

Theory and Design of Two-Dimensional Spectral Shearing Interferometry for Few-Cycle Pulse Measurement

Jonathan R. Birge*, Helder M. Crespo, and Franz X. Kärtner

*Department of Electrical Engineering and Computer Science, Massachusetts
Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02138, USA*

**Corresponding author: birge@mit.edu*

We provide a detailed explanation of our recently developed pulse characterization technique for few cycle pulses, called Two-Dimensional spectral Shearing Interferometry (2DSI). Based on SPIDER, 2DSI is relatively simple to implement, but requires no interferometer delay or scan calibration. We show simulations of 2DSI in the presence of noise and find that it retains the favorable noise performance of spectral shearing methods, performing similarly to standard SPIDER for a given measurement time. Experimental considerations when building and operating a 2DSI are provided, with experimental results shown for a 4.9 fs pulse, demonstrating the accuracy and precision of 2DSI. © 2007 Optical Society of America

OCIS codes: 320.7100, 120.5050.

1. Introduction

Spectral shearing interferometry (pioneered by the SPIDER method of Iaconis and Walmsley [1]) is unique among pulse measurement techniques in that it unambiguously measures the spectral phase of an optical signal by directly interfering neighboring frequency components. Furthermore, by encoding group delay information on the phase of a spectral domain fringe, SPIDER produces a frequency modulated signal and thus enjoys excellent immunity from noise and nonlinear optical bandwidth effects. This gives spectral shearing an advantage for measuring very wide bandwidths, relative to the “amplitude modulated” signals characteristic of other methods, differences in efficiency notwithstanding. Consequently, it has become one of the principal methods used to measure few-cycle pulses [2], alongside FROG [3] (and the many variants of both). There are, however, a few challenges with standard SPIDER that are relevant for the measurement, and especially optimization, of ultrahigh bandwidth sources approaching a single-cycle.

All spectral shearing methods, by their nature, involve measuring spectral group delay by observing the interference of two spectrally shifted copies of the pulse being measured. As will be explained in more detail in Section 2, any linear phase (delay) that occurs between the two components will be interpreted as a quadratic phase (dispersion) in the reconstruction. It turns out that the resulting error in measured pulse width scales linearly with the unaccounted for delay, multiplied by a factor proportional to the resolution of the measurement. At the

very least, this implies that great care must be taken in a SPIDER measurement to ensure that any linear phase is calibrated out.

Here, we present a method recently developed that seeks to ensure that there is no possibility for uncalibrated phase, by eliminating the possibility for a delay by robustly encoding the spectral phase measurement in a series of spectra of a *single* output pulse. While the requirement to take multiple measurements prohibits single-shot measurements, the method has the advantage that it involves no delay calibration, produces intuitively interpreted data, and requires only a low resolution spectrometer. We believe that 2DSI is thus a cost effective and efficient method for accurately and reliably measuring few- and even single-cycle pulses.

Our technique requires only the non-critical calibration of the shear frequency and does not perturb the pulse before up-conversion. Rather than encode the spectral group delay in a fringe in the spectral domain, 2DSI encodes phase in a pure cosine fringe along a completely independent dimension, by scanning the relative phase of the two spectrally sheared components. This reduces the resolution demands on the spectrometer to that required for proper sampling of the pulse itself, and allowing for complex phase spectra to be measured with high accuracy over extremely large bandwidths, potentially exceeding an octave.

2. Background and Motivation for Delay-Free Method

In any spectral shearing method, the net result is the creation of two spectral copies of the pulse under measure, with a small frequency shift Ω disposed between them. For sake of simplicity, in this paper we will ignore the overall offset in frequency

caused by the nonlinear upconversion, and will express everything in terms of the lowest upconverted pulse spectrum, $A(\omega) = |A(\omega)|e^{i\phi(\omega)}$. In the conventional SPIDER method, the two output pulses are delayed with respect to one another by a time τ , leading to a dense fringe in the spectral domain upon which the group delay spectrum is encoded,

$$I(\omega) = |A(\omega)|^2 + |A(\omega - \Omega)|^2 + 2|A(\omega)A(\omega - \Omega)|\cos[\tau\omega + \phi(\omega) - \phi(\omega - \Omega)]. \quad (1)$$

If the delay is large enough, the argument to the cosine can be isolated by signal processing of the fringe in the Fourier domain. The contribution of the delay is then subtracted out, leaving the finite difference of the spectral phase. More details are provided in [1] and [12].

Any unaccounted for delay is mistaken for a quadratic phase term, potentially one which results in underestimation of the true pulse width. As explained in [4], any error in estimating the interpulse delay results in an *absolute* error in the extracted pulse width. This measurement error is proportional to the delay error multiplied by the ratio of the shear to the bandwidth; this ratio is essentially the number of points over which we sample our spectral phase. The bottom line, however, is that because the error is absolute, the *relative* measurement error caused by a given interpulse delay uncertainty is proportional to the square of the pulse bandwidth [4],

$$\frac{\delta t}{t} \approx \frac{\delta \tau}{\Omega} \Delta \omega^2 \quad (2)$$

where $\delta\tau$ is the uncertainty in the interpulse delay, and $\Delta\omega$ is the pulse bandwidth.

The shear is essentially our sampling interval in the spectral domain. By the sampling theorem, it is inversely proportional to the temporal window over which we are measuring. As we will argue further in Section 6, the minimum time window we must resolve (i.e. the spectral sampling “rate”) ceases to scale inversely to bandwidth once pulse widths become shorter than the characteristic time scale of intracavity dispersion oscillations (such as those caused by reflections off the surface of dispersion compensating mirrors). The scaling of the delay uncertainty is more complicated, but by similar arguments, it either scales with the temporal measurement window (in the case of standard SPIDER) or is a constant (in the case of zero delay variants). As such, beyond a certain bandwidth the measurement uncertainty scales as the square of the bandwidth.

As the single cycle limit is approached, the effect of even tens of attoseconds of unknown delay becomes a concern. Calibrating and maintaining a delay on the order of hundreds of femtoseconds to within tens of attoseconds is not trivial, especially if the experimental configuration must be changed to measure the delay, as in those systems where the delay is measured by the spectral interference between the second harmonics of the pulses. At the very least, it is certainly necessary to recalibrate the device before each measurement.

The preferred way of calibrating a standard SPIDER fringe, introduced by Dorrer [4], does not involve explicitly computing a delay. In this approach, two SPIDER measurements with different shears are subtracted from each other to

obtain the SPIDER phase minus any common linear phase due to the pulse delay. (This assumes that it can be done so such that the delay is invariant between the two measurements. One can ensure this is the case by taking them simultaneously, as in [6].) When the spectrometer and delay calibrations are mixed together in this way, one then does not speak of “delay” calibration but rather a general baseline phase calibration that handles simultaneous calibration of the delay and any spectrometer nonlinearity. However, whether implicit or not, a measurement of delay is inherent, and as the main component of the phase and the leading term in any calibration error, it makes sense to cast any discussion of sensitivity in terms of the effective delay uncertainty. Even in the absence of systematic delay miscalibration, simple measurement noise will contribute some effective $\delta\tau$ that will be the dominant source of pulse width uncertainty. This is borne out in the noise simulations discussed in Section 8.

A second difficulty with standard SPIDER is the need to split the measured pulse into two copies. In practice, doing so without the introduction of additional dispersion is extremely difficult, especially for pulses with spectra exceeding an octave. The only type of beamsplitter that could potentially avoid significant dispersion is an extremely thin beam splitter, such as a pellicle. Unfortunately, this then introduces the problem of multiple reflections, introducing phase oscillations into the measurement.

The splitting issue can be avoided by the use of two chirped pulses, as pioneered by Zero Additional Phase SPIDER (ZAP-SPIDER) [7]. ZAP-SPIDER

produces a standard SPIDER fringe, but does so without requiring the interferometric splitting of the short pulse, adding no additional dispersion from the beamsplitting operation, hence the name. However, the calibration issue still remains in ZAP-SPIDER, as the two chirped pulses are disposed along separate paths to allow for the upconverted components to be separated and delayed with respect to one another.

3. Principle of 2DSI Operation

Our solution to the above issues is to use two collinear chirped beams such that the output is a single beam, in essence a single pulse with a complicated spectrum. Refer to Figure 2 for a frequency domain block diagram of the conceptual process, and Figure 1 for a prototypical experimental layout. This arrangement should eliminate the potential for a delay to occur between the two output components. That this is the case is suggested, at the least, by our results in Section 9.B. However, we must admit that the precise extent to which this is true cannot be proven conclusively down to the level of attoseconds. In the absence of a perfectly known “reference pulse,” a numerical spatiotemporal simulation of the noncollinear upconversion would have to be performed to determine the effects of phase matching on the output beam properties.

First, a highly chirped pulse is created by picking off a portion of the pulse to be measured, and dispersing it. This chirped pulse is sent through an interferometer to create two copies of the chirped pulse. The two chirped pulses are then mixed with the original short pulse in a Type-II $\chi^{(2)}$ crystal in a noncollinear

geometry (see Figure 2) such that the sum frequency generation (SFG) signal can be isolated from both the fundamental and second harmonic terms in the output. Presuming the chirped pulses are sufficiently dispersed, the short pulse under test effectively sees only two single frequency components, one from each chirped pulse. The frequencies are determined by the total delay through each path of the interferometer, with shear thus determined by the difference in the delay of each arm.

In the crystal, because the short pulse is mixed with two essentially pure frequencies the output spectrum consists of two spectrally sheared copies of the original spectrum. If the chirped pulses are collinear, the two frequency components will have the same propagation vector and thus we can regard the output as a single pulse. This neglects the small difference in transverse photon momentum between the two chirped pulses. However, for small deviations such that nonidealities in the imaging optics can be neglected, there should be no possibility for delay to arise. A given wavelength is imaged to the same pixel regardless of output angle, and thus no phase difference can arise by Fermat's principle.

The upconversion frequencies can be independently and arbitrarily chosen to suit the pulse characteristics and maximize the overall signal to noise ratio. A discussion of how best to do so is taken up in Section 6. This ability also allows for a wide range of pulse widths and complexities to be measured by the same setup. Given that a large range is likely to be encountered in the course of optimizing a laser, it is useful to have this flexibility. In contrast, the shear in standard SPIDER

is constrained by the dispersive element and the required delay, limiting its versatility. The collinearity of the chirped beams in 2DSI also means that the shear can be changed without affecting the alignment of the system.

The pulse under measurement only experiences a few reflections, and is therefore relatively unperturbed before measurement. The Michelson interferometer that splits the chirped pulses, on the other hand, can be highly dispersive and even unbalanced, allowing for the use of simple splitting optics. The low distortion experienced by the short pulse is extremely important for single-cycle pulses, as any spurious dispersion inherent to the apparatus will invariably end up (in the opposite sign) in the pulse after optimization.

The fact that the output is essentially one pulse, the spectrum of which encodes the spectral group delay, is the key to the stability and reliability of 2DSI. In theory, all the information is present in a single-shot spectrum if we were willing to trust the spectral amplitude information. However, such operation would negate the primary strengths of SPIDER methods: the encoding of group delay as a fringe phase, lending immunity to phase matching bandwidth effects. It would also introduce a time direction ambiguity.

To robustly observe the spectral phase difference between the two pulse components at a given wavelength (which we recall is proportional to the local group delay) the relative phase ϕ of one of the quasi-CW beams is scanned over a few optical cycles. This allows one to observe the phase difference between the two spectral components as a function of wavelength, providing a direct and immediate

measurement of spectral group delay. The phase scan is typically done by vibrating one of the mirrors in the Michelson interferometer, though it could also be achieved with a phase modulator, such as a liquid crystal or electro-optic device. The spectrum of the up-converted signal is then recorded as a function of the phase φ , yielding a 2-D intensity function that is given by

$$I(\omega, \varphi) = 2 |A(\omega)A(\omega - \Omega)| \cos[\varphi + \underbrace{\phi(\omega) - \phi(\omega - \Omega)}_{\tau_g(\omega - \Omega/2)\Omega + O(\Omega^2)}] + \text{D.C.}, \quad (3)$$

where, as before, $A(\omega)$ is the low upconverted pulse spectrum, and $\phi(\omega)$ is its spectral phase. The under-bracketed expression can be viewed as a finite difference approximation of the group delay scaled by the shear. This term is what all SPIDER variants measure, though in 2DSI it is available directly without any filtering, as discussed in the next section, since there are no other terms dependent on the frequency. A simple two-dimensional raster plot of $I(\omega, \varphi)$ reveals the shifted pulse spectrum along the ω -axis, with fringes along the φ -axis that are locally shifted in proportion to the group delay at a given frequency (as illustrated in Figure 3). The user can thus immediately ascertain salient properties of the complex spectrum simply by looking at the raw data: the cosine fringe at each wavelength is vertically shifted in proportion to its actual delay in time, with the fringe amplitude roughly proportional to the power spectral density (neglecting bandwidth effects). The ability to use the raw spectrometer data when optimizing a laser yields information not available from processed data from an inversion algorithm alone, such as measurement noise and laser stability.

Comparison of equations (1) and (3) shows that the fringes produced by 2DSI and conventional SPIDER are mathematically identical, except for the fact that the fringe in 2DSI is produced by a phase occurring in a separate domain (the φ term), whereas in SPIDER the fringe phase oscillates in the spectral domain. In both, the fringe can be viewed as creating sidebands in the respective Fourier domain of the fringe; with SPIDER we get sidebands in the pseudo-time domain and in 2DSI the sidebands are pulled into the “pseudo-frequency” domain relative to the phase delay. In either case, the purpose of the fringe is to pull the sidebands away from the central DC term so that they don’t interfere, rendering the phase extraction insensitive to the amplitude of the fringe. In SPIDER, the fact that the sidebands are in the optical frequency domain results in a significant increase in required spectrometer resolution over that needed to simply resolve the fundamental pulse spectrum. In 2DSI, the extra dimension means that the spectrometer resolution required is simply that required by the Nyquist limit for the pulse being measured (i.e. determined by the time window that must be resolved), enabling larger time bandwidth products to be measured than with standard SPIDER.

4. Relation to Other Spectral Shearing Methods

SEA-SPIDER, a method developed by Ellen Kosik [8] and successfully adapted for few-cycle pulses by Adam Wyatt [9], follows ZAP-SPIDER in using two chirped pulses. By putting a tilt between the two upconverted beams and measuring the output with an imaging spectrometer, a spatial fringe is created in one axis while spectrally resolving the other. 2DSI is similar to SEA-SPIDER in many

respects, with the fringe encoded in terms of an upconversion phase, rather than mixed with the pulse profile and encoded in space,

$$I(\omega, x) = 2|A(\omega, x)A(\omega - \Omega, x)|\cos[2\omega \cos \theta x/c + \phi(\omega) - \phi(\omega - \Omega)] + \text{D.C.}, \quad (4)$$

where x is the transverse spatial dimension and θ is the half angle of intersection between the two upconverted beams.

SEA-SPIDER's use of spatial encoding allows for single-shot pulse characterization, including some kinds of spatiotemporal measurement along one axis. However, the use of separate output paths from the crystal in SEA-SPIDER makes the method more susceptible to a delay occurring between the point of upconversion and interference on the CCD detector. It is necessary to calibrate both the intersection angle of the two beams in SEA-SPIDER (the nominal fringe spacing) as well as the angle of the spectrometer grating axis (i.e. the nominal fringe angle) and inspection of (4) and comparison to (1) shows that miscalibration of either angle is equivalent to an erroneous delay in SPIDER, manifesting as discussed in Section 2.

Because a nonlinear upconversion operation does not commute with a linear delay, the fact that a 2DSI system involves a scanning of the chirped pulse delay *prior* to upconversion makes it fundamentally different from SEA-SPIDER, which involves delays applied after. The result is that in SEA-SPIDER there is a varying time delay between pulse copies (with an associated spectral dependence), whereas in 2DSI there is a changing pure phase that is constant in frequency. We believe

varying the chirped pulse phase before upconversion results in a more robust measurement than changing a delay after upconversion, as our approach results in a simpler interferogram that requires less calibration. Because the output of 2DSI is essentially a single pulse, it is essentially impossible for the two components to experience further phase shifts.

More recently, Ian Walmsley *et al* introduced a method which simultaneously records multiple SPIDER fringes over a range of different shears [10]. This allows for the calibration of the delay through a consistency requirement on all of the fringes, and was recently demonstrated for a 70 fs pulse. Because it implicitly relies on a thick crystal, however, applying this method to few- or single-cycle pulses may not be possible.

The orthogonality of the temporal fringe to the spectral domain is also why the fringe period does not need to be known in 2DSI. At each wavelength, the fringe encodes only one piece of information (the group delay), and we are only concerned with the relative shift between wavelength. We thus do not care about the frequency of the fringe or even whether or not it is constant, only needing to know its relative phase. This eliminates many potential avenues for measurement error in a real system, such as scan linearity and calibration. In fact, the only calibration needed by 2DSI is for the up-conversion frequencies which produce the shear. Fortunately, this is a relatively non-critical calibration, as the relative pulse measurement uncertainty is proportional to the relative uncertainty in the shear [1].

However, there is a cost for 2DSI's simplicity in reconstruction and experiment, which is the loss of single-shot capability. Nonetheless, while the requirement for scanning in the present version of 2DSI renders single-shot operation impossible, it is still capable of video-rate operation, which is essential for laser tuning or automated pulse compression. The scanning mirror only needs to move a few microns, at most, and thus the system is theoretically capable of operation at kilohertz rates. In practice, signal to noise generally puts a lower bound on the required integration time to yield sufficient accuracy. This requirement will often exceed the pulse repetition rate, especially given the low efficiency of spectral shearing techniques (save for M-SPIDER [11]). This means that 2DSI will be appropriate for measuring a relatively stable oscillator, but a single-shot-capable method may be a better choice for an amplified pulse, especially given the extra pulse-to-pulse noise induced by amplification.

5. Physical Layout and Operation

As with any method, there are several ways to implement the optical operations required for 2DSI. In Figure 2, we provide a schematic of one approach. We have found that this arrangement provides for a cost effective, robust measurement.

To begin with, roughly four percent of the short pulse under test is picked off by the Fresnel reflection from the glass cube beamsplitter (C) used in the Michelson interferometer. This can be done with a wedge bonded to the beamsplitter, or by simply operating the interferometer at slightly shallower than a right angle geometry. The remainder of the pulse is split in the interferometer, where a one

inch glass beamsplitter provides sufficient chirping to measure a pulse in the few-cycle regime (see Section 18). If longer pulses need to be measured, an additional glass block can be placed in the output path of the interferometer.

After the interferometer, the polarization of each chirped beam is rotated by a half-wave plate. (The polarization must be rotated since we are using Type II upconversion.) A simple low-order half-wave plate is sufficient since the rotation only needs to occur at two frequencies that are separated by the shear, which is typically on the order of 5 to 10 THz.

The short pulse to be measured and the chirped pulses are made parallel and then focused by an off-axis parabolic mirror into a thin (roughly 30 μm) Type II BBO crystal. As pointed out by Walmsley [12], BBO is rather fortuitous as an upconversion medium for spectral shearing methods, as its Type II phase matching curve can be engineered to have octave-spanning bandwidths in one polarization, with narrow bandwidths in the other. The collinear output of 2DSI preserves the favorable phase-matching of standard SPIDER, as illustrated in Figure 6. Coupled with the fact that bandwidth effects don't impair spectral shearing measurements beyond reducing the signal to noise ratio, this renders 2DSI capable of precise self-referenced measurements of pulses down to a single-cycle [11]. The use of noncollinear input beams allows the SFG beam to be isolated from the SHG beams. Finally, a standard glass objective focuses the output onto a grating spectrometer.

Using a collinear input (where the chirped pulses and the pulse under test form a single beam) would simplify alignment, but it would present several

difficulties and is not recommended. First, combining the beams would involve loss, and would be impossible to do without introducing either significant loss (in the case of a Fresnel reflection) or phase distortion (in the case of a coating or pellicle) on the pulse to be measured. Second, separating the SFG from the fundamental would be difficult in the case of a single-cycle pulse, especially for a grating spectrometer. In addition, SHG light would likely be present, though this could be filtered out by proper signal processing during reconstruction.

Mirror (B) is controlled manually by a translation stage with a relatively long travel (at least millimeters). This mirror is used to control the delay between the chirped pulse copies that determines the shear. The other mirror (A) is controlled by a short throw piezoelectric translation stage. Since this scanning mirror will only need to be translated a few microns, at most, a flexure stage can be used to maximize stability and reduce noise during the scan.

A third delay stage (D) is used to adjust the temporal overlap between the short pulse and the chirped pulses, which determines the overall frequency of the upconversion. This stage should be rather long, on the order of a centimeter, at least, to provide for ease in alignment. Together, stages (B) and (D) provide sufficient degrees of freedom that both upconversion frequencies can be independently chosen.

While 2DSI appears on the surface to have an experimental complexity similar to IAC or FROG, due to the scanning, this is not the case given the short distance of the scan and the lack of any calibration needed for it. The scan itself

may be handled in an open loop fashion, and a basic function generator may be used to ramp the mirror while spectra are taken at regular intervals. (This is the method we employed to provide the experimental demonstrations shown in the following section, in fact.)

Our simulations show that nonlinearity of the fringe has no effect on the accuracy, though it does affect the reconstruction in the presence of noise by broadening the fringe in the frequency domain. Empirically, we found that even 50% nonlinearity (defined as the maximum relative deviation from the nominal scan rate) results in only a halving of the signal to noise ratio. Nonlinearities below 10% were not found to have an appreciable effect on the noise performance. The linearity of the scan will thus not be an issue for most commercial stages.

Before a reconstruction is performed, the upconversion frequencies must be determined. This is done by alternately blocking one of the arms of the Michelson and recording the spectrum of each individual upconverted component. By cross-correlating the upconverted spectrum with the fundamental pulse spectrum (taken by a separate OSA), the individual upconversion frequencies can be computed, and from this the shear, providing all information necessary for the reconstruction of the pulse. The calibration of the upconversion frequencies is not particularly sensitive; once this calibration is done, it does not need to be repeated for subsequent measurements unless the configuration is changed.

To take a measurement, a computer controlling the piezoelectric stage moves the (A) mirror (A) over a total range of roughly 1 micron, recording a dozen or so

spectra during the scan. Due to the effective monochromaticity of the chirped pulses, this is equivalent to scanning a pure zeroth-order phase of the corresponding upconverted spectral component. As discussed in Section 6, only four spectra are actually needed for reconstruction, but taking more results in a more intuitively understandable trace and one that allows for a simplified reconstruction algorithm.

6. Design Considerations

The construction and alignment of 2DSI presents no challenges beyond that required of for any Michelson interferometer and sum frequency generation. The experimental setup is essentially the same as a conventional SPIDER, with the main differences being the location of the dispersion and the addition of a motor to one delay stage. Care should be taken to ensure relative collinearity of the chirped beams with respect to each other. Fortunately, any deviation from this should not result in a spurious spectral phase, as discussed in Section 3. In fact, all conceivable misalignments, to the best of our knowledge, will simply result in attenuation of signal, not errors.

When putting together a 2DSI setup, there are essentially six design issues:

- i. How much shear to use
- ii. How far to scan and how linear it must be made
- iii. How much dispersion to use for chirping
- iv. How often to sample the spectrum during the scan
- v. Spectrometer resolution
- vi. Nonlinear crystal thickness

These six issues are not independent, and are all determined by the bandwidth and temporal extent of the pulse. Each is discussed individually in the following subsections.

A. *Shear Frequency*

The shear frequency ultimately becomes the “sampling period” of the phase in the spectral domain, and most of the other parameters can be derived from it. By the Shannon sampling theorem, this determines the temporal window over which we can reliably measure. In the remainder of this paper will refer to this time period as the *shear Nyquist duration*. While it may be tempting to simply pick a shear sufficient to include a region of interest, such as the main pulse, this is not sufficient. Any satellite pulses or pedestal structure will still be measured, but will be “aliased” into our chosen temporal window, resulting in errors. Thus, the shear must be chosen so as to allow resolution of all spectral features. Moreover, satellite pulses which may seem negligible can have significant effects when aliased onto the main pulse, due to the fact that aliases add in field. Our numerical simulations have shown that a satellite pulse with an intensity of only 2% of the main pulse can change the FWHM of the main pulse by up to 5.5% should it be aliased on top of it.

The required temporal window is not known a priori, by definition, in a pulse that we are seeking to characterize. It is tempting to simply conclude the sampling was sufficient by taking a measurement and verifying that the reconstruction is well contained with the time window. However, the nature of aliasing is that once sampling occurs, one cannot tell the difference between correctly or insufficiently

sampled signals. The only reliable way to assure that the sampling rate is sufficient is to take measurements at a series of decreasing shears, and verify that the measurements converge to required precision.

In practice, for sub-two-cycle pulses produced by oscillators with dispersion compensating mirrors, we have found that a shear of around 4–5 THz is required to sufficiently resolve the satellite pulses and pedestal. Unfortunately, this appears to be independent of the final pulse bandwidth, as one would expect given that the mechanisms for creating them are independent of bandwidth.

B. Chirping Dispersion

The signal of the final measurement is inversely proportional to the amount of dispersion used to create the quasi-CW beams, and thus the chirp of the ancillary pulses should be the minimal amount required to result in an accurate measurement. Over-chirping results in an unnecessarily weak signal, and under-chirping results in a complicated “blurring” of the measurement in the spectral domain as the upconversion occurs with a range of interacting wavelengths. Having determined the required spectral resolution with the shear, the chirping should be selected such that the associated blur is smaller than this resolution. Requiring that the chirped pulses’ instantaneous frequency does not change by more than Ω over the entire temporal window conservatively implies that the required second-order dispersion be

$$D_2 \geq \frac{1}{2\Omega^2} \tag{5}$$

For example, if 5 THz of shear is used, the dispersion required is 20,000 fs², obtained by roughly 10 cm of SF11. This is quite a bit of dispersion, and will significantly impair the signal power.

However, equation (5) presumes that we care about properly measuring the field over the entire temporal window. In most cases, we only care about accurately measuring a main pulse. (This is consistent with our previous statement that all satellite structure must be resolved; while we may not be interested in the accurate measurement of satellite pulses, we still must ensure that they do not alias.) If a well-separated satellite pulse sees shifted local frequencies of the chirped pulses, it will simply result in a *local* error in the reconstruction of that feature in proportion. Thus, a more reasonable criterion is simply for the frequency of the CW beam to change by no more than the shear over the temporal extent of the main pulse. Taking this width to be T, this gives

$$D_2 \geq \frac{T}{\Omega} \quad (6)$$

Since the optimal shear cannot be known a priori (as explained above) neither can the dispersion be determined without knowing the temporal extent of the pulse, exactly the thing we seek to measure. Ideally, one would iteratively increase the chirping as with the shear, until the measurement converged to some satisfactory precision. However, this is generally not feasible, as variable sources of dispersion with wide variability are not easily found and would be expensive regardless. As such, the chirp will have to be chosen somewhat conservatively when

the 2DSI system is built, with the worst-case pulse width in mind and a conservative estimate for the smallest shear likely to be used. A safe configuration for a few-cycle laser can be found by assuming a shear of 5 THz, and a pulse width of no more than 25 fs, yielding a dispersion of 5000 fs². This is a small enough amount that it can be provided by the cube beamsplitter used in the chirped beam interferometer (as in Figure 1).

C. Delay Scan Length

The scan must be long enough such that the sidebands shown in Figure 4 are well separated from the central DC peak. This distance will vary depending on whether or not we are performing windowing, and how accurately the scan can be matched to the fringe period. In general, one will be fine as long as at least three fringes are visible, as derived in Section 6. Beyond this consideration, the length of the scan actually does not matter. This implies scan lengths on the order of a micron or two, at most, allowing very stable short-throw piezoelectric stages to be used, and enabling high scan rates limited only by signal levels. In our setup, we have been able to achieve scan rates of several Hertz, limited by the readout speed of our spectrometer.

D. Scan Sample Rate

The number of points over which the scan is sampled is largely irrelevant so long as they are sufficiently sampled. Beyond that, doubling the number of samples but halving the integration time of each will result in identical measurements. That

is to say, the only thing that matters is the total measurement time. The exception to this is when the measurement is limited by readout noise, in which case keeping the samples to a minimum is advised. In such a situation, it is actually possible to get by with only four measurements.

Given that the scan rate doesn't affect the measurement, a rate high enough to yield a human-readable spectrogram can be useful, such as that shown in Figure 3. This allows the user to visually gauge the functioning of both the laser and the 2DSI system.

E. Spectrometer Resolution

The required spectrometer resolution is simply that consistent with the sampling rate set by the shear. Any more will only result in excess dark current noise. Of course, the spectrometer resolution is not a readily changed variable, so we recommend choosing a spectrometer with a resolution of at least 2 THz, to ensure the apparatus can be used to measure any reasonable time-bandwidth product.

F. Nonlinear Crystal Thickness

Thanks to the phase encoding inherent to spectral shearing, the bandwidth of the nonlinear crystal only affects the measurement from a signal-to-noise standpoint. The uniformity of the upconversion is thus not an issue, and the ideal crystal is such that the minimum conversion efficiency is maximized. In most cases, this entails choosing the crystal angle that maximizes the minimum conversion over

the expected bandwidth, regardless of the resulting variation of efficiency across the spectrum.

As mentioned earlier, a unique aspect of spectral shearing interferometry is that the nonlinear operation is between a signal with high bandwidth (the pulse under test) and one which is nearly monochromatic (the chirped beams). This allows one to take advantage of the inherent asymmetry in Type II upconversion, wherein one polarization will have greater bandwidth than the other. In the case of BBO, the dispersion works especially favorably for Type II upconversion in the NIR, such that bandwidths exceeding an octave can be efficiently upconverted with crystals of reasonable thickness. This, combined with the natural noise resilience of spectral shearing, goes a long way to making up for the relatively low optical efficiency of SPIDER methods. To optimize the conversion efficiency, a full non-collinear phase matching curve must be considered for the specific wavelength ranges to be used. An example of such a curve is shown in Figure 6.

7. Reconstruction Algorithm

A. Fringe Phase Extraction

The inversion process for a 2DSI fringe is significantly simpler than that required for FROG, or even SPIDER. The only information we need to extract from the 2D interferogram is the phase of the fringes along the direction of the scan delay (i.e. the vertical direction in Figure 3). Precise quantitative determination of the fringe phase, and thus the group delay, at each wavelength, can be obtained in one

of several ways. The most direct method of extracting the fringe phases is to simply fit a cosine to the fringe at each wavelength. While this approach is potentially quite accurate, it is relatively slow and will not always converge in the presence of noise.

If the fringe period is known exactly, the phase can be recovered using a variation of the Takada algorithm [13], analogously to what is done in SPIDER and SEA-SPIDER [8]. However, this would require having a calibrated linear scan. It would complicate the experiment greatly to require the scan to be known and linear, necessitating either feedback control or accurate measurement of the mirror displacement and linearization via signal processing. Fortunately, the fact that no information is encoded in each fringe other than the spectral group delay means that we can use a very simple reconstruction algorithm.

An efficient and direct way to access the fringe phase is to compute a series of 1D FFTs along the fringe axis. According to (2), the fringes generate well-separated sidebands in the continuous Fourier domain (see Figure 4) and their phase angle represents the spectral group delay term in (1). However, depending on the length of the scan and how close it is to a multiple of the CW wavelength, in practice windowing effects may broaden the sidebands such that they and the central DC term interfere, perturbing the phase. This can be effectively dealt with by applying a windowing function to eliminate the broadening. We have empirically found a hamming window to perform quite well if the number of fringes initially visible is at least three. This method comes with the disadvantage of throwing away information

at the edges of the scan, however, which will necessarily reduce the SNR of the measurement. Nonetheless, we mention it as it is by far the simplest approach to take, and will work in virtually any situation.

In practice, fortunately, windowing isn't necessary so long as a sufficient number of fringes are observed. Because we are actually measuring group delay, any constant offset is meaningless, which mitigates the effects of windowing. The worst-case relative error in the extracted pulse spectral phase can be shown (see Appendix A) to be approximately equal to

$$\frac{\phi_{\text{rec}}}{\phi_{\text{act}}} = \frac{\Delta\nu}{n + \Delta\nu}, \quad (7)$$

where n is the index of the FFT component (the harmonic number) and $\Delta\nu$ is the difference in frequency between the fringe and the harmonic (normalized to the fundamental, such that the worst case scenario is $\Delta\nu = 0.5$). As can be seen, the error will be relatively small if the fringe frequency is close to the FFT harmonic, or if n is large. As a rule of thumb, so long as three or more fringes are visible, or the scan length is within 10% of the fringe period, the worst possible error will be no more than a few percent and windowing will not be necessary.

It is often advisable to filter in the wavelength domain to suppress noise outside the measurement time window. This filtering can be efficiently combined with the reconstruction by computing a 2D FFT of the interferogram, and selecting a subset of the single line of “pixels” representing the 2D sideband within the shear

Nyquist limit. Wavelength domain filtering is especially important when the detector is significantly oversampling the spectrum.

B. Finite Difference Inversion

Having determined the phase of the fringe at each wavelength, multiplying by the shear yields the finite difference of the spectral phase, as shown in (3). The final step, then, is to compute the spectral phase of the measured pulse from these finite differences. This step is common to all spectral shearing methods, and has been discussed elsewhere in the literature.

The most straightforward way is to simply treat the finite differences as proportional to the spectral group delay, numerically integrating them using standard methods and accepting the $O(\Omega^2)$ error. However, this does not result in the most accurate measurement possible, as the truncation error due to the finite difference is actually reversible. One can regard a finite difference as approximating the continuous derivative operator with a two term “discrete time” FIR filter. The amplitude transfer function of this effective FIR filter is sinusoidal in the pseudotime domain, peaking at half the inverse period of the finite difference (half the shear Nyquist rate). Thus, it underestimates the magnitude of phase oscillations, with the underestimation increasing as the oscillation period decreases. In terms of pulse reconstruction, this means that the amplitude of satellite pulses will be underestimated the further away they are from the main pulse.

A rigorous approach to the phase reconstruction is to compute the inverse of the finite difference operation, yielding the phase without any errors other than

those caused by noise. To do so, the data must first be antialiased by filtering out all “frequencies” above $1/\Omega$. Then, a sinc interpolation can be used to compute the fringe phase φ on a regular grid of points spaced by Ω . A simple cumulative sum will then yield the spectral phase of the pulse. It may seem that this approach is not optimal in terms of noise performance, as the data will be concatenated at a lower resolution than that provided by the spectrometer, seemingly throwing out data points that are being skipped over. However, the antialiasing filter step provides the averaging in this case. By suppressing all noise beyond the shear Nyquist limit, the final filtered phase data will be made internally consistent such that the concatenation operation will yield the same result regardless of the starting point.

8. Sensitivity to Noise

As illustrated in Figure 4, the sidebands in 2DSI are spectrally compact. The fringe is simply an impulsive line in the 2D Fourier domain, the sharpness of which is limited by the scan length. This spectral compactness means that the bulk of any noise will not interfere with the signal, maximizing SNR. A significant amount of noise immunity is gained by isolating the sideband (as discussed in Section 6) assuming the noise is uniformly distributed over all frequencies. Furthermore, the information is contained in the phase of the sideband, not its intensity, further reducing the sensitivity to noise.

An important question to answer is whether or not 2DSI’s need to scan over a delay renders it more or less sensitive to noise than standard SPIDER. Intuitively, one might expect that for a given total integration time, they would be comparable

in light of our previous assertion that 2DSI simply takes the SPIDER sidebands and moves them in a new direction. To illustrate the noise immunity of 2DSI (and in fact spectral shearing in general) we simulated the measurement of a pulse with a rectangular spectrum of 300 THz bandwidth, whose transform-limited pulse width is roughly 3 fs. The SPIDER method was simulated using the same noise source, with an integration time equal to the total measurement time of the 2DSI trace. The SPIDER calibration method used was that of [4], where a calibration measurement with exactly zero shear was assumed to be available that exactly preserved the delay.

The results are shown in Figure 5. For a given measurement time, 2DSI has half the variance of SPIDER. The difference between the two is entirely due to uncertainty introduced into the SPIDER calibration by the detector noise. Furthermore, it is apparent from the middle plot in Figure 5 that the SPIDER error is predominantly composed of a second-order term, validating our consideration of the calibration as largely manifesting as an effective delay uncertainty.

If the SPIDER calibration noise issue is ignored, the two methods perform identically, regardless of the type of noise used (i.e. additive or shot). This is to be expected, given 2DSI simply takes the SPIDER sidebands and moves them into another dimension. From a signal processing point of view, a SPIDER measurement is equivalent to a 2DSI measurement where the upconversion phase is constant and the sidebands are created by a temporal delay. It is thus to be expected that the two methods would perform similarly in the absence of calibration issues.

9. Experimental Demonstrations

A. Precision Test

To gauge the relative precision of the method, a few-cycle (5 fs FWHM) pulse from a prismless Ti:sapphire laser was measured using a 2DSI setup similar to that shown in Figure 1, using a shear of 18 THz. The pulse was measured both before and after dispersion by a one mm fused silica plate, Figure 7(b). It is apparent from the spectral GD curves that the pulse is initially slightly negatively chirped, and the positive dispersion introduced by the glass plate is evident. The sharp roll-off in the GD below 650 nm is genuine, and caused by phase distortion from the output coupler. Oscillations in the spectral group delay, caused by the chirped mirrors, are also clearly visible. Despite these perturbations in the individual spectra, the oscillations completely cancel in the difference between the phase of the two measurements, which matches well to that predicted by the known Sellmeier equations for fused silica, as shown in Figure 7(c). In fact, in terms of phase delay, the 2DSI system measured the glass dispersion to within 30 attoseconds of phase delay over a bandwidth from 600 to 1000 nm, Figure 7(d). This precision was achieved despite the absolute phase delay of each measurement ranging over more than 40000 as.

While these results suggest that the 2DSI apparatus is capable of precise measurements, they do not rule out the possibility of constant errors occurring that are consistent between measurements. For example, if there were an unknown linear phase creeping in the measurement somehow (as discussed earlier in Section

7) such a systematic error would not be evident from these dispersion measurements.

B. Accuracy Test

To qualitatively demonstrate the absolute accuracy of the system and rule out the existence of systematic errors, we recently performed a measurement on an octave spanning sub-two-cycle pulse [14,15] from another unamplified Ti:sapphire oscillator, and compared the 2DSI measurement to that obtained with a standard IAC (Figure 8).

In Figure 8(d), we show the reconstructed pulse envelope and phase for the pulse, measured to have a full width half maximum of 4.9 fs. To our knowledge, this is the shortest pulse measured with a spectral shearing method directly from an oscillator (i.e. without use of an external amplifier). As shown in Figure 8(b), the measured IAC and that predicted by the 2DSI measurement show fairly close agreement. We attribute most of the difference between the two to band-limiting effects on the IAC, which is not well suited to measuring a 4.9 fs pulse (note, for example, the lack of symmetry in the IAC trace). Thus, we do not present this measurement as further evidence of the precision of 2DSI, as an IAC is not a particularly reliable measurement of fine detail. However, the fact that the overall pulse widths predicted by both are consistent does suggest that we are correct in assuming that no appreciable hidden linear phase can occur in 2DSI. This validates the assertion that 2DSI does not require a separate calibration step. This measurement has also been corroborated by detailed simulations of the laser cavity

in question, which also predict a 4.9 fs pulse with a large sub-pulse 9 fs away from the peak (see [16] for details).

10. Future Work

The geometry of 2DSI is unique in that the pulse to be measured never encounters a dispersive element, and yet the arrangement is still collinear. Were the spectrometer replaced by an imaging spectrometer, a spatially resolved 2DSI measurement could be taken along one axis. Moreover, if a full 2D imaging spectrometer were used, such as that available with imaging Fourier transform spectroscopy or grisms, one could make a full 2DSI measurement at a 2D array of points along both transverse axes. By spatially filtering the chirped beams so that they were spatially coherent, they would provide a constant phase reference across the beam profile, enabling a full spatiotemporal reconstruction up to a trivial constant and quadratic spatial phase (focusing). This would allow for the first self-referenced 3D measurement of few-cycle pulses.

By using a nanostructured stepped mirror, a single-shot version of 2DSI might be implemented. The CW phase would then be encoded as a function of space, and the fringe could be read using an imaging spectrometer. This arrangement would result in the mixing of the spatial profile with the fringe, essentially creating a collinear variation of SEA-SPIDER with a spatially varying zeroth-order phase as opposed to linear phase. This could provide similar advantages to SEA-SPIDER, but with potentially simpler calibration and alignment.

11. Conclusion

Two-dimensional spectral shearing interferometry involves a relatively simple optical setup with little calibration required, and yet is capable of spectral phase measurements accurate to within tens of attoseconds of phase delay over octaves of bandwidth. The lack of dispersion on pulse to be measured, the stable absence of delay between the sheared pulses, and the relaxed spectrometer resolution requirements make 2DSI extremely well suited for the measurement of wide-bandwidth pulses, including those with potentially complicated spectral phase.

12. Acknowledgements

The authors acknowledge funding from DARPA and the AFOSR, and would also like to thank Dr. Richard Ell, who designed and built the 5 fs laser used in the precision testing.

13. References

- [1] C. Iaconis and I. Walmsley, “Self-referencing spectral interferometry for measuring ultrashort pulses,” *IEEE J. Quantum Electron.* 35, 501–509 (1999).
- [2] L. Gallmann, D. Sutter, N. Matuschek, G. Steinmeyer, and U. Keller, “Techniques for the characterization of sub-10-fs optical pulses: a comparison,” *Appl. Phys. B* 70, S67–S75 (2000).
- [3] D. Kane and R. Trebino, “Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating,” *IEEE J. Quantum Electron.* 29, 571–579 (1993).

- [4] J. R. Birge and F. X. Kärtner, "Analysis and mitigation of systematic errors in spectral shearing interferometry of pulses approaching the single-cycle limit [Invited]," *J. Opt. Soc. Am. B* **25**, A111-A119 (2008).
- [5] C. Dorrer, "Influence of the calibration of the detector on spectral interferometry," *J. Opt. Soc. Am. B* **16**, 1160–1168 (1999).
- [6] C. Dorrer, "Implementation of spectral phase interferometry for direct electric-field reconstruction with a simultaneously recorded reference interferogram," *Opt. Lett.* **24**, 1532–1534 (1999).
- [7] P. Baum, S. Lockbrunner, and E. Riedle, "Zero-additional-phase SPIDER: full characterization of visible and sub-20-fs ultraviolet pulses," *Opt. Lett.* **29**, 210–212 (2004).
- [8] E. Kosik, A. Radunsky, I. Walmsley, and C. Dorrer, "Interferometric technique for measuring broadband ultrashort pulses at the sampling limit," *Opt. Lett.* **30**, 326–328 (2005).
- [9] A. S. Wyatt, I. A. Walmsley, G. Stibenz, and G. Steinmeyer, "Sub-10 fs pulse characterization using spatially encoded arrangement for spectral phase interferometry for direct electric field reconstruction," *Opt. Lett.* **31**, 1914–1916 (2006).
- [10] S.-P. Gorza, P. Wasylczyk, and I. A. Walmsley, "Spectral shearing interferometry with spatially chirped replicas for measuring ultrashort pulses," *Opt. Express* **15**, 15168–15174 (2007).

- [11] R. Morita, M. Hirasawa, N. Karasawa, S. Kusaka, N. Nakagawa, K. Yamane, L. Li, A. Suguro, and M. Yamashita, *Meas. Sci. Technol.* 13, 1710 (2002).
- [12] I. A. Walmsley, "Characterization of ultrashort optical pulses in the few-cycle regime using spectral phase interferometry for direct electric-field reconstruction," in "Few-Cycle Laser Pulse Generation and its Applications," Kärtner, ed. (Springer, 2004), pp. 265–291.
- [13] M. Takeda, H. Ina, and S. Kobayashi, "Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry," *J. Opt. Soc. Am.* 72, 156-160 (1982).
- [14] H. M. Crespo, J. R. Birge, E. L. Falcão-Filho, M. Y. Sander, A. Benedick, and F. X. Kärtner, "Nonintrusive phase stabilization of sub-two-cycle pulses from a prismless octave-spanning Ti:sapphire laser," *Opt. Lett.* 33, 833-835 (2008)
- [15] J. R. Birge, H. M. Crespo, M. Sander, and F. X. Kärtner, "Non-Intrusive Sub-Two-Cycle Carrier-Envelope Stabilized Pulses Using Engineered Chirped Mirrors," in *Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies*, OSA Technical Digest (CD) (Optical Society of America, 2008), paper CTuC3.
- [16] M. Y. Sander, H. M. Crespo, J. R. Birge, and F. X. Kärtner, "Modeling of Octave-Spanning Sub-Two-Cycle Titanium:Sapphire Lasers: Simulation and Experiment," *Conference on Ultrafast Phenomena*, (European Physical Society and Optical Society of America), paper THU11c (2008).

14. Appendix A: Derivation of FFT-computed fringe phase error

[I'm having second thoughts about whether or not this is overkill. Franz, do you think this is a good idea to include?]

15. List of Figure Captions

Figure 1. Experimental schematic of 2DSI setup.

Figure 2. Frequency domain block diagram of 2DSI process.

Figure 3. (Color online) Raw 2DSI traces from octave spanning laser, with (a) extracted spectral group delay overlaid to demonstrate the interpretation of fringe offset, and (b) the same pulse with 1 mm of fused silica. The presense of extra dispersion is evident in the raw trace without any need for reconstruction.

Figure 4. Illustration of the different sideband schemes between 2DSI (top) and SPIDER (bottom), showing the two schemes for pulling the information containing sidebands out of the DC term.

Figure 5. Simulated 2DSI spectrogram (top) measured with 64 phase steps for a sinc pulse with second- and third-order dispersion and a satellite, measured in the presence of additive and shot noise, such that the resulting SNR per spectrum is 0.5. A sample reconstruction, including comparison with SPIDER is shown in the middle frame. The bottom frame shows the standard deviation of the phase measurement for both SPIDER and 2DSI, showing that the lack of delay calibration in 2DSI yields a factor of two improvement in noise performance.

Figure 6. (Color online) top: 2DSI Phase matching plot for Type II sum frequency generation for BBO cut to measure a typical few-cycle Ti:sapphire laser. The lined areas denote the phasematched regions, with each line denoting increased efficiency by 10 percent. bottom: Slices of the phase matching curves for two upconversion

wavelengths separated by 6 THz, showing the efficiency of upconversion for the two spectrally sheared components.

Figure 7. (a) Spectrum of 5 fs laser used in test; (b) extracted group delay both with and without glass slide; (c) Phase of glass slide as measured by 2DSI and as predicted by known glass dispersion; (d) Net phase delay error in glass dispersion measurement.

Figure 8. (Color online) (a) Raw 2DSI data; (b) comparison of IAC and that predicted from the 2DSI measurement; (c) Extracted spectral phase (dashed); (d) Reconstructed pulse (solid), simulated pulse (dotted) and temporal phase (dotted).

16. Figures

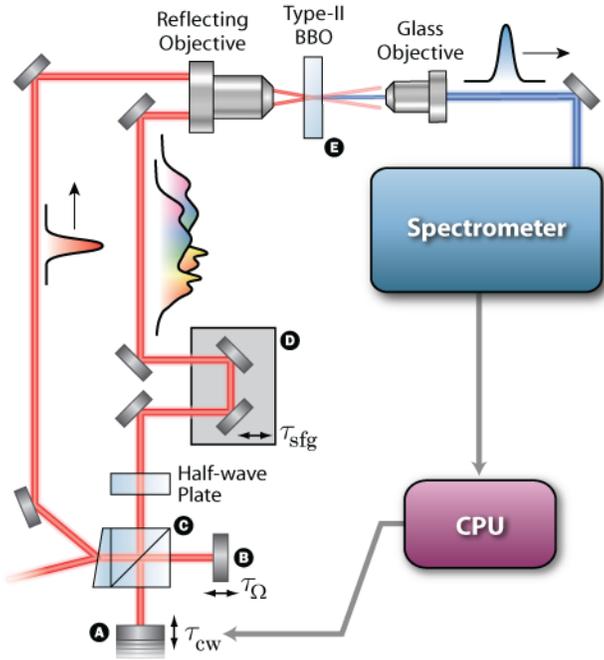


Figure 1. Experimental schematic of 2DSI setup.

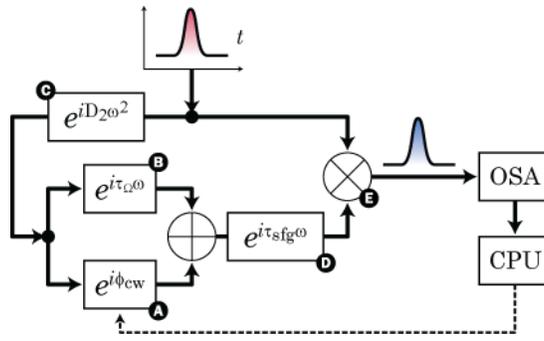


Figure 2. Frequency domain block diagram of 2DSI process.

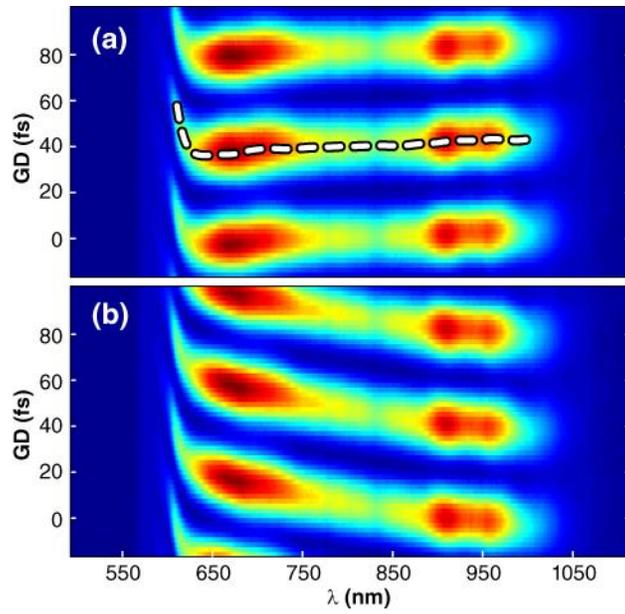


Figure 3. (Color online) Raw 2DSI traces from octave spanning laser, with (a) extracted spectral group delay overlaid to demonstrate the interpretation of fringe offset, and (b) the same pulse with 1 mm of fused silica. The presence of extra dispersion is evident in the raw trace without any need for reconstruction.

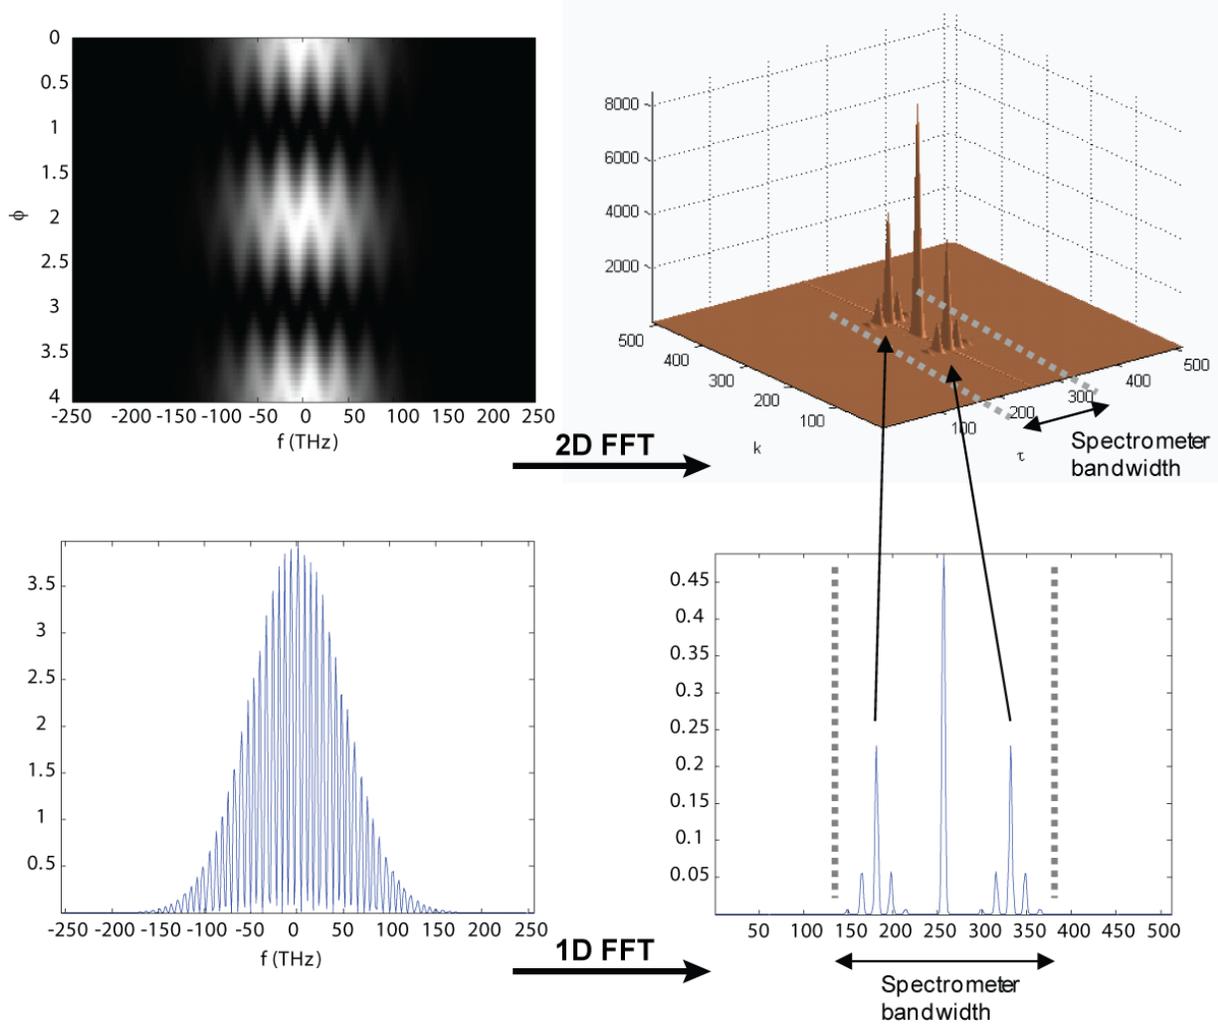


Figure 4. Illustration of the different sideband schemes between 2DSI (top) and SPIDER (bottom), showing the two schemes for pulling the information containing sidebands out of the DC term.

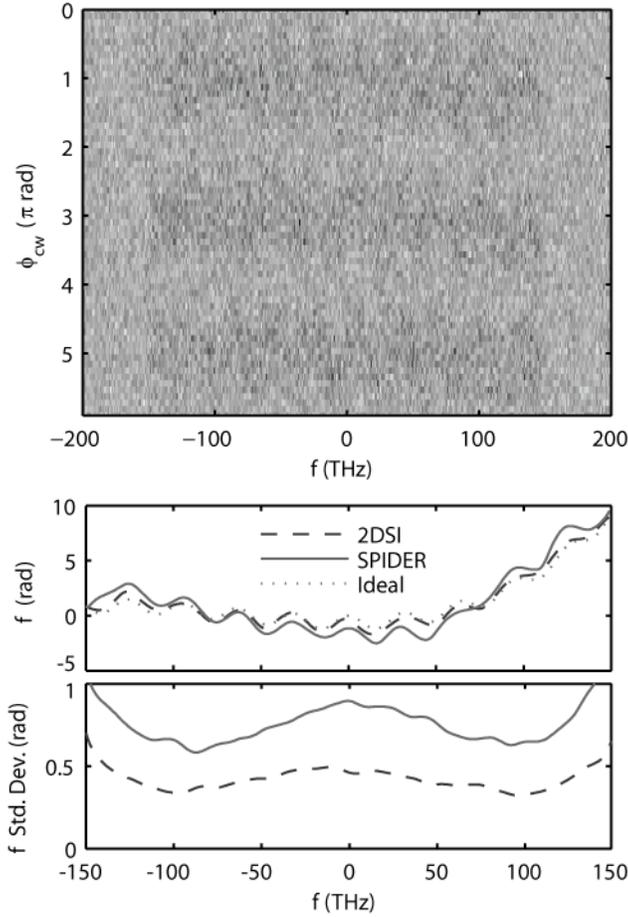


Figure 5. Simulated 2DSI spectrogram (top) measured with 64 phase steps for a sinc pulse with second- and third-order dispersion and a satellite, measured in the presence of additive and shot noise, such that the resulting SNR per spectrum is 0.5. A sample reconstruction, including comparison with SPIDER is shown in the middle frame. The bottom frame shows the standard deviation of the phase measurement for both SPIDER and 2DSI, showing that the lack of delay calibration in 2DSI yields a factor of two improvement in noise performance.

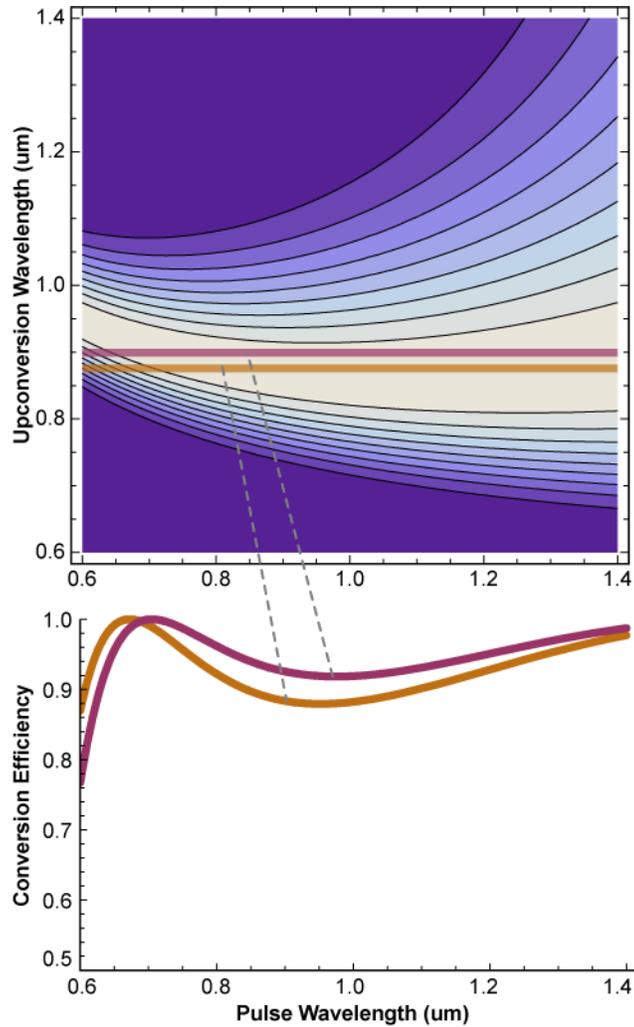


Figure 6. (Color online) top: 2DSI Phase matching plot for Type II sum frequency generation for BBO cut to measure a typical few-cycle Ti:sapphire laser. The lined areas denote the phasematched regions, with each line denoting increased efficiency by 10 percent. bottom: Slices of the phase matching curves for two upconversion wavelengths separated by 6 THz, showing the efficiency of upconversion for the two spectrally sheared components.

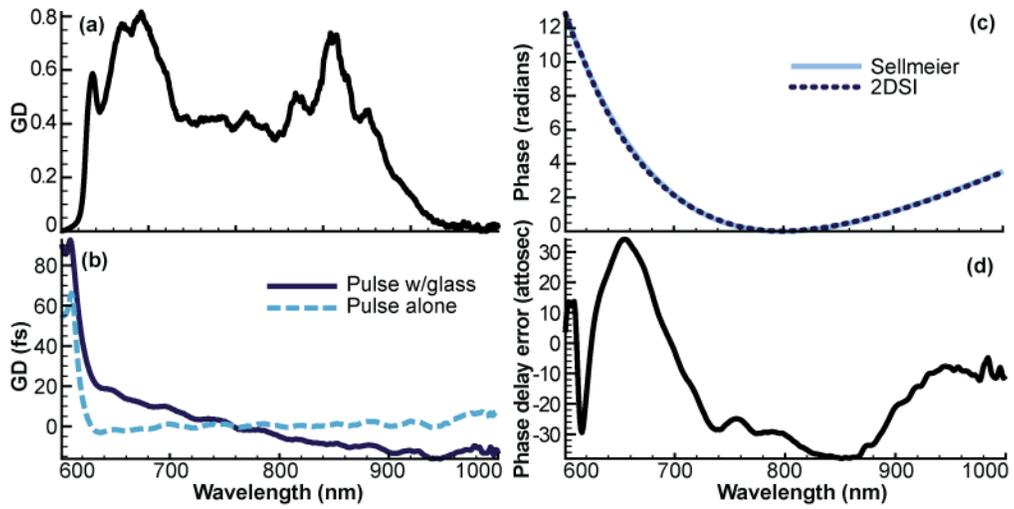


Figure 7. (a) Spectrum of 5 fs laser used in test; (b) extracted group delay both with and without glass slide; (c) Phase of glass slide as measured by 2DSI and as predicted by known glass dispersion; (d) Net phase delay error in glass dispersion measurement.

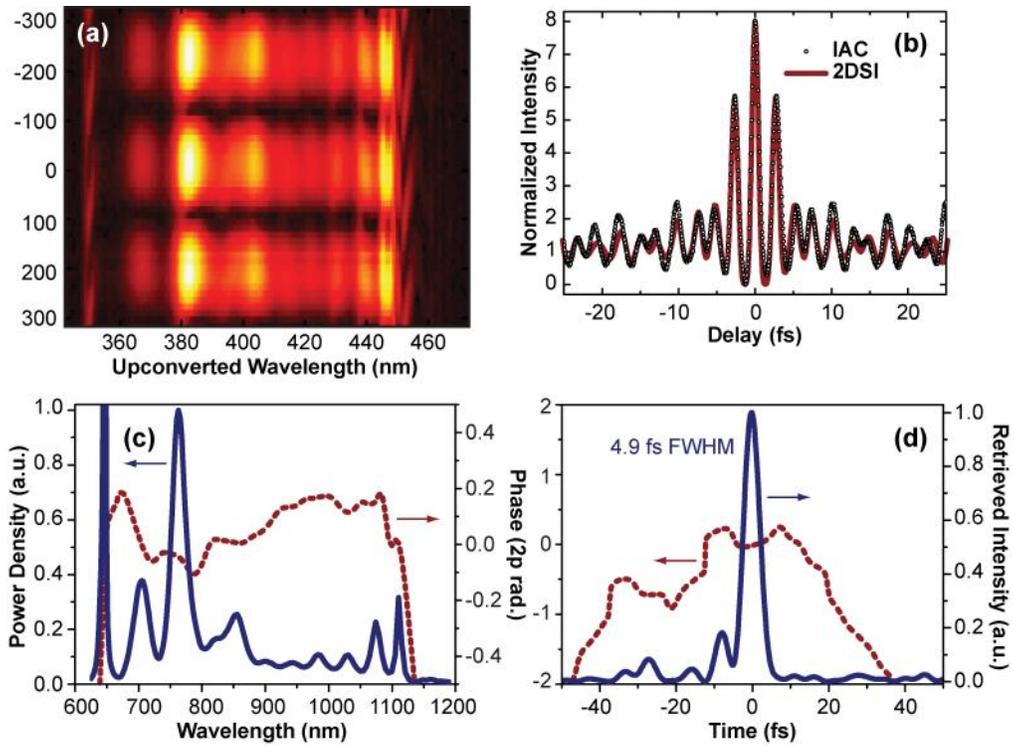


Figure 8. (Color online) (a) Raw 2DSI data; (b) comparison of IAC and that predicted from the 2DSI measurement; (c) Extracted spectral phase (dashed); (d) Reconstructed pulse (solid), simulated pulse (dotted) and temporal phase (dotted).