Simulations of frequency-resolved optical gating for measuring very complex pulses

Lina Xu,* Erik Zeek, and Rick Trebino

School of Physics, Georgia Institute of Technology, Atlanta, Georgia 30332, USA
**Corresponding author: gth665y@mail.gatech.edu*

Received November 16, 2007; revised February 20, 2008; accepted February 24, 2008;
posted March 17, 2008 (Doc. ID 89740); published May 6, 2008

Frequency-resolved optical gating (FROG) and its variations are the only techniques available for measuring complex pulses without a well-characterized reference pulse. We study the performance of the FROG generalized-projections algorithm for retrieving the intensity and phase of very complex ultrashort laser pulses [with time–bandwidth products (TBPs) of up to 100] in the presence of noise. We compare the performance of three versions of FROG: second-harmonic-generation (SHG) FROG, polarization-gate (PG) FROG, and cross-correlation FROG (XFROG), the last of which requires a well-characterized reference pulse. We found that the XFROG algorithm converged in all cases on the first initial guess. The PG FROG algorithm converged for all moderately complex pulses, for 99% of the pulses we tried, and for more than 95% of even the most complex pulses (TBP~100). The SHG FROG algorithm converged for 95% of the pulses we tried and for over 80% of even the most complex pulses. We found no additional ambiguities in any of these techniques.

OCIS codes: 320.7100, 120.5050.

1. INTRODUCTION

The shaping of ultrashort laser pulses into complex intensities and phases versus time is finding many applications, including coherent control [1], telecommunications [2], micromachining [3], and multiphoton imaging [4]. In addition, several commercial pulse shapers [5–8] have become available and can generate pulses with time–bandwidth products (TBPs) up to ~100. Complex pulse shapes also occur in continuum generation [9]. Unfortunately, methods for measuring the actual shapes (intensity and phase versus time) of such complex pulses have not received much attention.

Most techniques for measuring ultrashort laser pulses either do not yield the complete time-dependent intensity and phase (e.g., autocorrelation [10]) or can at best only measure simple pulses [e.g., spectral phase interferometry for direct electric-field reconstruction (SPIDER) [11]]. A method has recently been introduced [multiphoton intrapulse interference phase scanning (MIIPS) [12]] that simultaneously shapes and measures pulses, but it uses the same pulse shaper to both shape and measure the pulse and so cannot be said to constitute an independent measurement of the shaped pulse, and it has not been tested on complex pulse shapes. If a well-characterized reference pulse is available, linear spectral interferometry is, in principle, capable of measuring complex pulses, but most versions of it have artificially limited spectral resolution and so have not been able to do so. Also, spectral interferometry suffers from extreme alignment sensitivity and so is difficult to use. We recently introduced a simplified version of spectral interferometry [spatially encoded arrangement for temporal analysis by dispersing a pair of light e-fields (SEA TADPOLE) [13]], which has approximately an order of magnitude better spectral resolution than conventional versions, and it avoids the debilitating alignment sensitivity. We have demonstrated it for measuring shaped pulses with TBPs of several hundred. It also can measure the complete spatiotemporal intensity and phase at and near the focus. Like all other versions of spectral interferometry, however, it requires a well-characterized reference pulse whose spectrum contains that of the shaped pulse to be measured. Fortunately, when pulse shaping, such a pulse is generally available in the form of the unshaped pulse. Thus SEA TADPOLE is ideal for measuring shaped pulses, but such measurements require the use of two separate devices, one to measure the unshaped pulse and the SEA TADPOLE to measure the shaped pulse. The reference-pulse requirement prevents spectral interferometry from measuring continua.

Currently, the most commonly used method for measuring shaped and complex pulses (and the simplest) is frequency-resolved optical gating (FROG) [14], which does not require a reference pulse, and so the same device can be used to measure both the unshaped and shaped pulses. Since its introduction in 1991, FROG (see Fig. 1) and its many variations have been used to measure the full intensity and phase of a wide range of ultrashort laser pulses. FROG has measured the intensity and phase of few-femtosecond pulses, pulses over many wavelength regions, and single pulses. Variations on it [cross-correlation FROG (XFROG)] have even measured attosecond pulses [15], but, more importantly, for our purposes herein, XFROG has measured the most complex pulses—continua—ever measured [9]. XFROG, like spectral interferometry, requires a well-characterized reference pulse, but the XFROG reference-pulse spectrum need not contain that of the shaped pulse. The unshaped pulse is also

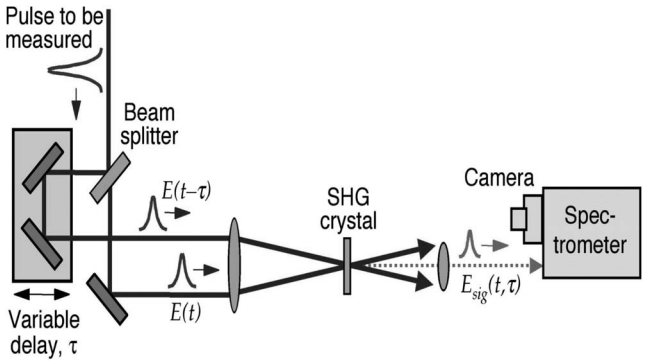

Fig. 1. Schematic of a SHG FROG apparatus. A pulse is split into two, one pulse gates the other in the SHG crystal, and the relative delay is varied. The nonlinear-optical signal pulse spectrum is then measured versus delay. In PG FROG, the nonlinearity is polarization gating in glass, and crossed polarizers are used. In XFROG, an independent (previously measured) reference pulse is used instead of one of the unknown pulses. Other FROG geometries exist and use other nonlinear-optical processes.

ideal for XFROG measurements of the shaped pulse, as well. Nevertheless, the version of FROG most commonly used to measure shaped pulses is second-harmonic-generation (SHG) FROG. SHG FROG is self-referenced, using the pulse to measure itself. Like its fellow FROG techniques, it has many advantages for measuring shaped pulses, including built-in independent checks on its measurements, geometries for single-shot operation, and versatile, often very simple, arrangements. A single SHG FROG–or any self-referenced FROG—device can be used to measure both the unshaped and shaped pulses. The various FROG iterative algorithms [15] are known for their reliability and robustness in measuring relatively simple pulses (TBP$<\sim$5) where they have been used most of the time. However, because pulse-shaping applications can involve rather complex pulses, it is important to check the various FROG algorithms' theoretical performance for complex pulses. This has never been done. Indeed, since in this case such measurements would involve the use of a complex pulse to measure a complex pulse, we have never expected SHG FROG to perform well for such pulses and have always recommended spectral interferometry or XFROG, which use a simple (unshaped) pulse to measure the complex pulse, for such measurements. But, as this advice has generally been ignored, and SHG FROG is already in common use for measuring complex pulses, we consider it in this contribution. Indeed, it is interesting to see how well the various FROG methods measure such complex pulses.

Specifically, we test three FROG iterative algorithms. These include those of SHG FROG; polarization-gate (PG) FROG, which is also self-referenced; and XFROG. They are all based on the same generalized-projections (GP) approach [14]. To do so, we generated complex test pulses in both the time and frequency domains and then added 1% Poisson noise to all the resulting traces. We find that the XFROG algorithm converges with 100% reliability on the first initial guess, which is in agreement with an existing proof that (noise-free) spectrogram (XFROG) inversion should always succeed [16]. On the other hand, we find, as expected, that PG and SHG FROG are not 100% reliable in the presence of noise, but, surprisingly, are much closer to perfection than expected. PG FROG achieves 100% convergence for pulses with TBP$<$30 and $\sim$95% for more complex pulses (including TBP$\sim$100). SHG FROG achieves $>$80% effectiveness at retrieving even the most complex pulses, provided that several initial guesses are allowed if the first fails to yield convergence. Also, when the algorithm fails to converge, it nevertheless succeeds in retrieving the approximate length and general shape of the pulse, failing only in the details, which could still be adequate for many purposes. Thus, while SEA TADPOLE and XFROG remain preferable, if not ideal, for measuring shaped pulses (in our opinion), SHG FROG and, in particular, PG FROG should also provide adequate and relatively robust measurements of such complex pulses, especially if the user desires to use only one device for measurements of both the unshaped and shaped pulses or a reference pulse is not available.

2. FROG, THE GENERALIZED-PROJECTIONS ALGORITHM, AND SIMULATION DETAILS

The general FROG apparatus is shown in Fig. 1.

The expression for the SHG FROG, PG FROG, or XFROG trace is

$$I_{\mathrm{FROG}}(\omega,\tau)=\left|\int_{-\infty}^{\infty}E(t)E_g(t-\tau)\exp(-i\omega t)\mathrm{d}t\right|^2,$$

where $E(t)$ is the unknown input-pulse electric field that we are trying to measure. These simple versions of FROG are distinguished by their gate pulses: in SHG FROG, $E_g(t)=E(t)$; in PG FROG, $E_g(t)=|E(t)|^2$; and in XFROG, $E_g(t)$ is an independently measured pulse. XFROG is mathematically equivalent to the well-known spectrogram, and SHG FROG and PG FROG yield autospectrograms (the pulse gates themselves). Like all time-frequency-domain methods, all three FROG methods involve measurements of intensity versus two variables, frequency (ω) and delay (τ). In other words, SHG FROG and PG FROG are spectrally resolved autocorrelations [17], and XFROG is a spectrally resolved cross correlation.

All versions of FROG were shown to be examples of mathematical problems called two-dimensional phase-retrieval and hence are able to yield essentially unique solutions as long as the entire FROG trace is nearly completely contained in the data set (that is, not significantly cropped at its edges) and the trace details are properly resolved in time and frequency [14]. The phase-retrieval algorithm usually used to retrieve the pulse from the measured trace is the GP algorithm [14,18]. It involves one-dimensionally Fourier transforming the signal field, $E_{\mathrm{sig}}(t,\tau)=E(t)E_g(t-\tau)$, back and forth between the t and ω domains, effectively alternately iteratively projecting onto two constraint sets (corresponding to the two equations above; see Fig. 2), eventually leading to the intersection of these two constraints—the final solution. Early versions of the FROG code were reliable, but slow; fortunately, more recent versions are much faster due to the use of faster programming languages (e.g., C), faster code, and

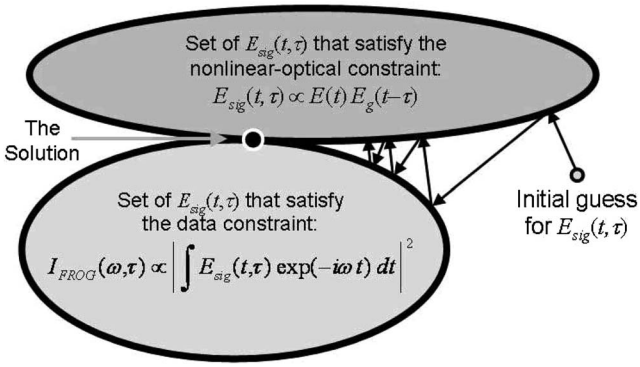

Fig. 2. Generalization-projections algorithm for FROG.

the realization that the generalized projections algorithm is so reliable that additional algorithms included for extra reliability were unnecessary and so have been eliminated. Commercial FROG codes now routinely achieve 20 pulse retrievals per second and are very reliable [19].

The initial guess can affect the convergence of the GP algorithm. When retrieving very simple pulses, the GP algorithm is generally not sensitive to the particular choice of an initial guess. But, when measuring complex pulses, the initial guess could, in principle, become more important. Early on, random noise was found to be the best initial guess for the FROG algorithm for simple pulses and is generally used as the initial guess in FROG programs. Nevertheless, in preparation for this study, we compared random noise and a Gaussian flat-phase pulse as initial guesses, and we found that random noise is generally a better initial guess when measuring complex pulses also. Thus, in all of our simulations, we have used random noise for the initial guess.

To test the performance of the GP algorithm in SHG FROG, PG FROG, and XFROG for measuring complex pulses, we generated a large set of complicated pulses. We generated each pulse by starting with a sequence of random complex numbers and multiplying it by a Gaussian pulse in the time domain. We then Fourier transformed the resulting pulse and multiplied its frequency-domain version by a Gaussian spectrum. We chose the widths of the Gaussians to yield the desired complexity. This yielded a complex pulse in both intensity and phase and in both time and frequency, ideal for testing the ability of these methods to measure generally complex pulses.

Although specifying the Gaussian widths yielded a

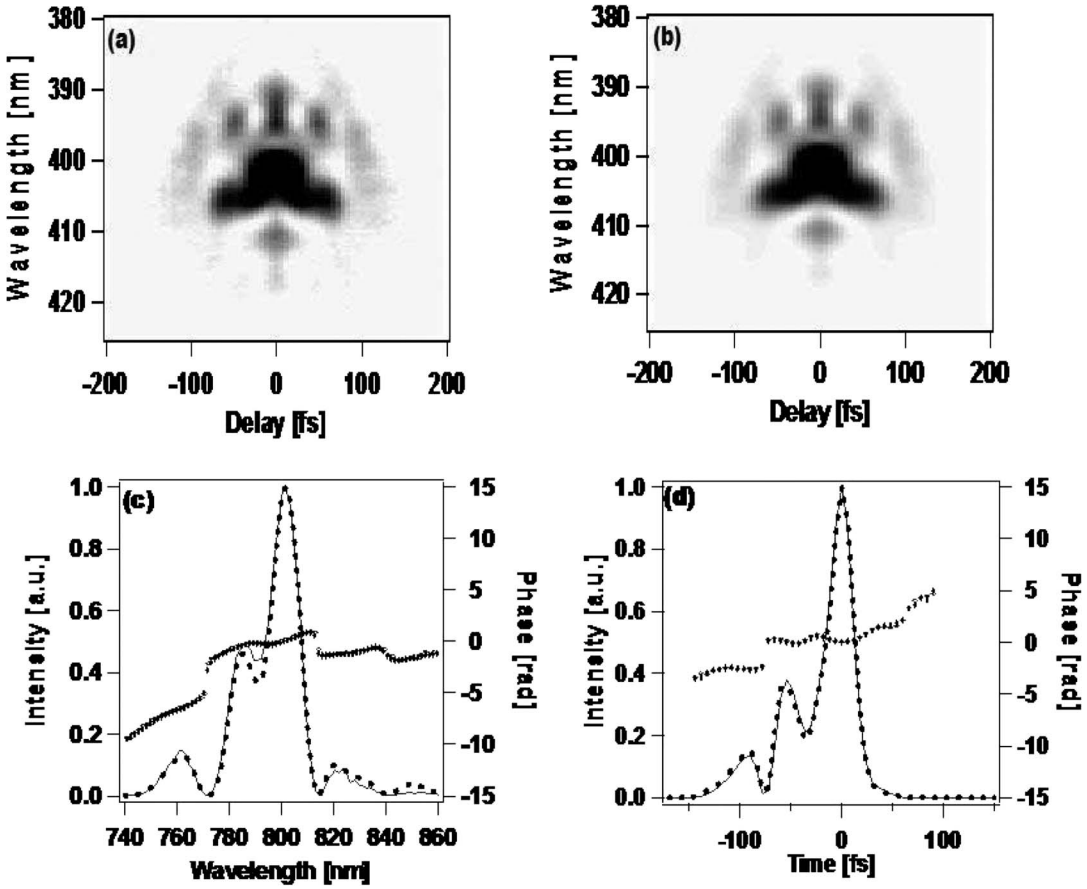

Fig. 3. Moderately complex pulse with TBP=4.7. (a) Original SHG FROG trace with noise, (b) retrieved trace, (c) generated and retrieved spectral intensity and phase, (d) generated and retrieved temporal intensity and phase. In (c) and (d) (and in all subsequent analogous figures), the generated pulse is indicated by curves and the retrieved pulse by dots. Good convergence has occurred here. The FROG trace is well contained in the image window. For better display, the shown FROG traces were cut from the generated one, which had a wavelength range from 327.675 to 514.385 nm and a delay range from −768 to 762 fs. The maximum value of three rows and columns along the perimeter of the image window is 0.88% of the peak value after background subtraction, so the trace was only slightly cropped.

good estimate of the pulse TBP, it is not exact. So, for each pulse, we computed the precise root-mean-square (rms) TBP as our measure of the complexity of the pulse:

$$\mathrm{TBP}_{\mathrm{rms}} = t_{\mathrm{rms}}\omega_{\mathrm{rms}},$$

$$\text{where} \quad t_{\mathrm{rms}}^2 = \langle t - \langle t\rangle^2\rangle = \langle t^2\rangle - \langle t\rangle^2,$$

$$\langle t^2\rangle = \int_{-\infty}^{\infty} t^2 I(t)\mathrm{d}t,$$

$$\omega_{\mathrm{rms}}^2 = \int_{-\infty}^{\infty} A'(t)^2\mathrm{d}t + \int_{-\infty}^{\infty} A(t)^2\phi'(t)^2. \tag{1}$$

In the above expressions, $I(t)$ is the normalized intensity, t_{rms} is the rms temporal width, and ω_{rms} is the rms spectral width. $A(t)$ is the real amplitude, and $\phi(t)$ is the temporal phase. The prime indicates the derivative [20].

We computed the FROG traces for the generated complex pulses. To simulate the experimental environment, we also added 1% additive Poisson noise to each FROG trace. This additive noise yields pixel-to-pixel signal variations independent of the FROG trace intensity. The reason for using Poisson noise is that it approximates the noise from practical noise sources, such as dark current. Such additive noise is also more challenging for the algorithm than multiplicative noise, which necessarily goes to zero in the wings of the trace. In this approach, the measured trace with the additive noise at each pixel will be [21]

$$I_{\mathrm{FROG}}^{(\bar{\eta})}(\omega_i, \tau_j) = I_{\mathrm{FROG}}(\omega_i, \tau_j) + \eta_{ij}\alpha/\bar{\eta}, \tag{2}$$

where η_{ij} is a pseudorandom number drawn from a Poisson distribution of mean $\bar{\eta}$, and a is the noise fraction, which was set to 0.01, and the mean of the Poisson distribution was 5 counts.

Suppressing background noise is important in FROG measurements. Any nonzero average background (due to noise) in a FROG trace implies spurious nonzero intensity at large times and with high frequency in the pulse, that is, spurious pulse wings with high frequency noise. Thus, in practice, before running the pulse retrieval program, background subtraction is always performed. Several methods are available, and they include Fourier low-pass

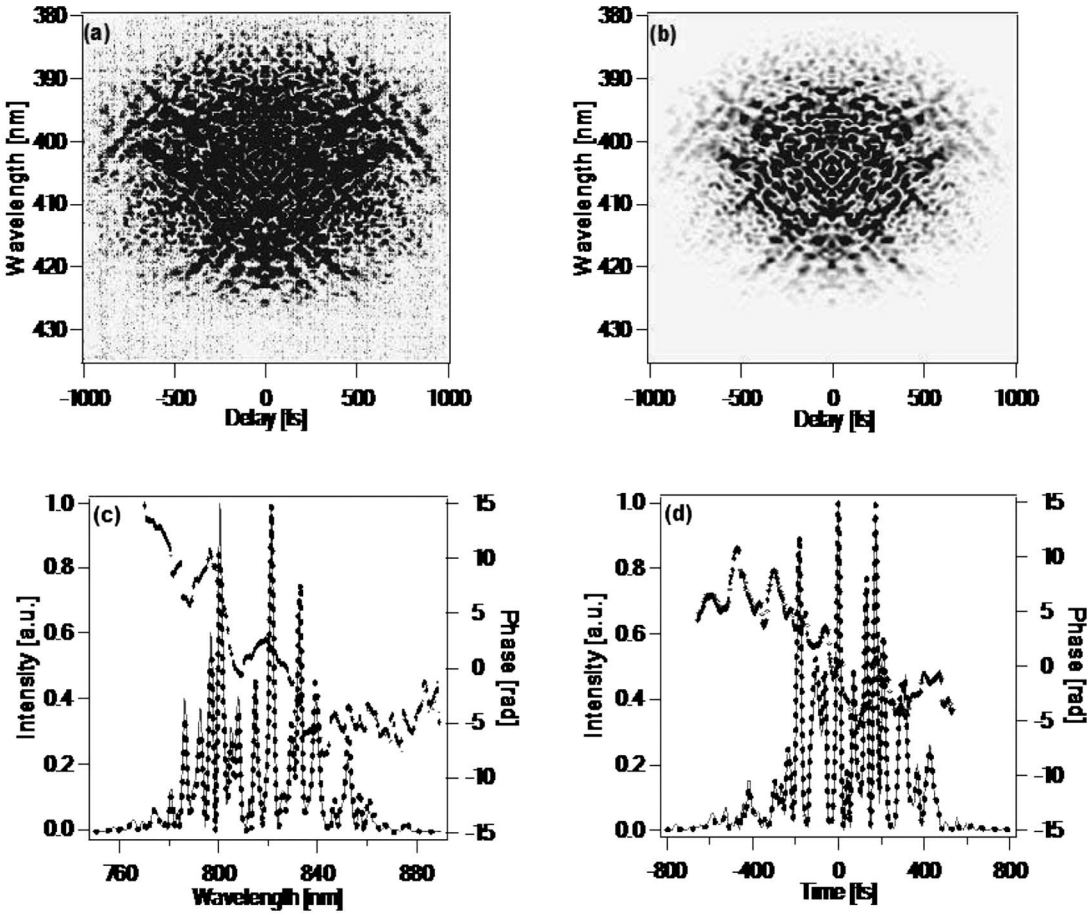

Fig. 4. Very complex pulse with TBP=40.6. (a) Original SHG FROG trace with noise, (b) retrieved trace, (c) generated and retrieved spectral intensity and phase, (d) generated and retrieved temporal intensity and phase. Good convergence has occurred here. The generated trace (here and in later figures also) appears somewhat darker due to the additive noise applied to it (and whose mean has been subtracted prior to running the algorithm); this subtraction and the algorithm combine to remove most of the added noise. Indeed, note the identical structure in both pulses and traces. For better display, the shown FROG trace is cut from the generated one with a wavelength range from 326.503 to 512.02 nm and a delay range from −1536 to 1530 fs. The FROG trace is well contained in the image window. The maximal value of the three rows and columns along the perimeter of the image window is 1.25% of the peak value after background subtraction.

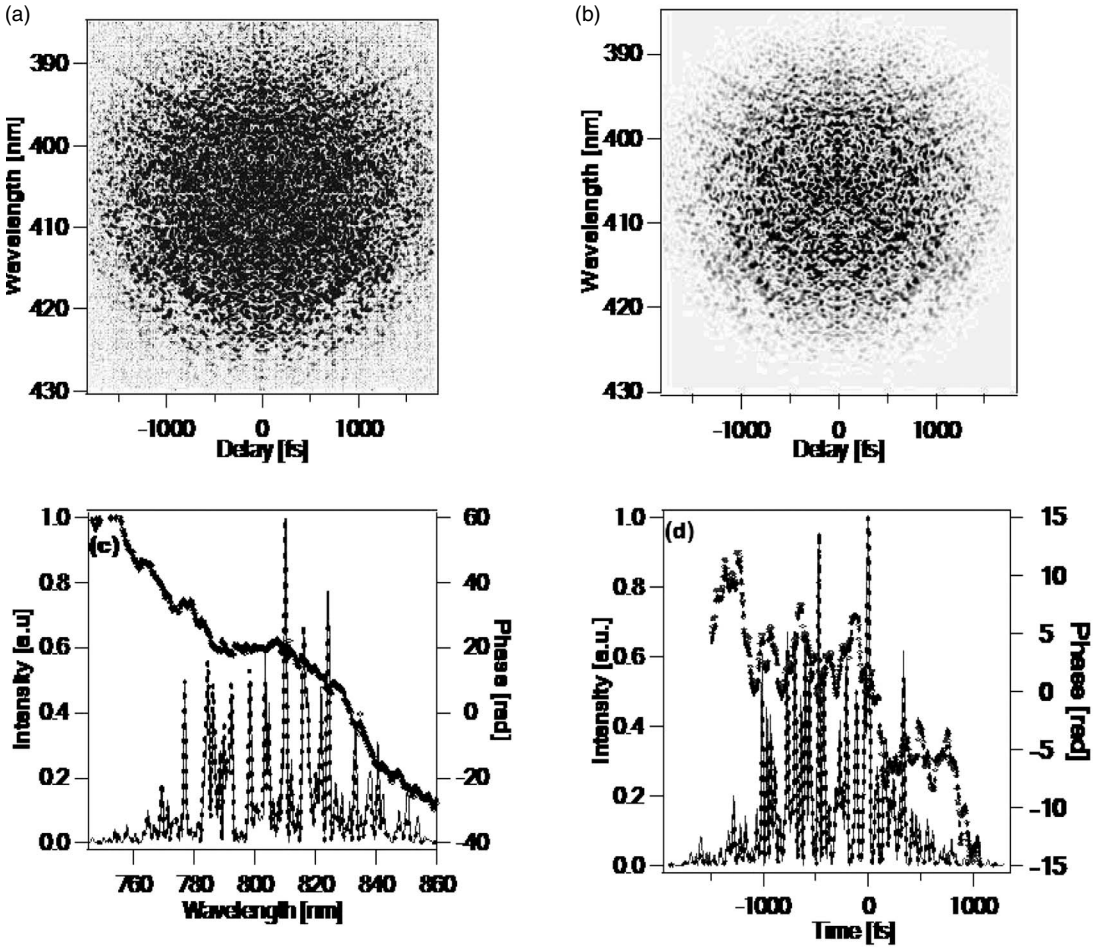

Fig. 5. Extremely complex pulse with TBP=94.3. (a) Original SHG FROG trace, (b) retrieved FROG trace (with noise), (c) generated and retrieved spectral intensity and phase, (d) generated and retrieved temporal intensity and phase. Good convergence has occurred here. The FROG trace is well contained in the image window. For better display, the shown FROG trace is cut from the generated one with a wavelength range from 327.348 to 514.385 nm and a delay range from −3072 to 3066 fs. The maximal value of three rows and columns along the perimeter of the image window is 1.11% of the peak value after background subtraction.

filtering, corner suppression, and mean background subtraction. In our simulations, we chose to perform only simple mean-background subtraction (although performing the others as well would likely have further improved the performance beyond what we observe). The mean of the noise was obtained by averaging the values in the 5 $\times$5 pixel squares in the four corners of the FROG trace (i.e., far from the center of the trace, where the most important pulse information is located). After subtracting this constant background from all points in the trace, we set all the resulting negative points to zero (as is usually done).

Another point worth mentioning is that, for the best retrieval, the FROG trace should be well contained in the imaging window. In our experience, most traces that have yielded retrieval problems have been severely cropped and so should not be expected to yield good convergence (since they do not contain all the pulse information). Thus we were careful to use arrays with delay and spectral ranges large enough to fully contain the trace. Specifically, it is best if the maximum trace value at the edges of the array is less than 1% of the peak value of the FROG trace, although by chance some of the traces in this study exceeded this value (and, we later discovered, proved to be more likely to exhibit convergence problems). Also, it is important to have sufficient temporal and spectral resolution, which simply means that the number of points be sufficient for the relevant trace complexity. These issues should be kept in mind when performing experiments, where the experimental spectral resolution must be sufficient (sufficient temporal resolution is generally easy) and the temporal and spectral ranges must also be sufficient, which could be a challenge for extremely complex pulses. We dealt with these issues by increasing the size of the array as the TBP of the pulse increased, using a 256 $\times$256 array for the simplest pulses up to 1024$\times$1024 for the most complex pulses we studied.

Figure 3 gives an example of a moderately complex pulse with a TBP of 4.7 and its SHG FROG trace (with noise) and the resulting retrieved pulse. The retrieved intensity and phase of the pulse and the generated one agree well with each other. A measure of the success of a pulse retrieval is the FROG error (the rms difference between the input and retrieved traces), which is 0.00341 for this 256$\times$256 grid, which is very good. Other measures of the retrieval success are available when performing simulations, and they are the mean rms intensity and phase errors between the generated pulse and the re-

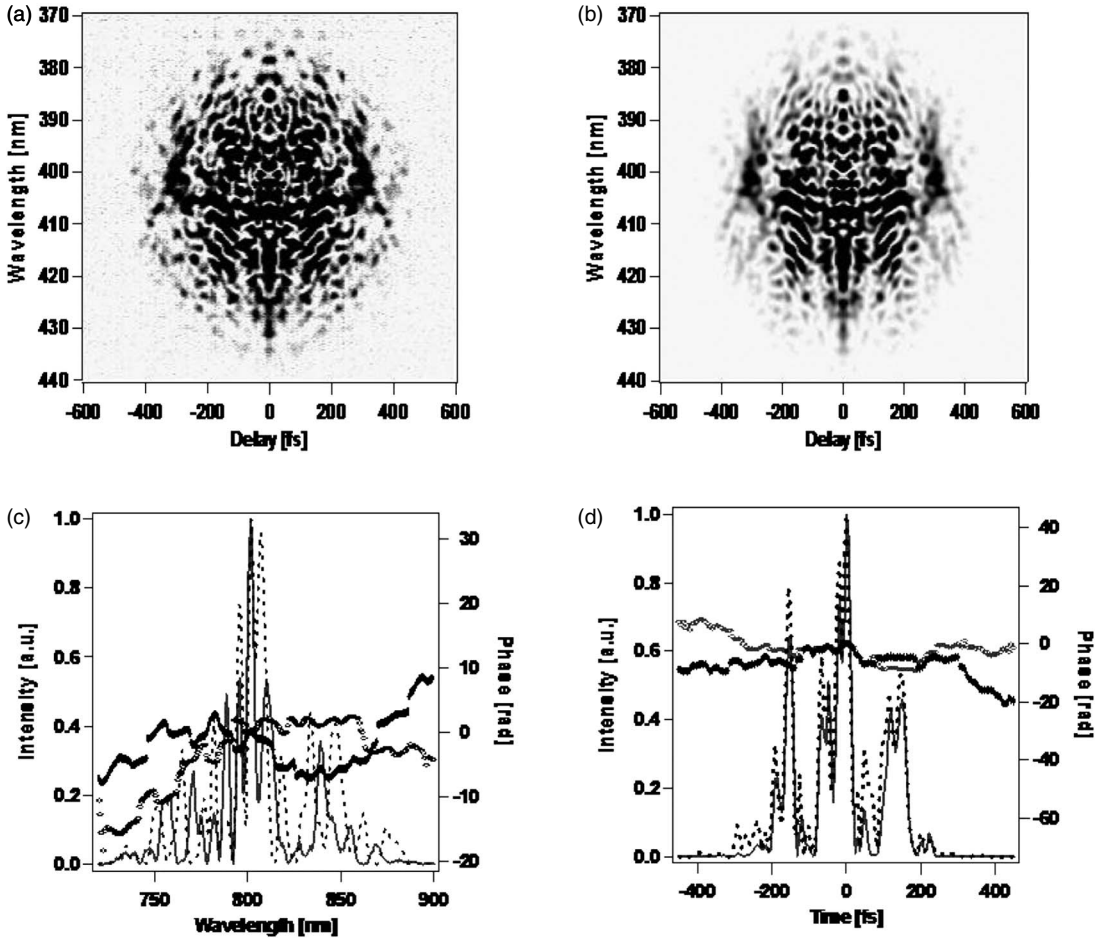

Fig. 6. A pulse for which convergence has not been achieved. (a) Generated SHG FROG trace with noise, (b) retrieved trace, (c) generated and retrieved spectral intensity and phase, (d) generated and retrieved temporal intensity and phase. Note the discrepancies between the generated and retrieved pulses. For better display, the shown FROG trace is cut from the generated one with a wavelength range from 327.464 to 514.385 nm and a delay range from −1530 to 1536 fs. The FROG trace is not as well contained in the image window as in the previous examples, perhaps the reason for the poor convergence. The maximal value of the three rows and columns along the perimeter of the image window is 1.51% of the peak value after background subtraction.

trieved pulse [21]. For this pulse, we found them to be 0.007 and 0.0209, respectively. The excellent pulse retrieval in this case is not surprising because the SHG FROG GP algorithm is known to work very well when measuring such moderately complex pulses.

A considerably more complex pulse, with a TBP of 40.6, and its corresponding SHG FROG trace are shown in Fig. 4. The retrieved intensity and phase of the pulse agree very well with the generated curves. The FROG error is 0.0052 for this 512×512 grid. The mean rms intensity and phase errors are 0.0274 and 0.0219, respectively.

An extremely complex pulse, with a TBP of 94.3, and its corresponding traces are shown in Fig. 5. The retrieved pulse agrees very well with the generated pulse here as well. The FROG error is 0.4% for the 1024×1024 grid. The mean rms intensity and phase errors are 0.0337 and 0.0452, respectively.

Because the mathematical constraints in SHG and PG FROG are not purely convex, convergence of the GP algorithm is not necessarily guaranteed, and we find that convergence cannot be achieved for all traces.

In practice, convergence of the GP algorithm is generally indicated by the FROG error. In this work, as in previous work, we find that, for FROG data with ~1% additive noise, convergence is achieved when the FROG error is less than ~1%. When the FROG error is greater than 1%, the GP algorithm can be seen to have generated a pulse that is visibly different from the generated pulse. See Fig. 6, which shows an example for which convergence has not been achieved (TBP_{rms} of 38.9). The SHG FROG error for this case is 1.6% for this 512×512 grid size.

In Fig. 7, we show the distribution of the SHG FROG error for 30 pulses with TBP values from 30 to 40. We found for all converging cases that the FROG error is much less than 0.5%, and for the nonconverging cases, the FROG error is greater than 1.5%. We confirmed the converging and the nonconverging cases by visual inspection of the generated and retrieved traces, intensities, and phases.

As a result, in this analysis, we defined convergence to have failed for a given initial guess if a FROG error of $<1\%$ has not been achieved, and we define convergence to have failed in general if a FROG error of less than 1% has not been achieved after ten runs of the algorithm using ten different randomly generated initial guesses. We

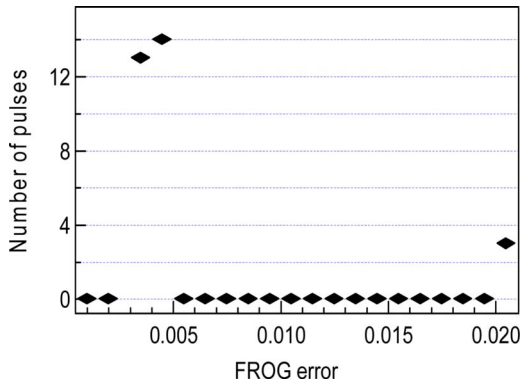

Fig. 7. (Color online) Histogram of FROG errors for 30 pulses with a TBP value from 30 to 40, showing a clear delineation between converging (FROG error <1%) and nonconverging (FROG error >1%) cases.

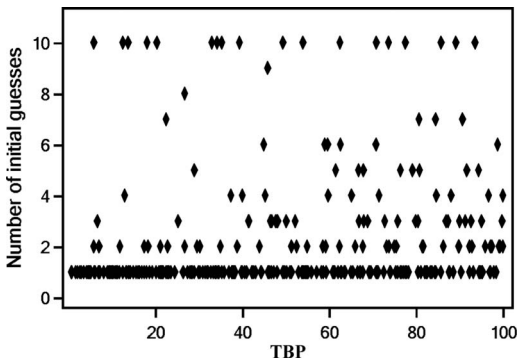

Fig. 8. Number of initial guesses required for correct pulse retrieval in SHG FROG versus TBP for the pulses in our analysis. Note that most pulses can be retrieved in SHG FROG using only a few initial guesses, but some (shown as requiring ten pulses) cannot.

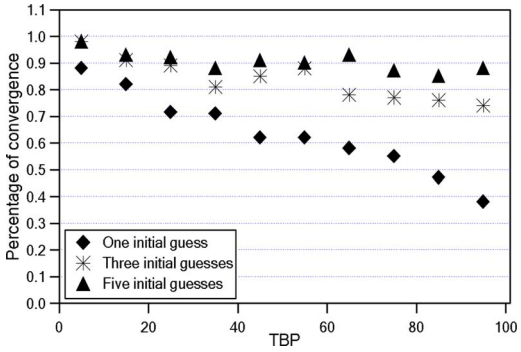

Fig. 9. (Color online) Statistical analysis of the performance of the GP algorithm in SHG FROG. In most cases, when convergence is not achieved after one initial guess, convergence is achieved after a few more tries, but not always.

should point out that, at first glance, it would seem that use of the rms intensity and phase errors would be more rigorous and hence more appropriate. However, use of these definitions is actually significantly complicated by several factors. First, one would need to take into account SHG FROG's ambiguity in the direction of time and the need to curve fit each intensity and phase curve to the precise center of time, peak intensity, and absolute phase—quantities not measured by FROG (or any other method for measuring ultrashort pulse shapes). Also, phase-unwrapping issues as well as the meaninglessness of the phase when the intensity approaches zero (and the associated arbitrariness in the definition of the phase error) further complicate the problem. As a result, the FROG error is a far better approach for automating this analysis, which involves more than 1000 pulses. However, we also visually inspected a large fraction of the retrieved intensities and phases to verify convergence and also to confirm the absence of ambiguities (possibly very different pulses with the same traces) beyond the trivial ones mentioned above.

Using this approach, we studied the general performance of the various FROG GP algorithms for measuring complex pulses with TBPs up to 100. We generated 350 pulses for each statistical analysis. The TBP's of these pulses ranged from 1 to 100. For the purposes of displaying our results, we binned sets of 35 pulses evenly into intervals of TBP ranging from 1 to 10, 11 to 20, etc.

3. TESTING THE SHG FROG GP ALGORITHM

The results of our analysis for SHG FROG are shown in Fig. 8. We find that, in general, the more complex the pulse, the more initial guesses on average are needed for convergence. In other words, the SHG FROG algorithm is more sensitive to the initial guess for more complex pulses.

Figure 9 shows the percentage convergence for pulses as a function of the pulse rms TBP. Of course, with more initial guesses, the percentage of convergence increases. In general, we find that SHG FROG works remarkably well for such complex pulses, given the complexity of the problem. When five initial guesses are allowed, more than 80% of the pulses, and occasionally 90%, can be retrieved, even for extremely complicated pulses with TBPs of ~100. In most cases, even when convergence is not achieved after five initial guesses, convergence is achieved after a few more tries. Only 5% of the pulses failed to achieve convergence after ten initial guesses.

4. TESTING THE PG FROG GP ALGORITHM

We performed an analogous simulation for PG FROG. A typical example is shown in Fig. 10. The TBP of the generated pulse is 15.5. The retrieved intensity and phase versus time and wavelength agree with the generated ones. The FROG error is 0.0034 for this 512×512 grid. The mean rms intensity and phase errors are 0.0236 and 0.0293, respectively.

We generated 300 new random pulses with TBPs from 1 to 100 to test the general performance of the PG FROG GP algorithm. Thirty pulses were generated for each TBP interval. The results are shown in Figs. 11 and 12. We found that the PG FROG GP algorithm works extremely well for retrieving complex pulses with TBPs less than 40. Even for extremely complex pulses, the PG FROG GP algorithm converged most of the time. Only one initial guess was needed to obtain the correct pulse in most cases. Only three nonconverging cases occurred, and it is

interesting that no traces yielded convergence for four to nine initial guesses, implying that the nonconvergent traces were pathological in some way. Indeed, we believe that, by chance, these PG FROG traces were not well contained in the array. We plan to rerun this analysis later, checking for such cropping and eliminating such cropped traces from the analysis (or enlarging the array before running the algorithm).

5. TESTING THE XFROG GP ALGORITHM

We performed an analogous simulation of complex pulse measurement using XFROG. We used a reference pulse that was a simple Gaussian pulse with a FWHM of 50 fs and zero phase. Although a great deal has been written about choosing optimal pulses for generating spectrograms, we made no effort to optimize this pulse for optimal results. A typical example pulse and trace are shown in Fig. 13. The TBP value of this generated pulse is 66. The retrieved phase and intensity versus time and versus wavelength agree with the generated ones. Note that this trace is simpler than SHG and PG FROG traces for similarly complex pulses due to the simplicity of the gate pulse in XFROG. The FROG error is 0.003 for this 512 $\times$512 grid. The mean rms intensity and phase errors are 0.0411 and 0.0314, respectively.

We also generated 350 new pulses with TBPs from 1 to 100 to test the general performance of the XFROG GP algorithm. As before, 35 pulses were generated for each TBP interval from 1 to 10, 11 to 20, etc. The results are shown in Fig. 14. We found that the XFROG GP algorithm works extremely well for retrieving even the most complex pulses in the presence of noise, converging for every pulse on the first try. Only one initial guess was needed to obtain the correct pulse for every case, and the FROG error was always less than 0.01. Our results verify that the spectrogram [14] always yields the correct pulse and that the GP algorithm is an excellent algorithm to do so.

6. DISCUSSION AND ADDITIONAL OBSERVATIONS

The three FROG algorithms we considered in this study performed quite well. XFROG performed perfectly, PG FROG performed very well, and SHG FROG performed reasonably well. The performance order scales with the complexities of the gate pulses in these techniques.

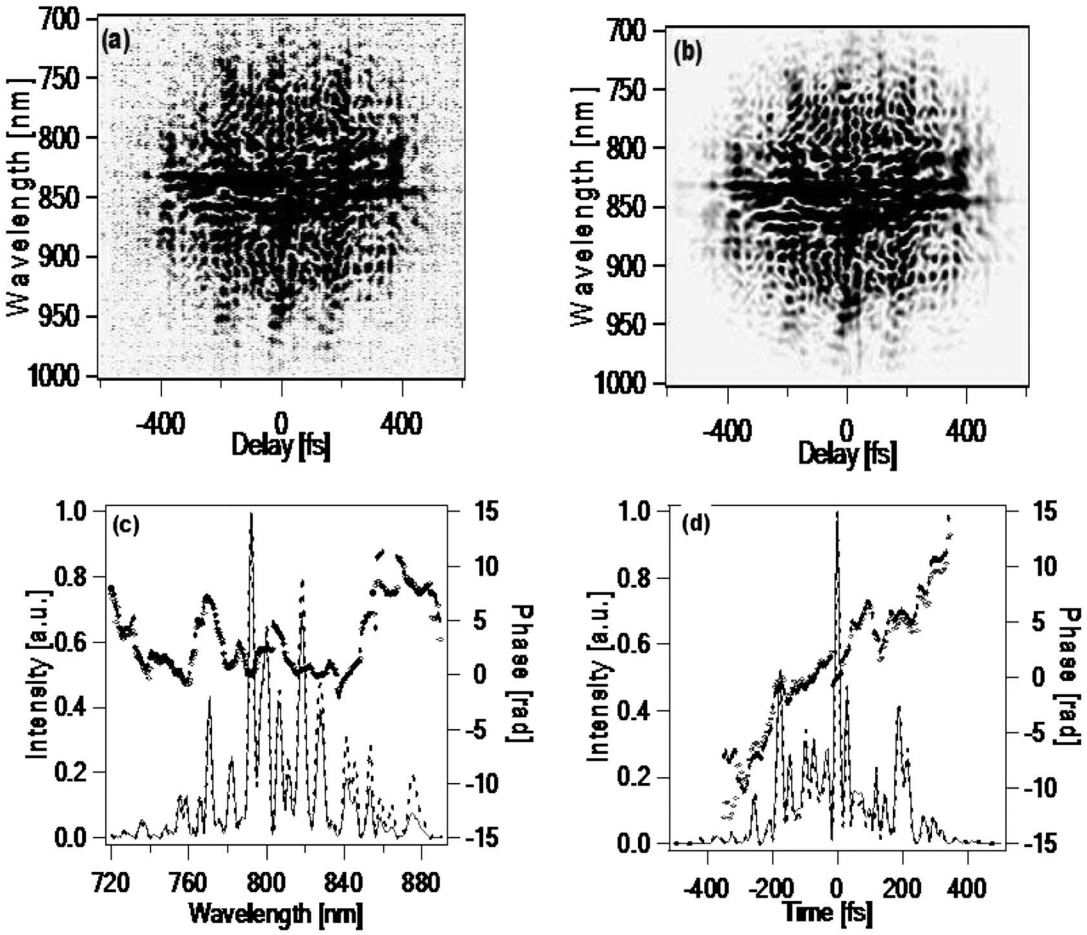

Fig. 10. Example of PG FROG for measuring a complex pulse (here a pulse with TBP=15.5). (a) Generated FROG trace with noise, (b) retrieved trace, (c) generated and retrieved spectral intensity and phase, (d) generated and retrieved temporal intensity and phase. For better display, the shown FROG trace is cut from the generated one with a wavelength range from 327.464 to 514.385 nm and a delay range from −1536 to 1530 fs. The PG FROG trace is well contained in the image window. The maximal value of three columns along the perimeter of the image window is 0.89% of the peak value after background subtraction.

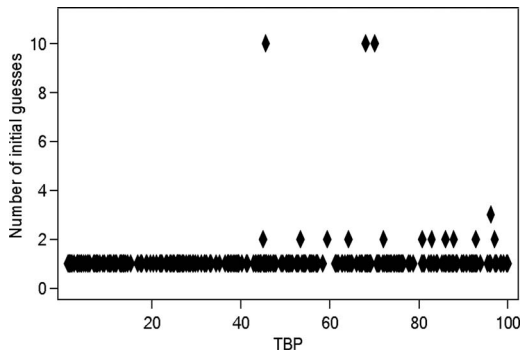

Fig. 11. Number of initial guesses required for correct pulse retrieval in PG FROG versus TBP. Note that most pulses can be retrieved using only one initial guess, and nearly all can be retrieved after two or three.

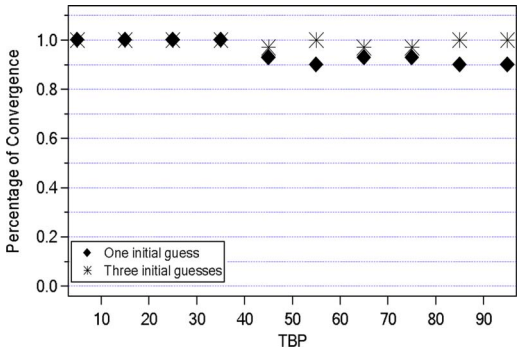

Fig. 12. (Color online) Statistical analysis of the performance of the GP algorithm in PG FROG.

XFROG generally uses a very simple gate pulse, and ours was particularly simple; PG FROG uses the intensity of the unknown pulse, and SHG FROG uses the unknown pulse intensity and phase. The simpler the gate pulse, the better the performance.

The question that naturally arises is why the SHG and PG FROG algorithms converge so reliably for most extremely complex pulses but not for some considerably less complex pulses. In the course of working with the various FROG techniques over time, we have noted that one way to yield poor convergence is to crop the trace in time or frequency. We have noticed that, in this study, the traces that failed to yield convergence were those with the largest nonzero values along their perimeters. The simple solution in this case would be to scan further in delay and frequency. In a future study, we plan to further investigate this issue, and we believe that we will achieve even better convergence then. Also, it would be interesting to see if the use of additional noise-filtering techniques and

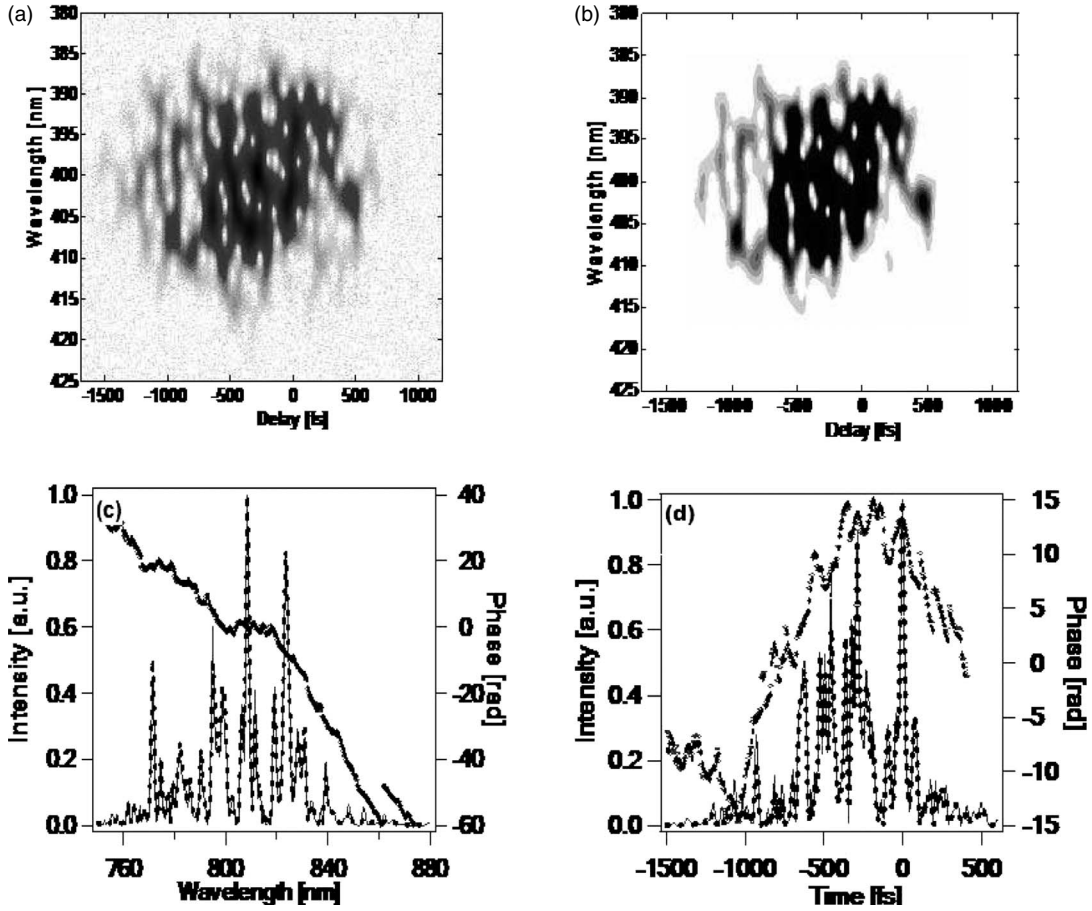

Fig. 13. Example of XFROG for measuring complex pulses (here a pulse with TBP=66). (a) Generated FROG trace with noise, (b) retrieved trace, (c) generated and retrieved spectral intensity and phase, (d) generated and retrieved temporal intensity and phase. For better display, the shown FROG trace is cut from the generated one with a wavelength range from 343.02 to 480.065 nm and a delay range from −2048 to 2040 fs. The XFROG trace is well contained in the image window. The maximal value of three columns along the perimeter of the image window is 1.013% of the peak value.

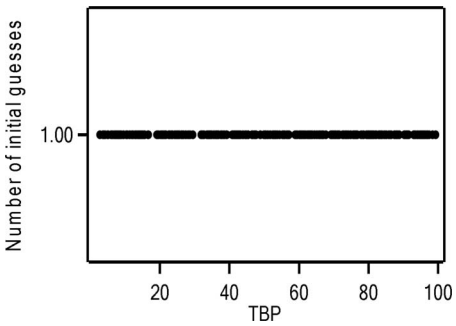

Fig. 14. Statistical analysis of the performance of the XFROG GP algorithm. Convergence is always achieved after only one initial guess, even for extremely complicated pulses.

additional algorithmic techniques, such as those used in previous versions of the FROG algorithm but abandoned when the GP approach proved to work so well, would help.

Another question that our study addressed is the issue of the possibility that other, previously unknown, ambiguities exist in these techniques. It is well known that SHG FROG possesses a (trivial) ambiguity in the direction of time (which is easily removed by a second measurement with additional glass in the beam), and most pulse measurement techniques, including FROG, do not measure the absolute phase and the pulse arrival time and also have (trivial) ambiguities in the relative phases of well-separated pulses and modes (although a properly designed XFROG measurement lacks these). But our study of complex pulses allowed us to search for additional previously unknown ambiguities (trivial or otherwise) that could in principle occur for complex pulses. In view of the fact that we added noise to the traces, it was also possible that we could find approximate ambiguities, that is, additional pulses whose traces are not identical to that of the correct pulse, but, due to the presence of noise, are within the experimental error of the trace of the correct pulse and yet are quite different from the correct pulse. However, in the ~1000 randomly generated complex pulses studied in this effort, we found no pulse whose retrieved trace was equal to or very similar to its original trace but whose retrieved intensity and phase differed significantly from the original pulse. Thus, we found no new exact—or approximate—ambiguities in SHG FROG, PG FROG, or XFROG. This confirms that these techniques should work well for the measurement of simple and complex pulses.

It should also be mentioned that, while the FROG algorithm is very fast for relatively simple pulses (typically using 64×64 traces and requiring a fraction of a second), it is much slower for complex pulses, such as those of this study, scaling as $N^2 \ln N$, where $N\times N$ is the size of the array.

Another issue is the accuracy with which FROG can measure zeros and weak regions of the pulse. Logarithmic plots of one of our SHG FROG simulations, Figs. 4(c) and 4(d), are shown in Fig. 15. Only the intensity is shown. The generated and retrieved pulses agree very well in the lower intensity regions, with some slight discrepancies, usually well below 1%. In view of our addition of 1% noise to the trace, this represents very good performance. Thus, FROG could be expected to measure relatively weak regions of the pulse with good accuracy.

Finally, we should point out that, when the FROG algorithm does not converge, this fact is made clearly evident by the relatively large FROG error, so one always knows when an additional initial guess is required.

7. CONCLUSIONS

We have simulated the performance of the generalized-projections algorithm for retrieving very complex pulses (with TBP up to 100) from SHG FROG, PG FROG, and XFROG traces in the presence of additive noise.

For SHG FROG, we find that, even if the pulse is extremely complicated, the intensity and phase of the pulse can usually be retrieved. But more than one initial guess is often needed for such complicated pulses in the presence of noise.

The PG FROG algorithm performance is considerably better. It always converged for simple and moderately complex pulses with TBPs up to 30. Overall, it could retrieve 99% of the pulses we tried, and when the TBP value is greater than 30, approximately 95% of the complex pulses could be retrieved using only one initial guess.

XFROG worked perfectly, retrieving all complex pulses on the first initial guess, even in the presence of noise.

We conclude that, if a suitable reference pulse is available, and the user is willing to build two separate devices,

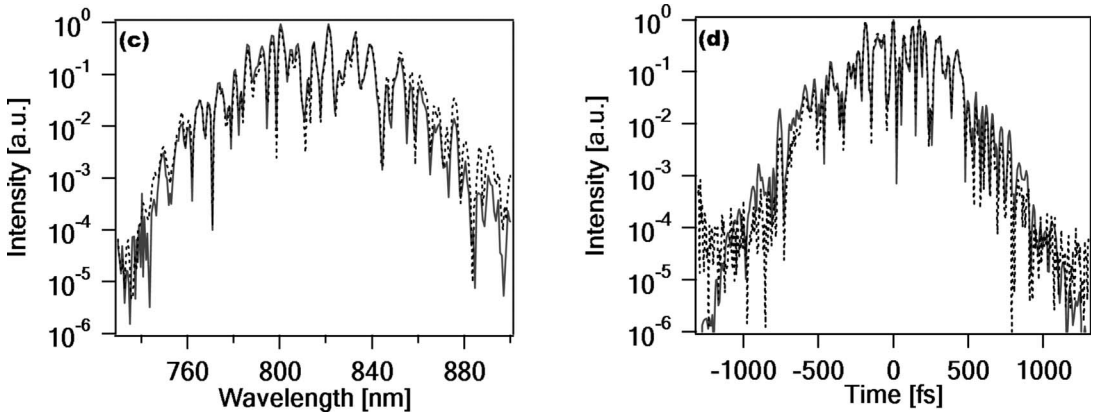

Fig. 15. Logarithmic plots of the generated and retrieved pulses from Fig. 4(c) (the spectrum of the pulse) and Fig. 4(d) (the temporal intensity of the pulse).

the XFROG and SEA TADPOLE techniques remain the best choices for measuring complex pulses, but, if not, PG FROG is an excellent choice, and SHG FROG is a reasonable choice also.

REFERENCES

1. J. P. Ogilvie, D. Débarre, X. Solinas, J.-L. Martin, E. Beaurepaire, and M. Joffre, "Use of coherent control for selective two-photon fluorescence microscopy in live organisms," Opt. Express **14**, 759–766 (2006).
2. A. M. Weiner, J. P. Heritage, and J. A. Salehi, "Encoding and decoding of femtosecond pulses," Opt. Lett. **13**, 300–302 (1988).
3. I. H. Chowdhury, X. Xu, and A. M. Weiner, "Laser machining using temporally controlled ultrafast pulses," in *Conference on Lasers and Electro-Optics/International Quantum Electronics Conference and Photonic Applications Systems Technologies*, Technical Digest (CD) (Optical Society of America, 2004), paper CThD5.
4. C. J. Bardeen, V. V. Yakovlev, J. A. Squier, K. R. Wilson, S. D. Carpenter, and P. M. Weber, "Effect of pulse shape on the efficiency of multiphoton process: implications for biological microscopy," J. Biomed. Opt. **4**, 362–367 (1999).
5. A. M. Weiner, "Femtosecond pulse shaping using spatial light modulators," Rev. Sci. Instrum. **71**, 1929–1960 (2000).
6. E. Zeek, K. Maginnis, S. Backus, U. Russek, M. Murnane, G. Mourou, H. Kapteyn, and G. Vdovin, "Pulse compression by use of deformable mirrors," Opt. Lett. **24**, 493–495 (1999).
7. W. Hillegas, J. X. Tull, D. Goswami, D. Strickland, and W. S. Warren, "Femtosecond laser pulse shaping by use of microsecond radio-frequency pulses," Opt. Lett. **19**, 737–739 (1994).
8. E. Frumker, E. Tal, Y. Silberberg, and D. Majer, "Femtosecond pulse-shape modulation at nanosecond rates," Opt. Lett. **30**, 2796–2798 (2005).
9. X. Gu, L. Xu, M. Kimmel, E. Zeek, P. O'Shea, A. P. Shreenath, R. Trebino, and R. S. Windeler, "Frequency-resolved optical gating and single-shot spectral measurements reveal fine structure in microstructure-fiber continuum," Opt. Lett. **27**, 1174–1176 (2002).
10. K. Sala, G. Kenney-Wallace, and G. Hall, "CW autocorrelation measurements of picosecond laser pulses," IEEE J. Quantum Electron. **16**, 990–996 (1980).
11. C. Iaconis and I. A. Walmsley, "Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses," Opt. Lett. **23**, 792–794 (1998).
12. V. V. Lozovoy, I. Pastirk, and M. Dantus, "Multiphoton intrapulse interference IV. Ultrashot laser pulse spectral phase characterization and compensation," Opt. Lett. **29**, 775–777 (2004).
13. P. Bowlan, P. Gabolde, A. Shreenath, K. McGresham, R. Trebino, and S. Akturk, "Crossed-beam spectral interferometry: a simple, high-spectral-resolution method for completely characterizing complex ultrashort pulses in real time," Opt. Express **14**, 11892–11900 (2006).
14. R. Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic, 2002).
15. A. Baltuska, M. S. Pshenichnikov, and D. Wiersma, "Amplitude and phase characterization of 4.5 fs pulses by frequency-resolved optical gating," Opt. Lett. **23**, 1474–1476 (1998).
16. R. A. Altes, "Detection, estimation, and classification with spectrograms," J. Acoust. Soc. Am. **67**, 1232–1246 (1980).
17. 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).
18. K. W. Delong, D. N. Fittinghoff, R. Trebino, B. Kohler, and K. Wilson, "Pulse retrieval in frequency-resolved optical gating based on the method of generalized projection," Opt. Lett. **19**, 2152–2154 (1994).
19. D. J. Kane, "Real-time measurement of ultrashort laser pulses using principal component generalized projections," IEEE J. Sel. Top. Quantum Electron. **4**, 278–284 (1998).
20. L. Cohen, *Time-Frequency Analysis* (Prentice-Hall, 1995).
21. D. N. Fittinghoff, K. W. DeLong, R. Trebino, and C. L. Ladera, "Noise sensitivity in frequency-resolved optical-gating measurements of ultrashort pulses," J. Opt. Soc. Am. B **12**, 1955–1967 (1995).