).
[40]R. A. Altes, J. Acoust. Soc. Am. **67**(4), 1232–1246 (1980).
[41]S. H. Nawab, T. F. Quatieri, and J. S. Lim, IEEE Trans. Acoust. Speech Signal Process. **31**(4), 986–998 (1983).
[42]R. Trebino and D. J. Kane, IEEE J. Quant. Electron. **29**(2), 571–579 (1999).
[43]R. Trebino, K. W. DeLong, D. N. Fittinghoff, J. N. Sweetser, M. A. Krumbügel, and D. J. Kane, Rev. Sci. Instrum. **68**(9), 3277–3295 (1997).
[44]D. J. Kane and R. Trebino, Opt. Lett. **18**(10), 823–825 (1993).
[45]D. J. Kane and R. Trebino, IEEE J. Quantum Electron. **29**(2), 571–579 (1993).
[46]K. W. DeLong and R. Trebino, J. Opt. Soc. Am. A **11**(9), 2429–2437 (1994).
[47]K. W. DeLong, R. Trebino, and D. J. Kane, J. Opt. Soc. Am. B **11**(9), 1595–1608 (1994).
[48]K. W. DeLong, R. Trebino, J. Hunter, and W. E. White, J. Opt. Soc. Am. B **11**(11), 2206–2215 (1994).
[49]K. W. DeLong, D. N. Fittinghoff, R. Trebino, B. Kohler, and K. Wilson, Opt. Lett. **19**(24), 2152–2154 (1994).
[50]K. W. DeLong, D. N. Fittinghoff, and R. Trebino, IEEE J. Quant. Electron. **32**(7), 1253–1264 (1996).
[51]G. Taft, A. Rundquist, M. M. Murnane, I. P. Christov, H. C. Kapteyn, K. W. DeLong, D. N. Fittinghoff, M. A. Krumbügel, J. N. Sweetser, and R. Trebino, IEEE J. Sel. Top. Quantum Electron. **2**(3), 575–585 (1996).
[52]T. Bendory, P. Sidorenko, and Y. C. Eldar, IEEE Signal Process. Lett. **24**(5), 722–726 (2017).
[53]K. W. DeLong, C. L. Ladera, R. Trebino, B. Kohler, and K. R. Wilson, Opt. Lett. **20**(5), 486–488 (1995).
[54]R. Jafari, T. Jones, and R. Trebino, Opt. Express **27**(3), 2112–2124 (2019).
[55]R. Jafari and R. Trebino, IEEE J. Quantum Electron. **55**(4), 1–7 (2019).
[56]R. Jafari and R. Trebino, IEEE J. Quantum Electron. **56**(1), 1–8 (2020).
[57]B. A. Richman, K. W. DeLong, and R. Trebino, "Temporal characterization of the Stanford mid-IR FEL by frequency-resolved optical gating," in *Conference on Lasers and Electro-Optics*, edited by T. Deutsch, J. Goldsmith, D. Killinger, and G. Valley (Optical Society of America, 1995), Vol. 15 of OSA Technical Digest, paper No. CWF21.
[58]B. Kohler, V. V. Yakovlev, K. R. Wilson, J. Squier, K. W. DeLong, and R. Trebino, Opt. Lett. **20**(5), 483–485 (1994).
[59]T. S. Clement, A. J. Taylor, and D. J. Kane, Opt. Lett. **20**(1), 70–72 (1995).
[60]D. J. Kane, A. J. Taylor, R. Trebino, and K. W. DeLong, Opt. Lett. **19**(14), 1061–1063 (1994).
[61]K. Michelmann, T. Feurer, R. Fernsler, and R. Sauerbrey, Appl. Phys. B **63**(5), 485–489 (1996).
[62]S. Backus, J. Peatross, Z. Zeek, A. Rundquist, G. Taft, M. M. Murnane, and H. C. Kapteyn, Opt. Lett. **21**, 665–667 (1996).
[63]I. Thomann, A. Bahabad, X. Liu, R. Trebino, M. M. Murnane, and H. C. Kapteyn, Opt. Express **17**(6), 4611–4633 (2009).
[64]F. Quéré, Y. Mairesse, and J. Itatani, J. Mod. Opt. **52**(2–3), 339–360 (2005).
[65]P. Bowlan and R. Trebino, Opt. Express **19**(2), 1367–1377 (2011).
[66]S.-D. Yang, A. M. Weiner, K. R. Parameswaran, and M. M. Fejer, Opt. Lett. **30**(16), 2164–2166 (2005).
[67]J. Zhang, A. P. Shreenath, M. Kimmel, E. Zeek, R. Trebino, and S. Link, Opt. Express **11**(6), 601–609 (2003).
[68]X. Gu, L. Xu, M. Kimmel, E. Zeek, P. O'Shea, A. P. Shreenath, R. Trebino, and R. S. Windeler, Opt. Lett. **27**(13), 1174–1176 (2002).
[69]M. M. Malley and P. M. Rentzepis, Chem. Phys. Lett. **7**, 57–60 (1970).
[70]A. Brun, P. Georges, G. LeSaux, and F. Salin, J. Phys. D **24**, 1225–1233 (1991).
[71]S. Szatmári, F. P. Schäfer, and J. Jethwa, Rev. Sci. Instrum. **61**(3), 998–1003 (1990).
[72]G. Stibenz and G. Steinmeyer, Opt. Express **13**(7), 2617–2626 (2005).
[73]J. Hyyti, E. Escoto, and G. Steinmeyer, J. Opt. Soc. Am. B **34**(11), 2367–2375 (2017).
[74]J. Hyyti, E. Escoto, and G. Steinmeyer, Rev. Sci. Instrum. **88**, 103102 (2017).
[75]S. Akturk, C. D'Amico, and A. Mysyrowicz, J. Opt. Soc. Am. B **25**(6), A63 (2008).
[76]L. Xu, E. Zeek, and R. Trebino, J. Opt. Soc. Am. B **25**(6), A70–A80 (2008).
[77]M. Rhodes, Z. Guang, and R. Trebino, Appl. Sci. **7**(40), 137601 (2017).
[78]M. Rhodes, M. Mukhopadhyay, J. Birge, and R. Trebino, J. Opt. Soc. Am. B **32**(9), 1881–1888 (2015).
[79]M. Rhodes, G. Steinmeyer, J. Ratner, and R. Trebino, Laser Photonics Rev. **7**(4), 557–565 (2013).
[80]M. Rhodes, G. Steinmeyer, and R. Trebino, Appl. Opt. **53**(16), D1–D11 (2014).

[81]J. Ratner, G. Steinmeyer, T. C. Wong, R. Bartels, and R. Trebino, Opt. Lett. **37**(14), 2874–2876 (2012).
[82]E. Escoto, R. Jafari, R. Trebino, and G. Steinmeyer, Opt. Lett. **44**(12), 3142–3145 (2019).
[83]S. Linden, H. Giessen, and J. Kuhl, Phys. Status Solidi B **206**(1), 119–124 (1998).
[84]D. Reid, P. Loza-Alvarez, C. Brown, T. Beddard, and W. Sibbett, Opt. Lett. **25**(19), 1478–1480 (2000).
[85]B. Richman, M. Krumbügel, and R. Trebino, Opt. Lett. **22**(10), 721–723 (1997).
[86]A. Lanin, A. Voronin, A. Fedotov, and A. Zheltikov, Sci. Rep. **4**, 6670 (2014).
[87]P. O'Shea, M. Kimmel, X. Gu, and R. Trebino, Opt. Lett. **26**(12), 932–934 (2001).
[88]C. Radzewicz, P. Wasylczyk, and J. S. Krasinski, Opt. Commun. **186**(4–6), 329–333 (2000).
[89]A. G. Akmanov, A. I. Kovrigin, and N. K. Podsotskaya, Radio Eng. Electron Phys. **14**, 1315 (1969).
[90]D. H. Auston, Opt. Commun. **3**, 272 (1971).
[91]S. Akturk, M. Kimmel, P. O'Shea, and R. Trebino, Opt. Express **11**(5), 491–501 (2003).
[92]S. Akturk, M. Kimmel, P. O'Shea, and R. Trebino, Opt. Express **11**(1), 68–78 (2003).
[93]J. Cohen, D. Lee, V. Chauhan, P. Vaughan, and R. Trebino, Opt. Express **18**(16), 17484–17497 (2010).
[94]S. Akturk, M. Kimmel, P. O'Shea, and R. Trebino, Opt. Lett. **29**(9), 1025–1027 (2004).
[95]M. Van Noort, L. Der Voort, and M. Löfdahl, Sol. Phys. **228**(1), 191–215 (2005).
[96]R. J. Steriti and M. A. Fiddy, Opt. Lett. **19**(8), 575–577 (1994).
[97]T. C. Wong and R. Trebino, J. Opt. Soc. Am. B **30**(11), 2781–2786 (2013).
[98]T. C. Wong, J. Ratner, V. Chauhan, J. Cohen, P. M. Vaughan, L. Xu, A. Consoli, and R. Trebino, J. Opt. Soc. Am. B **29**(6), 1237–1244 (2012).
[99]T. C. Wong, J. Ratner, and R. Trebino, J. Opt. Soc. Am. B **29**(8), 1889–1893 (2012).
[100]C. Froehly, A. Lacourt, and J. C. Viénot, Nouv. Rev. D'Opt. **4**, 183–196 (1973).
[101]D. N. Fittinghoff, J. L. Bowie, J. N. Sweetser, R. T. Jennings, M. A. Krumbügel, K. W. DeLong, R. Trebino, and I. A. Walmsley, Opt. Lett. **21**(12), 884–886 (1996).
[102]B. Alonso, M. Miranda, Í. J. Sola, and H. Crespo, Opt. Express **20**(16), 17880–17893 (2012).
[103]C. Dorrer, N. Belabas, J.-P. Likforman, and M. Joffre, J. Opt. Soc. Am. B **17**(10), 1795–1802 (2000).
[104]P. Bowlan, U. Fuchs, R. Trebino, and U. D. Zeitner, Opt. Express **16**(18), 13663–13675 (2008).
[105]P. Bowlan, P. Gabolde, M. A. Coughlan, R. Trebino, and R. J. Levis, J. Opt. Soc. Am. B **25**(6), A81–A92 (2008).
[106]P. Bowlan, P. Gabolde, A. Shreenath, K. McGresham, R. Trebino, and S. Akturk, Opt. Express **14**(24), 11892 (2006).
[107]P. Bowlan, P. Gabolde, and R. Trebino, Opt. Express **15**, 10219–10230 (2007).
[108]J. P. Geindre, P. Audebert, A. Rousse, F. Falliès, J. C. Gauthier, A. Mysyrowicz, A. Dos Santos, G. Hamoniaux, and A. Antonetti, Opt. Lett. **19**(23), 1997–1999 (1994).
[109]S. D. Gennaro, Y. Sonnefraud, N. Verellen, P. Van Dorpe, V. V. Moshchalkov, S. A. Maier, and R. F. Oulton, Nat. Commun. **5**, 3748 (2014).
[110]P. Bowlan and R. Trebino, U.S. patent 8,953,166 (10 February 2015).
[111]J. Cohen, P. Bowlan, V. Chauhan, and R. Trebino, Opt. Express **18**(7), 6583–6597 (2010).
[112]J. Cohen, P. Bowlan, V. Chauhan, P. Vaughan, and R. Trebino, Opt. Express **18**(24), 24451–24460 (2010).
[113]J. Cohen, P. Bowlan, and R. Trebino, IEEE J. Sel. Top. Quantum Electron. **18**(1), 218–227 (2012).
[114]J. Cohen, P. Bowlan, V. Chauhan, P. Vaughan, and R. Trebino, Opt. Commun. **284**, 3785–3794 (2011).
[115]C. Iaconis and I. A. Walmsley, Opt. Lett. **23**(10), 792–794 (1998).
[116]T. Oksenhendler, S. Coudreau, N. Forget, V. Crozatier, S. Grabielle, R. Herzog, O. Gobert, and D. Kaplan, Appl. Phys. B Lasers Opt. **99**(1), 7–12 (2010).
[117]M. Takeda, H. Ina, and S. Kobayashi, J. Opt. Soc. Am. **72**(1), 156 (1982).
[118]S. L. Hahn, *Hilbert Transforms in Signal Processing* (Artech House, Boston, 1996).
[119]J. Bethge and G. Steinmeyer, Rev. Sci. Instrum. **79**, 073102 (2008).
[120]G. Stibenz and G. Steinmeyer, Rev. Sci. Instrum. **77**, 073105 (2006).
[121]P. Baum, S. Lochbrunner, and E. Riedle, Opt. Lett. **29**(2), 210–212 (2004).
[122]T. Witting, D. R. Austin, and I. A. Walmsley, Opt. Lett. **34**(9), 881 (2009).
[123]J. R. Birge, R. Ell, and F. X. Kärtner, Opt. Lett. **31**(13), 2063–2065 (2006).
[124]E. M. Kosik, A. S. Radunsky, I. A. Walmsley, and C. Dorrer, Opt. Lett. **30**(3), 326 (2005).
[125]S. Akturk, X. Gu, P. Bowlan, and R. Trebino, J. Opt. **12**, 093001 (2010).
[126]Z. Bor and Z. L. Horváth, Opt. Commun. **94**(4), 249–258 (1992).
[127]Z. L. Horváth and Z. Bor, Opt. Commun. **100**(1–4), 6–12 (1993).
[128]Z. L. Horváth and Z. Bor, Opt. Commun. **108**(4–6), 333–342 (1994).
[129]Z. L. Horváth, K. Osvay, and Z. Bor, Opt. Commun. **111**(5–6), 478–482 (1994).
[130]M. Kempe and W. Rudolph, Opt. Lett. **18**, 137–139 (1993).
[131]M. Kempe and W. Rudolph, Phys. Rev. A **48**(6), 4721–4729 (1993).
[132]U. Fuchs, U. D. Zeitner, and A. Tünnermann, Opt. Express. **13**(10), 3852–3861 (2005).
[133]T. Tanabe, H. Tanabe, Y. Teramura, and F. Kannari, J. Opt. Soc. Am. B **19**(11), 2795–2802 (2002).
[134]Y. Teramura, M. Suekuni, and F. Kannari, J. Opt. A Pure Appl. Opt. **2**(1), 21–26 (2000).
[135]Z. Bor, J. Mod. Opt. **35**(12), 1907–1918 (1988).
[136]A. Federico and O. Martinez, Opt. Commun. **91**(1), 104–110 (1992).
[137]P. Bowlan and R. Trebino, J. Opt. Soc. Am. B **27**(11), 2322–2327 (2010).
[138]D. Meshulach, D. Yelin, and Y. Silberberg, J. Opt. Soc. Am. B **14**(8), 2095–2098 (1997).
[139]A. Börzsönyi, A. P. Kovács, M. Görbe, and K. Osvay, Opt. Commun. **281**(11), 3051–3061 (2008).
[140]D. J. McCabe, A. Tajalli, D. R. Austin, P. Bondareff, I. A. Walmsley, S. Gigan, and B. Chatel, Nat. Commun. **2**, 447 (2011).
[141]P. Bowlan, H. Valtna-Lukner, M. Lõhmus, P. Piksarv, P. Saari, and R. Trebino, Opt. Lett. **34**(15), 2276–2278 (2009).
[142]M. Lõhmus, P. Bowlan, R. Trebino, H. Valtna-Lukner, P. Piksarv, and P. Saari, Lith J. Phys. **50**(1), 69–74 (2010).
[143]P. Saari, P. Bowlan, H. Valtna-Lukner, M. Lõhmus, P. Piksarv, and R. Trebino, Laser Phys. **20**(5), 948–953 (2010).
[144]P. Piksarv, H. Valtna-Lukner, A. Valdmann, M. Lõhmus, R. Matt, and P. Saari, Opt. Express **20**(15), 17220–17229 (2012).
[145]E. Betzig, M. Isaacson, and A. Lewis, Appl. Phys. Lett. **51**(25), 2088–2090 (1987).
[146]J. Trägårdh and H. Gersen, Opt. Express **21**(14), 16629–16638 (2013).
[147]P. Bowlan and R. Trebino, J. Opt. Soc. Am. B **29**(2), 244–248 (2012).
[148]P. Saari and K. Reivelt, Phys. Rev. Lett. **79**(21), 4135 (1997).
[149]I. Alexeev, K. Kim, and H. Milchberg, Phys. Rev. Lett. **88**(7), 073901 (2002).
[150]H. Valtna-Lukner, P. Bowlan, M. Lõhmus, P. Piksarv, R. Trebino, and P. Saari, Opt. Express **17**(17), 14948–14955 (2009).
[151]P. Saari, P. Bowlan, H. Valtna-Lukner, M. Lõhmus, P. Piksarv, and R. Trebino, Opt. Express **18**(11), 11083 (2010).
[152]A. Valdmann, P. Piksarv, H. Valtna-Lukner, and P. Saari, Opt. Lett. **39**(7), 1877–1880 (2014).
[153]A. Valdmann, P. Piksarv, H. Valtna-Lukner, and P. Saari, J. Opt. **20**(9), 095605 (2018).

[154]J. Extermann, L. Bonacina, F. Courvoisier, D. Kiselev, Y. Mugnier, R. Le Dantec, C. Galez, and J.-P. Wolf, Opt. Express **16**(14), 10405–10411 (2008).
[155]P. Gabolde and R. Trebino, Opt. Express **14**(23), 11460 (2006).
[156]P. Gabolde and R. Trebino, J. Opt. Soc. Am. B **25**(6), A25–A33 (2008).
[157]Z. Guang, M. Rhodes, and R. Trebino, Appl. Opt. **54**(22), 6640–6651 (2015).
[158]P. Zhu, R. Jafari, T. Jones, and R. Trebino, Opt. Express **25**(20), 24015 (2017).
[159]Z. Guang, M. Rhodes, M. Davis, and R. Trebino, J. Opt. Soc. Am. B **31**(11), 2736–2743 (2014).
[160]Z. Guang, M. Rhodes, and R. Trebino, J. Opt. Soc. Am. B **33**(9), 1955–1962 (2016).
[161]M. Rhodes, Z. Guang, J. Pease, and R. Trebino, Appl. Opt. **56**(11), 3024–3034 (2017).
[162]P. Bowlan, "Measuring the spatiotemporal electric field of ultrashort pulses with high spatial and spectral resolution," Ph.D. thesis (Georgia Institute of Technology, 2009), available at https://smartech.gatech.edu/handle/1853/28188.
[163]J. Hyyti, E. Escoto, G. Steinmeyer, and T. Witting, Opt. Lett. **42**, 2185 (2017).