Appl. Phys. B 70 [Suppl.], S85–S93 (2000) / Digital Object Identifier (DOI) 10.1007/s003400000316

Applied Physics B
Lasers
and Optics

Invited paper

The effects of noise on ultrashort-optical-pulse measurement using SPIDER

M.E. Anderson*, **L.E.E. de Araujo, E.M. Kosik, I.A. Walmsley**

The Institute of Optics, University of Rochester, Rochester, NY 14627, USA
(Fax: +1-716/244-4936, E-mail: slambo@optics.rochester.edu)

Received: 2 November 1999/Revised version: 12 March 2000/Published online: 24 May 2000 –

Abstract. We have studied the accuracy of ultrashort-pulse reconstruction using SPIDER (spectral phase interferometry for direct electric-field reconstruction) with noisy data. Specifically, we have looked at the effects of additive noise, multiplicative noise, and quantization of the interferogram on the reconstruction of the intensity and phase, and discovered that the inversion routine is relatively insensitive to noise. In particular, with 10% additive noise, SPIDER is able to recover the optical phase to an accuracy of approximately 0.04 radians and recover the intensity pulse shape to an accuracy of approximately 1.5%. Further, we have identified the optimal parameters for pulse reconstruction from noisy data with which SPIDER should operate.

PACS: 07.60.Ly; 42.65.Re; 42.30.Rx

In the last few years, progress in the measurement of ultrashort optical pulses has led to the development of several techniques that can be used to characterize laser pulses completely, easily, and rapidly: FROG [1], SPIDER [2], DOSPM [3], ENSTA [4], XPM [5], MIFROG [6], IAC-Spectrum [7], etc. Several of these methods have been known to yield the complete time-dependent electric field, usually given as amplitude and phase, of pulses approaching a few fs in duration [8, 9]. Such information is an invaluable tool for experimentalists in a variety of applications, not least of which is the optimization of the laser system itself. Indeed, complete information is critical for laser designers hoping to control higher order dispersion and reach the single-cycle regime. But full characterization is likewise important for applications such as quantum control, where the spectral or temporal components and their corresponding phases play a critical role in controlling, for example, the production rates of chemical reactions [10] or the rate of two-photon ionization [11].

Of critical importance to the usefulness of any measurement tool is how well it performs in real laboratory settings. It is therefore a useful exercise to examine the merits of an ultrafast measurement scheme under less-than-ideal conditions. Important questions are: How well can it reconstruct the field in the presence of noise? How well does it perform with limited precision? What are the optimum parameters for the technique in the presence of noise?

In this paper we address these questions in the context of spectral phase interferometry for direct electric-field reconstruction (SPIDER) [2]. The SPIDER method is a self-referencing interferometric technique for measuring ultrashort optical pulses, and is characterized by being conceptually straightforward and easy to implement. Moreover, it uses a noniterative inversion algorithm. SPIDER is extremely flexible: It can be used over a wide range of wavelengths, from infrared to blue [12]; it can measure pulse durations from roughly a ps down to less than 6 fs [9]; it can measure pulses from amplifiers or oscillators [2, 13]; and it can run in real time with 20-Hz update rates [13] or in a single-shot configuration [14].

Our study is modeled after the thorough investigation of the performance of frequency-resolved optical gating (FROG) in the presence of noise presented in a paper by Fittinghoff et al. [15]. We have adopted their strategies for modeling additive noise, multiplicative noise, and quantization, and for measuring the fidelity of pulse reconstruction. (Others have recently suggested different fidelity measures. See [16].)

Aside from testing the robustness of the inversion algorithm in the presence of noise, we also address the issue of optimizing the parameters involved in configuring SPIDER (namely the pulse delay, the spectral shear, and the window width). In this paper we will not address systematic errors such as beam misalignment, inaccurate SPIDER calibration traces, or spatial chirp. Other sources of error, such as miscalibration of the spectrometer, have been studied recently, and it was shown that SPIDER is insensitive to these errors [17]. It is important to note that we also assume that the spectrum of the input pulse can be measured to high accuracy since it requires no nonlinear process. In this study we measure the fundamental spectrum to 8-bit precision, although 6 bits is also adequate. Therefore, the noise is applied only to the interferogram itself, which is consistent with experimental implementations of SPIDER.

The rest of this paper is organized as follows. Section 1 will present the SPIDER inversion routine and a description

*Corresponding author.

of its experimental implementation. Section 2 will describe the five different pulse shapes we will consider for this investigation. Section 3 will address the issue of optimizing the parameters for the SPIDER inversion. Section 4 will explain the procedures used for simulating noise in this study. Section 5 will present the results of the noise simulations. Section 6 will discuss other issues such as one-bit quantization, interferogram averaging, and a qualitative discussion of the reconstruction. Finally Sect. 7 will offer some concluding remarks.

1 SPIDER apparatus and inversion routine

The SPIDER apparatus and inversion algorithm are fairly simple to execute, and they are described in this section. We will briefly outline the implementation of upconversion SPIDER (for measuring infrared pulses) to provide a framework for the discussion that follows. However, it should be kept in mind that these results are general and apply equally well to any incarnation of the SPIDER method. (For a much more detailed description of SPIDER, see [2].) We will examine the inversion routine in detail as this provides important insights into the role of noise and the parameter adjustment.

SPIDER is an embodiment of the technique of spectral-shearing interferometry, in which two pulsed fields of identical temporal shape but different center frequency are spectrally resolved [18, 19]. The key ingredient for spectral-shearing interferometry to work with ultrafast pulses is the ability to generate a shear of the required magnitude. Typically, ultrafast pulses have bandwidths of many THz and the required shear is of the order of a few THz. Since current state-of-the-art electro-optic and acousto-optic modulators cannot shift carrier frequencies by this amount, SPIDER uses nonlinear optics to achieve the necessary shear. In the upconversion mode of SPIDER, two replica pulses are combined with a third stretched (chirped) pulse and upconverted in a type-II second-harmonic crystal. The resulting two second-harmonic pulses have a spectral shear that is dictated by the time delay between the incident pulse pair and the amount of chirp on the stretched pulse. This method provides the required spectral shear, which is in fact adjustable over a wide range from essentially zero to the available bandwidth in the incident pulse. The two second-harmonic pulses are then spectrally resolved in a spectrometer and an interferogram is recorded. This interferogram contains the phase information, encoded into the positions of the peaks and valleys in the interferogram. Extraction of this information is direct using Fourier analysis.

An apparatus for implementing upconversion SPIDER in the laboratory is shown in Fig. 1. The setup provides a method of generating the pulse pair and a means for stretching a third replica. The input pulse is first split into two beams. One beam is sent through a device that will create two identical pulses, separated in time by τ. This can be achieved either with a Michelson interferometer, or from the first and second facet reflections from a flat uncoated etalon (the multiple reflections that follow are considerably lower in energy and may be neglected). The other beam is sent through a stretcher, either a grating pair stretcher or, in the case of very short pulses (<10 fs), a piece of glass. These two orthogonally polarized beams are then recombined and focused into a type-II nonlinear crystal. The upconverted pulses propagate to a spectrometer, at the output of which is a 512-element photodiode array. The signal from the photodiode array is read and analyzed by a computer.

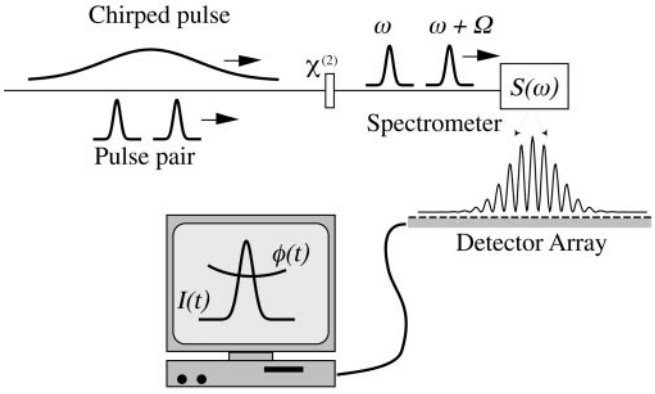

Fig. 1. Experimental apparatus for upconversion SPIDER. The chirped pulse and pulse pair are upconverted in a $\chi^{(2)}$ nonlinear crystal. The resulting sheared pulses are measured with a spectrometer and detector array. The measured interferogram is sent to the computer for analysis

The interferogram $S(\omega)$ has a form given by

$$S(\omega) = I(\omega) + I(\omega+\Omega) + 2\sqrt{I(\omega)I(\omega+\Omega)} \times \cos\{\phi(\omega+\Omega) - \phi(\omega) + \omega\tau\}, \quad (1)$$

where $I(\omega)$ is the input power spectrum, Ω is the spectral shear, and τ is the pulse delay. The phase information we wish to recover is $\phi(\omega)$. One can see that knowledge of the spectral shear and the pulse delay is key to recovering this spectral phase. Both of these quantities are easily measured in the laboratory [2]. For the purposes of this study, interferograms are simulated by sampling $S(\omega)$ calculated from $I(\omega)$ and $\phi(\omega)$ using (1) on a 512-element array of real numbers.

The inversion routine is shown in Fig. 2. The frequency-domain interferogram is first Fourier transformed to a pseudo-time domain. This procedure yields a function of three peaks (in time), separated by the pulse delay τ. The peak at $t = +\tau$ is selected using a "super-duper" Gaussian window (an eighth-order super Gaussian). This peak is then inverse-Fourier transformed back to the frequency domain. The argument of this function is the spectral phase of the interferogram (which is discontinuous, due to the principal-value calculation [20], but can be easily unwrapped). We subtract the linear $\omega\tau$ term (obtained in the laboratory from a separate [2] or simultaneous [14] calibration trace), leaving the phase difference between two different frequency components of the input pulse,

$$\theta(\omega) = \phi(\omega) - \phi(\omega+\Omega). \quad (2)$$

To recover the phase $\phi(\omega)$ a concatenation procedure is used. Setting the phase at the center frequency to zero (it is impossible to measure the "absolute" phase using SPIDER), we can write the phase at a discrete set of frequencies $\{(\omega_c \pm n\Omega)$, n an integer number$\}$ as

$$\vdots$$
$$\phi(\omega_c - 2\Omega) = \theta(\omega_c - 2\Omega) + \theta(\omega_c - \Omega)\,,$$

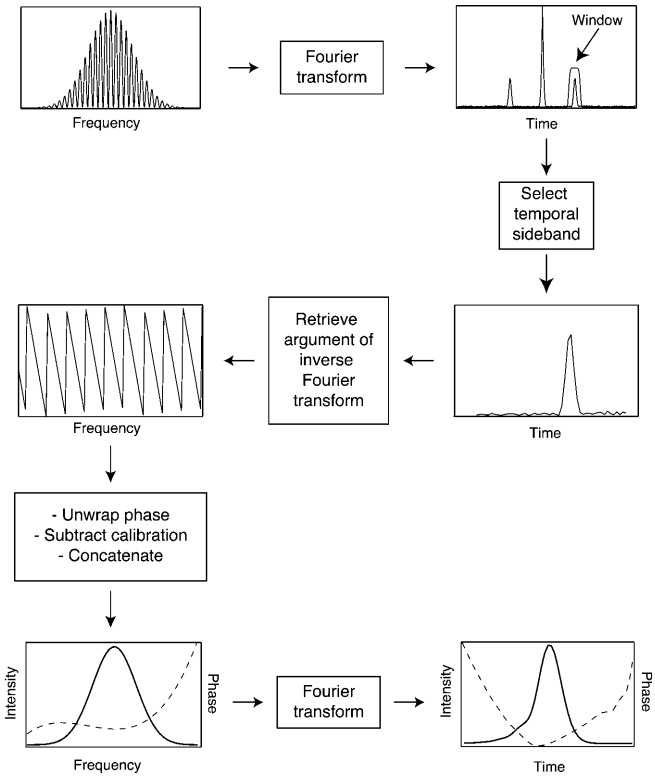

Fig. 2. SPIDER inversion algorithm. The intensity and phase (in frequency or time) can be reconstructed from the measured interferogram by means of simple Fourier transforms

$$\begin{aligned}
&\phi(\omega_c - \Omega) = \theta(\omega_c - \Omega)\,,\\
&\phi(\omega_c) = 0\,,\\
&\phi(\omega_c + \Omega) = -\theta(\omega_c)\,,\\
&\phi(\omega_c + 2\Omega) = -\theta(\omega_c + \Omega),
\end{aligned} \tag{3}$$

where ω_c is the "center" frequency of the pulse spectrum. The measured phase differences $\theta(\omega)$ are summed to recover the true spectral phase $\phi(\omega)$ at points separated by the spectral shear. Obviously, the choice of this spectral shear is important for the reconstruction, a topic discussed in a later section.

The final step of the SPIDER inversion algorithm is returning to the temporal domain. Now that we have the spectral phase, we merely need the spectrum $I(\omega)$ to recover the complete pulse information. The temporal electric field is simply given by the Fourier transform of the spectral electric field, namely

$$E(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \mathrm{d}\omega \left\{\sqrt{I(\omega)}\mathrm{e}^{\mathrm{i}\phi(\omega)}\right\}\mathrm{e}^{\mathrm{i}\omega t}. \tag{4}$$

There is one point that needs consideration, however. The spectrum and spectral phase are defined on a 512-element array which typically covers a spectral range four times the bandwidth of the input power spectrum $\Delta\omega$. This means that a direct Fourier transform will yield a temporal pulse of resolution $\approx 1/(4\Delta\omega)$. Although the sampled spectrum provides a complete specification of the pulse, this is not high enough temporal resolution to give an intuitive time-domain picture for some pulse shapes. Thus we adopt the usual method of padding the spectral electric-field array with zeros before taking the Fourier transform. We generally increase the array to 2048 elements before going to the time domain. This gives an adequate temporal resolution even for fairly complicated pulse shapes.

2 Test pulse shapes for reconstruction simulations

Now that we have described SPIDER's operation, let us turn our attention to the five different pulse shapes we wish to consider in our noise studies. Since SPIDER naturally returns the spectral phase of the pulse, we will consider pulses that differ by their input spectral phase. This is also consistent with the most commonly used pulse-shaping apparatuses, such as the devices that use a spectrally resolved liquid-crystal [21] or acousto-optic [22] modulator. These pulse shapers operate directly on the spectral phase and amplitude of the pulse.

The pulses we will consider are shown in Fig. 3a–e, where we have plotted the temporal intensity and phase. The inset of each graph shows the corresponding spectral intensity and spectral phase. All five pulses have Gaussian spectrums with FWHM given by 10.3 THz ($\Delta\omega = 65 \times 10^{12}$ rad/s, $\Delta\lambda =$ 24 nm), corresponding to a transform-limited temporal duration (FWHM) of 43 fs. The phases are given by a fourth-order polynomial:

$$\phi(\omega) = a(\omega-\omega_c)^2 + b(\omega-\omega_c)^3 + c(\omega-\omega_c)^4, \tag{5}$$

where the coefficients a, b, and c indicate the quadratic, cubic, and quartic components of the spectral phase (the zeroth order term is the so-called "absolute" phase and the first-order term gives rise to a time shift of the pulse, neither of which is recoverable from a SPIDER measurement). The first pulse (3a) has a flat phase (transform-limited pulse, $a = b =$

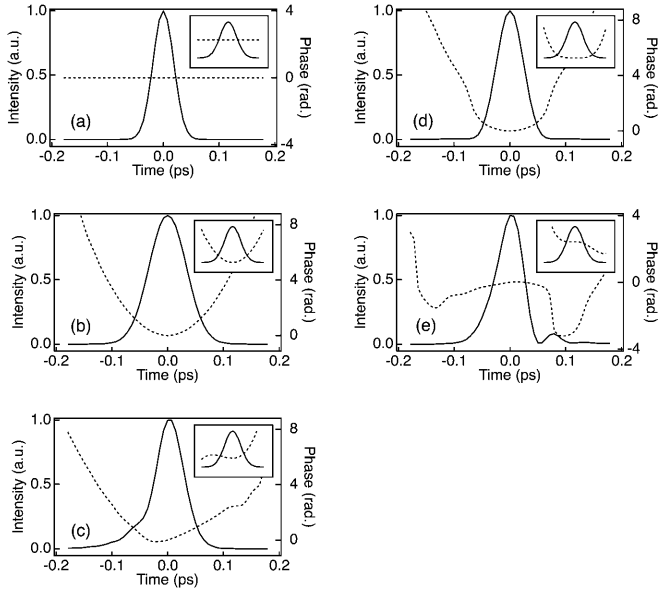

Fig. 3a–e. Pulse shapes used in this study. Graphs show intensity (*solid line*) and phase (*dashed line*) as a function of time (inset shows corresponding spectrum and spectral phase). All pulses have a Gaussian spectrum with: **a** zero phase; **b** quadratic phase; **c** quadratic and cubic phase; **d** quartic phase; and **e** quadratic, cubic and quartic phase

$c = 0$), the second (Fig. 3b) is quadratic ($a = 5 \times 10^{-4}\,\mathrm{ps}^2$, $b = c = 0$), the third (Fig. 3c) is quadratic plus cubic ($a = 1 \times 10^{-4}\,\mathrm{ps}^2$, $b = 5 \times 10^{-6}\,\mathrm{ps}^3$, $c = 0$), the fourth (Fig. 3d) is quartic ($a = b = 0$, $c = 6 \times 10^{-8}\,\mathrm{ps}^4$), and the fifth (Fig. 3e) is quadratic plus cubic plus quartic ($a = -3 \times 10^{-4}\,\mathrm{ps}^2$, $b = -1 \times 10^{-5}\,\mathrm{ps}^3$, $c = 9 \times 10^{-8}\,\mathrm{ps}^4$). This complicated phase on the fifth pulse leads to interesting behavior in the temporal domain as is seen in Fig. 3e. The intensity shows a strong second lobe and a large phase jump. It is this last pulse which will be the benchmark for the robustness of the SPIDER inversion.

3 Choice of optimum parameters

There are two free parameters in the SPIDER apparatus and one in the inversion algorithm that require careful consideration. The first parameter is the pulse delay τ. The requirements on this parameter are twofold. First, τ must be large enough that the resulting interferogram has a decent number of interference fringes (we typically operate with 20 to 30 fringes). Second, τ must be small enough that the resolution in the spectrometer and detector array is able to satisfy the Nyquist sampling criterion [23]. Namely, the detection scheme must be able to "resolve" the interferogram, with greater than two points per fringe (in the lab one would arrange the spectrometer and photodiode array to record about 15 points per fringe).

To find a good value for τ, we ran our numerical simulation of SPIDER under the following conditions. We added 10% noise to our interferograms (as described in the next section) for five different input pulse shapes (as described in the previous section) and ran the SPIDER inversion routine. We reconstructed the pulse shapes in time, compared them to the original input pulse shapes, and calculated the average error in temporal intensity $\bar{\varepsilon}_\mathrm{I}$ (also discussed later). We then changed the pulse pair delay τ and repeated the error calculation. Since the noise was random on each realization, we ran the calculation repeatedly and averaged the ensuing pulse shape errors. The results are shown in Fig. 4, where we have plotted the average temporal intensity error $\bar{\varepsilon}_\mathrm{I}$ (in percentage) versus pulse delay τ. There is a broad minimum from 0.2 ps to 6 ps, with the error increasing at either extreme. This finding is consistent with what we would expect from the Nyquist criterion: For large τ the spectral resolution is not high enough to record two points per fringe and for small τ the number of fringes drops dramatically, leading to the intermingling of the Fourier transform peaks at $t = +\tau$ and $t = 0$.

The second free parameter is the spectral shear Ω. This shear is determined by the overlap of the chirped pulse and the pulse pair. For example, in an experimental apparatus using a grating pair stretcher, the shear would be given by $\Omega = -\tau/(2\phi_2)$, where ϕ_2 is the second-order dispersion in the stretcher arm [2, 24]. Although Ω is dependent on the pulse delay τ, it can be adjusted *independently* since the chirped pulse can be stretched to an arbitrary value (ϕ_2 can be chosen at will). The choice for Ω is a little trickier than the choice for τ, since the shear must be chosen such that the requirements of the Whittaker–Shannon sampling theorem are met [23]. The situation is as follows. According to Whittaker–Shannon, only a pulse that is contained entirely within a time

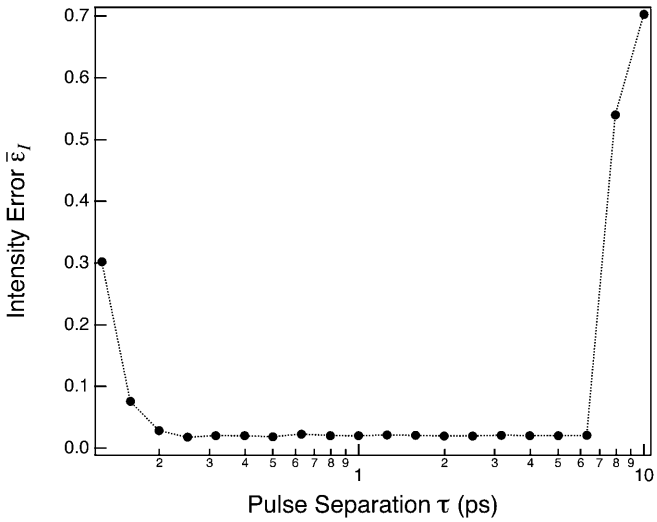

Fig. 4. Error in temporal intensity reconstruction as a function of pulse pair separation

interval $T \leq 2\pi/\Omega$ can be accurately reconstructed. Accordingly, one should choose Ω such that T is considerably longer than the input pulse duration τ_p. Generally, we use a spectral shear such that T is at least ten times greater than the transform-limited input pulse duration, i.e. $T \geq 10\tau_\mathrm{p}$. However, as discussed earlier, the reconstructed spectral phase is recovered only at spectral phase points separated by the spectral shear Ω. Hence a choice for Ω that is too small will likely lead to errors in the reconstruction in the presence of noise, since the difference between the spectral phases at $\phi(\omega)$ and $\phi(\omega + \Omega)$ will tend to zero as Ω decreases. In this case only pulses whose spectral phase varies greatly with ω would be reconstructed. For complicated pulse shapes, this can lead to errors. If Ω is too small, the interferogram's fringes will not vary much with respect to the nominal ($\approx 1/\tau$) spacing. In the presence of noise, this will lead to reconstruction errors.

To find an optimal spectral shear we ran the algorithm as described above, but this time with a fixed τ and a variable spectral shear. The results are shown in Fig. 5, where we have plotted the average temporal intensity error $\bar{\varepsilon}_\mathrm{I}$ versus relative spectral shear $\Omega/\Delta\omega$ (given in terms of the fraction of the input pulse spectrum). Also, we stepped the shear in integer multiples of pixel frequency separation since the phase is reconstructed at steps of the shear. The plot shows a fairly flat region from $\Omega/\Delta\omega \approx 4\%$ to $\Omega/\Delta\omega \approx 35\%$ with a minimum at $\Omega/\Delta\omega \approx 23\%$. This corresponds to support in the temporal domain of 425 fs, which is roughly ten times the duration of our (zero phase) input pulse, consistent with the Whittaker–Shannon sampling theorem. The error increases at higher spectral shears because the support in the time domain becomes shorter than the actual pulse. The error also increases at lower spectral shears, due to the noise. In our noise studies, we chose a spectral shear of 13% which is small enough to give accurate reconstruction of our most complicated test pulses.

The third and final parameter that can be optimized is in the inversion algorithm itself: The temporal width of the window applied to the transform of the interferogram (see Fig. 2). This window must be centered at one of the sidebands $\pm\tau$, but its width may be chosen independently. The choice is rather

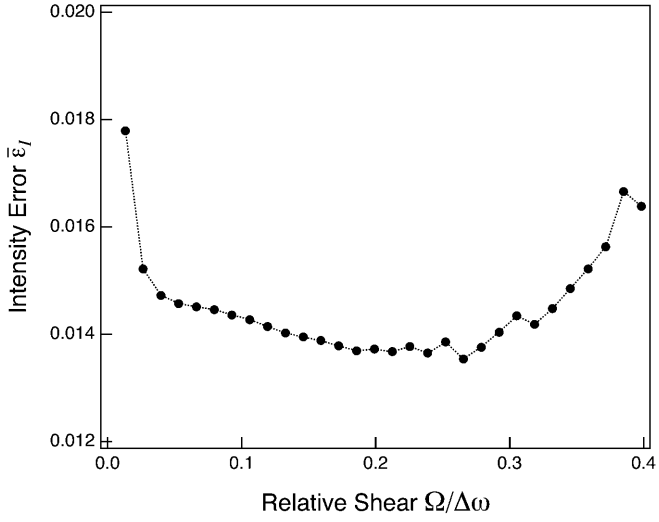

Fig. 5. Error in temporal intensity reconstruction as a function of relative shear between the uopconverted pulses

intuitive, however. It should be narrow enough to not overlap the strong peak at $t = 0$, but broad enough to completely enclose the peak at $t = +\tau$. If the window is too narrow, the peak near $t = +\tau$ will be distorted, and the fidelity of reconstruction will decrease. On the other hand, if the window is too wide, the total noise power may be sufficient to reduce the recontruction fidelity. To test this, we ran the SPIDER inversion as described above with varying window width. The results of these calculations are shown in Fig. 6 as plots of the temporal intensity error $\bar{\varepsilon}_I$ versus window width. There is a broad minimum in the error which rises sharply at either extreme of the window width, i.e. in the cases where the window either includes the $t = 0$ term or cuts off part of the $t = +\tau$ term. One might expect the optimal window width to depend on the signal-to-noise ratio of the interferogram, yet surprisingly we have found it to be fairly independent of the noise value. This indicates that the choice for window width is not critical; we typically operate with a width $\approx \tau/3$.

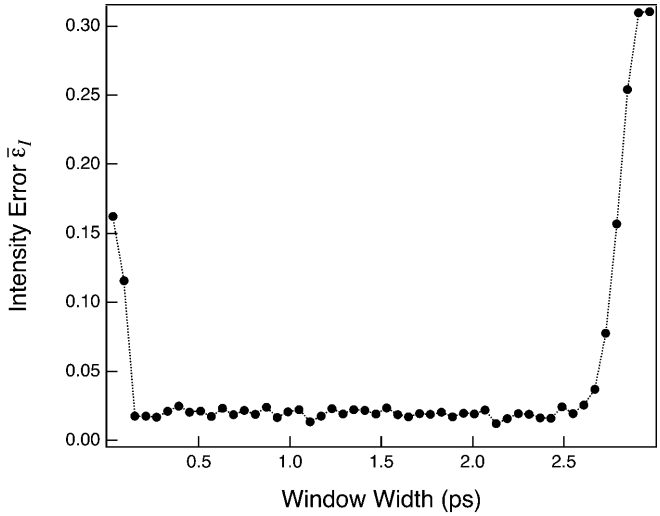

Fig. 6. Error in temporal intensity reconstruction as a function of window width in selection of temporal sideband in the SPIDER inversion algorithm

4 Sources of error: noise and quantization

There are three sources of error we consider for this study: additive noise, multiplicative noise, and quantization error. Additive noise can arise from electronic or thermal noise in the detector. Multiplicative noise is indicative of pixel-to-pixel gain variation in the photodiode array. Quantization error comes about because of the limited precision of the detection system (photodiode array and A/D board). Since the interferogram is the primary measured entity in SPIDER, it is this signal to which we apply the noise and quantization discretization.

Additive noise is the first source of error we consider. To simulate it, we simply add a pseudorandom number to each element of our interferogram and run this noisy interferogram through the SPIDER algorithm. Following Fittinghoff et al. [15], we use noise that is Poisson distributed since this resembles thermal noise and has the benefit of being positive definite (we do not expect negative voltages from our photodiodes). We are using a 512-pixel array, and the resulting noisy interferogram has a form given by

$$S^{\text{add}}_{\text{noisy}}(\omega_i) = S_{\text{input}}(\omega_i) + \frac{\alpha}{n}\eta_i, \tag{6}$$

where $S^{\text{add}}_{\text{noisy}}(\omega_i)$ is the interferogram resulting from additive noise, ω_i is the optical frequency associated with pixel element i, $S_{\text{input}}(\omega_i)$ is the input noiseless interferogram, α is the noise fraction, and η_i is a pseudorandom number obtained from a Poisson distribution of mean value n. In this study, $n = 5$. An example of an interferogram with 10% ($\alpha = 0.1$) additive noise is shown in Fig. 7a.

For multiplicative noise, we simply multiply the input interferogram by a pseudorandom number m_i that varies around unity by an amount given by the noise fraction α. The noisy interferogram is then given by

$$S^{\text{mult}}_{\text{noisy}}(\omega_i) = S_{\text{input}}(\omega_i)(1 + \alpha m_i), \tag{7}$$

where $S^{\text{mult}}_{\text{noisy}}(\omega_i)$ is the interferogram arising from multiplicative noise, and m_i is a pseudorandom number drawn from a Gaussian distribution with zero mean and unit variance. An example of an interferogram with 10% ($\alpha = 0.1$) multiplicative noise is shown in Fig. 7b.

Quantization error results from the limited precision of the detection apparatus, which includes the photodiode array and A/D computer board. The output voltages from the photodiodes or CCD array is discretized into one of $2^{\text{"bitdepth"}}$ levels. The speed of A/D conversion decreases as "bitdepth" increases, so that in many cases, one does not want to use 16-bit precision if it is not necessary. In fact, most CCD arrays are limited to 8 bits, which corresponds to 256 voltage levels. The resulting quantization of the interferogram will lead to errors, and it behooves us to determine how great these are. To model this effect, we begin with a noiseless interferogram and then "quantize" it to the prescribed bit depth. The new interferogram is then given by

$$S^{\text{quant}}(\omega_i) = \frac{1}{(2^{\text{bitdepth}} - 1)} \left\langle S_{\text{input}}(\omega_i)(2^{\text{bitdepth}} - 1) \right\rangle, \tag{8}$$

where $S^{\text{quant}}(\omega_i)$ is the quantized interferogram resulting from a detection system of precision bitdepth, and $\langle x \rangle$ represents

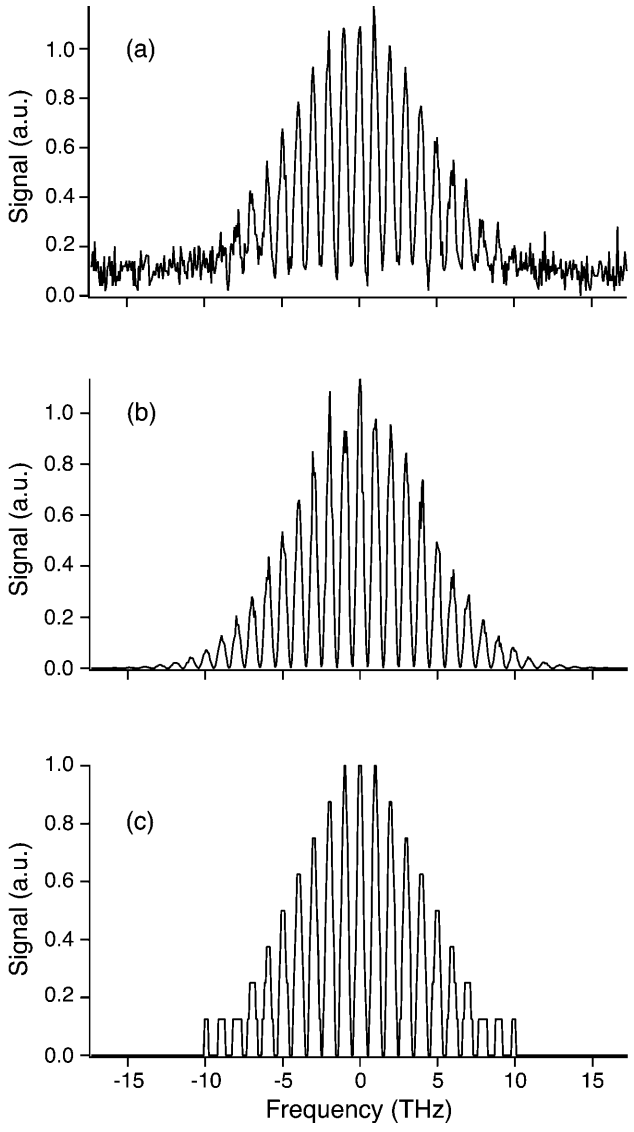

Fig. 7a–c. Examples of interferograms used in this study, with: **a** 10% additive noise, **b** 10% multiplicative noise, and **c** 3-bit quantization

x rounded to the nearest integer. The available levels are then given by 2^{bitdepth}. The prescription is to multiply the input interferogram by the maximum level $(2^{\text{bitdepth}} - 1)$ and then round this value to the nearest integer (for example, a bit depth of 2 will give us integers 0, 1, 2, 3). By dividing again by the maximum level we scale the peak of the interferogram to unity. An example of quantization is given in Fig. 7c for a bit depth of 3 (eight levels, which is unrealistic for a real detection system, but illustrates the procedure).

To quantify the errors in the recovered fields induced by level quantization, we evaluated the reconstruction fidelity of the temporal intensity and phase. In each case, we used error definitions similar to those of Fittinghoff et al. [15]. The temporal intensity error is defined by

$$\varepsilon_{\text{I}} = \sqrt{\frac{1}{N}\sum_{j=1}^{N}\left[I_{\text{input}}(t_j) - I_{\text{SPIDER}}(t_j)\right]^2}, \tag{9}$$

where N is the number of elements in our (padded) temporal list, $I_{\text{input}}(t_j)$ is the input temporal intensity for the j^{th} element of the list corresponding to time slot t_j, and $I_{\text{SPIDER}}(t_j)$ is the recovered temporal intensity from SPIDER for the j^{th} element. The input pulse was scaled so that the peak had unit amplitude. The error ε_{I} may therefore be regarded (approximately) as a percentage error. Since the SPIDER algorithm returns a temporal pulse that has no defined absolute temporal position or peak height, we have shifted and scaled $I_{\text{SPIDER}}(t_j)$ to minimize the error.

The phase error is somewhat more difficult to define, because it is impossible to define a percentage error when the input phase is zero. It is useful, therefore, to use a rms error for the phase. However, since the phase is not defined where the intensity is zero, it is not useful to simply calculate a rms phase error across the entire range of elements. Instead we multiply the phase error by the intensity at the corresponding points to construct a "weighted" error, $I(t)[\phi_{\text{input}}(t) - \phi_{\text{SPIDER}}(t)]$. The rms phase error is then given by

$$\varepsilon_{\phi} = \frac{\sqrt{\frac{1}{N}\sum_{j=1}^{N} I_{\text{input}}^2(t_j)[\phi_{\text{input}}(t_j) - \phi_{\text{SPIDER}}(t_j)]^2}}{\sqrt{\frac{1}{N}\sum_{j=1}^{N} I_{\text{input}}^2(t_j)}}, \tag{10}$$

where $I_{\text{input}}(t_j)$ and $\phi_{\text{input}}(t_j)$ are the input pulse temporal intensity and phase for the j^{th} element, respectively, and $\phi_{\text{SPIDER}}(t_j)$ is the recovered SPIDER phase for the j^{th} element. Note that the phase error has units of radians. Furthermore, there is an equivalence principle (Parseval's theorem) between the representation of the electric field in frequency and in time which mandates that the recorded errors have similar magnitudes in either representation. Thus we will report the errors in temporal phase, but find similar errors for the weighted spectral phase.

5 Results

We studied the effects of additive noise, multiplicative noise, and quantization for the five pulse shapes shown in Fig. 3. We ran the simulations for the five different pulse shapes individually, with the error at each noise fraction averaged over one hundred runs, each of which used a new noise realization. The error can vary significantly from shot to shot since the noise is random. Because the reconstruction fidelity was nearly identical for the five different pulses studied, we averaged the error from the five pulse shapes to produce "global" error figures $\bar{\varepsilon}_{\text{I}}$ and $\bar{\varepsilon}_{\phi}$. These are shown in Figs. 8–10 and give a flavor of the expected performance of SPIDER in the presence of experimental noise, regardless of pulse shape.

Figure 8 shows the intensity error $\bar{\varepsilon}_{\text{I}}$ and the phase error $\bar{\varepsilon}_{\phi}$ versus noise fraction for the average of the five pulses in the presence of additive noise. Although we averaged over one hundred noise runs, there is still some spread in error values due to the randomness of the noise. This graph gives a good indication of the range of errors one might expect with a given amount of noise. The noise fraction ranges from 0.001 (0.1% noise) to 0.5 (50% noise). The results show that SPIDER performs well over a broad range of values, with an intensity error of $\bar{\varepsilon}_{\text{I}} \approx 0.15\%$ at the lowest noise to $\bar{\varepsilon}_{\text{I}} \approx 3\%$ at the highest. The phase error ranges from $\bar{\varepsilon}_{\phi} \approx 0.07$ radians at the lowest noise to $\bar{\varepsilon}_{\phi} \approx 0.1$ radians at the highest. A good testbed

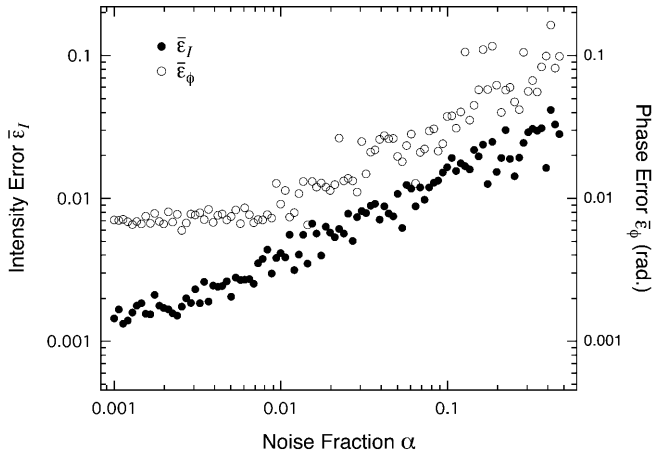

Fig. 8. Error in reconstructed temporal intensity (*solid circles*) and phase (*open circles*) versus noise fraction for the case of additive noise

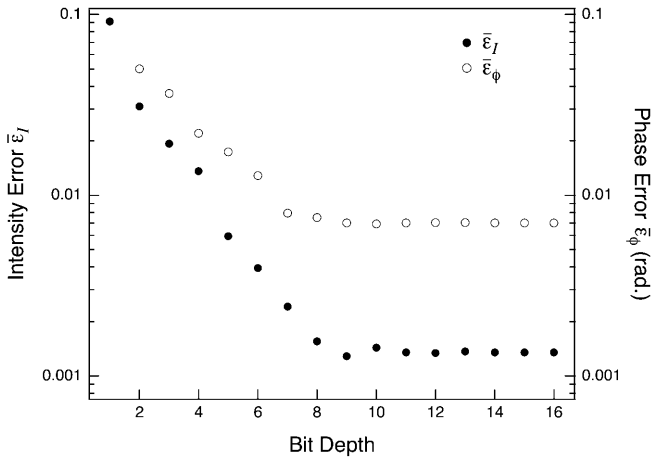

Fig. 10. Error in reconstructed temporal intensity (*solid circles*) and phase (*open circles*) versus quantization bit depth

for the performance of SPIDER is given by the error at 10% noise. Here we find that the recovered intensity is accurate to $\bar{\varepsilon}_I \approx 1.5\%$ and the phase is accurate to $\bar{\varepsilon}_\phi \approx 0.04$ radians. This is an acceptable level of accuracy for most applications.

The results for multiplicative noise are shown in Fig. 9, where we have again plotted $\bar{\varepsilon}_I$ and the phase error $\bar{\varepsilon}_\phi$ versus noise fraction for the average of the five pulses. These results indicate that SPIDER is much more resilient to this type of noise. In fact for noises up to 50% the recovered intensity is still accurate to $\bar{\varepsilon}_I \approx 1\%$. The main reason is that the signal-to-noise ratio is constant across the spectrum, so that the noise in the wings of the interferogram is significantly less than that when additive noise is used. Instead, multiplicative noise acts primarily in the central region. Since SPIDER is essentially insensitive to the amplitude of the interferogram, but is instead concerned with the spacing of the fringes, multiplicative noise causes little error in reconstruction.

The results of discretization are shown in Fig. 10. At eight-bit resolution the reconstructed intensity is accurate to $\bar{\varepsilon}_I \approx 0.15\%$ and the phase is recovered to $\bar{\varepsilon}_\phi \approx 0.07$ radians. Higher precision is unnecessary since there is no significant improvement in the errors for higher bit depths. Clearly these results show that eight-bit resolution (256 gray levels) is adequate for accurate pulse reconstruction in most cases. This is an important result, since most CCD arrays are nominally eight-bit and at this precision can operate with reasonably fast readout.

It is important to note that these results were obtained without prefiltering the raw data. The inversion routine was performed with the noisy interferograms directly. However, the SPIDER algorithm has an implicit "filtering" procedure built into it, namely the windowing of the interferogram's Fourier transform. In this study, the noise in the interferogram is taken to be uncorrelated from pixel to pixel and to be present across the entire array. As a consequence, noise in the time domain is also uniform and uncorrelated. Therefore, because the filtering occurs in the time domain, much of the noise in the interferogram is removed. It is only the noise that occupies the temporal region around the window that is detrimental.

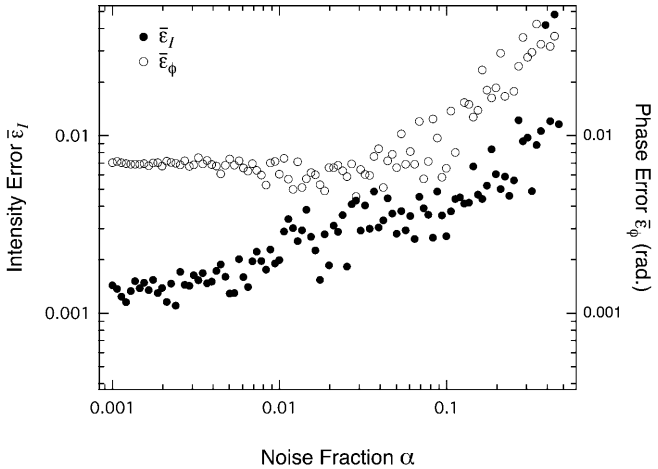

Fig. 9. Error in reconstructed temporal intensity (*solid circles*) and phase (*open circles*) versus noise fraction for the case of multiplicative noise

6 One-bit SPIDER, interferogram averaging, and pulse reconstruction

An interesting feature of Fig. 10 is that the reconstruction error for a bit depth of one is still only about 10%. This is quite startling considering that this implies that each pixel in the array be either "on" [$S(\omega_c) = 1$] or "off" [$S(\omega_c) = 0$]. There are no gray levels. This result begs explanation.

The answer lies in the power of the interferometry itself. The SPIDER algorithm is not concerned with the amplitude of the interferogram. It is concerned only with the *spacing* of the interference fringes. Provided there are a sufficient number of pixels per fringe, a one-bit signal will still have fringes that are spaced correctly. Referring back to (8), which describes the procedure for quantizing the interferogram, it is clear that a one-bit interferogram will have non-zero values across only the FWHM of the pulse spectrum since the "threshold" is set at 50%. (By threshold we mean the level above which an individual pixel turns on.) However, one can actually do much better by setting the threshold to a different level. (This is equivalent, for example, to changing the input light level.) Setting the threshold lower effectively extends the spectral range of the quantized interferogram. In fact, the threshold should

be set optimally at about twice the noise amplitude. For instance, we added 10% additive noise to our interferogram, set the threshold for the pixel "on" state to 20%, and were able to recover the temporal intensity to $\bar{\varepsilon}_I \approx 2\%$ accuracy and the phase to $\bar{\varepsilon}_\phi \approx 0.04$ radians. So pulse reconstruction with high fidelity is possible even with a one-bit detector!

Of course this is not quite the whole story: to recover the complete temporal envelope and temporal phase requires the spectrum of the incident pulse. We assume that one can measure this spectrum to moderately high accuracy (six bits or better). However, if one is only interested in the spectral phase, which is often useful for applications such as laser adjustment, then knowledge of the spectrum is not required and a one-bit detector is adequate. This is a dramatic reduction in the amount of information that needs to be collected and processed, and could lead to vast improvements in SPIDER's update rates.

In practical terms the small size of the SPIDER data set means that another avenue to improved performance is possible: averaging several interferograms. Because SPIDER has no moving components, data can be acquired very rapidly. In this study we have considered data acquisition for the noise present on a single measurement of the SPIDER interferogram. Since SPIDER is rather quick (it takes a few ms to read out a single interferogram), it is possible to numerically average several interferograms before processing in order to reduce the noise. If the noise is truly random from shot to shot, the signal-to-noise ratio of the interferogram is improved by roughly $\sqrt{M}$, where M is the number of interferograms averaged. This averaging would correspond to a slower update rate. For instance, in the demonstrated real-time SPIDER at 20 Hz [13], the data acquisition time was 10 ms and the algorithm inversion took 40 ms. Averaging five interferograms would increase the data acquisition time to 50 ms, but the inversion time stays the same. The total time is then 90 ms, corresponding to an update rate of 11 Hz. For example, consider the case of 10% additive noise. In Fig. 11 we have plotted the intensity error and phase error for 10% additive noise as a function of the number of interferograms averaged. We find that with 10% additive noise, averaging five interferograms (lowering the update rate to 11 Hz), we can lower the intensity error to $\bar{\varepsilon}_I \approx 0.8\%$ (an error reduction of nearly a factor of two) and reduce the phase error to $\bar{\varepsilon}_\phi \approx 0.014$ radians (a factor of three improvement).

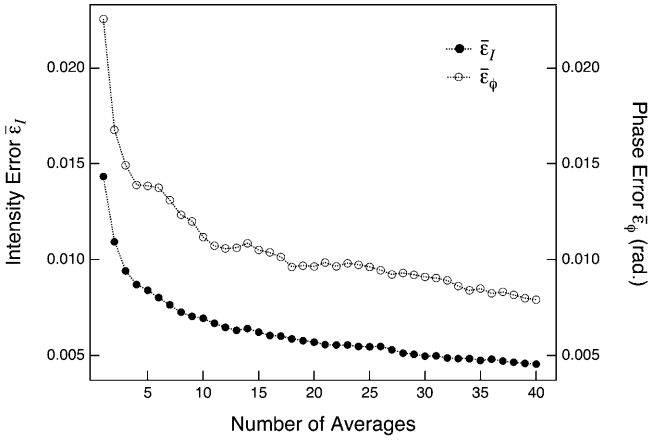

Fig. 11. Error in reconstructed temporal intensity (*solid circles*) and phase (*open circles*) versus number of interferogram averages for the case of 10% additive noise

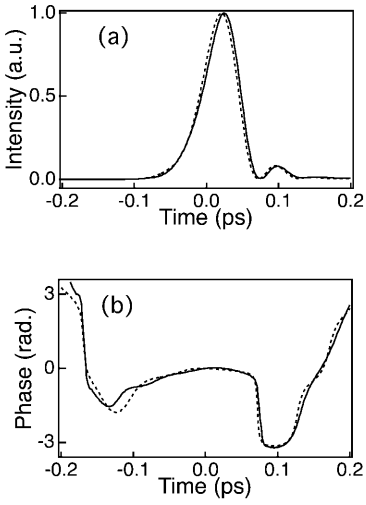

Fig. 12a,b. Reconstructed pulse shape for the pulse of Fig. 3e with 10% additive noise: **a** temporal intensity for original (*solid line*) and reconstructed (*dashed line*), **b** temporal phase for original (*solid line*) and reconstructed (*dashed line*)

While the noise results presented above are a quantitative answer to the question of SPIDER's robustness in the presence of noise, a more useful indicator is a qualitative one. What does the reconstruction "look" like? The benchmark pulse in Fig. 3e with 10% additive noise is plotted in Fig. 12, where we show the original and reconstructed intensity and phase. Visually it appears that SPIDER does quite well at recovering the profile, including the intensity node and the second lobe. The phase is also reconstructed quite well, including the phase jump. For this particular example, $\bar{\varepsilon}_I \approx 2\%$, and $\bar{\varepsilon}_\phi \approx 0.03$ radians.

7 Conclusions

In this paper we have analyzed the fidelity with which ultrashort optical pulses can be reconstructed from noisy SPIDER interferograms. We did this by simulating noisy interferograms and using the standard SPIDER algorithm to reconstruct the pulse shape. In order for SPIDER to operate effectively it is necessary to ensure that the shear Ω and pulse delay τ are correctly adjusted, and the effect of varying these parameters on accurate pulse reconstruction was also studied numerically. To optimize the parameter settings, we ran simulations to measure the reconstructed pulse accuracy with 10% additive noise. We found that the error vs. temporal pulse delay τ can be quite uniform over a wide range of values of τ, with a value of $\tau = (2\pi/\Delta\omega)(N_s/N_f)$ giving excellent results, where $\Delta\omega$ is the FWHM of the input pulse spectrum, N_s is the number of pixels across the FWHM, and N_f is the number of pixels per fringe. For the pulses we considered, $\tau = 1$–2 ps gave good results. The error with varying spectral shear Ω has a minimum at about half the Whittaker–Shannon limit of $\Omega/\Delta\omega = 1/d$, where d is the degree above transform limit of the pulse duration. For our pulses, this corresponds

to near $\Omega/\Delta\omega \approx 1/4$. Finally, the algorithm was found to be fairly insensitive to window width, with a good value given by $\approx \tau/3$.

The role of noise on pulse reconstruction in SPIDER has been studied numerically. We looked at the effects of additive noise, multiplicative noise, and quantization on the reconstructed intensity and phase. The most deleterious source of error came from additive noise, primarily due to the increasingly low signal-to-noise in the wings of the interferogram. Typically for the case of 10% noise we found that the intensity profile can be reconstructed to $\bar{\varepsilon}_I \approx 1.5\%$ and the phase can be recovered to $\bar{\varepsilon}_\phi \approx 0.04$ radians. While these results are quite acceptable in most circumstances, a much better performance can be achieved by averaging the noisy interferograms at the expense of slower update rates. For example, with 10% additive noise an average of only five interferograms improves the recovered intensity to an accuracy of $\bar{\varepsilon}_I \approx 0.8\%$ and the recovered spectral phase improves to $\bar{\varepsilon}_\phi \approx 0.014$ radians.

Acknowledgements. We would like to thank Peter Parker for insightful discussions. This work was supported by a grant from the National Science Foundation under the GOALI program.

References

1. D.J. Kane, R. Trebino: IEEE J. Quantum Electron. **QE-29**, 571 (1993); R. Trebino, D.J. Kane: J. Opt. Soc. Am. A **10**, 1101 (1993); K.W. DeLong, R. Trebino: J. Opt. Soc. Am. A **11**, 2429 (1994); K.W. DeLong, D.N. Fittinghoff, R. Trebino, B. Kohler, K. Wilson: Opt. Lett. **19**, 2152 (1994)
2. C. Iaconis, I.A. Walmsley: Opt. Lett. **23**, 792 (1998); C. Iaconis, I.A. Walmsley: IEEE J. Quantum Electron. **QE-35**, 501 (1999)
3. K.C. Chu, J.P. Heritage, R.S. Grant, K.X. Liu, A. Dienes, W.E. White, A. Sullivan: Opt. Lett. **20**, 904 (1995)
4. M.A. Franco, H.R. Lange, J.-F. Ripoche, B.S. Prade, A. Mysyrowicz: Opt. Commun. **140**, 331 (1997)
5. H.R. Lange, M.A. France, J.F. Ropoche, B.S. Prade, P. Rousseau, A. Mysyrowicz: J. Select. Top. Quantum Electron. **4**, 295 (1998)
6. C.W. Siders, J.L.W. Siders, F.G. Omenetto, A.J. Taylor: IEEE J. Quantum Electron. **QE-35**, 432 (1999)
7. A. Baltuska, A. Pugzlys, M.S. Pshenichnikov, D.A. Wiersma, B. Hoenders, H. Ferwerda: unpublished
8. G. Taft, A. Rundquist, M.M. Murnane, I.P. Christov, H.C. Kapteyn, K.W. DeLong, D.N. Fittinghoff, M.A. Krumbugel, J.N. Sweetser, R. Trebino: IEEE J. Select. Top. Quantum Electron. **QE-2**, 575 (1996)
9. L. Gallmann, D.H. Sutter, N. Matuschek, G. Steinmeyer, U. Keller, C. Iaconis, I.A. Walmsley: Opt. Lett. **24**, 1314 (1999)
10. A. Assion, T. Baumert, M. Bergt, T. Brixner, B. Kiefer, V. Seyfried, M. Strehle, G. Gerber: Science **282**, 919 (1998)
11. D. Meshulach, Y. Silberberg: Nature **396**, 239 (1998)
12. C. Iaconis, I.A. Walmsley: OSA Annual Conference (Baltimore, MD 1998)
13. T.M. Shuman, M.E. Anderson, J. Bromage, C. Iaconis, L. Waxer, I.A. Walmsley: Opt. Express **5**, 134 (1999)
14. C. Dorrer, B. de Beauvoir, C. Le Blanc, S. Ranc, J.P. Rousseau, P. Rousseau, J.P. Chambaret, F. Salin: Opt. Lett. **24**, 1644 (1999); C. Dorrer: Opt. Lett. **24**, 1532 (1999)
15. D.N. Fittinghoff, K.W. DeLong, R. Trebino, C.L. Ladera: J. Opt. Soc. Am. B **12**, 1955 (1995)
16. S. Yeremenko, A. Baltuska, M.S. Pshenichnikov, D.A. Wiersma: unpublished
17. C. Dorrer: J. Opt. Soc. Am. B **16**, 1160 (1999)
18. V.A. Zubov, I.I. Kuznetsova: Sov. J. Quantum Electron. **21**, 1286 (1991)
19. V. Wong, I.A. Walmsley: Opt. Lett. **19**, 287 (1994)
20. M. Takeda, H. Ina, S. Kabayashi: J. Opt. Soc. Am. **72**, 1 (1982)
21. A.M. Weiner, D.E. Leaird, J.S. Patel, J.R. Wullert II: IEEE J. Quantum Electron. **QE-28**, 908 (1992)
22. M.A. Dugan, J.X. Tull, W.S. Warren: J. Opt. Soc. Am. B **14**, 2348 (1997)
23. J.W. Goodman: *Introduction to Fourier Optics* (McGraw-Hill, New York 1988)
24. S. Kane, J. Squier: J. Opt. Soc. Am. B **14**, 1237 (1997)