Measuring ultrashort optical pulses in the presence of noise: an empirical study of the performance of spectral phase interferometry for direct electric field reconstruction

Steven Jensen and Matthew E. Anderson

We have measured the performance of a real spectral phase interferometry for direct electric field reconstruction (SPIDER) apparatus operating under suboptimal conditions. We analyzed the errors in SPIDER's measurements of the temporal phases and intensities of 50-fs ultrashort laser pulses as a function of the additive noise in the detected signal. It was found that SPIDER performs exceptionally well, particularly in the case of additive noise. Specifically, a signal with 10% noise yields a pulse that has a mere 2% error in its intensity profile and a phase that differs from the nominal value by 0.2 rad. Furthermore, we quantified SPIDER's performance with limited detector resolution and as a function of signal averaging. © 2004 Optical Society of America

OCIS codes: 320.5540, 320.7100.

1. Introduction

Ultrashort laser pulses, with durations of only tens of femtoseconds, have enabled exciting research in many areas of physics. With femtosecond laser pulses, interferometric second-harmonic spectroscopy can be performed to identify critical point features in modern solid-state materials,[1] generate x rays,[2] initiate fluorescence in studies of biological tissues,[3] drill holes and cut steel in micromachining applications,[3] excite particles in laser-plasma studies,[4] and manipulate electron states in quantum control experiments.[5] These are only a few examples from a rapidly growing list of applications that use ultrashort laser pulses. Of critical importance to these studies is a firm knowledge of the ultrashort pulse itself, namely, its amplitude and phase, in either the temporal or the frequency domain. Until recently, this was a daunting task.

Fortunately, in the past decade many different devices have been developed to perform measurements on ultrashort pulses. These techniques have moved well beyond the simple spectrometer and autocorrelator,[6] which yielded the intensity envelopes in the spectral or temporal domains. Currently, full-field characterization is possible by way of several different methods. By far the most popular technique is frequency-resolved optical gating (FROG),[7] a beautiful approach that has borne out many variations (and inspired a new era of acronymous invention), including second-harmonic generation FROG,[8,9] twin recovery of electric field envelopes by use of FROG (TREEFROG),[10] temporal analysis by dispersion of a pair of light *e* fields (TADPOLE),[11] multipulse interferometric FROG (MI-FROG),[12] and grating-eliminated no-nonsense observation of ultrafast incident laser light *e* fields (GRENOUILLE).[13] Other pulse measurement techniques include cross-phase modulation,[14] phase and intensity from cross correlation and spectrum only (PICASO),[15] direct optical spectral phase measurement,[16] sonograms,[17] spectral interferometry,[18] and spectral phase interferometry for direct electric field reconstruction (SPIDER).[19,20]

SPIDER uniquely combines two advantageous elements found in pulse measurement devices. SPIDER requires no moving parts and uses a direct, noniterative retrieval algorithm that produces a unique and unambiguous phase and intensity profile for the measured pulse. In addition, SPIDER has been shown to be accurate,[19–22] fast[23,24] (it was recently operated at 1 kHz), and capable of measuring

When this research was performed the authors were with the Department of Physics, San Diego State University, 5500 Campanile Drive, San Diego, California 92182-1233. M. E. Anderson's e-mail address is matt@sciences.sdsu.edu.

S. Jensen is now with Photonics Division, General Atomics, 10240 Flanders Court, San Diego, California 92121.

Received 17 April 2003; revised manuscript received 23 October 2003; accepted 27 October 2003.

0003-6935/04/040883-11$15.00/0

© 2004 Optical Society of America

single pulses as well as coherent pulse trains.[25] SPIDER has measured pulse durations down to less than 6 fs,[26] it can operate over a wide range of frequencies from infrared to blue, and it costs relatively little. SPIDER is gaining recognition as an excellent instrument for applications that require the complete characterization of an ultrashort laser pulse.

Of course, understanding the performance of any metrology instrument is a precursor to using it confidently in a laboratory setting. To this end, several studies have addressed the performance of SPIDER. Dorrer investigated the effects of spectrometer calibration error on the performance of SPIDER.[27] He found that problems associated with calibration error were completely eliminated through use of reference phase differences. Yeremenko *et al.* tested several pulse measurement techniques head to head and suggested new metrics to quantify their performance.[28] Anderson *et al.* performed numerical investigations on the SPIDER inversion routine to determine its response to nonideal data.[21] In this last study, the authors programmed computer simulations to determine how well SPIDER could reconstruct the electric field of an arbitrary pulse in the presence of noise, how well SPIDER performed with an input signal of limited resolution, and what effect averaging had on the accuracy of the measured field. Anderson *et al.* determined the optimum parameters for SPIDER devices and showed that SPIDER was relatively insensitive to noise and the resolution of the detected signal. They also showed that the accuracy of SPIDER improved significantly when several signals were averaged prior to executing the retrieval algorithm. Dorrer and Walmsley[22,29] proposed a new metric to quantify the performance of SPIDER, based on the rms electric field error, and numerically investigated this error as a function of signal-to-noise ratio (SNR). These studies illustrate a simple and physically reasonable estimate of SPIDER's ability to reconstruct optical pulses, and the results show good agreement with the Anderson *et al.* study.[21] All these studies were numerical, however, and what remains to be seen is how these results will translate to a real laboratory setting. It is precisely this endeavor that we address in this paper.

2. Theoretical Basis

At the heart of SPIDER's theoretical underpinnings is spectral-shearing interferometry. The fundamental principle is as follows. An ultrashort-pulse pair that has a separation in time τ and a separation in center frequency Ω is measured with a spectrometer. The resulting interference pattern, or spectral interferogram, contains information about the spectral phase $\phi(\omega)$ of the original ultrashort pulse. It is precisely this information that the SPIDER algorithm retrieves. Combining the spectral phase with the electric field spectral amplitude $E(\omega)$ (given by the square root of the intensity spectrum), the pulse is uncovered. Namely,

$$E(\omega)\exp[-i\phi(\omega)] \tag{1}$$

contains the requisite information to completely define our pulse.

The experimental complexities of SPIDER are detailed in Section 3. For now, however, assume we have generated two pulses that are temporally and spectrally sheared and given by

$$\begin{aligned} E_1(\omega) &= |E_1(\omega)|\exp[-i\phi(\omega)], \\ E_2(\omega-\Omega) &= |E_2(\omega-\Omega)|\exp\{-i[\phi(\omega-\Omega)+\omega\tau]\}, \end{aligned} \tag{2}$$

where τ represents the temporal delay between pulses and Ω represents the spectral shear.

When the two pulses of Eqs. (2) spectrally interfere in the spectrometer, they produce fringes. The intensity of the fringes is dictated by the squared modulus of the resultant electric field, which equals

$$\begin{aligned} S(\omega) = {} & |E_1(\omega)|^2 + |E_1(\omega)||E_2(\omega-\Omega)|\exp\{i[\phi(\omega) \\ & - \phi(\omega-\Omega)+\omega\tau]\} + |E_1(\omega)||E_2(\omega-\Omega)| \\ & \times \exp\{-i[\phi(\omega)-\phi(\omega-\Omega)+\omega\tau]\} \\ & + |E_2(\omega-\Omega)|^2 . \end{aligned} \tag{3}$$

A computer records this interferogram and executes the SPIDER retrieval algorithm. To characterize the incident pulse, the retrieval algorithm needs to determine $E(\omega)$ and $\phi(\omega)$ for the incident pulse. The spectral amplitudes of the two pulses are nearly equal, i.e., $E_1(\omega) = E_2(\omega) = E(\omega)$, and can easily be obtained from any spectrographic measurement of one pulse or even two pulses with a spectral shear of zero. We typically measure $E(\omega)$ in a calibration measurement before calculating $\phi(\omega)$.

When the SPIDER retrieval algorithm calculates the phase, it begins with Eq. (3). A Fourier transform of Eq. (3) is given by

$$\begin{aligned} \Im\{S(\omega)\} = {} & \int_{-\infty}^{\infty} d\omega[\exp(i\omega t)][|E_1(\omega)|^2 + |E_2(\omega-\Omega)|^2] \\ & + \int_{-\infty}^{\infty} d\omega\{\exp[i\omega(t+\tau)]\}|E_1(\omega)||E_2(\omega \\ & - \Omega)|\exp\{i[\phi(\omega)-\phi(\omega-\Omega)]\} \\ & + \int_{-\infty}^{\infty} d\omega\{\exp[i\omega(t-\tau)]\}|E_1(\omega)||E_2(\omega \\ & - \Omega)|\exp\{-i[\phi(\omega)-\phi(\omega-\Omega)]\} . \end{aligned} \tag{4}$$

The top line of Eq. (4) contains the amplitude information (the spectral intensities) that we already know; the next two lines contain the phase information. Lines four and five of Eq. (4) are time-reversed replicas of the same information; one is centered about $t = +\tau$ and one about $t = -\tau$ (the amplitude information is centered about $t = 0$). We use a square window filter to isolate the temporal sideband containing $+\tau$.

Next we take the inverse complex Fourier transform of the filtered signal. This returns the original

interference term, containing $+\tau$, minus amplitude information as

$$\exp\{i[\phi(\omega) - \phi(\omega - \Omega) + \omega\tau]\}. \quad (5)$$

The value for τ is measured independently and easily removed from Eq. (5) with multiplication by $\exp(-i\omega\tau)$. This leaves us, as the argument of the exponential, $[\phi(\omega) - \phi(\omega - \Omega)]$.

The phase can be extracted from this argument with a concatenation process. This process samples the phase at values of the frequency separated by Ω, the spectral shear. For spectral shears that are small relative to the structure of the spectral phase, direct integration can be used to uncover the spectral phase. Namely, if we let

$$\theta(\omega) = [\phi(\omega) - \phi(\omega - \Omega)] \approx \Omega\frac{d\phi(\omega)}{d\omega}, \quad (6)$$

then

$$\phi(\omega) = \frac{1}{\Omega}\int\theta(\omega)d\omega. \quad (7)$$

The sampling theorem ensures that, as long as aliasing does not occur in our original measurement of the signal, the calculated phase will be unique. The spectral phase can now be combined with the spectral amplitude to yield the complete electric field in the spectral domain. A simple Fourier transform yields the electric field in the temporal domain:

$$E(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty}|E(\omega)|\exp[-i\phi(\omega)]\exp(i\omega\tau)d\omega. \quad (8)$$

Measuring the spectral phase is the experimentally challenging component of this endeavor. In Section 3 we detail the experimental implementation.

3. Experimental Methods

A. Precision Versus Accuracy

In our experiment we were interested in how our measured temporal phase and temporal intensity changed as experimental conditions varied; e.g., we wanted to know to what degree the measured temporal intensity deviated from its optimal value if the noise in the signal increased. This illustrates an important difference between our experiments and the numerical simulations. The computer simulations began with an exactly known pulse shape; they compared this pulse shape to the pulse extracted from the inversion routine to determine the accuracy of SPIDER. We have no absolutely known pulse shape (if we already knew the pulse shape, we would not need SPIDER!). Rather than use an absolutely accurate known temporal phase and intensity, we use a temporal phase and intensity measured under optimal conditions as our reference. We define optimal conditions as interferograms recorded with the highest possible SNR and the averaging of at least 50 signals. This condition is similar to how the SPIDER device is calibrated, thus leading to good accuracy.[22]

The obvious question then arises: Is this a good way to characterize the input pulse? Namely, will averaging noisy interferograms ultimately lead to the correct pulse shape, or are there systematic errors within the SPIDER algorithm that prevent an accurate reconstruction? According to the numerical studies, the following appears to be the case. There is a baseline error for perfectly noiseless interferograms introduced by the finite resolution of the detection system. That error amounted to a temporal envelope mismatch of roughly 0.15% and a temporal phase mismatch of roughly 0.007 rad, both exceedingly minute quantities.

So in the presence of noise, inevitably introduced by any experimental system, what happens if we average multiple interferograms? A well-known result in statistical error analysis says that the precision of a measured quantity will improve as the square root of the number of measurements.[30] Thus, averaging multiple noisy interferograms before the inversion routine should improve the precision of the reconstructed pulse shape. In the numerical studies, that was indeed found to be the case. In fact, a fit to our data for errors versus interferogram averaging has a nearly perfect $1/\sqrt{N}$ dependence if a baseline error is included (our estimate of this baseline error is close to the minimum errors reported earlier). Thus it appears that, in our experiments, use of a pulse reconstructed from a high-SNR, multiply averaged interferogram will suffice as our reference pulse profile. We then compare the reconstructed pulse from a noisy interferogram to this reference pulse. In essence, we are measuring the precision of SPIDER, relying on the demonstration of previous studies of SPIDER's accuracy.[19–23,25–27]

B. Apparatus

Our experimental apparatus is shown in Fig. 1. Our light source, a mode-locked Ti:sapphire laser,[31] generated 50-fs pulses at 80 MHz with an average output power of approximately 100 mW. The other elements of our system included a collimating telescope, an external-cavity paired-prism pulse compressor (for imparting negative chirp to selected pulses), fused-silica plates (for imparting positive chirp to selected pulses), an automated beam alignment system, and a SPIDER apparatus. The SPIDER apparatus followed a design by Dorrer[32] except that a type I β-barium borate crystal was used to eliminate the need for a polarization-rotating half-wave plate in the pulse-stretching arm.

The experimental protocol is as follows. An ultrashort optical pulse enters the SPIDER apparatus through a pair of alignment pinholes. The incident pulse strikes an etalon at near-normal incidence. The two reflections from the front and back sides of the flat glass constitute two nearly identical copies of the original pulse (the slight differences in intensity and chirp do not cause significant problems). Only a

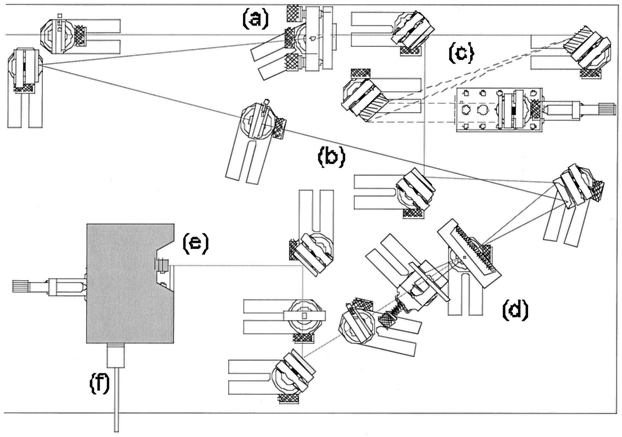

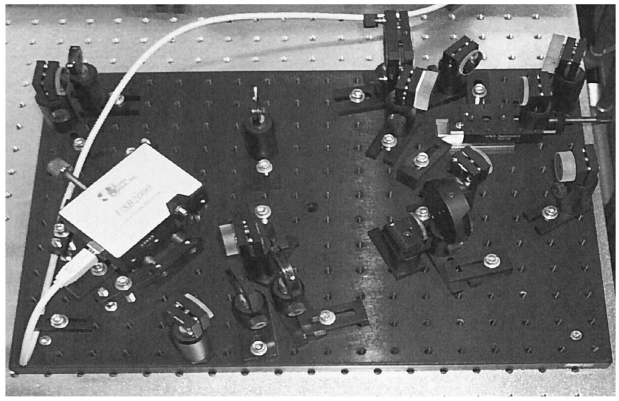

(g)

Fig. 1. SPIDER apparatus (a) An ultrashort laser pulse enters through a pinhole in the top left corner and strikes an etalon where it is separated into three pulses. (b) Two of these pulses travel down one arm of SPIDER to the type I β-barium borate crystal. (c) The third pulse travels to a pair of diffraction gratings and is temporally stretched before meeting the pulse pair at the crystal. (d) The crystal spectrally shears the pulse pair and doubles their frequencies before they enter the spectrometer. (e) Inside the spectrometer the pulses spread out and interfere with each other. (f) The interference pattern, or interferogram, is detected by an array of photodiode detectors and exported to a computer for mathematical analysis. (g) The entire apparatus fits on a bread board that is 12 in. (31 cm) wide by 18 in. (46 cm) long.

temporal shift, by a time τ, distinguishes the two. The temporal delay depends on the thickness of the etalon and the pulse's angle of incidence. The temporal delay for our SPIDER was 1.77 ps. This corresponds to approximately 20 fringes across the FWHM of our pulse, well within the optimal operating window as defined in Ref. 21. The light that propagates through the etalon is sent into a grating-pair stretcher (actually a compressor in the typical ultrafast jargon, but equivalent to a stretcher with opposite sign of chirp). The diffraction gratings (1200 lines/mm blazed at 750 nm) impart a large second-order dispersion to this pulse. In our particular SPIDER the second-order dispersion equaled 348 kfs^2.

The pulse pair is mixed with the stretched pulse in a type I β-barium borate nonlinear crystal, where

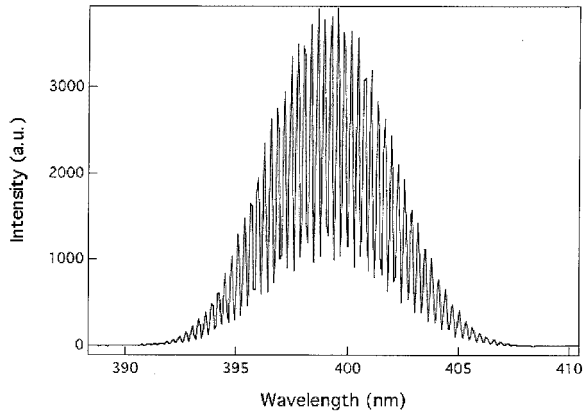

Fig. 2. Measured interferogram with the highest possible SNR.

upconversion takes place. Because we are using a type I crystal, we employ a noncollinear phase-matching geometry to minimize doubling of the individual pulses. With a pinhole, we can further isolate the mixed signal. Because the stretched pulse is long relative to the durations of the individual pulses, the upconversion process produces two blue pulses that replicate the phase of the incident red pulse, separated in time by τ. And because the stretched pulse is chirped, the two blue pulses have a spectral shear Ω between them, given by the relationship between τ and the amount of chirp. For our experiments, Ω was approximately 1/20 the FWHM bandwidth of the incident pulse.

The spectrometer then records the interference pattern between these two pulses. We employed an Ocean Optics 2000-USB spectrometer with a grating blazed at 750 nm, 1200 lines/mm, and a 2048 array of photodiode detectors in which each detector had 12-bit resolution. This spectrometer yielded a spectral resolution of better than 0.25 nm. The spectral interference pattern of light and dark bands created by SPIDER, which the spectrometer's photodiodes detect as an array of periodically varying intensity levels, forms an interferogram or record of the interference. The interferogram contains complete phase information for the input pulse. An example of an interferogram obtained with our system is shown in Fig. 2. The SPIDER algorithm then extracts the spectral phase information. However, it should be remembered that the interferogram is made from blue pulses; our original pulse was near infrared. Thus we have to convert the information at frequency ω back to frequency $\omega/2$. Furthermore, the spectrometer is linear in wavelength, not frequency, so our algorithm interpolates to convert $S(\lambda)$ to $S(\omega)$.

There are also two calibration steps that need to be performed to complete the pulse reconstruction. The first is to measure the delay τ. This is achieved when we look at the interferogram that results from the two pulses reflecting off the etalon. Namely, we block the stretched arm of SPIDER, remove the blue filter, and let some near-infrared light leak into the

spectrometer. We extract τ from the fringe spacing of the resultant interferogram. This light is also used to record the fundamental spectrum. By Fourier transforming this interferogram, keeping instead the $t = 0$ term, and transforming back, we obtain the spectrum (and spectral amplitude). These steps are performed before the real-time SPIDER begins operating (although there is a way to measure the fundamental spectrum dynamically by letting some near-infrared light leak into the spectrometer.[33])

C. Noise Sources

There are two detrimental effects we wish to consider. The first is additive noise primarily due to electronic noise in the detector. In the case of low-light levels, this can significantly degrade the quality of the interferogram. The noisy interferogram can be modeled by

$$S_{\text{noisy}}(\omega_i) = S_{\text{input}}(\omega_i) + \sigma_{\text{elec}}(\omega_i), \tag{9}$$

where $S_{\text{noisy}}(\omega_i)$ is the resultant noisy interferogram, ω_i is the optical frequency of pixel element i, $S_{\text{input}}(\omega_i)$ is the normalized ideal (low-noise) interferogram for the input pulse, and $\sigma_{\text{elec}}(\omega_i)$ is the electronic noise due to the detection system.

The second detrimental effect is due to the finite (gray-scale) resolution of the detector itself. By artificially quantizing the resultant interferogram, we can simulate the effects of a low-bit-depth detector. The interferogram is modeled by

$$S_{\text{quant}}(\omega_i) = \frac{1}{(2^{\text{bit depth}} - 1)} \langle (2^{\text{bit depth}} - 1) S_{\text{input}}(\omega_i) \rangle, \tag{10}$$

where $S_{\text{quant}}(\omega_i)$ is the quantized interferogram, bit depth is the detector resolution, and the brackets $\langle\rangle$ represent rounding to the nearest integer. Because $S_{\text{input}}(\omega_i)$ is normalized to a peak value of unity, this procedure gives us $2^{\text{bit depth}}$ gray-scale levels in our interferogram.

D. Sample Preparation

To prepare our interferogram samples, we first performed a background subtraction to eliminate the effects of room light and inconsistencies between bias levels and sensitivities in the spectrometer's detectors. We next truncated our data array to 1024 elements; our measured interferograms occupied only approximately 600 elements of the spectrometer, so we could easily perform the truncation without losing any information from the signal. From the 1024-element interferogram, the SPIDER inversion algorithm extracted a 512-element spectral phase. To convert the calculated spectral phase and the spectral intensity (obtained directly in a separate calibration measurement) into a temporal phase and temporal intensity required only a simple fast Fourier transform. We padded the spectral phase and intensity arrays with zeros prior to executing this fast Fourier transform to ensure that our temporal phase and intensity arrays had 1024 elements.

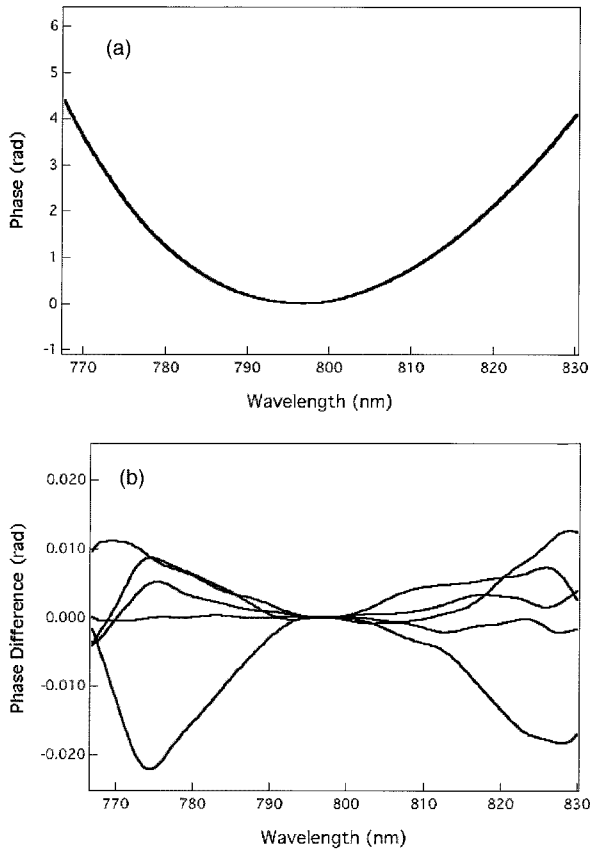

Fig. 3. Five separate measurements of the spectral phase of an ultrashort laser pulse. (a) The five measurements are indistinguishable on this scale. (b) The difference between each individual phase and the average of all five phases.

E. Intensity and Phase Errors

When operating under high-SNR conditions, the SPIDER apparatus processes extremely precise measurements with only minute random errors. For example, Fig. 3(a) shows five different spectral phase measurements made by SPIDER of the same pulse train under optimal conditions; the similarity of these precise measurements renders them virtually indistinguishable. To see the difference between phases, Fig. 3(b) is a plot of the difference between each phase and the average phase, plotted over the bandwidth of the pulse. The vertical scale indicates the minute differences between measurements, a testament to SPIDER's repeatability.

To quantify the precision of SPIDER we adopt the method introduced by Fittinghoff *et al.*[34] with a slight modification following from above. (The rms error estimate introduced by Dorrer and Walmsley[22,29] has unique benefits; however, we chose the following metric to provide a direct comparison with previous research on both FROG[34] and SPIDER.[21]) We calculated the temporal intensity error ε_I according to

$$\varepsilon_I = \left\{ \frac{1}{N} \sum_{j=1}^{N} [I_{\text{ref}}(t_j) - I(t_j)]^2 \right\}^{1/2}, \tag{11}$$

where N indicates the total number of elements in the array (1024); j represents the index of the jth pixel in the array; $I_{\mathrm{ref}}(t_j)$ is the reference temporal intensity for the jth pixel, as determined by use of optimum conditions; and $I(t_j)$ is the measured temporal intensity of the jth pixel. The array $I(t_j)$ was normalized such that it equaled unity at its maximum value. Therefore ε_I nearly equals the percent error of the temporal intensity. Because the SPIDER algorithm returns a temporal pulse that has no defined absolute temporal position or peak height, we scaled and shifted $I(t_j)$ to minimize ε_I. We accomplished this by walking the reference and measured temporal intensities past one another in steps of one pixel. At each step we calculated ε_I. When the two intensity peaks were properly aligned, the error was minimized.

We could not calculate a percentage error for the temporal phase because this would be undefined when the input phase was zero; instead, we used a rms error. In addition, the temporal phase error was undefined for times at which the temporal intensity was zero. Consequently, the rms phase error was weighted by the temporal intensity such that, for times when the temporal intensity was nonexistent, the contribution to the phase error was zero. The phase error ε_ϕ was calculated according to

$$\varepsilon_\phi = \frac{\left\{\frac{1}{N}\sum_{j=1}^{N} I^2(t_j)[\phi_{\mathrm{ref}}(t_j) - \phi(t_j)]^2\right\}^{1/2}}{\left[\frac{1}{N}\sum_{j=1}^{N} I^2(t_j)\right]^{1/2}}, \qquad (12)$$

where N equals the total number of elements in the input array from our detector (1024), j represents the index of the jth pixel of the detector, $I(t_j)$ is the measured temporal intensity of the jth pixel, $\phi_{\mathrm{ref}}(t_j)$ indicates the reference temporal phase for the jth pixel, and $\phi(t_j)$ is the measured phase for the jth pixel. Unlike the percentage error of the temporal intensity, the error in the temporal phase has units of radians.

4. Experimental Results

We measured the performance in three specific cases; we analyzed the dependence of the errors on the noise in the sample, the resolution of the detector, and the number of signals averaged prior to running the inversion routine.

In our experiments involving noise, we consider only additive noise in our system although the numerical study looked at additive and multiplicative noise. There are two reasons for this. Multiplicative noise, such as that arising from slight differences in the biasing level of our spectrometer photodiodes, was easily removed from the experiment and consequently had no effect on our precision; second, there was no practical way to introduce reproducible, controllable, and quantifiable multiplicative noise into our experimental system.

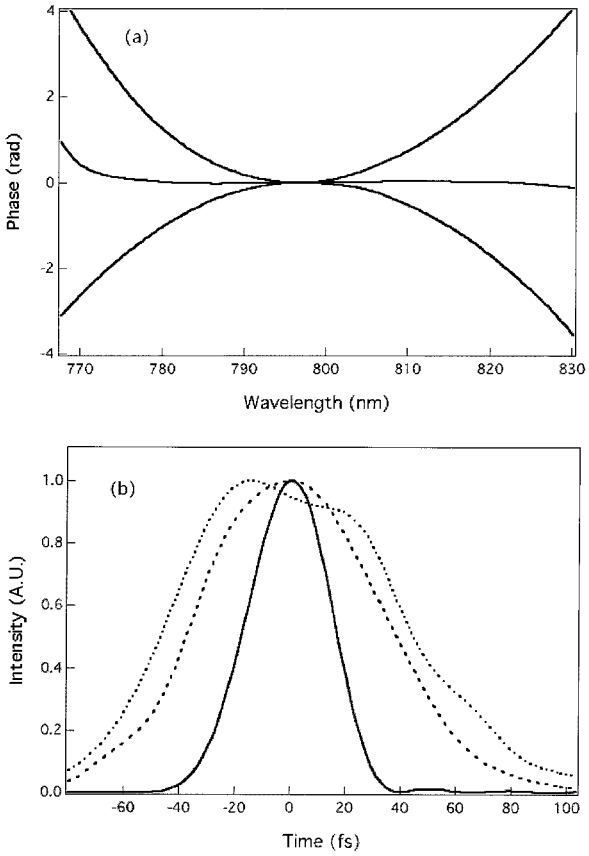

Fig. 4. Three pulses used in this study. (a) The spectral phases were positive quadratic, flat, and negative quadratic. (b) The corresponding temporal profiles for positive quadratic (dashed curve), flat (solid curve), and negative quadratic (dotted curve). (The positive quadratic pulse actually contained some cubic terms as well, which leads to the interesting temporal profile.)

A. Temporal Error Versus Noise in the Interferogram

We employed three different readily available pulse shapes for the noise studies: pulses that were near the transform limit, pulses that had a positive chirp, and pulses that had a negative chirp. The near-transform-limited pulses were obtained directly from the laser. To create pulses with positive chirp, the laser pulses were passed through 5.0 cm of fused silica prior to entering SPIDER. And to create pulses with negative chirp, the laser pulses were sent through an external-cavity prism-pair compressor. The spectral phases for the three different pulse types are shown in Fig. 4(a), and their temporal intensities are shown in Fig. 4(b).

To accomplish our first task, determining the temporal error as a function of noise in the interferogram, we first had to determine a relationship between the real additive noise in our system and the theoretical noise fraction used in computational studies. In the numerical study,[21] the additive noise was modeled by $(\alpha)(n)/\eta_i$, where α is the noise fraction and η_i is a pseudorandom number taken from a Poisson distribution of mean value n. (Note that in the numerical studies, $n = 5$.) By setting the rms value of the

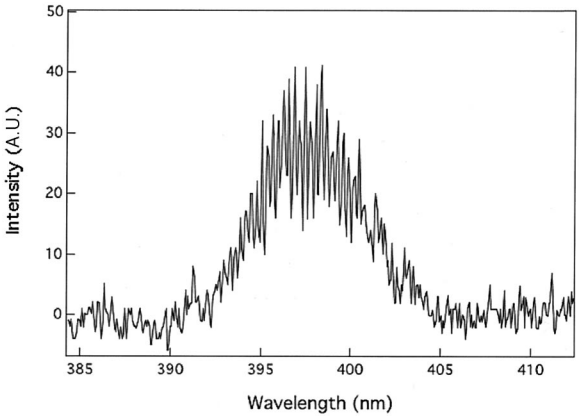

Fig. 5. Measured interferogram with 10% noise.

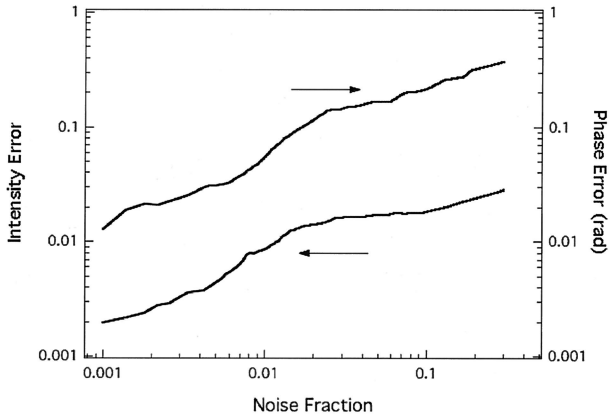

Fig. 6. Average error in reconstructed temporal intensity (lower curve) and phase (upper curve) versus noise fraction.

electronic noise equal to the theoretically modeled noise,

$$\mathrm{rms}(\sigma_{\mathrm{elec}}) = \mathrm{rms}\left(\frac{\alpha}{n}\,\eta_i\right), \tag{13}$$

we determined the noise fraction α. The noise fraction was inversely proportional to the measured SNR and scaled linearly with the peak intensity of the interferogram. A typical interferogram with 10% noise fraction is shown in Fig. 5. Clearly this appears as a rather poor-looking interferogram. This is primarily due to the limited fringe resolution that makes this noisy interferogram appear to have a much lower SNR than it actually does.

Having established a protocol for determining the temporal errors and having defined the noise fraction, the experiment proceeded rather straightforwardly. We began with a maximum SNR and determined the temporal intensity errors and the temporal phase errors as the noise was increased. Initially we increased the noise fraction by decreasing the integration time on the spectrometer. When the integration time reached its lowest value (3 ms), we increased the neutral density in our beam to further diminish the signal. For each input pulse, we took 100 measurements at each noise fraction and report the average of those values. (The error is averaged over all three pulse shapes.)

The results for this first part of the experiment are shown in Fig. 6, where we plotted the average intensity and phase errors versus noise fraction. The noise fraction was increased from 0.001 (low noise) to 0.3 (high noise). The intensity error ranged from approximately 0.2% at the low end to 3% at the high end. The phase error ranged from approximately 0.015 rad at the low end to 0.4 rad at the high end. It also appears that both errors increase roughly as the square root of the noise fraction.

These results agree well with those found in the numerical simulations. The slopes of the lines and relative positions of the intensity and phase errors are similar. Our results indicate slightly higher magnitudes of error, however. This is not unexpected when a simulated interferogram and an actual interferogram are compared. The theoretical fringe depth was almost 100% whereas the actual fringe depth in our experiment is closer to 60%. This arises from several factors, including slight differences in the energies of the interfering pulses, imperfect overlap of the interfering pulses, and the finite width of a detector pixel compared with the width of a fringe (there are approximately 4.5 pixels per fringe). Because the experimental fringe depth is just over half of the peak intensity, the SNR is actually approximately twice as high as calculated; this trend is reflected in the data of Fig. 6.

To get a better visual idea of its performance, Fig. 7 shows an actual pulse shape calculated during this run of experiments. The reference temporal intensity and phase are given by the dashed curves. The measured temporal intensity and phase with 10% noise are given by the solid curves. The temporal intensity error is approximately 2% and the temporal phase error is approximately 0.04 rad.

B. Temporal Errors Versus Detector Resolution

In the second part of our experiment we looked at the temporal errors as a function of the resolution of the

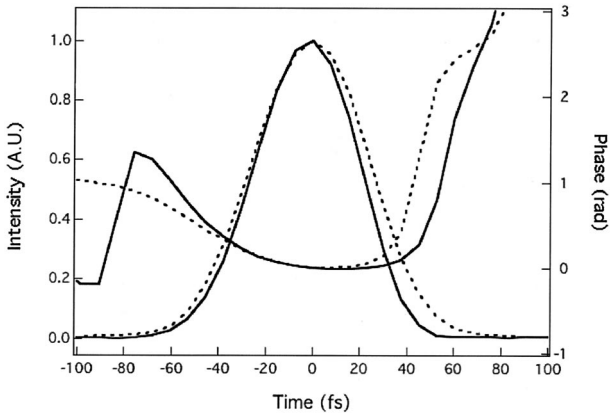

Fig. 7. Reconstructed pulse intensity and phase from interferograms recorded under optimum conditions (dashed curves) and with 10% noise fraction (solid curves).

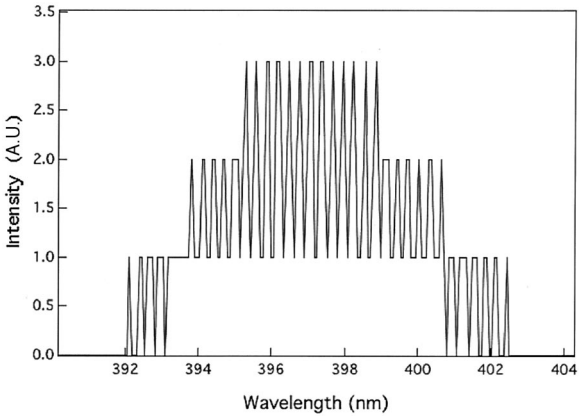

Fig. 8. Interferogram rounded to 2 bits (four levels).

photodiode detectors. Because the spectrometer had a fixed resolution, we artificially lowered the resolution by digital processing in the computer according to Eq. (10). All interferograms used for rounding were acquired under optimal conditions. Figure 8 shows an interferogram rounded to four levels. After the levels were rounded, the interferograms were processed by the SPIDER algorithm. The temporal intensity and phase errors were calculated according to Eqs. (11) and (12).

Again, for each pulse shape at each resolution, 100 trials were run for a total of 300 trials per resolution level. The average value of the measurements is shown in Fig. 9. The intensity errors range from roughly 3% at low quantization to 0.2% at high quantization. The phase errors range from 0.5 rad down to 0.004 rad. These results are close to those previously obtained by numerical simulations. However, in this case, our measured phase errors for high resolutions are slightly smaller than those found in the numerical study. The reason is that the reference phase and intensity we used are calculated in the same manner as the measured values; in the computational study, processed to unprocessed temporal profiles are compared. Inevitably there will be small errors in the processed phase, particularly because the interference fringes are only sampled at a small integer number of points. Because our reference phase goes through the same processing as the measured phase, it is not surprising that it is more similar to our reference than a computer-determined phase is to an actual, unprocessed phase.

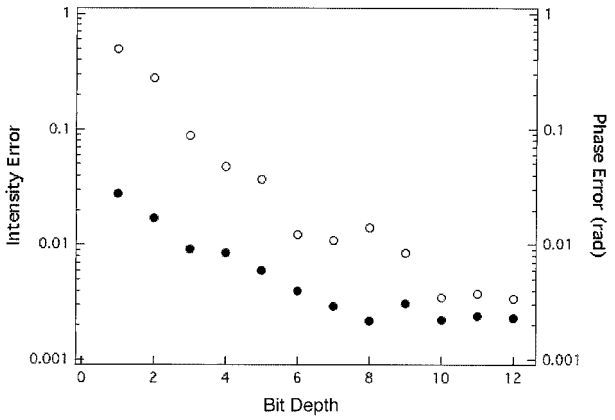

Fig. 9. Average error in reconstructed pulse intensity (solid circles) and phase (open circles) versus quantization bit depth.

The most important information recoverable from Fig. 9 is that the temporal errors remain at the same level for 10-bit and higher resolution. (The intensity error appears to bottom out around 8 bits.) A spectrometer with 12 bits performs exactly as well as a spectrometer with 10 bits of resolution. For detectors with resolutions of less than 10 bits, the error varies inversely with the number of levels in the detected interferogram. And it appears that good reconstructions are still recoverable from interferograms with resolutions as low as 3 and 4 bits.

In fact, as pointed out previously,[21] it is possible to recover reasonably precise phase information from a 1-bit interferogram. This counterintuitive feature is due, in essence, to an interesting benefit of the retrieval algorithm. Because the primary facet that defines the spectral phase is the spacing of the fringes in the interferogram, a 1-bit interferogram with the proper fringe spacing will still yield a correct spectral phase. Of course, to reconstruct the complete pulse, knowledge of the spectrum is required with high resolution. However, in many experimental situations (such as the aligning of stretchers or compressors), the spectrum does not change, and the spectral phase measurement provides the necessary information. For example, a flat spectral phase indicates an optimized compressor.

Because one is often interested in speeding up the inversion algorithm and the entire SPIDER reconstruction,[24] a 1-bit interferogram is quite useful. Specifically, the computer can execute Fourier transforms on an array of 1-bit numbers much faster than it can on an array of 12-bit numbers. A 1-bit interferogram might also be easier to record with low signal systems by use of high-gain amplifiers such as avalanche photodiode detectors. The problem with a 1-bit interferogram is twofold when our rounding scheme is used; first, all the fringes with an amplitude less than half of the maximum value are lost. This usually means that almost 50% of the fringes are lost. Second, especially for an experimental interferogram in which the fringe depth is only slightly larger than 50%, the apparent width of the fringes will vary greatly. This is apparent in Fig. 10 where a dashed horizontal line marks the 1-bit round off. Near the center of the interferogram the 1-bit fringes will alternate between thin and thick.

Instead of simple rounding, another possible way to generate a 1-bit interferogram is to threshold with a suitable Gaussian. Signal levels above the Gaussian become unity, those below become zero. Of course this requires knowledge of the spectrum to properly define this Gaussian threshold, but as argued above, this can be accomplished with one measurement as part of the calibration steps. The benefits to this rounding scheme are then rather ob-

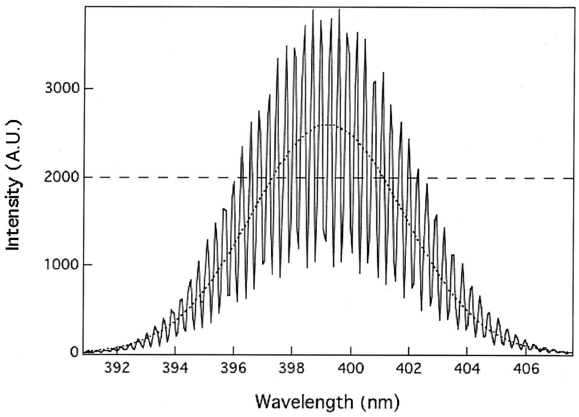

Fig. 10. Measured interferogram (solid curve), 1-bit round off (dashed horizontal line), and Gaussian threshold (dotted curve).

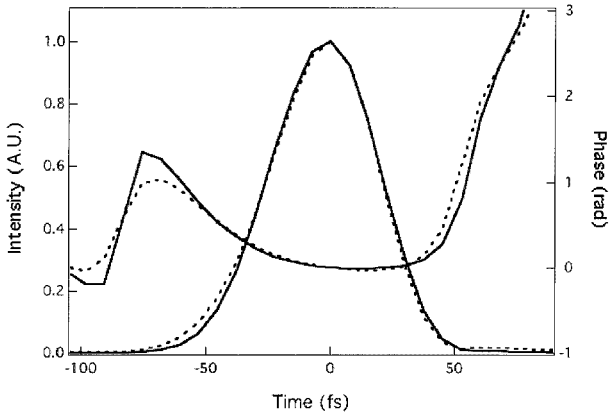

Fig. 12. Reconstructed pulses from an optimal interferogram (solid curves) and the Gaussian threshold interferogram of Fig. 11(a) (dashed curves).

vious. Because all the fringes will cross the rounding line at their middles, there will be no distortion of fringe width. A Gaussian threshold for the interferogram of Fig. 10 is shown as the dotted curve arcing through the center of the interferogram. Figure 11(a) shows the interferogram of Fig. 8 rounded to a 1-bit interferogram. Figure 11(b) shows the same interferogram converted to a 1-bit interferogram by the Gaussian threshold technique. Clearly, this approach yields many more fringes than its bit-rounded counterpart.

We analyzed the error in 1-bit interferograms created by the rounding and thresholding. For each type of 1-bit interferogram we measured the average of 300 trials, 100 with each pulse shape, and all recorded under optimal conditions. The results show the temporal phase error in the thresholded 1-bit interferogram to be 75% lower than in the rounded 1-bit interferogram!

Equally impressive, the temporal intensity error was 67% lower in the thresholded 1-bit interferogram than the rounded 1-bit interferogram. Figure 12 shows a pulse reconstructed from a thresholded 1-bit interferogram [like Fig. 11(b)] as solid curves and a regular 12-bit interferogram, recorded under optimal conditions, as dotted curves. These graphs are surprisingly close considering that one of them was obtained from a 1-bit interferogram; the temporal intensity error for the reconstructed pulse of Fig. 12 is 0.5% and the temporal phase error is 0.02 rad.

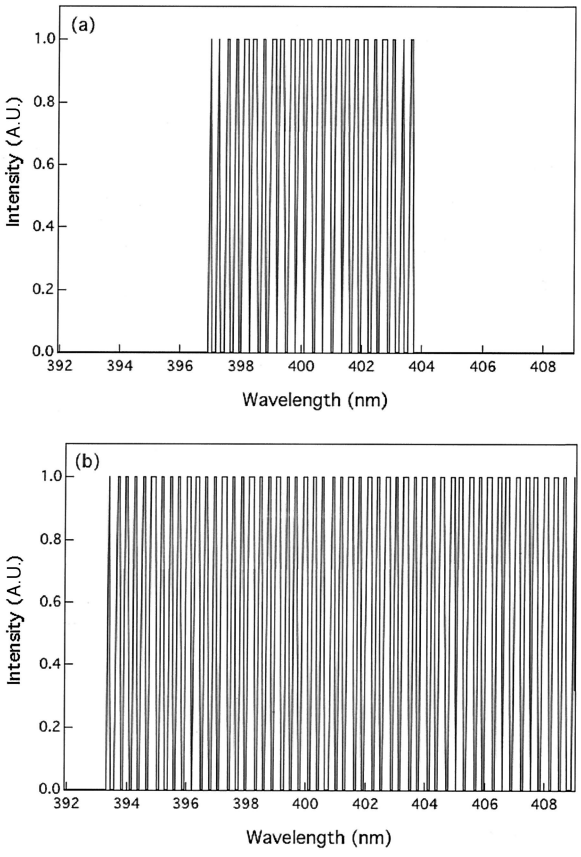

Fig. 11. Two techniques used to retrieve 1-bit interferograms: (a) bit rounding and (b) Gaussian thresholding.

C. Temporal Error Versus Averaging

In the third part of our experiment we analyzed the effect of improving pulse reconstructions by averaging interferograms with high initial noise levels. If the noise is truly random additive noise, averaging multiple noisy interferograms should uncover the buried signal. Conversely, if there are nonrandom (systematic) effects present in these cases of low SNR that were not present in the reference measurement case of high SNR, averaging multiple interferograms would not improve the signal.

To simplify the experimental procedure, we proceeded in two steps. First, we recorded 1000 interferograms with a noise fraction between 15 and 20%. Second, we looked at the effects of averaging by choosing random interferograms from this pool of 1000. We chose one interferogram from this group at random and calculated its intensity and phase errors. We then chose another interferogram at random and calculated its intensity and phase errors.

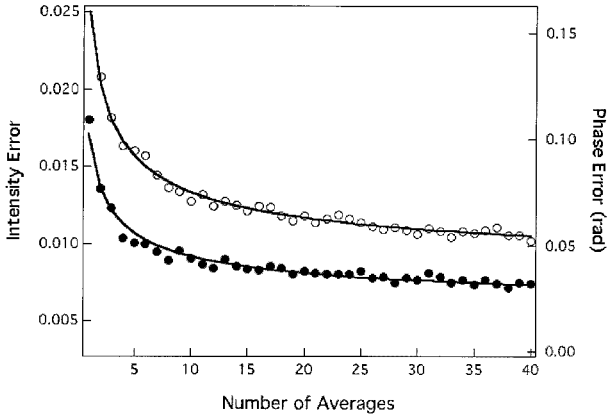

Fig. 13. Average error in reconstructed pulse intensity (solid circles) and phase (open circles) versus number of interferogram averages. The solid curves are fits to Eq. (14).

We repeated this 100 times and averaged the resultant intensity and phase errors. This gave us values for the errors when no averaging is present. Then we randomly selected two interferograms, averaged them, ran them through the inversion algorithm, and calculated the intensity and phase errors. We did this 100 times to obtain an average value of the errors that exists when two interferograms are averaged prior to running the inversion algorithm. We proceeded in this manner until we obtained error data for up to 40 interferograms averaged prior to processing.

Our results appear in Fig. 13 where we plotted the errors as a function of interferogram averaging. The data are indicated by the circles and the curves indicate least-square fits of the functional form

$$\varepsilon = \gamma + \beta/\sqrt{N}, \tag{14}$$

where γ represents the baseline error and β is a scaling factor. For the intensity error, $\gamma = 0.0055$ and $\beta = 0.0116$. For the phase error, $\gamma = 0.034$ and $\beta = 0.131$. These results indicate baseline errors $\varepsilon_I \approx 0.55\%$ and $\varepsilon_\phi \approx 0.034$ rad, which are slightly higher but still in reasonable agreement with those found in the numerical studies. Considering that our noise levels were higher, and our fringe visibility was limited, these baseline errors represent decent marks for SPIDER's performance. Furthermore, these results show that, averaging as few as four or five noisy signals, we can reduce the reconstruction error by 50% or more.

Signal averaging may be important for low-light applications when, even at maximum integration times, the signal strength is still low. To determine just how low a signal could be measured with our incarnation of SPIDER, we used neutral-density filters to reduce our incident beam strength to 1.1 mW. With our detector on the highest possible integration time, 65535 ms, we recorded an interferogram and found that it had a temporal intensity error of only 0.9% and a temporal phase error of 0.08 rad. The noise fraction for this signal was 1.7%; this corresponds to a peak intensity of 300 units whereas the peak intensity under optimal conditions is close to 4000 units. A spectrometer with a longer integration time could record a higher peak intensity and, consequently, a lower noise fraction and reconstruction error.

5. Conclusions

We have demonstrated that a real-time, compact, working SPIDER device is capable of excellent precision measurements even in cases of excessive noise. We have demonstrated the measurement of 50-fs pulses from a mode-locked Ti:sapphire laser in three different dispersion regimes and verified many theoretical predictions of SPIDER's performance based on computer modeling. We have shown that, even with 10% noise in its signal, SPIDER can reproduce a pulse's temporal profile with only 2% error in the temporal intensity and 0.2 rad of error in the temporal phase. We found that a SPIDER device need not use a detector with greater than 10 bits of resolution. We also confirmed that a signal detected with a 1-bit detector can reconstruct the temporal intensity to within 2% and the temporal phase to within 0.2 rad of their nominal values. We verified that, when four signals are averaged prior to application of SPIDER's retrieval algorithm, the error in the reconstructed temporal phase and intensity drop by approximately 50% from their levels for a single measurement. We also found that SPIDER accurately measures pulse trains with very low average powers; for example, SPIDER determined the temporal intensity and phase of a pulse train with an average power of only 1.1 mW to within 0.9% and 0.08 rad, respectively.

These same types of test could be repeated for more complicated pulse shapes, which might prove a greater challenge to the retrieval algorithm, especially when noisy. We do not anticipate complex pulses creating problems, however; computational studies have used both simple phases and complicated phases containing quadratic, cubic, and quartic terms with similar errors reported for all.[21,22,29] We have reproduced the results for the simple cases easily and do not anticipate any major problems with complicated phases. (The cases we studied were the only ones experimentally available to us.) It would also be interesting to see the resolution studies performed with detectors that actually have low resolutions rather than artificially lowering the resolutions of better detectors. We would especially be interested to see if it is possible to reproduce the thresholded 1-bit interferogram results using masks or filters laid directly on the detector array, instead of carrying it out purely by digital computer processing.

This study confirmed many predictions made by the numerical simulations investigated previously. The interesting component, however, is in the application to a real laboratory setting. It was not entirely clear how the theoretical modeling would translate to a real device with real limitations (for example, the finite spectrometer resolution leading to an interferogram with limited visibility). Fortu-

nately, SPIDER appears robust even under extreme experimental circumstances.

We gratefully acknowledge Christophe Dorrer for the SPIDER design, programming assistance, and for many helpful discussions.

References

1. P. T. Wilson, Y. Jiang, O. A. Aktsipetrov, E. D. Mishina, and M. C. Downer, "Frequency-domain interferometric second-harmonic spectroscopy," Opt. Lett. **24,** 496–498 (1999).
2. H. Kapteyn and M. Murnane, "Ultrashort light pulses: life in the fast lane. Researchers can now routinely use lasers to generate femtosecond light pulses in the lab, making it possible to study processes ranging from photosynthesis to the switching of electronic circuits. How the latest advances could lead to the generation of even shorter pulses," Phys. World **12,** 31–35 (1999).
3. J.-M. Hopkins and W. Sibbett, "Ultrashort-pulse lasers: big payoffs in a flash-imaging, microelectronic manufacturing, fiber optics and industrial chemistry are eagerly adopting lasers that emit light in powerful bursts lasting only quadrillionths of a second," Sci. Am. **283,** 72–79 (2000).
4. C. W. Siders, S. P. L. Blanc, D. Fisher, T. Tajima, M. C. Downer, A. Babine, A. Stepanov, and A. Sergeev, "Laser wake-field excitation and measurement by femtosecond longitudinal interferometry," Phys. Rev. Lett. **76,** 3570–3573 (1996).
5. D. Meshulach and Y. Silberberg, "Coherent quantum control of two-photon transitions by a femtosecond laser pulse," Nature (London) **396,** 239–241 (1998).
6. I. Walmsley and R. Trebino, "Measuring fast pulses with slow detectors," Opt. Photon. News **7,** 23–28 (1996).
7. D. J. Kane and R. Trebino, "Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating," IEEE J. Quantum Electron. **29,** 571–579 (1993).
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. G. Taft, A. Rundquist, M. M. Murnane, H. C. Kapteyn, K. W. DeLong, R. Trebino, and I. P. Christov, "Ultrashort optical waveform measurements using frequency-resolved optical gating," Opt. Lett. **20,** 743–745 (1995).
10. K. W. DeLong, R. Trebino, and W. E. White, "Simultaneous recovery of two ultrashort laser pulses from a single spectrogram," J. Opt. Soc. Am. B **12,** 2463–2466 (1995).
11. D. N. Fittinghoff, J. L. Bowie, J. N. Sweetser, R. T. Jennings, M. A. Krumbugel, K. W. DeLong, R. Trebino, and I. A. Walmsley, "Measurement of the intensity and phase of ultraweak, ultrashort laser pulses," Opt. Lett. **21,** 884–886 (1996).
12. C. W. Siders, J. L. W. Siders, F. G. Omenetto, and A. J. Taylor, "Multipulse interferometric frequency-resolved optical gating (invited paper)," IEEE J. Quantum Electron. **35,** 432–440 (1999).
13. R. Trebino, P. O'Shea, M. Kimmel, and X. Gu, "Measuring ultrashort laser pulses just got a lot easier," Opt. Photon. News **12,** 22–25 (2001).
14. M. A. Franco, H. R. Lange, J. F. Ripoche, B. S. Prade, and A. Mysyrowicz, "Characterization of ultra-short pulses by cross-phase modulation," Opt. Commun. **140,** 331–340 (1997).
15. J. W. Nicholson, J. Jasapara, W. Rudolph, F. G. Omenetto, and A. J. Taylor, "Full-field characterization of femtosecond pulses by spectrum and cross-correlation measurements," Opt. Lett. **24,** 1774–1776 (1999).
16. K. C. Chu, J. P. Heritage, R. S. Grant, K. X. Liu, A. Dienes, W. E. White, and A. Sullivan, "Direct measurement of the spectral phase of femtosecond pulses," Opt. Lett. **20,** 904–906 (1995).
17. I. G. Cormack, W. Sibbett, R. Ortega-Martinez, and D. T. Reid, "Ultrashort pulse characterization using a scanning Fabry-Perot etalon enabling rapid acquisition and retrieval of a sonogram at rates up to 1.52 Hz," Rev. Sci. Instrum. **72,** 4071–4079 (2001).
18. L. Lepetit, G. Cheriaux, and M. Joffre, "Linear techniques of phase measurement by femtosecond spectral interferometry for applications in spectroscopy," J. Opt. Soc. Am. B **12,** 2467–2474 (1995).
19. C. Iaconis and I. A. Walmsley, "Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses," Opt. Lett. **23,** 792–794 (1998).
20. C. Iaconis and I. A. Walmsley, "Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses," IEEE J. Quantum Electron. **35,** 501–509 (1999).
21. M. E. Anderson, L. E. E. de Araujo, E. M. Kosik, and I. A. Walmsley, "The effects of noise on ultrashort-optical-pulse measurement using SPIDER," Appl. Phys. B **70,** S85–S93 (2000).
22. C. Dorrer and I. A. Walmsley, "Accuracy criterion for ultrashort pulse characterization techniques: application to spectral phase interferometry for direct electric field reconstruction," J. Opt. Soc. Am. B **19,** 1019–1029 (2002).
23. T. M. Shuman, M. E. Anderson, J. Bromage, C. Iaconis, L. Waxer, and I. A. Walmsley, "Real-time SPIDER: ultrashort pulse characterization at 20 Hz," Opt. Exp. **5,** 134–143 (1999), http://www.opticsexpress.org.
24. W. Kornelis, J. Biegert, J. W. G. Tisch, M. Nisoli, G. Sansone, C. Vozzi, S. D. Silvestri, and U. Keller, "Single-shot kilohertz characterization of ultrashort pulses by spectral phase interferometry for direct electric-field reconstruction," Opt. Lett. **28,** 281–283 (2003).
25. C. Dorrer, B. de Beauvoir, C. LeBlanc, S. Ranc, J. P. Rousseau, P. Rousseau, J. P. Chambaret, and F. Salin, "Single-shot real-time characterization of chirped-pulse amplification systems by spectral phase interferometry for direct electric-field reconstruction," Opt. Lett. **24,** 1644–1646 (1999).
26. L. Gallmann, D. H. Sutter, N. Matuschek, G. Steinmeyer, U. Keller, C. Iaconis, and I. A. Walmsley, "Characterization of sub-6-fs optical pulses with spectral phase interferometry for direct electric-field reconstruction," Opt. Lett. **24,** 1314–1316 (1999).
27. C. Dorrer, "Influence of the calibration of the detector on spectral interferometry," J. Opt. Soc. Am. B **16,** 1160–1168 (1999).
28. S. Yeremenko, A. Baltuska, M. S. Pshenichnikov, and D. A. Wiersma, "The criterion of pulse reconstruction quality based on Wigner representation," Appl. Phys. B **70,** S109–S117 (2000).
29. C. Dorrer and I. A. Walmsley, "Precision and consistency criteria in spectral phase interferometry for direct electric-field reconstruction," J. Opt. Soc. Am. B **19,** 1030–1038 (2002).
30. P. R. Bevington and D. K. Robinson, *Data Reduction and Error Analysis for the Physical Sciences*, 2nd ed. (McGraw-Hill, Boston, Mass., 1992).
31. See, for example, Kapteyn-Murnane Laboratories at www.kmlabs.com.
32. C. Dorrer, Lucent Technologies, Murray Hill, N.J. 07974-0636 (personal communication, 2001).
33. C. Dorrer, "Implementation of spectral phase interferometry for direct electric-field reconstruction with a simultaneously recorded reference interferogram," Opt. Lett. **24,** 1532–1534 (1999).
34. 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).