Adaptive Laser Pulse Shaping using an Acousto-Optic Programmable Dispersive Filter and a Genetic Algorithm

Master's thesis
by
Per Johnsson

Lund Reports on Atomic Physics, LRAP-291
Lund, November 2002

Abstract

In this work an Acousto-Optic Programmable Dispersive Filter has been used to temporally shape femtosecond pulses from the oscillator of the T^3-laser in Lund. The actual shaping has been done through a feedback loop, controlled by a genetic algorithm implemented as part of the work. The goal was to optimize the efficiency for second harmonic generation, thus expecting to minimize the pulse length. The diagnostics have been made through measurements of the pulse length using a multi-shot autocorrelator and measurements of the spectral phase using Spectral-Phase Interferometry for Direct Electric field Reconstruction (SPIDER). The experimental investigations have shown that the algorithm works, finding pulses with a pulse length of a few percent above the transform-limit. During the work experience has been gained on the pulse shaper as well as on the practical application of a genetic algorithm in a feedback loop.

Contents

Chapter 1

Introduction

At the present time laser technology has advanced to the point where laser pulses in the femtosecond range are available from easy-to-use table-top devices. While laser researchers now strive to reach the attosecond domain, femtosecond technology constantly finds new applications in both research and real-world systems. Previously liquid dye lasers were most often used for femtosecond pulse generation. Advances in solid-state laser technology have now made these lasers more common. The advantages of solid-state lasers over their dye laser counterparts are among other things an increase in both output power and stability.

The attractiveness of ultrashort laser pulses lies first in the possibility to observe and manipulate processes on a femtosecond time scale. Ten femtoseconds correspond to around five optical cycles in the visible regime and the geometrical length of such a pulse is only 3 μm in vacuum. For comparison the period for the motion of an electron in a hydrogen atom, excited to the state with quantum number $n = 10$, is around 150 fs in a semi-classical picture [1]. With the short duration of a femtosecond pulse, it is possible to probe very fast processes. An important application is femtochemistry, where ultrashort pulses are used to diagnose and control chemical reactions. The large bandwidths associated with short pulses are also expected to be useful in data-transfer technology.

In a femtosecond laser pulse, energy can be concentrated into a very short time interval, causing pulses with rather modest pulse energies to exhibit high peak powers. As an example a 50 fs laser pulse with an energy of 1 mJ exhibits a peak power of 20 GW. If this pulse is focused to a 100 μm^2 spot, the peak intensity will be 10 PW/cm^2, corresponding to an electric field strength of around 3 GV/cm. For comparison a typical inner-atomic field has a strength of 1 GV/cm. The high peak-powers associated with ultrashort laser pulses are thus comparable to typical inner-atomic fields, implying interesting applications within strong-field atomic physics. If an ultrashort laser pulse is focused onto a solid-state target, i.e. a metal, it is possible to generate broadband X-rays which may be useful in for example various medical applications [2]. If instead the pulses are focused into a jet of inert gas, high harmonic generation can be achieved, an application

for which this work can be said to be a pre-study.

In many of the applications mentioned above, it is important not only to have short pulses, but also to be able to measure and control their characteristics. If an intense laser pulse suffers from pre-pulses, say only with a peak intensity of a few percent of the main pulse, the energy in those may still be high enough to for example ionize an atom, making any examination of the atom in question impossible. Many experiments, especially those using solid targets, require elimination of the pre-pulses. In femtochemistry, where a laser pulse is used to control a chemical reaction, the pulse shape determines which chemical reaction that takes place. Thus with the ability to control the pulse characteristics, it is possible to synthetically produce molecules of a certain kind. For the high harmonic generation experiments mentioned above, it has been shown [3] that the pulse characteristics of the exciting pulse have influence on the phase of the generated harmonics. Other experiments have demonstrated that the pulse shape also affects the relative intensities of the different harmonic orders [4]. Through constructive interference between the different harmonics, it is believed that attosecond pulses can be generated. Thus the control of the phase of the harmonics, through pulse shaping of the exciting pulse, is an important topic in current attosecond pulse research. Pulse shaping is a growing field within short pulse laser research and its applications.

Pulse shaping can be done in many ways. Traditionally it is done by a technique known as spatial masking, where a dispersive element is used to transform the frequency spectrum of a pulse into a spatial distribution, which can be filtered by a spatial mask. This setup is often called a zero dispersion-line, since if no mask is used, the pulse characteristics will be unchanged. The spatial mask used as the frequency filter may be of different types. First of all it is possible to use a fixed mask, affecting the amplitude (by variation of the transmission) or the phase (by variation of the refractive index) for the spectral components of the pulse. It is also possible to do a combination of amplitude and phase filtering. Recent experiments have payed increasing attention to the use of programmable filters, referred to as adaptive pulse shaping. Those make it possible to change the filtering during the experiment, allowing also for automatic control based on feedback from the experiment. The most common programmable filters today are implemented as liquid-crystal arrays, using the same technique as for liquid-crystal displays. By applying an electric voltage to a specific array element (or pixel), the transmission or the refractive index of the array element may be varied. To obtain both amplitude and phase filtering a combination of LC-arrays has to be used. Another common implementation is with an Acousto-Optic Modulator (AOM) operated in the Bragg-regime. However such a device can not be used in a zero dispersion-line as described above and the setup has to be slightly modified.

In 1997, P. Tournois [5] proposed a novel design for a pulse shaper using acousto-optic interaction in a crystal to obtain collinear pulse shaping. Collinear does in this case mean that the shaped pulses will propagate in the same direction as the input pulses. This device does not require a zero dispersion-line to function and by the collinear arrangement it is fairly easy to implement the pulse shaping in an existing setup. The device, called the Dazzler, is an Acousto-Optic

Programmable Dispersive Filter (AOPDF) and will be used in the present study.

This work will investigate the possibilities to use the Dazzler AOPDF for adaptive pulse shaping. In the future the device is to be used at the 1 kHz laser system in Lund, shaping the pulses used to generate high harmonics in a gas jet. The approach is to control the pulse shaper using a genetic algorithm, which is a search method belonging to the field of evolutionary programming. The algorithm will perform the search based on some feedback from the experimental setup. Previously Reitze *et al.* has reported optimization of high order harmonic tuning using a genetic algorithm [6].

This pre-study will however not be made on the 1 kHz-system, but instead the pulses from an oscillator producing laser pulses in the femtosecond range will be used. The goal is to optimize the efficiency for second harmonic generation, which is in general equivalent to minimizing the pulse length. Hence it is expected that as short pulses as possible will be achieved from the optimization.

This report is organized as follows. In chapter 2 the theory of short laser pulses is presented. Emphasis will be put on the frequency-domain description, useful to understand the concept of pulse shaping. The process of second harmonic generation will be discussed, concentrating on the efficiency for second harmonic generation, since this is used as a feedback signal during this work. This chapter also presents some techniques used in the study to characterize ultrashort laser pulses. In chapter 3 the theory of genetic algorithms is presented, and the actual implementation done for this work is described. Finally, chapter 4 focuses on the experimental work done. After a description of the setup and the experimental methods, the main results will be presented and discussed.

Chapter 2

Short laser pulses

2.1 Mathematical description of short pulses

This chapter will be concerned with the mathematical description of short laser pulses. It will also treat nonlinear interaction, characterization of short pulses and pulse shaping.

2.1.1 Time-domain description

A laser pulse is often described by its complex representation $\tilde{E}(t)$ expressed as the product of an amplitude function and a phase term:

$$\tilde{E}(t) = A(t)e^{i\Gamma(t)} \tag{2.1}$$

The real electric field can then be written as $E(r) = \text{Re}\ \tilde{E}(t)$. In most cases the frequency components of the pulse are centered around a carrier frequency ω_L so that

$$\tilde{E}(t) = A(t)e^{i\varphi(t)}e^{i\omega_L t} = \tilde{A}(t)e^{i\omega_L t} \tag{2.2}$$

where $\varphi(t)$ is the time dependent phase of the pulse. $A(t)$ is called the field envelope and $\tilde{A}(t)$ the complex field envelope. The instantaneous frequency $\omega(t)$ of the pulse is given by the time derivative of the phase function $\Gamma(t)$:

$$\omega(t) = \frac{d\Gamma(t)}{dt} = \omega_L + \frac{d\varphi(t)}{dt} \tag{2.3}$$

If the phase across the pulse is constant , $d\varphi/dt = 0$, the instantaneous frequency is equal to the carrier frequency. If the phase of the pulse is linearly increasing (or decreasing) with time, the instantaneous frequency is still constant but no longer equal to the carrier frequency. A higher order phase implies that $d\varphi/dt$ is a function of time leading to a varying instantaneous frequency along the pulse, called a chirp. If $\omega(t)$ is increasing the chirp is said to be positive, and if it is decreasing, the chirp is negative.

2.1.2 Frequency-domain description

The temporal and spectral properties of a laser pulse are related through a Fourier-transform relation. More precisely, the spectral electric field $\tilde{E}(\omega)$ is the Fourier transform of the temporal electric field $E(t)$. Since $E(t)$ is a real quantity, the properties of the Fourier transform imply that $\tilde{E}(\omega) = \tilde{E}^*(-\omega)$ making it possible to describe the spectral electric field as $\tilde{E}(\omega) = \tilde{A}(\omega) + \tilde{A}^*(-\omega)$ where $\tilde{A}(\omega)$ is the complex spectral amplitude, being nonzero only for positive frequencies. The complex spectral amplitude is given by the Fourier transform of the complex temporal electric field

$$\tilde{A}(\omega) = \int_{-\infty}^{+\infty} \tilde{E}(t) e^{-i\omega t} dt \tag{2.4}$$

while the complex temporal electric field is in turn given by the inverse Fourier transform of the complex spectral amplitude:

$$\tilde{E}(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \tilde{A}(\omega) e^{i\omega t} d\omega \tag{2.5}$$

As for the complex field envelope, the complex spectral amplitude $\tilde{A}(\omega)$ might be written as the product of an amplitude function and a phase term

$$\tilde{A}(\omega) = A(\omega) e^{-i\varphi(\omega)} \tag{2.6}$$

where $A(\omega)$ is the spectral amplitude and $\varphi(\omega)$ is the spectral phase. Here it can be seen that a constant phase will only lead to a phase shift, not affecting the pulse in any way, while a phase function that is linearly dependent on frequency will displace the pulse in time, not affecting the pulse shape. If the spectral amplitude is large over a sufficiently broad range, it is necessary to regard higher-order frequency-dependence of the phase function expanded around the carrier frequency ω_L as $\varphi(\omega) = \varphi_0 + \varphi_1 (\omega - \omega_L) + \frac{\varphi_2}{2} (\omega - \omega_L)^2 + \ldots$. Different spectral regions of the pulse are then affected differently, changing the shape of the pulse envelope and introducing a chirp. From this discussion one can also infer that the situation in which all high-order coefficients $(\varphi_n, n \geq 2)$ are zero, represent the shortest pulse achievable for a given spectral amplitude distribution $A(\omega)$.

2.1.3 Pulse propagation

Propagation along the z-axis is included by generalizing (2.5) as

$$\tilde{E}(t, z) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \tilde{A}(\omega) e^{i(\omega t - k(\omega) z)} d\omega \tag{2.7}$$

where $k(\omega) = n(\omega)\omega/c$ is the propagation constant with $n(\omega)$ as the frequency-dependent index of refraction and c the velocity of light in vacuum. The effect of the propagation constant is that it introduces an accumulated phase $k(\omega)z$ to the pulse as it propagates. In cases when the spectral amplitude is sharply peaked

around the central frequency ω_L it is possible to approximate the wavenumber $k(\omega)$ with a first-order Taylor expansion around ω_L:

$$\begin{aligned} k(\omega) &= k(\omega_L) + \left(\frac{\partial k}{\partial \omega}\right)_{\omega_L} (\omega - \omega_L) + ... \\ &= k_L + k'_L (\omega - \omega_L) + ... \end{aligned} \tag{2.8}$$

Inserting this expansion for $k(\omega)$ into (2.7) yields

$$\tilde{E}(t,z) = \frac{e^{i(\omega_L t - k_L z)}}{2\pi} \int_{-\infty}^{+\infty} \tilde{A}(\omega) e^{i(\omega-\omega_L)\left(t - k'_L z\right)} d\omega \tag{2.9}$$

and since the integral in (2.9) is a function of the composite variable $t - k'_L z$ only, this in turn might be rewritten as

$$\tilde{E}(t,z) = \tilde{A}(t - k'_L z) e^{i(\omega_L t - k_L z)} \tag{2.10}$$

where $\tilde{A}(t)$ is the complex field envelope.

Group velocity

Comparing (2.10) with (2.2) it can be seen that the complex field envelope moves without changing its shape. The velocity of the pulse given by

$$v_g = \frac{dz}{dt} = \frac{1}{k'_L} = \left[\left(\frac{\partial k}{\partial \omega}\right)_{\omega_L}\right]^{-1} \tag{2.11}$$

is called the group velocity of the pulse and represents the energy transport of the pulse. In terms of the refractive index, the group velocity may be written as:

$$v_g = \frac{c}{n + \omega \left(\partial n / \partial \omega\right)} \tag{2.12}$$

Here it can be seen that only in the case when the refractive index is frequency-independent ($\partial n/\partial \omega = 0$), the group velocity of a pulse is the same as the phase velocities of the individual modes ($v_p = c/n$). When $\partial n/\partial \omega > 0$, a material is said to exhibit normal dispersion. This is the case for most materials and as a result, the group velocity is generally smaller than the individual phase velocities.

Group velocity dispersion

If the spectral bandwidth $\Delta\omega_L$ of the pulse is very broad[1], the expansion (2.8) may no longer be a good approximation for the dispersion relation. In this case different frequency components of the pulse will move with different group

[1] The spectral bandwidth $\Delta\omega_L$ is often chosen as the FWHM of the spectral intensity.

velocities, causing a temporal broadening of the pulse. By including the second order term in the approximation (2.8) it becomes

$$k(\omega) = k_L + k'_L(\omega - \omega_L) + \frac{1}{2}k''_L(\omega - \omega_L)^2 + ... \tag{2.13}$$

The time necessary for a pulse to go through a material with propagation constant $k(\omega)$ and length l is $\tau_p = l/v_g = lk'_L$ and thus the difference in time delay between the different frequency components contained in $\Delta\omega$ is

$$\Delta\tau_p \approx l\left|k''_L\right|\Delta\omega_L \tag{2.14}$$

describing the pulse broadening. The quantity $k''_L = \partial^2 k/\partial\omega^2$ is often referred to as the group velocity dispersion (GVD) of a material. It has the unit $[s^2/m]$ and describes the pulse broadening per unit bandwidth and unit material length for propagation of a pulse through a medium[2]. The pulse broadening originating from GVD is due to a spreading along the propagation axis of the different frequencies. For example, if k''_L is positive the group velocity of a low frequency component is higher than that of a high frequency component, causing the low frequency components to spread towards the trailing edge of the pulse. Because of this, the instantaneous frequency of the pulse will increase along the pulse giving a positive chirp (see section 2.1.1). Most materials have $k''_L > 0$, and are then said to exhibit positive dispersion.

Higher order effects

The effects of higher order coefficients in the dispersion relation (2.13) are generally dominated by the group velocity dispersion. However, for very short pulses (large bandwidths $\Delta\omega_L$) or if the GVD for some reason is very small, the quadratic approximation fails to describe the pulse propagation correctly. In these cases it might be necessary to also include third and fourth order dispersion (TOD and FOD). These effects do not simply broaden the pulse, but might also introduce some internal structure, like pre- and post-pulses.

2.1.4 Gaussian pulses

To get a better understanding of the properties of a laser pulse it is useful to consider a known pulse envelope. A Gaussian pulse shape is often used as an approximation, mostly because it is preserved under Fourier transformation. In many actively mode-locked systems, the generated pulses have a shape very close to Gaussian [7]. In passively mode-locked systems, the steady-state solution for the pulse shape is a hyperbolic secant function [8], somewhat harder to treat than the Gaussian.

Considering a Gaussian pulse with a complex field envelope of the form

$$\tilde{A}(t) = A(t)e^{i\varphi(t)} = e^{-at^2}e^{ibt^2} \tag{2.15}$$

[2]The concept of GVD is straightforward only for a homogenous medium. When talking about compound optical elements exhibiting dispersion, instead the group delay dispersion (GDD) is used, having the unit $[s^2]$.

the pulse length τ_p, chosen as the FWHM of the temporal intensity, is related to the parameter a by

$$\tau_p = \sqrt{\frac{2\ln 2}{a}} \tag{2.16}$$

while the phase function $\varphi(t) = bt^2$ introduces a quadratic temporal phase. Using equation (2.3) it is seen that the instantaneous frequency of this pulse is given by $\omega(t) = \omega_L + 2bt$. Thus b is a chirp parameter, determining the magnitude of the linear chirp of the pulse. To obtain the complex spectral amplitude of this pulse, one has to take the Fourier transform of (2.15) (see equation (2.4)). In this case, this is particularly easy, since a Gaussian simply transforms into another Gaussian. Ignoring a scaling constant and a constant phase term the complex spectral amplitude gets the form:

$$\tilde{A}(\omega) = e^{-\frac{(\omega-\omega_L)^2}{4(a-ib)}} = e^{-\frac{a(\omega-\omega_L)^2}{4(a^2+b^2)}} e^{-i\frac{b(\omega-\omega_L)^2}{4(a^2+b^2)}} = A(\omega)e^{-i\varphi(\omega)} \tag{2.17}$$

From the spectral amplitude $A(\omega)$ an expression for the bandwidth $\Delta\omega_L$ can be obtained. The bandwidth is normally chosen as the FWHM of the intensity spectrum.

$$\Delta\omega_L = \sqrt{\frac{8\,(a^2+b^2)\ln 2}{a}} = \sqrt{8\ln 2}\sqrt{a\big(1+\frac{b^2}{a^2}\big)} \tag{2.18}$$

From this expression it is seen that for a given pulse length determined by the parameter a according to (2.16), the presence of a linear chirp determined by the parameter b requires a wider spectral bandwidth $\Delta\omega_L$. In the same way, for a given spectral bandwidth $\Delta\omega_L$, the presence of a linear chirp causes the pulse to spread out in time increasing the pulse length τ_p.

From the spectral phase in expression (2.17) it is seen that a linear frequency chirp, caused by a quadratic temporal phase, also implies a quadratic spectral phase.

The time-bandwidth product

Something that is often considered when it comes to ultrashort laser pulses is the time-bandwidth product. If (2.16) and (2.18) are multiplied together one obtains

$$\tau_p \cdot \Delta\omega_L = 2\pi \cdot 0.441 \cdot \sqrt{1+\frac{b^2}{a^2}} \geq 2\pi \cdot 0.441 \tag{2.19}$$

which shows that there is a minimum value for the product between the pulse duration and the spectral bandwidth. A pulse for which the equality in (2.19) holds is indeed the shortest possible pulse given that spectral bandwidth. Such a pulse exhibits no frequency chirp and is said to be transform-limited.

The properties of the time-bandwidth product holds for all pulse shapes, although the minimum value differs. As shown above, for a Gaussian pulse the value is 0.441. For the other common pulse shape, the hyperbolic secant, the corresponding value is 0.315 [9].

2.2 Nonlinear optical processes

The propagation of electromagnetic waves in material is governed by Maxwell's equations [10] from which it is possible to obtain wave equations for both the electric and the magnetic field. When an electric field is present in a medium, a displacement of charges within the material is obtained. This displacement is described by the electric polarization P. The electric polarization acts as a source term in the wave equation for the electric field, causing radiation to be emitted at the frequency of the polarization wave. Under moderate field strengths the relationship between the polarization and the electric field can be assumed to be linear, causing emission of electromagnetic radiation at the same frequency as the original field. This is the process behind the propagation of electromagnetic waves in a medium, with the refractive index determined by the dependence of the polarization on the electric field.

When the field is sufficiently strong, the linear approximation of the electric polarization fails, making it necessary to consider higher order effects.

2.2.1 Second harmonic generation

One particular nonlinear process, commonly used in various applications of laser optics, is second harmonic generation (SHG) or frequency doubling. To study this, one may consider the first nonlinear term of the electric polarization

$$P^{NL} = \varepsilon_0 \chi^{(2)} E^2 + ... \tag{2.20}$$

where $\chi^{(2)}$ is called the second order nonlinear susceptibility and ε_0 is the permittivity of vacuum. Consider an incident one dimensional plane wave of frequency ω

$$E_\omega(t,z) = \frac{\tilde{A}_\omega}{2} e^{i(\omega t - k_\omega z)} + c.c. \tag{2.21}$$

where $c.c.$ denotes the complex conjugate. Using relation (2.20) this results in a polarization wave

$$P^{NL}(t,z) = \frac{\varepsilon_0 \chi^{(2)} \tilde{A}_\omega^2}{4} \left[e^{i(2\omega t - 2k_\omega z)} + 1 \right] + c.c. \tag{2.22}$$

Apart from the DC term, equation (2.22) describes a polarization wave $P_{2\omega}$ with frequency 2ω propagating at the phase velocity $v_p^P = \omega / k_\omega$. Since the origin of the electric polarization is a displacement of the electric charges within the material, at each point the displaced charges oscillates and radiates with a frequency 2ω. The electric field produced by the polarization at one particular point in the material can then be described as

$$E_{2\omega}(t,z) = \frac{\tilde{A}_{2\omega}}{2} e^{i(2\omega t - k_{2\omega} z)} + c.c. \tag{2.23}$$

which is a plane wave of frequency 2ω propagating with the phase velocity $v_p^E = 2\omega / k_{2\omega}$. This shows that in the presence of a second order polarization, a

sufficiently strong elecric field will give rise to an electric field oscillating at the double frequency.

When doing frequency doubling, another important factor needs to be considered. If the phase velocities of the polarization wave and the frequency-doubled electric wave differs, the radiated fields at different positions along the propagation direction will be out of phase. This causes the generated frequency doubled light to increase and decrease in intensity during its passage through the medium. At a position where $P_{2\omega}$ is exactly out of phase by a whole cycle with $E_{2\omega}$ at the entrance of the medium, the frequency-doubled light will cancel itself out.

To solve this problem, and get an efficient frequency doubling, it is necessary to match the two phase velocities v_p^P and v_p, a procedure known as phase-matching. Since the phase velocity is given by $v_p = c/n$, it is possible to express the phase matching condition as:

$$n_\omega = n_{2\omega} \tag{2.24}$$

This condition is not possible to obtain in a material with normal dispersion where the refractive index increases with frequency. However, by using an anisotropic medium, the frequency doubled light may be generated with a different polarization, experiencing a different refractive index than the incident light, making it possible to match them. In an anisotropic crystal the relation (2.20) is replaced by a tensor relation, since in this case an electric field may introduce a polarization that is not parallel to the direction of the electric field.

Short pulse phase-matching

When it comes to frequency doubling of very short pulses, implying a broad range of frequencies around the central frequency ω, it is hard to achieve phase-matching over the whole spectrum of the pulse, because of the dependence of the refractive index on frequency. This will give a larger doubling-efficiency for the phase-matched central frequencies than for the wings of the spectrum which will be suppressed. This causes a narrowing of the frequency-doubled spectrum giving a longer pulse (see section 2.1.4). Another equivalent way to describe this is that the group velocity of the frequency-doubled pulse differs from that of the fundamental pulse and thus the position where a frequency-doubled pulse is created will not overlap with the position of previously generated second harmonic pulses.

2.2.2 SHG-efficiency

Under the assumption of scalar fields in one dimension, the electromagnetic wave-equation may be solved for the fundamental and second harmonic fields (see Appendix A). This results in two coupled differential equations for the field variables A'_ω and $A'_{2\omega}$, which are defined in such a way that $|\tilde{A}'_\omega|^2$ and $|\tilde{A}'_{2\omega}|^2$ are directly proportional to the intensity of the fundamental and second harmonic

fields respectively.

$$\frac{\partial \tilde{A}'_\omega}{\partial z} = -\kappa \tilde{A}'^*_\omega \tilde{A}'_{2\omega} e^{-i\Delta k z} \tag{2.25}$$

$$\frac{\partial \tilde{A}'_{2\omega}}{\partial z} = \kappa \tilde{A}'^2_\omega e^{i\Delta k z} \tag{2.26}$$

In these equations κ is the coupling constant and $\Delta k = k_{2\omega} - 2k_\omega$ describes the phase mismatch between the modes. Here the solutions to these equations are presented for some particular cases in order to get a measure of the efficiency of the second harmonic generation.

Non-depleted input beam

In the case when $\Delta k z \gg 1$, the situation is far from phase matched. In this case the conversion efficiency is so low that there is negligible depletion of the input beam at frequency ω. Then $\tilde{A}'_\omega(z) = \tilde{A}'_\omega(0)$ is constant throughout the whole crystal, making it possible to integrate (2.26) and obtain an expression for $\tilde{A}'_{2\omega}(L)$ at the output of the crystal:

$$\begin{aligned} \tilde{A}'_{2\omega}(L) &= \int_0^L \kappa \tilde{A}'^2_\omega(0) e^{i\Delta k z} dz = \kappa \tilde{A}'^2_\omega(0) \int_0^L e^{i\Delta k z} dz \\ &= i\frac{\kappa \tilde{A}'^2_\omega(0)}{\Delta k}\left[1 - e^{i\Delta k L}\right] \end{aligned} \tag{2.27}$$

The efficiency for second harmonic generation η is defined as the ratio between the power of the generated second harmonic beam and the power of the incident fundamental beam.

$$\eta(L) = \left|\frac{\tilde{A}'_{2\omega}(L)}{\tilde{A}'_\omega(0)}\right|^2 = \kappa^2 \left|\tilde{A}'_\omega(0)\right|^2 L^2 \left(\frac{\sin(\Delta k L/2)}{\Delta k L/2}\right)^2 \tag{2.28}$$

When $\Delta k z \gg 1$, the efficiency vares periodically with crystal length, having its first maximum when $L = \pi/\Delta k$, a distance known as the coherence length. This periodic transfer of energy is due to the fact that the fundamental and the second harmonic wave move with different phase velocities, causing the frequency doubled light generated at different positions in the crystal to interfere destructively with itself (see section 2.2.1).

In the phase matched case (2.28) is only valid for small efficiencies, since the derivation of this expression assumes a non-depleted input beam.

Depleted input beam

In the case of perfect phase-matching $\Delta k = 0$, equations (2.25) and (2.26) become:

$$\frac{\partial \tilde{A}'_\omega}{\partial z} = -\kappa \tilde{A}'^*_\omega \tilde{A}'_{2\omega} \tag{2.29}$$

$$\frac{\partial \tilde{A}'_{2\omega}}{\partial z} = \kappa \tilde{A}'^2_\omega \tag{2.30}$$

From these expressions it is seen that if $\tilde{A}'_\omega(0)$ is taken to be a real quantity, both $\tilde{A}'_\omega(z)$ and $\tilde{A}'_{2\omega}(z)$ will be real quantities for all z. From the Manley-Rowe relation (A.15) one can then infer that

$$A'^2_\omega(z) + A'^2_{2\omega}(z) = A'^2_\omega(0) \tag{2.31}$$

and from (2.30) one obtains a differential equation

$$\frac{\partial A'_{2\omega}}{\partial z} = \kappa \left[A'^2_\omega(0) - A'^2_{2\omega}(z)\right] \tag{2.32}$$

with the solution

$$A'_{2\omega}(z) = A'_\omega(0) \tanh\left(A'_\omega(0)\kappa z\right) \tag{2.33}$$

From this expression, the second harmonic conversion efficiency η can be calculated for a crystal length L:

$$\eta(L) = \left|\frac{A'_{2\omega}(L)}{A'_\omega(0)}\right|^2 = \tanh^2\left(A'_\omega(0)\kappa L\right) \tag{2.34}$$

This shows that it is actually possible to get a full conversion of light from the fundamental frequency to the second harmonic. In most practical cases however, there are other processes involved making it difficult to achieve.

2.2.3 Total conversion efficiency and pulse length

The dependence of second harmonic generation efficiency on input beam intensity, as discussed in the previous section, suggests that the amount of frequency doubled light obtained for a pulse of constant energy will be strongly dependent on the pulse length. Let us consider an incident Gaussian pulse of the form

$$I_{\omega,i}(t) = \frac{C_p}{\tau_p} e^{-4\ln 2\left(\frac{t}{\tau_p}\right)^2} \tag{2.35}$$

where τ_p is the pulse length and C_p is a constant proportional to the total energy of the pulse. The pulse energy W is obtained by integrating the intensity profile over the entire pulse[3], giving:

$$W_{\omega,i} = \int_{-\infty}^{+\infty} I_{\omega,i}(t)dt = \int_{-\infty}^{+\infty} \frac{C_p}{\tau_p} e^{-4\ln 2\left(\frac{t}{\tau_p}\right)^2} dt = C_p\sqrt{\frac{\pi}{4\ln 2}} \tag{2.36}$$

Assuming for simplicity a conversion coefficient $\eta(L)$ which is independent of the instantaneous frequency, the intensity of the frequency-doubled pulse at the output of the crystal is given by

$$I_{2\omega,o}(t) = \eta(L) I_{\omega,i}(t) \tag{2.37}$$

[3]Note that the pulse energies in this section is defined as pulse energy per unit area of the beam, i.e. with the unit [J/m^2]. Normally the intensity is also integrated over the spatial coordinates of the beam, but for this analysis this is not necessary.

Non-depleted approximation

Under the non-depleted approximation, holding for low conversion efficiencies, the conversion coefficient is given by (2.28) and is proportional to the input intensity. The pulse energy of the frequency-doubled pulse is then

$$\begin{aligned} W_{2\omega,o} &= \int_{-\infty}^{+\infty} \eta(L) I_{\omega,i}(t) dt = \frac{2\kappa^2 L^2}{\varepsilon_0 c} \int_{-\infty}^{+\infty} I_{\omega,i}^2(t) dt \\ &= \frac{2\kappa^2 L^2}{\varepsilon_0 c} \int_{-\infty}^{+\infty} \frac{C_p^2}{\tau_p^2} e^{-8\ln 2\left(\frac{t}{\tau_p}\right)^2} dt = \frac{2\kappa^2 L^2 C_p^2}{\tau_p \varepsilon_0 c} \sqrt{\frac{\pi}{8\ln 2}} \end{aligned} \tag{2.38}$$

where perfect phase-matching has been assumed ($\Delta k = 0$). Substituting C_p from equation (2.36) into equation (2.38) and dividing by the pulse energy of the input pulse gives a measure of the total conversion efficiency for the pulse:

$$\eta_{tot} = \frac{W_{2\omega,o}}{W_{\omega,i}} = \frac{2\kappa^2 L^2 W_{\omega,i}}{\tau_p \varepsilon_0 c} \sqrt{\frac{2\ln 2}{\pi}} \propto \frac{W_{\omega,i}}{\tau_p} \tag{2.39}$$

To conclude the results of this simplified calculation, for low conversion efficiencies the total pulse conversion efficiency is, not only proportional to the energy of the incident pulse, but also inversely proportional to the pulse length.

Depleted input pulse

In the depleted case, the conversion efficiency $\eta(L)$ is given by expression (2.34) and the pulse energy of the frequency doubled pulse is then

$$\begin{aligned} W_{2\omega,o} &= \int_{-\infty}^{+\infty} \eta(L) I_{\omega,i}(t) dt \\ &= \int_{-\infty}^{+\infty} I_{\omega,i}(t) \tanh^2\left(\kappa L \sqrt{\frac{2I_{\omega,i}(t)}{\varepsilon_0 c}}\right) dt \end{aligned} \tag{2.40}$$

From this expression, it is seen that it is not easy to find a good analytical expression for the total pulse conversion coefficient $\eta_{tot} = W_{2\omega,o}/W_{\omega,i}$ in the depleted case. This calls for a numerical approach and in figure 2.1 the total conversion efficiency is plotted for both the non-depleted and the depleted case, as a function of the fundamental pulse energy for a constant pulse length. As seen, the non-depleted approximation will only be valid for sufficiently low conversion efficiencies.

Under the assumption of a constant pulse energy, figure 2.2 shows the total conversion efficiency as a function of pulse length. From this figure it is seen that the total conversion efficiency for a pulse is strongly dependent on the pulse length. Hence, to maximize the total SHG-efficiency is equivalent to minimizing the pulse length. This will be used in the present study. In agreement with the previous plot, figure 2.2 clearly shows that the non-depleted approximation is in good agreement with the true result only for low conversion efficiencies.

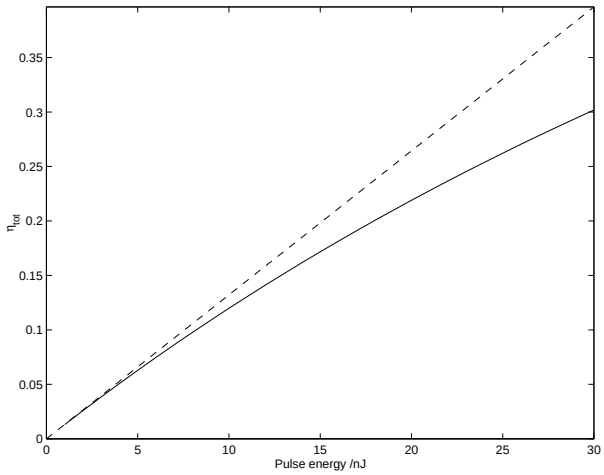

Figure 2.1: Variation of the total conversion efficiency with the input pulse energy. The dashed line shows the non-depletion approximation and the solid line is drawn for a depleted input pulse. The plot was made for a pulse length of 40 fs.

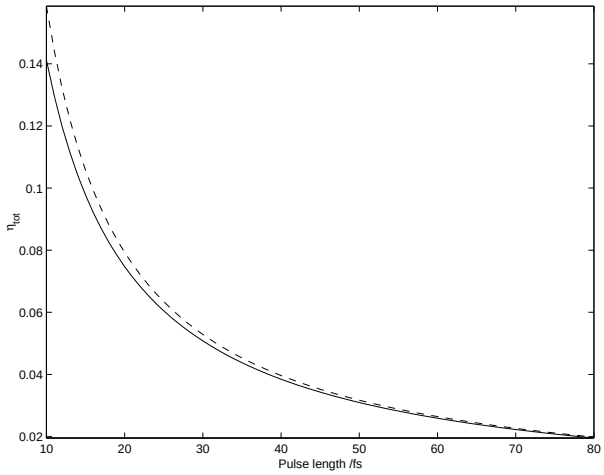

Figure 2.2: Variation of the total conversion efficiency with the input pulse length. The dashed line shows the non-depletion approximation and the solid line is drawn for a depleted input pulse. The plot was made for a pulse energy of 3 nJ.

The plots in figure 2.1 and figure 2.2 were made for a wavelength of 800 nm, using the estimated parameters for the experiments performed in this work. The crystal used was a 100 μm thick BaB_2O_4 (BBO), with an effective nonlinear coefficient $d_{eff} = 1.6 \cdot 10^{-12}$ m/V [11]. The refractive indices needed to calculate the coupling constant were calculated using the Sellmeier formulas for BBO [12], and the calculated value was $\kappa = 4.9 \cdot 10^{-6}$ V^{-1}. The spot size of the laser beam in the crystal was estimated to be 230 μm^2. The plot in figure 2.1 was made for a pulse length of 40 fs. Figure 2.2 was plotted for a pulse energy of 3 nJ.

2.3 Characterization of short pulses

In section 2.1 a mathematical description of short laser pulses was made. Concepts like amplitude and phase was introduced and discussed both in the time-domain and in the frequency-domain. An important issue that is involved in almost all practical situations, is how to measure these quantities. When dealing with ultrashort laser pulses with a time duration in the order of femtoseconds, it is obvious that traditional electronic equipment will not be adequate due to the long response time of electronics [4]. Measurements of the spectral amplitude of short pulses pose no additional problems as compared to continuous light, and an ordinary spectrometer may be used.

Most methods used today for time characterization of ultrashort laser pulses, uses a reference pulse as a probe for doing the measurement. However, to correctly resolve the time behaviour of a laser pulse, the reference pulse needs to be shorter than the measured pulse, which is a problem since the measured pulse is often the shortest pulse available. If however a shorter reference pulse is available, in turn its time behaviour needs to be well known, and we are back where we started. The solution to this is to use the pulse itself as the reference. In this way at least some properties of the time profile may be extracted. If such a measurement is combined with for example spectral data, the pulse shape can be fully characterized.

2.3.1 Autocorrelators

As the name implies, an intensity autocorrelator is used to record the intensity autocorrelation signal defined as

$$A(\tau) = \int_{-\infty}^{+\infty} I(t)I(t-\tau)dt \tag{2.41}$$

where $I(t)$ is the temporal intensity of the pulse. This signal, which is always symmetric and thus contains very little information of the pulse shape, can be used to measure the duration of an ultrashort laser pulse, under the assumption

[4] The fastest photodiodes today have a response time 100 ps, while for example a streak camera might be able to do measurements with a time resolution of a few hundred femtoseconds.

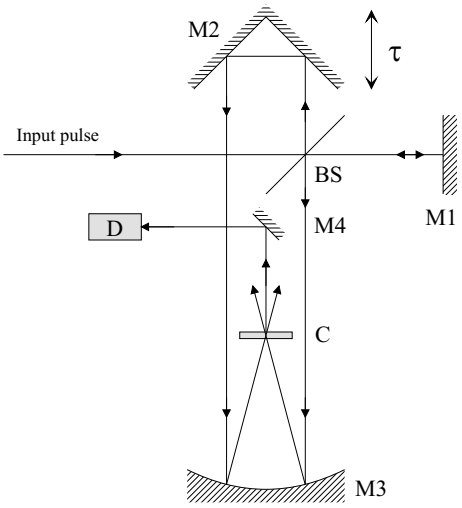

Figure 2.3: Multi-shot autocorrelator. The input pulse is split at the beamsplitter BS and one part is reflected at mirror M1 while the other part is reflected at the roof-top mirror M2, which is placed on a translation stage. Both parts are then focused by the mirror M3 and crossed in the crystal C. The autocorrelation signal is recorded by the photomultiplicator D.

of a known pulse shape. More precisely, the FWHM of the autocorrelation signal τ_{ac} and the pulse length τ_p are related by

$$\tau_{ac} = a\tau_p \tag{2.42}$$

where the constant a depends on the pulse shape. For a Gaussian pulse the constant is $a = 1.414$ while for a hyperpolic secant pulse shape, the constant is $a = 1.543$ [1]. Using the relation 2.42 and assuming a specific pulse shape, the pulse length of an ultrashort laser pulse can easily be retrieved from an autocorrelation trace. The pulse length is an important parameter and, when working with an experiment, often the first thing that needs to be measured. Furthermore, the autocorrelation is a rather fast and comparingly simple measurement to carry out, and due to this it is widely used in practice.

In figure 2.3 a possible design of a multi-shot autorrelator is shown. The incoming laser pulse is split using a beamsplitter, and the two replicas of the pulse takes different paths before they are both focused and crossed in a crystal. Due to the focusing, the intensity in the crystal is high, allowing for sum-frequency generation[5] to occur between the two pulses. Due to the phase-matching condition for the crossed beams, the sum-frequency signal is generated at the bisector of the two beams, and may then be detected by a photomultiplier. In this way the value of the overlap integral (2.41) may be measured and by translating the roof-top mirror M2, the time delay τ between the pulses may be scanned while the autocorrelation signal is recorded.

[5]Sum-frequency generation is caused by the same physical process as second harmonic generation, as described in section 2.2.1, with the difference that the second order term in (2.20) is now the product between the amplitudes of two different pulses.

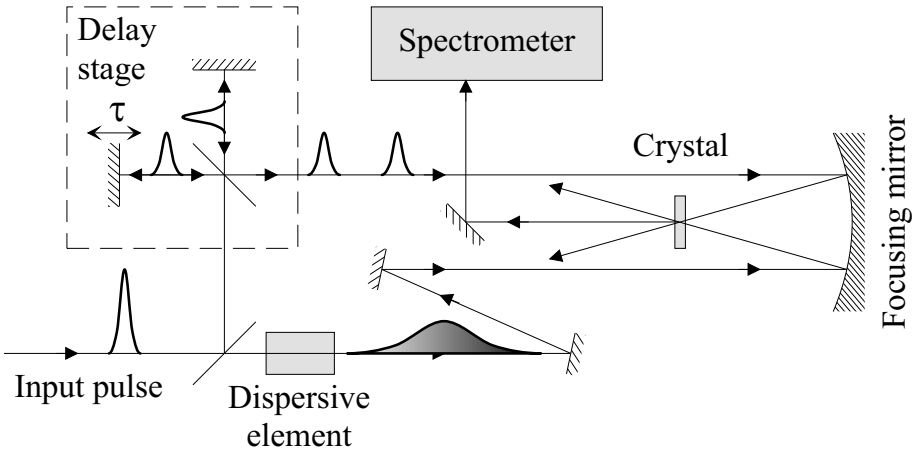

Figure 2.4: SPIDER-apparatus. The two time-delayed replicas are overlapped with a chirped third replica in the crystal. The spectrum of the sum-frequency pulses along the bisector between the crossed beams is detected using a spectrometer.

The intensity autocorrelator used in this work has a slightly more sophisticated design, where a rotating parallel mirror assembly is used to vary the time delay τ.

2.3.2 SPIDER

SPIDER stands for Spectral-Phase Interferometry for Direct Electric-field Reconstruction and is a technique based on frequency-domain interferometry with spectral shearing as described in Appendix B. The time-delayed replicas of the pulse are easily obtained in the same way as for the autocorrelator above (see figure 2.3), while the spectral shear between the two pulses corresponds to a linear temporal phase modulation, which is not obtainable by any electronic modulator today. To solve this, the non-linear process of sum-frequency generation may be used. By focusing the two time-delayed replicas into a crystal, together with a chirped third replica of the pulse, the sum-frequency generation will occur with different parts of the chirped pulse for the two time-delayed pulses. In this way, the nonlinear signal will contain both replicas, upshifted with different frequencies, and thus spectrally sheared relative to one another.

In figure 2.4 the SPIDER-apparatus used in this work is depicted, and in figure 2.5 the overlap between the two time-delayed pulses and the chirped pulse in the crystal is shown. Note that since the chirped pulse is considerably longer than the time-delayed pulses, it can be assumed that the frequency shift for each pulse is the same for all frequencies. After recording the SPIDER-signal, the time-delay can be measured by recording the spectrum of the direct-doubled time-delayed pulses and measuring the fringe distance (see Appendix B). The spectral shear Ω may finally be determined by blocking one of the replicas at a

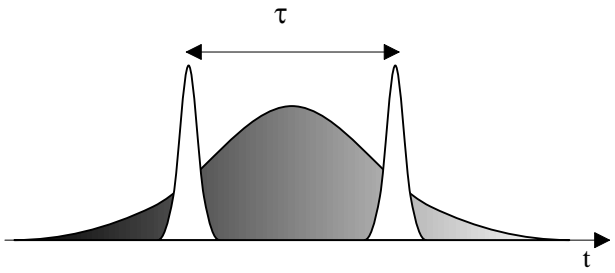

Figure 2.5: Overlap between the two time-delayed replicas and the chirped pulse in the crystal.

time, measuring the displacement of the spectrum.

When the spectral phase has been measured and calculated using the FTSI-method described in Appendix B, it is possible to reconstruct the temporal shape of the pulse through an inverse Fourier transform. To do this it is necessary to know the spectral intensity of the input pulse, readily measured by a spectrometer.

2.4 Pulse shaping

This section describes different techniques used to change the shape of laser pulses, starting with pure compression in the time-domain, continuing with some common pulse shaping techniques and finishing off with the description of the Acousto-Optic Programmable Dispersive Filter (AOPDF) that was used in this work.

2.4.1 Pulse compression techniques

As seen above in section 2.1.4, the shortest pulse for a given spectral bandwidth, the transform-limited one, is obtained when all higher-order phase terms are zero. If one wants to compress a pulse to a pulse length shorter then the transform limit, it is therefore necessary to increase the spectral bandwidth. This can be done, for example, by self-phase modulation in a hollow fiber [9]. The increase in bandwidth is however always accompanied by a chirp which needs to be compensated.

More generally, since most materials exhibit positive dispersion (see section 2.1.3), a laser pulse will always accumulate a second order phase term upon propagation. Even in air a laser pulse broadens as it propagates, gaining an increasing linear chirp. This calls for optical elements exhibiting a negative GDD, i.e. a group delay that decreases with increasing frequency.

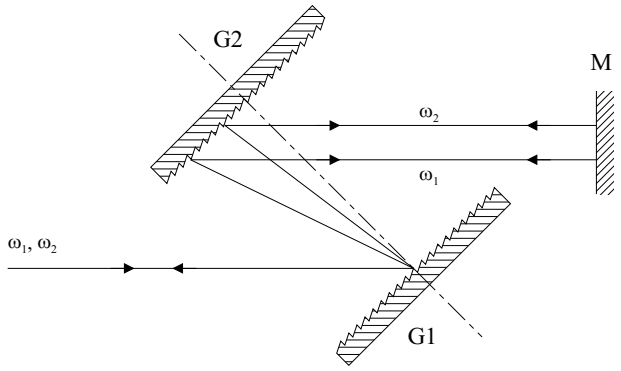

Figure 2.6: Parallel grating compressor. The incident beam at A is dispersed by the grating $G1$. After reflection at grating $G2$, the different frequency components are parallel but spatially spread and the component of frequency ω_2 has travelled a shorter distance than the component with frequency ω_1 when reflected at the mirror M.

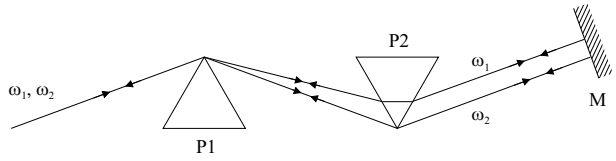

Figure 2.7: Prism compressor. The angular dispersion of prism $P1$ causes the different frequency components of the beam to pass through different amount of material in prism $P2$.

The grating compressor

A common approach is to make use of the angular dispersion of gratings, arranging them in such a way that low frequency components of a pulse will be delayed more than high frequency ones. In figure 2.6 an arrangement of two parallel gratings is depicted.

Let us consider the two frequency components ω_1 and ω_2 of the incident beam, with $\omega_1 < \omega_2$. Due to the angular dispersion of the first grating, the high frequency component ω_2 travels a shorter optical path $AB'C'B'A$ than the low frequency component ω_1 ($ABCBA$). In this way the grating pair introduces a negative GDD, compensating for a positive second order phase.

The prism compressor

By using a pair of prisms, arranged as in figure 2.7, a negative GDD may be achieved. In this case it is the angular dispersion of the first prism ($P1$) that is used to separate the frequency components of the beam. Depending on

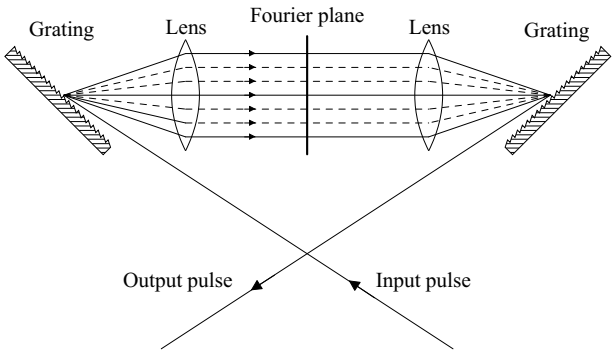

Figure 2.8: Zero dispersion-line used for pulse shaping in the Fourier plane.

frequency, the components will then pass through different amounts of material in the second prism ($P2$). More precisely, the low frequency components will experience less deviation by the first prism, thus passing through a larger amount of material at the second prism, being delayed relative to the high frequency components. By adjusting the distance between the prisms, it is possible to control the amount of negative GDD introduced by the compressor and by translating the first prism along the normal to its base plane, it is possible to control the amount of material that all frequency components go through, making it easy to do fine adjustments to the applied compression.

2.4.2 Common pulse shaping techniques

As seen from section 2.1, a laser pulse can be fully characterized by its complex spectral amplitude. A pulse can then be "shaped" by manipulating both its spectral amplitude and the phase. A common and straightforward way to achieve this, is to spatially spread the frequency components of the laser pulse, and then use filters to control the amplitude and phase for different spectral regions.

The basic apparatus used for this type of pulse shaping is referred to as a "zero dispersion pulse compressor" or a "zero dispersion-line" and is depicted in figure 2.8. In a zero dispersion-line the frequency components in the input beam are angularly dispersed by the first grating, and then collimated by the first lens. Components with different frequencies are in this way focused at different points in the back focal-plane of the lens. The effect of the lens is to transform the angular dispersion of frequencies into spatial separation at the focal plane, thus performing a Fourier transform (the focal plane is also called the Fourier plane). The second lens then performs the inverse transform back to angularly dipersed frequencies and then the last grating collimates the beam again.

By inserting a filter in the Fourier plane, the different frequency components may be arbitrarily modulated in amplitude or phase. If no filter is used, a well-aligned zero dispersion-line should have no effect on the pulse at all.

Using this type of setup, a variety of filters might be used, both fixed and programmable. In the latter case, the pulse shaping is said to be adaptive. The filters might be phase-only, affecting only the phase of the spectral components, or amplitude-only, changing the shape of the spectral amplitude. It is also possible to use filters which modify both the amplitude and the phase of the spectral components.

When adaptive pulse shaping is required, programmable masks have to be used. One common type of programmable mask is the Liquid Crystal Spatial Light Modulator (LC SLM), using electronically addressed liquid crystal displays. Modulators of this kind have the disadvantage of a finite resolution depending on the size of the smallest programmable point of the mask. Another approach is to use an Acousto-Optic Modulator (AOM), operated in the Bragg-regime [13], giving a lower transmission due to the acoustic diffraction efficiency, but offering today a factor of 2 better resolution than the LC SLM [14].

2.4.3 The Dazzler

The Dazzler[6] is a pulse shaping device proposed and designed by P. Tournois [5]. It is an Acousto-Optic Progammable Dispersive Filter (AOPDF) based on collinear[7] acousto-optic interaction in a crystal of Paratellurite (TeO_2). Paratellurite is a positive uniaxial crystal, and as such, it can be described by two indices of refraction, an ordinary n_o and an extraordinary n_e, corresponding to two different states of polarization. The dielectric tensor for a uniaxial crystal may be written as:

$$\varepsilon = \varepsilon_0 \begin{pmatrix} n_o^2 & 0 & 0 \\ 0 & n_o^2 & 0 \\ 0 & 0 & n_e^2 \end{pmatrix} \tag{2.43}$$

For a positive crystal the extraordinary index is larger than the ordinary one while for a negative crystal $n_e < n_o$.

Interaction between acoustic and optical waves

Acoustic waves are pressure waves, propagating through a medium at the speed of sound, generally much slower then the speed of light. This makes it possible to study the interaction of light with acoustic waves under the assumption that the acoustic wave is stationary or "frozen" during the short interaction time. As a pressure wave, an acoustic wave introduces strain into a medium, and as a result of the photoelastic effect [10] the refractive index will change with the amount of strain introduced. A common approach is to regard this change in refractive index as a perturbation to the dielectric tensor (2.43). In a medium where an acoustic wave is propagating in the z-direction the dielectric tensor

[6]The Dazzler is manufactured and sold by Fastlite, Ecole Polytechnique, XTEC Bâtiment 404, 911 28 Palaiseau, France.

[7]Collinear interaction means that the acoustic and optical waves travels in parallel through the material, thus maximizing the interaction length.

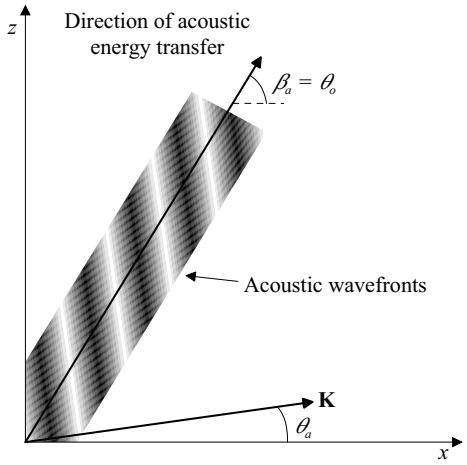

Figure 2.9: An acoustic wave launched in a Paratellurite crystal with its wavevector $\mathbf{K}$ directed at an angle θ_a from the principal x-axis will have its direction of energy transfer in an angle β_a from the x-axis.

may be written as

$$\varepsilon(z) = \varepsilon + \Delta\varepsilon(z) \tag{2.44}$$

where ε is the unperturbed dielectric tensor in 2.43 and $\Delta\varepsilon(z)$ is the perturbation caused by the acoustic wave, having off-diagonal elements.

As mentioned above, there are two possible modes of propagation in a unixial crystal, and light propagating as an ordinary or extraordinary wave will stay in its state during propagation. The effect of the pertubation $\Delta\varepsilon(z)$ is that it couples these modes of propagation, so that energy can be transferred between the two modes by interaction with the acoustic wave [10]. This transfer will only occur in the case of phase matching between the optical modes and the acoustic wave, a condition that may be written as

$$\mathbf{k_e} = \mathbf{k_o} \pm \mathbf{K} \tag{2.45}$$

where $\mathbf{k_o}$ and $\mathbf{k_e}$ are the wavevectors of the ordinary and extraordinary modes, and $\mathbf{K}$ is the acoustic wavevector with magnitude $K = 2\pi f/V$ where f is the acoustic frequency and V is the speed of sound in the material.

If an acoustic wave is launched in a paratellurite crystal, with its wavevector $\mathbf{K}$ directed in an angle θ_a from the principal x-axis, the actual direction of energy transfer is different from the direction of the $\mathbf{K}$-vector and at an angle β_a from the principal x-axis (see figure 2.9). This is due to the strong anisotropy of the crystal, causing the acoustic phase velocities V_x and V_z along the x and z-axis to differ.

The incident ordinary ray is directed at an angle $\theta_o = \beta_a$ from the principal x-axis. This choice is made in order to maximize the interaction length with the acoustic wave and thus minimize the acoustic power needed. Coupling between

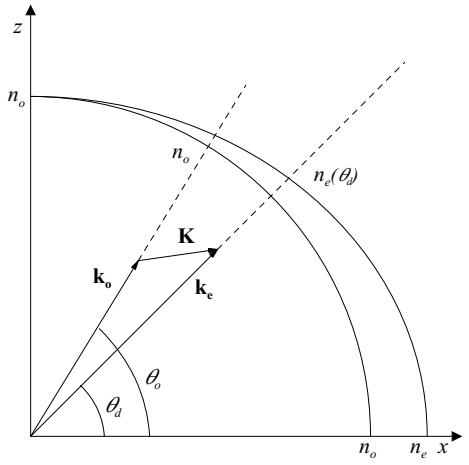

Figure 2.10: Phase-matching in a Paratellurite crystal between an ordinary wave $\mathbf{k}_o$ and an extraordinary wave $\mathbf{k}_e$ with an acoustic wave $\mathbf{K}$.

the ordinary and extraordinary waves and the acoustic wave is achieved when the phase-matching condition (2.45) is fulfilled and in figure 2.10 this situation is depicted.

From (2.45), using the fact that the optical anisotropy $\Delta n = n_e - n_o$ is generally small compared to n_0, an expression for the ratio α between the copuled acoustic and optical frequencies can be obtained (see [15]).

$$\alpha = \frac{f}{\nu} = \Delta n \frac{V(\theta_a)}{c} \frac{\cos^2 \theta_0}{\cos(\theta_0 - \theta_a)} \tag{2.46}$$

where $V(\theta_a)$ is the acoustic phase velocity along the direction θ_a.

For the device used in this thesis, the crystal is cut so that the angle of the acoustic wavevector is $\theta_a = 8°$. For maximized interaction length, the incidence angle of the ordinary wave is $\theta_o = 58.5°$ and for a wavelength of 800 nm $\alpha = 1.42 \cdot 10^{-7}$. In order to transfer light at the wavelength 800 nm, corresponding to a frequency of 375 THz, the acoustic wave must have a frequency of 53.2 MHz, easily obtainable with todays RF generators. Furthermore, due to the dispersion of the Paratellurite crystal, the optical anisotropy Δn and hence the ratio α is wavelength-dependent.

Amplitude and phase control of short pulses

This section will describe in a phenomenological way how an acoustic wave can be used to introduce a linear chirp to a laser pulse. For each frequency component in an incident ordinary laser pulse, it is possible to find an acoustic frequency that fulfills the phase-matching criterion, transfering part of the frequency component into an extraordinary mode of propagation. In the Dazzler, both the acoustic wave and the laser pulse propagate together through the length

of the crystal L. The acoustic wave is produced using a transducer excited by an RF generator. Since the acoustic wave propagates through the crystal with the velocity $V(\theta_a)$, its temporal profile is reproduced spatially in the crystal for a given moment of time. Due to the short duration of the laser pulse, and to the fact that the speed of light far exceeds the speed of the acoustic wave, the laser pulse can be assumed to encounter a fixed dielectric grating upon entering the crystal. Locally this dielectric grating has only one spatial frequency, related to the corresponding temporal frequency of the acoustic wave, and hence only one frequency component of the incident laser pulse will be phase-matched at a certain position in the crystal and coupled to the extraordinary mode. Depending on the time behaviour of the acoustic signal, it is then possible to transfer different frequency components of the incident pulse into the extraordinary mode at different positions in the crystal.

After the crystal, the extraordinary mode contains all the frequencies that have been phase-matched at different positions in the crystal. Depending on the amplitude of the acoustic wave for different frequencies, the coupling efficiency may be controlled individually for each frequency component, making it possible to directly modify the spectral amplitude of the laser pulse.

Since paratellurite is an anisotropic crystal, a frequency component will experience different indices of refraction when travelling in the different modes. Hence if the coupling takes place at a coordinate x along the crystal (x being zero at the front face), the phase shift of the component due to propagation through the whole length L of the crystal is

$$\Delta\varphi(x) = -k_o x - k_e(L - x) = -k_e L + \frac{\delta n}{c}\omega x \tag{2.47}$$

where k_o and k_e are the propagation constants for the ordinary and extraordinary waves and δn is the difference in refractive index between the extraordinary and the ordinary wave in the direction of propagation. The first term represents the normal effect on the phase and does not depend on where in the crystal the coupling takes place. The second term describes the phase that is controlled by the acoustic wave. Consider for example an acoustic signal with a negative linear chirp, that is with a linearly decreasing instantaneous frequency. Inside the crystal, this signal reproduces itself spatially with a positive linear chirp (as seen when travelling in the direction of the laser pulse). More specifically the low frequency components are at the beginning of the crystal and the high frequency components at the end. Lower optical frequency components will thus be coupled earlier in the crystal than higher frequency ones. Neglecting dispersion, assuming that α and δn are frequency-independent, x in equation (2.47) becomes proportional to ω. Thus the second term in the introduced phase shift is approximately quadratic in frequency, meaning that a negative linear chirp is introduced to the pulse (see section 2.1.4).

An illustration to this situation is depicted in figure 2.11 where an incident pulse with an initial positive chirp is compressed through the application of an acoustic wave with a negative linear chirp.

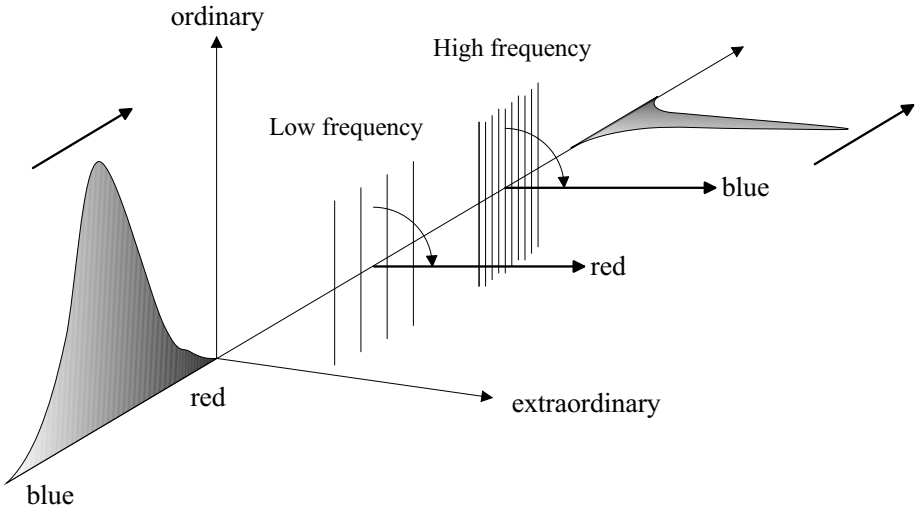

Figure 2.11: Compression of a pulse with positive chirp using an acoustic wave with negative chirp. The red components of the pulse gets coupled earlier in the crystal than the blue ones, thus travelling a longer distance experiencing the higher extraordinary refractive index. In this way the low frequency components are delayed relative to the high frequency components, causing a compression of the pulse.

Theoretical result for the Dazzler

A more thourough investigation shows that [5] the output optical pulse is proportional to the convolution between the input pulse and the scaled acoustic signal

$$E_{out}(t) \propto E_{in}(t) \otimes S(t/\alpha) \tag{2.48}$$

where S is proportional to the electric signal driving the piezoelectric transducer used to excite the acoustic wave in the crystal and α is the ratio between the acoustic and optical frequency as defined in (2.46). In the frequency-domain (2.48) becomes

$$E_{out}(\omega) \propto E_{in}(\omega) S(\alpha\omega) \tag{2.49}$$

Group delay compensation

As seen above the Dazzler can apply a negative GDD on a laser pulse. The maximum applicable GDD is dependent on the crystal length and the actual bandwidth of the pulse. The larger the bandwidth of the pulse, the longer crystal is needed to introduce a fixed amount of GDD. An important issue is that the crystal itself introduces a positive GDD of around 12500 fs^2. If this is to be compensated by the Dazzler itself, there is an upper bandwidth limit for which this is possible. At 800 nm, with a crystal length of $L = 25$ mm, the maximum

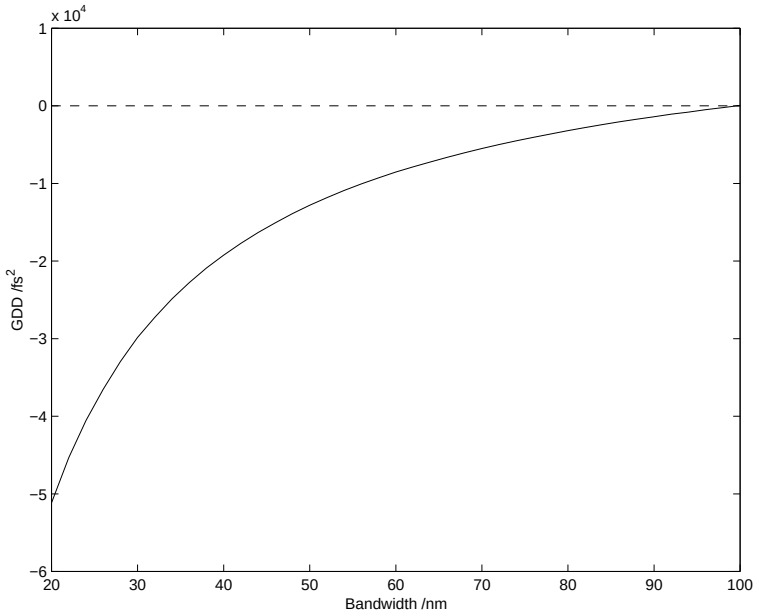

Figure 2.12: Limits for the amount of introduced negative GDD with the crystal dispersion included. Note that the limit for the introduced GDD is zero at the maximum self-compensation bandwidth of 100 nm.

total bandwidth of self-compensation called the cut-off bandwidth[8], is 100 nm. In order to preserve the wings of the spectrum it is therefore in practice not possible to have a bandwidth larger than about half of this maximum bandwidth. In figure 2.12 the maximum negative GDD that can be introduced by the Dazzler is shown as a function of the pulse total bandwidth. In this plot the crystal dispersion has also been included.

Another important consideration related to the introduced group delay is the interaction efficiency of the acoustic wave with the laser pulses. A single frequency treatment with coupled mode analysis [10] allows to relate the diffracted light intensity to the incident light intensity and to the acoustic power density in the interaction area. Due to material properties there is however a limit for the maximum allowed acoustic power density in the crystal. Because of this it is preferrable to spread out the local frequencies of the acoustic wave so that the whole crystal is used for diffraction, allowing for a larger total interaction efficiency. As discussed above, this spreading is controlled by the chirp of the acoustic wave, making it important to at least introduce some GDD in order to fill the crystal sufficiently with the acoustic wave.

To allow the Dazzler to operate within a dynamic range where none of the above limits are violated, it is often necessary to introduce a fixed GDD by external means. In this work a prism compressor was used before the Dazzler to achieve this.

[8]Note that this is truly the total bandwidth, as opposed to the usual FWHM-bandwidth.

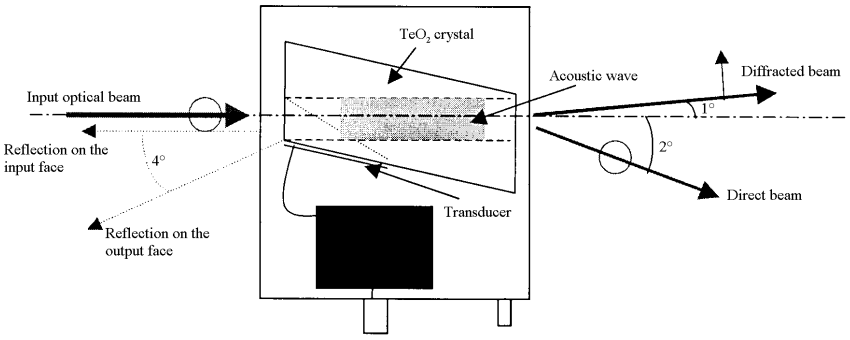

Figure 2.13: Topview of crystal unit. The crystal is cut so that both direct and diffracted beams deviate a little from the input beam. Picture from [16]

Wavelength resolution and calibration

The Dazzler used in this work is constructed to work at a central wavelength of 800 nm, allowing for a spectral bandwidth of 200 nm limited by the bandwidth of the piezoelectric transducer. For a collimated input beam, the spectral resolution is 0.6 nm, which means 340 individually programmable points within a bandwidth of 200 nm. A divergent incident beam degrades this resolution and the maximum allowed divergence is $\delta\theta_{1/2} = 0.04°$. The angle of incidence of the input beam is more tolerant to misalignment, causing only a degradation of the spectral calibration for small misalignments. This is due to the fact that the frequency ratio α in (2.46) is calculated for a specific angle of incidence and thus has to be changed in the software if another angle of incidence is required. A large misalignment of the input beam will however lead to poor interaction with the acoustic wave. The normal procedure for fine tuning of the incidence angle is to match the wavelengths set in the software with those of a spectrometer by tilting the crystal.

Dazzler setup and control software

The optical setup of the Dazzler crystal is shown in figure 2.13. The polarization of the input beam should be vertical, leading to a horizontal polarization of the diffracted beam, which has a deviation of 1° from the direction of the input beam. The acoustic wave is launched in the crystal by the transducer, recieving its signal from the RF generator via a coaxial cable.

The RF generator has a waveform memory, making it possible to operate the generator in a stand-alone mode, as soon as a wave form has been sent to it from the control computer. The launch of the waveforms may be triggered either internally or by an external signal. In internal trigger mode, it is possible to set the delay between subsequent launched acoustic signals. When externally triggered, it is possible to set the internal delay between the trigger signal and

the acoustic signal. In usual operation, working with a pulsed laser system, the RF generator is triggered externally by a signal from the laser system[9], to assure that the acoustic signal is fully present in the crystal when the laser pulses arrive. In order to get a higher throughput through the crystal, for example when aligning the setup, it is also possible to let the RF generator work in a continuous mode, where the next acoustic signal is sent directly following the previous one, allowing for an acoustic signal repetition rate of approximately 34 kHz[10].

The control software delivered with the Dazzler, allows the user to arbitrarily apply a filter to the laser pulse. For the amplitude, a supergaussian profile with a hole can be applied. The position and width of the profile and the position, width and depth of the hole can be adjusted by panel controls. For the phase, a polynomial with variable coefficients up to the fourth order can be applied. It is also possible to feed the program with two files, one for the amplitude and one for the phase, containing values of the amplitude and the phase for different wavelengths, written in a standard column format. When reading such files, the program makes a linear interpolation between the specified points.

When a filter function has been chosen, the software computes the corresponding RF signal in the frequency domain, using the scaling factor α as described above. The signal is then transferred to the time domain via an inverse Fourier transform, and sent to the RF generator.

In many applications, there is a need to control the applied pulse shape from an external program. For these cases the Dazzler software supports a loop-mode, in which the software is idle until a request is made through a request file. The request file is a text file containing the file path to a parameter file with the front panel settings described above. Upon finding a request file, the software reads the desired parameter file, calculates the corresponding RF waveform and sends it to the RF generator. When the waveform has been sent, the request file is deleted indicating that the requested pulse shape is now active.

Due to the maximum allowed acoustic power density in the crystal, caution has to be taken not to use a to high trigger rate. The maximum allowed trigger rate is dependent on the average power of the acoustic waveform. In normal operation, the control software calculates the maximum trigger frequency and sends it together with the waveform to the RF generator. If the trigger frequency at any moment exceeds this maximum value, the RF generation is stopped and the user is alerted. When operated in loop-mode, this safety feature is however not active since it would considerably slow down the switching between waveforms.

[9] In a laser system using a pulse picker to select the pulses, this signal is most conveniently used to trigger the RF generator, and the internal delay may then be adjusted to synchronize the acoustic signal with the arrival of the pulse.

[10] This repetition rate is due to the sampling frequency of the RF generator, $f_s = 70$ MHz, and the number of sampling points, $N = 2048$, causing the signal to repeat itself after $N/f_s = 29.3$ μs.

Chapter 3

Genetic algorithms

3.1 Theory

Genetic Algorithms are a class of search algorithms, belonging to the field of Evolutionary Programming, using natural evolution as the model for a specific problem[1]. The concept of genetic algorithms was developed by Holland *et al.* at the University of Michigan in 1975, in hope to transfer the important natural mechanisms of adaptation and exploration into artificial systems [17].

Genetic algorithms are parallel in the sense that the search is not done from a single point in the search space, but rather from a set of points called a population. Each search point in the population is called an individual. Each individual contains encoded information about the parameters representing the search point. In this way the genetic algorithm does not work directly with the parameters themselves. Each population, or each set of individuals, is referred to as a generation, and the strategy of the search procedure is to evolve the current population creating a new generation with better suited individuals.

An important concept of genetic algorithms is the fitness or objective function. The fitness function is needed to direct the search, describing what a good individual should look like. In fact, the fitness function may be regarded as a function mapping each possible individual to a fitness value and the higher the fitness, the better the individual. The process of calculating the fitness[2] for each individual in a population is an important step in a genetic algorithm and is often referred to as the evaluation of the population.

In figure 3.1 a flow-chart of genetic algorithm operation is shown. Starting from

[1] In many cases the term "genetic algorithm" is reserved for implementations using a bit encoding of the search parameters while implementations with a more complex encoding are called Evolutionary Algorithms. In the present work, a bit encoding was used and hence the term "genetic algorithm" will be used.

[2] In this case, "calculating" does not necessarily refer to a raw calculation, since the fitness of an individual might sometimes be determined by for example the result of an experiment or the performance of the individual in solving a specific problem.

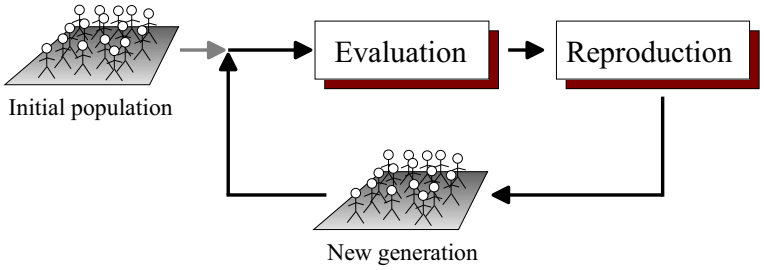

Figure 3.1: Flow-chart describing genetic algorithm operation.

an initial random population, the population is evaluated and reproduced to obtain a new generation. The new generation is then evaluated and the process goes on until some pre-defined stop-criteria is met. The completion of the entire process of optimization, starting from a random population and stopping when the criteria is met, is often referred to as a "genetic algorithm run" or just a "run". The reproduction step creates the new generation based on the fitness of the different individuals. Reproduction is often divided into three steps, or operators, called selection, crossover and mutation which are discussed in section 3.1.3.

3.1.1 Advantages over normal search procedures

Traditional calculus-based search methods may be divided into two main classes: indirect and direct. Indirect methods try to localize a maxima by solving the non-linear set of equations describing the problem, setting the gradient of the function to zero[3]. Direct methods on the other hand, search from a point in the search space, moving in the direction of the local gradient. Calculus-based methods are thus depending on the existence of local derivatives over the whole search-space, a condition that is hard to fulfill for search spaces with discontinuities and noise. Another problem with these methods is that they only search for local maxima, with a risk of missing the global maximum far away from the current point in the search space.

Another type of search method is the enumerative approach. As the name implies, these simple schemes uses brute force to calculate the function value at every possible point in the search space, keeping the best-so-far as they go on. The problem with such an approach is obvious. In a search space with a large number of dimensions, the number of possible points gets enormous, making it impossible to examine all of them in a reasonable time.

During the evaluation step, a genetic algorithm only needs to calculate the value of the fitness function for the individuals in the current population, making it independent of derivatives or other indirect information. Thus a genetic algorithm

[3]Here the goal of the search is described as finding a maximum value of a function. This is a true description in all searches if the function considered is the fitness function of the search space as described in section 3.1.

can be expected to perform well even in a search space with discontinuities and noise. As described above, a genetic algorithm does not search from a single point, but rather from a whole set of points, contained in the current population. Because of this it is less likely to get stuck on a local maximum, since it never ceases to explore other parts of the search space as well.

3.1.2 Parameter encoding

In a genetic algorithm, the parameters describing a point in the search space are not handled directly, but are encoded as a bit-string or chromosome. The individual bits of the string are often referred to as genes. Consider a chromosome with 10 genes

$$(a_9, a_8, a_7, a_6, a_5, a_4, a_3, a_2, a_1, a_0) \tag{3.1}$$

where the genes a_i may take on the values 0 and 1. The most common decoding of such a string is to regard the chromosome as an positive integer A expressed in the base 2:

$$A = \sum_{i=0}^{9} a_i \cdot 2^i \tag{3.2}$$

Generally for a chromosome with l genes, this decoding means that the unsigned integer A may take on 2^l different values, ranging from 0 to $2^l - 1$.

Mapping of real parameters

In most cases, the search space can not be described by an unsigned integer, but rather by real numbers. This makes it necessary to discretize the parameters and map them to the integer value of the chromosome. Considering the real search parameter x, allowed to vary between the values x_{min} and x_{max}, being mapped to a chromosome with l genes, a straightforward mapping is

$$x = x_{min} + \frac{\sum_{i=0}^{l-1} a_i \cdot 2^i}{2^l - 1} (x_{max} - x_{min}) \tag{3.3}$$

which is the mapping that was used in this work. The precision of the mapping can be calculated as

$$\Delta x = \frac{x_{max} - x_{min}}{2^l - 1} \tag{3.4}$$

showing that when the limits x_{min} and x_{max} have been set, the number of genes l has to be chosen so that the required precision is obtained. A multiparameter coding is easily obtained by concatenation of several single-parameter codings. If for example a point in the search space can be described by the three real parameters x_a, x_b and x_c and these are to be coded with l_a, l_b and l_c genes each, the multiparameter chromosome will look like

$$(a_{l_a-1}, ..., a_0, b_{l_b-1}, ..., b_0, c_{l_c-1}, ..., c_0) \tag{3.5}$$

with the total length $l = l_a + l_b + l_c$ and where the genes a_i encode the parameter x_a and so on.

3.1.3 Reproduction operators in a conventional genetic algorithm

The heart of the genetic algorithm is the process of reproduction, the aim of which is to preserve high-scoring individuals as well as to create new individuals to explore other parts of the search space. The first thing that needs to be done is to decide which individuals in the current population should be given offspring populating the next generation. This is done by the selection operator, and the choices are based on the fitness values assigned to the individuals in the evaluation step. When the selection is made, the chosen individuals are combined (or mated) to produce offspring for the next generation. By this mating, or crossover, the promising characteristics of one individual may be combined with some other characteristics of another individual. The last step is the mutation step, in which the individuals may be randomly changed by some usually small probability. This step prevents that the algorithm will converge prematurely, forcing it to try random and unexplored points of the search space once in a while.

This section will describe the standard set of operators commonly used in a conventional genetic algorithm. The particular choice of reproduction operators together with the associated parameters discussed below is commonly referred to as the reproductive plan.

Selection

In a population containing N individuals the fitness of a single individual is denoted f_i where i ranges from 1 to N. The task of the selection operator is to select a total of N individuals out of this population. The selection is done with replacement so that the same individual may be selected several times.

The most common way of doing this is to decide that the probability of a certain individual being chosen should be proportional to its fitness. A selection mechanism doing this type of proportionate selection can be implemented by the steps shown in figure 3.2. This method is often referred to as Roulette-Wheel Selection [18], since it may be described by a roulette wheel with circumference f_{sum} where each individual is assigned a sector occupying a length f_i of the circumference.

Crossover

For each two individuals that have been selected by the selection operator, there is a chance, determined by the crossover probability, that they will be mated before they are copied to the next generation. Working with binary strings as chromosomes there are many ways in which this can be done. The most common scheme is the Single-Point Crossover. Single-Point Crossover randomly chooses a crossover-point between two genes in the chromosome. The chromosomes of both individuals are then cut at this position and the first

1. Calculate the sum of all the individual fitnesses in the population f_sum.

$$f_{sum} = \sum_{i=1}^{N} f_i$$

2. Pick a random number f_s between 0 and f_{sum}
3. Select the first member in the population whose fitness, added to the sum of all previous fitnesses, exceeds or equals the random number f_s. If j is the smallest integer for which

$$\sum_{i=1}^{j} f_i \geq f_s$$

is true, the individual j should be selected.

4. Repeat from step 2 until N individuals has been selected

Figure 3.2: Roulette-Wheel Selection-scheme.

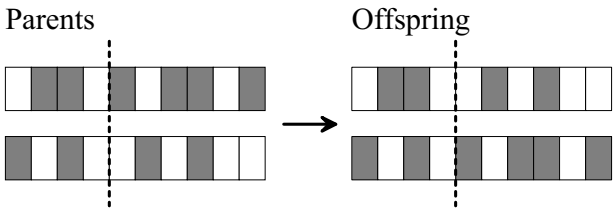

Figure 3.3: Single-Point Crossover. The crossover-point is shown as dashed lines.

part of one individual's chromosome is merged with the second part of the other individual's chromosome, and vice versa. Figure 3.3 shows a Single-Point-Crossover between two individuals with chromosomes of length 10.

Mutation

Of the three basic operators, mutation is the simplest to implement. After two individuals have been selected, and in some case mated, they are to be copied to the next generation. When copying the individuals, the mutation operator is introduced as a small probability that each of the copied genes will change its value, from 0 to 1 or from 1 to 0.

3.1.4 Modifications and improvements

Many modifications and improvements to the "Standard Genetic Algorithm" described above have been suggested. One of these, called fitness scaling, is used in this work. Other modifications as additional selection schemes and variants of the crossover operator can be found in [17] and [18].

Linear fitness scaling

The fitness scaling mechanism is introduced in order to keep a balance in the level of competition between different individuals in a population. In the beginning of an optimization run, the average fitness of the population is often low, due to the random choice of the initial population. In this case the few individuals with a rather high fitness will be selected with very high probability, causing it to dominate the other individuals and risking a premature convergence of the algorithm, possibly to a local maximum. To avoid this, there is a need to scale down the fitness early in a run. Later on in the run, there are many individuals with rather high fitness, and the average fitness of the population may be quite close to the fitness of the best individual, causing the probability that an average individual is selected to be almost as high as the probability of selecting the best individual. In this case the fitness needs to be scaled up, directing the search towards the most promising individuals.

Both of these situations can be solved by a technique known as linear scaling. With this technique the scaled fitness f_i' is calculated through $f_i' = af_i + b$, where f_i is the raw fitness and a and b are suitably chosen scaling constants. Normally the scaling constants are chosen so that the average scaled fitness f_a' is equal to the average raw fitness f_a and the maximum scaled fitness $f_{max}' = C \cdot f_a'$, where C is a constant, normally set to 2. In some cases it is not possible to scale the fitness as much as the constant C requires since it will lead to negative fitness values and in that case the fitness is scaled as much as is possible, so that the individual with the lowest fitness gets the scaled fitness 0. In figure 3.4 the raw and scaled fitnesses are shown for a real situation, early in the run, and later on when the algorithm has converged. In this case the scaling constant C was set to 2.

3.1.5 Measuring convergence by off-line performance

To quantify the performance of a genetic algorithm, some suitable measure is needed. In complex systems the goal is not necessarily only to obtain optimum performance, but rather to reach some satisfying performance quickly. To understand the differences between different performance measures, it is useful to assume that the algorithm is used to find the optimal control parameters for some application, where the fitness of an individual is determined by the actual performance of the application when the control parameters described by that individual are used. In 1975, De Jong introduced the off-line performance, which is a measure of convergence in the cases where not all evaluated individ-

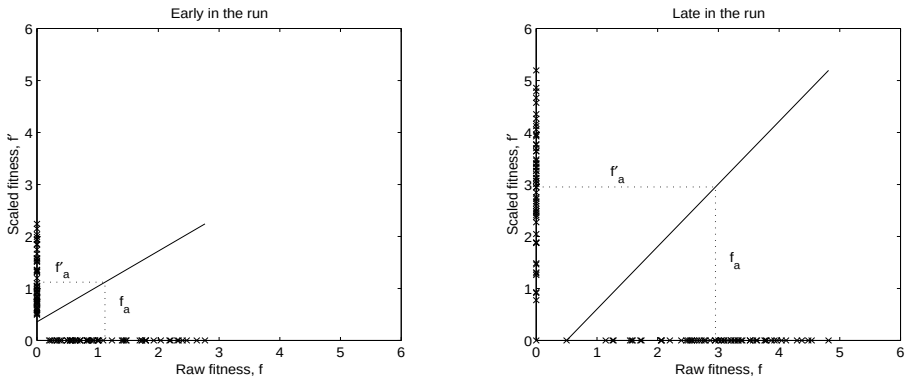

Figure 3.4: Linear fitness scaling. The left plot shows the raw and scaled fitness early in the run and the right plot shows the situation later on when the algorithm has converged. Each cross represents an individual. The dotted lines indicates that the average fitness is not changed when scaled. Note that early in the run the maximum fitness is scaled down while later in the run it is scaled up.

uals have to be used for the application. A typical case is when the individuals are evaluated using a simulation of the application, and the best-so-far individual[4] is used in the application. De Jong also introduced the on-line performance [17], measuring the ongoing performance, which is a better measure in the case where all evaluated individuals have to be evaluated by applying them to the actual application.

The state of a genetic algorithm run can be described using a variable g, being the number of the current generation with $g = 1$ for the initial generation. As a measure of an individual's performance, often the value of the fitness function is used, but in many applications it is often convenient to use some other quantity. In the current work, the goal is to maximize the SHG-efficiency for a laser pulse, and since this efficiency is inversely proportional to the pulse length (see section 2.2.3), the pulse length may be used to calculate the off-line performance. In this case good performance will correspond to a low performance value. If the performance of a single individual in generation g is denoted $u_i(g)$, with i as the number of the individual, the performance of the best individual in the generation may be denoted $u_b(g)$. Defining the "best performance so far" as

$$u_b^*(G) = \text{best}\left[u_b(g)\right],\ g = 1, 2, ..., G \tag{3.6}$$

the off-line performance in generation G is the average of the best performance so far for each generation up to and including generation G:

$$U^*(G) = \frac{1}{G}\sum_{g=1}^{G} u_b^*(g) \tag{3.7}$$

[4]The best-so-far individual is the individual that up to the current point of the run has shown the highest fitness.

By this definition the off-line performance depends not only on the best performance found, but also on whether it was found early in the run. Thus the off-line performance is a relevant measure of the convergence of a genetic algorithm.

3.2 Implementation

3.2.1 The LabVIEW programming environment

LabVIEW is a graphical programming language using dataflow programming, meaning that the execution of a program written in LabVIEW is determined by the flow of data rather then by the instructions. A LabVIEW program, called a Virtual Instrument (VI), consist of a front panel and a block diagram. The front panel represents the user interface, or if the VI is a subroutine, the inputs and outputs. The block diagram contains graphical representations of functions wired together so that it resembles a flowchart. The objects on the front panel are represented in the block diagram as data sources or sinks.

The graphical interface makes LabVIEW an ideal environment for writing programs that are easy to grasp, and thus easy to modify and debug. Its many predefined functions and its built-in capabilities of communication with peripherical equipment make it easy to write programs that work well in many experimental situations. The genetic algorithm used in this study was written in LabVIEW.

To make the implementation more abstract, and thus easier to modify later, the whole algorithm was written to work with a set of predefined structures or building blocks. These structures are either pure data types, describing an individual or a whole population, or templates, defining operators and functions by specifying their inputs and outputs.

In section 3.2.2 the connections between the genetic algorithm and the experimental setup are discussed, while the different building blocks used are described in Appendix C.

3.2.2 Evaluation process

So far only the genetic algorithm has been considered. There is however a problem-dependent part required to get the algorithm running, namely the evaluation step. First of all, the algorithm needs to know how to decode an individual's chromosome into some useful information and second, how to measure the fitness associated with this information.

In this work the individuals of a population will each represent a waveform to be sent to the Dazzler (see section 2.4.3). These waveforms are represented as parameters in a parameter file which is fed to the Dazzler control software, thus the decoder function needs to know how to map a string of bits into parameters for the Dazzler. The decoder function will be further discussed in the next

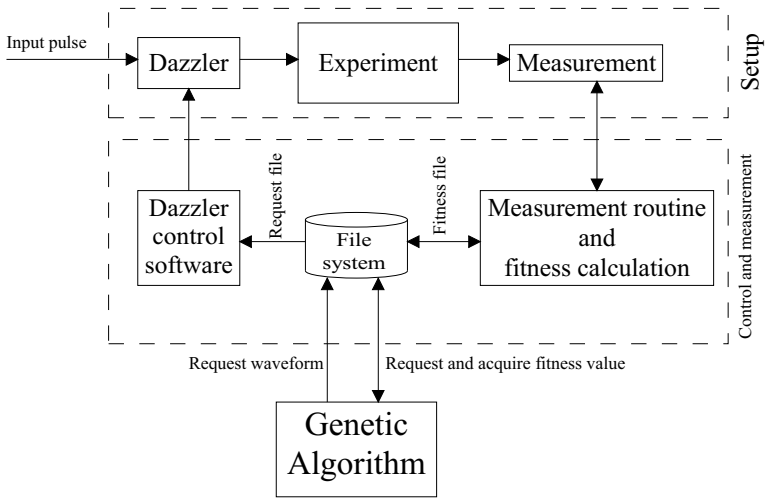

Figure 3.5: Sketch of the relationship beween the experimental setup, the control and measurement routines and the genetic algorithm. By letting all communication with the genetic algorithm pass through the file system, the actual algorithm implementation is independent of the practical details of the experiment.

section.

After a waveform has been requested, the algorithm needs some way of measuring the fitness of the current waveform. In order to get a general adaptive system, the measurement routine and the fitness calculation has been implemented separated from the genetic algorithm, working in the same way as the Dazzler control software in its loop-mode as described in section 2.4.3. This means that the fitness calculation routine is idle as long as no request for fitness is made. The communication with the routine is made through a specific file, hereafter called the fitness file. When this file is available, the fitness calculation routine stays idle. When the algorithm wants the fitness of the waveform that it has just requested, it deletes the file. When the file is deleted, the fitness routine does some measurement, calculates the fitness from the result of the measurement, and writes the fitness value back to the fitness file.

In figure 3.5 a sketch of the relationship between the actual experimental setup, the control and measurement routines and the genetic algorithm is shown. As can be seen, using this implementation, the genetic algorithm is not dependent or adapted to any specific type of measurement, since it really doesn't matter what the measurement routine actually does, as long as it handles the fitness file as the algorithm expects it to.

Chapter 4

Experimental methods and results

4.1 Setup

4.1.1 Laser source

In the experimental part of this work, laser pulses from the oscillator of the Lund T^3-laser[1] was used. Normally this oscillator is used to seed a chain of amplifiers operating at a pulse repetition rate of 10 Hz. To be able to pre-compensate amplitude- and phase-effects of the amplification, a Dazzler AOPDF triggered at 10 Hz is placed directly after the oscillator. By using the pulses not diffracted by the Dazzler, i.e. those that are not used for amplification, the experiments could be made without interfering with other ongoing experiments at the facility.

The active medium of the oscillator is a crystal of titanium-doped sapphire (Ti:Al_2O_3), pumped at a wavelength of 514.5 nm by an Argon-ion laser. The width of the gain profile of Ti:sapphire is large, about 230 nm with the peak at 800 nm, theoretically allowing for tranform-limited pulses of 4 fs duration (see section 2.1.4). To obtain pulsed operation of such a laser, the phases of the individual frequency-components have to be locked relative to one another, a technique known as mode-locking. Practically this can be achieved by modulating the losses of the oscillator cavity in such a way that one favors frequency-components with relative phases giving rise to short pulses [8]. In the oscillator used in the present work, mode-locking is achieved by a technique called Kerr-lens mode-locking. The optical Kerr-effect, or the quadratic electro-optic effect [10], causes the refractive index of a material to be intensity-dependent. For a laser beam with a gaussian spatial profile, this causes the material to act essentially as a lens [8] with a focal length being shorter for higher intensities. By letting the active medium act as a Kerr-lens and aligning the cavity in such a way that a harder focused beam experiences smaller losses, mode-locking is

[1] Table-Top Terrawatt laser

achieved. To compensate for cavity-dispersion, i.e. the dispersion introduced by the optical elements in the cavity, the oscillator has an intra-cavity prism-compressor (see section 2.4.1).

During the experiments described here, the bandwidth of the oscillator was generally around 30 nm, allowing transform-limited gaussian pulses of about 30 fs. The repetition-rate of the oscillator is 80 MHz with an average power of about 500 mW, giving a pulse energy of 6-7 nJ.

The experiments were performed in another room than the one of the T^3-laser and thus the beam had to be transferred a distance of about 15 m before entering the experimental setup. As mentioned in section 2.4.3 it is important to have a collimated beam at the Dazzler and to achieve that a telescope placed directly after the T^3-oscillator was aligned so that the beam waist was located at the Dazzler crystal.

4.1.2 Pulse preparation

In figure 4.1 the experimental setup is depicted with the input beam entering the optical table from the left.

Since the input beam was taken from the non-diffracted (zeroth order) beam of the T^3-Dazzler, it was necessary to compensate for at least 12500 fs^2 of GDD introduced by that Dazzler. The actual compensation needed will actually be somewhat larger due to the positive GDD introduced by the telescope and the propagation distance in air. There is also a small positive GDD in the beam leaving the oscillator, due to dispersion in a part of the titanium-sapphire crystal and in the outcoupling mirror. The total dispersion introduced by these elements can be estimated to be no higher than 1000 fs^2. As can be seen in figure 2.12 a compensation of this magnitude would give a cut-off bandwidth of less than 50 nm, which is too narrow, making an external compressor necessary. Therefore a four-pass prism compressor was used with prisms made of SF10 glass (see figure 4.1).

Figure 4.2 shows intuitively how to think when considering the necessary GDD to be introduced by the compressor. In this picture the dispersion effects of the used Dazzler are considered to be introduced by a dispersive element. The Dazzler itself is considered to exhibit zero dispersion. As discussed above the GDD of the pulses entering the optical table can be estimated to be $12500 + 1000 = 13500$ fs^2 due to the dispersion of the oscillator, the first Dazzler, the telescope and the propagation in air. Including the dispersion of the used Dazzler through the dispersive element, this leaves us with a total positive GDD of 26000 fs^2. As seen in the figure, the pulse entering the Dazzler is called the "input pulse", and from here on when talking about the input pulse of the Dazzler, the dispersive effects of the Dazzler has already been included. If one wants to introduce a GDD of -10000 fs^2 with the Dazzler, the GDD of the input pulse has to be 10000 fs^2. Hence the compressor has to be aligned so that it introduces -16000 fs^2 of GDD. As seen in section 2.4.3 it is possible to introduce a maximum negative GDD of 12500 fs^2 with a cut-off bandwidth of 100 nm, giving a rather large range for

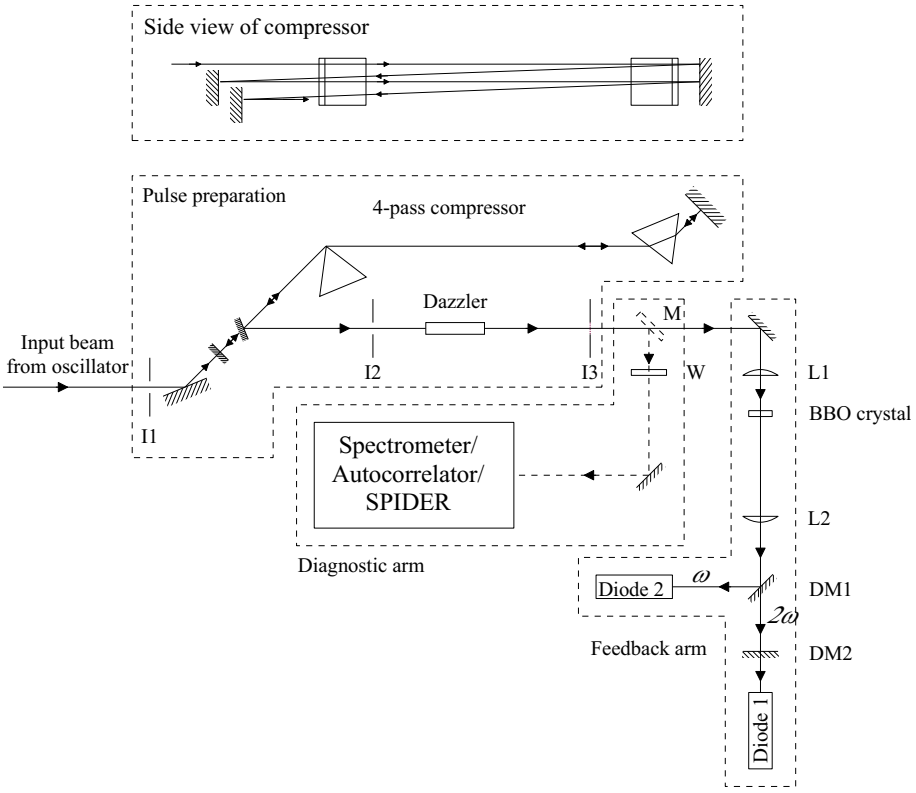

Figure 4.1: Experimental setup. The input beam passes the compressor before entering the Dazzler crystal. After the crystal the zeroth order beam is blocked and the diffracted pulses are sent either to the feedback arm or to a diagnostic arm by the insertion of the mirror M. In the feedback arm the beam is focused in the BBO crystal by the lens L1. The lens L2 is placed so that a focus is obtained at diode 1, maximizing the signal for the frequency-doubled light. DM1 and DM2 are dichroic mirrors, separating the fundamental component from the second order harmonic. Above the picture a side view sketch of the passage of the beam through the prism compressor is shown.

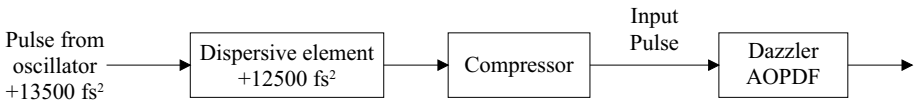

Figure 4.2: Picture to regard when considering the amount of GDD necessary to introduce by the compressor. The dispersive effects of the Dazzler are regarded as a dispersive element placed before the compressor.

second-order phase-corrections with the operating point chosen as -10000 fs^2. A theoretical calculation (see [22]) shows that the prism distance required to introduce this amount of GDD in four passes is around 0.8 m, and this was the distance used during the first experiments.

After the compressor the beam was sent through the Dazzler. The polarization of the input beam was in this case horizontal and the Dazzler was therefore aligned side-up to match this polarization (compare with figure 2.13). Thus the polarization of the diffracted beam was vertical, matching the required input polarization for the used autocorrelator. After the Dazzler an iris (I3) was used to block the zeroth order beam, letting the diffracted beam pass. In figure 4.1 two more irises can be seen (I1 and I2) placed at the entrance of the setup and before the Dazzler respectively. These were used as references when realigning the setup, assuring the correct beam path through the compressor and the mirrors before the Dazzler. This arrangement was necessary due to the large long-term spatial drift of the input beam, resulting from the long propagation distance and many reflections of the beam prior to the setup.

Finally the diffracted pulses were sent to the feedback arm or alternatively to a diagnostic arm by the insertion of the mirror (M).

Triggering and timing

When the Dazzler is used in a laser system, the triggering of the acoustic waveform is made at the same rate as the pulse picking mechanism in the system. In this work no pulse picker was used and the repetition rate for the input pulses was 80 MHz. The Dazzler was triggered at a constant rate of 1 kHz, in order to imitate the conditions of a complete laser system. With a trigger rate of 1 kHz, there is no risk to exceed the maximum allowed acoustic power density of the crystal, making it possible to use the device in loop-mode without the need to worry about applying an excessive power. A drawback of operating the Dazzler without a pulse picker is that not only the requested pulse shaping takes place. In figure 4.3 a time-domain picture of the result of such operation is shown.

Pulses that arrive when the entire acoustic wave is in the crystal will be diffracted as expected, while pulses arriving at times when only a part of the acoustic wave is present will only be partly diffracted, seen as the wings of the diffracted pulse envelopes in figure 4.3. This means that the feedback measurement should be made within the time $t_c - t_s$ indicated in the figure, measuring only on correctly diffracted pulses. In a system with a pulse picker, the used pulse should be picked from inside this time range. For the used crystal $t_c \approx 24\ \mu s$, meaning that for an acoustic wave filling 90% of the crystal, there will be about 200 useful pulses when using a repetition rate of 80 MHz. Due to the fact that the acoustic wave has a linear chirp, the pulses that are diffracted at the wrong times will get a narrower bandwidth with a spectrum built up only by the optical frequencies corresponding to the acoustical frequencies present in the crystal at that instant. One can thus expect the incorrect pulses to have a pulse length much larger than the correct ones (see section 2.1.4). The phase within this bandwidth will however still be the applied one.

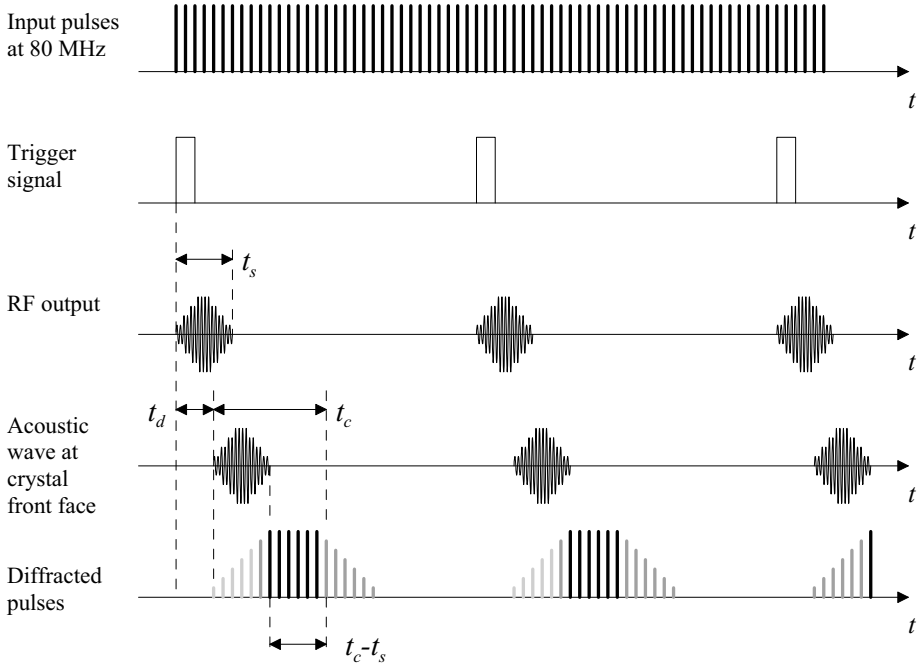

Figure 4.3: Time-domain picture of diffraction of a pulse train without the use of a pulse picker. t_s is the duration of the acoustic wave. $t_d \approx 5\ \mu s$ is the delay related to the propagation of the RF signal to the crystal. $t_c \approx 24\ \mu s$ is the propagation time for the acoustic wave through the crystal. Only pulses within a time range of $t_c - t_s$ are diffracted by the whole acoustic wave.

4.1.3 Feedback signal

The goal of this work is to use the genetic algorithm to optimize the SHG-efficiency since it in section 2.2 was concluded that the second harmonic generation is more efficient when the pulse is shorter. It was also seen that the SHG-efficiency is dependent on the pulse energy (see figure 2.1) and as discussed in section 2.4.3 the acousto-optic interaction efficiency, and hence the diffracted pulse energy, is dependent on how the acoustic wave fills the crystal. This makes the feedback signal dependent on the applied amplitude and phase through the pulse energy as well. However if one assumes that the diffracted pulse energy is only weakly dependent on the diffraction conditions around a maxima of the SHG-efficiency, the dependence of the feedback signal on pulse energy might be an advantage, helping the algorithm to avoid solutions with low pulse energy.

To measure the SHG-efficiency the diffracted pulses are focused by the lens L1 with a focal length of 50 mm into a BBO crystal (BaB_2O_4). The crystal is aligned so that second harmonic generation is achieved and by the dichroic mirrors DM1 and DM2 the fundamental and frequency-doubled pulses are separated and detected by two photodiodes. The second harmonic pulses are weak and to enhance the signal, a lens L2 with a focal length of 100 mm is used to focus the frequency-doubled light onto diode 1. During this work it was not practically possible to gate the measurement and hence the entire diffracted pulse train of figure 4.3 was detected. To get a strong enough signal the signals from the photodiodes had to be terminated with a 1 MΩ resistance, making it impossible to temporally resolve even the envelope of the diffracted pulses due to the associated slow response time. Instead the peak-to-peak value of the diode signals had to be used as a measure of the total energy in the diffracted envelope.

Another important issue is that the ratio of the diode signals is not the SHG-efficiency as defined in section 2.2.3 since the fundamental pulse energy should then be measured before the second harmonic generation. For low conversion efficiencies the measured signal will however be a very good approximation of the SHG-efficiency since the depletion of the fundamental beam can then be ignored. Using one of the acoustic waveforms resulting from an optimization, the SHG-efficiency was measured by a power meter to be lower than 3%. This means that the non-depleted approximation does indeed hold in the experiments presented here (see figure 2.1 and figure 2.2).

Practically the diode signals were sent to an oscilloscope triggered on the RF signal from the Dazzler RF generator. The built-in measurement function of the oscilloscope was used to determine the peak-to-peak values of the two signals and the results were transferred using GPIB[2] to the measurement routine as described in section 3.2.2.

[2]General Purpose Interface Bus

4.1.4 Diagnostics

In order to be able to get quantitative results from an optimization, a diagnostic arm was also present in the experimental setup. By inserting the mirror M in the diffracted beam from the Dazzler, the light was coupled to the diagnostic equipment. The window W as seen in figure 4.1 was made of the same material as the lens L1 having a thickness almost equal to the lens. The reason for introducing this additional material into the diagnostics arm is understood if one considers that the actual measurement of the SHG-efficiency is not done directly after the Dazzler but rather in the BBO crystal, after the lens L1. The effect of the window is then to assure that the phase measured by the diagnostic equipment is equivalent to the one for which the SHG-efficiency was measured during the optimization. During the experiments aimed at choosing a reproductive plan a multi-shot intensity autocorrelator and a grating spectrometer were used as diagnostics while the analysis of the applied spectral phase was done using SPIDER as described in section 2.3.

The autocorrelator used was a Femtochrome FR103-XL autocorrelator. In this device the time delay is achieved by a rotating parallel mirror assembly. The separate scans are made with a repetition rate of 10 Hz and the signal from the photodiode may be viewed on an oscilloscope triggered from the autocorrelator. As a first approach the autocorrelator was operated separately while the Dazzler was set to operate in continuous mode when measuring the pulse length. With a pulse train of identical pulses this method is often used successfully but by studying figure 4.3 one can understand why it does not work here. Even if the Dazzler is operated continuously there will always be pulses that have passed the crystal when the acoustic wave is not present in full. These pulses will, as discussed in section 4.1.2, have a pulse length longer than the one resulting from a correctly diffracted pulse and hence some of the autocorrelator scans will be incorrect and even asymmetric. Tests confirmed this and when averaging the autocorrelator signal over several scans the resulting trace was much wider than expected leading to an incorrectly measured pulse length.

In order to be able to measure the pulses correctly with the autocorrelator it is necessary to synchronize the autocorrelator scan with the diffracted pulse envelopes in such a way that the actual signal from the autocorrelator is recorded within the time range $t_c - t_s$ depicted in figure 4.3. As a result of this it will not be possible to measure the correct pulse length of pulses that are so long that they can not be scanned by the autocorrelator in the time $t_c - t_s$. For an acoustic wave filling the crystal well, as required for high interaction efficiency, the time range $t_c - t_s$ will be short and thus even shorter pulses are required if they are to be measured correctly.

To synchronize the autocorrelator with the diffracted pulse envelopes, it was necessary to trigger the RF generator on the autocorrelator trigger signal. A digital time delay was used for this signal and by adjusting the delay between the autocorrelator trigger and the launch of the acoustic waveform it was possible to diffract the pulse envelope just when the autocorrelator scan was at the position of overlap between pulses. By fine tuning the delay it was seen that there existed a maximum for the autocorrelation signal, as predicted by this discussion, just

Parameter	*Unit*	*Min*	*Max*	*Precision*	*Gene*
width	nm	0	90	0.09	70..79
hole position	nm	700	900	0.20	60..69
hole width	nm	0	70	0.07	50..59
hole depth	au	0	1	$1.0 \cdot 10^{-3}$	40..49
delay	fs	0	3600	3.5	30..39
order 2	fs^2	-20000	-5000	14.7	20..29
order 3	fs^3	-100000	100000	195.5	10..19
order 4	fs^4	-300000	300000	586.5	0..9

Table 4.1: Waveform parameters used for the genetic algorithm.

when the scan was including the correctly diffracted pulses. Due to the fact that the diffraction in the Dazzler crystal might take place at different positions, the diffracted pulse envelope may also be delayed in time, making it necessary to do the fine tuning over again for each pulse to be measured on.

From this discussion it is realized that the autocorrelator measurement will probably not give any reliable value for the pulse length, at least not for long pulses, since the measurement might, and probably will, be made partly on pulses that are not the correct ones. The problems described are solely due to the fact that there are no pulse picker in the used system.

4.2 Method

The reproductive plan is the particular choice of reproduction operators together with the associated parameters for those operators (see section 3.1.3). In this experiment the standard set of operators described in section 3.1.3 was implemented and used. The population size and the probabilities for crossover and mutation were varied in order to find a combination leading to good convergence measured by the off-line performance with respect to pulse length. To improve the effect of the selection mechanism, a linear fitness scaling mechanism as described in section 3.1.4 was used.

4.2.1 Encoding of waveform parameters

The waveform parameters corresponding to the Dazzler front panel controls was encoded using ten bits each, except for the position of the supergaussian amplitude which was fixed to the central wavelength of the spectrum. In table 4.1 the used parameters are shown together with their minimum and maximum values and the precision obtained. In figure 4.4 a picture of the actual chromosome with 80 genes is shown and the location of each parameter is indicated. In the table it can be seen that the precision of the three first parameters is unnecessarily good, since the actual wavelength resolution of the Dazzler is around 0.6 nm (see section 2.4.3). This however is not a problem and simply means

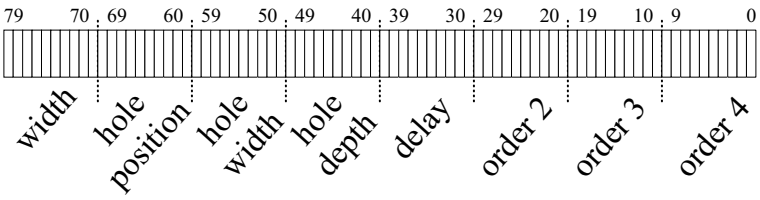

Figure 4.4: The chromosome with 80 genes used to encode the parameters.

that some of the genes in the chromosome are not important in the search.

4.2.2 Description of a genetic algorithm run

As mentioned in section 3.1, a genetic algorithm run corresponds to a single optimization, starting from a random population and stopping when some predefined stop-criteria is met. The stop-criteria might be that a certain performance is achieved or that a certain number of evaluations have been made. In this study several algorithm runs were to be made, some of them with different population sizes. To be able to compare those it is important that one measures the progress in terms of the number of evaluations rather than in terms of generations. This is due to the fact that for different population sizes, a different number of evaluations have to be made for each generation. Hence the stop-criteria was here chosen to be a certain number of evaluations, namely 1200.

Due to the long-term drift of the input laser beam, the fitness value, i.e. the measured SHG-efficiency, could not be used for reliable comparison between different generations in the run. Hence the pulse length was measured when the run was completed as discussed in section 4.1.4 and the off-line performance was then calculated with respect to the measured pulse length (see section 3.1.5). To calculate the off-line performance, the pulse length of the best-so-far individual, i.e. the individual representing the shortest pulse length up to the current time, has to be calculated for each generation in the run. It was however practically impossible to measure the pulse length of every single individual in each generation[3], and hence only the pulse length for the best individual in terms of fitness was measured in each generation. In figure 4.5 the best-so-far pulse length is plotted together with the calculated off-line performance in terms of pulse length.

It can be noted that in the experiments presented here, the "true" pulse length will always be smaller than the value of the off-line performance. Firstly because the off-line performance is a running average, including the pulse lengths from all previous generations, and secondly because of the limitations on the autocorrelator measurements discussed in section 4.1.4. Being an important point, the conclusion from section 3.1.5 is repeated here: The off-line performance in

[3]Doing this would require 1200 pulse length measurements for each single run, and for the current experiment involving 60 separate runs, a total of 72000 single measurements.

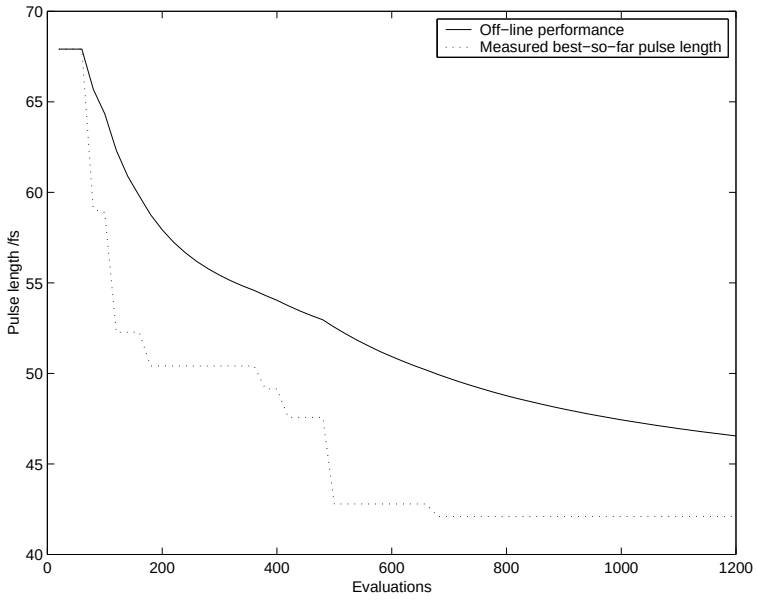

Figure 4.5: The best-so-far pulse length plotted together with the calculated off-line performance in terms of pulse length for a single genetic algorithm run.

terms of pulse length is a better measure of convergence than the actual pulse length, since it does not only depend on the current pulse length, but also on whether it was reached early or late in the run.

4.2.3 Choice of a reproductive plan

The choice of the parameters N, p_c and p_m which is the population size, the crossover probability and the mutation probability respectively, is not straightforward. This is due to the fact that the optimal parameters are critically dependent on the nature of the problem itself. There exist however some recommendations based on experiences which are useful as a starting point.

When it comes to the choice of the population size a larger population leads to a decreased number of reproductive steps for a fixed number of fitness evaluations, causing the evolution and thereby the convergence to slow down. A too small number of individuals makes the population unable to maintain its diversity and the algorithm's ability to simultaneously search around different maxima is degraded [19]. This would then make the algorithm more likely to converge to a local maximum rather than to the global one. In this test population sizes of 20, 40, 60 and 100 were used. The mutation is important in making the search try unexplored parts of the space while a too high mutation rate makes the search analog to a random search. In general mutation rates between 0.001 and 0.01 are recommended. In the present study mutation rates of 0.001, 0.005, 0.01, 0.05 and 0.1 were tested.

For the crossover probability, in [20] it is suggested a value of 0.6 for large population sizes while in [21] it is shown that for smaller populations it is reasonable to increase the crossover probability to 0.9 in order to prevent premature convergence. In the experiments performed here, a crossover probability of 0.8 was used while the population size and mutation rate was varied. For the chosen N and p_m, the crossover rate was then varied between 0.5 and 1.0 in steps of 0.1.

To be able to compare the different runs the optimizations were made for a fixed number of 1200 evaluations. With the population sizes 20, 40, 60 and 100 this means that the algorithm was run for 60, 30, 20 and 12 generations. For each combination of parameters three independent runs were made and the results were averaged in order to suppress the noise introduced by day-to-day variations of the input pulses.

Choice of population size and mutation rate

In figure 4.6 the off-line performance is plotted for the different combinations of population size N and mutation rate p_m. These results give an indication of which population size and mutation rate that is best adapted to the problem at hand. An indication of lack of statistics can be seen by regarding the values of the off-line performance for the initial population, as seen leftmost in each plot of figure 4.6. Since the initial population is selected at random the initial off-line performance should be the same for a given population size independently of the mutation rate. This makes difficult the comparison between the different curves, but however there is a clear converging trend in all of them.

The left plot of figure 4.7 shows the average initial off-line performance as a function of population size. As can be seen the performance of the initial population decreases with population size, which is expected since a larger number of individuals means a higher probability that at least one of them will be a high scoring individual. However this strong dependence might also indicate that the problem posed is actually not complicated enough to benefit from a genetic algorithm approach. This becomes especially clear for the population of 100 individuals where the initial guess is not at all far from the result of the search after 1200 evaluations.

The right plot of figure 4.7 shows the final off-line performance for each population size N plotted against the mutation rate. Due to the low number of averages taken these results are not that clear but one can see a tendency of higher mutation rates for the smaller populations. For $N = 20$ the best results are obtained with the highest mutation rate of 0.1 while for $N = 40$ a mutation rate of 0.05 seems to be a better choice. This can be expected since it is hard to maintain the diversity in small populations and hence mutation helps to reintroduce gene values that might otherwise be lost due to the selection mechanism. For the population of 60 individuals the best performance is achieved for a mutation rate of 0.01, but here it is also seen that the highest mutation rate leads to high performance. The population of 100 individuals seems to be rather independent on the mutation rate with just a slight tendency for higher probabilities. This is probably due to the good initial guess obtained due to

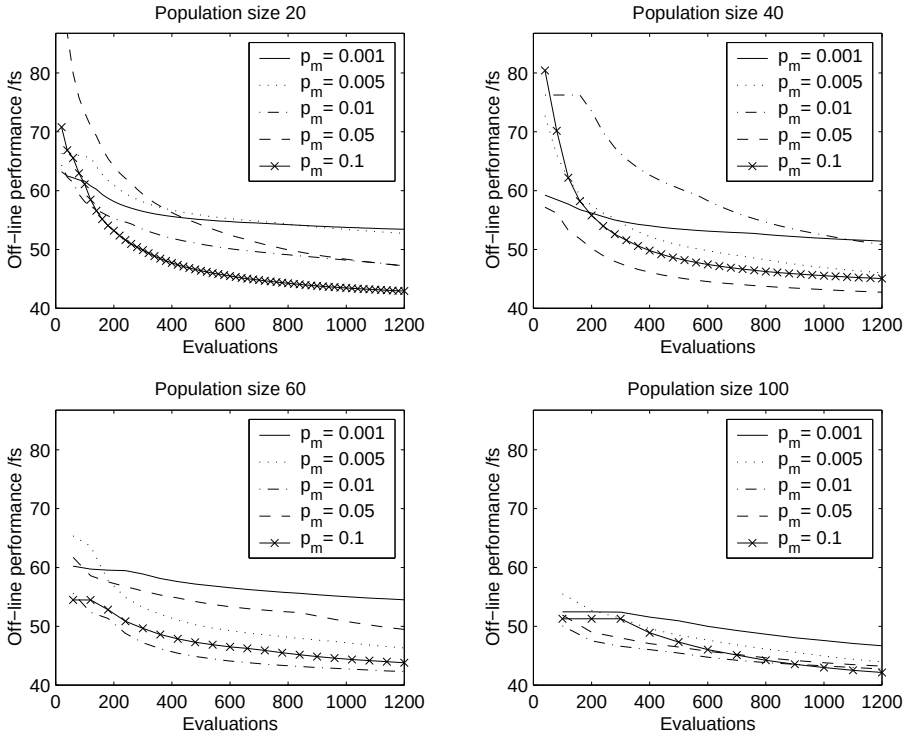

Figure 4.6: Off-line performance averaged for three independent optimization runs involving 1200 evaluations each. The different plots are for different population sizes N and for each N is shown the off-line performance for different mutation probabilities p_m.

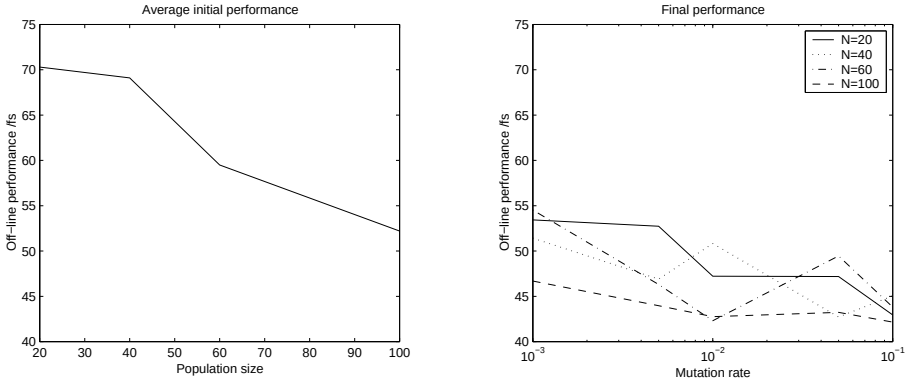

Figure 4.7: Average initial and final off-line performance plotted against the population size and the mutation rate respectively.

the lack of complexity as described above. For the analysis of all these cases caution has to be taken when regarding the highest mutation rate. A probability of 0.1 for mutation is high compared with the recommended ones and is very likely to in principle cause a random search. This is especially true for the larger populations where less reproductive steps are taken for a fixed number of evaluations.

It is difficult to draw reliable conclusions from the runs made. The best performance measured was for a population size of 100 individuals and a mutation rate of 0.1. The mechanism behind the results in those cases is however very likely to be more random than evolutionary based on the good initial guess made for all mutation rates with this population size. If that reproductive plan was then to be applied to a more complex problem, it could not be expected to perform better than a random search. For the population of 60 individuals there is however a minimum value for a mutation rate of 0.01, with almost as good performance as for the largest population. This choice of population size and mutation rate is also in good agreement with the ones recommended in the literature [18], [21].

Choice of crossover probability

With the chosen values for the population size and the mutation rate two independent runs were made for different values of the crossover probability and the resulting off-line performance is shown in figure 4.8. The best performance was obtained for a crossover probability $p_c = 0.9$, but it is also seen that the differences in performance are rather small except for $p_c = 0.5$ and $p_c = 0.8$ which suffers from a poor initial off-line performance due to the small number of averages taken. The weak dependence of the performance on the crossover probability is disturbing in the sense that it implies that the search is not so strongly dependent on the evolutive process as one may expect. This is however in agreement with the above suspicions that the problem is not complex enough to benefit from a genetic approach.

4.3 Results

4.3.1 Pulse length analysis

Here the main results based on pulse length measurements with the autocorrelator are presented and analyzed. For this three independent runs of 2400 evaluations each were made using the reproductive plan obtained in section 4.2.3, i.e. $N = 60$, $p_m = 0.01$ and $p_c = 0.9$. The resulting average off-line performance is shown in figure 4.9 together with the average of the shortest measured pulse lengths so far. As discussed in section 4.1.4, the pulse length can probably not be trusted to be correct early in the run, but the plot still shows that the initial decrease in pulse length is rather fast, finding pulses with the length 41 fs after around 600 evaluations. From there on, the convergence

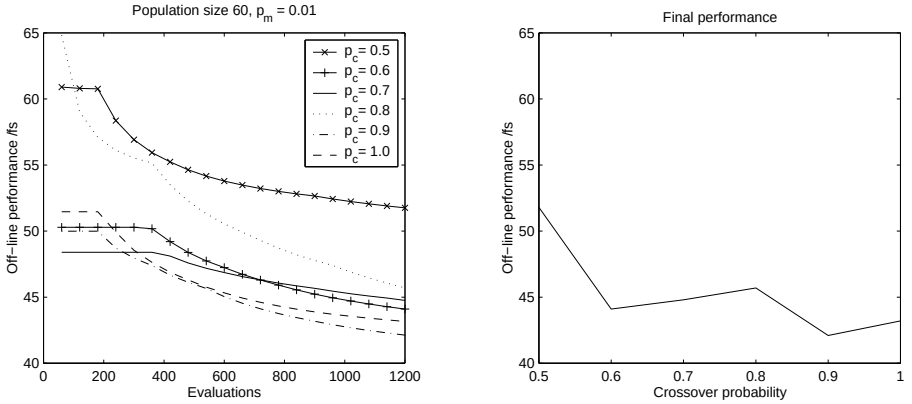

Figure 4.8: Performance for $N = 60$ and $p_m = 0.01$ for different crossover probabilities p_c. To the left is the off-line performance for 1200 evaluations and to the right the final off-line performance is plotted against the crossover probability.

is slower, but still the pulse length decreases steadily along the run.

To study the actual effect of the optimization one of the runs can be studied more closely. The optimization leading to the shortest pulse length resulted in an individual with a pulse length of 39 fs after 1800 evaluations. The applied amplitude filter for this individual is shown in figure 4.10 together with the measured spectrum both before and after optimization. The bandwidth of the spectrum is 31 nm, and the transform-limited pulse length obtained by the inverse Fourier transform is 36 fs.

As can be seen from the figure the amplitude filter is chosen in such a way as to make the spectrum more symmetric and thus increasing the FWHM-bandwidth. An inverse Fourier transformation of the measured spectra assuming a flat phase shows that the decrease in pulse length obtained by the amplitude filter is about 1 fs. This form of the amplitude filter was observed in almost all of the runs, meaning that the algorithm is able to apply an amplitude filter in such a way as to increase the spectral bandwidth and thus decrease the pulse length.

The spectral phase determines the pulse shape as seen in section 2.1.2 and in figure 4.11 the coefficients of the phase polynomial is plotted for each individual in the population throughout the run. The second order phase is crucial for the pulse length and as can be seen, it quickly finds its value around -11500 fs^2 and stays there for the rest of the run. The convergence of the second order phase is then responsible for the first rapid decrease in pulse length seen in figure 4.9. The GDD introduced by the compressor was in section 4.1.2 chosen so that the input pulse should have a positive GDD of 10000 fs^2 and thus the found optimal second order phase deviates with -1500 fs^2 from the expected value. The theoretical calculation made was only aimed at finding an approximate prism distance so the actual GDD of the input pulse will probably deviate a little from the expected value. That the error would be as large as 1500 fs^2 is however hard

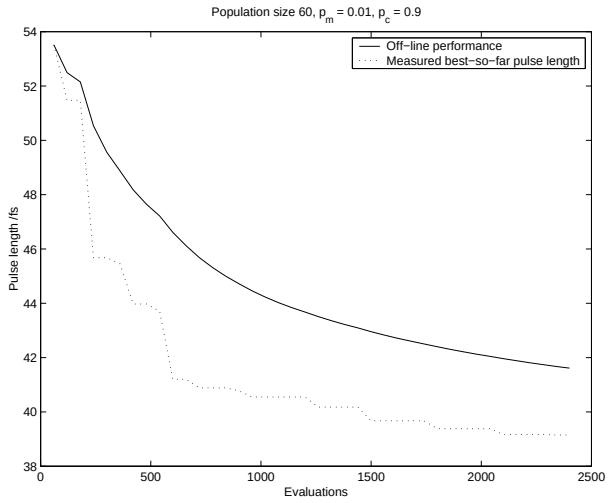

Figure 4.9: Performance for $N = 60$, $p_m = 0.01$ and $p_c = 0.9$. In the plot both the off-line performance and the shortest pulse length so far is shown for a genetic algorithm run of 2400 evaluations. The plot shows the average of three independent runs.

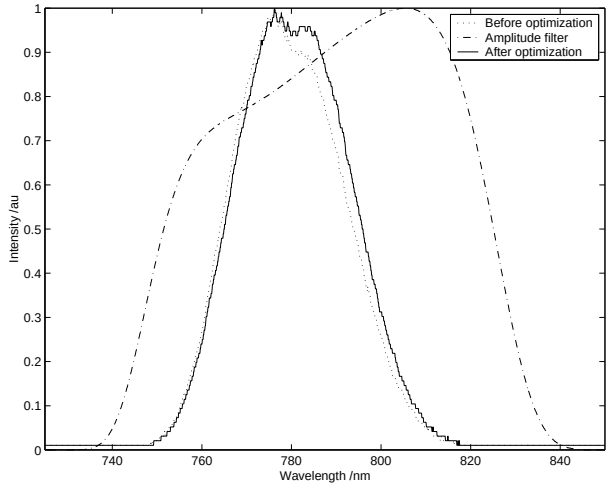

Figure 4.10: Spectrum measured both before and after the optimization together with the amplitude filter function.

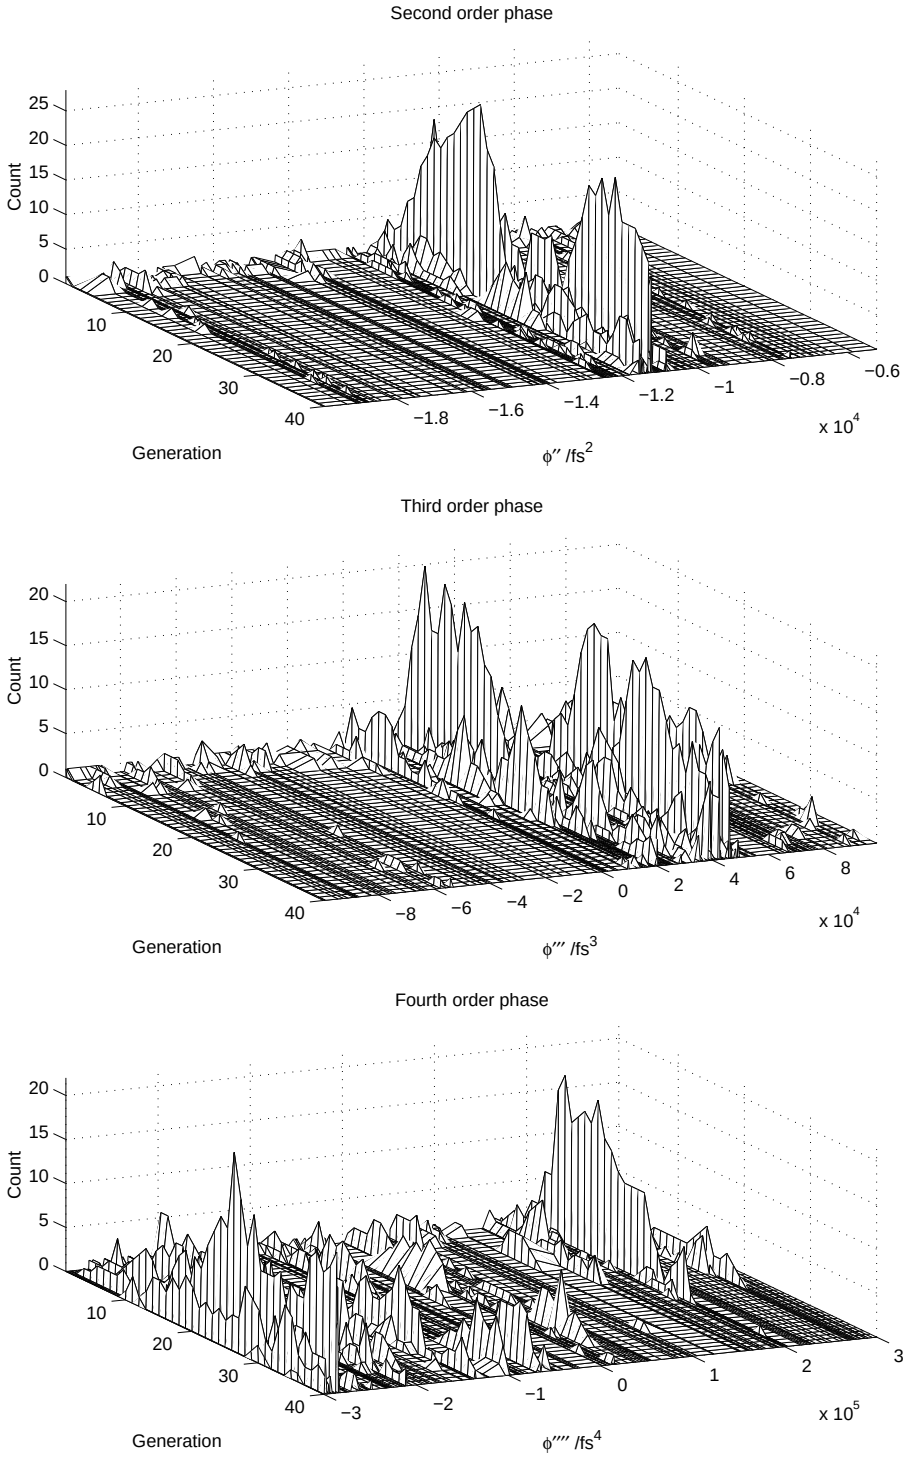

Figure 4.11: Evolution of the coefficients of the phase polynomial for each individual throughout the run.

Parameter	*value*
width	80 nm
hole position	762 nm
hole width	65 nm
hole depth	0.30
delay	1601 fs
order 2	-11466 fs^2
order 3	45845 fs^3
order 4	-139296 fs^4

Table 4.2: Waveform parameters of the individual with a pulse length of 39 fs found after 1800 evaluations.

to believe, and even if it was so the optimization should still be able to find an optimal pulse length closer to the Fourier limit of 36 fs. The conclusion must then be that at least some of the deviation is due to a difference between the shortest pulse and the pulse showing the highest fitness.

There are several possible explanations to this difference. First of all it might have to do with the differences between measuring the SHG-signal and measuring the pulse length. As seen in figure 4.1 a window W is used to introduce the same dispersion for the pulses in the diagnostic arm as the lens L1 does in the feedback arm. This compensation might not be exact and thus the pulse length is measured for a pulse with a slightly different phase than the pulse for which the feedback signal was measured during the run. The difference in introduced GDD between the window and the lens will however be very small (around some hundred fs^2). Another issue, briefly mentioned in section 2.2.1, is the phase-matching conditions for short pulses in the BBO-crystal. With the broad spectra of short pulses the phase-matching conditions will not be the same for all frequencies, and thus it is not guaranteed that the shortest pulse at the entrance of the SHG-crystal will yield the highest SHG-efficiency.

Another possible explanation is based on the dependence of the SHG-efficiency on the pulse energy as seen in section 2.2.3. Since the diffracted pulse energy in the Dazzler crystal is strongly dependent on the applied second order phase it is possible, and even probable, that the maximum of the SHG-efficiency will occur for an applied GDD that is a trade-off between a flatter phase and an increased pulse energy, thus with a tendency toward a larger negative GDD.

The third order phase converges as well, but is not so well localized as the second order term. For the fourth order term convergence can not be concluded from the figure. The reason is probably that the pulse length is not so sensitive to the fourth order term and hence the algorithm has difficulties in finding an optimum value. The waveform parameters of the individual with a pulse length of 39 fs found after 1800 evaluations is shown in table 4.2.

It is useful to investigate how the different phase coefficients affect the pulse shape. If one assumes that the found optimal individual has a flat phase[4], i.e.

[4]This is of course not entirely true since the pulse length of a transform-limited pulse would

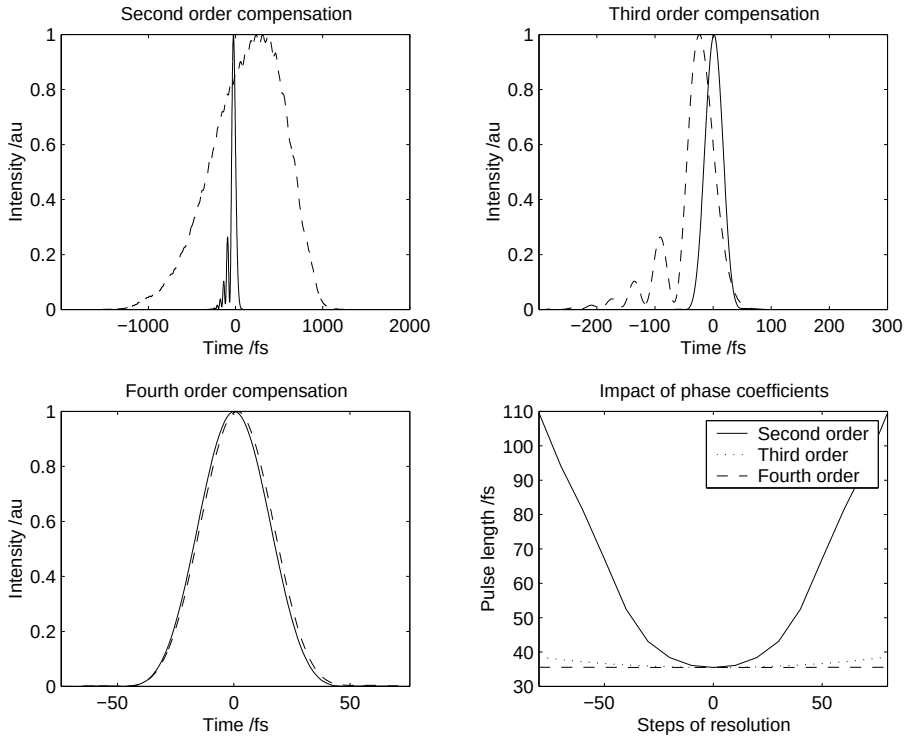

Figure 4.12: Results from simulation. The first third plots shows what happens when the second, third and fourth phase orders are corrected one at a time. The dashed line shows the pulse shape before the correction is applied and the solid line shows the pulse shape after the correction (Note that the time scales are different). The bottom right plot shows the resulting pulse length when the different orders are varied around zero in units of their resolution during the optimization run.

all phase coefficients zero, the input pulse[5] will have the opposite of the phase seen in table 4.2. In figure 4.12 the result of a simulation using the measured spectra is shown. In the top left plot, the temporal profile of the input pulse is shown as the dashed line, while the solid line shows the pulse shape after the second order phase correction of -11466 fs^2 has been applied. In the top right plot the dashed line shows again the pulse shape after the second order correction, this time on another time scale, and the solid line shows the temporal profile after the third order phase correction of 45845 fs^3 has been applied. The bottom left plot in the same way shows the effect of the fourth order phase correction of -139296 fs^4.

It can be seen that the second order compensation shortens the pulse significantly but is unable to correct for the asymmetry of the pulse which is taken care of by the third order compensation. The second order compensation short-

be 36 fs and the measured pulse length is 39 fs or 1.08 times the transform-limit.

[5]Remember the definition of the input pulse from section 4.1.2 and figure 4.2.

Parameter	*300 evaluations*
delay	1548 fs
order 2	-10220 fs^2
order 3	38807 fs^3
order 4	-165396 fs^4

Table 4.3: Phase coefficients for the individual with the highest fitness after 300 evaluations.

ens the pulse from the initial pulse length of 1022 fs down to 51 fs. The third order compensation makes the pulse more symmetric and shortens it further down to 36 fs. The impact of the fourth order term is much smaller, leading to shortening of less than 0.5 fs.

In the bottom right plot the effect of the different phase coefficients on the pulse length is plotted against the units of resolution for the current optimization (see table 4.1). Generally the treatment of the different terms should not be done independently, since the value of one may determine the impact of another. In this case it is however useful to see that around an optimum as determined by the second order term, the pulse length is much less sensitive to the higher order terms on the scale used during this optimization. This explains the increased diversity for the higher order phase terms among the individuals in figure 4.11. It also suggests that the limits for at least the fourth order phase coefficient should have been wider.

4.3.2 Spectral phase analysis using SPIDER

The previous results suggest that the optimized output pulse was not transform-limited, having a duration of 3 fs above the transform-limit. To get a closer look on the solution found by the algorithm, access to a SPIDER-apparatus as described in section 2.4 was possible for a few days, thus making it possible to fully characterize the spectral phase of the optimized pulses. As described in section 4.1.4 the SPIDER replaced the autocorrelator in the experimental setup.

Prior to this analysis the T^3-oscillator was adjusted in order to get a broader spectrum, and since the optimized second order phase in the previous runs were close to the limit for a cut-off bandwidth of 100 nm, a slight adjustment was made to the compressor just to be on the safe side. This shows up as the second order phase found in this run being different from the one found in the previous runs. Previously the optimized second order phase was approximately -11500 fs^2 and here it is around -10500 fs^2.

For this experiment the same reproductive plan as in section 4.3.1 was used and a run was made for 300 evaluations. In this run only the phase was modified and the amplitude was set to a supergaussian of width 80 nm without a hole. In table 4.3 the phase coefficients of the individual with the highest fitness after 300 evaluations are shown.

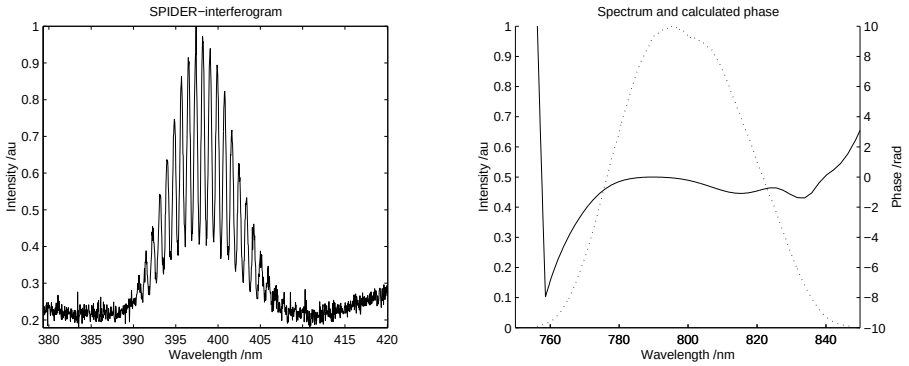

Figure 4.13: SPIDER measurement. To the left the recorded SPIDER-interferogram is shown and in the right plot the measured pulse spectrum is shown as the dashed line together with the calculated spectral phase (solid line).

The bandwidth of the optimized pulse was 46 nm and the corresponding transform-limited pulse length was calculated to be 28 fs. In figure 4.13 the measured SPIDER-interferogram is shown in the left plot and to the right the pulse spectrum is shown together with phase calculated with the Fourier-Transform Spectral Interferometry method described in Appendix B. As can be seen, there are still uncompensated phase terms remaining and the pulse length was calculated to be 35 fs, or about 1.25 times the transform-limit.

The uncompensated phase-terms indicates that the algorithm had not fully converged after 300 evaluations. A continuation of the same run up to 600 evaluations also showed that this was the case and the change of the phase coefficients obtained with the continued run was -308 fs^2, -17204 fs^3 and -189247 fs^4 for the second, third and fourth order respectively. At this point it was unfortunately not possible to do any additional SPIDER-measurements and thus the change in the phase was applied mathematically. Figure 4.14 shows the expected phase after 600 evaluations obtained by adding the polynomial phase difference to the phase measured after 300 evaluations.

As can be seen the result of the continued evaluation is to further flatten the phase and the calculated pulse length after 600 evaluations turns out to be 29 fs or only 1.03 times the transform-limit. In figure 4.15 the calculated pulse shape after 300 evaluations is shown together with the estimated pulse shape after 600 evaluations. As can be seen, the continued evaluation leads to a much better correction for the third order phase, making the pulse more symmetric.

This analysis suggests that it would be very interesting to perform a more thorough investigation of the spectral phase resulting from an optimization run. It also shows that it is very important to assure that the algorithm has fully converged before making quantitative diagnostics.

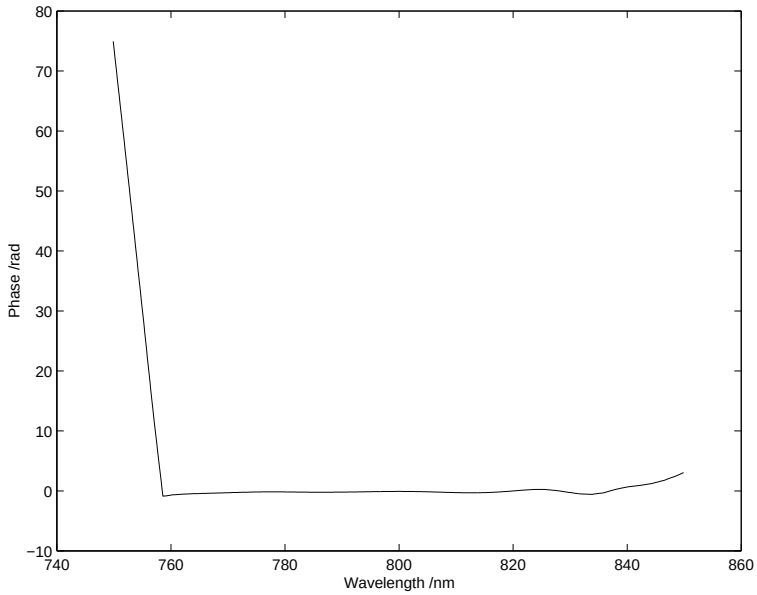

Figure 4.14: Expected phase after 600 evaluations obtained by adding the polynomial phase difference to the phase measured after 300 evaluations.

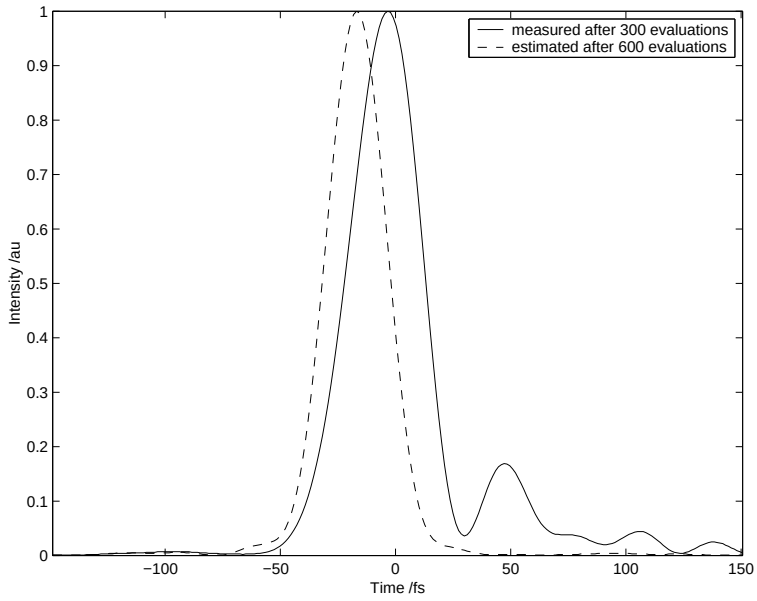

Figure 4.15: Pulse shapes. The measured pulse shape after 300 evaluations (solid line) is shown together with the estimated pulse shape after 600 evaluations (dashed line).

Chapter 5

Summary and outlook

In this study an Acousto-Optic Programmable Dispersive Filter has been used to temporally shape femtosecond pulses from the oscillator of the T^3-laser in Lund. The actual shaping has been done through a feedback loop, controlled by a genetic algorithm, in order to optimize the efficiency for second harmonic generation, thus expecting to minimize the pulse length. The experimental investigations have shown that the algorithm works, finding pulses with a pulse length of a few percent above the transform-limit. During the work experience has been gained on the pulse shaper as well as on the practical application of a genetic algorithm in a feedback loop. Below some of the conclusions reached are summarized.

5.1 Pulse shaper

An important issue when the Dazzler is used in an experimental setup is that the second order phase may have to be externally compensated so that the GDD that has to be introduced by the Dazzler lies within its dynamic range of compensation. This range is determined by the required energy throughput and the required spectral bandwidth. Since the shaped pulse is created by diffraction due to acousto-optic interaction in the crystal, the actual length of interaction between the acoustic and optical waves is crucial for the diffraction efficiency, and hence determines the energy throughput. The interaction length is dependent on the spread of the acoustic wave, which in turn is strongly dependent on the applied GDD. Thus a larger (absolute value) of the applied GDD will lead to a higher energy throughput. The spectral bandwidth limit is due to the finite length of the crystal itself. If a too large GDD is applied, the acoustic wave will spread so that it does not fit in the crystal, and is thus cut off, resulting in some acoustic frequencies not being present in the crystal. The optical frequency components that would otherwise be diffracted by the missing acoustic frequencies will hence be missing in the diffracted pulse spectrum. This was not a severe problem in the present work, but if the Dazzler is to be used in a laser-system with an amplifier, attention has to be payed to this issue.

This is due to the fact that an amplifier is sensitive to the bandwidth of the pulses. If pulses with a too narrow bandwidth are used to seed an amplifier, all gain will be concentrated into a smaller spectral range, with the risk of reaching intensities damaging the gain medium or associated optics. The need for external compensation in this work led to the addition of a compressor to the setup, while in a complete laser system some sort of compressor is always present, reducing the compensation necessary for the Dazzler to a simple adjustment.

In this work the Dazzler was operated in a system without a pulse picker. As has been discussed in section 4.1.2 it implies that not only the expected pulses are obtained, but also pulses diffracted at times when the acoustic wave is not present in full in the crystal. In the performed experiments this introduced some uncertainties to the optimization process and led to some experimental difficulties for the diagnostics. In a real laser system, like the 1 kHz-laser in Lund, this will however not cause any problems, since a pulse picker is present in this system.

In conclusion it can be said that the Dazzler showed excellent performance in shaping the laser pulses for the performed experiments. An advantage of the Dazzler over other pulse shaping devices is its capability to apply a filter to both the amplitude and the phase with a single operation. Often one thinks of the optimization of SHG-efficiency as a matter of introducing a flat phase only. In this work it was however seen that even an amplitude filter is useful to symmetrize the amplitude spectrum and thus increasing the bandwidth. The possibility to use an amplitude filter is especially important in a complete laser system with an amplifier. In amplifiers an effect known as "gain-narrowing" causes the bandwidth of the spectrum to decrease on amplification, leading to longer pulses. By suppressing the amplitude of the central frequencies prior to amplification, this effect can be minimized and thus shorter pulse lengths can be reached. Finally the Dazzler is very "user-friendly" through its control software, allowing it to be controlled in a "loop-mode" as described in section 2.4.3. It is also easy to introduce the Dazzler into an existing setup, since there is no need for a zero dispersion-line and since the shaped pulses are travelling in the same direction as the input pulses.

5.2 Optimization algorithm

On the whole, the genetic algorithm worked well, achieving pulses down to a few percent above the transform-limit. It was seen that the optimization process was even able to find an amplitude filter with the effect of making the spectral amplitude more symmetric and thus having a larger bandwidth. The brief spectral phase analysis made with SPIDER indicated that the optimized waveforms did indeed flatten the phase as required to obtain short pulses. The results did however show deviation from the transform-limited case and a possible reason for this might be the dependence of the SHG-efficiency on pulse energy. Two approaches on solving this are presented here:

- Instead of using the SHG-efficiency directly one may use the ratio between

the SHG-efficiency and the pulse energy in the fundamental. This would compensate for the dependence of the feedback signal on the pulse energy, at least under the non-depleted approximation (see section 2.2.3).

- In the Dazzler software there is a special function allowing to use a constant gain for the acoustic wave. Normally the power of the acoustic wave is set by a power value between 0 and 1, where a value of unity means that the maximum amplitude of the acoustic wave will be normalized to the maximum voltage allowed. If the constant gain function is activated, the same scaling factor that is used for the current acoustic wave will be used to scale all acoustic waves. This means that if for example a second order phase is applied, leading to a spread of the acoustic wave, the maximum amplitude of the acoustic wave decreases. In this way it will be possible to keep the diffraction efficiency almost constant. Unfortunately, the use of this function was beyond my knowledge when the present experimental work started.

After this work was finished, I got the opportunity to make a visit to Fastlite, the company that designed the Dazzler. It was during this visit that I got explained to me the effect of using the constant gain function in the Dazzler software as explained above. During this trip I also visited a laboratory at CEA (Commissariat à l'Energie Atomique), where they have a laser system similar to that in Lund. I was able to try my genetic algorithm on this system, optimizing on the SHG-efficiency of the pulse after amplification. The algorithm converged but due to technical problems it was not possible to use the SPIDER to characterize the optimized pulses. However the results were compared with the results of an optimization made at the same time using a simplex method search, showing excellent agreement for the found phase coefficients. The simplex method however showed much faster convergence, thus confirming what has been discussed in the experimental section, that the problem posed here is probably too simple to benefit from an evolutionary approach.

The implementation of the genetic algorithm used in this work was on purpose made as abstract as possible. It is for example not dependent on the use of some particular feedback, since the fitness is calculated by an external measurement routine as described in section 3.2.2. It is also easy to run the algorithm with different encodings of the waveform parameters due to the use of a function template for the decoder routine as described in Appendix C. The operators for selection, crossover and mutation have also been built on templates, making it easy to change these as well. It now remains for the genetic algorithm to be applied to a more difficult problem, where it can be more useful compared to other optimization algorithms, for example in the high harmonic generation research towards attosecond pulses, currently performed at the 1 kHz laser system in Lund.

Acknowledgements

First of all I would like to thank my supervisor *Anne L'Huillier* for giving me the opportunity to do my diploma project within the basic research group, and for her valuable help through questions, suggestions and discussions, encouraging me to find and explore the physics behind all the genetics.

I would also like to thank the rest of the basic research group for all their help and for all fruitful discussions. Especially I would like to thank *Anders Persson* and *Johan Mauritsson* for all their help with the experimental equipment. Johan also deserves an extra credit for helping me and participating in the SPIDER-analysis, making it possible to include this in the report.

The people at *Fastlite* I would like to thank for all support with the Dazzler, and for answering many of my questions when I was there for a visit. A special thanks also to *Olivier Gobert and colleagues* at the CEA, for letting me visit their facility and even daring to let me try my algorithm on their system.

At last I would like to thank my girlfriend *Maria*, and the rest of my family, for putting up with me even during the most intense periods of night-work, but also for reminding me now and then that it is sometimes useful to focus on other things than work.

Bibliography

[1] J.-C. Diels, and W. Rudolph, *Ultrashort laser pulse phenomena*, Academic Press, Inc., 1996.

[2] A. Sjögren, *Laser-matter interactions at extreme irradiance, X-ray generation and relativistic channelling*, Doctoral Thesis, Lund Reports on Atomic Physics, LRAP-288, 2002.

[3] J. Norin, J. Mauritsson, A. Johansson, M.K. Raarup, S. Buil, A. Persson, O. Dühr, M.B. Gaarde, K.J. Schafer, U. Keller, C.-G. Wahlström, and A. L'Huillier, *Time-frequency characterization of femtosecond XUV pulses*, Physical Review Letters, 88(19), 2002.

[4] R. Bartels, S. Backus, E. Zeek, L. Misoguti, G. Vdovin, I.P. Christov, M.M. Murnane, and H.C. Kapteyn, *Shaped-pulse optimization of coherent emission of high-harmonic soft X-rays*, Nature, 406:164-166,2000.

[5] P. Tournois, *Acousto-optic programmable dispersive filter for adaptive compensation of group delay time dispersion in laser systems*, Optics Communications, 140:245-249, 1997.

[6] D. Reitze, F. Weihe, S. Kazamias, G. Mullot, D. Douillet, F. Auge, J.P. Chambaret, D. Hulin, and P. Balcou, *Adaptive optimization of high order harmonic generation in a free propagation geometry*, Preprint.

[7] A.E. Siegman, *Lasers*, University Science Books, 1986.

[8] O. Svelto, *Principles of Lasers*, Plenum Press, 1998.

[9] J. Mauritsson, *Generation of ultrashort laser pulses using gas-filled hollow waveguides*, Lund Reports on Atomic Physics, LRAP-247, 1999.

[10] A. Yariv, and P. Yeh, *Optical waves in crystals*, John Wiley & Sons, Inc., 1984.

[11] O. Prakash, S.K. Dixit, and R. Bhatnagar, *On the role of the coherence width and its evolution in a short-pulse fundamental beam in second harmonic generation from Beta-Barium Borate*, IEEE Journal of Quantum Electronics, 38(6):603-613, 2002.

[12] F.-T. Wu, and W.-Z. Zhang, *Consideration of angular acceptance angle in BBO crystal on a highly efficient second harmonic generation*, Optics & Laser Technology, 30:189-192, 1998.

[13] S.-G. Pettersson, S. Borgström, and H. Hertz, *Advanced Optics*, Department of Physics, LTH, 1999.

[14] A.M. Weiner, *Femtosecond pulse shaping using spatial light modulators*, Review of Scientific Instruments, 71(5):1929-1960, 2000.

[15] D. Kaplan, and P. Tournois, *Theory and performance of the Acousto-Optic Programmable Dispersive Filter used for femtosecond laser pulse shaping*, J. Phys. IV France, 12(Pr5):69-75, 2002.

[16] *Dazzler System Operating Manual*, Fastlite, V310 July 2001.

[17] D.E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley Publishing Company, Inc., 1989.

[18] K.F. Man, K.S. Tang, and S. Kwong, *Genetic Algorithms*, Springer-Verlag London, 1999.

[19] E. Belmont-Moreno, *The role of mutation and population size in genetic algorithms applied to physics problems*, International Journal of Modern Physics C, 12(9):1345-1355, 2001.

[20] K.A. Dejong, and W.M. Spears, *An analysis of the interacting roles of population size and crossover in genetic algorithms*, Proc. First Workshop Parallel Problem Solving from Nature, Springer Verlag, 38-47, 1991.

[21] J.J. Grefenstette, *Optimization of control parameters for genetic algorithms*, IEEE Transactions on Systems, Man, and Cybernetics, SMC-16(1):122-128, 1986.

[22] R.L. Fork, O.E. Martinez, and J.P. Gordon, *Negative dispersion using pairs of prisms*, Optics Letters, 9(5):150-152, 1984.

Appendix A

Coupled-mode analysis of second harmonic generation

When analyzing second harmonic generation, it is useful to confine to the scalar case even if the crystal is anisotropic. This is possible if one instead of a tensor relation for the electric polarization uses a scalar relationsship on the same form as (2.20), but with an effective nonlinear coefficient d_{eff}

$$P^{NL} = 2\varepsilon_0 d_{eff} E^2 \tag{A.1}$$

where the scalars P^{NL} and E now refers to the field component along the appropriate axis. In this case, the wave equation for the electric field may be written as [8]

$$\frac{\partial^2 E}{\partial z^2} - \mu_0 \varepsilon \frac{\partial^2 E}{\partial t^2} = \mu_0 \frac{\partial^2 P^{NL}}{\partial t^2} \tag{A.2}$$

where z denotes a coordinate along the propagation direction. Only considering the components of frequencies ω and 2ω, the electric field and the polarization field may be expressed as

$$E(t,z) = \frac{1}{2}\left[\tilde{A}_\omega e^{i(\omega t - k_\omega z)} + \tilde{A}_{2\omega} e^{i(2\omega t - k_{2\omega} z)}\right] + c.c. \tag{A.3}$$

$$P^{NL}(t,z) = \frac{1}{2}\left[P^{NL}_\omega e^{i(\omega t - k_\omega z)} + P^{NL}_{2\omega} e^{i(2\omega t - k_{2\omega} z)}\right] + c.c. \tag{A.4}$$

Using (A.3) for the electric field in relation (A.1), and carrying out the multiplication, the two amplitudes for the polarization components in (A.4) can be identified as

$$P^{NL}_\omega = 2\varepsilon_0 d_{eff} \tilde{A}^*_\omega + \tilde{A}_{2\omega} e^{-i\Delta k z} + c.c. \tag{A.5}$$

$$P^{NL}_{2\omega} = \varepsilon_0 d_{eff} \tilde{A}^2_\omega e^{i\Delta k z} + c.c. \tag{A.6}$$

where $\Delta k = k_{2\omega} - 2k_\omega$. Since $k = n\omega/c$, it can be seen that by defining Δk in this way, the phase-matching condition (2.24) becomes $\Delta k = 0$.

The wave equation (A.2) defines a relationship between the two fields (A.3) and (A.4). Inserting the field components of frequency ω into (A.2), gives a second order differential equation for the electric field amplitude $\tilde{A}_\omega$:

$$2ik_\omega \frac{\partial \tilde{A}_\omega}{\partial z} - \frac{\partial^2 \tilde{A}_\omega}{\partial z^2} = \frac{2d_{eff}\omega^2}{c^2}\tilde{A}^*_\omega \tilde{A}_{2\omega} e^{-i\Delta kz} \tag{A.7}$$

A similar equation can be obtained for $\tilde{A}_{2\omega}$ by inserting the field components of frequency 2ω into (A.2). The two obtained differential equations couples the component amplitudes causing energy transfer to occur between the two modes during propagation through the crystal. On a wavelength scale, this energy transfer is rather slow, and one can expect that the energy of a mode does not change significantly over a propagation distance in the order of one wavelength, or that:

$$\left| 2ik_\omega \frac{\partial \tilde{A}_\omega}{\partial z} \right| \gg \left| \frac{\partial^2 \tilde{A}_\omega}{\partial z^2} \right| \tag{A.8}$$

Because of this one may neglect the second order derivatives of the field amplitudes, an approximation called the slowly varying amplitude approximation. Doing this, the second order differential equations are replaced by two coupled first order differential equations describing the evolution of the field amplitudes for the two modes.

$$\frac{\partial \tilde{A}_\omega}{\partial z} = -i\frac{d_{eff}\omega}{n_\omega c}\tilde{A}^*_\omega \tilde{A}_{2\omega} e^{-i\Delta kz} \tag{A.9}$$

$$\frac{\partial \tilde{A}_{2\omega}}{\partial z} = -i\frac{d_{eff}\omega}{n_{2\omega} c}\tilde{A}^2_\omega e^{i\Delta kz} \tag{A.10}$$

Talking about energy transfer, and the efficiency of energy transfer, it is useful to transfer the coupled equations to hold for intensity rather than field amplitude. According to Poynting's theorem [10], the intensity of an electromagnetic field is proportional to the refractive index, which suggests the definition of two new field variables, $\tilde{A}'_\omega = \sqrt{n_\omega}\tilde{A}_\omega$ and $\tilde{A}'_{2\omega} = i\sqrt{n_{2\omega}}\tilde{A}_{2\omega}$. With these new variables, the equations (A.9) and (A.10) can be written as:

$$\frac{\partial \tilde{A}'_\omega}{\partial z} = -\frac{d_{eff}\omega}{n_\omega\sqrt{n_{2\omega}} c}\tilde{A}'^*_\omega \tilde{A}'_{2\omega} e^{-i\Delta kz} \tag{A.11}$$

$$\frac{\partial \tilde{A}'_{2\omega}}{\partial z} = \frac{d_{eff}\omega}{n_\omega\sqrt{n_{2\omega}} c}\tilde{A}'^2_\omega e^{i\Delta kz} \tag{A.12}$$

Note that $|\tilde{A}'_\omega|^2$ and $|\tilde{A}'_{2\omega}|^2$ are directly proportional to the intensity of the fields, with a proportionality constant independent of frequency. By introducing a coupling constant $\kappa = \frac{d_{eff}\omega}{n_\omega\sqrt{n_{2\omega}} c}$ with units [V^{-1}], equations (A.11) and (A.12) reduces to:

$$\frac{\partial \tilde{A}'_\omega}{\partial z} = -\kappa\tilde{A}'^*_\omega \tilde{A}'_{2\omega} e^{-i\Delta kz} \tag{A.13}$$

$$\frac{\partial \tilde{A}'_{2\omega}}{\partial z} = \kappa\tilde{A}'^2_\omega e^{i\Delta kz} \tag{A.14}$$

By transposing (A.13) and multiplying both sides by $\tilde{A}'_{\omega}$, and multiplying both sides of (A.14) with $\tilde{A}'^{*}_{2\omega}$, one obtains a relationship known as the Manley-Rowe relation, which simply says that energy is conserved in the conversion between the two frequencies.

$$\frac{\partial \left|\tilde{A}'_{2\omega}\right|^2}{\partial z} = -\frac{\partial \left|\tilde{A}'_{\omega}\right|^2}{\partial z} \tag{A.15}$$

Appendix B

Frequency-domain interferometry

Frequency-domain interferometry is a widely used method with many applications. This Appendix will focus on the ability to characterize laser pulses using these techniques. As the name implies, frequency-domain interferometry is an interferometric technique, based on the addition of electric fields, where the measurements are done in the frequency-domain often using an ordinary spectrometer. Figure B.1 depicts a setup where two pulses are added using a semi-transparent mirror. The spectrum of the superposition of these two pulses is then recorded using a spectrometer. The pulse $E_R(t-\tau)$ is delayed a time τ relative to the pulse $E_C(t)$ to be characterized. To begin with we assume that $E_R(t-\tau)$ is a reference pulse with a known spectral phase $\varphi_R(\omega)$ and with the same carrier frequency ω_L as $E_C(t)$. The spectral envelopes of the two pulses may then be written as

$$\tilde{A}_C(\omega) = A_C(\omega)e^{i\varphi_C(\omega)} \quad \text{(B.1)}$$

$$\tilde{A}_R(\omega) = A_R(\omega)e^{i\varphi_R(\omega)} \quad \text{(B.2)}$$

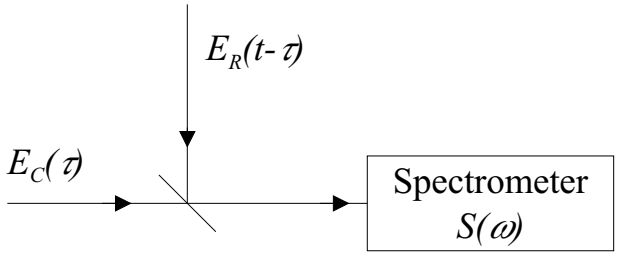

Figure B.1: In a setup for frequency-domain interferometry, two pulses are superpositioned and the resulting spectra is recorded using a spectrometer. The pulse $E_R(t-\tau)$ is delayed a time τ relative to the pulse to be measured on, $E_C(t)$.

and the time delay of the reference pulse will show up as an additional linear phase term in the spectral envelope (see section 2.1.2). The spectrometer measures the spectral intensity of the superposition of the two pulses which is proportional to

$$\begin{aligned} S(\omega) &= \left|\tilde{A}_C(\omega) + \tilde{A}_R(\omega)e^{-i\omega\tau}\right|^2 \\ &= |A_C(\omega)|^2 + |A_C(\omega)|^2 + 2A_C(\omega)A_R(\omega)\cos\left[\omega\tau + \Delta\varphi(\omega)\right] \end{aligned} \tag{B.3}$$

where $\Delta\varphi(\omega) = \varphi_C(\omega) - \varphi_R(\omega)$ is the phase difference between the two pulses at the frequency ω. The two first terms are a superposition of the individual spectra of the two pulses, while the last term is an interference term, causing fringes to appear on top of the spectrum. The average distance between the fringes is $2\pi/\tau$, depending on the time delay between the two pulses, while the difference in spectral phase between the two pulses causes the fringe distance to vary.

Assuming that the phase of the reference pulse is flat, which means that $\varphi_R(\omega)$ is constant, it can be seen from (B.3) that if for example the phase of the measured pulse is quadratic, $\varphi_C(\omega) = C\omega^2$, the fringe distance in the recorded spectrum will be $\frac{2\pi}{\tau+2C\omega}$, increasing or decreasing with frequency.

B.1 Fourier-Transform Spectral Interferometry

As described above, the phase difference $\Delta\varphi(\omega)$ can be retrieved from the interferogram by extracting the phase of the interference fringes. One way to do this is to localize the maxima of the fringe pattern and calculate the varying fringe distance. Another method which uses Fourier transforms to extract the phase goes under the name Fourier-Transform Spectral Interferometry (FTSI).

Taking the inverse Fourier transform of the measured spectral intensity in equation (B.3) results in:

$$\begin{aligned} \mathcal{F}^{-1}\{S\}(t) = \; & \textstyle\int_{-\infty}^{+\infty} A_C(t')A_C^*(t'-t)dt' + \int_{-\infty}^{+\infty} A_R(t')A_R^*(t'-t)dt' + \\ & f(t+\tau) + f^*(-t+\tau) \end{aligned} \tag{B.4}$$

The two first terms in this equation is recognized as the field autocorrelations of each pulse (compare with (2.41)) and as such they are symmetric and centered around $t = 0$. The function

$$f(t) = \int_{-\infty}^{+\infty} A_C(t')A_R^*(t'-t)dt' \tag{B.5}$$

is the cross-correlation function between the two pulses, centered around $-\tau$ for the third term and around τ for the last term. With a large enough value of the time delay τ, the correlation signal centered around τ will not overlap with the central autocorrelation signals, and can therefore be filtered out. Transforming the filtered signal back to the frequency domain yields a complex function $\tilde{S}_f(\omega)$:

$$\tilde{S}_f(\omega) = A_C(\omega)A_R(\omega)e^{-i(\omega\tau+\Delta\varphi(\omega))} \tag{B.6}$$

By using the argument of this complex function, the spectral phase difference $\Delta\varphi(\omega)$ may be calculated, since the time delay τ is known and the term $\omega\tau$ thus may be subtracted.

B.2 Spectral shearing

From equation (B.6) it is seen that it is not only important to have a reference pulse with a known spectral phase, but also that the bandwidth of the reference pulse must be larger or equal to the bandwidth of the measured pulse in order to be able to reconstruct the phase over the whole spectra. When working with ultrashort laser pulses this is often not possible since a wider spectra corresponds to a shorter pulse, and the shortest pulse available is almost always the pulse to be characterized. To be able to characterize such short pulses, it is then necessary to use the pulse itself as a reference, an approach that is made possible using a technique called spectral shearing. In spectral shearing, the spectrum of for example the reference pulse is displaced with a displacement Ω, making it possible to write the spectral envelope of the reference pulse as

$$\tilde{A}_R(\omega) \quad = \quad A_C(\omega+\Omega)e^{i\varphi_C(\omega+\Omega)} \tag{B.7}$$

giving for the filtered spectral signal

$$\tilde{S}_f(\omega) = A_C(\omega)A_C(\omega+\Omega)e^{-i(\omega\tau+\Delta\varphi_\Omega(\omega))} \tag{B.8}$$

where $\Delta\varphi_\Omega(\omega) = \varphi_C(\omega) - \varphi_C(\omega+\Omega)$. From the argument of (B.8) it is thus possible to obtain the phase difference between the component with frequency ω and the component with frequency $\omega+\Omega$. By setting the phase to zero for some frequency (often the central frequency), the phase can then be calculated for the whole spectrum with a resolution in frequency of Ω.

Appendix C

LabVIEW building blocks for the genetic algorithm

C.1 Templates

A template in LabVIEW is simply a specification of a subroutine by specifying its inputs and outputs. In LabVIEW a subroutine call may be made dynamically from a reference of a subroutine template. This means that the actual subroutine called is not determined until it is actually called, and that the user may specify the wanted subroutine at run-time by specifying its filename. Defining templates for the functions and operators used by the genetic algorithm allows the user to always use the same algorithm engine, choosing the wanted operators at each run. The templates of the standard operators select, crossover and mutate are shown in figure C.1.

The selection operator is defined as taking an array of real fitness values together with the pre-calculated fitness statistics. The result of the selection should the be given as the index of the chosen individual in the fitness array.

The crossover operator takes two individuals together with the crossover parameters. It returns the two offspring individuals together with the number of mutations that has occured. It also returns a boolean indicating if crossover took place or not.

The mutation operator is the simplest one, taking a gene and the probability for mutation as input, and returning the possibly mutated gene together with a boolean indicating if the gene was mutated.

In addition to the three standard operators, there is a need for two function templates: decode and scale. These are shown in figure C.2.

The decoder function takes an individual as input, decodes it, and outputs the wave parameters for the corresponding waveform. It also returns the number

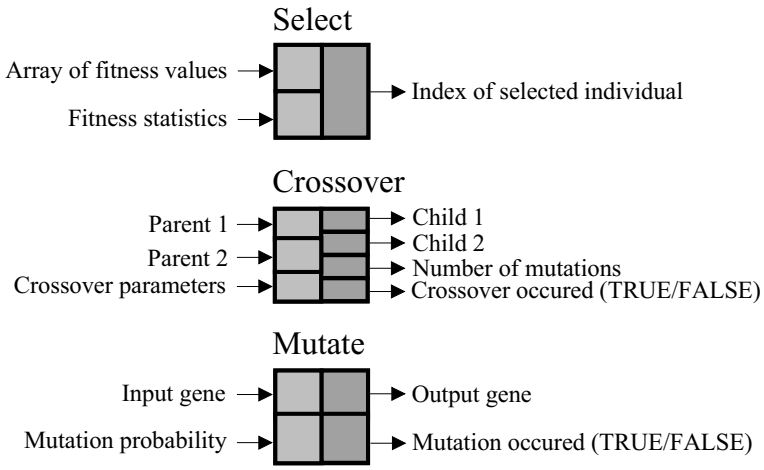

Figure C.1: Templates for the three standard operators select, crossover and mutate.

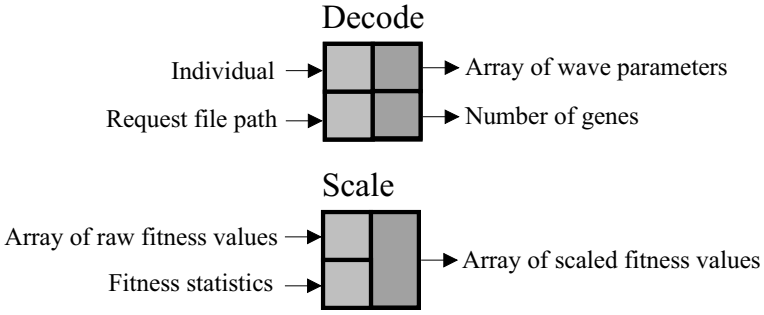

Figure C.2: Templates for the two functions decode and scale.

of genes it is designed to handle. To support arbitrary filters (see section 2.4.3) the decoder also takes the path of the request file in order to be able to generate arbitrary amplitude or phase files and put them in the correct location.

The scaling function is used prior to the selection operator, making it possible to implement fitness scaling as described in section 3.1.4. It takes as inputs an array of raw fitness values together with the pre-calculated fitness statistics, and outputs the corresponding array of scaled fitness values.

C.2 Data types

A data type in LabView is simply a front panel without any block diagram, representing compound data-carrying objects. The advantage of defining different data types is that if one for example has defined a datatype representing the whole population, whole populations may be passed as a single parameter in the data-flow of the block diagram.

The basic datatype for the genetic algorithm is the individual, consisting of a chromosome represented as an array of booleans together with the fitness value of the individual.

The different parameters needed for a genetic algorithm run are grouped so that related parameters are grouped together. The four parameter groups are:

1. *Fitness parameters:* Parameters needed for calculation of the fitness of an individual.
    - *Decode reference:* Used to obtain the wave parameters described by an individual.
    - *Request file path:* Needed to be able to make a request for a waveform.
    - *Fitness file path:* Needed to obtain the fitness from the measurement routine.
2. *Reproduction parameters:* Parameters needed to evolve the population to obtain the next generation.
    - *Scale reference:* Needed to scale the raw fitness prior to selection.
    - *Select reference:* Used to select individuals for reproduction.
    - *Crossover reference:* Used to combine two individuals before they are copied to the next generation.
3. *Crossover parameters:* Parameters passed to the crossover operator.
    - *Crossover probability:* The probability that a crossover will occur.
    - *Mutation probability:* The probability of a single gene being mutated.
    - *Mutate reference:* Used to mutate the genes during the copy operation.

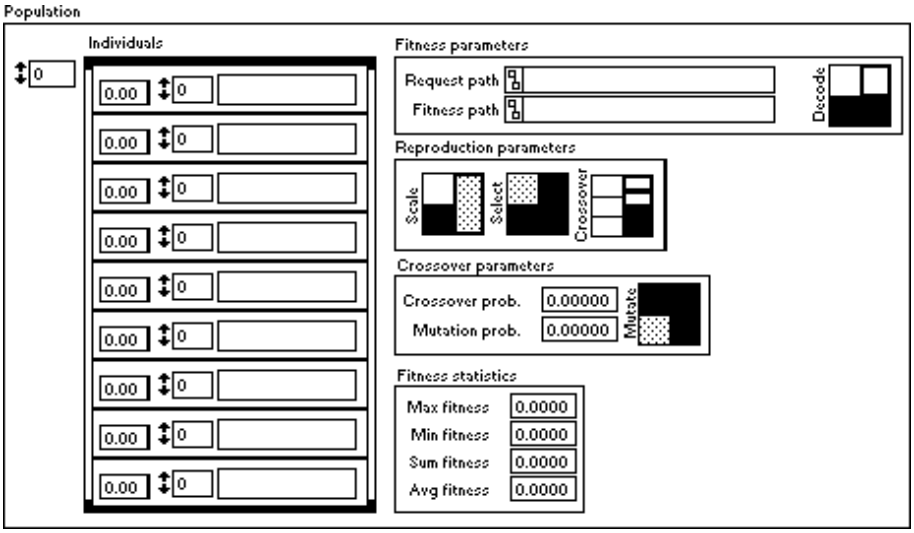

Figure C.3: The front panel object of the population data type.

4. *Fitness statistics:* Some statistics over the fitnesses in a population. Used by the selection operator and the scaling function during reproduction.
   - *Maximum fitness*
   - *Minimum fitness*
   - *Average fitness*
   - *Sum of all fitnesses*

All of these datatypes are finally combined in a compound data type representing a population. Figure C.3 shows the actual front panel object of the population datatype.