

Determining error bars in measurements of ultrashort laser pulses

Ziyang Wang, Erik Zeek, and Rick Trebino

Georgia Institute of Technology, School of Physics, Atlanta, Georgia 30332

Paul Kvam

Georgia Institute of Technology, School of Industrial and Systems Engineering, Atlanta, Georgia 30332

Received February 10, 2003; revised manuscript received July 23, 2003; accepted July 23, 2003

We present a simple and automatic method for determining the uncertainty in the retrieved intensity and phase versus time (and frequency) due to noise in a frequency-resolved optical-gating trace, independent of noise source. It uses the “bootstrap” statistical method and also yields an automated method for phase blanking (omitting the phase when the intensity is too low to determine it). © 2003 Optical Society of America
OCIS codes: 320.7100, 320.5550.

1. WHY ERROR BARS ARE NOT TYPICALLY REPORTED

How does one place error bars on a measurement of an ultrashort laser pulse? For decades, about the only available measure of an ultrashort laser pulse was the autocorrelation. Unfortunately, an autocorrelation trace typically corresponds to many, often quite different, intensities, so even a perfect noise-free measurement of the autocorrelation results in a large, and unknown, uncertainty in the shape of the pulse’s intensity versus time.^{1–3} In addition, the autocorrelation yields no phase-versus-time information at all. As a result, it is not possible to determine a pulse intensity or phase from an autocorrelation. Even when additional information, such as the spectrum, is included, the autocorrelation measurements typically correspond to many significantly different intensities and phases, even in the absence of noise.⁴ Such methods thus have a type of unquantifiable “internal noise” that occurs even in the absence of measurement error, and it makes no sense to attempt to place error bars on such reconstructions of the pulse’s intensity and phase.

Of course, it is now possible to measure a pulse’s intensity and phase versus time (and frequency) essentially unambiguously. The most commonly used method, frequency-resolved optical gating¹ (FROG), retrieves the full pulse intensity and phase versus time (and frequency) without the need for any assumptions about the pulse. And recent advances in the FROG technique have expanded its capabilities significantly. It has been extended to a wide range of wavelengths, from the IR to the x-ray, and FROG measurements can now be made with an extremely simple beam geometry that make it the simplest pulse-measurement technique¹ available. FROG’s cousin, cross-correlation FROG, is now being used to measure extremely complex pulses, such as ultrabroadband continuum originating from a microstructured fiber with a time–bandwidth product in excess of 1000.^{5,6} And in addition, with new nonlinearities, cross-correlation FROG can now measure ultraweak pulses with <100 photons

each.⁷ FROG is also a very accurate technique, with accuracy limited only by the accuracy of the measured trace. Numerous simulations and experiments have both corroborated this fact.¹

But just how accurate is a given FROG measurement of a pulse? Some indication of the measured-pulse accuracy is available from the “FROG error,” the rms difference between the measured and retrieved FROG traces, which tells us how well the retrieved pulse’s trace matches the measured trace (and usually reveals the presence of systematic error in the measurement). However, this measurement depends on the trace size, and it gives no quantitative information regarding the error in the intensity and phase at a given time or frequency. Singular-value decomposition-based algorithms provide an alternative measure of systematic error, but, like the FROG error, it also tells little about the error in the derived intensity and phase points.⁸ What FROG, and pulse measurement, in general, need, is a method for determining the uncertainty in each of the retrieved intensity and phase points, that is, error bars.

Unfortunately, direct computation of error bars in any experiment involves a tedious, and often questionable, accounting of all the known sources of error. Worse, when the relationship between the measured data and retrieved values is complex, as is the case in FROG, it is not clear how such sources of error in the data translate to actual uncertainties in the retrieved values. As a result, error bars are not often reported in measurements and are essentially never reported in pulse measurements.

Therefore we present here a simple, robust, and general technique for placing error bars on the intensity and phase retrieved in a FROG measurement. It operates automatically, makes no assumptions about the pulse or error, and requires no extra measurements or analysis, instead operating with only a single measured trace.

A related issue in the measurement of ultrashort pulses (and measurements of phase, in general) is the tendency of the phase to become meaningless as the intensity goes

to zero. This is obvious, but the problem is this: At what point should we stop plotting the phase? Omitting phase points for which the intensity is below some threshold is often called “phase blanking.” Usually, in deciding the threshold for phase blanking, a judgment must be made, and it is often made based on aesthetics rather than science.

Here we present an automated method that makes this decision objectively and appropriately. Once error bars are determined, the problem of phase blanking is quite simple: When the phase error exceeds or equals 2π , then the phase is clearly undetermined, and phase blanking is appropriate. Note that, once the technique for the determination of error bars is automated, phase blanking is also. Moreover, it requires no arbitrary judgments on the part of the user.

2. BOOTSTRAP METHOD

The technique that solves the phase-blanking and error-bar problems is the “bootstrap” method, a well-established statistical resampling tool for determining uncertainty.^{9–11} The bootstrap method involves taking the data set of M points, sampling a new set of M points from it, with replacement (resulting in some points possibly occurring more than once and others not at all), and then running the relevant parameter-determination algorithm on this new set of points. This process is then repeated numerous times on the data set (generating many resamplings of M points and hence many computations of the desired parameters), so that each desired parameter will be estimated numerous times.

It has been shown that, when the data overdetermine the parameters (i.e., when there are many more data points than parameters), the set of parameter values obtained by this procedure approximates the sampling distribution for the original parameter estimate.^{10,11} In particular, the mean and standard deviation of this distribution are consistent estimators for the measured parameter and the error bar for that parameter, respectively. This method is called the bootstrap method because it appears that one is getting something for nothing (“pulling oneself up by one’s own bootstraps”), but this is not the case, and, in fact, the method can be quite computation intensive. With the use of modern computers, however, the required computations are not prohibitive, and the computations we perform require at most a few minutes on a personal computer. The bootstrap method has the additional advantage that it is “nonparametric,” that is, it assumes nothing about the distribution of the noise and parameters.

The bootstrap method is quite general and so has been successfully applied to a wide range of problems.^{10,11} In particular, its application to positron emission tomography^{12,13} and single-photon emission computed tomography¹² closely resembles our application, in which a complex algorithm operates on multidimensional data to determine the desired quantities.

Applying this approach to ultrashort-pulse measurement simply involves running the FROG retrieval algorithm on approximately 100 resamplings of the measured FROG trace, tabulating the statistics of the retrieved intensity and phase values obtained during these runs (see

Fig. 1). The mean intensity and phase values for each time and frequency are then the measured values. More importantly, the standard deviations yield the error bars. This works because the FROG trace overdetermines the pulse, that is, contains many more points than the resulting intensity and phase.

To resample the points according to the bootstrap procedure, we take the original, measured FROG trace and select, at random and with replacement, a number of points equal to the original number. This resampling with replacement allows data points to be selected more than once, and typically about 2/3 of the points occur at least once in this new trace. Data points not selected are simply ignored in the generalized-projections FROG algorithm¹ (in the magnitude replacement step), although other versions of the FROG algorithm¹ can accommodate duplicate points. Running the algorithm with only a fraction of the points does not harm its accuracy for this purpose. Indeed, the degree to which the solution varies when points are removed is the desired measure of the error. And, as mentioned above, it has been shown that, in general, using this procedure, the statistics of the retrieved values accurately approximate the actual statistics of the derived parameters,^{10,11} in this case, the intensity and phase values at the various times (and frequencies). We then take the resulting mean intensity (or phase) for each time as the actual intensity (or phase), and plus or minus one standard deviation at each point as the uncertainty, although one could use plus or minus two or three times the standard deviation depending on one’s tolerance for error.

A minor point: While we use the means of the bootstrap retrieved intensities and phases, the retrieved value for each intensity and phase obtained using only the original trace, without bootstrap resampling, can be shown to be a slightly better estimate than the mean values and can therefore be used instead. This decision does not affect the bootstrap standard deviations and error bars.

3. DEALING WITH AMBIGUITY

This process involves some subtlety in its application to pulse measurement, however, because no pulse-

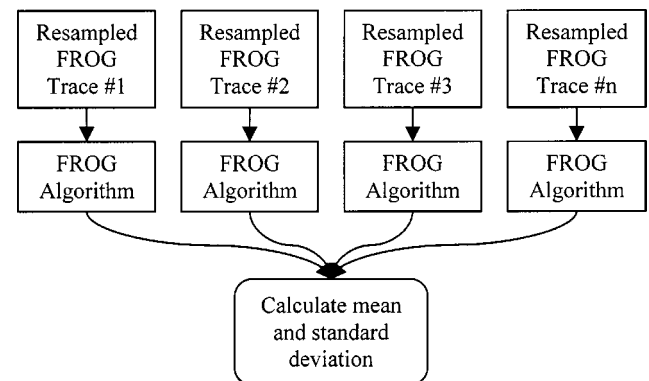


Fig. 1. Schematic of the bootstrap process for application to FROG measurement of ultrashort pulses. Each of the ~ 100 resampled traces (obtained from a single measured trace as described in the text) is run through the FROG algorithm, and the mean and standard deviation of each point of the retrieved pulse intensity and phase versus time and frequency is calculated.

measurement technique actually determines all pulse parameters. In particular, FROG has what are known as “trivial ambiguities.” These are ambiguities in the pulse peak intensity, I_0 , the absolute phase, ϕ_0 (the zeroth-order term in the temporal- and spectral-phase Taylor series), and the pulse arrival time, t_0 (which is also the first-order term in the spectral-phase Taylor expansion). Usually nonmeasurement of the absolute phase and arrival time are advantageous, eliminating the need for tedious stabilization of irrelevant path lengths. If, however, the absolute phase is allowed to float, the phase-versus-time and phase-versus-frequency curves will float randomly over the full 2π range in the retrievals required for application of the bootstrap procedure, yielding large phase errors, even in the absence of noise in the trace! Similarly, nonmeasurement of the peak intensity and arrival time will cause excessive errors in the pulse intensity and spectrum. Thus we must carefully fix these parameters at the same (arbitrary) constants in each bootstrap retrieval in order to properly apply the bootstrap method. Indeed, our ability to fix these parameters provides the ultimate limit on how small a meaningful error bar may be accurately determined.

Here we only consider polarization-gate FROG (PG FROG) in order to limit the number of ambiguities. Other versions of FROG (as well as other techniques) have additional ambiguities to consider. For example, second-harmonic-generation FROG (SHG FROG) has a well-known direction-of-time ambiguity. Thus in SHG FROG, each of the retrieved pulses must have its direction of time fixed. While we have devised a method for experimentally removing this ambiguity from SHG FROG by placing an etalon in the beam,¹⁴ it is important that this ambiguity be considered. There is also a relative-phase ambiguity in SHG FROG (and other methods) for well-separated pulses in time and frequency. In this paper, however, we will restrict our attention to the case of PG FROG and the three undetermined parameters mentioned above. We will treat the case of known, and potentially unknown, ambiguities in a future publication; the method will be a generalization of this approach.

Fortunately, it is easy to remove the trivial ambiguities before performing bootstrap computations. In order to fix the delay of the pulse’s arrival time, we simply center the pulse in time. Specifically, we set the first moment of the intensity to be zero. The first moment is

$$t_0 = \frac{\sum_i t_i I(t_i)}{\sum_i I(t_i)}, \quad (1)$$

in which $I(t_i)$ is the intensity of the pulse at time t_i . We simply shift the pulse so that its first moment is at $t_0 = 0$.

To account for the other two trivial ambiguities, I_0 and ϕ_0 , we simply use

$$\hat{E}(t) = \frac{E(t)}{E(t_0)}, \quad (2)$$

where the electric field of the pulse is $E(t)$ and $E(t_0)$ is the electric field at the time t_0 . This single procedure

not only normalizes the intensity at t_0 to unit intensity, but it also sets the phase at t_0 to 0. It is important to note that this (artificially) removes all uncertainty in the pulse field at the point at t_0 . (One could estimate the uncertainty in this point by averaging that of its neighboring points or simply not plotting it, as it is automatically set to 1.)

4. BOOTSTRAP PROCEDURE IN THE ABSENCE OF NOISE

A test of (1) our ability to fix these parameters; (2) the FROG algorithm and technique in general; and (3) the bootstrap procedure for this application is whether we obtain error bars of zero length in the absence of noise. We tested our approach using a complex triple pulse with a phase jump, whose noise-free trace is shown in Fig. 2. To choose which points are included in bootstrap samples (and later to add noise to traces), we used the built-in random-number generator, function `random()`, in the Delphi language and compiled it with the Delphi 5 compiler. We then ran the commercial Femtosoftware FROG code (modified to resample the trace 88 times as described above) on the trace. Figure 3 shows the retrieval of this pulse for a noise-free, polarization-gate (PG) FROG (the version of the FROG method that uses the polarization-gate beam geometry¹) trace including error bars.

We defined the intensity error as simply the integrated error of the intensity over the trace,

$$S_I = \frac{\sum_{i=1}^n \sigma_i^I}{nI_{\max}}, \quad (3)$$

and the intensity-weighted phase error as

$$S_\phi = \frac{\sum_{i=1}^n \sigma_i^\phi I_i}{nI_{\max}}, \quad (4)$$

where σ_i^I and σ_i^ϕ are the mean intensity and phase standard deviations at the i th time or frequency, I_i is the intensity at the i th time or frequency, and I_{\max} is the maxi-

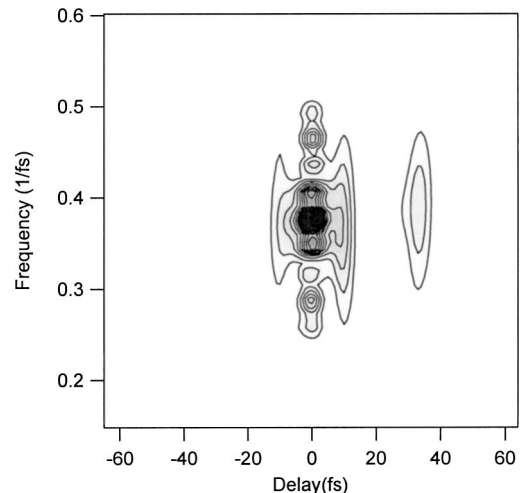


Fig. 2. Polarization-gate (PG) FROG trace of the pulse used in these simulations.

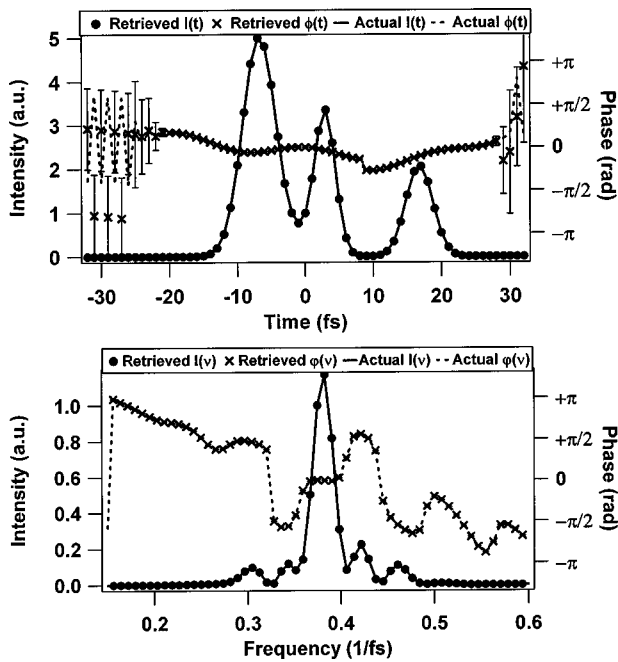


Fig. 3. Retrieved intensity and phase for the noise-free FROG trace in Fig. 2. Error bars have been computed with the bootstrap method as described in the text to determine whether error bars have zero length in the noise-free case, as required. Solid curves are the actual intensity and phase. In time, the integrated intensity error was 1.8×10^{-6} and the integrated intensity-weighted phase error was 2.6×10^{-6} , and in frequency, these errors were 5.7×10^{-6} and 8.7×10^{-8} , respectively. (In this and all other plots, error bars smaller than the size of the data point are omitted.)

imum intensity versus time or frequency. We weight the phase uncertainty by the intensity because the phase and its uncertainty are meaningless when the intensity is zero.

With these definitions, the integrated intensity error for this noise-free trace was 1.8×10^{-6} , and the intensity-weighted phase error was 2.7×10^{-6} (5.7×10^{-6} and 8.7×10^{-8} , respectively, in frequency). These error values for this complex pulse and trace are not just measures of the error due to the bootstrap method, but are in fact the sum of the errors due to our normalization procedure, the numerical round-off error of our personal computer, and the FROG algorithm itself. The low values achieved above show that all of these processes work very well.

5. BOOTSTRAP PROCEDURE IN THE PRESENCE OF NOISE

Having checked our implementation of the bootstrap method for traces without noise, we can now place error bars on the intensity and phase obtained from a trace with noise. To do this, we added 1% additive noise to each point in the above trace to simulate experimental noise. Figure 4 shows the retrieved intensity and phase of the same theoretical pulse, but now with error bars determined by use of the bootstrap method. The error bars represent the ± 1 standard deviation points about the mean value of each retrieved intensity or phase value for each time. Note that the resulting intensity errors are of

the order of 1%, but vary somewhat with intensity. The phase noise is large in the pulse wings, of course, because the intensity goes to zero there, and thus the phase there is meaningless. Note also that $\sim 60\%$ of the actual points fall within the error range, which indicates that this procedure is reasonable.

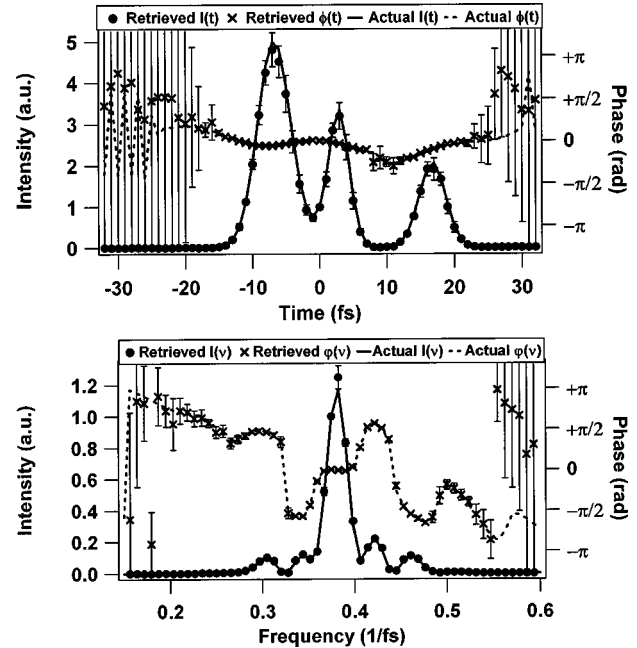


Fig. 4. Retrieved intensity and phase of a theoretical pulse with 1% additive noise introduced numerically to the FROG trace. The intensity error was 9.3×10^{-3} and the phase error was 1.2×10^{-2} , and in frequency, the errors were 2.5×10^{-3} and 3.3×10^{-3} .

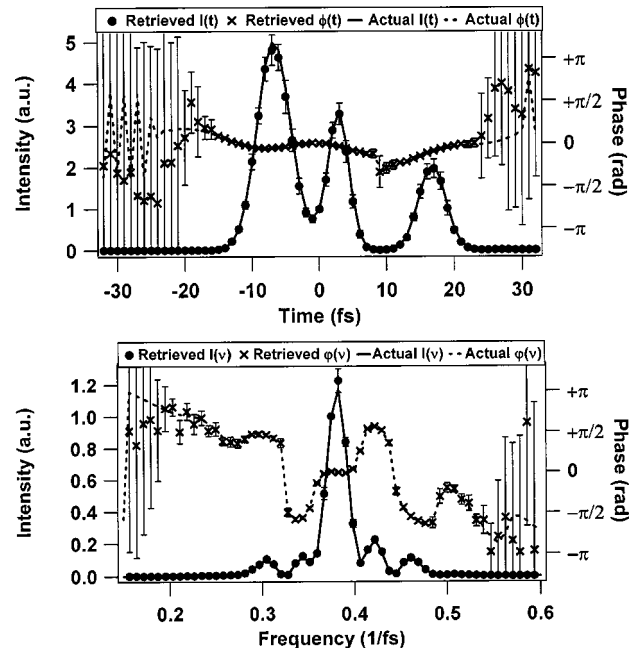


Fig. 5. Retrieved intensity and phase of a theoretical pulse with a different realization of 1% additive noise. Here the intensity error was 9.8×10^{-3} , and the phase error was 1.1×10^{-3} (2.3×10^{-3} and 3.6×10^{-3} in frequency), essentially identical to the retrieval in Fig. 4.

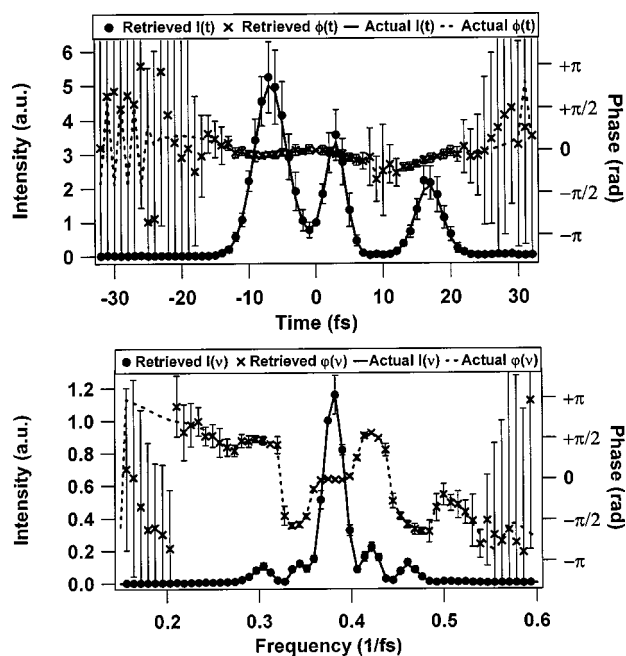


Fig. 6. Pulse retrieved from the same FROG trace, but now with 10% additive noise added. The error bars are about an order of magnitude larger, and the integrated errors are also larger, in time, 2.2×10^{-2} for intensity and 4.5×10^{-2} for phase (in frequency, the numbers are 5.9×10^{-3} and 1.6×10^{-2}).

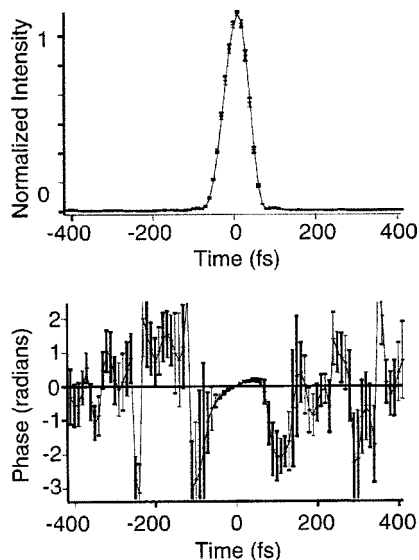


Fig. 7. Error bars with the bootstrap method for an experimentally measured FROG trace.

How do we know these error bars are correct? First, the bootstrap method is supported by a vast array of theoretical statistical analysis.^{10,11} While it does fail for some applications, it is mainly only because too few bootstrap samples were used. Nevertheless, it is important to check that its results are reasonable in this application. This is a bit tricky because there is not currently an established method for the determination of error bars for any pulse-measurement technique!

One test is to generate additional traces using the same average noise but with a different realization of the noise

(a different set of random numbers). Then we can retrieve pulses from these new traces and check whether the distribution of retrieved pulses in this simulation matches those retrieved from the first set. Figures 4 and 5 show examples of two such retrievals. Both begin with the same FROG trace, and each uses the same level of noise, but the random-number generator used a different seed for each trace. This yields different “noise,” but with the same magnitude. That the two sets of computed error bars are nearly identical confirms the validity of the bootstrap approach.

A second test is to show an increasing uncertainty when more noise is added to the trace. There should be a simple monotonic relationship between the computed error bars and the error added to the FROG trace. As can be seen from Fig. 6, the error bars in the retrieval are appropriately longer than those in the 1% noise cases.

We have also run the bootstrap method for experimental FROG data for an 800-nm regeneratively amplified Ti:sapphire laser pulse that yielded a polarization-gate FROG trace. The trace exhibited mainly multiplicative noise, so the retrieved pulse is well defined in the wings with very low error. The resulting measured pulse and its error bars are shown in Fig. 7. These error bars are also quite reasonable.

We have also run this procedure for other noise levels and types and for a variety of pulses and FROG variations, and we have found it to yield reasonable results in all cases.

6. PHASE BLANKING

The bootstrap method also allows us to objectively phase blank. Figure 8 shows our implementation of phase blanking. Here we have taken the pulse from Fig. 6 and applied our phase-blanking technique, which, as mentioned previously, involves omitting the phase when its error bar exceeds 2π in length. The result is much easier to view, without the mass of meaningless information in the wings of the pulse in the original. Importantly, this process is automatic and requires no aesthetic judgments. The only subtlety remaining to resolve is when (and whether) to phase unwrap (forcing the phase to be continuous by adding the appropriate multiple of 2π) and when not to. We find that phase unwrapping during the bootstrap procedure is required (whether phase blanking is desired or not), or else the phase error never exceeds

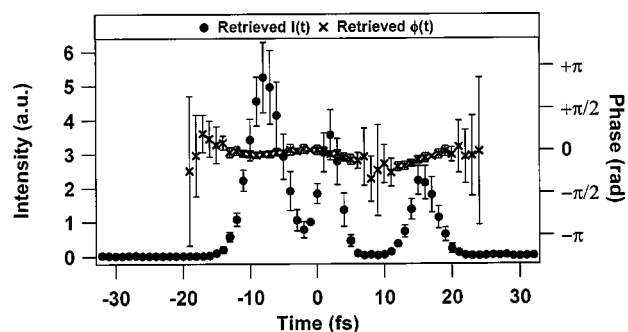


Fig. 8. Pulse from Fig. 6, but with phase blanking applied. Note how the removal of the extra (meaningless) phase points simplifies the plot.

2π . After the bootstrap procedure, one can phase unwrap or not, according to taste.

Phase blanking is especially useful when phase unwrapping. In this case, the phase error can be very large ($>2\pi$) and essentially meaningless, even though it appears to be very small because it is only a tiny fraction of the full range of the phase. Phase blanking using the bootstrap approach provides a quantitative method for determining just how far to plot the phase.

7. CONCLUSIONS

The bootstrap method is a rigorous and general approach to calculating uncertainty in measurements on ultrashort laser pulses. It makes no assumptions about the pulse or the noise on the trace. It is also very accurate, capable of computing error bars as small as a few millionths. In addition, it is also relatively fast. It may seem that, because it requires numerous runs of the FROG algorithm, it could be quite slow. However, due to advances in computer technology and the FROG algorithm, pulse retrieval typically requires from 0.1 s to a few seconds (rarely more than a minute, even for complex pulses) on a PC or Macintosh computer. So implementing the bootstrap procedure should not take too much time in most cases.

Finally, this process is extremely convenient: It is completely automated and easily implemented, especially within the FROG code, and, unlike other error analyses, does not require a careful analysis of the experimental apparatus. Indeed, it requires only the measured trace and no additional measurements. And it allows automatic and objective phase blanking.

ACKNOWLEDGMENTS

The authors gratefully thank Michael Munroe for his code writing and ideas in the early phases of this work. This work was supported by the National Science Foundation (contract ECS-9988706).

REFERENCES

1. R. Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Kluwer Academic, Boston, Mass., 2002).
2. E. J. Akutowicz, "On the determination of the phase of a Fourier integral, I," *Trans. Am. Math. Soc.* **83**, 179–192 (1956).
3. E. J. Akutowicz, "On the determination of the phase of a Fourier integral, II," *Trans. Am. Math. Soc.* **84**, 234–238 (1957).
4. J.-H. Chung and A. M. Weiner, "Ambiguity of ultrashort pulse shapes retrieved from the intensity autocorrelation and power spectrum," *IEEE J. Sel. Top. Quantum Electron.* **7**, 656–666 (2001).
5. 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).
6. J. M. Dudley, X. Gu, L. Xu, M. Kimmel, E. Zeek, P. O'Shea, R. Trebino, S. Coen, and R. S. Windeler, "Cross-correlation frequency resolved optical gating analysis of broadband continuum generation in photonic crystal fiber: simulations and experiments," *Opt. Express* **10**, 1215–1221 (2002).
7. J.-Y. Zhang, A. P. Shreenath, M. Kimmel, E. Zeek, R. Trebino, and S. Link, "Measurement of the intensity and phase of attosecond femtosecond light pulses using optical-parametric-amplification cross-correlation frequency-resolved optical gating," *Opt. Express* **11**, 601–609 (2003).
8. D. J. Kane, F. G. Omenetto, and A. J. Taylor, "Convergence test for inversion of frequency-resolved optical gating spectrograms," *Opt. Lett.* **25**, 1216–1218 (2000).
9. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge University, Cambridge, UK, 1995).
10. B. Efron and R. J. Tibshirani, eds., *An Introduction to the Bootstrap* (CRC, Boca Raton, Fla., 1993).
11. A. C. Davison and D. V. Hinkley, *Bootstrap Methods and Their Application* (Cambridge University, Cambridge, UK, 1997).
12. I. Buvat, "A non-parametric bootstrap approach for analysing the statistical properties of SPECT and PET images," *Phys. Med. Biol.* **47**, 1671–1775 (2002).
13. M. Dahlbom, "Estimation of image noise in PET using the bootstrap method," *IEEE Trans. Nucl. Sci.* **49**, 2062–2066 (2002).
14. E. Zeek, A. P. Shreenath, P. O'Shea, M. Kimmel, and R. Trebino, "Simultaneous automatic calibration and direction-of-time removal frequency-resolved optical gating," *Appl. Phys. B* **74**, S265–S271 (2002).