Practical Issues in Ultrashort-Laser-Pulse Measurement Using Frequency-Resolved Optical Gating

Kenneth W. DeLong, *Member, IEEE,* David N. Fittinghoff, and Rick Trebino

Abstract— **We explore several practical experimental issues in measuring ultrashort laser pulses using the technique of frequency-resolved optical gating (FROG). We present a simple method for checking the consistency of experimentally measured FROG data with the independently measured spectrum and autocorrelation of the pulse. This method is a powerful way of discovering systematic errors in FROG experiments. We show how to determine the optimum sampling rate for FROG and show that this satisfies the Nyquist criterion for the laser pulse. We explore the low- and high-power limits to FROG and determine that femtojoule operation should be possible, while the effects of self-phase modulation limit the highest signal efficiency in FROG to 1%. We also show quantitatively that the temporal blurring due to a finite-thickness medium in single-shot geometries does not strongly limit the FROG technique. We explore the limiting time-bandwidth values that can be represented on a FROG trace of a given size. Finally, we report on a new measure of the FROG error that improves convergence in the presence of noise.**

I. Introduction

THE USEFULNESS of a new scientific technique is determined by the details of its implementation. Many clever techniques that are otherwise rigorous and sound fail to become useful workhorse laboratory techniques because of difficulties in experimental implementation, noise sensitivity, cumbersomeness, or other more fundamental limitations. Often these details go unreported, resulting in a loss of valuable time and resources for research groups that attempt to use the technique. It is therefore incumbent upon the developers of a new experimental technique to determine whether practical limitations will render the technique less than ideal, and it is crucial that this information reach the scientific community. In this work, we shall examine one such recently developed technique for measuring the intensity and phase of an ultrashort laser pulse.

Recently, there have been major advances in the measurement of ultrashort laser pulses. Several techniques that give the complete time-dependent intensity and phase of an ultrashort laser pulse have been developed [1]–[5]. One of these techniques, frequency-resolved optical gating (FROG), is mathematically rigorous and can operate on a single-shot basis [6]. FROG involves a simple experimental setup coupled with the use of a phase-retrieval-based algorithm [7]–[9] to retrieve the intensity and phase of laboratory laser pulses. It has been demonstrated in experiments from the mid-infrared [10] to the ultraviolet [11], [12], and on pulse lengths from sub-10 fs to several picoseconds. FROG has been used to align and characterize chirped-pulse amplifier systems [13], to study high-intensity light-matter interactions, and to explore the nature of modelocking in ultrashort-pulse laser oscillators [14].

In our experience developing FROG, we have identified many of the practical issues involved in the everyday use of FROG. We have also worked with several different research groups that were implementing FROG in their laboratories, allowing us to add the benefit of their experience to our own. The questions discussed in this paper are ones that were commonly asked of us by research groups attempting to use FROG as a tool, as well as ones that we asked of ourselves. These include questions such as "How can we tell if our FROG data is valid?" "How densely should we sample the FROG trace?" "What limits the minimum and maximum intensities that can be used with FROG?" "What is the effect of the finite thickness of the nonlinear medium?" "Can a pulse of arbitrary complexity be represented on a FROG trace?" It is the purpose of this paper to answer practical questions such as these.

After reviewing the basic concepts of the FROG technique for the various experimental geometries, we present a simple and powerful technique to check the consistency of the FROG data with an independently measured spectrum and autocorrelation of the pulse. We show how these checks can uncover systematic errors in the measurement apparatus that might otherwise go undiscovered. We then discuss the data-sampling rate for FROG and show that FROG satisfies the Nyquist-sampling rate, ensuring that full information about the pulse is available. We also determine the upper and lower limits to the sampling rate in order for the FROG data to be properly sampled. We then explore some fundamental limits to the FROG technique. We explore the low-power limits to FROG and find that in its most sensitive incarnation FROG may break the 1 pJ barrier for 100-fs pulses. On the other end of the scale, we find that the upper limit to allowable FROG-signal efficiency is about 1% and is limited in $\chi^{(3)}$ geometries

Manuscript received June 26, 1995; revised March 14, 1996. This work was supported in part by the U.S. Department of Energy, Office of Basic Energy Sciences, Chemical Sciences Division, the Cooperative Research and Development Agreement program, and the Laboratory-Directed Research and Development program of Sandia National Laboratories.

K. W. DeLong was with Sandia National Laboratories, Livermore, CA 94551–0969 USA. He is now with Lawrence Livermore National Labs, Livermore, CA 94550 USA.

D. N. Fittinghoff and R. Trebino are with Sandia National Laboratories, Livermore, CA 94551-0969 USA.

Publisher Item Identifier S 0018-9197(96)05021-X.

by self- and cross-phase modulation-induced distortion. We then calculate the effect of the finite thickness of the medium used to generate the FROG signal and find that this is not a severe limitation to the measurement of pulses using FROG. We also compute the maximum time-bandwidth product for pulses that can be properly represented on a FROG trace of a given size. This maximum (rms) time-bandwidth product ranges anywhere from 3 to 20 for a FROG trace of 128×128 pixels and varies strongly with the form of the pulse itself and the size of the trace. Finally, we present a slightly modified form of the FROG error that allows for more robust pulse retrieval in the presence of noise.

II. Basics of FROG

The FROG technique consists of two parts: an experimental apparatus that generates a "FROG trace" of a laboratory laser pulse and a phase-retrieval-based algorithm that takes this FROG trace as an input and yields the time-dependent intensity and phase of the pulse that generated the trace.

Experimentally, FROG data are obtained by taking the spectrum of a signal field, generated by a nonlinear autocorrelation signal of two replicas of the pulse to be measured, for all relevant values of the time delay between the two replicas. To date, four FROG geometries have been demonstrated: polarization-gate (PG) [6], self-diffraction (SD) [12], [15], second-harmonic generation (SHG) [16], [17], and third-harmonic generation (THG) [18], although other geometries are possible [19]. In PG FROG, where the signal field is generated by Kerr-effect polarization rotation, as in polarization spectroscopy, the signal field has the form

$$E_{\mathrm{sig}}^{\mathrm{PG}}(t,\tau) = E(t)|E(t-\tau)|^2. \tag{1}$$

When self-diffraction is used, the signal field is

$$E_{\mathrm{sig}}^{\mathrm{SD}}(t,\tau) = E^2(t)E^*(t-\tau) \tag{2}$$

while in SHG FROG, the signal field is

$$E_{\mathrm{sig}}^{\mathrm{SHG}}(t,\tau) = E(t)E(t-\tau) \tag{3}$$

and for third-harmonic generation we have

$$E_{\mathrm{sig}}^{\mathrm{THG}}(t,\tau) = E^2(t)E(t-\tau). \tag{4}$$

This signal field is then frequency-resolved by a spectrometer to yield the FROG trace

$$I_{\mathrm{FROG}}(\omega,\tau) = \left| \int_{-\infty}^{\infty} dt E_{\mathrm{sig}}(t,\tau) \exp\left(i\omega t\right) \right|^2. \tag{5}$$

The FROG trace is a spectrogram of the pulse, a type of time-frequency distribution [20] that contains all relevant information about the pulse.

This experimentally measured FROG trace is then used as input to an algorithm that inverts the trace to yield the time-dependent intensity and phase of the pulse. The FROG algorithm is based on the iterative-Fourier-transform algorithm pioneered in the field of phase-retrieval [21]. As first reported, the FROG algorithm was useful only for PG FROG [7]. However, the application of powerful techniques such as the method of generalized projections has resulted in an algorithm that works well for PG, SD, SHG, and THG FROG [8], [9], [18].

III. Redundancy and Error Detection in FROG

We have used the FROG algorithm to retrieve thousands of pulses, both simulated and experimental, and we wish to stress that for valid data, the algorithm has never failed to converge. In our experience, the only times that the algorithm does not converge is for undersampled (aliased) pulses or for experimental data that is excessively corrupted by noise, incorrectly calibrated, truncated, distorted, or contains some other systematic or large random error. This is in fact a unique advantage of the FROG method: if the algorithm does not converge, one can be sure that an error exists in the data-collection apparatus, while if the algorithm converges well it nearly guarantees that the correct result has been found. The reason for this is that an electric field sampled at N points has $2N$ degrees of freedom (N points of both magnitude and phase), but corresponds to a FROG trace with N^2 pixels. Thus, there is great redundancy in the FROG trace (it is overdetermined), meaning that there are many more possible "FROG traces"—configurations of N^2 pixels—than are allowed by a physically realizable electric field. Thus, a real FROG trace with some systematic or random error added to it most likely will not correspond to a physically valid FROG trace, leading to nonconvergence of the algorithm, while convergence of the algorithm nearly guarantees that the measurement was accurate and free of artifacts.

Typically, the algorithm converges to the closest physically valid FROG trace to the input trace. In the presence of random noise, the algorithm often converges to a trace that is closer to the noise-free trace than the original (noisy) input data [22]. In the case of systematic noise, the distorted FROG trace is almost always quite unlike any physically valid trace, leading to nonconvergence of the algorithm. Fortunately, there are ways to detect this type of systematic error, even before running the algorithm, as we shall see below.

The redundancy in the FROG trace allows for several types of checks on the accuracy of the data. When analyzing experimental FROG data, in the majority of cases we have discovered systematic errors in the data-collection apparatus simply from inspection of the FROG trace and its marginals (see Section IV), allowing us to "debug" experiments remotely. In methods that do not provide such checks, these errors would probably have gone undiscovered. Far from being a disadvantage, the two-dimensional, redundant nature of the FROG data set gives us the powerful ability to detect systematic errors in the experimental setup.

This sensitivity to systematic experimental errors is a crucial part of the diagnostic power of FROG. The FROG technique consists of a pump-probe-type process followed by the measurement of the spectrum. If one is unable to do these two experiments correctly (due to spatial chirp in the beam, incorrect calibrations, beam distortions at the focus, etc.), then it is unlikely that any other experimental results generated by the apparatus will be valid, as almost all experiments require

a pump-probe and/or spectral measurement. For example, an autocorrelation made using this apparatus will be in error, although (unlike in FROG) no indication of this will generally be available. Thus FROG allows an extremely germane check on the performance of the experimental apparatus.

IV. Self-Consistency Checks for FROG Data: The Marginals

The fact that the FROG trace contains redundant data (i.e., it is overdetermined) allows for some fairly simple, yet powerful, checks on the consistency of the experimental data. These checks involve the marginals, which are the one-dimensional curves obtained by integrating the FROG trace over one of its coordinates. The marginals can be compared to quantities involving the pulse's spectrum and autocorrelation. When these easily computed quantities agree, one can be fairly sure of the consistency of the data. The more useful of the two marginals is the frequency marginal

$$M(\omega) = \int_{-\infty}^{\infty} d\tau I_{\mathrm{FROG}}(\omega, \tau) \tag{6}$$

obtained by integration of the FROG trace over the delay variable.

The form of the marginals was explored in an earlier work [19]. As an example, we consider here the case of SHG FROG, which is perhaps the easiest to understand. In this case, the frequency marginal should have a functional form identical to the autoconvolution of the pulse spectrum $I(\omega)$

$$M_{\mathrm{SHG}}(\omega - 2\omega_0) = I(\omega - \omega_0)^* I(\omega - \omega_0) \tag{7}$$

where ω_0 is the carrier frequency, the asterisk denotes convolution, and we have used (3), (5), and (6). Therefore, by simply measuring the fundamental pulse spectrum, one can easily check the consistency of the FROG data. If the SHG FROG frequency marginal does not agree with the autoconvolution of the pulse spectrum, one can be assured that there is a systematic experimental error somewhere in the system. Some typical possibilities are incorrect wavelength or temporal calibrations of the FROG data, a spectrometer/camera response that varies with wavelength, insufficient doubling-crystal bandwidth, spatial chirp, spatio-temporal distortions of the pulse at the focus, etc. Under these conditions, it is illogical to expect that the spectrum of the field retrieved by the FROG algorithm should match the experimentally measured spectrum—the failure of (7) indicates that the experimental SHG FROG data are inconsistent with the measured pulse spectrum.

In the PG and SD FROG geometries, the same sorts of consistency checks are available. The exact form of the frequency marginal varies with the type of geometry, however. In the case of PG FROG, the frequency marginal has the form

$$M_{\mathrm{PG}}(\omega - \omega_0) = I(\omega - \omega_0)^* \mathcal{T}\{A^{(2)}(\tau)\} \tag{8}$$

where $\mathcal{T}\cdot$ indicates Fourier transform and $A^{(2)}(\tau)$ is the usual second-order intensity autocorrelation. Thus in PG FROG, it is possible to check the consistency of the FROG data with both the spectrum and the autocorrelation of the pulse through a simple convolution of these two quantities. In SD FROG, the frequency marginal is

$$M_{\mathrm{SD}}(\omega - \omega_0) = I(-\omega + \omega_0)^* I_{\mathrm{SH}}(\omega - 2\omega_0) \tag{9}$$

while for THG FROG we find

$$M_{\mathrm{THG}}(\omega - \omega_0) = I(\omega - \omega_0)^* I_{\mathrm{SH}}(\omega - 2\omega_0) \tag{10}$$

where $I_{\mathrm{SH}}(\omega)$ is the spectrum of the second harmonic of the pulse. When using the marginals as a check on the data consistency, it is important to remember to center all spectra on the same carrier frequency ω_0 and to interpolate laboratory spectra to a constant frequency spacing. In the case of extremely broadband light, in addition to merely changing the abscissa of the data, this also involves a nonlinear scaling factor

$$I_\omega(\omega) = \frac{\lambda^2}{2\pi c} I_\lambda(\lambda)|_{\lambda = 2\pi c/\omega} \tag{11}$$

to change from a spectrum sampled on a constant-wavelength interval (as most spectrometers yield) to one sampled on a constant-frequency interval [as needed for a fast Fourier-transform (FFT)].

In order to illustrate the usefulness of the marginals, we multiplied the SHG FROG trace of a self-phase-modulated pulse by a Gaussian spectral-response function. This sort of spectral distortion is often encountered in practice, since spectrometers optimized for flat response at the fundamental frequency rarely perform well at the second-harmonic frequency. The frequency marginal of both the original trace and the distorted trace are plotted in Fig. 1(a), along with the spectral-response function. When this distorted FROG trace was input into the algorithm, a distorted pulse was retrieved. The results are seen in Fig. 1(b), where we see that the retrieved spectrum, not surprisingly, does not match the measured fundamental spectrum. The fact that the retrieved spectrum does not match the measured spectrum would perhaps be mistaken for a failure of the algorithm, when in fact it is a systematic error in the data. The algorithm does not converge to arbitrarily low FROG error for a distorted trace such as this; however, *a priori* it is difficult to tell whether the failure to converge is the result of excessive experimental noise (which leads to a nonzero FROG error [22]) or a systematic error in the data. A comparison of the frequency marginal and the autoconvolution of the spectrum, as in Fig. 1(a), shows unambiguously that the problem lies with the data (and thus in the experimental apparatus) rather than with the algorithm. Thus, the marginals provide a powerful check on the experimental apparatus.

V. FROG Sampling Rate

The issue of sampling in a time-frequency technique such as FROG is a subtle one. Before we directly address this issue, there are some basic terms and concepts that must be established.

The FROG trace is considered properly sampled when the data are not truncated, i.e., all the nonzero FROG data lie within the FROG-trace grid. Of course, in a strict mathematical sense, FROG-trace data never go to zero: the FROG trace of

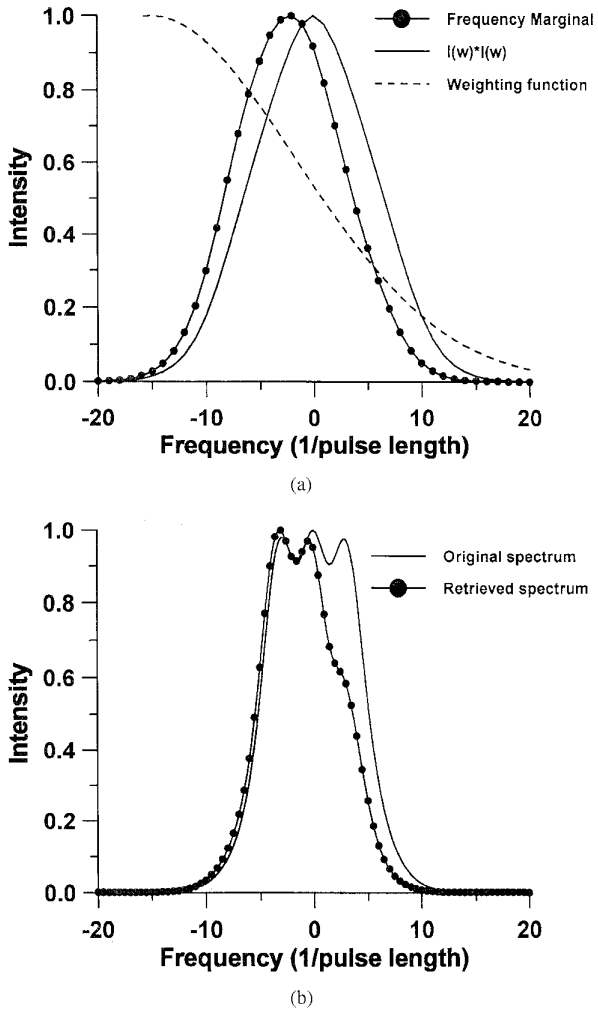

Fig. 1. (a) The autoconvolution of the fundamental spectrum and the frequency marginal of the FROG trace after multiplication with an exponentially varying spectral response (dashed line). (b) The true fundamental spectrum of the pulse and the spectrum of the pulse retrieved from the distorted SHG FROG trace.

a Gaussian pulse, for example, decays like a Gaussian in all directions. As a practical criterion, we will consider FROG-trace data to be properly sampled when all of the data points with an intensity 10^{-4} of the peak of the trace or greater are contained on the FROG-trace grid. A data set that satisfies this criterion is said to satisfy the FROG sampling rate (FSR). A FROG trace that satisfies the FSR completely determines the intensity and phase of the pulse that created it, as we shall show.

A similar criterion that arises in signal processing is the Nyquist criterion [23]. A band-limited function is properly sampled if the sampling rate is at least as high as the so-called Nyquist rate, where the highest frequency occurring in the signal is sampled at least twice per period. In this case, the discrete samples of the signal contain all the information about the signal; sampling at a higher rate, noise considerations aside, produces no new information. Note that a signal sampled at the Nyquist rate is not necessarily aesthetically pleasing; a sine wave sampled twice per period looks like a sawtooth, while a Nyquist-sampled Gaussian has only one point above the half-maximum and only five points larger than 1% of the peak! (Note that a Gaussian can never be a band-limited function, as it has infinite extent in both time and frequency. However, in analogy with the FROG trace, we can consider the 10^{-4} points of the spectrum as the "band limit.") Nevertheless, despite their jagged appearance, these sampled signals contain complete information about the original waveform.

The FSR is a stricter criterion than the Nyquist rate. A pulse sampled at the Nyquist rate will have FROG-trace data that are truncated at significant energies: the data are not fully contained on the FROG-trace grid, and thus the FSR is not satisfied. In order to get all of the data on the grid and satisfy the FSR, the pulse must be sampled at a higher rate than the Nyquist rate. This means that any FROG-trace data that satisfies the FSR automatically contains all the information about the pulse. No new information is gained (except perhaps noise immunity) by sampling at higher rates.

The FSR actually comprises two limits to the sampling of the data. First, the data must be sampled with a small enough temporal step Δt so that the data do not extend off of the FROG-trace grid in the frequency direction (the temporal step size and the full frequency width of the FROG-trace grid are inversely related). For a Gaussian spectrum and PG FROG, this occurs when

$$\Delta t \leq \frac{1}{6.3 f_p} \approx \frac{\lambda_0^2}{6.3 c \lambda_p} \tag{12}$$

where λ_0 is the central wavelength and $\lambda_p(f_p)$ is the FWHM of the spectrum in wavelength (frequency). If the pulse is transform limited, this becomes $t_p/\Delta t \geq 2.78$, where t_p is the FWHM of the pulse temporal width. The second limit inherent in the FSR is that the temporal-sampling step size be large enough that the FROG-trace data do not extend off the grid in the time-delay direction. For a Gaussian pulse in time, this evaluates to

$$\Delta t \geq \frac{4.5 t_p}{N} \tag{13}$$

for PG FROG. Thus, for a Gaussian transform-limited pulse of 100-fs FWHM, the temporal step size used to generate the FROG-trace data must lie between 7 fs and 36 fs for a 64 $\times$ 64 grid.

In general, then, the procedure for selecting the sampling rate and grid size for an arbitrary pulse is as follows. First, the upper limit for the temporal step size is set by (12). The minimum grid size N is then selected so that the extent of the FROG-trace grid in the delay direction contains all of the nonzero FROG-trace data. If the pulse is complex enough that both limits cannot be satisfied simultaneously, the grid size N must be increased. The numerical factors in (12) and (13) were calculated for Gaussian pulses; other pulse and spectral shapes will have slightly different FSR limits. Differing FROG geometries will also have slightly different FSR's.

Sampling a pulse close to both limits of the FSR, so that the data nearly fill the FROG-trace grid, allows one to use a very small-size FROG-trace grid. (The FSR can be satisfied

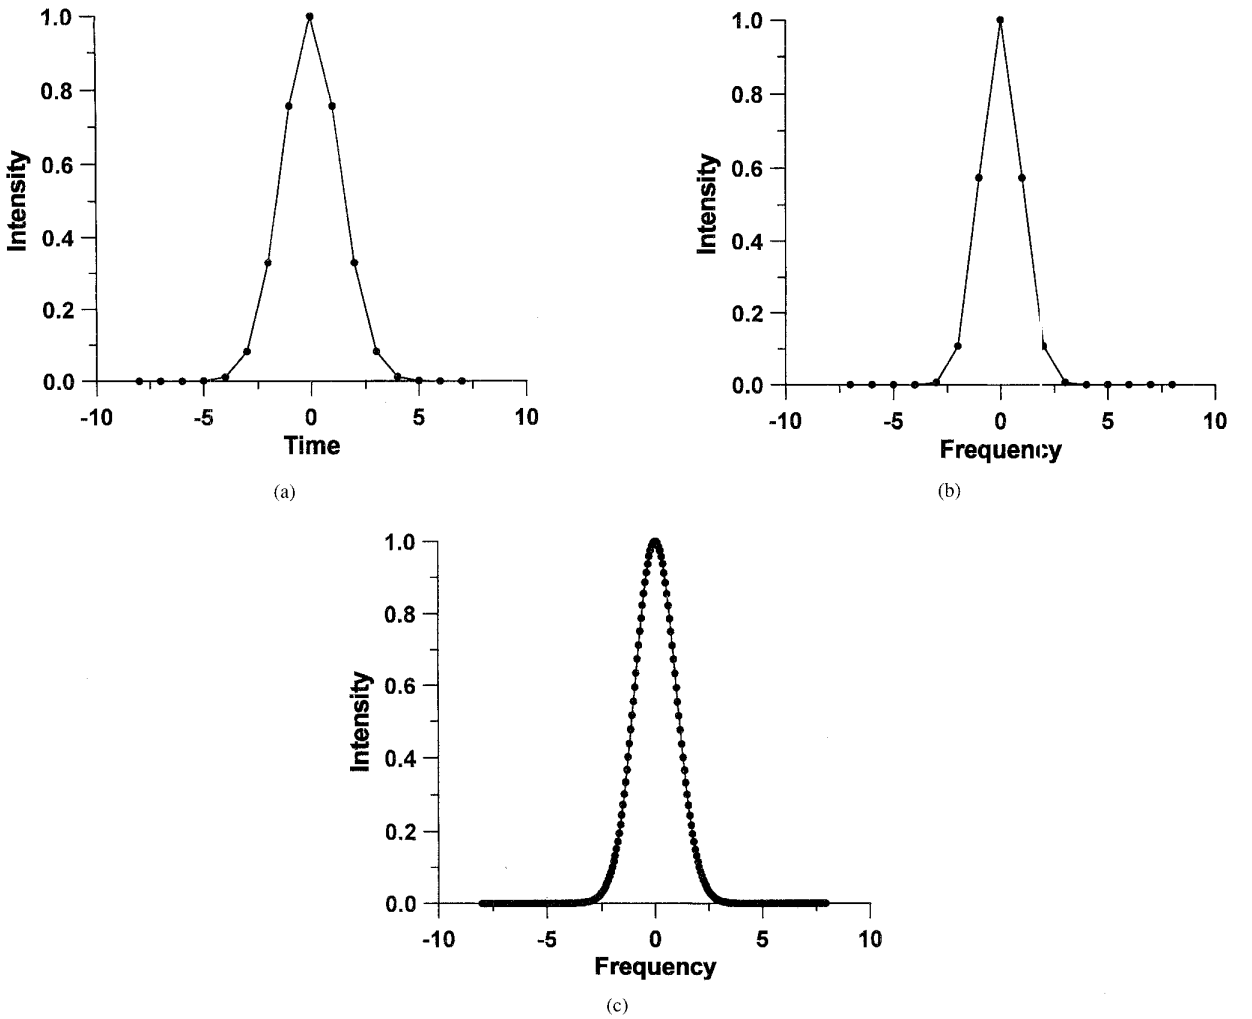

Fig. 2. The intensity as a function of (a) time and (b) frequency for a transform-limited Gaussian pulse sampled near the FROG sampling-rate limit. The pulse had a FWHM of 3.16 pixels in the time domain on a 16 × 16 pixel grid. In both the time and frequency domain, the pulse profile is jagged and not aesthetically pleasing. (c) The time-domain field in (a) was padded with zeros to fit in a 256-element array and then Fourier-transformed to give this spectrum. The spectrum is now densely sampled and looks quite Gaussian. Although the field in (a) and (b) looks undersampled, the information content is equivalent to that in (c).

for a transform-limited pulse in PG FROG on a 16 × 16 pixel trace!) Using a smaller number of pixels is of great advantage in retrieving the pulse: the FROG algorithm slows down as $N^2 \ln(N)$. However, the resulting retrieved pulse will look jagged and decidedly un-Gaussian in both the time and frequency domain.

One way to avoid the jaggedness in the output field is to sample at a greater rate than the FSR. This necessarily involves going to a larger grid size for the FROG trace. If the temporal sampling rate is kept constant while moving to a larger grid, the spectrum of the pulse will be sampled at a higher rate (because the frequency increment is the reciprocal of the temporal range). If the temporal sampling rate is increased commensurate with the increase in the size of the FROG trace (keeping the temporal range the same), then the sampling rate in the frequency domain will remain constant (while the temporal rate increases). However, this method of increasing the grid size extracts its toll in the slowing down of the FROG algorithm.

A better solution to this problem is to sample the pulse at a rate near the FSR: use a FROG-trace grid that is small enough that the FROG-trace data nearly fill the grid. Pulse retrieval using the algorithm will be swift on a small grid, but the output fields will look jagged. The jaggedness of the retrieved field can be eliminated, however, by taking advantage of the completeness of a Nyquist-sampled field. Specifically, in order to get a higher point density in the frequency domain, simply pad the $E(t)$ returned by the algorithm with zeros on the left and right and put it into a larger array before Fourier-transforming $E(t)$ to obtain $E(\omega)$. A high-point density in the time domain can be achieved by similarly padding the field in the frequency domain before inverse Fourier-transforming to the time domain. This is a valid procedure, at least to the order of the approximations that we made for band-limited functions above (i.e., truncation at or below the 10^{-4} level). Thus, it is straightforward to obtain as many data points as desired in each domain.

This procedure is illustrated in Fig. 2. Fig. 2(a) shows a Gaussian transform-limited pulse sampled so that its intensity FWHM is 3.16 pixels. This pulse satisfies the FSR for a 16 × 16 pixel grid. Both the temporal and spectral intensity [Fig. 2(b)] are quite jagged, yet are sampled at a higher rate than the Nyquist rate. Fig. 2(c) shows the spectrum of the exact same field, where the field in the time domain was padded out to 256 pixels before the Fourier transform to obtain $E(\omega)$. We see that the spectrum is now extremely smooth and densely sampled. Fig. 2(b) and (c) display the same waveform and the same information but have a quite different appearance.

One particularly pleasing sampling rate for FROG data is one that makes the sampling rates roughly equal in both domains. Equal sampling in both the time and frequency domains could be defined as having the ratio of the delay step size to the temporal FWHM of the pulse be the same as the ratio of the frequency step size to the spectral FWHM of the pulse. Therefore, equal sampling in both the time and frequency domains is satisfied when the delay step size is $\Delta t = t_p/M$ and the frequency step size $\Delta f = f_p/M$. On a FROG trace of $N \times N$ pixels, the frequency step size is set to $\Delta f = 1/(N\Delta t)$ by the fast Fourier transform.[1] A simple calculation yields

$$M = \sqrt{t_p f_p N} \approx \sqrt{\frac{t_p \lambda_p N c}{\lambda_c^2}}. \tag{14}$$

For a transform-limited Gaussian pulse, the spectral FWHM is defined in terms of the temporal FWHM t_p as $f_p = 2\ln(2)/\pi t_p$, so that $M = \sqrt{2N\ln(2)/\pi}$. Note that for this sampling, the SHG FROG-trace contours describe perfect circles, while in PG FROG the contours are elliptical. For pulses farther from the transform limit, the delay step size can be decreased while still maintaining a reasonable frequency-domain sampling rate.

Fig. 3 demonstrates this effect. In Fig. 3(a), we see the (64 × 64) PG FROG trace of a Gaussian transform-limited pulse sampled in time at a rate of $\Delta t = t_p/20$. The trace extends over only a few pixels in frequency. In Fig. 3(b), we see the same pulse sampled at the optimum rate of $\Delta t = t_p/5.3$. In this case, the trace (and hence the resulting fields) are sampled evenly in both time and frequency.

It should be stressed that as long as the pulse satisfies the FSR, the FROG trace contains equivalent (and complete) information regardless of the sampling rate. It is only the visual appeal of the trace and the resulting fields that is affected by the choice of sampling rates. Also, intuitively we expect that a trace of the form of Fig. 3(b) will have a more robust retrieval in the algorithm than a trace like that in Fig. 3(a).

[1] Fundamentally, the delay coordinate and the frequency coordinate are not constrained by the fast Fourier-transform (FFT): the FFT relates the *temporal* coordinate t to the frequency coordinate. In principle, one could use a different step size when sampling the delay than when sampling the time t. However, this will lead to many practical difficulties in the programming of the algorithm, so that for convenience's sake the step size in both delay τ and time t are usually taken to be identical. This leads to an effective coupling, through the FFT, of the delay and frequency axes. We assume this convention throughout all of our papers.

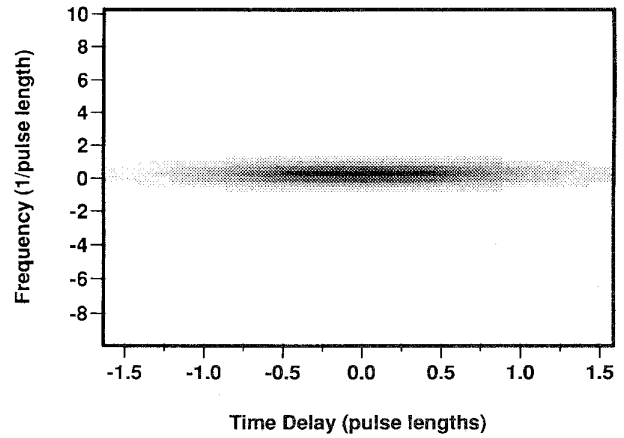

(a)

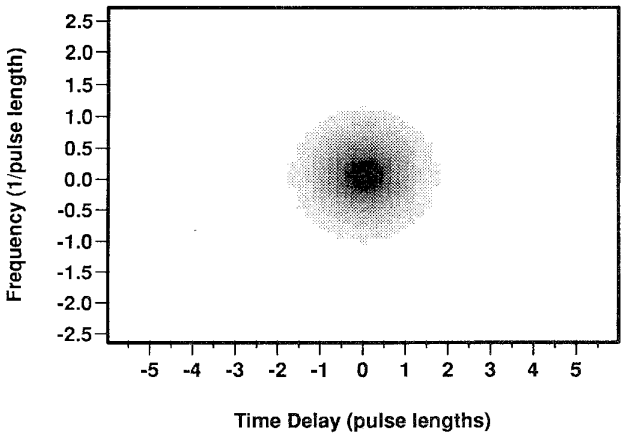

(b)

Fig. 3. Two PG FROG traces (64 × 64) of the same transform-limited Gaussian pulse. The field in (a) was sampled at 20 points per temporal FWHM, while the field in (b) was sampled at the optimum 5.3 points per temporal FWHM. Although the two traces contain equivalent information, the trace in (a) extends over only a few pixels in the frequency direction, while in (b) there is a more even distribution between time and frequency.

VI. Limitations to the FROG Technique

In this section, we shall explore some experimental limitations to the FROG technique. As FROG uses a nonlinear process to generate the signal field, there is a limit to the lowest intensity pulse that can be measured with a given geometry. We shall quantify some of these limits. Also, we will discuss the issues of maximum allowable diffraction efficiency in the $\chi^{(3)}$ experiments, the effects of a finite thickness for the medium that generates the FROG signal, and the maximum obtainable time-bandwidth products for a FROG trace of fixed size.

A. Low-Power Limits

Because a nonlinear process generates the FROG signal field, FROG requires a certain minimum power in the pulse to be measured. The single-shot version of FROG requires more power in the pulse than the multishot version, because in the single-shot version the pulse is focused to a line focus instead

TABLE I
LOW-POWER LIMITS TO FROG

FROG geometry	multishot	signal avg. single-shot	true single-shot
SHG	190 W (15 pJ)	25 kW (2 nJ)	3 MW* (240 nJ)
PG	6 MW (480 nJ)		45 MW (3.6 μJ)

* indicates values extrapolated from other entries in the table (see text).

of a point focus as in the multishot version. For a pulse of given power, this results in a lower intensity at the nonlinear medium for the single-shot case. Also, the minimum pulse power necessary to generate a usable FROG trace will vary with the geometry used. Second-order effects (SHG) require lower intensities at the nonlinear medium than third-order effects (PG, SD, and THG), so that SHG FROG can function with lower pulse powers.

Multishot SHG FROG is the most sensitive of all the FROG geometries. In our experiments, we used a KDP crystal as the nonlinear medium. Pulses from a Ti:Sapphire oscillator operating at 860 nm with 80-fs FWHM were split into two beams by a beamsplitter and focused with a 200-mm focal-length spherical lens onto a 300-μm-thick KDP sample. The noncollinear SHG signal beam was collected and recollimated with another spherical lens, dispersed by a grating, and focused onto a Pulnix 8-b CCD camera. We were able to record SHG FROG traces at a 40 : 1 peak signal-to-dark-current-noise ratio with our setup for peak pulse powers as low as 190 W (15 pJ for our pulses) in each of the two beams. The 96-MHz pulse train is averaged over the camera read-out time of 16 ms. The spectrum covered approximately 50 pixels on the camera.

The apparatus was not optimized, and further gains in sensitivity could be realized through the use of a crystal with a higher nonlinearity, tighter focusing, or using a more sensitive or less noisy camera (our 8-b camera produced about 5 counts of dark current per pixel over a measurement period). With these improvements, we expect that multishot SHG FROG traces could be made for pulses in the femtojoule range when the pulse length is 100 fs or shorter.

In the single-shot SHG FROG configuration, we were able to record FROG traces for input-pulse powers as low as 25 kW (2 nJ for our pulses) when using signal averaging of the 96-MHz pulse train over the 16-ms camera read-out time. From these numbers, we anticipate that true single-shot operation should be available for pulse powers of 3 MW.

In PG FROG, a multishot geometry using 1 mm of UV-grade fused silica as the nonlinear medium produced acceptable signals for a total input power of 300 nJ for 45-fs pulses (6-MW peak power). This experimental setup used a 150-cm focal-length lens to focus the beams into the sample. Single-shot PG FROG was reported in [6].

Our results for the low-power limits to FROG are summarized in Table I.

More slowly responding nonlinearities usually have much stronger responses than instantaneous nonlinearities; however, the use of very slowly responding nonlinearities, such as photorefractives, is contraindicated in FROG. It has been shown that for these very slow nonlinearities, the only information available is the frequency spectrum of the pulse [24].

B. Signal Efficiency Limit

In PG and SD FROG, the signal efficiency of the process used to generate the FROG signal field cannot be made arbitrarily high, even if the pulse energy is high, because the same fundamental mechanism that is responsible for generating the signal field is also responsible for self- and cross-phase modulation, which distort the pulse spectrum and hence the FROG trace. Here, we make a simple estimate (ignoring propagation effects) of the maximum signal efficiency available without significant distortion of the pulse being measured. We calculate the FROG signal efficiency as a function of the maximum gate-beam-induced phase change $\Delta\phi$ of the probe beam. We then calculate the amount of pulse distortion for a given amount of cross-phase modulation, which is also measured by $\Delta\phi$. In this way we can place an upper limit to the FROG signal efficiency.

In PG FROG, the signal efficiency is easily found to be $\eta = (\Delta\phi)^2/4$, where $\Delta\phi = k_0 n_2 I L$, and k_0 is the vacuum wavevector, n_2 is the nonlinear refractive index, I is the intensity of the gate pulse, and L is the length of the medium. To estimate the effects of cross-phase modulation for a given value of $\Delta\phi$, we examine a Gaussian pulse with a time-dependent phase of $\Phi(t) = \Delta\phi I(t)$, where $I(t)$ is the intensity of the pulse normalized to a peak of unity. The results of numerical calculations are summarized in Table II. Because the time-domain intensity profile is unaffected, the amount of spectral broadening is equal to the increase in the time-bandwidth product.

We see that in order to keep the rms spectral broadening less than 1.5%, $\Delta\phi$ must be kept below 0.2 radians, which corresponds to a peak signal efficiency of $(0.2)^2/4$ or 1%. This is a much higher limit to the signal efficiency than (erroneously) previously reported [6]. In SD FROG, the argument is essentially identical. The limits to maximum signal efficiency in SHG FROG (again, about 1%) stem from pump-depletion considerations and were discussed earlier [16].

C. Finite-Thickness Medium

In the single-shot configuration of FROG, the time resolution is limited by the finite thickness of the medium used to generate the FROG signal. Due to the differing propagation directions of the two beams, they experience differing delays

TABLE II
SIGNAL-EFFICIENCY LIMITS IN THIRD-ORDER FROG

$\Delta\phi$	peak signal efficiency	Spectral broadening	
		rms	FWHM
0.2 rad	1 %	1.5 %	0.16 %
0.6 rad	9 %	13 %	1.4 %

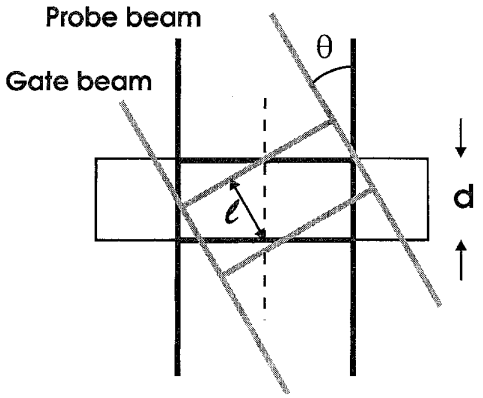

Fig. 4. Origin of the time "blurring" in single-shot FROG measurements. The probe beam (black) propagates a distance d through the medium, while the gate beam (gray) propagates a distance l. This results in a blurring of the time delay at any given spatial coordinate (dotted line) due to the finite thickness of the medium.

as they propagate through the medium [25] (see Fig. 4). This leads to a "blurring" of the time delay. This effect will increase with the thickness of the medium. Here we calculate the amount of blurring and demonstrate its effects on the FROG pulse-retrieval process.

In the single-shot PG FROG geometry, the probe beam is typically incident on the nonlinear medium in a perpendicular fashion, while the gate beam propagates at an (internal) angle θ with respect to the probe beam. Although the time delay is nominally encoded along the transverse spatial coordinate, the differing propagation distances (d for the probe beam, and ℓ for the gate beam, as seen in Fig. 4) of the two beams through the finite thickness medium cause each spatial point to sample a range of blurred time delays. For a medium of refractive index n and thickness d, the blurring of the time delay is

$$\begin{aligned}\Delta\tau &= \frac{nd}{c}(1-\cos\theta)\\ &= \frac{2nd}{c}\sin^2\frac{\theta}{2}\\ &\approx \frac{nd\theta^2}{2c} \quad \text{for } \theta \ll 1 \end{aligned} \tag{15}$$

where c is the speed of light. The signal field leaving the medium will then be blurred such that

$$E_{\text{sig}}^{\text{blurred}}(t,\tau) = \int_{\tau-\Delta\tau/2}^{\tau+\Delta\tau/2} d\tau' E_{\text{sig}}(t,\tau') \tag{16}$$

where $E_{\text{sig}}(t,\tau)$ is the signal field from (1).

The blurring of the time delay is equivalent to a blurring of the gate pulse. Note that we can write the signal field as

$$E_{\text{sig}}(t,\tau) = E(t)G(t-\tau) \tag{17}$$

where $G(t)$ is the gate function, which takes on various forms depending on the FROG geometry used according to (1)–(4). Then (16) becomes

$$E_{\text{sig}}^{\text{blurred}}(t,\tau) = E(t)\int_{\tau-\Delta\tau/2}^{\tau+\Delta\tau/2} d\tau' G(t-\tau') \tag{18}$$

so that the effect of the blurring along the delay coordinate is equivalent to a blurring of the gate function.

In order to quantify the effect of this blurring on the PG FROG pulse-retrieval process, we used numerically simulated pulses and integrated the signal field according to (16) before generating a FROG trace. We then used this distorted FROG trace as input to the pulse-retrieval algorithm and monitored the distortion of the retrieved pulse as a function of the amount of delay blurring $\Delta\tau$. The results are seen in Table III.

The effect of the blurring is remarkably small. For a transform-limited Gaussian pulse of intensity FWHM t_p, the retrieved pulse was broadened by a factor of roughly $0.1\Delta\tau/t_p$. In other words, a $\Delta\tau$ that was 10% of the FWHM of the pulse resulted in only a 1% broadening in the retrieved pulse. This relationship held for values of $\Delta\tau/t_p$ of up to 0.6. Even for $\Delta\tau = t_p$, the resultant broadening in the retrieved pulse was only 15%. For a linearly chirped Gaussian with an rms time-bandwidth product of over 1.5 (transform limit is 0.5), the results were essentially the same, with less than 1% distortion for a blurring of $\Delta\tau < 0.3t_p$. For this pulse, a blurring of $\Delta\tau = t_p$ led to a 12% broadening of the FWHM and a 9% error in the quadratic-phase coefficient of the retrieved pulse. Thus we can conclude the limitations imposed by blurring of the temporal delay in a finite-thickness medium are not severe.

Fig. 5 shows a pulse retrieved by the FROG algorithm when the signal field was blurred by an amount equal to half the original pulse width, or $\Delta\tau = 0.5t_p$. The very small amount of broadening in the retrieved pulse illustrates the relatively small effect of the finite medium thickness in FROG.

This remarkable insensitivity to the blurring along the delay coordinate is reminiscent of the insensitivity of SHG FROG to blurring along the temporal coordinate (in this case due to group-velocity mismatch between the fundamental and second harmonic), reported in [16]. There also it was found that a blurring of the temporal coordinate by a window with a width equal to the FWHM of the pulse resulted in only a 10%

TABLE III
EFFECTS OF FINITE MEDIUM THICKNESS

delay blurring $\Delta\tau$	transform-limited	linearly-chirped
$0.2t_p$	1.3 %	0.9 % (< 1 %)
$0.4t_p$	3.2 %	1.6 % (2%)
$0.6t_p$	6.3 %	5.3 % (3.8 %)
$0.8t_p$	10.4 %	7.5 % (6.3%)
$1.0t_p$	15.3 %	11.5 % (9.3 %)

The numbers in parentheses represent the error in the quadratic coefficient of the retrieved temporal phase. t_p is the pulse length.

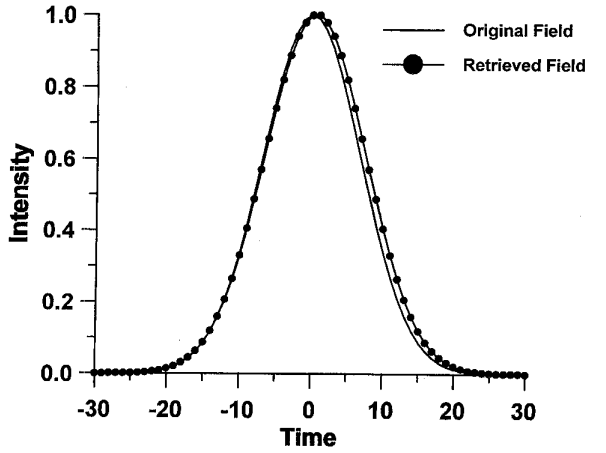

Fig. 5. Effect of time-delay blurring on the retrieval of pulses. The solid line is the original field, while the dotted line is the field retrieved by the FROG algorithm after the signal-field time delay was blurred by a $\Delta\tau$ of one-half the FWHM of the original pulse. The very small amount of distortion in the retrieved pulse indicates that FROG is highly insensitive to the effect of a finite-thickness medium. The retrieved phase (not shown) is flat.

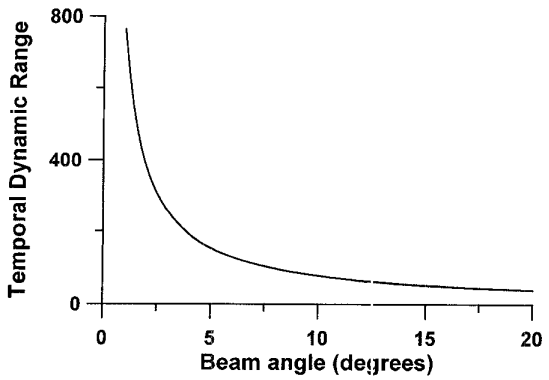

Fig. 6. The temporal-dynamic range for a beam diameter of $D = 1$ cm, and a medium with a thickness $d = 1$ mm, and an index of refraction $n = 1.5$. The temporal-dynamic range is the ratio of the maximum allowable pulse length (limited by the finite size of the beam) to the minimum allowable pulse length (limited by the finite thickness of the medium).

broadening of a transform-limited pulse. Evidently the phase-retrieval algorithm at the heart of the FROG technique is quite robust against these common systematic errors encountered in experimental work.

Another effect that arises due to the finite thickness of the nonlinear medium is group-velocity dispersion (GVD). This will affect multishot as well as single-shot geometries. As the pulse being measured propagates through the nonlinear medium, it acquires a spectral phase, which is dominated by the quadratic term (linear chirp), which causes the pulse to broaden. From [27] and [28], we can calculate straightforwardly the amount of broadening of the pulse due to GVD. If we begin with a transform-limited Gaussian pulse and require that the pulse is broadened by less than ε, then the maximum allowable thickness of the medium is

$$L_{\max} = \frac{\sqrt{2\varepsilon + \varepsilon^2}}{4\beta_2 \ln(2)} \tau_0^2 \tag{19}$$

where τ_0 is the original pulse FWHM and β_2 is the dispersion parameter of the nonlinear medium. For an allowable pulse broadening of 1% ($\varepsilon = 0.01$), an original pulsewidth of 100 fs at a wavelength of 800 nm, and using BK7 glass as a nonlinear medium ($\beta_2 = 495$ fs^2/cm), the maximum allowable length of the medium is 1.03 cm. This is not a very strict limitation, although the effect scales with pulsewidth squared, so that this requirement becomes stricter for shorter pulses.

D. Temporal-Dynamic Range of FROG

The temporal resolution of the single-shot FROG technique, as roughly given by (15), should be considered in conjunction with the maximum time-delay range available from the same beam geometry. For a beam with a usable (flat intensity) width D on the sample, the maximum time-delay range is easily found from simple geometry to be

$$\Delta T = \frac{D \tan\theta}{c}. \tag{20}$$

In order to get an accurate measurement of a pulse, its temporal length must fall inside the range specified by $\Delta\tau$ and ΔT, so that the ratio

$$\frac{\Delta T}{\Delta \tau} = \frac{D}{d} \frac{n \tan\theta}{2\sin^2(\theta/2)} \tag{21}$$

TABLE IV
MAXIMUM TIME-BANDWIDTH PRODUCTS FOR VARIOUS FROG TRACES

Pulse type	PG	SD	THG	SHG
Linear chirp	8.8 (7.8)	4.9 (4.3)	4.8 (4.2)	5.1 (4.5)
Self-phase mod.	20 (20)	7.5 (7.6)	6.6 (6.7)	8.8 (9.0)
Spectral cubic phase	3.4 (0.89)	3.0 (0.86)	2.6 (0.81)	3.3 (0.87)

forms a kind of "temporal-dynamic range" for the technique. In Fig. 6, we have plotted the temporal-dynamic range for the typical values $D = 1$ cm, $d = 1$ mm, and $n = 1.5$. The temporal-dynamic range increases with decreasing crossing angle and is 153 at a beam-crossing angle of 5°, where $\Delta\tau = 19$ fs and $\Delta T = 2.9$ ps. These values will be different for different beam diameters D and medium thicknesses d, of course. Note that multishot FROG geometries have no inherent upper limit to the time delay (which is obtained through delay lines) and thus sport an infinite temporal-dynamic range.

E. Maximum Time-Bandwidth Limit

A given digitized FROG trace of size N cannot represent a pulse with an arbitrarily large complexity. For pulses with too large a time-bandwidth product (the time-bandwidth product forms a convenient measure of pulse complexity), the trace will have significant intensity off the edge of the FROG-trace grid, and the pulse will be aliased in the time and/or frequency domains (as discussed in the section on sampling). Therefore, there is an upper limit to the time-bandwidth product of a pulse that can be properly represented on a FROG trace of a given size (number of pixels). We have numerically determined these upper limits for FROG traces of several sizes. (In making this calculation, we continue with our convention of generating an $N \times N$ FROG trace from a field N elements long.)

The time-bandwidth products reported in this work are usually quoted using both the full-width at half maximum (FWHM) and the root-mean-square (rms) measures. The FWHM (rms) time-bandwidth products utilize the linear (angular) frequency, such that a transform-limited Gaussian has a FWHM (rms) time-bandwidth product of 0.441 (0.5). It should be noted that an rms time-bandwidth product of 0.5 is an absolute minimum for all waveforms, while the FWHM time-bandwidth product can range quite low (a transform-limited Lorentzian has a FWHM time-bandwidth product of 0.221, while its rms time-bandwidth product is 2.47). The rms time-bandwidth product is a better measure for theoretical and numerically generated pulses, while for experimental pulses it can suffer from undue influence of noise in the wings of the pulse.

We found that the upper limit for the time-bandwidth product varies strongly with the form of the pulse. Therefore, we used pulses with linear chirp, self-phase modulation, and spectral-cubic phase. The results for a 128×128 FROG trace are seen in Table IV for the four main FROG geometries. The criterion chosen was that the intensity at the edge of the FROG-trace grid was 10^{-4} or less of the peak intensity. For a FROG trace of $N \times N$, we found that the maximum available time-bandwidth product scales roughly with N, except for the case of self-phase-modulated pulses, where it appears to scale roughly like $N^{1.25}$. The algorithm was able to retrieve all of the pulses with maximum time-bandwidth product. However, it should be noted that these were noise-free, simulated pulses, and the presence of noise may make these pulses more difficult to retrieve (the amount of redundancy is minimal for data that fills the FROG-trace grid). Therefore, with experimental data, it is probably best to use a larger array when the time-bandwidth product of the pulse approaches the maximum value.

VII. A NOTE ON THE FROG ERROR

The FROG error G is an important quantity, as it is used in the algorithm to guide and monitor the retrieval and also to measure the quality of the retrieved FROG trace. The FROG error is computed as an rms average across the entire trace

$$G = \sqrt{\frac{1}{N^2}\sum_{i,j=1}^{N} |I_{\mathrm{FROG}}(\omega_i,\tau_j) - \alpha I_{\mathrm{FROG}}^{(k)}(\omega_i,\tau_j)|^2} \quad (22)$$

of the difference between the experimental FROG trace $I_{\mathrm{FROG}}^{(k)}$ and the retrieved or reconstructed FROG trace I_{FROG}. I_{FROG} is always normalized to a peak of unity. In the case of numerically generated data (i.e., noise free), the scaling parameter α is selected so that $I_{\mathrm{FROG}}^{(k)}$ is also normalized to a peak of unity. In the case of experimental data, this is not appropriate. The reason is that if the highest intensity pixel in I_{FROG} is corrupted by noise, the normalization of I_{FROG} will be skewed, thus biasing the calculation of G and leading to incorrectly retrieved pulses.

For experimental data, we use the following procedure. The experimental data I_{FROG} are as usual normalized to a peak of unity. On each iteration (labeled by k), the algorithm calculates a new estimate for the retrieved FROG trace $I_{\mathrm{FROG}}^{(k)}$. When calculating the error G, the normalization of the retrieved trace $I_{\mathrm{FROG}}^{(k)}$ is allowed to vary to a value that gives the minimum value for the error G. In other words, G is minimized with respect to α on each iteration. This procedure reduces the

sensitivity of the algorithm to the effects of noise-distorted normalization in experimental data. Although the calculation of the error on each iteration takes slightly longer than before (due to the minimization), we find that the algorithm generally converges in fewer steps, so that the amount of real time taken is less.

The algorithm acts to minimize the value of G. In practice, two things can limit the lowest achievable G for an experimental trace. In the first case of noise-free data (no random additive or multiplicative noise), the lowest retrievable G will be limited by distortions, systematic errors, stray light, and other deviations from a physically valid [one that obeys (1)–(5)] FROG trace. It is difficult to quantify these sorts of deviations or predict the effect on the retrieved pulse. In the second case of perfect FROG data that is corrupted with random noise (for example, the random dark-current background of a CCD camera), the algorithm may closely retrieve the FROG trace, but the value of G reported will be close to the noise level. This case is amenable to quantitative analysis and was extensively investigated [22]. In both cases, a visual comparison of the experimental and retrieved traces, coupled with the knowledge of G, is the best indicator of the quality of the retrieved field.

As noted above, G provides the only quantitative measure of the convergence of the FROG algorithm. Therefore, whenever experimental FROG results are given, the value of G should always be quoted. In our experience, errors of 0.005 or less result from accurate retrieval of low-noise data (128×128 pixel trace). In SHG FROG, errors of 0.002 are readily obtainable. A visual comparison of the experimental and retrieved traces is also enlightening, and should be considered essential: noise in experimental data will always raise the value of G, even if the experimental and retrieved traces are exactly the same shape. We suggest plotting the square root of the FROG-trace intensities to enhance the low-level features of the trace.

The size of the FROG-trace grid should also be reported along with the error. Due to the varying number of pixels in the trace, we find that the magnitude of the FROG error G scales like $N^{-1/2}$ for an $N \times N$ trace in the case of noise-free data. Thus, if an error between two traces generated by a pair of distinct pulses is equal to 0.01 on a 64×64 trace, it becomes an error of 0.007 07 on a 128×128 trace. Thus, care must be taken when comparing errors between traces of varying sizes.

VIII. Conclusion

We have illustrated several experimental details in the technique of FROG, used for measuring ultrashort laser pulses. It is essential to investigate these sorts of details in order for any technique to be useful. We have shown a simple yet powerful way to check the consistency of the FROG data with the measured spectrum and/or autocorrelation of the pulse. This procedure is very effective at uncovering systematic errors in the data-collection apparatus. We have addressed the subtle issue of the proper sampling rate for a laser pulse being measured by FROG and show that a properly sampled FROG trace guarantees that the laser-pulse field is Nyquist-sampled. We have quantified several limitations to the FROG technique, specifically the maximum signal efficiency allowed by self- and cross-phase modulation, and the maximum medium thickness allowed by temporal blurring of the signal field and broadening due to group-velocity dispersion. In both cases, we find that the restrictions are extremely lax, allowing a maximum diffraction efficiency of up to 1% and a temporal blurring due to finite medium thickness of up to one-half of the FWHM of the pulse. We have also quantified the maximum time-bandwidth product of a pulse that can be properly represented by a FROG trace. This value varies from 3 to 20 for a 128×128 FROG trace and also depends strongly on the type of pulse and the particular FROG geometry under consideration. In all, these results show that FROG is indeed an extremely useful technique, and we hope that these results will aid researchers who need to accurately characterize ultrashort laser pulses.

Acknowledgment

The authors would like to thank Dr. B. Kohler for providing his results on low-power PG FROG, Dr. J. Fourkas for comments concerning the thick medium calculation, and R. Jennings for assistance in the laboratory.

References

[1] J. L. A. Chilla and O. E. Martinez, "Direct determination of the amplitude and the phase of femtosecond light pulses," *Opt. Lett.*, vol. 16, pp. 39–41, 1991.

[2] K. C. Chu, J. P. Heritage, R. Grant, K. Liu, A. Dienes, W. E. White, and A. Sullivan, "Direct measurement of the spectral phase of femtosecond pulses," *Opt. Lett.*, vol. 20, no. 8, pp. 904–906, 1995.

[3] V. Wong and I. A. Walmsley, "Iterative and deterministic methods for ultrashort pulse-shape reconstruction from frequency-domain phase measurements," presented at the *Generation, Amplification, and Measurement of Ultrashort Laser Pulses II,* San Jose, CA, 1995.

[4] B. S. Prade, J. M. Schins, E. T. J. Nibbering, M. A. Franco, and A. Mysyrowickz, "A simple method for the determination of the intensity and phase of ultrashort optical pulses," *Opt. Commun.*, vol. 113, no. 1, pp. 79–84, 1995.

[5] F. Reynaud, F. Salin, and A. Barthelemy, "Measurement of phase shifts introduced by nonlinear optical phenomena on subpicosecond pulses," *Opt. Lett.*, vol. 14, no. 5, pp. 275–277, 1989.

[6] D. J. Kane and R. Trebino, "Single-shot measurement of the intensity and phase of an arbitrary ultrashort pulse by using frequency-resolved optical gating," *Opt. Lett.*, vol. 18, no. 10, pp. 823–825, 1993.

[7] R. Trebino and D. J. Kane, "Using phase retrieval to measure the intensity and phase of ultrashort pulses: frequency-resolved optical gating," *J. Opt. Soc. Amer. A*, vol. 10, no. 5, pp. 1101–1111, 1993.

[8] K. W. DeLong and R. Trebino, "Improved ultrashort pulse-retrieval algorithm for frequency-resolved optical gating," *J. Opt. Soc. Amer. A*, vol. 11, pp. 2429–2437, 1994.

[9] K. W. DeLong, D. N. Fittinghoff, R. Trebino, B. Kohler, and K. Wilson, "Pulse retrieval in frequency-resolved optical gating using the method of generalized projections," *Opt. Lett.*, vol. 19 no. 24, pp. 2152–2154, 1994.

[10] B. A. Richman, K. W. DeLong, and R. Trebino, "Temporal characterization of the Stanford mid-IR FEL micropulses by 'FROG'," *Nucl. Instrum. Methods Phys. Res. A*, vol. 358, pp. 268–271, 1995.

[11] D. J. Kane, A. J. Taylor, R. Trebino, and K. W. DeLong, "Single-shot measurement of the intensity and phase of a femtosecond UV laser pulse," *Opt. Lett.*, vol. 19, no. 14, pp. 1061–1063, 1994.

[12] T. Sharp-Clement, A. J. Taylor, and D. J. Kane, "Single-shot measurement of the amplitude and phase of ultrashort laser pulses in the ultraviolet," *Opt. Lett.*, vol. 20 no. 1, pp. 70–72, 1995.

[13] B. K. Kohler, V. V. Yakovlev, K. R. Wilson, J. Squier, K. W. DeLong, and R. Trebino, "Phase and intensity characterization of femtosecond pulses from a chirped-pulse amplifier by frequency-resolved optical gating," *Opt. Lett.,* vol. 20, no. 5, pp. 483–485, 1995.
[14] G. Taft, A. Rundquist, M. M. Murnane, H. C. Kapteyn, K. W. DeLong, R. Trebino, and I. P. Christov, "Ultrashort optical waveform measurements using frequency-resolved optical gating," *Opt. Lett.,* vol. 20, no. 7, pp. 743–745, 1995.
[15] D. J. Kane and R. Trebino, "Charaterization of arbitrary femtosecond pulses using frequency-resolved optical gating," *IEEE J. Quantum Electron.,* vol. 29, pp. 571–579, 1993.
[16] K. W. DeLong, R. Trebino, J. Hunter, and W. E. White, "Frequency-resolved optical gating using second-harmonic generation," *J. Opt. Soc. Amer. B,* vol. 11, pp. 2206–2215, 1994.
[17] J. Paye, M. Ramaswamy, J. G. Fujimoto, and E. P. Ippen, "Measurement of the amplitude and phase of ultrashort light pulses from spectrally resolved autocorrelation," *Opt. Lett.,* vol. 18, no. 22, pp. 1946–1948, 1993.
[18] T. Tsang, M. A. Krumbugel, K. W. DeLong, D. N. Fittinghoff, and R. Trebino, "Frequency-resolved optical gating using third-harmonic generation," in *Ultrafast Phenomena X*, P. Barbara, W. Knox, J. G. Fujimoto, Eds. New York: Springer-Verlag, 1996.
[19] K. W. DeLong, R. Trebino, and D. J. Kane, "A comparison of ultrashort-pulse frequency-resolved-optical-gating traces for three common beam geometries," *J. Opt. Soc. Amer. B,* vol. 11, pp. 1595–1608, 1994.
[20] L. Cohen, "Time-frequency distributions," *Proc. IEEE,* vol. 77, pp. 941–981, 1989.
[21] J. R. Fienup, "Phase retrieval algorithms: A comparison," *Appl. Opt.* vol. 21, no. 15, pp. 2758–2769, 1982.
[22] D. N. Fittinghoff, K. W. DeLong, R. Trebino, and C. L. Laders, "Noise sensitivity in frequency-resolved optical gating measurements of ultrashort pulses," *J. Opt. Soc. Amer. B,* vol. 12, no. 12, pp. 1955–1967, 1995.
[23] W. H. Press, W. T. Vetterling, and S. A. Teukolsky, *Numerical Recipes in C,* 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1992.
[24] A. M. Levine, E. Ozizmir, R. Trebino, C. C. Hayden, A. M. Johnson, and K. L. Tokuda, "Induced-grating autocorrelation of ultrashort pulses in a slowly responding medium," *J. Opt. Soc. Amer. B,* vol. 11, no. 9, pp. 1609–1618, 1994.
[25] J. T. Fourkas, L. Dhar, K. A. Nelson, and R. Trebino, "Spatially encoded, single-shot ultrafast spectroscopies," *J. Opt. Soc. Amer. B,* vol. 12, no. 1, pp. 155–165, 1995.
[26] A. E. Siegman, "Lasers," in *Lasers.* Mill Valley, CA: Univ. Science Books, 1986, ch. 9.
[27] G. P. Agrawal, *Nonlinear Fiber Optics.* San Diego: Academic, 1989.

Kenneth W. DeLong (S'73–M'95), was born in Allentown, PA, in 1962. He received the B.S. in physics from Penn State University, University Park, PA, in 1984 and the Ph.D. degree in optical sciences in 1990 from the University of Arizona, Tucson.

From 1990 to 1992, he worked at NTT Basic Research Laboratories in Tokyo, Japan, on femtosecond spectroscopy of semiconductor microcrystallites, soliton interactions, and dynamics of discrete nonlinear equations. While at Sandia National Laboratory, Livermore, CA, he was engaged in the characterization of ultrashort laser pulses. Currently, he is with Lawrence Livermore National Labs, Livermore, CA.

David N. Fittinghoff received the B.S. in physics and the M.S. and Ph.D. degrees from the University of California at Davis in 1985, 1990, and 1993, respectively.

After receiving the B.S. degree, he participated in research on the generation of high-power microwaves at Physics International. He is currently a Post-Doctoral Appointee at Sandia National Laboratories' Combustion Research Facility, Livermore, CA. His interests include ultrashort pulse lasers and their applications, ultrashort pulse measurement, optical inverse scattering, and neural networks.

Rick Trebino was born in Boston, MA, on January 18, 1954. He received the B.A. degree from Harvard University, Cambridge, MA, in 1977, and the Ph.D. degree from Stanford University, Palo Alto, CA, in 1983.

His dissertation research involved the development of a technique for the measurement of ultrafast events in the frequency domain using long-pulse lasers by creating moving gratings. He continued this research during a three-year term as a Physical Sciences Research Associate at Stanford. In 1986, he joined Standia National Laboratories, Livermore, CA, where he has studied higher order wave-mixing effects, nonlinear-optical perturbation theory using Feynman diagrams, and ultrashort-laser-pulse techniques with application to chemical dynamics measurements and combustion diagnostics. His latest work has been the development of a technique for the measurement of the intensity and phase of ultrashort laser pulses.

Dr. Trebino is a member of the Optical Society of America, the American Physical Society, and the American Association for the Advancement of Science.