

Noise sensitivity in frequency-resolved optical-gating measurements of ultrashort pulses

David N. Fittinghoff, Kenneth W. DeLong, Rick Trebino, and Celso L. Ladera*

*Combustion Research Facility, Sandia National Laboratories,
M.S. 9051, P.O. Box 969, Livermore, California 94551*

Received September 30, 1994; revised manuscript received May 15, 1995

Frequency-resolved optical gating (FROG), a technique for measuring ultrashort laser pulses, involves producing a spectrogram of the pulse and then retrieving the pulse intensity and phase with an iterative algorithm. We study how several types of noise—multiplicative, additive, and quantization—affect pulse retrieval. We define a convergence criterion and find that the algorithm converges to a reasonable pulse field, even in the presence of 10% noise. Specifically, with appropriate filtering, 1% rms retrieval error is achieved for 10% multiplicative noise, 10% additive noise, and as few as 8 bits of resolution. For additive and multiplicative noise the retrieval errors decrease roughly as the square root of the amount of noise. In addition, the background induced in the wings of the pulse by additive noise is equal to the amount of additive noise on the trace. Thus the dynamic range of the measured intensity and phase is limited by a noise floor equal to the amount of additive noise on the trace. We also find that, for best results, a region of zero intensity should surround the nonzero region of the trace. Consequently, in the presence of additive noise, baseline subtraction is important. We also find that Fourier low-pass filtering improves pulse retrieval without introducing significant distortion, especially in four-noise cases. We show that the field errors in the temporal and the spectral domains are equal. Overall, the algorithm performs well because the measured trace contains N^2 data points for a pulse that has only $2N$ degrees of freedom; FROG has built in redundancy.

© 1995 Optical Society of America

1. INTRODUCTION

Frequency-resolved optical gating (FROG) is a recently developed technique that uses a phase-retrieval algorithm to retrieve the intensity $I(t)$ and the phase $\phi(t)$ from a measured spectrogram of an ultrashort laser pulse.^{1–3} The spectrogram, or FROG trace, is a time–frequency representation⁴ of the pulse that is produced by frequency resolution of a nonlinear autocorrelation-type signal generated by two variably delayed replicas of the pulse. Any nonlinear-optical interaction can be used to make FROG measurements, and experiments and theoretical studies have been performed that used polarization gating (PG),^{5–7} second-harmonic generation (SHG),^{8,9} and self-diffraction (SD).^{1,6} Previous simulations showed that, for noise-free data, the FROG retrieval algorithm^{2,3,5,10} retrieves the correct intensity and phase for all pulses attempted, including those with complex intensity and phase structure.

In practice, however, noise is present in actual FROG traces, and several important questions remain to be answered: How well does the FROG retrieval algorithm retrieve the pulse in the presence of such noise? Does it always converge? If so, then what errors can be expected in the retrieved pulse intensity and phase versus time for a given noise level in the experiment FROG trace? Finally, what additional techniques can be incorporated to improve the retrieval?

It is particularly important to determine whether the algorithm has converged. In the absence of noise this is generally a straightforward task: either the error between the retrieved FROG trace and the measured trace goes to zero or it does not. Unfortunately, in the presence

of noise this task is not so straightforward. In general, the algorithm proceeds until it reaches a finite minimum error. Is the resulting inevitable error in the recovered pulse intensity and phase simply the error that is due to measurement error? Or is the error an indication of a lack of convergence? This distinction is important because, in the former case, the result is simply the best estimate of the pulse obtainable from the available data, whereas in the latter case the resulting pulse may bear no relation to the actual pulse and hence is *meaningless*. It is thus crucial to define the concept of convergence in the presence of noise, and we do. Using it, we find that the algorithm essentially always converges, even for complex pulses and in the presence of massive noise. Surprisingly, in most cases the algorithm converges to a more accurate FROG trace than the original trace. In other words, using only the noisy trace as input, the algorithm retrieves a pulse whose FROG trace better approximates the noise-free trace than the original noisy trace. Because the algorithm constraints the result only to the magnitude of the data and to the mathematical form of the nonlinear interaction, there is no averaging in the usual sense that forces a result that averages through the noise. This result is due instead to the oversampling and the redundancy that are naturally built into the FROG trace: although the pulse has only $2N$ degrees of freedom, the FROG trace has N^2 degrees of freedom. Thus the overwhelming majority of mathematically constructable two-dimensional images do not correspond to FROG traces of possible pulses. So the addition of noise produces a trace that does not correspond to a physically realizable pulse, and the algorithm must then find a trace that does and is also reasonably close to the input trace.

In so doing, it finds a somewhat averaged, generally more accurate, trace.

Several image-processing techniques can improve the retrieval. When we consider additive noise (dark current), which provides an effective nonzero Poisson distributed background, we find that it is necessary to subtract off the mean of such background before running the algorithm on the trace. Even constant background at large delay or frequency offsets is a problem. Nonzero background at times far from the peak of the pulse produces background on the retrieved intensity. Nonzero background at frequencies far from the pulse's central frequency produces high-frequency spikes when Fourier transformed and leads to high-frequency noise in the retrieved intensity and phase. Constant background at both large delay and frequency offsets therefore yields background with high-frequency noise in the retrieved intensity (and phase). Spurious background, whether constant or noisy, is thus extremely undesirable. In other words, the region of nonzero values in the FROG trace must form an island in a sea of zero values. We note that background problems are different from the problem of cropping the FROG trace in either the delay or the frequency direction; if the trace continues off the grid, information about the trace is missing. We assume in this study that the trace is not cropped in this manner. We use a variety of methods to ensure that any spurious background in the periphery of the trace is removed. Corner suppression, or multiplying the trace by a function to preferentially reduce the values at the edges of the trace, is also useful for suppressing the background. Another method for removing noise is simple Fourier low-pass filtering of the trace, which removes high-frequency noise without significantly affecting the pulse-intensity and -phase information present in the trace. We use an extended example involving a complex pulse with intensity substructure and a phase jump (chosen to be a serious challenge to the algorithm even in the absence of noise) and show how its retrieval is affected by the above filters and the various types and quantities of noise. We also calculate the mean intensity and phase errors induced by the filters themselves. We find that such filtering, although it is not essential for achieving convergence, is extremely helpful for noisy traces and improves the retrieval significantly in high-noise situations, and we give a guide as to how and when to use such filtering. Note that, although we concentrate on the intensity and the phase versus time throughout the paper, we find that we can retrieve the spectral intensity and phase with equal accuracy. That this must be the case can be shown by a simple Parseval's theorem argument.

Convergence of the algorithm does not necessarily mean, however, that the intensity and the phase of the retrieved pulse will be less noisy than the trace from which they have been retrieved. This is an entirely different question, one that we study in detail. To do so, we numerically simulate several types of noise for the two most commonly used FROG geometries, PG FROG and SHG FROG. We consider additive and multiplicative noise at each pixel to simulate dark current and pixel-to-pixel gain variation, respectively. We also model the effects of quantization error on the FROG retrieval algorithm's performance and hence determine the price to be paid in

accuracy by using, say, an 8-bit camera. We find that, for multiplicative noise, 10% noise in the trace results in 1% rms error in the retrieved pulse intensity and phase. (It should be kept in mind that 10% multiplicative noise corresponds to 10% noise near the peak of the pulse but to 0% in the wings of the pulse—a rms error of $\sim 1\%$ for the traces used in this study.) Additive noise, on the other hand, is a more severe problem because it distorts the zero as well as the nonzero region of the trace. However, using the aforementioned simple filtering techniques to remove such noise, we find that 10% additive noise also yields 1% rms error in the retrieved pulse intensity and phase. (Unlike 10% multiplicative noise, 10% additive noise is 10% noise everywhere in the trace; this excellent value of noise in the resulting intensity and phase is partly due to noise reduction as the result of filtering and partly due to the redundancy in the FROG trace). Finally, we find that an 8-bit video camera records FROG traces accurately enough to yield 1% rms errors in the retrieved intensity and phase. No filters were used in this last study.

This paper has the following structure. In Section 2 we discuss the basics of FROG and describe our current phase-retrieval algorithm. We should point out here that the algorithm essentially consists of a simple generalized-projections approach¹⁰; the additional methods published previously are almost inconsequential in comparison.^{2,3} Then, in Section 3, we describe the types of noise that we use and our methods for simulating the noise. In Section 4 we describe the trace preparation and image processing (Fourier low-pass filtering) that we use to improve retrieval and demonstrate the effects on the test pulse and its FROG trace. In Section 5 we define the intensity and the phase errors used to measure the quality of the retrieved pulses. In Section 6 we discuss the effects of additive and multiplicative noise on the retrieved intensity and phase and describe optimized filtering of the FROG trace. Then, in Section 7, we define a convergence criterion and describe the convergence properties of the algorithm. In Section 8 we discuss quantization noise and give examples showing the ability of the algorithm to retrieve the pulse accurately from 8-bit data. We then summarize our conclusions in Section 9.

2. FROG BASICS

Consider a laser pulse with a complex electric field $E(t)$ described by $E(t) = A(t)\exp[i\phi(t)]$, where $A(t)$ is an amplitude and $\phi(t)$ is a phase. Making a FROG measurement¹⁻³ of this pulse involves splitting the pulse into two replicas and crossing the replicas in a nonlinear medium to produce a signal field $E_{\text{sig}}(t, \tau)$. Here t is time and τ is the time delay between the arrival times of the two replicas at the medium. Measuring the spectrum of the signal field as a function of τ gives the FROG trace:

$$I^{\text{FROG}}(\omega, \tau) \propto \left| \int_{-\infty}^{\infty} E_{\text{sig}}(t, \tau) \exp(i\omega t) dt \right|^2. \quad (1)$$

The FROG trace is a spectrogram of the pulse and contains essentially all information about $E(t)$.^{2,4,11-13} The time response of the nonlinear medium must be rapid compared with the pulse length, and we consider the

response to be quasi-instantaneous. Although FROG measurements can use any quasi-instantaneous nonlinear effect, the PG^{1,5,6,14,15} and the SHG^{6,8,9} geometries are the most heavily researched geometries to date. PG FROG relies on the electronic Kerr effect to rotate the polarization of a probe pulse passing between crossed polarizers and produces a signal field of the form

$$E_{\text{sig}}(t, \tau) \propto E(t)|E(t - \tau)|^2. \quad (2)$$

SHG FROG relies on SHG and produces a signal field of the form

$$E_{\text{sig}}(t, \tau) \propto E(t)E(t - \tau). \quad (3)$$

Although much intuitive information is immediately visible in the FROG trace, we use a phase-retrieval algorithm to extract the intensity and the phase from the FROG trace. One can see that a phase-retrieval algorithm retrieves the intensity and the phase from a FROG trace by rewriting relation (1) in terms of $E_{\text{sig}}(t, \Omega)$, the Fourier transform of $E_{\text{sig}}(t, \tau)$ with respect to τ . Relation (1) then becomes precisely the two-dimensional phase-retrieval problem, a solved problem from image science that is well known to have an essentially unique solution.^{16–19} We use an iterative Fourier-transform algorithm^{17,19–22} as is commonly done in image science problems, which in this case yields the intensity and the phase of the pulse.^{2,3,10} The algorithm transforms the signal field $E_{\text{sig}}(t, \tau)$ back and forth between the t and the ω domains. In the time domain the algorithm constrains the signal field to the mathematical form of, for example, relation (2) or (3). In the frequency domain the algorithm constrains the squared magnitude of the Fourier-transformed signal field $E_{\text{sig}}(\omega, \tau)$ to the FROG trace intensity, as given by relation (1).

Our current algorithm^{3,10} is significantly more robust than the algorithm described in our original study.² It now incorporates the method of generalized projections,^{10,16} a powerful technique used for phase-retrieval problems in image science. Our original algorithm used a generalized projection for implementation of the constraint in the frequency domain but not for the constraint in the time domain and hence tended to stagnate for complex pulses. Interestingly, when it does converge, our original algorithm converges quickly, even though it is less robust than the algorithm based on generalized projections. We therefore include both methods in the computer program. We also use several other methods,^{2,3} including a full least-squares minimization in which we treat all values of the field as minimization parameters. However, in essentially all cases the generalized projections algorithm with noise as the initial guess is the only algorithmic technique necessary.

The goal of the algorithm is to find the best solution, which is defined as the pulse with the minimum FROG error,

$$\epsilon_{\text{FROG}} = \left\{ \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N [\tilde{I}^{\text{FROG}}(\omega_i, \tau_j) - I_n^{\text{FROG}}(\omega_i, \tau_j)]^2 \right\}^{1/2}, \quad (4)$$

Here \tilde{I}^{FROG} is the retrieved FROG trace (the FROG trace calculated from the retrieved intensity and phase) and

I_n^{FROG} is the experimental FROG trace including any noise. The summations are over the N frequency and N delay points in the FROG traces. The FROG error indicates the degree to which the retrieved FROG trace reproduces the experimental FROG trace. The use of this error definition is necessary because no other error is available in practice. It is also reasonable because the FROG trace essentially uniquely determines the pulse, and similar traces yield similar pulses. Thus the least-squares distance from the FROG trace is a good (the best available) measure of the distance of the retrieved pulse from the actual pulse. The computer program changes to a different technique when the error has decreased by less than 0.5% of the value 10–15 iterations before. To be certain that the error is the minimum obtainable, the program cycles through the different techniques twice, even though the program typically produces the minimum error on the first cycle at the generalized projection step. When the algorithm stops, the program outputs the pulse with the smallest FROG error as the best estimate of the pulse that generated the FROG trace.

3. SIMULATION OF NOISE IN FROG TRACES

We used five test pulses in this study, and they are shown in Fig. 1. These pulses are representative of many experimental pulses. The pulses are shown in order of increasing complexity from Pulse 1, a transform-limited Gaussian, in Fig. 1(a) to Pulse 5, a double pulse with phase distortions that include linear chirp, spectral cubic phase, and self-phase modulation, in Fig. 1(e). As an example of a FROG trace, Fig. 2 shows the PG FROG trace of the pulse shown in Fig. 1(e).

Inasmuch as the FROG trace is the measured quantity, we let all noise occur on the FROG trace, which for this research is a 64×64 array of data values representing the response of camera pixels. We consider three types of noise that are representative of experimental noise: multiplicative noise, additive noise, and quantization error. Multiplicative noise describes pixel-to-pixel signal variations that are proportional to the intensity at the pixel. Gain variation at each pixel of the CCD camera is a common example of multiplicative noise for a single-shot experiment. We simulate multiplicative noise at each pixel by defining the value of the noisy FROG trace at frequency ω_i and time delay τ_j as

$$I_n^{\text{FROG}}(\omega_i, \tau_j) = I^{\text{FROG}}(\omega_i, \tau_j)(1 + m_{ij}\alpha). \quad (5)$$

Here m_{ij} is a pseudorandom number drawn from a zero-mean unit-variance Gaussian distribution and α is the noise fraction. The noise fraction is a convenient measure of the amount of noise, and we use it as such throughout this paper. We normalize the maximum value of the noiseless FROG trace to 1. It should be remembered, however, that the rms noise in the trace contaminated by multiplicative noise is considerably less than α , generally $\sim \alpha/10$ for traces in this paper.

Additive noise describes pixel-to-pixel signal variations independent of the FROG intensity at the pixels. A common example of additive noise is thermal noise that occurs in the CCD cameras typically used for single-shot FROG

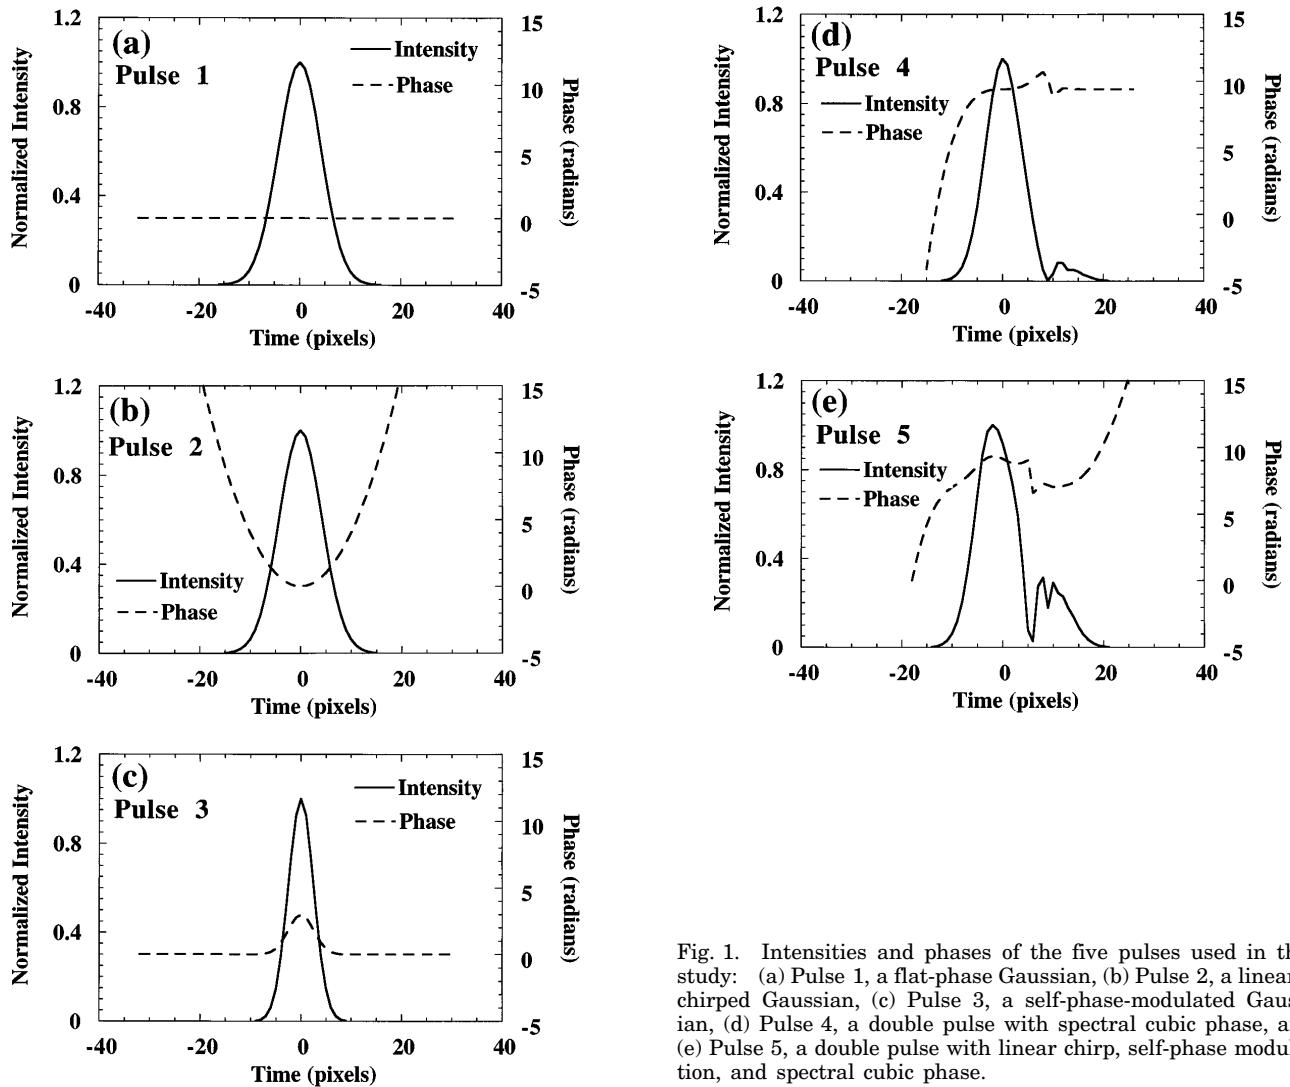


Fig. 1. Intensities and phases of the five pulses used in this study: (a) Pulse 1, a flat-phase Gaussian, (b) Pulse 2, a linearly chirped Gaussian, (c) Pulse 3, a self-phase-modulated Gaussian, (d) Pulse 4, a double pulse with spectral cubic phase, and (e) Pulse 5, a double pulse with linear chirp, self-phase modulation, and spectral cubic phase.

measurements. We simulate additive noise at each pixel by defining

$$I_n^{\text{FROG}}(\omega_i, \tau_j) = I^{\text{FROG}}(\omega_i, \tau_j) + \eta_{ij}(\alpha/n). \quad (6)$$

Here η_{ij} is a pseudorandom number drawn from a Poisson distribution of mean n . We use a Poisson distribution for additive noise because this distribution probably best describes thermal noise. We have also performed a few example simulations for a uniform noise distribution and a half-Gaussian noise distribution. These distributions gave results similar to those for the Poisson distribution, indicating that our choice of a Poisson distribution is sufficiently general. For additive noise, unlike for multiplicative noise, the rms noise in a trace contaminated by such noise is approximately α .

In this study we ignore large-scale systematic errors, such as a wash of stray light across a camera image. Large-scale systematic error is inherently a problem for the algorithm (as it necessarily is in any scientific measurement technique) because it represents a deviation from a physically realizable FROG trace. Our observations and simulations show that such noise rapidly degrades the performance of the algorithm, and removal of any large-scale systematic error is necessary to en-

sure accurate retrieval, as it must be in any method. Many types of such noise can be removed easily. For example, one can subtract off stray light from the probe pulse, biases in dark current, and incoherent polarizer leakage (in PG FROG) by blocking the gate pulse and recording the background levels, provided that these levels are repeatable from shot to shot. Incoherent polarizer leakage yields a delay-independent pulse-spectrum-shaped baseline across the trace. One can remove it by sampling the few columns of data at the maximum and minimum delays, averaging these values, and subtracting the computed spectrum from the trace for all values of the delay. Coherent polarizer leakage yields fringes that cannot be simply subtracted. Thus, for PG FROG, using polarizers with high extinction ratios ($\geq 10^5$) is desirable.

4. TRACE PREPARATION: BACKGROUND SUBTRACTION AND IMAGE PROCESSING

We find that some preprocessing of the experimental trace is quite helpful, and in this section we describe background subtraction, multiplication of the trace by a function that is unity near the center of the trace but that falls to zero at the edges, typically a super-Gaussian (what we call corner suppression), and Fourier low-pass filtering.

We have found that it is important that any trace used as input to the FROG algorithm have regions of zero signal intensity around the perimeter. Significant nonzero background at large delay or frequency offset tends to cause algorithm stagnation. There are probably several reasons for this. First, even if such background were physically valid, its failing to fall to zero would indicate a truncation of the trace. Without the full pulse data the algorithm cannot be expected to reproduce the correct pulse accurately. Also, in image science problems, phase-retrieval algorithms are well known to operate best under conditions of finite support, that is, finite extent of the image. When it is not valid, background at large time delays in the FROG trace leads to incorrect nonzero background intensity, and background at large frequency offsets in the FROG trace leads to incorrect high-frequency noise fluctuations in the recovered intensity and phase. Thus a constant background everywhere in a FROG trace leads to noisy background everywhere in the retrieved pulse. Unfortunately, noise in the perimeter acts as background. As a result, the methods that we




Fig. 2. Polarization-gate FROG trace of Pulse 5 [shown in Fig. 1(e)].

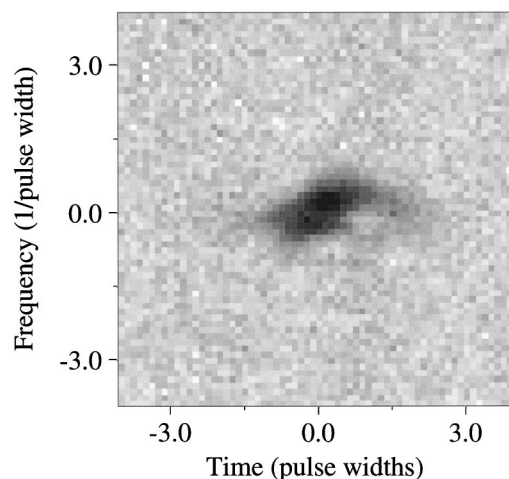


Fig. 3. PG FROG trace of Pulse 5 after inclusion of 10% additive noise. The noise is Poisson distributed, with $n = 5$. No image processing has been used on the trace. The noise at large time delays leads to nonzero background intensity, and the noise at large frequency offsets leads to high-frequency fluctuations in the recovered intensity and phase.

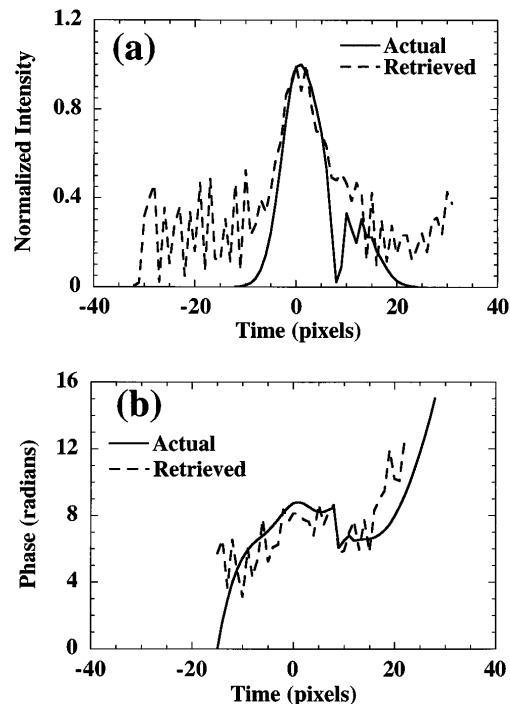


Fig. 4. Retrieved pulse for the FROG trace of Fig. 3 without mean subtraction. The retrieved intensity exhibits a large background intensity, and the secondary peak is unresolved. Both the retrieved intensity and phase exhibit high-frequency fluctuations. (a) The actual and the retrieved intensities. The rms intensity error defined in Eq. (8) is 15%. (b) The actual and the retrieved phases. The rms phase error defined in Eq. (9) is 0.65 rad.

describe in this section are aimed mainly at suppressing noise in the perimeter of the trace, although one method (Fourier low-pass filtering) also addresses the problem of noise in the central region of the trace.

For the purposes of presentation we consider the test pulse of Fig. 1(e), whose trace is shown in Fig. 2, and show how these methods improve pulse retrieval. Figure 3 shows the FROG trace from Fig. 2 with additive Poisson-distributed noise included. The noise fraction shown here is quite large: 0.1 (i.e., 10% of peak FROG signal). The mean of the Poisson distribution, n , was chosen to be 5 counts, and the peak of the trace corresponds to 50 counts. The noise does not have a zero mean, and this introduces an effective background offset with a magnitude approximately equal to the mean of the noise. Comparison of Figs. 2 and 3 indicates that the structure in the surrounding regions of the trace in Fig. 2 is now buried in the noise in Fig. 3, and even some of the main components of the trace are difficult to resolve, which raises serious doubts that the algorithm can reasonably be expected to reproduce the pulse. This is a particularly challenging case, which we chose specifically for this reason.

We expect that traces such as that in Fig. 3 will yield less than satisfying results if the algorithm operates on the trace as is. As an illustration of the expected results of such a trace, see Fig. 4, which shows the actual intensity and phase and the retrieved phase to which the algorithm has converged for this trace. Note that, surprisingly, the algorithm does roughly retrieve the major features of the intensity and the phase. The pulse length is about right, and the general shape of the phase

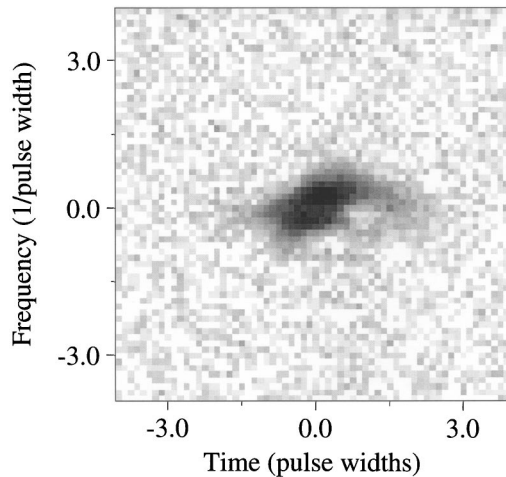


Fig. 5. PG FROG trace of the test pulse with 10% additive noise in the trace after subtraction of the mean of the noise. Subtracting the mean of the noise lowers the unphysical values at large time delays and frequencies, which is crucial for accurate pulse retrieval.

is correct. Quite surprisingly, the retrieved phase contains the phase jump near $t = 5$, but there is sufficient noise elsewhere in the phase that it is difficult to be confident in the jump. Worse, the secondary intensity peak cannot be resolved, and there is a large background intensity with high-frequency fluctuations, as expected because of the large background in the trace.

We obtain significantly better algorithm performance in all cases, and especially in this one, simply by subtracting the background off. In particular, before running the pulse-retrieval algorithm we can process the trace by subtracting off the mean of the noise. Any negative points that result from the subtraction are set to zero. We estimate the mean by averaging the data in the 8×8 pixel squares in the corners of the FROG trace. In practice, other methods might be preferred.²³

Figure 5 shows the resulting FROG trace after this simple procedure for our example trace. Figure 6 shows the retrieved intensity and phase for the FROG trace of Fig. 5. Note that the algorithm now resolves the secondary intensity peak, and the amplitude of the high-frequency fluctuating background has been reduced by nearly an order of magnitude. The phase jump near $t = 5$ is still accurately retrieved, and the remainder of the phase behavior is much more accurately obtained, which yields more confidence in the phase jump. Significant phase deviations remain only at very low intensity values. The improvement is significant and quite impressive, given the amount of noise present in the original trace. Nevertheless, even after mean subtraction the tail of the Poisson distribution of the noise causes pixels at the edges of the trace to have nonnegligible values. This leads to the residual noise in the wings of the pulse seen in Fig. 6.

As a result, because the FROG trace is essentially a two-dimensional image, we also considered several image-processing techniques to improve the retrieval, some of which did not help. We tried—and rejected—out-range pixel smoothing and median filtering.²⁴ These methods involve replacing the actual pixel values with the average or median value of the surrounding pixels.

This had the effect of broadening the traces in both the time and the frequency dimensions, thus unphysically increasing the apparent pulse time–bandwidth products. As these methods did not introduce additional features, such as chirp and structure, which would be required for the increased time–bandwidth products, the traces no longer resembled those of physically valid pulses. This proved to confuse the algorithm sufficiently that we never obtained a clear improvement in the retrieved intensities and phases.

For additive noise we found that retrieval could be improved by what we call corner suppression, that is, simply multiplying the FROG trace by a radially symmetric super-Gaussian of the form

$$G(\omega_i, \tau_j) = \exp\{-\beta[(\omega_i - N/2)^2 + (\tau_j - N/2)^2]/d^4\}. \quad (7)$$

Here d is the full width at half-maximum of the super-Gaussian and $\beta = 16 \ln(2)$. Figure 7 shows the FROG trace of Fig. 3 after mean subtraction and corner suppression with $d = 45$. The corner suppression reduces the values and also the noise at the edges, and especially at the corners, of the trace without significantly distorting the nonzero regions of the trace. Corner suppression is reasonable because the trace ought to be zero in these regions in the first place.

Figure 8 shows the retrieved intensity and phase from the trace in Fig. 7. Note the significant improvement in this retrieval over that in Figs. 4 and 6, especially in the wings of the intensity. After mean subtraction and corner suppression, the algorithm is able to extract the pulse intensity and phase remarkably well, considering the initial noise.

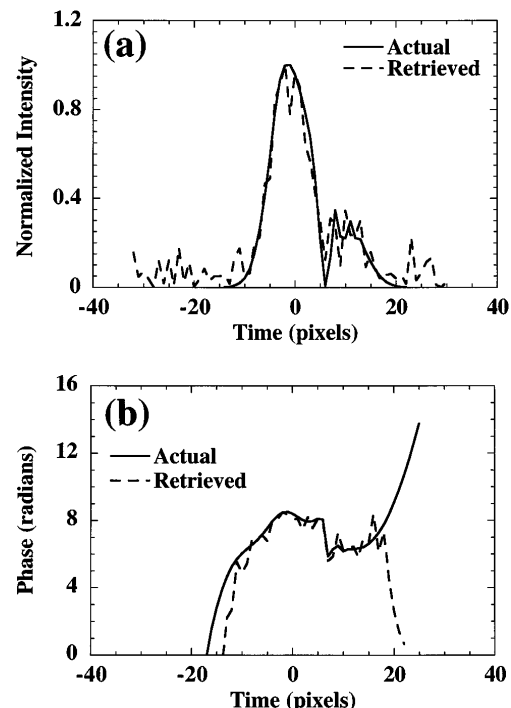


Fig. 6. Retrieved pulse for the FROG trace of Fig. 5. Subtracting the mean greatly reduces the background intensity and high-frequency fluctuations. (a) The actual and the retrieved intensities. (b) The actual and the retrieved phases.




Fig. 7. PG FROG trace of the test pulse with 10% additive noise after subtraction of the mean of the noise and corner suppression with $d = 45$ pixels. Super-Gaussian corner suppression forces the values at the perimeter, and especially at the corners, to zero.

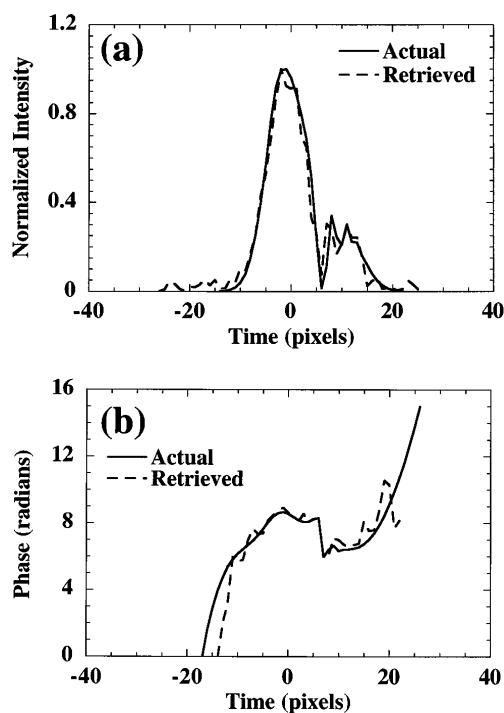


Fig. 8. Retrieved pulse for the FROG trace of Fig. 7. Note that the use of corner suppression lowers the background intensity in the retrieved pulse. Also, the reduction in the high-frequency noise allows the algorithm to resolve the two peaks of the pulse cleanly. (a) The actual and the retrieved intensities. (b) The actual and the retrieved phases.

It should be remembered that corner suppression, although it is extremely useful for high-noise cases, can distort the trace somewhat, and the errors produced by the filter limit the accuracy of the retrieval. Thus corner suppression is quite useful for noisy traces when some pulse distortion can be tolerated, but it is of less utility for clean traces. It is also more useful for traces with additive noise than for those with multiplicative noise, in which latter case corner suppression is generally unnecessary.

Of all the image-processing techniques that we have considered, Fourier low-pass filtering of the FROG trace has proved to be the most useful. We implement the low-pass filter by two-dimensional Fourier transforming the FROG trace, multiplying the transformed trace by a top-hat function of radius, in pixels, $\rho N/2$ (setting all values outside $\rho N/2$ equal to zero), and transforming the result back from the transform space to the image space. Note that when $\rho = 1$ the top-hat function has a diameter of N pixels and excludes only the corners of the transformed trace. The effect of low-pass filtering on the FROG trace is to remove the higher spatial frequencies and smooth through the noise but without significantly broadening the trace. $\rho = 1$ corresponds to suppression of one half of the fluctuations between adjacent points. Values of ρ less than 1 yield suppression of lower-spatial-frequency noise in the trace. One should be careful, when using values of ρ less than 1, not to filter out real spatial frequencies in the trace corresponding to actual pulse fluctuations. Because the noise that we are simulating (and that usually occurs) is generally at higher spatial frequencies than are contained in the FROG trace, it is generally possible to improve the retrieval by low-pass filtering with an appropriate choice of ρ .

Figure 9 shows the smoothing that occurs on the FROG trace of Fig. 7 when a low-pass filter with $\rho = 0.5$ (in addition to background subtraction and corner suppression with $d = 45$) is used. Figure 10 shows the improvement in the retrieval provided by the filtering. In particular, the retrieval of the phase is very good, except for slight ringing induced by the low-pass filter at the sharp phase change near $t = 5$. Also, the noise-induced prepulses evident in Fig. 8(a) are reduced. Note that the retrieved pulse phase disagrees significantly with the actual phase only in regions of very low intensity, where the phase is not well defined.

The choice of how tightly to low-pass filter, that is, the choice of the value of filter radius ρ , depends on the quantity of noise and the type of noise present. Like all other filtering operations, low-pass filtering distorts the trace and introduces errors of its own. We find that low-pass




Fig. 9. PG FROG trace of the test pulse with 10% additive noise after subtraction of the mean of the noise, corner suppression with $d = 45$ pixels, and low-pass filtering with $\rho = 0.5$. The low-pass filtering removes the high-spatial-frequency noise from the entire trace. The resulting smoothing effect is apparent.




Fig. 10. Retrieved pulse for the FROG trace of Fig. 9. The retrieved intensity is very good, clearly reproducing the peaks of the pulse with minimal background or high-frequency noise. The phase is also remarkably good for such an initially high amount of noise and does not strongly deviate from the actual phase, except, as expected, at times when the intensity is below 1% of the peak intensity. (a) The actual and the retrieved intensities. The rms intensity error is 4%. (b) The actual and the retrieved phases. The rms phase error is 0.14%.

filtering, while critical for traces with large amounts of noise, is not so useful for low levels of noise, which yield much better results anyway. To give an idea of the magnitude of the distortion involved, we have applied low-pass filtering to a noise-free trace. Figure 11 shows a plot of the mean intensity and phase errors (as defined in Section 5) as a function of filter radius induced by low-pass filtering noise-free PG FROG traces of the five test pulses. Note that, as expected, as the radius of the filter decreases, the errors introduced by the low-pass filter increase. When minimal noise is present in the trace, distortions induced by the low-pass filter are quite noticeable and hence undesirable. On the other hand, when large amounts of noise are present in the trace the error introduced by filtering may be negligible compared with the noise reduction obtained by its use. Thus the decision to use the low-pass filter will in practice depend on the type and the amount of noise. The filter is essential in high-noise cases but detrimental in low-noise cases, and it should be used accordingly.

A similar noise-suppression distortion trade-off exists for corner suppression. Figure 12 shows a plot of the mean intensity and phase errors induced as a function of the corner-suppression filter diameter. Just as for low-pass filtering, as the radius of the filter decreases, the errors introduced increase. In this case, however, the errors are nearly constant and moderate in magnitude, even for relatively large diameters such as $d = 60$.

In view of these results we will define the concept of optimized filtering, which will involve heavy filtering in high-noise cases and weaker filtering in low-noise cases,

where it is less necessary and our standards for pulse distortion are higher. We will discuss the specific details of optimized filtering in Section 6 as we discuss the effects of noise on the retrieval of the intensity and phase.

Before going to Section 5 we note that additional, more-sophisticated image-processing techniques, such as Wiener filtering²⁴ and, in particular, wavelet noise reduction,²⁵ may also prove to be useful for removing noise from FROG traces.

5. INTENSITY AND PHASE ERRORS

The performance of the algorithm can be quantitatively measured in two ways: how well it retrieves the original intensity and phase and how well it retrieves the original FROG trace. In experiments in which the original intensity and phase are unknown, only the latter error is available. As a result, it plays a central role in the algorithm, as discussed above. In this section we will also define intensity and phase errors to give us an idea of how well the algorithm retrieves the pulse itself, rather than only its FROG trace, in simulations. All these measures are related because of the one-to-one correspondence of

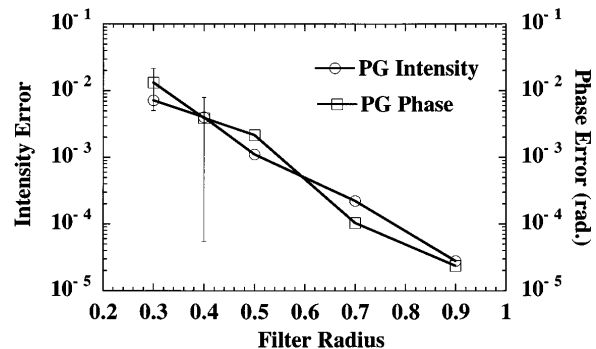


Fig. 11. Mean rms intensity and phase errors for PG FROG induced by low-pass filtering the FROG trace for the five test pulses. The data are plotted as a function of filter radius. The filter radius ρ in the transform space is given as a fraction of $N/2$, where N is the number of pixels in each of the two dimensions of the FROG trace. Error bars indicating one standard deviation from the mean error for the intensity and the phase are shown at filter radii of 0.4 and 0.3, respectively.

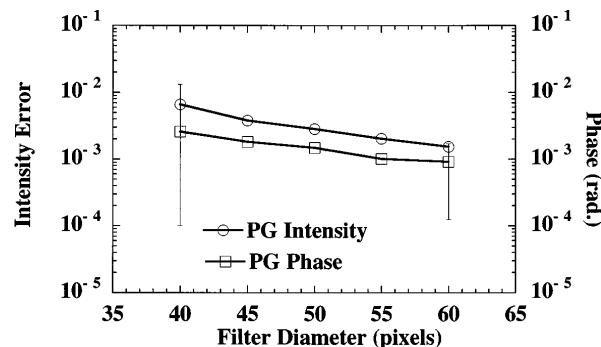


Fig. 12. Mean rms intensity and phase errors for PG FROG induced by the use of corner suppression on the FROG trace for the five test pulses. The data are plotted as a function of filter radius. The filter is centered in the center of the FROG trace and has a diameter given in pixels in each of the two dimensions of the FROG trace. Error bars indicating one standard deviation from the mean error for the intensity and the phase are shown at filter diameters of 40 and 60, respectively.

pulses to possible FROG traces, but it is not *a priori* obvious that rms intensity and phase errors correspond to a given rms FROG trace error. We find that the former errors are of the order of the latter error.

To measure how closely the algorithm retrieves the original intensity we use the rms intensity error in the obvious manner:

$$\epsilon_I = \left\{ \frac{1}{N} \sum_{j=1}^N [\tilde{I}(t_j) - I(t_j)]^2 \right\}^{1/2}, \quad (8)$$

where I and \tilde{I} are the actual and the retrieved intensities, respectively. Here the peak of $I(t)$ is normalized to 1, so we often quote this error as a percentage. The summation is over the N discrete time points.

The choice of error to use as a monitor of the algorithm's ability to retrieve the phase is not so obvious. An unweighted rms phase error is inappropriate because the phase of the pulse is meaningless when the intensity is near zero. The retrieved phase can thus exhibit large—but meaningless—phase variations for low intensities that skew the rms phase error. Therefore we would not want to include such phase points in our error calculation. A more appropriate choice of the error to use to monitor the retrieval of the phase is the intensity-weighted rms phase error:

$$\epsilon_\phi = \left\{ \frac{1}{N} \sum_{j=1}^N I^2(t_j) [\tilde{\phi}(t_j) - \phi(t_j)]^2 \right\}^{1/2} / \left[\frac{1}{N} \sum_{j=1}^N I^2(t_j) \right]^{1/2}. \quad (9)$$

Here ϕ and $\tilde{\phi}$ are the actual and the retrieved phases, respectively. Note that ϵ_ϕ has units of radians.

When calculating both of these errors we minimize the errors with respect to the peak intensity, a temporal shift, and the absolute phase to account for the fact that the FROG trace is invariant to overall shifts in time and phase in the pulse and that the algorithm yields a pulse of no particular peak intensity. As an example of the relative magnitude of the errors, the intensity and the phase errors obtained for Figs. 4(a) and 4(b) were 15% and 0.65 rad, respectively, whereas the intensity and the phase errors obtained for Figs. 10(a) and 10(b) were only 4% and 0.14 rad, respectively.

We could also define and calculate the errors in the spectra intensity and phase rather than the temporal quantities. A simple argument, however, shows that the spectral errors contain no additional information. Consider the error in the field as a function of time, ϵ_t , defined by

$$\epsilon_t^2 = \sum_{j=1}^N |\tilde{E}(t_j) - E(t_j)|^2 / \sum_{j=1}^N |E(t_j)|^2. \quad (10)$$

Here $\tilde{E}(t)$ is the reconstructed field at time t , and the summation is over N discrete times. The error in the field as a function of frequency may be similarly defined by

$$\epsilon_\omega^2 = \sum_{j=1}^N |\tilde{E}(\omega_j) - E(\omega_j)|^2 / \sum_{j=1}^N |E(\omega_j)|^2. \quad (11)$$

Here $E(\omega)$ and $\tilde{E}(\omega)$ are the actual and the reconstructed fields, respectively, as a function of frequency ω , and the summation is over the N discrete frequencies. By Parseval's theorem, $\epsilon_t^2 = \epsilon_\omega^2$. Thus, as the error in the

field versus time equals the error in the field versus frequency, the errors in the spectral intensity and phase are similar to those in the temporal intensity and phase.

6. EFFECTS OF ADDITIVE AND MULTIPLICATIVE NOISE ON RETRIEVAL

We now consider the behavior of the intensity and the phase errors in the presence of additive and multiplicative noise. We use the same noise (i.e., set of η_{ij}) in all cases; only its magnitude α changes. If we use other sets of η_{ij} we obtain results similar to those presented here. For additive noise Fig. 13 shows the mean PG and SHG FROG retrieval errors for the five pulses versus noise fraction, α . The mean of the noise background was subtracted from the FROG traces before retrieval, but no other filtering was performed. In general, filtering is strongly advised; this initial study without filtering is mainly for the sake of illustration. The error bars indicate one standard deviation from the mean for the five pulses. The errors for PG and SHG FROG are comparable and decrease roughly as the square root of the noise fraction. For 0.1% additive noise ($\alpha = 10^{-3}$) the algorithm achieves intensity and phase errors of the order of 1% and 0.01 rad, respectively. For 10% additive noise the algorithm achieves intensity and phase errors of the order of 10% and 0.1 rad, respectively.

It is interesting to consider where in the pulse the noise resides. We find that it occurs mainly where the intensity is highest and is much lower in the wings of the pulse. Figure 14 shows, plotted on a log scale, the actual intensity for Pulse 1 and the retrieved intensities for additive noise with $\alpha = 0.001$ and $\alpha = 0.0001$ for PG FROG without filtering. The rms intensity errors for these noise levels are 0.0071 and 0.00145, respectively, but the background levels or noise floors of the retrieved intensities are much lower: roughly equal to the amount of noise. This result holds for all the pulses and the noise levels and shows that FROG measurements can be made with a dynamic range limited only by (and equal to) the

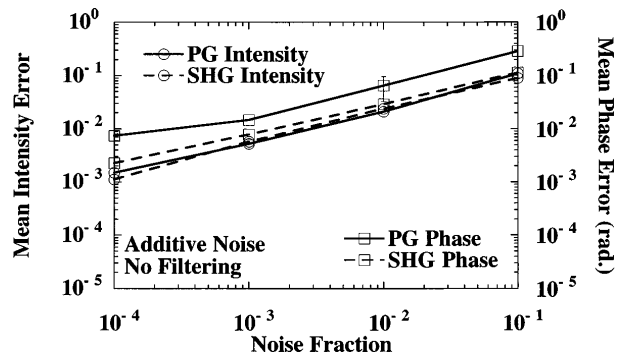


Fig. 13. Mean rms intensity and phase errors for PG and SHG FROG for additive noise with $n = 5$. The mean of the noise background was subtracted from the FROG traces before retrieval, but no other filtering has been performed. The data are plotted as a function of noise fraction. The errors decrease roughly as the square root of the noise. An error bar indicating one standard deviation from the mean phase error for PG FROG is shown at a noise fraction of 10^{-2} . The standard deviation for the PG intensity at a noise fraction of 10^{-3} is contained within the circular marker. These high noise levels in the retrieved intensity and phase reveal the importance of filtering (see Fig. 15 below).




Fig. 14. Actual and retrieved intensities of Pulse 1 for additive noise with $\alpha = 0.001$ and $\alpha = 0.0001$ for PG FROG. The noise floor in the wings of the retrieved intensities is roughly equal to, and decreases linearly with, the amount of additive noise, α .

level of noise. FROG thus effectively finds weak satellite pulses, but it is less adept at revealing slight distortions in the high-intensity regions of the pulse.

Significantly better results can be obtained by use of what we call optimized filtering, which involves selective corner suppression and low-pass filtering, mostly in cases when noise is large and some pulse distortion can be tolerated as the price to be paid for noise reduction. Specifically, in view of the results of Figs. 11 and 12, we define optimized filtering for additive noise to include low-pass filtering and corner suppression with filter radii chosen such that the mean induced errors that are due to the filters themselves are at the level of error desired. The optimized filter thus uses smaller filter radii in high-noise cases, where some distortion can be tolerated, and larger radii in low-noise cases, where such distortions will be noticed. In this study we chose the desired errors to be an order of magnitude lower than the mean errors obtained without filtering that are shown in Fig. 13. Thus the optimized corner-suppression radii are $d = 60, 41,$ and 35 for additive noise with $\alpha = 10^{-3}, 10^{-2},$ and 10^{-1} , respectively. For $\alpha = 10^{-4}$, corner suppression is not used. The optimized low-pass filter radii for additive noise are $\rho = 0.58, 0.58, 0.4,$ and 0.3 for $\alpha = 10^{-4}, 10^{-3}, 10^{-2},$ and 10^{-1} , respectively.

Figure 15 shows the results of optimized filtering of the additive-noise-contaminated traces before the algorithm for PG FROG was run. Note the significant improvement obtained, which is consistent with the extended example of Section 4. Now, for 10% additive noise, the rms intensity and phase errors are only of the order of 1% and 0.01 rad, respectively. The results for SHG FROG are similar. It is thus clear that filtering should play a significant role in the pulse-retrieval process in FROG in the presence of additive noise. As for the unfiltered case, the noise resides mainly where the intensity is highest and is lower in the wings of the pulse. The noise floors of the retrieved intensity for filtered traces are roughly equal to the noise that remains in the wings of the trace after filtering.

Figure 16 shows the PG and the SHG FROG retrieval errors versus noise fraction for multiplicative noise. Multiplicative noise as high as 10% produces errors of 1%, indicating that, of course, the algorithm deals with multiplicative noise better than it deals with additive noise. This is not surprising because multiplicative noise

is significant only where the signal intensity is large, so multiplicative noise leaves the perimeter of the trace unchanged. In addition, the average error in the trace for a given value of α is less for multiplicative noise than for additive noise. We find that, as happens for additive noise, the errors decrease with decreasing noise fraction. We used optimized filtering for multiplicative noise cases but observed no significant improvement in the reconstructions. The errors without the filters are much smaller than for additive noise and require such large filter radii (to prevent inducing even larger errors with the filters) that the filters prove unnecessary where only multiplicative noise exists.

7. CONVERGENCE CRITERION

As mentioned above, it is important to define a convergence condition to determine whether error in the retrieved pulse is simply the noise resulting from the measurement or instead is an indication of algorithm stagnation. In the former case the retrieved pulse is meaningful, whereas in the latter it is not.

For this study, in which we know the actual pulse, we could simply consider how well the algorithm retrieves the original pulse intensity and phase. However, it is not

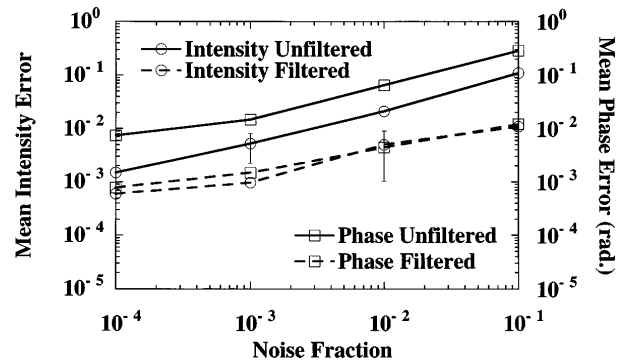


Fig. 15. Intensity and phase errors for PG FROG for additive noise with $n = 5$ with and without optimized filtering. Filtering lowers the retrieved errors dramatically. Error bars indicating one standard deviation from the mean intensity error for the filtered and the unfiltered cases are shown at noise fractions of 10^{-2} and 10^{-3} , respectively.

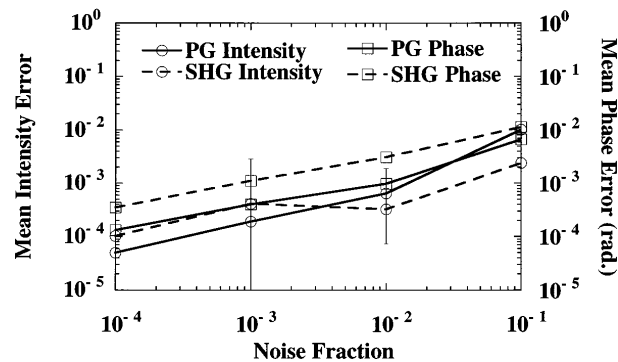


Fig. 16. Intensity and phase errors for PG and SHG FROG for multiplicative noise. The traces were not filtered before the pulse was retrieved. The data are plotted as a function of noise fraction. Error bars indicating one standard deviation from the mean phase error for PG and SHG FROG are shown at noise fractions of 10^{-2} and 10^{-3} , respectively.

clear how to define a quantitative criterion because noise in the FROG trace has different units from that in the phase, for example. And noise in the intensity is not related to noise in the FROG trace either, even though they have the same units, because the two quantities live in different functional spaces. Instead, we must attack this issue by asking how well the algorithm retrieves the desired (i.e., noise-free) FROG trace, which is directly related to the noise in the noisy FROG trace. This is reasonable in view of the uniqueness of FROG traces.

Thus we consider the algorithm to have converged to the original trace if the error between the retrieved FROG trace and the actual noise-free FROG trace is at most only slightly larger than the error between the noisy FROG trace and the actual FROG trace. Formally, we consider convergence to have occurred when

$$R = \frac{\epsilon[\tilde{I}^{\text{FROG}}(\omega_i, \tau_j), I^{\text{FROG}}(\omega_i, \tau_j)]}{\epsilon[I_n^{\text{FROG}}(\omega_i, \tau_j), I^{\text{FROG}}(\omega_i, \tau_j)]} < 2, \quad (12)$$

where

$$\epsilon[C, D] = \left[\frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N (C_{ij} - D_{ij})^2 \right]^{1/2} / \left[\frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N D_{ij}^2 \right]^{1/2} \quad (13)$$

is the rms error of FROG trace C with respect to FROG trace D. If R is less than 2 we consider the retrieved trace to be a reasonable representation of the original trace. We use 2, and not 1, because the noisy trace is in some sense displaced in function space from the original noise-free trace, and the retrieved trace is then expected to be displaced somewhat from this trace still. So a number greater than 1 seems appropriate, and 2 seems reasonable. If, as often happens, R is less than 1, the retrieved FROG trace is a *better* representation of the correct FROG trace than the noisy trace used as input to the algorithm. The algorithm has in some sense ignored the noise and found a solution that is better than the input! This apparent ability of the algorithm to separate the noise from the true FROG trace occurs because the noise does not satisfy the mathematical constraint imposed by the nonlinear interaction on the signal field [relation (2) or (3)] and because the FROG trace has built-in redundancy. The FROG trace has N^2 data points, but only $2N$ data points are required for a description of the intensity and the phase.

Figure 17 shows the mean convergence criterion versus noise fraction for additive and multiplicative noise for PG FROG and SHG FROG. We used no filters for these data. In all cases R is significantly less than 2, and we conclude that the algorithm converges. For the individual pulses the algorithm gives convergence ratios below 1, and in some cases the convergence ratios are lower even than 0.1. As stated above, these convergence ratios indicate that the algorithm smooths through the noise to produce a representation of the pulse that is more accurate than the original trace with noise. In addition, the results for optimized filtering are numerically similar and the algorithm converges.

We should mention that initially we found that, for a few very low-noise cases ($\alpha \approx 10^{-4}$), the algorithm gave a value of R slightly above 2. Because the noise was so low in these cases, the intensity and phase errors for these cases were typically still of the order of 0.1%. We found that for such cases the algorithm could get trapped in a local minimum very close to the global minimum. We have found that simply altering the noise used for the initial guess for the intensity and the phase avoids the local minimum and allows convergence for cases that give such unexpectedly high FROG errors as the result of a local minimum. The best initial guesses for the intensity

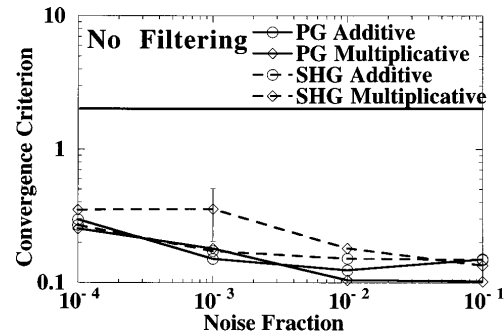


Fig. 17. Convergence ratio for PG and SHG FROG for additive and multiplicative noise without filtering. For additive noise $n = 5$. The solid horizontal line indicates the convergence limit where $R = 2$. The algorithm converges for all cases. Error bars indicating one standard deviation from the mean convergence ratio for multiplicative noise for PG and SHG FROG are shown at noise fractions of 10^{-2} and 10^{-3} , respectively. The upper error bar for PG FROG is contained within the diamond marker.

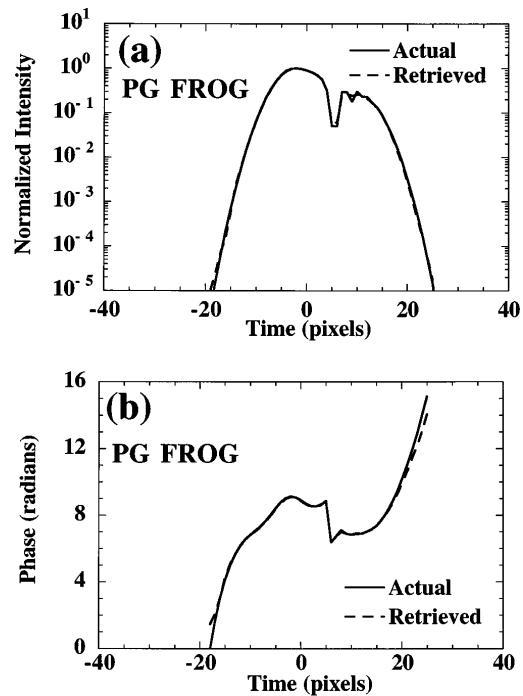


Fig. 18. Retrieved test pulse for PG FROG after quantization to 8 bits. The errors in the retrieved intensity and phase are small and become apparent mainly at intensities below 10^{-3} of the peak intensity. (a) The actual and the retrieved intensities. The rms intensity error is 1.1%. (b) The actual and the retrieved phases. The phase error is 0.0044 rad.

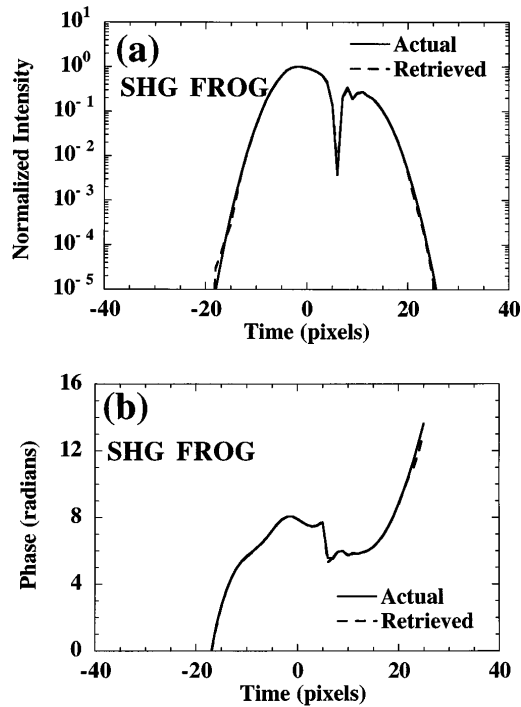


Fig. 19. Retrieved test pulse for SHG FROG after quantization to 8 bits. The errors in the retrieved intensity and phase are small and become apparent mainly at intensities below 10^{-3} of the peak intensity. (a) The actual and the retrieved intensities. The rms intensity error is 0.63%. (b) The actual and the retrieved phases. The phase error is 0.00216 rad.

and phase were noise uniformly distributed from 0 to 1. We used this technique for avoiding local minima in the above results; it was necessary mainly for the most complex pulse with the phase jump and only for such very small amounts of noise. For pulses with less structure or higher levels of noise, the algorithm always converged for the first choice of initial-guess noise.

8. QUANTIZATION NOISE

Digitizing a FROG trace, as is inevitable when one is transferring the data to a digital computer, leads to inaccuracies in retrieval. The data assume discrete values and no longer represent the pulse accurately. A standard camera has 8 bits of resolution, and much more expensive cameras can have 16 bits of resolution. We focus here on 8-bit digitization. For PG and SHG FROG, Figs. 18 and 19 show the retrieved intensity and phase for Pulse 5 after the FROG traces were reduced to 8 bits of resolution. For PG FROG the intensity and the phase errors are 1.1% and 0.0044 rad, respectively, and for SHG FROG the intensity and phase errors are 0.63% and 0.00216 rad, respectively. In both cases the deviations of the retrieval become apparent only for intensities near 10^{-3} of the peak intensity. This is not surprising, because the FROG trace has 256 levels. These traces are representative of all our quantization results, and thus we find that a standard 8-bit camera is probably sufficient for measurements with $\sim 10^2$ dynamic range in the intensity of the FROG signal. Switching to a 14-bit camera will increase the number of quantization levels to 16384 and will give a dynamic range up to $\sim 10^4$. Making mea-

surements with dynamic ranges of 10^5 or more will require separate calibrations of a FROG device for different intensity regimes. We should mention that no filtering was used in this analysis, but it seems reasonable that an optimized filter would yield some improvement (although not so much as in the additive-noise case).

9. CONCLUSIONS

We numerically simulated up to 10% additive, multiplicative, and quantization noise for polarization-gate and second-harmonic-generation frequency-resolved optical gating. We find that our phase-retrieval algorithm is capable of retrieving pulses even with the largest values of noise that we considered. The errors in the intensity and the phase for PG and SHG FROG are comparable and decrease roughly as the square root of the amount of noise. We find that optimal preparation of the trace for input into the algorithm includes requiring the nonzero region of the trace to be surrounded by a region of zeros and elimination of nonrandom backgrounds. For traces with additive noise we find also that mean-background subtraction and optimized application of super-Gaussian corner suppression and low-pass filters greatly improves retrieval of the intensity and the phase. Using these techniques, we obtained intensity and phase errors at most of order 1% and 0.01 rad, respectively, for 10% additive and multiplicative noise. In addition, we showed that the spectral and the temporal field errors are equal, so the errors in the spectrum and the spectral phase will be similar to the errors in the intensity and the phase. We defined a convergence criterion and showed that in general the algorithm converges to a solution that is a better representation of the original than the noisy input trace. Finally, we showed that, in the presence of quantization noise and when the trace has a finite number of discrete levels, the dynamic range of the retrieved pulse is close to that of the FROG trace itself. Therefore an 8-bit camera is sufficient for retrieving pulses with a 10^2 dynamic range.

ACKNOWLEDGMENTS

We acknowledge the support of the U.S. Department of Energy, Office of Basic Energy Sciences, Chemical Sciences Division, and the Laboratory-Directed Research and Development Program of Sandia National Laboratories.

*Permanent address, Department de Fisica, Universidad Simon Bolivar, Apartado 89000, Caracas 1080A Venezuela.

REFERENCES AND NOTES

1. D. J. Kane and R. Trebino, "Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating," *IEEE J. Quantum Electron.* **29**, 571–579 (1993).
2. 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. Am. A* **10**, 1101–1111 (1993).
3. K. W. DeLong and R. Trebino, "Improved ultrashort-pulse retrieval algorithm for frequency-resolved optical gating," *J. Opt. Soc. Am. A* **11**, 2429–2437 (1994).
4. L. Cohen, "Time-frequency distributions—a review," *Proc. IEEE* **77**, 941–981 (1989).

5. 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.* **18**, 823–825 (1993).
6. K. W. DeLong, R. Trebino, and D. J. Kane, "Comparison of ultrashort-pulse frequency-resolved-optical-gating traces for three common beam geometries," *J. Opt. Soc. Am. B* **11**, 1595–1608 (1994).
7. B. 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.* **20**, 483–485 (1995).
8. 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.* **18**, 1946–1948 (1993).
9. K. W. DeLong, R. Trebino, J. R. Hunter, and W. E. White, "Frequency-resolved optical gating with the use of second-harmonic generation," *J. Opt. Soc. Am. B* **11**, 2206–2215 (1994).
10. 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.* **19**, 2152–2154 (1994).
11. W. Koenig, H. K. Dunn, and L. Y. Lacy, "The sound spectrograph," *J. Acoust. Soc. Am.* **18**, 19–49 (1946).
12. S. H. Nawab, T. F. Quatieri, and J. S. Lim, "Signal reconstruction from short time Fourier transform magnitude," *IEEE Trans. Acoust. Speech Signal Process.* **ASSP-31**, 986–998 (1983).
13. R. A. Altes, "Detection, estimation, and classification with spectrograms," *J. Acoust. Soc. Am.* **67**, 1232–1246 (1980).
14. D. J. Kane, A. J. Taylor, R. Trebino, and K. W. DeLong, "Single-shot measurement of the intensity and phase of femtosecond UV laser pulse using frequency-resolved optical gating," *Opt. Lett.* **19**, 1061–1063 (1994).
15. B. Kohler, V. V. Yakovlev, and K. R. Wilson, "Characterization of amplified pulses for quantum control," in *Generation, Amplification and Measurement of Ultrashort Laser Pulses*, R. P. Trebino and I. A. Walmsley, eds., *Proc. Soc. Photo-Opt. Instrum. Eng.* **2116**, 360–364 (1994).
16. H. Stark, ed. *Image Recovery: Theory and Application* (Academic, Orlando, Fla., 1987).
17. R. P. Millane, "Phase retrieval in crystallography and optics," *J. Opt. Soc. Am. A* **7**, 394–411 (1990).
18. D. Israelevitz and J. S. Lim, "A new direct algorithm for image reconstruction from Fourier transform magnitude," *IEEE Trans. Acoust. Speech Signal Process.* **ASSP-35**, 511–519 (1987).
19. J. R. Fienup, "Reconstruction of a complex valued object from the modulus of its Fourier transform using a support constraint," *J. Opt. Soc. Am. A* **4**, 118–123 (1987).
20. R. G. Lane, "Phase retrieval using conjugate gradient minimization," *J. Mod. Opt.* **38**, 1797–1813 (1991).
21. J. H. Seldin and J. R. Fienup, "Iterative blind deconvolution algorithm applied to phase retrieval," *J. Opt. Soc. Am. A* **7**, 428–433 (1990).
22. J. R. Fienup, "Phase retrieval algorithms: a comparison," *Appl. Opt.* **21**, 2758–2769 (1982).
23. For real experimental data we prefer the following background-removal techniques. In a multishot experiment, in which spectra of the signal field are taken for various values of the relative delay time, we generally subtract off a background spectrum taken under dark conditions from all the measured spectra. Similarly, in a single-shot experiment we subtract the spectrum measured at large delay times to remove the contributions of incoherently scattered light.
24. J. S. Lim, *Two-Dimensional Signal and Image Processing*, (Prentice-Hall, Englewood Cliffs, N.J., 1990), Chap. 9.
25. D. L. Donoho, "Wavelet shrinkage and W. V. D-A ten-minute tour," Stanford University Tech. Rep. 416 (Stanford University, Stanford, Calif., January 1993).