Appl. Phys. B 70 [Suppl.], S133–S141 (2000) / Digital Object Identifier (DOI) 10.1007/s003400000282

Applied Physics B
Lasers
and Optics

Invited paper

Minimization of dispersion in an ultrafast chirped pulse amplifier using adaptive learning

A. Efimov[1], M.D. Moores[1], B. Mei[1], J.L. Krause[2], C.W. Siders[3], D.H. Reitze[1]

[1] Physics Department, University of Florida, Gainesville, FL 32611, USA
(Fax: +1-352/392-3591, E-mail: reitze@phys.ufl.edu)
[2] Quantum Theory Project, University of Florida, Gainesville, FL 32611, USA
[3] Department of Chemistry and Biochemistry, University of California San Diego, La Jolla, CA 92093-0339, USA

Received: 2 October 1999/Revised version: 7 February 2000/Published online: 24 May 2000 –

Abstract. Minimizing residual frequency dispersion that accompanies pulse stretching, amplification, and recompression is an important consideration in ultrashort chirped-pulse amplifiers. Here we show how an adaptive learning algorithm can be used in conjunction with a pulse shaper to compensate for higher-order and nonlinear dispersion in a chirped-pulse amplifier. Using spectral blueshifting as a sensitive diagnostic for pulse shape, we implement a 'learning loop' comprised of the pulse shaper, strong field laser ionization, and a genetic algorithm to minimize dispersion through the amplifier. We verify our optimization results using frequency-resolved optical gating (FROG) measurements and also show theoretically and experimentally that spectral blueshifting is indeed a sensitive diagnostic for pulse shape, and specifically, for higher-order dispersion.

PACS: 42.65.R; 42.50.H; 33.80.R

Computer-based optimization strategies [1] coupled directly with femtosecond pulse-shaping methods [2] are becoming increasingly popular within the ultrafast laser community for controlling dispersion in lasers and chirped-pulse amplifiers. Numerous groups have reported the use of computer optimization routines coupled directly with a fs pulse-shaper for minimizing pulse duration [3–8] and for generating waveforms of arbitrary temporal shape [9, 10]. (Indeed, the reader will find many articles in this special edition devoted to adaptive learning methods in ultrafast optics.) The reasons for this increasing popularity are twofold. First, the optimization algorithms used in adaptive learning (for example, simulated annealing, evolutionary algorithms, and genetic algorithms) are able to rapidly search through a large parameter space for the optimal pulse shape. Second, these algorithms function efficiently in the presence of experimental noise such as laser intensity and phase fluctuations. The latter characteristic is particularly important in applications involving chirped-pulse amplifiers, where fluctuations in intensity and bandwidth may play a dominant role.

In this paper, we report on the use of adaptive learning for correcting higher-order dispersion in a chirped-pulse amplifier without any intentional predistortion of the pulse. We use a highly sensitive feedback signal, the spectral blueshift [11, 12] from strong field ionization, to detect changes in spectral phase of the pulse as it is being optimized. Unlike other diagnostic methods, spectral blueshifting is quite sensitive to changes in the leading edge of the pulse which accompany frequency dependent dispersion. The paper is organized as follows. In Sect. 1 we discuss in detail the experimental techniques and the implementation of the adaptive control. Our experimental phase optimization results are presented in Sect. 2, in which we display how the pulse phase evolves for different initial conditions. In Sect. 3, we present pulsewidth and phase diagnostic measurements (second-harmonic generation frequency-resolved optical gating) which corroborate our optimization results. In Sect. 4, we present a detailed discussion of the use of spectral blueshifting as a diagnostic of pulse phase. In particular, we show that the blueshift signal is highly sensitive to the sign of the cubic phase and presents an unambiguous signature for an optimized (transform limited) pulse. We conclude in Sect. 5.

1 Experimental details

1.1 Experiment description

Our experiments were performed using a multipass Nd:YAG-pumped, 10-Hz repetition rate chirped-pulse amplifier system using a hybrid double-pass pulse stretcher/shaper capable of phase compensation and pulse shaping [13]. For these experiments, 2-nJ, 15-fs pulses were selected from a Ti:sapphire oscillator and stretched to 20 ps. To provide higher order phase control, a computer-programmable, 128-pixel spatial light modulator (SLM) was placed in the Fourier plane of the stretcher. The pulses were then amplified to 2 mJ through 10 passes and then compressed to 42 fs (in the absence of any phase compensation). Combined losses from the stretcher optics and SLM resulted in a 10% throughput, thus an overall gain of 10^7 was required by the amplifier. Gain narrowing reduced the spectral width to 30–35 nm depending on amplifier saturation and a theoretical transform-limited output pulse of 28–33 fs. Spatial filters placed in each pass of the amplifier ensured minimal amplified spontaneous emission and good

beam quality at the amplifier output. After compression, 1 mJ pulse energies with 6%–10% pulse-to-pulse energy fluctuations were obtained.

To monitor the dispersion, we measured the magnitude of the spectral blueshift during strong field ionization in ambient air, nitrogen or argon. As shown in the experiments of Wood, et al., the self-phase modulation of the laser pulse resulting from the rapid creation of an electron-ion plasma leads to a time-dependent and monotonically *decreasing* index of refraction $\partial n/\partial t = -(2\pi e^2/m_e\omega^2)\partial N_e(t)/\partial t$ with a corresponding frequency shift toward the *blue* side of the spectrum [11]:

$$\Delta\omega = -\frac{\omega_0}{c}\int_0^L \frac{\partial n}{\partial t}(\omega, t, z)\,\mathrm{d}z. \tag{1}$$

For intense pulses ($> 10^{15}\,\mathrm{W/cm^2}$), time-resolved experiments have shown that the ionization predominantly occurs at early times on the leading edge of the pulse [14]. Since the magnitude of the blueshift depends on the *ionization rate* (1), we expect heuristically that optimizing the leading edge of the pulse should result in a larger ionization rate and, consequently, a larger shift in the spectrum toward the blue. It is well known that the presence of any residual higher-order phase on the pulse compromises the wing structure of the pulse on both the leading and trailing edges of the pulse [15]. Thus, monitoring the self-phase modulated spectrum of the light subsequent to an ionization event should provide a sensitive diagnostic for the presence of higher-order phase.

Linearly polarized pulses from the amplifier were focused by a 10 cm-focal-lens off-axis parabolic mirror in air or argon at atmospheric pressure to a spot diameter of approximately 20 μm and a maximum pulse intensity of $> 10^{16}\,\mathrm{W/cm^2}$. The light from the interaction region was collected using a large collection lens and sent to a diffraction grating for spectral filtering. To obtain an unambiguous signal, wavelength components in the range 450 ± 10 nm were collected and measured using a photo-multiplier tube (PMT). In the absence of any phase compensation, no light was detected in this spectral region. To minimize noise in the data due to pulse-to-pulse fluctuations in energy, spectrum, and/or temporal duration, a small amount of light was picked off before focusing and sent to a reference PMT. Because the plasma breakdown is a highly nonlinear process, we normalized the signal pulse by the square of the reference pulse energy. A schematic of the experimental apparatus is shown in Fig. 1.

Subsequent to the optimization experiments, second-harmonic generation frequency-resolved optical gating measurements [16] (SHG-FROG) were performed to assess the amount of phase compensation that resulted from the optimization. A custom single-shot apparatus was built to rapidly acquire FROG traces, autocorrelations, or spectra. Two replicas of the input beam were cylindrically focused in a vertical plane and crossed in a 100-μm KDP crystal at a 15° angle. The back surface of the crystal was imaged with magnification onto the input slit of an imaging spectrometer. All FROG traces were 5-shot averaged. Typical reconstruction errors were in the range $0.001 - 0.004$.

1.2 Phase optimization using a genetic algorithm

During each experiment, the phase is optimized by programming the SLM to provide a parameterized phase compensation of the form:

$$\Phi_{\mathrm{compenasation}}(\omega) = \frac{1}{2!}\frac{\mathrm{d}^2\Phi}{\mathrm{d}\omega^2}(\omega-\omega_0)^2 + \frac{1}{3!}\frac{\mathrm{d}^3\Phi}{\mathrm{d}\omega^3}(\omega-\omega_0)^3 + \frac{1}{4!}\frac{\mathrm{d}^4\Phi}{\mathrm{d}\omega^4}(\omega-\omega_0)^4 + \Phi_{\mathrm{NL}}(\omega-\omega_0), \tag{2}$$

where the pulse phase is defined as $\exp[-i\Phi(\omega)]$. The goal of each experiment is to determine the set of coefficients that produces the shortest pulse. The first three coefficients,

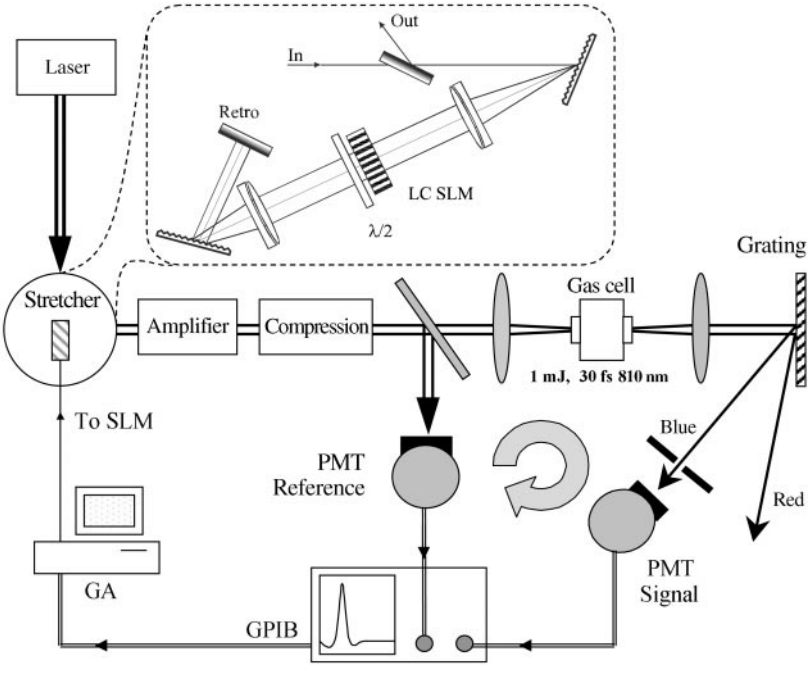

Fig. 1. Schematic layout of the phase optimization experiment. The learning loop consists of the Ti:sapphire laser, the pulse stretcher which incorporates a programmable spatial light modulator, the amplifier, the compressor, the ionization experiment, the spectral filter and detector, and the computer-based genetic algorithm which reprograms the SLM

$d^{(i)}\Phi/d\omega^{(i)}$, are the standard quadratic, cubic, and quartic higher-order dispersion terms. Numerical analysis of our amplifier leads us to believe that cubic dispersion is the dominant dispersion term. Nevertheless, we expect some quartic and nonlinear phase to be present, so we include them as well as quadratic phase in the optimization.

The last term in (2), $\Phi_{\mathrm{NL}}(\omega-\omega_0)$, reflects nonlinear phase shifts arising from the propagation of the amplified pulse through transmissive elements in the amplifier chain [17]. To a good approximation, the large chirp imposed by the stretcher linearly maps the frequency of the pulse onto time such that the temporal intensity profile of the stretched pulse closely approximates the frequency intensity profile, $I_{\mathrm{stretch}}(t) \sim I(\alpha\omega)$, where α is the quadratic chirp parameter. Thus, the self-phase modulation resulting from the amplification takes the form:

$$\Phi_{\mathrm{NL}}(t) = \frac{2\pi n_2 L}{\lambda} I(\alpha\omega), \tag{3}$$

where n_2 is the nonlinear index of refraction and L is propagation length. Here, we assume a Gaussian dependence for the frequency spectrum,

$$\Phi_{\mathrm{NL}}(\omega-\omega_0) = \Phi_{\mathrm{NL,max}} \exp\left[\frac{-(\omega-\omega_0)^2}{\Delta\omega^2}\right], \tag{4}$$

validated by our experimental measurements of amplified pulse spectra and parameterize the constants in (3) by a single phase $\Phi_{\mathrm{NL,max}}$.

The optimization was performed using GENESIS, a genetic algorithm (GA) [18]. The GA is a global optimization method that mimics the paradigms of biological evolution to efficiently search a large dimensional parameter space. The optimization begins with a random sample population of individuals (trials) comprised of the phase coefficients to be optimized, in this case the set $\{d^2\Phi/d\omega^2,\ d^3\Phi/d\omega^3,\ d^4\Phi/d\omega^4,\ \Phi_{\mathrm{NL,max}}\}$ or a subset thereof. Each coefficient is parameterized to 1024 possible values. In the GA, each parameter is represented by a "gene", an N-bit binary string that uniquely specifies its value (shown in Fig. 2). Each individual in our optimization is composed of 4 genes. A population of 50 randomly chosen individuals is selected to begin the optimization. After evaluating each individual in the initial population in the test function (in our case, by measuring the signal from the experiment), individuals are chosen to propagate to the next generation by a random selection process such that the expected number of times an individual appears in the next generation is proportional to the strength, or fitness, of that individual. A large fraction of individuals in the new generation is 'crossed' with another individual in the population. The crossover operation takes a subset of N-m bits from one parent and m bits from the second parent to form one offspring. The remaining m bits from parent 1 and N-m bits from parent 2 are combined to form a second offspring. Because particularly strong genes will be replicated many times in a subsequent generation, they may survive intact. A small fraction of the genes are mutated, i.e., a bit is flipped randomly as the propagation occurs. Mutation insures genetic diversity and guards against a particularly fit individual dominating the entire population. It also prevents the optimization from becoming trapped in a local minimum.

2 Optimization of pulse phase

2.1 Full phase optimization

In the first experiment, we programmed the laser system to optimize the full set $\{d^2\Phi/d\omega^2,\ d^3\Phi/d\omega^3,\ d^4\Phi/d\omega^4,\ \Phi_{\mathrm{NL,max}}\}$ of the parameters in the range ($\pm1,000\,\mathrm{fs}^2$, $\pm50,000\,\mathrm{fs}^3$, $\pm1,000,000\,\mathrm{fs}^4$, ±5 rad). The entire parameter space corresponds to approximately 10^{12} possible phase combinations. For each trial, four laser shots were averaged, with the GA seeking to maximize the amount of light spectrally shifted into the 440–460 nm wavelength band. The optimization proceeded until either 1000 trials had taken place (corresponding to 20 generations) or the GA determined that an optimal solution had been achieved. Each run took up to 30 min.

In Fig. 3, we display how the blueshift signal evolves as the experiment progresses. In this case, the achievement measures an increase in the blueshift signal. The squares denote the best single result in a particular generation, i.e, the largest blueshift signal corresponding to a specific set of phase

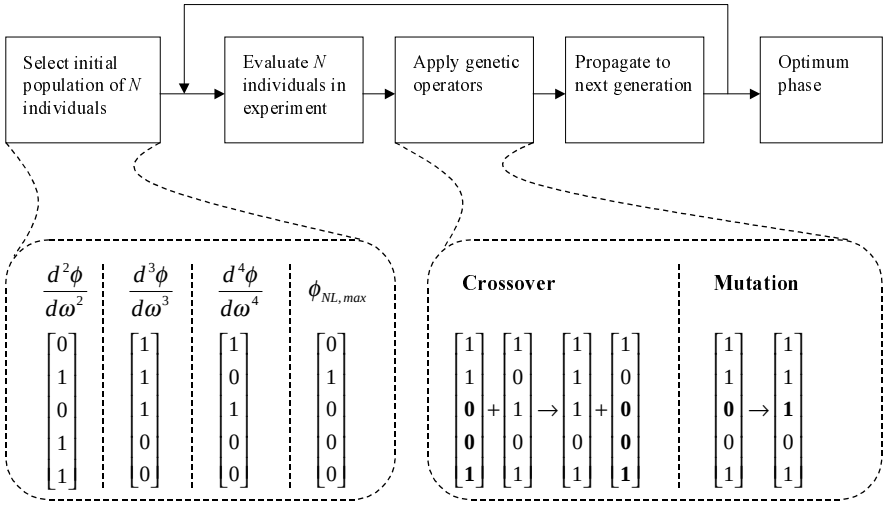

Fig. 2. The basic operation of the genetic algorithm. The phase coefficients are parameterized in N-bit strings and 'evaluated' in the test function. The genetic algorithm generates a new set of individuals using the crossover and mutation operations. The evaluation and genetic operations are continued until an optimal solution is reached

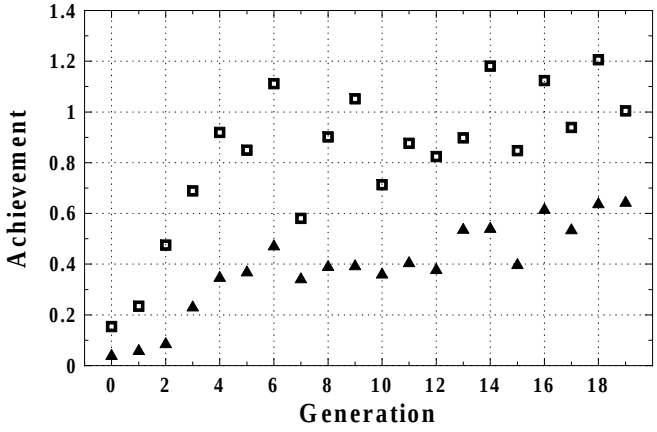

Fig. 3. The evolution of the blueshift signal versus generation when optimizing the parameter set $\{d^2\Phi/d\omega^2,\ d^3\Phi/d\omega^3,\ d^4\Phi/d\omega^4,\ \Phi_{NL,max}\}$. The best result in each generation is plotted as *squares*. The average of the blueshift signal for all individuals in the generation is plotted as *triangles*

coefficients. The triangles represent the average value of the blueshift signal for all of the trials in that generation. The average achievement monitors the rate of convergence of the optimization, increasing as the 'strong' individuals in each trial are subsequently propagated to the successive generation. In generation 0 of the experiment, one of the randomly selected individuals has increased the blueshift to a non-zero level. (No blueshift signal is observed in the absence of any phase compensation.) As the experiment progresses through 5 generations, the best blueshift signal rapidly increases in magnitude to 8 times its initial value. After the 5th generation, the blueshift increases at a slower rate with pronounced fluctuations from generation to generation. The average value also shows the same trend. After 1000 trials, the experiment is terminated by the GA, although the average achievement has not yet reached a plateau.

Figure 4 displays the evolution of the phase coefficients during the experiment. The points in each graph represent the average value of a particular phase coefficient in each generation. The error bars display the statistical standard deviation of the data set and represent the dispersion of values for a particular generation. For all of the phase coefficients, the standard deviation decreases rapidly in the first three or four generations, indicating that the GA has converged on a particular parameter that maximally enhances the blueshift signal. However, only the quadratic phase has converged to a well-defined value ($-800\,\mathrm{fs}^2$) by the termination of the experiment. All of the other terms in (2), particularly the nonlinear phase, are still varying somewhat by the end of the experiment. Mutations are evident throughout the course of the experiment, as evidenced by an increase in the standard deviation. A mutation is especially prominent in generation 15 where the changes in the quadratic, cubic, and nonlinear coefficients result in a decrease in both the best and average values.

2.2 Cubic and nonlinear phase optimization

Since we expect the cubic term to dominate the dispersion in our amplifier, a second experiment was performed

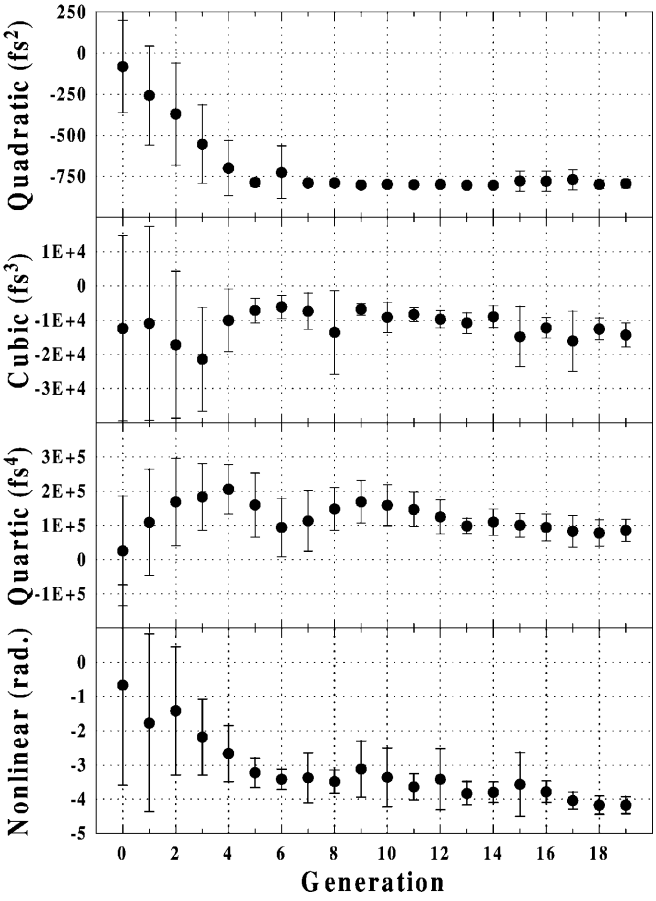

Fig. 4. The average values (*circles*) and standard deviations (*error bars*) of each of the four phase coefficients versus generation. As the experiment progresses, the standard deviation decreases significantly, as evidenced by the reduction in the error bars, indicating that the GA has converged to a particular individual (set of phase coefficients)

in which only the cubic and nonlinear phase are optimized. The best blueshift signal in this case increases to its maximum value (≈ 1.6) by generation 4 and remains constant for the duration of the experiment (shown in Fig. 5). The

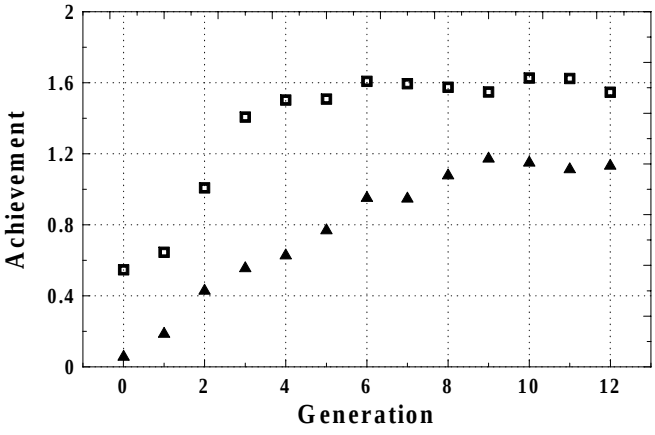

Fig. 5. The evolution of the blueshift signal versus generation when optimizing the reduced parameter set $\{d^3\Phi/d\omega^3, \Phi_{NL,max}\}$. The best result reaches a plateau at generation 4; the average value is optimized at generation 9. Both the best and average values exceed those obtained during the optimization of the set $\{d^2\Phi/d\omega^2,\ d^3\Phi/d\omega^3,\ d^4\Phi/d\omega^4,\ \Phi_{NL,max}\}$

average value also increases, reaching a plateau at generation 9. In addition, final achievement is greater than in the full phase optimization experiment (1.6 versus 1.2). The experiment terminates after 12 generations (or 600 trials) when the GA stagnates. The accompanying evolution of the cubic and nonlinear phase (shown in Fig. 6), behaves similarly to the full phase optimization case, displaying a rapid convergence to a specific value. However, unlike the previous experiment, the phase values remain relatively constant ($\mathrm{d}^3\Phi/\mathrm{d}\omega^3 \approx -22,\ 500\,\mathrm{fs}^3,\ \Phi_{\mathrm{NL,max}} \approx -1.7\,\mathrm{rad}$) after generation 4. Moreover, these values are not the same as those obtained in the full phase optimization ($\mathrm{d}^3\Phi/\mathrm{d}\omega^3 \approx -14,\ 350\,\mathrm{fs}^3,\ \Phi_{\mathrm{NL,max}} \approx -4.2\,\mathrm{rad}$).

The variability from generation to generation is markedly decreased during the reduced phase experiment when compared with the full phase experiment. There are several possible explanations for this. First, the parameter space was reduced from 10^{12} discrete points (for the full phase experiment) down to 10^6 (for the reduced phase experiment). Also, as our FROG measurements show below, cubic and nonlinear phase are the dominant dispersion terms in our amplifier. Thus in the reduced phase experiment, the GA is allowed to focus on those parameters that most affect the pulse shape. In addition, the Gaussian functional form of the nonlinear phase $\Phi_{\mathrm{NL}}(\omega-\omega_0)$ couples to the quadratic and quartic phase terms through a Taylor series expansion consisting of even powers of $\omega-\omega_0$: $\exp[-(\omega-\omega_0)^2/\Delta\omega^2] \approx 1-(\omega-\omega_0)^2/\Delta\omega^2+(\omega-\omega_0)^4/2\Delta\omega^4+\ldots$). Any changes in $\Phi_{\mathrm{NL,max}}$ can affect the quadratic and quartic phase parameters in a complicated fashion. In addition, oppositely signed quartic and quadratic terms can partially compensate each other. Thus in the full phase optimization experiment, the even-order parameters compete against each other during the optimization, possibly trapping the algorithm in a local minimum and reducing the effectiveness of the GA to converge rapidly to a global minimum. In the reduced phase optimization experiment, the parameters are independent of each other and therefore can be efficiently optimized.

2.3 Recovered CPA dispersion

The results of the experiments above enable us to reconstruct the higher-order phase inherent in the CPA, and thus provide a way of quantifying the dispersion as well as minimizing it. The frequency-dependent phase recovered for each experiment is shown in Fig. 7, which plots both the full phase optimization (open circles) and cubic/nonlinear optimization (filled circles). The amplified pulse spectrum is also shown for reference. In both cases, the experimentally determined phase has a characteristic cubic dependence, indicating that third-order phase is the dominant dispersion term. This is consistent with numerical ray-tracing results for our CPA. The intrinsic dispersion in the CPA is the additive inverse of the phases shown in Fig. 7.

It is interesting to observe that the recovered phases are quite similar in each case despite the fact that the final coefficients and, more significantly, the functional form of the optimization differed in each experiment. This may be a consequence of the relatively simple nature of the parameter space, with exactly one minimum corresponding to constant phase across the frequency spectrum. In fact, the optimal

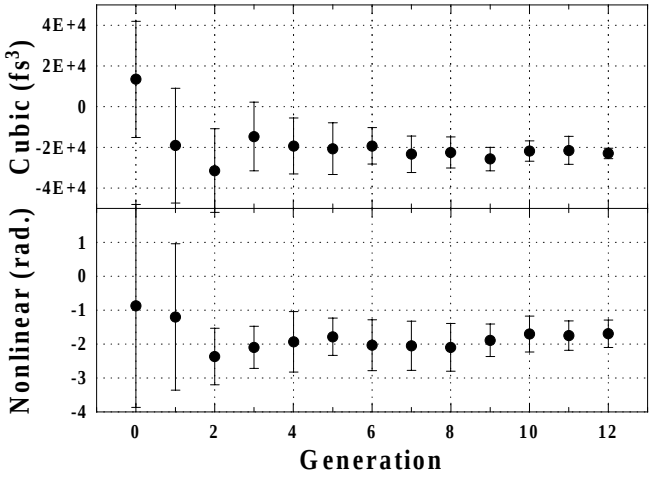

Fig. 6. The average values (*circles*) and standard deviations (*error bars*) of the cubic and nonlinear phase coefficients versus generation

phase shows a functional phase form quite similar to that of the full phase optimization, indicative of the predominance of the cubic term.

The fact that these two different runs give essentially the same results for the phase indicates that the optimization is robust. In general, cubic dispersion is dependent upon the wavelength-dependent optical path in the amplifier, which does not change significantly on a day-to-day basis. Once the optimization has been performed for a specific configuration of the amplifier, it need not be performed again until the amplifier configuration has changed or the input pulse phase has been grossly distorted.

3 Pulse width diagnostics

While the results of the prior section have shown that the GA can 'learn' a pulse phase that enhances the blueshift signal, it remains to be demonstrated that the maximal blueshift signal correlates with the shortest possible pulse. For example, because the blueshift signal is sensitive to the leading edge of the pulse, it is conceivable that optimization using strong

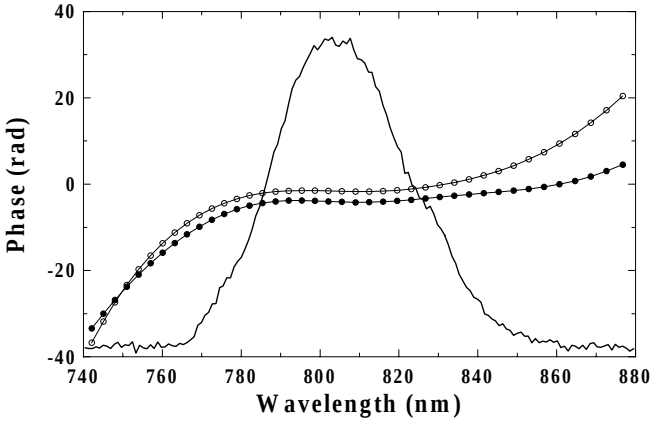

Fig. 7. The optimal phase profile that the amplifier has 'learned' from the experiment by optimizing the full set of phase coefficients (*open circles*) and cubic and nonlinear phase (*filled circles*). The amplified intensity spectrum is shown for reference

field ionization as a diagnostic merely changes the sign of the cubic phase term and optimizes the leading edge of the pulse at the expense of the trailing edge. Below, we argue, based on experimental data and theoretical modeling that optimizing the self-phase modulation during ionization indeed produces the shortest possible pulse and optimal phase for a given spectrum. In this section, we show *directly* that the solution determined by the GA satisfies our goal of producing a shorter pulse. Previously, we used slow-scan autocorrelations to demonstrate that the shorter pulses were produced subsequent to the optimization [6]. In those measurements, a reduction in pulse duration from 37 fs to 30 fs was observed for a pulse with a theoretical transform-limited duration of 28 fs. Here, we present second-harmonic generation frequency-resolved gating [16] (SHG FROG) measurements of the pulse before and after optimization which confirm that the solution found by the GA substantially improves the phase profile of the pulse.

Figure 8a shows an experimental FROG trace of an unoptimized amplified pulse with a general shape characteristic of mainly cubic phase distortions. When the optimal phase found by the GA feedback experiment is loaded into the SLM the FROG trace shown in Fig. 8b results. The reconstructed time dependence of the pulse intensity before (dashed line) and after the optimization (solid line), plotted on a logarithmic scale in Fig. 9a, illustrates the improvement in the temporal profile of the pulse after the optimization. Prior to optimization, temporal oscillations are observed on the wings of the pulse, consistent with the presence of excess cubic phase. After the pulse is compensated, a substantial improvement in the pulse shape is evident. As a result of the optimization, the duration of the pulse decreases from 42 fs to 36 fs. For these measurements, the transform limit corresponds to 33 fs, as a result of the narrower amplified spectrum than in our previous autocorrelation measurements. Figure 9b displays the reconstructed spectral phase corresponding to the unoptimized pulse (dashed line) and the optimized pulse (solid line). The intensity spectrum is shown for reference. The flattening of the phase profile after optimization is seen clearly.

Generally, the spectral phase profile as recovered from the experimental FROG traces always shows a certain degree of ripple structure not conveniently described via the Taylor series expansion (see, for example, the unoptimized phase in Fig. 9b). Compensation for this phase structure through adaptive learning is possible but will require an individual control of each pixel of the SLM increasing considerably the complexity of the parameter space for optimization, such as demonstrated by Yelin et al. [3] and Baumert et al. [4].

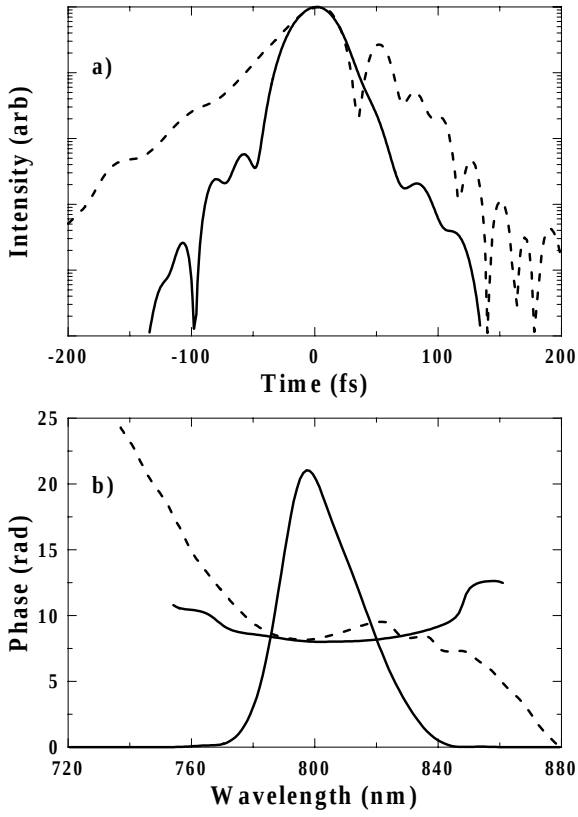

Fig. 9a,b. Reconstructed temporal intensity (**a**) and spectral phase (**b**) for the initial unoptimized pulse (*dashed curves*) and after optimization with the GA learning loop (*solid curves*). Recovered pulse intensity spectrum is superimposed onto the phase curves for reference

4 Discussion

While the FROG measurements above vindicate our selection of blueshifting as the diagnostic in this experiment, it is natural to ask why the choice was made of such an experimentally complex and potentially ambiguous measure of

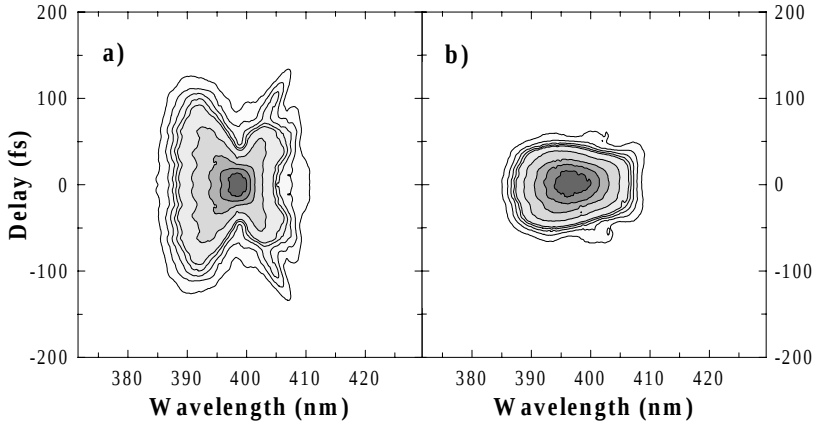

Fig. 8a,b. SHG FROG spectrograms of **a** the initial unoptimized pulse and **b** the pulse optimized by the GA feedback loop. *Contour lines* correspond to the intensity levels of 2%, 4%, 6%, 8%, 10%, 20%, 40%, 60% and 80%

pulse phase (as opposed to second-harmonic generation for instance) and what is the nature of the correlation between the molecular ionization and pulse phase that allows these experiments to work? Here, we discuss the considerations in selecting blueshifting as a useful diagnostic.

4.1 Adaptive phase compensation in CPAs: the role of noise

Under ordinary operation, a chirped-pulse amplifier produces pulses that are quadratically compensated. Pulse broadening in a CPA occurs due to the lack of compensation of higher-order phase terms. Whereas the magnitude of the temporal broadening depends on the specific configuration of the amplifier, it is typically a factor of two or less for ~ 50-fs pulses. For the CPA used in these experiments, the difference between the measured pulse duration before phase compensation and the transform-limited pulse duration is approximately 10 fs, or a 25% increase. Thus, we are primarily seeking to produce a small improvement by optimizing the cubic and higher-order phase terms.

In the absence of experimental noise, a simple diagnostic such as SHG works quite well. However, noise sources (such as pulse-to-pulse fluctuations in energy and spectrum) 'pollute' the diagnostic signal and compete against the desired optimization variable. Learning algorithms such as the GA *cannot* discriminate against experimental fluctuations if they have significant influence on the signal. Under ideal conditions (where there are no experimental fluctuations present), the increase in SHG signal is approximately $(\tau_{\text{initial}}/\tau_{\text{optimized}})^2$ where $\tau_{\text{optimized}}$ is the pulse duration after the optimization and τ_{initial} is the pulse duration at the start of the experiment. Under our experimental conditions (initial pulse duration of 42 fs, transform-limited pulse duration of 33 fs after amplification), we estimate that a SHG signal is limited to a maximum increase ≈ 1.6 with no competing noise sources present. This is in marked contrast to situations where the initial pulse is stretched to many times its transform limit [3, 4]. When noise is included, the situation changes dramatically. Pulse-to-pulse fluctuations in energy ΔE and spectral width $\Delta\nu$ also contribute to the magnitude of the SHG signal, since both affect the intensity of the pulse, $I \approx (E_{\text{pulse}} \pm \Delta E/2)(\nu_{\text{pulse}} \pm \Delta\nu/2)$, and therefore the SHG efficiency ($\approx I^2$). Here ν_{pulse} is the pulse bandwidth. RMS fluctuations of 10% in energy and 15% in bandwidth result in variations of the SHG signal by 1.6. Thus, the algorithm is trying to find the optimal SHG signal in a noisy background that fluctuates *as much as* the maximum achievable signal enhancement. Such large background fluctuations render SHG useless in this case. This was verified experimentally: we found that *SHG optimization stagnated* with no observed improvement in the SHG signal after optimization.

Since the ionization occurs well before the peak [14], the origin of the blueshifted signal makes it much more sensitive to the steepness of the pulse leading edge rather than the peak intensity of the pulse. Therefore, fluctuations in the amplifier output energy do not obscure the signal variation due to phase optimization as much as they do in the SHG experiment where both the fluctuations and the desired signal are tied to the peak intensity. Finally, the blueshift experiment as we implement it is a zero background measurement (i.e., we detect a part of the spectrum such that no signal is observed unless some phase correction is present).

4.2 The correlation between ionization and dispersion

While the above arguments suggest that blueshifting is a better optimization candidate than SHG for our experiments, we now show directly that it correlates with the dispersion of the optical pulse. In general, blueshifting by strong field ionization is a highly complex phenomenon that depends on the pulse energy and pulse shape as well as on the specific properties (composition and density) of the atomic or molecular species. There is also a spatial dependence to the self-phase shift arising from the transverse spatial profile of the beam in the interaction region which further complicates the interpretation of the data and obfuscates the connection between phase and blueshift [12]. To examine the effect of dispersion on the blueshifted spectrum, we experimentally and theoretically examine how cubic phase in particular influences the nature of the blueshifted spectrum.

Figure 10 displays the experimentally measured spectra for varying amounts of higher-order cubic phase taken under the same conditions as the optimization experiments (air at atmospheric pressure; each spectrum is an average of 100 laser shots). The cubic phase is measured with respect to the optimized phase. We stress that only the cubic phase was varied in these experiments; the pulse energy and focal characteristics were held constant to within the pulse-to-pulse fluctuations of the amplifier.

Several features are apparent in the data. In particular, the spectral weight shifts furthest toward shorter wavelengths when no residual cubic phase is present on the pulse. The presence of relatively small amounts of cubic phase *substantially* reduces the magnitude of the spectral shift toward the short wavelength end of the spectrum (defined here as the integrated photon intensity in a specific spectral bandwidth). For $\mathrm{d}^3\Phi/\mathrm{d}\omega^3 = -10\,000\,\mathrm{fs}^3$ (which broadens a 30-fs Gaussian pulse to 35 fs FWHM), the spectral weight has diminished at wavelengths shorter than 580 nm. To more clearly illustrate the recession of the shift with increasing phase, equi-intensity contours are also plotted on the graph. At each of the wavelengths, the maximum intensity corresponds to exactly the optimal cubic phase. This experimental observation further validates the use of strong field ionization as a pulse phase diagnostic. In addition, the *character* of the spectrum

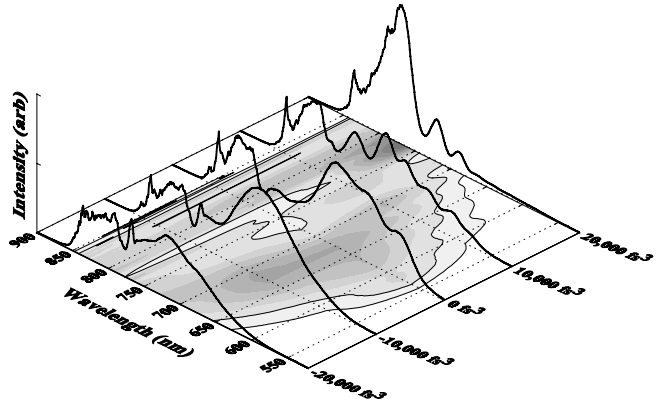

Fig. 10. The experimentally obtained intensity spectrum of the blueshifted light in air versus wavelength and the cubic phase. The optimally compressed pulse (here referred to as having zero cubic phase) displays the largest blueshift, as evidenced by the equi-intensity contours

depends qualitatively upon the *sign* of the cubic phase. For negative cubic phase, significant spectral weight has shifted away from the initial spectrum and forms a smooth, continuous distribution below 760 nm. For positive cubic terms, we observe significant oscillatory structure in the shifted spectral components with a dramatic reduction in overall magnitude of the spectral shift. With $+20\,000\,\mathrm{fs}^3$ phase on the pulse, the spectral shift has diminished by a factor of 4 at 680 nm.

These experimental results are corroborated by numerical modeling of the ionization dynamics. Recent relativistic simulations by Rau et al. using longer pulses and greater cubic dispersion (up to $10^6\,\mathrm{fs}^3$) showed qualitatively different character for blueshifted spectra depending on the sign of the cubic phase [19]. Here, we use a simpler one-dimensional model based on Maxwell's equations for pulse propagation [20] and Ammosov–Delone–Krainov (ADK) ionization [21] to model our specific experiment. Figure 11 presents numerically simulated frequency shifts resulting from the ionization of and concomitant propagation through molecular nitrogen by ultrashort laser pulses as a function of cubic phase for our pulse parameters. While the detailed quantitative nature of the simulated spectra differs from the experimental data in terms of their shifted wavelengths and exact spectral shapes, the overall agreement between these simulations and the experiment is good. (In particular, this model does not include diffraction and self-lensing effects in the ionization region caused by the transverse spatial intensity variation, which have been shown to affect the character of the spectra [12]). Three features are worth noting in comparing the numerical and experimental plots. First, the magnitude of the shift decreases significantly for larger amounts of cubic phase, which is consistent with the experimental data. For shorter wavelengths (near 600 nm), the greatest spectral weight accompanies the transform-limited pulse. Second, the presence of positive cubic phase on the pulse leads to noticeable oscillatory structure in the spectrum, again in agreement with the experimental data. These oscillations in the ionization spectra have been attributed to the temporal interference of equally spaced frequency components above and below the center frequency which have the same group delay advance. The interference of these frequency components in the time domain on the leading edge of the pulse is transferred back onto the spectrum of the pulse during the ionization process. Finally, the spectra associated with negative cubic phase exhibit a broad shifted feature at the shortest wavelengths and a sharp spike nearer to the original wavelength band. All of these features are also observed in the experimental spectra.

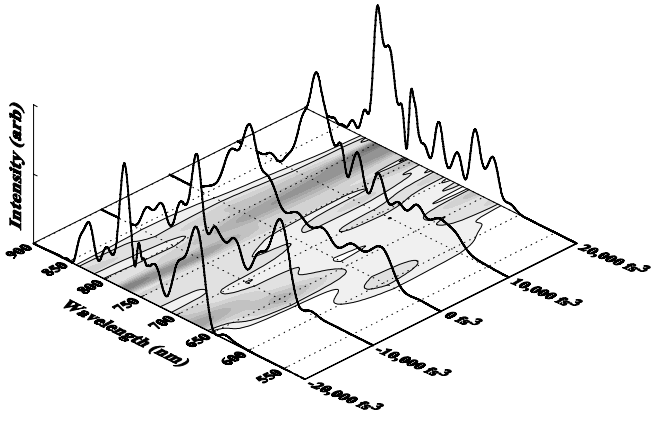

Fig. 11. The theroretical blueshift spectrum versus cubic phase for N_2 ionization and the conditions used in the optimization experiments. The spectra were computed using a one-dimensional propagation and ADK ionization code

The previous results demonstrate that if the propagation length does not change dramatically as the cubic phase is changed over a range of $40\,000\,\mathrm{fs}^3$ (as happens in our experiments), the detection of blueshifted light at the very edge of the spectrum provides a sensitive diagnostic of spectral phase. The question then becomes: how much do the propagation dynamics change? It could be argued that pulse shaping distorts the spatial mode of the pulse, resulting in spatial profiles that change with the amount of phase present on the pulse. Such effects have been examined extensively by Wefers et al. [22] and shown to be important when the shaped pulse is of significantly longer duration than the input pulse. We have recently shown, however, that the spatial mode of the pulse does not change for small amounts ($< 80\,000\,\mathrm{fs}^3$) of cubic phase imparted by the SLM in the stretcher [13]. Thus, any change in interaction length must be due to the intrinsic strong field ionization dynamics. This could possibly occur if the small change in pulse duration resulting from the phase compensation leads to a subsequent ionization event (for example, $N_2^j \rightarrow N_2^{j+1}$) which changes the radial profile of the index of refraction and thus the propagation length. However, at our intensities ($I > 10^{16}\,\mathrm{W/cm}^2$), the molecules are already at a relatively high ionization state. Thus, an ionization event at these intensities causes a small perturbation on the index of refraction and therefore on the defocusing and interaction length. We believe therefore that the propagation dynamics are not significantly affected.

5 Conclusions

We have demonstrated that adaptive learning algorithms can be used for compensating higher-order dispersion in chirped-pulse amplifiers. Using a sensitive diagnostic signal, spectral blueshifting, to monitor the character of the pulse, we have shown that it is possible to decrease the amplified pulse duration and compensate for higher-order dispersion. We anticipate that learning algorithms will play an important role in controlling the shape of ultrashort amplified fs laser pulses.

Acknowledgements. The authors would like to especially thank F. Omenetto for useful discussions and advice on the subtleties of phase retrieval from FROG codes. This research was supported by the NSF through the National High Magnetic Field Lab and by the Research Corporation.

References

1. See, for example, R.L. Haupt, S.E. Haupt: *Practical Genetic Algorithms* (Wiley, New York 1998)
2. A.M. Weiner, D.E. Leaird, J. Patel, J. Wullert: IEEE J. Quantum Electron. **QE-28**, 908 (1992)
3. D. Yelin, D. Meshulach, Y. Silberberg: Op. Lett. **22**, 1793 (1997)
4. T. Baumert, T. Brixner, V. Setfried, M. Strehle, G. Gerber: Appl. Phys. B **65**, 779 (1997)
5. D. Zeidler, T. Hornung, M. Motzkus: Appl. Phys. B, Suppl. to **70**, 125 (2000)
6. A. Efimov, M.D. Moores, N.M. Beach, J.L. Krause, D.H. Reitze: Opt. Lett. **23**, 1915 (1998)

7. F.G. Omentto, B.P. Luce, A.J. Taylor: Appl. Phys. B, Suppl. to **70**, 143 (2000)
8. T. Brixner , M. Strehle, G. Gerber: Appl. Phys. B **68**, 281 (1999)
9. D. Meshulach, D. Yelin, Y. Silberberg: J. Opt. Soc. Am. B **15**, 1615 (1998)
10. T. Brixner, A. Oerhlein, M. Strehle, G. Gerber: Appl. Phys. B, Suppl. to **70**, 119 (2000)
11. M. Wood, C.W. Siders, M.C. Downer: IEEE Trans. Plasma Sci. **21**, 20 (1993)
12. S.P. LeBlanc, R. Sauerbrey: J. Opt. Soc. Am. B **13** 72 (1996)
13. A. Efimov, D.H. Reitze: Opt. Lett. **23**, 1612 (1998)
14. M.C. Downer, W.M. Wood, J.J. Trisnadi: Phys. Rev. Lett. **65**, 2832 (1990)
15. A. Sullivan, W.E. White: Opt. Lett. **20**, 192 (1995)
16. K.W. Delong, R. Trebino, J. Hunter, W.E. White: J. Opt. Soc. Am. B **11**, 2206 (1994)
17. A. Braun, S. Kane, T. Norris: Opt. Lett. **22**, 615 (1997)
18. J.D. Grefenstette: GENESIS v5.0, Copyright ©1990, available at http://www.aic.nrl.navy.mil/galist/src/
19. B. Rau, C.W. Siders, S.P. Le Blanc, D.L. Fisher, M.C. Downer, T. Tajima: J. Opt. Soc. Am. B **14**, 643 (1997)
20. B.M. Penetrante, J.N. Bardsley, W.M. Wood, C.W. Siders, M.C. Downer: J. Opt. Soc. Am. B **14**, 2032 (1992)
21. M.V. Ammosov, N.B. Delone, V.P. Krainov: Sov. Phys. JETP **64**, 1191 (1986)
22. M.M. Wefers, K.A. Nelson: J. Opt. Soc. Am. B **12**, 1343 (1995)