a 2D Spatial Light Modulator for spatio-temporal shaping

Daan Sprünken

a 2D Spatial Light Modulator for spatio-temporal shaping

A Master's thesis in which a two dimensional spatial light modulator is implemented, characterised, and calibrated. Spatio-temporal pulse shaping is demonstrated through a novel in situ FROG setup that contains no moving parts. Phase parameterisations used in evolutionary algorithms are studied and improved

For the title of Master of Science

Daan P. Sprünken

August 29, 2008

Applied Physics - Optical Sciences

Faculty of Science and Technology
University of Twente
Enschede
The Netherlands

Graduation committee

Prof. Dr. Jennifer L. Herek
Dr. Allard P. Mosk
Dr. ir. Herman L. Offerhaus
Dr. Aliakbar Jafar Pour

University of Twente
Enschede - The Netherlands

Contents

1 Spatio-temporal phase shaping

Light is our primary means of observing the universe. Sunlight scatters off the objects around us, showing us the world we live in. At night, on the open sea, one is treated to the most beautiful celestial views, with light coming from stars at unintelligible distances away from us. Now we want to use light as a tool, not in its raw form, but harnessed and sculptured, to see things previously beyond our reach. We lit up the dark when we learnt to harness fire, and nowadays we can control light in more fundamental ways, to illuminate the universe on a molecular scale. In this thesis, we will show how light can be manipulated and shaped into the extreme forms required for this, and we will demonstrate some of the applications.

Shaping light is done with a Spatial Light Modulator (SLM). For this project a novel two dimensional SLM was bought and implemented. Chapter 2 discusses the different types of two dimensional Spatial Light Modulators available and the considerations made in choosing the newly bought SLM. The implementation and calibration of the SLM is described in Chapters 3 and 4, in the spatial and spectral domain, respectively. Several new calibration techniques are introduced, including a pixel-by-pixel characterisation of the SLM. In Chapter 5 we present a unique application of a two dimensional SLM with an in situ FROG setup that contains no moving parts. Finally, Chapter 6 focusses on the algorithms underlying the optimisation experiments in which SLMs are often used. New insights in the coupling between the SLMs and the driving algorithms are presented, including a novel parameterisation that improves the optimisations.

This chapter will continue with an introduction in the world of shaping light. We will start with the spectral domain, in which we manipulate different colours separately. Then the spatial domain is discussed, and we end with some potential applications of the combination of spatial and spectral shaping of light.

1.1 Spectral shaping

Molecular processes occur on a time range of femtoseconds to nanoseconds. These processes include such things as the interaction between molecules in chemical reactions, the folding of proteins into different configurations, and the excitation and subsequent fluorescent emission of dyes. In order to influence these systems we must interact with them at these time scales. Pulses of light only a few tens of femtoseconds long, allow us to manipulate and steer molecular interactions. However, even the most simple molecules exhibit intricate behaviour that places extreme demands on the nature of the light with which it interacts. In order to effect a desired response, light must be tailored very specifically to manipulate the interaction of light and matter.

It is the field of coherent control that busies itself with controlling physical systems on a quantum mechanical level through interaction with light. By shaping the light so that its temporal signature matches the development of the system under scrutiny, it is possible to steer processes and gain insight in their nature.[1]

For one, we can use this control that we have over molecules to effect a certain result. As an example of this we can look at photodynamic therapy,[2] wherein cancer cells are selectively treated by a photosensitizer: a drug that only becomes active through excitation with light. By specifically tailoring the excitation light to the dynamics of this photosensitizer, we can increase its efficiency and reduce unwanted side effects. For this, precise control over the temporal profile of the excitation light is a necessity.

Secondly, we can use the interaction of light with molecules to gain further insight into the workings of the molecules. While these systems are too complex to understand a priori, we can get a deeper understanding of the mechanisms involved in the interactions when we can

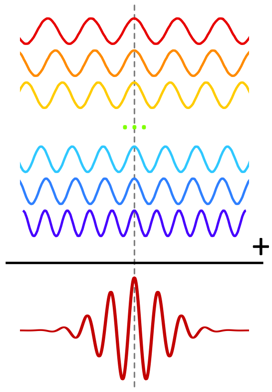

Figure 1.1: The addition of sinusoidal waves with varying frequencies. If a broad spectrum of frequencies is combined, an ultrashort pulse is the result.

excite the molecules with specific and well defined light pulses.[3]

Evidently, we need precise control over the temporal characteristics of our light pulses in order to manipulate systems on a femtosecond or picosecond timescale. This is known as spectral (or temporal) shaping of light. For an extensive treatise on spectral shaping the reader is referred to the review article of Weiner.[4] Here, only an attempt is made to give the reader an intuitive understanding of spectral shaping.

1.1.1 Femtosecond light pulses

Forgoing a debate on the wave-particle duality of light, let us picture light as a wave with a certain amplitude and frequency. The amplitude gives the intensity of the light while the frequency defines its colour. A single colour would then give us one continuous sinusoidal wave, but what happens if we combine a broad range of colours, as in Fig. 1.1?

Where the waves of all the colours simultaneously reach their peak value, we find a peak in our intensity. However, when we look at our beam of light slightly before or after this peak intensity, we see the intensity drop quickly. As the frequencies of all colours differ slightly, the total contribution of the different colours does not amount to much at several tens of femtoseconds from the peak intensity. Some are at the peak of a vibration, while others are in a trough. Summed up, we see that there is only a short pulse in time wherein all the colours interfere constructively. In fact, the more colours we add, the shorter this pulse becomes.

If we thus want to make a pulse that is as short as possible — and remember that we need a very short pulse to interact with the fast vibrations of the molecules — then we must have a broad spectrum of colours. The converse is also true: if we have a pulse that is very short in

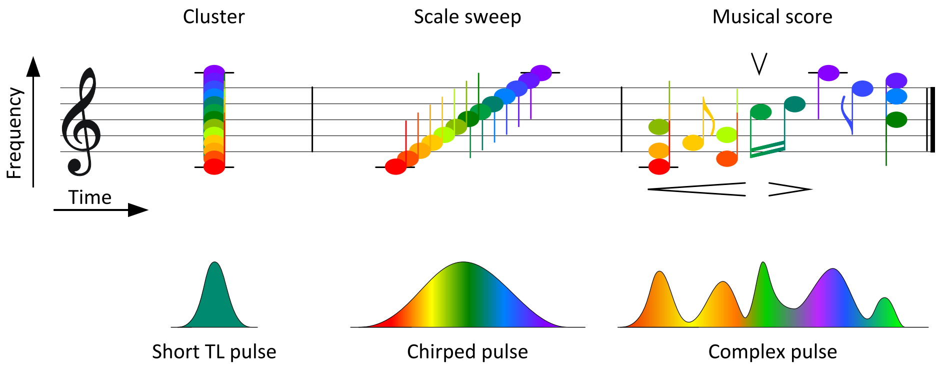

Figure 1.2: Analogy between the acoustic and optical domain. The staff describes the time and frequency characteristics of the music. Similarly, we can describe a pulse of light in time and frequency. Figure adapted from Ref. 3.

time, then it will consist of many different colours.

1.1.2 A melody of light

We now know that ultrashort pulses have a broad spectrum (i.e. it consists of many different colours). So how does this permit us to shape this pulse in time? This question is most easily answered with a musical analogue to light.

While light is an electromagnetic vibration where different frequencies represent different colours, sound is a vibration of air where each pitch has a unique frequency. The representation of a sinusoidal wave with an amplitude and frequency holds in both cases.

In music, each frequency is represented by a note on a musical score. Fig. 1.2 shows such a musical score, with the pitch varying as we move up on the staff and time progressing as we read the notes from left to right.

The femtosecond pulse that was described before, can be represented in the musical analogue as hitting all the keys on a piano at once. This will produce a short burst of sound with a lot of energy, but very little melody. If we want to make music, then we have to shift the notes with respect to each other.

The simplest "music" that we might play, would be to slide a hand across the piano, striking all the keys one after another. On a musical score, this is represented by placing the notes one after another, going up in frequency as we move from left to right across the keys. With our ears, we will also hear the pitch go up as we move up the musical scale. Crucial, is the delay between the notes which shifts the low and high pitch in time.

A real music composition is more complicated. A composer creates an intricate pattern of notes in which he plays with the length of the notes, their combinations, their amplitudes, and most importantly, the delays between them. If this is done right, then Beethoven's Fifth symphony or the national anthem of the Republic of Palau will emerge.

With light we can do the same. As said, key is the delay between the different notes or, better said, between the different frequencies. If we look back to Fig. 1.1, then we see that a specific point in time was chosen where all colours are simultaneously at the peak of their vibration. The phase (delay) between the different colours is zero and we obtain a single burst of light, just like when hitting all the keys of a piano at once.

Playing with the relative phases of the different colours allows us to create not a musical melody, but a "melody" of light. A sweep across the musical scale is a chirped pulse in the optical domain: the red light precedes the blue, just as the low pitch precedes the higher notes or vice versa. More complicated phase differences can give rise to higher order chirp, pulse trains, or light pulses that are specifically tailored to specific molecules.

Let us consider the musical analogue once more. Occasionaly, a particularly skilful musician may compose an exceptionally good song. In fact, it is even possible to alter someone's mood and to play on someone's emotions with the right music. Maybe you go out dancing and get in a party mood with the latest disco hit or you wind down after a stressful day with your favourite mellow pop song. It requires a very specific composition of notes to elicit an emotional response. We can say that this music is tailored to a specific "vibration" and provides control over your emotional processes.

With light we now aim to do the same thing. Not with emotions, but rather with molecules. While we may not be able to make a molecule cry or giddy with joy, we can influence a chemical reaction, break an intramolecular bond, or selectively excite a certain molecular marker. It is towards this goal that we employ spectral phase shaping.

1.1.3 Experimental setup

One question is left unanswered: how do we manipulate the phase of light? For this, we have a spatial light modulator (SLM). We want to apply a delay in time (corresponding to a shift in phase) to a ray of light. Two ways exist of doing this, but before we look into this we must first separate the spectral components of the light.

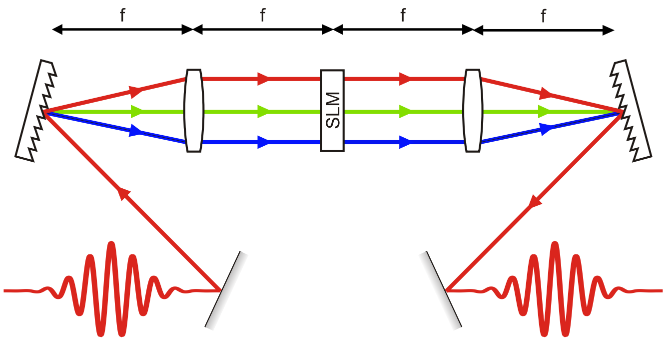

Figure 1.3: A 4f zero dispersion pulse compressor. The incident pulse is dispersed on a grating, separating the colours. Each colour is focussed onto an SLM pixel by the positive lens. In a transmissive setup as displayed here, the right half exactly mirrors the left half, thereby recombining the pulse after the second grating.

4f zero dispersion pulse compressor

Figure 1.3 shows a 4f zero dispersion pulse compressor.[4] The operation of this setup is three-fold.

The incident femtosecond pulse is first divided into its separate frequency components with a grating. The workings of a grating are slightly different, but the end result is similar to that of a prism: every colour component is reflected off the grating at a slightly different angle, thus converting a spectral difference (frequency) into a spatial difference (angle).

The second optical element is a positive lens which has two functions. One is to collimate the beam coming off the grating so that it can later be recombined into one pulse. More importantly, it focusses each colour onto a separate pixel on the SLM. This will allow us to then manipulate each colour individually. Precisely how this manipulation is achieved we will find out in the next section.

The third aspect of this setup is that it recombines all the colours back into one pulse. This is achieved by the right half of Fig. 1.3, which is the precise mirror image of the left half. Shown in Fig. 1.3 is a so-called transmissive setup wherein the light is manipulated as it is transmitted through the SLM. Alternatively, a reflective setup can be used, using an SLM that reflects the light. In that case, the light is automatically recombined as it retraces its path through the focussing lens and past the grating.

If we remove the SLM from the setup, then we in no way alter the light as it passes through the setup and we obtain the exact same pulse at the output as we put in. Also of note is the distance between the various optical elements, which should be precisely the focal length f of the lens. Hence the name 4f zero dispersion pulse compressor.

So how does the SLM work? We have separated the spectrum and each colour is incident on a different part of the SLM. With what magic is modulation of phase achieved?

Liquid crystal displays

A liquid crystal display (LCD) — as used in flatscreen computer monitors, telephone displays, and digital alarm clocks — is a very commonplace example of a spatial light modulator. Light from a backlight is first polarized and then passed through a liquid crystal (LC) which rotates the polarization. The amount of rotation is regulated by applying a voltage over the LC. A second polarizer at the front of the display then makes the intensity dependent on the driving voltage.

For our application we are looking for phase modulation and not amplitude modulation. Pure phase modulation can also be obtained using LCs, though some customization is required. In a normal LCD the liquid crystal is a so called twisted nematic LC which *twists* the polarization. An alternative — and more expensive — form of nematic LCs is the parallel aligned nematic variety. Here the polarization is left intact and applying a voltage over the LC will change the refractive index for one polarization direction. By slowing down the light, we can thus apply a phase shift as a function of a driving voltage. This gives us a phase-only spatial light modulator.

Liquid crystal based SLMs can be used both in transmissive and reflective setups. The electrodes in LCDs are transparent, thus allowing light to pass through. A reflective coating can be added if a reflective setup is desired. Liquid crystal on silicon (LCoS) devices are only usable in reflective setups as they have an opaque silicon backplane.

Deformable mirrors

An alternative method of phase modulation makes use of deformable mirrors. In a reflective 4f setup, we can also induce phase delay by physically changing the path length of the various colours. In a deformable mirror, the surface of the mirror is deformed, thereby increasing or decreasing the distance traversed by different parts of the spectrum.

1.2 Spatial shaping

Apart from manipulating light in time, we can also shape it in space using a spatial light modulator. While we are concerned with the phase delay as a function of frequency in spectral shaping, it is easiest to consider a monochromatic beam when discussing spatial shaping.

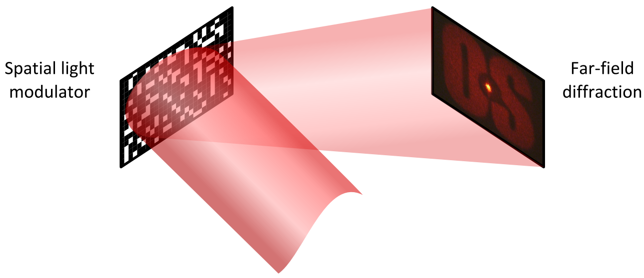

Figure 1.4: With spatial shaping, the SLM shapes the phase of the wavefront. Far from the SLM (in the far-field) this transforms into an amplitude modulation. With a non-intuitive phase mask, complicated figures can be made, such as the Optical Sciences logo. So, what is the point?

Spatial shaping is nothing other than the manipulation of a beam profile. A glass lens is a simple example of a two dimensional “phase mask” used for spatial shaping: different parts of a beam incident on a lens will travel through different lengths of glass, thereby focussing the beam onto the focal plane.

A more complicated spatial shaping process occurs on a grating. Diffraction off the regular pattern of the grating causes a collimated incoming beam to be sent off into multiple diffraction orders. In general, this is the process that applies to all forms of spatial shaping. The familiar terms of Fresnel and Fraunhofer diffraction apply.

An example is given in Fig. 1.4. On the SLM, a non-intuitive phase mask is displayed, which manipulates the wavefront of the incident light. The intensity is left unchanged. Far from the SLM — in the aptly named far-field — we obtain a complicated diffraction pattern. Given the proper phase mask, we can even create highly complex images.

In short, we can create complicated scattering patterns on a two dimensional SLM to spatially modulate a beam. An incoming monochromatic beam can be transformed into a wildly scattered, inhomogeneous beam profile by applying a complicated phase mask. The inverse is then also true: highly complex spatial distortions can be compensated using an SLM. This is the concept by which adaptive optics in astronomical telescopes can compensate for atmospheric distortions. It is even possible to compensate the incalculable distortions occurring in highly turbid media such as bone, teeth, or egg shells.[5]

1.3 Applications of spatio-temporal phase shaping

A few recent examples of spatio-temporal shaping — the combination of shaping in the spectral and spatial domains — will be discussed here, as well as some yet untested ideas.

The use of two dimensional (liquid crystal) SLMs for spatio-temporal shaping is spearheaded by the research groups of Nelson and Silberberg. They have recently presented some proof-of-concept experiments wherein one dimension of the SLM is used for spectral shaping of femtosecond pulses, while the second dimension is given various different uses.

Feurer and co-workers [6] presented proof-of-concept spatio-temporal shaping already in 2002. Across the spatial dimension of their SLM they created different pulse trains in time, resulting in an image appearing in an unusual space-time plot. The limitation of the spatial shaping to one dimension appears to severely limit applications of this technique to the creation of funny images, as fervently shown by Vaughan. [7]

A more serious application of a two dimensional SLM was presented by Gundogdu and co-workers [8] in 2007, in the dubiously named paper “Multidimensional coherent spectroscopy made easy”. In this work, they physically divide one incident beam into four beams on the SLM, spreading the spectrum in one dimension with a grating, and separating the four beams along the spatial dimension. The four separately shaped beams are then combined in a BOXCARS setup to provide two dimensional spectra.

Frumker and Silberberg [9] used an identical device as used in the work of this thesis to shape both amplitude and phase. By applying a grating in the spatial dimension of the SLM, they can selectively attenuate their spectrum. The spectrum diffracts part of the incoming light into higher orders which is not coupled back at the output of their shaper setup. In this way, they can use a phase-only SLM for both phase and amplitude modulation. In related work, [10] they added a scanning mirror that walked the beam across the spatial dimension of the SLM. By applying different phase masks along the spatial dimension, they could greatly increase the perceived refresh rate of their SLM. They report refresh rates of about 140 kHz, three orders of magnitude greater than achievable with regular liquid crystal SLMs.

The main advantages named in the aforementioned applications all stem from the fact that a single incident beam can be used for the same purposes where several separate laser beams were originally required. By subdividing one beam on the SLM, one obtains a setup that is easy to align, has an inherent phase relation between the various beams, and allows for arbitrary waveform generation in all beams. Another possible application of two dimensional SLMs that offers these same advantages is presented in this work in Chapter 5. There, we demonstrate a single beam FROG setup with no moving parts which allows for quick and easy characterisation of complicated waveforms at the sample position.

All applications of spatio-temporal shaping with a single two dimensional SLM are restricted to one spatial dimension. This is sufficient for simple setups to separate beams, but

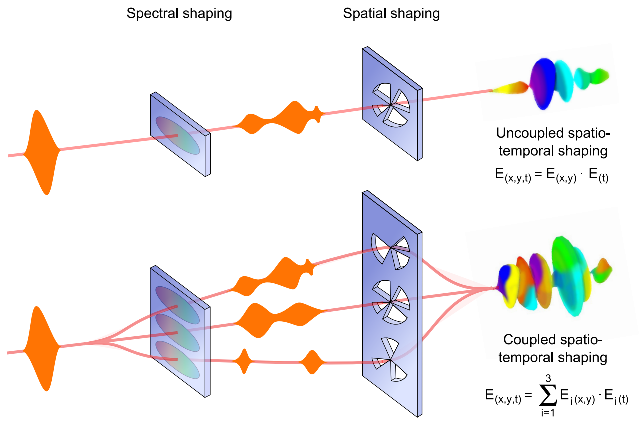

Figure 1.5: Spatio-temporal pulse shaping can be either (a) uncoupled or (b) coupled.

is generally insufficient if the spatial characteristics of the system being studied are relevant. However, systems that operate in only one spatial dimension do exist, and some (limited) applications of spatio-temporal shaping may exist here. One might consider, for example, the excitation and propagation of surface plasmons, or the propagation of light in photonic waveguides. Whether spatio-temporal shaping has any real applications in the world of surface plasmons remains to be seen, as — to the knowledge of the author — no such research has been performed.

1.3.1 Smart CARS

Real spatio-temporal shaping where the light is shaped in two spatial dimensions to accommodate the spatial characteristics of the sample is more easily done using separate shapers for both the spatial and spectral domains. An example of such a system[11], would be the combination of chemically selective imaging through shaped CARS[12] (coherent anti-Stokes Raman spectroscopy) with resolution enhancement and compensation for wavefront distortions with spatial shaping.[5]

By tailoring the excitation pulses in CARS to the vibrations of specific molecules, or possibly even the specific folding-configurations of proteins, excellent chemical selectivity in microscopy can be obtained.[12] This may in the future allow for imaging cell processes in vivo.

Spatial shaping has been proven to allow for compensation of severe scattering in turbid media and can even provide focussing within these media.[5] This should then also be possible for less homogeneous samples such as skin tissue.

Combining the spatial and spectral shaping may in the future allow for chemical selective imaging of biologically relevant processes within cell tissue or, in other words, deep within human tissue. Apart from imaging, very local excitation of photosensitizers (previously mentioned in Section 1.1) may lead to unintrusive treatment of skin ailments.

The precise implementation of spatio-temporal shaping using multiple shapers will require significant study. As shown in Fig. 1.5, the straightforward combination of spectral and spatial shaping will lead to uncoupled spatio-temporal shaping wherein the spatial and spectral shaping are always separable. Preferably, the pulse can be shaped not just both in time and space, but uniquely shaped in time at every point in space, and shaped in space at every unique point in time. This will require coupling of the spatial and spectral shaping, for example by combining multiple beams.

2 2D Spatial Light Modulators

Several design parameters were defined for the selection of a two dimensional spatial light modulator. First, we had the need for a very large bandwidth, preferably from 400 up to 900nm. Secondly, only phase-only modulators were considered with as little parasitic amplitude modulation as possible. Finally, a minimum of a few hundred uniquely addressable pixels in each dimension were required to assure sufficient resolution for the creation of more complicated pulse shapes.

This chapter lists the main manufacturers of two dimensional liquid crystal based SLMs, as well as deformable mirrors. The given information is accurate per September 2007. Considering the rapid advances in SLM technology — especially in the field of deformable mirrors — the use for the reader of this information in selecting a new SLM may be limited.

2.1 Liquid crystal displays

Liquid crystal based two dimensional SLMs are produced by only a few companies. The devices that we looked at are from Holoeye, Hamamatsu, and Boulder Nonlinear Systems. The different products will be discussed here separately. The information given in this chapter comes from both the websites of the manufacturers, as well as from telephone and e-mail conversations with their engineers and sales people. Their contact details are given in Appendix A.

2.1.1 Hamamatsu

In 2007, Hamamatsu introduced a new two dimensional SLM labelled the Programmable Phase Modulator (PPM). The PPM consists of a parallel aligned nematic liquid crystal (PAL-SLM) which is optically addressed with an LCD. The LCD at the back obviously has pixels, but the LC at the front of the device which applies the phase modulation, has no pixel structure. This means that there is less diffraction and you get a continuous shaping of your light. Whereas a normal LCD is driven electrically, the PAL-SLM changes its optical properties when a write-light is shone on it from the back. The pixelated light coming from the LCD is therefore slightly defocused and shone on the PAL-SLM where it is converted into a phase modulation. It should be noted that even though there are no pixels on the LC, there is a limit on the resolution of the device. The LC cannot be modulated over shorter ranges than 10–20 μm, regardless of the write-light setup used. Moreover, Hamamatsu does not look favourable upon replacing the original write light setup with a custom built one.

The PPM has a resolution of 1024 by 768 pixels and has standard bandwidths of 100 nm in a range from 350 to 1600 nm for which a 2π phase shift is guaranteed. Custom devices can be made with a maximum bandwidth of up to 350 nm. The restriction on the bandwidth is caused by the need to have both a good anti-reflection coating at the front as well as a mirror at the back of the device. The delivery time for a custom device is around 4 months for a price exceeding €22.000,=. The Hamamatsu device is the only 2D LC-SLM with a 100% fill factor.

2.1.2 Boulder Nonlinear Systems

Boulder Nonlinear Systems (BNS) offers the XY Phase Series spatial light modulators in 5 varieties. Given our pulses with a central wavelength of 800 nm, only the model P512-0785 is relevant. This device has a wavelength range of 760 to 865 nm. This range can potentially be extended slightly, but at a loss of phase-resolution for the extended part of the spectrum. Unique for the BNS device is the possibility of adding a dielectric mirror coating on top of the electrical contacts at the back of the LC. This extends the fill factor to 100% and thus mitigates the negative side-effects of a pixelated device. However, this extra coating does reduce the wavelength range slightly. The resolution of the XY Phase Series SLMs is 512 by 512 pixels, with

50 addressable levels per pixel for a 2π phase shift. The refresh rate is around 18 Hz, which is three times less than the other available phase shapers which operate at 60 Hz. Delivery times of this SLM are up to 6 weeks, though reduced delivery times of 2 weeks are occasionally possible. The estimated cost of a device was €20.000,=.

2.1.3 Holoeye

Holoeye has a large range of spatial light modulators, but only in the second half of 2007 did they introduce a phase only modulator. This is the HEO1080p spatial light modulator with an amazing resolution of 1920 by 1080 pixels with 8 bit addressing. This is only slightly smaller than the largest SLM on the market which has 1920 by 1200 pixels and is also produced by Holoeye. This high resolution does come at a cost: the fill factor is only 87%. While this is good for pixelated devices, it will lead to significant amplitude loss which will not occur with the devices of Hamamatsu and BNS. The bandwidth of the HEO1080p is officially given as 420–810 nm, but the device should also still be useful at slightly higher wavelengths. Other groups have already used this SLM with 30 fs pulses centred at 800 nm. A 2π phase shift is guaranteed across the whole advertised wavelength range. The Holoeye device has a delivery time of 4 weeks at a price of €15.000,=.

2.2 Deformable mirrors

In the field of astronomy, deformable mirrors are the norm for wavefront correction. The surface of a mirror is actively shape, thereby changing the path length of reflecting light. This translates into a phase shift and thus these devices can be used for all forms of spatial and spectral phase shaping. Both pixelated devices, as well as mirrors with a continuous flexible membrane are available.

The mechanical nature of these shapers limit the number of actuators to a few hundred or thousand, although rapid advances in microelectromechanical systems (MEMS) have led to the recent introduction of a 200 by 240 pixel micro mirror device by the Fraunhofer Institute for Photonic Microsystems. This same company also sells a 1 megapixel device for use in industrial lithography, but this SLM is only usable in the UV. The pixel stroke — the maximum displacement of each pixel — is also limited to 180 (nm), which would equate to less than half a vibration for 800 nm light.

2.3 Our choice

Given the wide bandwidth range required, only the Hamamatsu and Holoeye SLMs were considered as serious candidates. While the Hamamatsu shaper had the advantage of being unpixelated, the excessive delivery time coupled with the high price made us opt for the Holoeye HEO1080P instead. The HDTV resolution (1920 by 1080 pixels) of this SLM was an added benefit.

The main advantage of deformable mirrors is their very high refresh rate of several kilohertz, compared to the typical 50–60 Hz of LC based devices. When the pixel count of deformable mirrors is increased through further development in MEMS fabrication, then these devices will be a serious alternative to the high resolution LC SLMs currently on the market. At the start of this project, however, the available two dimensional deformable mirrors were considered too low resolution to be useful for spatio-temporal shaping purposes.

3 Spatial shaping

The easiest setup imaginable with a two dimensional SLM is to aim a monochromatic laser at it, and observe the far field diffraction when different phase masks are applied. Using simple variations on this basic recipe, the Holoeye HEO1080P SLM is calibrated and characterised. We will determine the SLM input vs. phase delay relationship, as well as the phase delay as a function of position on the shaper. In the process, a novel pixel-by-pixel characterisation technique is introduced, and several limitations of the SLM are detailed.

3.1 Implementation

The Holoeye HEO1080P Phase Only Modulator is connected to a computer through a DVI cable and is set up as a second monitor. The Windows desktop can then be extended to the second monitor and anything that is displayed on the extended desktop will be displayed on the SLM.

In practice, displaying phase masks is done either by manually selecting a 2944x1080 (monitor size + SLM size) pixel wallpaper (set to *tile*) in Windows, or automatically through LAB-

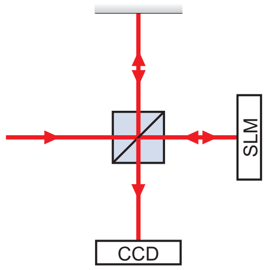

Figure 3.1: Michelson interferometer with an SLM in one of the arms.

VIEW. In LABVIEW, the image is put on the second screen by displaying either a matrix or image on a borderless, full-screen window that is appropriately offset. While the device only uses the green colour channel, it is often easier to use a greyscale image which of course works equally well. In the rest of this thesis, the term *SLM input* will be used to indicate the values set on the computer and sent to the driving electronics of the SLM. The SLM has 8 bit addressing, meaning that each pixel can be given a green or gray colour level from 0 to 255.

The driver software of the device can be adjusted using the software supplied by Holoeye Corp. The device can be connected with a serial cable, after which new calibration data can be loaded. The procedures for this are well explained in the manuals on the cd-rom that is delivered with the SLM.

3.2 Calibration methods

The SLM converts a driving electric signal into an optical phase shift for every one of the 2 million pixels in the LCD. The relation between the SLM input (256 levels as set on the computer) and the output optical phase shift (up to 2π, or more) has to be determined. Several different ways of calibrating the SLM input vs. phase relationship will be discussed here.

Frumker and Silberberg[9] make use of a Michelson interferometer (Fig. 3.1). The SLM area is divided into two parts: one half is left unchanged while the other half is scanned through the phase range. By analysing the resulting interference pattern with a camera, the input/phase relation can be retrieved.

A similar method is suggested by Holoeye in the documentation of the SLM. Using an opaque plate with two slits, they split up the beam into two distinct beams as in Fig. 3.2. These two beams are reflected off the SLM and focussed onto one spot where an interference pattern can be observed with a camera. Since the two beams are incident on different halves of the SLM, the phase difference between the two beams can easily be controlled. If a varying

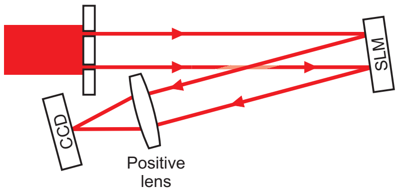

Figure 3.2: Interference of two parallel beams. The lens overlaps the two beams on the CCD.

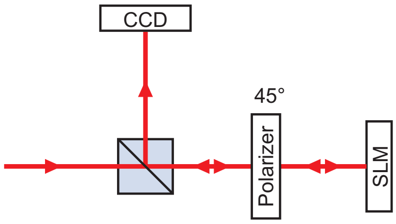

Figure 3.3: Calibration setup using a polarizer to convert phase modulation into amplitude modulation.

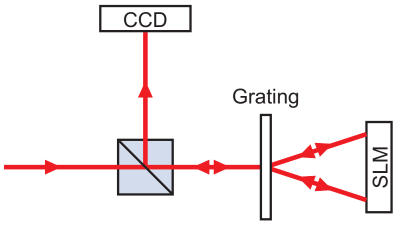

Figure 3.4: Calibration setup using a transmissive grating to split the beam. On the CCD interference fringes can be seen.

phase delay is applied to one of the beams, the interference fringes will shift. This shift is a good measure of the phase delay and thus directly related to the driving voltage. The main advantage of this method over the setup of Frumker and Silberberg, is the simple alignment. Furthermore, Frumker and Silberberg use a beam splitter which can give unwanted reflections and additional interference effects that complicate measuring the movement of the fringes.

Vaughan[7] gives two alternative methods. In the first method, a polarizer is placed before the SLM at 45° (Fig. 3.3). The LCD is birefringent and a phase change will only be applied to one polarisation; the polarisation will thus change as a function of the input. By passing the modulated beam back through the polarizer, the amplitude is modulated. The amplitude change can be observed with a camera, and the desired relation can be retrieved.

The second method used by Vaughan makes use of a (transmissive) grating (Fig. 3.4). A monochromatic beam is split into two beams using a grating in such a way that the two beams strike the SLM at distinctly different locations. By applying a phase shift to one of the two

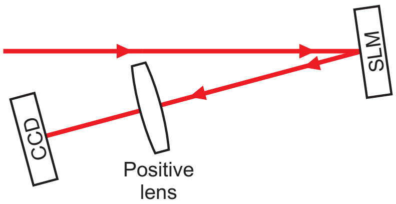

Figure 3.5: Calibration using the far field diffraction pattern of simple phase masks. The lens applies the Fourier transform to get the far field diffraction.

beams, the interference pattern of the returning beams is changed. This is essentially similar to the Michelson interferometer, but has only one beam that has to be aligned. This is offset by the need for a proper grating. The method proposed by Holoeye is a simplification of this scheme. In that setup the two beams are not collinear as in the Vaughan setup, but the lack of a prism greatly simplifies the implementation.

A final method by Chimento[13] uses the far-field diffraction pattern (Fig. 3.5). The far-field diffraction pattern for simple phase plates can be solved algebraically, and the measured diffraction can thus be readily compared. Chimento uses a Heaviside function as a phase pattern: a jump in phase across the middle of the SLM with a flat (but different) phase on both halves. This is actually the same phase pattern as used in the aforementioned setups, but with a better name. In this configuration, the relation between voltage and phase can be found by analysing the heights and depth of the two peaks and one trough in the far field diffraction pattern associated with the Heaviside function. The setup does require good alignment of the beam on the centre of the phase mask, but this has the advantage of having a well aligned beam for subsequent experiments.

The Michelson interferometer of Frumker and Silberberg; and the setup of Vaughan using the polarizer were primarily considered for calibration of our SLM. Eventually, the method suggested by Holoeye proved to give better results with a simpler setup. The method used by Chimento was found later and not applied to our SLM.

There is most likely a dependence of the SLM input to phase shift correlation on the incoming light's wavelength. When broadband pulses are shaped, this correlation has to be re-established as is done in Section 4.4.

3.3 Calibration results

Initially, a simplified version of the polarisation setup of Vaughan (Fig. 3.6) was constructed to determine the SLM input vs. phase delay calibration.

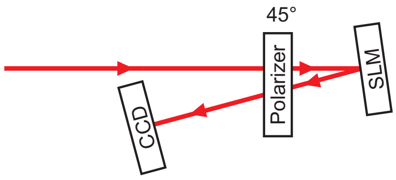

Figure 3.6: Simplified setup with a polarizer to allow amplitude modulation. The beam is passed twice through a polarizer at 45° while undergoing a change in polarisation on the SLM that depends on the applied phase delay. On the second pass through the polarizer, the polarisation modulation is converted into amplitude modulation.

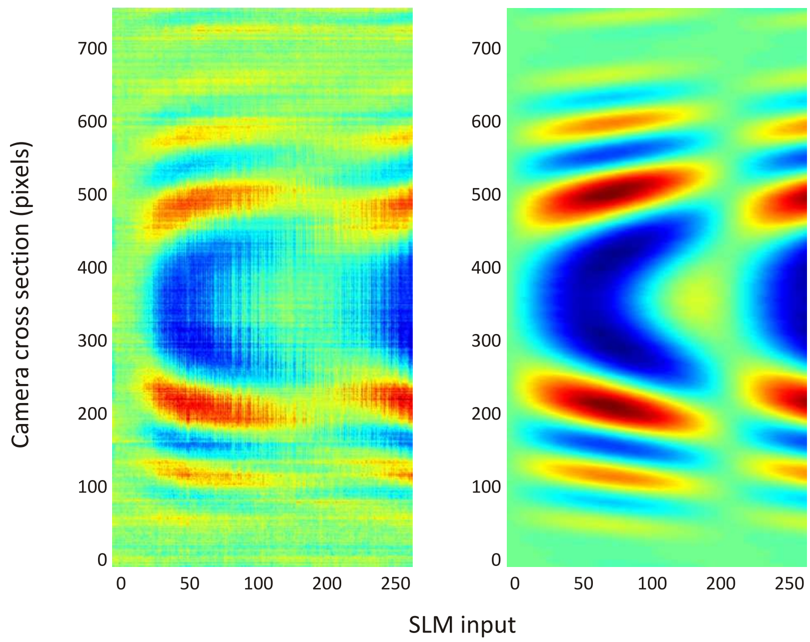

Figure 3.7: Intensity plot of measured (left) and calculated (right) near field diffraction pattern for a phase mask with one horizontal bar of varying phase delay. The horizontal axis gives the phase delay as the input value for the SLM; the vertical axis is a cross section of the CCD.

With the polarizer in place, a checkerboard pattern was applied to the SLM with varying sizes of the squares. The resulting intensity distributions differed greatly from the expected patterns. For specific checkerboard sizes and phase delays the unshaped parts of the beam would also display some amplitude modulation. To simplify the setup, only a single vertical band of light was given a phase delay. On the camera, this again caused variations in intensity over a much larger region of the beam. A numerical analysis of the near field diffraction offered the explanation as can be seen in Fig. 3.7. The large distance between the SLM and the detector (approximately 0.8m) greatly enlarged the relatively small diffraction off of the edges of the

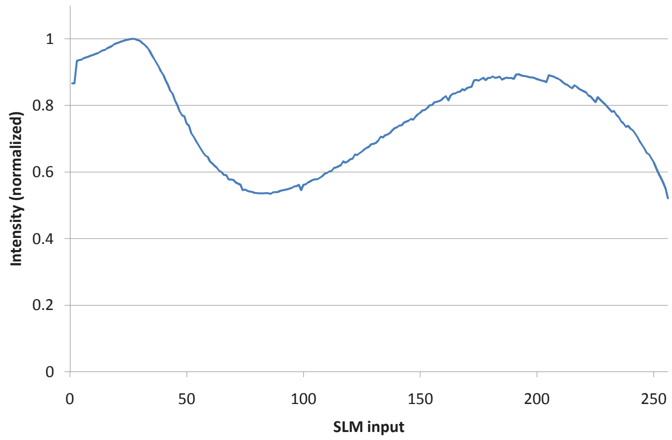

Figure 3.8: Measured intensity for varying phase delays using a polarizer at 45° as in Fig. 3.6. The SLM input is defined by the green colour value of the pixels sent to the SLM.

phase steps in the checkerboard pattern. This also explains the observed intensity variation in the unshaped parts of the checkerboard patterns.

There is one important feature in Fig. 3.7 that should be recognized. There is a slight curvature in the lobes of the measured intensity plot which is not present in the simulations. This curvature is caused by the non-linear calibration of the SLM. Ideally, a linear increase in input value (horizontal axis) leads to a linear increase of the phase delay. In reality, this relationship is more complicated and it is precisely this relation that we are trying to determine in this part of the project.

Moving the detector closer to the SLM reduces the extent of the diffraction pattern as the Fraunhofer limit is approached:

$$R > \frac{a^2}{\lambda} \approx 30\,\text{cm}, \tag{3.1}$$

wherein R the distance to an aperture with width a. However, for complicated phase masks with small features (a in Eq. 3.1), the effects are still noticeable.

By averaging over many measurements, it can also suffice to apply a phase delay to the whole SLM and thus the whole laser beam. This does mean that we cannot normalise the intensity with an unmodulated beam. Since the near field diffraction will induce intensity fluctuations in the unmodulated beam anyway, this is a small price to pay. The laser intensity has also proven to be sufficiently constant across the duration of the experiments to do away with this normalisation.

The summed intensity on the CCD as a function of the SLM input is given in Fig. 3.8. On the positive side, the curve resembles a sinusoid if we take into account the aforementioned nonlinearity. However, two odd features stand out: the intensity initially rises for an increasing phase delay and at high input values there are some jumps in intensity. We will at first ignore these oddities and attempt to map this curve to a proper sinusoid with suitable amplitudes and attenuation. After all, the expected plot of intensity as a function of the phase delay is a neat sinusoid. A short explanation follows, referencing the elements of the setup as shown in Fig. 3.6.

After the light passes the polarizer for the first time, it will have a polarisation direction $\vec{r}$. We can describe this beam as a sinusoid with an amplitude A_{in} and split it into two parts with orthogonal polarisations, of which only one will be shaped by the birefringent SLM. For simplicity, we will assume a polarisation angle of 45°.

$$\begin{aligned} A_{in}\sin(\frac{2\pi c}{\lambda}t)\vec{r} &= A_{in}\sin(\frac{\pi}{4})\sin(\frac{2\pi c}{\lambda}t)\vec{s} + A_{in}\cos(\frac{\pi}{4})\sin(\frac{2\pi c}{\lambda}t)\vec{p} \\ &= A_s\sin(\frac{2\pi c}{\lambda}t)\vec{s} + A_p\sin(\frac{2\pi c}{\lambda}t)\vec{p} \end{aligned} \tag{3.2}$$

Here, the senkrecht polarisation direction $\vec{s}$ (with amplitude A_s) is the polarisation perpendicular to the optical table. It is this polarisation that is shaped by the SLM. The other polarisation direction $\vec{p}$ (with amplitude A_p) is parallel with respect to the optical table and reflected by the SLM without modulation. Upon reflection off the SLM, a phase shift ϕ is imparted on the parallel part of the beam, leaving us with two sinusoids that have different amplitudes and a phase shift.

$$\rightarrow A_s\sin(\frac{2\pi c}{\lambda}t + \phi_{shaped})\vec{s} + A_p\sin(\frac{2\pi c}{\lambda}t)\vec{p} \tag{3.3}$$

Now we have an elliptically polarized beam that goes back through the polarizer at 45°. If we still consider the two beams separately, then the beams after polarisation are described by the following equation.

$$\begin{aligned} &\rightarrow \sin(\frac{\pi}{4})A_s\sin(\frac{2\pi c}{\lambda}t + \phi_{shaped})\vec{r} + \cos(\frac{\pi}{4})A_p\sin(\frac{2\pi c}{\lambda}t)\vec{r} \\ &= A_{s,out}\sin(\frac{2\pi c}{\lambda}t + \phi_{shaped})\vec{r} + A_{p,out}\sin(\frac{2\pi c}{\lambda}t)\vec{r} \end{aligned} \tag{3.4}$$

To calculate the intensity as we measure it with the CCD, we merely have to square and inte-

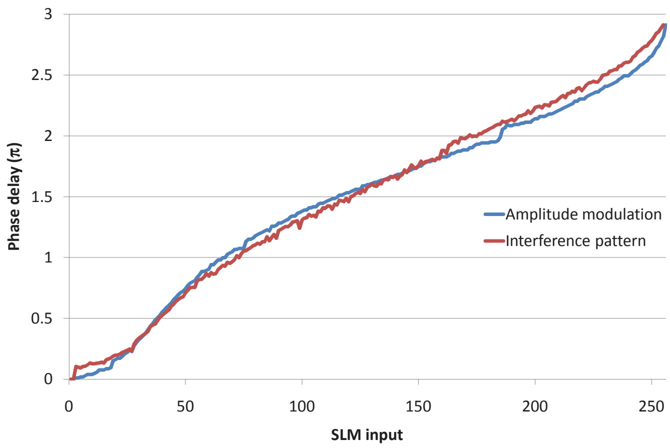

Figure 3.9: Calibration curves for the phase delay as a function of input level. The blue curve was obtained using amplitude modulation; the red curve is obtained using an interference pattern.

grate these sinusoids.

$$\begin{aligned} I &= \int_0^{2\pi} \left(A_{s,out} \sin\left(\frac{2\pi c}{\lambda} t + \phi_{shaped}\right) + A_{p,out} \sin\left(\frac{2\pi c}{\lambda} t\right) \right)^2 dt \\ &= A \cos(\phi) + B \end{aligned} \tag{3.5}$$

Apparently, all we are left with is a cosine with an offset and amplitude. For the case of a polarizer at 45° we have $A = B = \frac{1}{2}\pi$, but Eq. 3.5 actually holds for all polarizer directions. For a polarisation angle α, Eq. 3.5 becomes:

$$I = \int_0^{2\pi} A_{in} \left(\sin\left(\frac{2\pi c}{\lambda} t + \phi_{shaped}\right) \sin(\alpha)^2 + \sin\left(\frac{2\pi c}{\lambda} t\right) \cos(\alpha)^2 \right)^2 dt. \tag{3.6}$$

Figure 3.8 should thus be mapped to a cosine, albeit with a phase shift to compensate for the initially rising intensity. If this is done correctly, then Fig. 3.9 is found. There are three jumps discernable in this curve — most notably at an input of 180 — exactly there where the measured intensity has its maxima and minimum. The accuracy with which two cosines can be mapped around their maxima or minima is of course relatively small and as such this method is restricted in its accuracy, at least for a part of the input values.

The initial rise of the intensity and the jumps at high input values have not yet been explained. The rise in intensity is probably caused by an inherent birefringence of the SLM. The parallel polarisation will then always be given a phase delay compared to the orthogonal polarisation. Initially, an applied phase delay will compensate this inherent refractive index

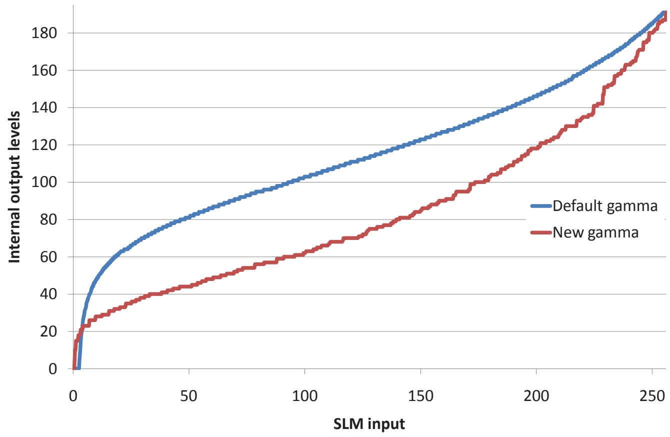

Figure 3.10: Default and new gamma curve for a 5:5 bitplane configuration.

difference and cause the intensity to rise. The jumps in intensity at higher input values can, however, not be easily explained. They do appear to be reproducible across measurements, but are probably caused by alignment issues as later experiments with different setups did not exhibit this phenomenon.

Since the calibration of the SLM is limited in accuracy when using amplitude modulation, the alternative method as suggested by Holoeye (Fig. 3.2) has also been tried. A double slit was custom made from a business card and the resulting interference pattern was imaged onto the CCD via two lenses. The second lens served to enlarge the image for better analysis. Half of the SLM was kept at a fixed input level of 0, while the other half was scanned across the 256 different shaping levels. The shifts of the interference peaks were recorded and from this the phase shift — measured in radians — was derived. The resulting relation between input level and phase delay is compared with the amplitude modulation in Fig. 3.9. A more detailed explanation of this calibration technique is given in Appendix B.

When the relation between SLM input level and phase delay has been determined, the SLM software can be updated with an improved *gamma curve*. The gamma curve is the internal look up table of the SLM which converts the input of the user to the appropriate voltage for the liquid crystal. Ideally, a linear increase in SLM input leads to a linear increase in phase delay. Adjusting the gamma curve based on the calibration experiments explained above, allows us to obtain this linear input-phase relation.

For the current configuration (see Section 3.4), there are 192 discrete voltage levels available to drive the SLM. We thus have to map 256 (8 bit) input values to 192 voltages in such a way that we get a linear relation between our input and the applied phase delay. The de-

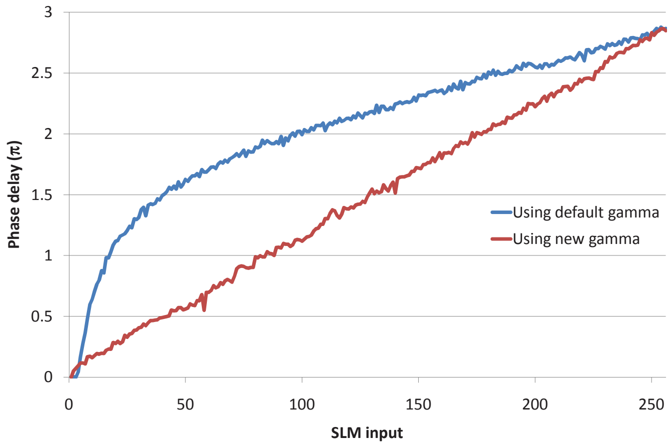

Figure 3.11: Measured phase delay as a function of input level for a 5:5 bitplane configuration using the default gamma curve (blue) and the optimized gamma curve (red).

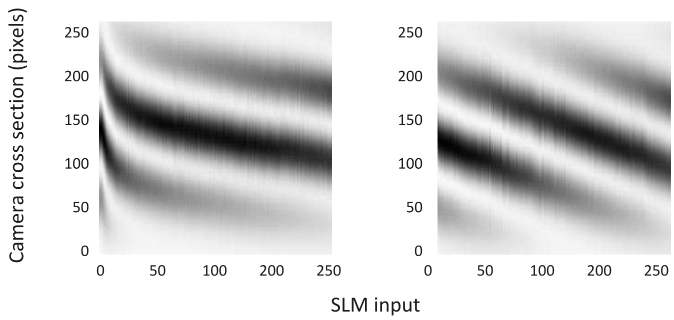

Figure 3.12: Observed interference pattern shift for increasing phase delays. Shift with default gamma curve (left) and optimized gamma curve (right). With the new gamma curve a linear relation between the SLM input and the associated retardance is obtained.

fault gamma curve is shown in Fig. 3.10. With this curve loaded in the SLM, we obtain the phase delays as in Fig. 3.11. Using the default gamma and the associated phase delays, we can find a new gamma curve that should give a linear phase delay. The observed interference patterns before and after the calibration are shown in Fig. 3.12 and it is clear that this calibration method allows us to obtain a linear relation between the input values and the resulting phase delay.

3.4 Bitplane configuration

After obtaining the improved gamma curve, a very simple helical phase mask was applied to obtain a so called Laguerre-Gaussian beam profile: a donut mode. While the donut-shape

Figure 3.13: Laguerre-Gaussian beam profile showing radial asymmetry.

was clearly visible, a strong vibration was also present in the beam profile. This vibration was not caused by mechanical vibrations in the system, or by instabilities in the laser. Holoeye confirmed that this flickering is caused by the SLM, and they consider it normal behaviour: "Due to the limited viscosity of the LC molecules these molecules flicker a little bit around the mean value." They suggested reconfiguring the SLM to use a 5:5 bitplane configuration, instead of the default 22:6 bitplane configuration.

By changing the bitplane configuration, the addressing frequency of the liquid crystals can be increased. While the driving electronics only accept 60 frames per second from the computer, internally they can duplicate these frames. Each frame is then sent several times to the liquid crystals within the refresh rate of 60Hz. However, as more updates are sent to the crystals each second, less information can be contained in each. As a result, some resolution is lost in the driving voltage of the SLM. With a 22:6 bitplane configuration 1472 voltage levels are accessible (through 256 input levels), while with a 5:5 bitplane configuration only 192 voltage levels are accessible (through 256 input levels).

After setting the bitplane configuration to 5:5 some reduction of the flickering was found. Even with the 5:5 bitplane configuration, a strong temporal dependence of the phase shaping effectiveness is found. Frumker and Silberberg report a similar modulation with a characteristic frequency of 300 Hz.[9] They assign this flickering to the electronic modulation scheme and state that this is not a physical limitation of the liquid crystals.

3.5 Spatial shaping results

Some very simple beam profiles and diffraction patterns have been created with spatial shaping. The Laguerre-Gaussian beam mentioned in Section 3.4 was made by applying a helical phase mask that sweeps through a whole 2π phase shift. In the centre of the beam at the focal

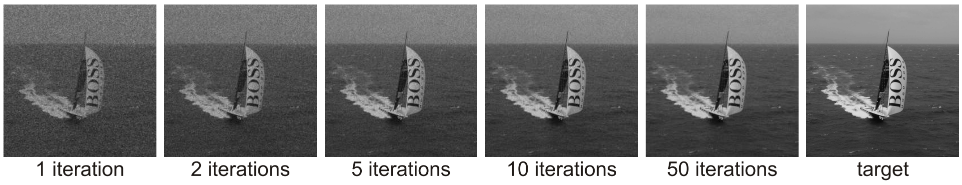

Figure 3.14: Optimisation of phase mask for a complicated target using the Gerchberg-Saxton algorithm. After only a few iterations a good approximation of the target is obtained.

point, the beam will then destructively interfere. In this way, a donut mode is created.

Apart from the jittering mentioned earlier, there is a radial asymmetry in the resulting beam profile as can be seen in Fig. 3.13. A similar effect is observed by Chimento using a MEMS device.[13] The asymmetry is not dependent on the orientation of the phase mask, and changing the position of the beam on the SLM only has a small influence on the intensities of the two lobes but not on their positions. The asymmetry can be ascribed to the spatial deformation of the SLM as well as the non-uniform effective phase delay discussed in Section 3.6.

More complicated beam patterns require the calculation of unintuïtive phase masks. The Gerchberg-Saxton (GS) algorithm[14] gives a way of iteratively optimising a phase mask to approximate a known target image in the focal plane. For most problems a good approximation of the target is obtained after only a few iterations, (Fig. 3.14). If strong restraints — such as binary phase steps — are added, convergence is less good. Many alternatives to the GS algorithm are available[15–17], but none share its simplicity and speed. Of course, this is offset by less noise, sharper features, and less zero-order light. In our work, the speed of the GS algorithm was preferred over the accuracy of its alternatives. With proper optimisations, phase mask calculations at video-rate may even be possible.

In the GS algorithm (Fig. 3.15), an initial random phase mask is applied to the known (Gaussian) amplitude profile of the incident beam and the result is Fourier transformed to obtain the Fraunhofer diffraction pattern at the focal plane. While the calculated phase is maintained, the amplitude is replaced by the desired amplitude of the target image and the inverse Fourier transform is calculated. The resulting phase is the new phase mask, and the amplitude should be replaced with the known amplitude of the incident beam. This process can then be repeated ad infinitum.

In principle, the GS algorithm could also be used with a feedback loop. If the far field image — containing the actual amplitude profile as opposed to the calculated target A_d— is captured with a camera, this can be used to approximate the additional phase shaping from the optical elements between the SLM and the camera, as well as any irregularities in the beam profile and

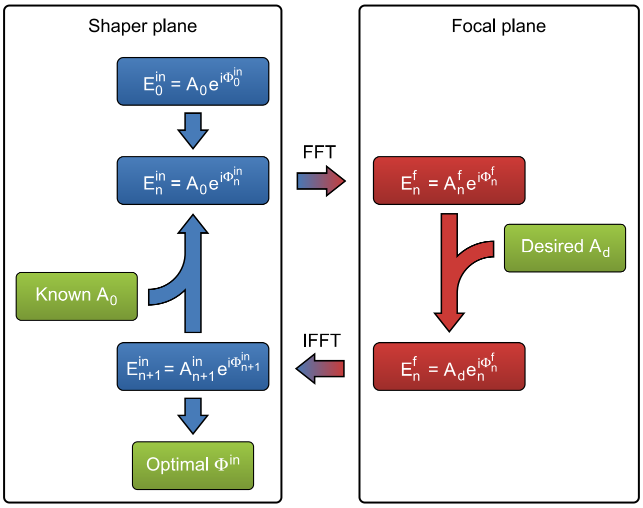

Figure 3.15: The Gerchberg-Saxton algorithm. (Ref. 14)

in the SLM itself. With this additional shaping in mind, a new phase mask can be calculated in order to get an improved image on the camera. When this process is repeated, a noticeable improvement can be obtained. A slight experimental complexity that has to be overcome is the mapping of the measured beam profile to the matrices in the algorithm. In order to get good resolution both on the SLM and in the focal plane, the desired image has to fill the whole matrix used in the GS algorithm. This means that in the experimental setup one also gets a very large image which could not be projected onto the available camera. Ordinarily, the Fourier transform could be applied to a larger matrix padded with zeroes in order to obtain a smaller image with equal resolution. However, given that the SLM already has 2 megapixels, this makes the calculations infeasible. A better solution would be to sacrifice some resolution in both domains and use some appropriate lenses to project the image onto the camera.

The large diffraction patterns calculated with the GS algorithm are all easily viewed with the naked eye. The vibrations of the liquid crystals do not have a noticeable effect in that case due to the inherent temporal filtering of the human eye.

3.6 Effective phase delay

In the specification sheet of the SLM, Holoeye discusses the optical flatness of the display. The listed maximal deviation of the SLM surface is 3–4 waves at 633 nm. This is much larger than the customary flatness of mirrors in the order of tens of nanometers. Since the deformation is mostly spherical, it can be compensated for directly in a 4f setup (see Section 4.2) by moving

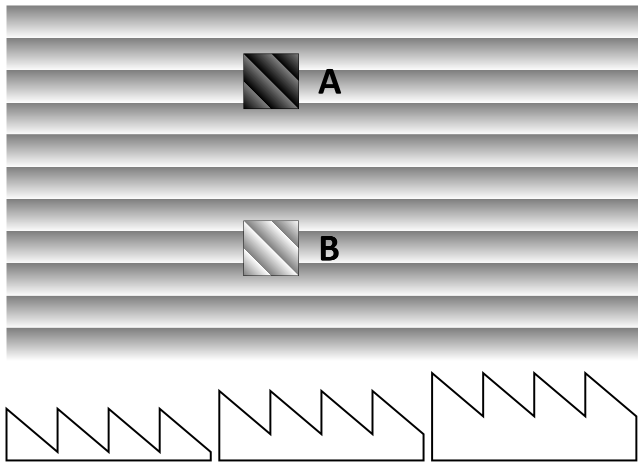

Figure 3.16: Phase mask applied to the SLM for obtaining the effective phase delay per superpixel. The light from A and B is steered off from the main bundle via a diagonal grating. At A, an additional, variable, flat phase is added on top of the sawtooth grating in order to obtain the effective phase delay of the pixels in A (bottom).

the SLM slightly out of focus. Additionally, a phase mask can be applied to the SLM that further compensates the deformations. However, the varying thickness of the LC panel will also influence the phase delay range. In order to analyze the variation in effective phase delay, a novel calibration method has been developed.

Because the HEO 1080P SLM is a two dimensional display, gratings can be displayed that can steer incident light off into higher diffraction orders. By applying such a grating to a small part of the display, a part of the light can be diffracted onto a camera. We use this to very selectively analyze small parts of the SLM. The effective phase delay is then determined using the same method as in Section 3.3.

All the light is steered off vertically into a higher diffraction order, as shown in Fig. 3.16. To two small squares of 32 by 32 pixels a different grating is applied that diffracts a part of the light away from the main beam. By using a diagonal grating, we obtain the least interference from the rest of the light. The applied grating has a phase range of 0 to 128 input levels. This leaves another 127 input levels for a constant phase delay. One of the two squares is thus given an additional variable phase delay and the resulting shift of the interference pattern is recorded with a camera. This process is repeated for the whole SLM, subdivided into superpixels of 32 by 32 pixels. The maximal phase delay per superpixel is then plotted to give a map of the maximal retardance as a function of position on the LC display. A characteristic plot is given in Fig. 3.17. The distribution of the maximal phase delay roughly matches a spherical deformation. An ad-

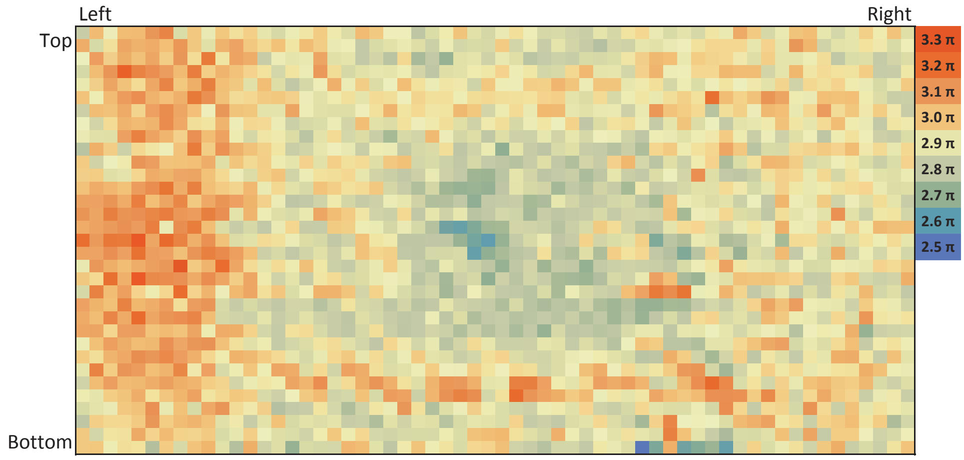

Figure 3.17: Maximal phase delay per superpixel as a function of position on the SLM. The variation per superpixel is up to 10% per measurement; the overall shape, however, is constant.

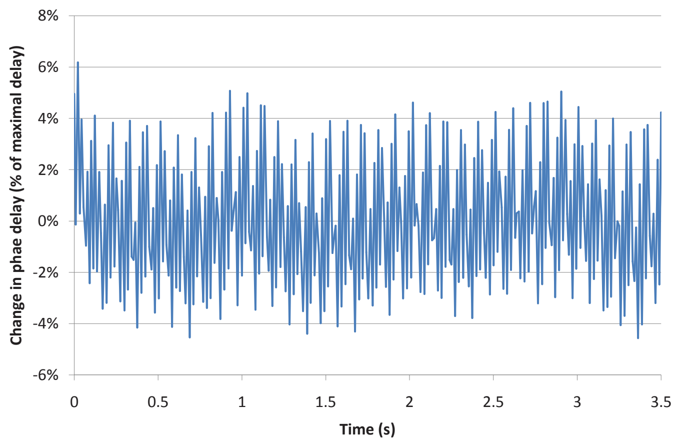

Figure 3.18: Fast fluctuations in the maximal phase delay of one superpixel in the middle of the SLM. The vibrations are undersampled due to the response time of the CCD.

ditional thickening is observed at the left side of the SLM. The measured delay differences also correspond well with the 1.8–2.4 μm maximal thickness difference reported by Holoeye. Given a common LC thickness of roughly 10 μm, a 2 μm thickness difference results in a maximal phase delay bandwidth of 0.75π, similar to the results shown in Fig. 3.17.

Despite long integration times of 30 ms, a strong time dependence was observed with maximal variations of 10%. First of all, there are very fast fluctuations as in Fig. 3.18, where the maximal phase delay for one superpixel is shown as a function of time. These variations have previously been observed by Frumker and Silberberg[9] and found to have a frequency of 300 Hz

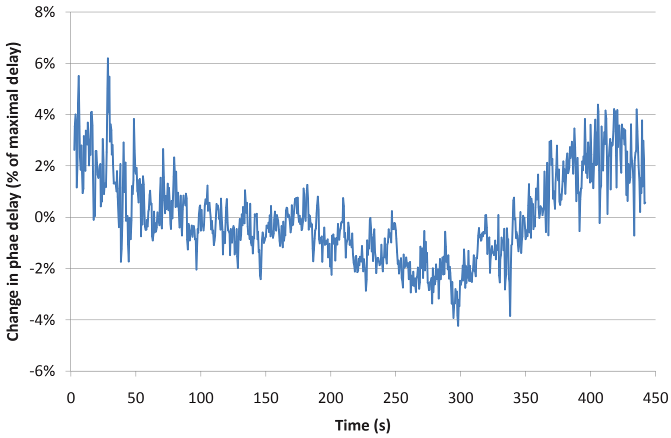

Figure 3.19: Slow fluctuations of the maximal phase delay of one superpixel in the middle of the SLM.

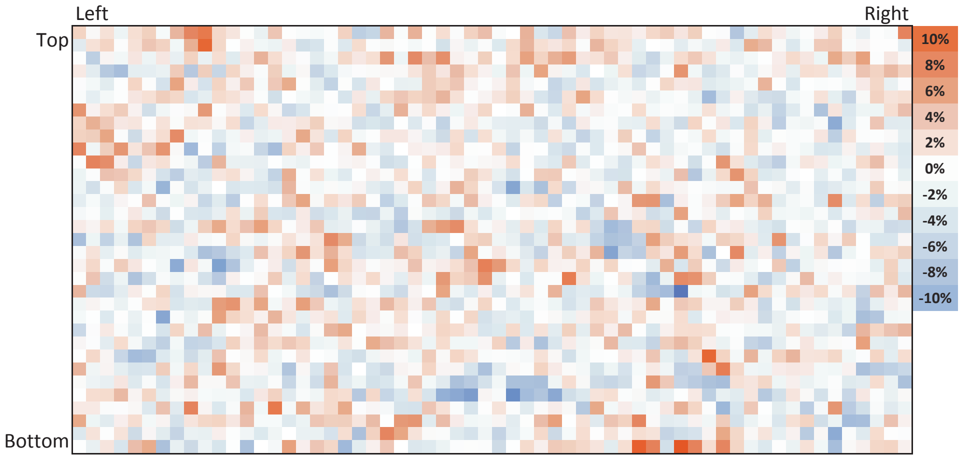

Figure 3.20: Difference in shaping effectiveness between two consecutive measurements. Fig. 3.17 gives one of the two original measurement series.

(in the measurements these vibrations were undersampled due to the finite response time of the CCD). Apart from the high frequency vibrations, a much slower time dependence is also observed. By increasing the integration time of the CCD, these variations can be easily observed. Fig. 3.19 shows the change in maximal phase delay over the space of several minutes. When two consecutive characterisations of the SLM are done, the influence of the long term fluctuations becomes more apparent. Fig. 3.20 shows the difference between two measurement series. The total measurement time per superpixel is about 3 seconds and the time between the two measurement series is several hours. Strong variations of up to 10% are observed across the SLM. The overall positional dependence of shaping effectiveness, however, is comparable between different measurements. Real pixel-by-pixel calibration is thus unfeasible, but a good profile can be obtained by applying a smoothing filter to the measured profile at the cost of some resolution.

4 Spectral shaping

With the basic calibration of the spatial light modulator described in Chapter 3, we now turn to the spectral domain. First, the setup and alignment of a folded 4f zero dispersion pulse compressor setup is explained, after which a quick and simple method for wavelength versus pixel calibration is introduced. The spectral shaping capabilities of the SLM are demonstrated using a second harmonic generation experiment.

4.1 Femtosecond laser source

A Coherent Legend Elite Ti:Sapphire laser amplifier is used to generate femtosecond pulses at 800 nm with a FWHM bandwidth of 36 nm. The amplifier is pumped with a tuneable Coherent Micra Ti:Sapphire laser, which in turn is pumped by an integrated Verdi laser. The pulses emerge with a power of 3 W and a repetition rate of 5 kHz. Since the amplifier is needed to power a Coherent OPerA Solo Optical Parametric Amplifier, only a part of the 3 W is sent into the shaper setup. As the maximum permissible power for the SLM is not known to us or the manufacturer, this ensures that we err on the side of caution.

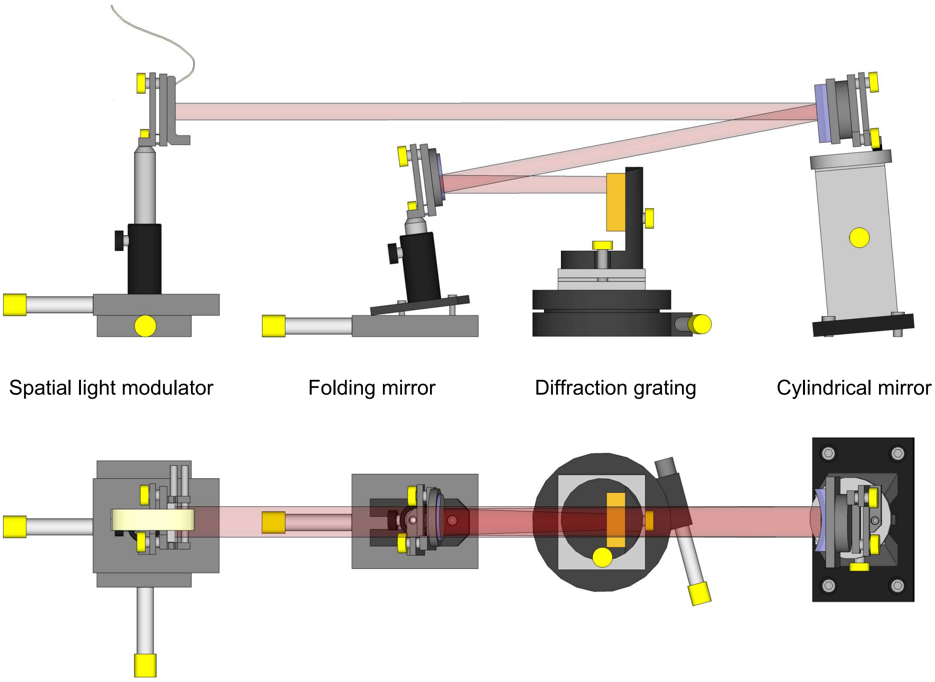

Figure 4.1: Spectral shaping setup; the yellow control knobs are used for fine adjustment of position and rotation of the various elements in the setup.

4.2 Shaper setup

For shaping in the spectral domain, the SLM is placed in a 4f zero dispersion pulse compressor.[4] Since the Holoeye HEO1080P is a reflective SLM, the incoming and outgoing pathways of the shaper setup are identical.

A plan of the setup is given in Fig. 4.1. The broadband pulse is initially dispersed on a gold coated ruled diffraction grating from Edmund Optics with 830 lp/mm and a blaze angle of 19° for 800 nm. After passing a folding mirror, the first order diffraction is focussed by a 25 cm cylindrical mirror which was gold coated in-house.

The folding mirror and SLM are placed on linear stages with micrometer screws for fine adjustment of the distances to the cylindrical lens. A second linear stage under the SLM moves it perpendicular to the propagation plane for alignment of the spectrum on the SLM. Both the cylindrical mirror and the grating are affixed to AMOLF-built mounts that allow fine adjustment for all three axes of rotation.

4.2.1 Alignment

The alignment of the shaper setup is a multistep process which is complicated by the small form factor of the setup as well as the — partially invisible — wavelength range of the laser source.

First and foremost, the grating has to be aligned so that all diffraction orders are in the same plane as the incident beam. This ensures that the first diffraction order emerges parallel to the table from the grating.

The initial alignment steps are done in the zero order. In those instances, a mirror should be substituted for the SLM to prevent accidental burning of pixels in the focus of the zero order.

- One after another, the grating, folding mirror, and cylindrical lens are placed into position, ensuring that the beam always emerges straight along the same path.
- The replacement mirror is placed at the position of the SLM. The rotation of this mirror must be such that the incoming and outgoing beams on the grating overlap horizontally. This can be checked easily with an IR viewer. Vertically, some tilt has to be applied in order to separate the two beams spatially on the grating. The distance between the replacement mirror and the cylindrical lens is adjusted until a collimated beam is obtained at the output of the shaper.
- The grating is rotated to the first order. The proper rotation is obtained if a collimated beam again emerges from the shaper and the two spots on the grating overlap.
- The mirror is replaced with the SLM and the distance between the SLM and the lens is adjusted until a collimated outgoing beam is again obtained.
- If the pulse length can be measured with an autocorrelator, then the folding mirror position can be adjusted by optimising for a minimal pulse length after the shaper.

After these steps, the spatial dispersion in the emerging beam must be minimized. This is most easily done by applying a grating to the SLM such that most light is diffracted away from the main beam, apart from narrow bands both at the far right and far left of the spectrum. By alternating between the red and blue part of the spectrum, the precise position of the different colours can be determined.

The vertical separation between the left and right part of the spectrum is adjusted by rotating the cylindrical lens around the axis in line with the beam path. Horizontal spatial chirp is

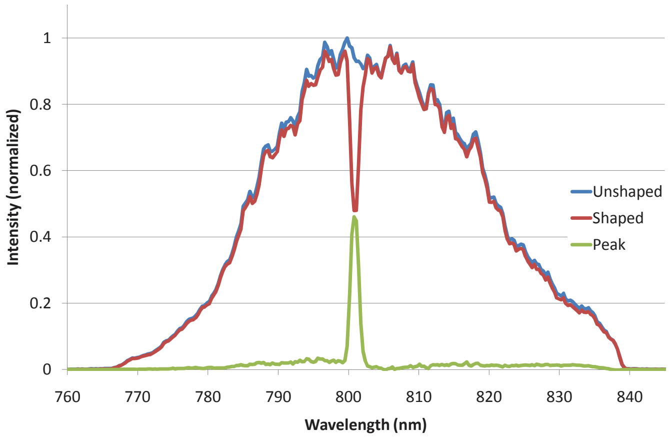

Figure 4.2: Spectral dip observed when a 32 pixel wide grating is applied at pixel 743. By scanning a grating across the SLM, the wavelength versus pixel relation can be obtained.

compensated for by rotating the grating and SLM in unison. If the initial alignment has been done properly, this should not be necessary.

4.3 Wavelength versus pixel calibration

To determine which colour is focussed on which pixel, the two dimensional shaping capabilities of the SLM can be employed. By applying a grating in the vertical direction of the SLM, the beam can be realigned spatially. If this grating is applied selectively to a band of a few pixels wide, it is possible to steer off a part of the *spectrum* of the beam. By analyzing the resulting spectrum, the wavelength versus pixel calibration is easily obtained.

Figure 4.2 shows a typical measurement as the grating is scanned along the LCD. After an initial determination of the peak position, the accuracy in the wings of the spectrum is improved by extrapolating the well-defined peak positions in the centre of the spectrum to the less intense wings. In this way, an accurate determination of the wavelength per SLM pixel can be determined as shown in Fig. 4.3.

4.4 Wavelength versus effective phase delay

As previously explained in Section 3.6, we can determine the maximum retardance as a function of location on the SLM. This allows us to determine the relation between wavelength and the effective maximal phase delay. For this purpose only a scan in the spectral dimension of the SLM is necessary.

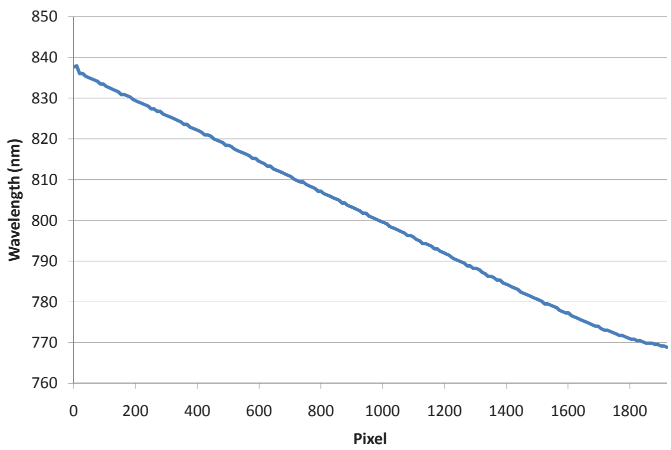

Figure 4.3: Pixel versus wavelength calibration. In the wings of the spectrum the accuracy of the calibration is reduced due to the low intensities.

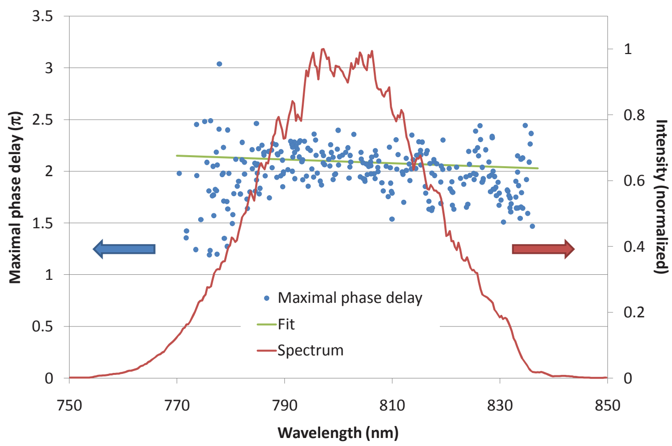

Figure 4.4: Retardance (blue dots) as a function of wavelength for maximal SLM input. A linear fit (green line) is drawn through the points where the intensity (red line) was sufficiently high. In the wings of the spectrum, the error in the maximum retardance is too great for an accurate fit.

By determining the effective phase delay, we can also verify the SLM input vs. phase delay relation found in Section 3.3. For all wavelengths of the laser system used, we still find a linear SLM input vs. phase relation. There is thus no need to create look up tables to convert a desired phase delay to the appropriate SLM input for each wavelength, but the single gamma curve within the SLM suffices.

Figure 4.4 shows the maximum retardance as a function of wavelength. The input spectrum is overlaid for comparison. In the wings of the spectrum the observed interference pattern was too weak to accurately determine the retardance; this results in a broader distribution of the measurements.

The large spread in the measured effective phase delays can be ascribed to the temporal variations described in Section 3.6. Nevertheless, a linear fit can be made through the points where the maximum retardance could be determined with high confidence. A near constant effective phase delay is found across the whole spectrum of the femtosecond pulses. Since the temporal fluctuations of the SLM are much larger than the influence of the wavelength on the effective phase delay, it is deemed unnecessary to compensate for this minor change in later experiments.

The results in Fig. 4.4 are not compensated for the spatial distortion found in Section 3.6. Apparently, the spatial distortion of the SLM and the wavelength dependent refractive index of the liquid crystals, combined with the alignment of the SLM in the 4f configuration, result in a near constant effective phase delay across the whole SLM. In later experiments with spectral shaping we can thus also forego compensating for the spatial inhomogeneity of the SLM.

4.5 Second harmonic generation

One of the simplest experiments that applies phase shaping of ultrashort pulses, is the optimisation of second harmonic generation (SHG) in a non-linear crystal. In such a crystal, the incoming pulses are converted into light with half the wavelength — the second harmonic — with an efficiency that is dependent on the intensity of the pulse. SHG is a χ_2 process, as in Eq. 4.1.[18,19] A shorter pulse of equal energy will result in higher SHG efficiency as the peak power is greater. By optimising the SHG signal, the pulse duration is optimised to approach a transform limited pulse (the shortest possible pulse). In the spectral representation this corresponds to a flat phase profile across the spectrum.

$$\vec{P}(t) = \epsilon_0 \left(\chi_1 \vec{E}(t) + \chi_2 \vec{E}^2(t) + \chi_3 \vec{E}^3(t) + \ldots \right) \tag{4.1}$$

A detailed description of the algorithm used to find the flat phase profile is given in Chapter 6. In short, the SHG signal of a random sample of phase profiles is measured, after which a computer algorithm suggests a new set of phase profiles of which the SHG signals are then again measured. With each loop through the optimisation — one generation — the SHG signal slightly improves, and the optimisation is terminated after a fixed number of generations, or when a stable phase mask is obtained.

With the SLM working in the spectral domain, the optimisation looks for the phase pro-

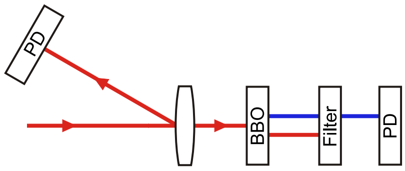

Figure 4.5: Setup for measuring SHG. Two photodiodes measure the SHG signal as well as a reference signal from the reflection from the focussing lens.

file that compensates best for phase distortions in the input pulse as well as for the effects of elements in the beam path such as dispersion in lenses. Important to note is that distortions occurring after the shaper are also compensated for.

SHG optimisation occurs purely in the spectral domain and thus does not make use of the two dimensional characteristic of the SLM. If the second (spatial) direction would also be optimised, the spatial and spectral optimisation could not be decoupled and no information can be extracted from the optimisation. Moreover, any spatial optimisation would be purely related to alignment of the SHG experiment, while the spectral optimisation gives insight into the pulse characteristics. Most other experiments call for transform limited pulses which is the main motivation to do an SHG optimisation experiment. Optimising the spatial dimension may result in higher SHG efficiency, but does not improve the beam characteristics. Most probably, a spatial optimisation would be restricted to improved focussing of the beam onto the non-linear crystal.

4.6 SHG optimisation setup

After being shaped, the beam is focussed onto a thin Beta Barium Borate (BBO) crystal of 25 μm with a 3.5 cm lens as in Fig. 4.5. The SHG signal is filtered with a bandpass filter, detected with a photodiode and enhanced via a transimpedance amplifier. A second photodiode measures the reflectance from the lens as a reference.

The measured second harmonic signal is divided by the reference to compensate for fluctuations in both laser- and shaper stability. The result of this is normalized with the SHG obtained with no phase mask applied to the SLM. The resulting amplification factor is used as the fitness factor in the CMA evolutionary algorithm.

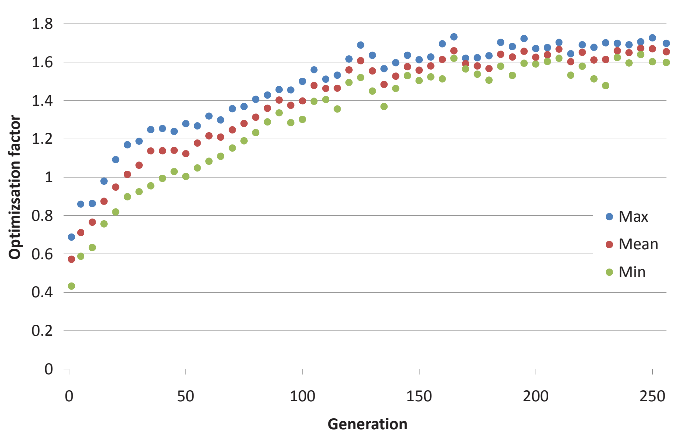

Figure 4.6: Best, mean, and worst fitness values per generation of a characteristic SHG optimisation run. This optimisation took two hours to complete with about 10,000 phase masks evaluated.

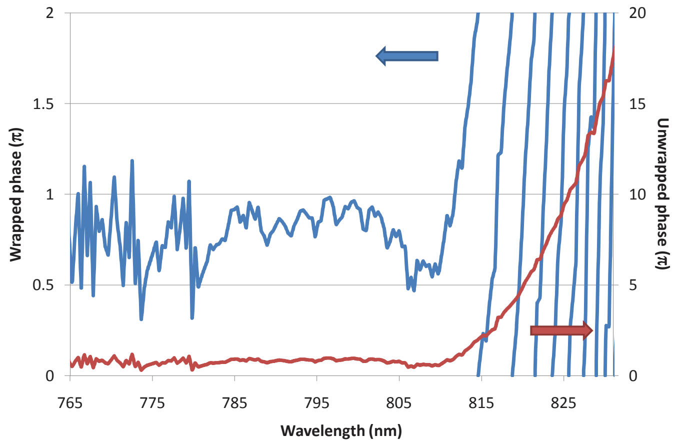

Figure 4.7: Best phase mask using a photodiode to measure SHG signal. The wrapped phase (blue line, left axis) can be unwrapped to obtain the red line (right axis). The phase at low wavelengths could not be unwrapped.

4.7 SHG optimisation results

A typical optimisation run is shown in Fig. 4.6. After starting off worse than the unoptimised pulse, the algorithm quickly improves in the first 50 generations. Then, a slow but steady rise continues until a maximum optimisation factor of 1.6 is reached. The final distribution of fitness values (as indicated by the maximal, minimal and mean fitness) is purely caused by the noise in the experiments. Longer optimisations will not show an additional reduction of this spread.

Fig. 4.7 shows the best phase mask at the end of the optimisation. For longer wavelengths

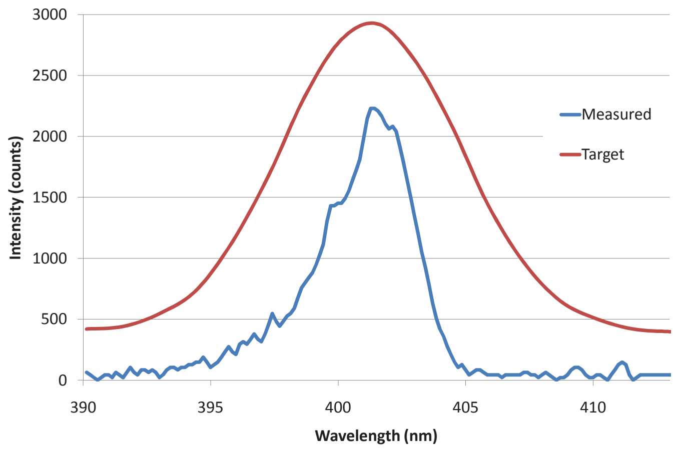

Figure 4.8: The measured SHG spectrum (blue) is optimised towards a target spectrum (red). The target spectrum is intentionally broader to increase the importance of the wings of the spectrum.

the phase could be neatly unwrapped and fit with a third order polynomial. The phase at low wavelengths, however, was not as neatly defined and no clear pattern is discernable.

It must be noted that the overall phase pattern, including the oscillations at the center wavelength — was accurately reproduced in multiple runs and with both the free optimisation, as well as the moving nodes parameterisation (see Section 6.3.3).

A trick employed in the experiments by Savolainen and co-workers[3] is to optimise towards a spectrum, rather than integrated SHG intensity. The SHG light is measured with a spectrometer and compared to a target spectrum. The target spectrum is intentionally chosen to be somewhat broader than the expected SHG spectrum (Fig. 4.8). This ensures that more effort is put into optimising the wings of the SHG spectrum, further improving the optimisation. This technique can also be applied to this optimisation, and the results are shown in Fig. 4.9. It is apparent that the wings of the spectrum are better resolved using this method.

Ultimately, the goal of the SHG optimisation is to obtain a transform limited pulse. In an autocorrelator trace, we should thus find the shortest possible pulse after the optimisation. In Fig. 4.10 the autocorrelation of the optimised pulse is compared to the pulse before the shaper. Some clear higher order distortions are apparent in the autocorrelation trace of the optimised pulse. Several explanations for this can be given.

First and foremost, the measured autocorrelation is not precisely that of the optimised pulse. As explained before, the SHG optimisation will compensate for all phase distortions in the shaper setup, as well as for the distortions of all optical elements thereafter. A thick focussing lens is used to focus the beam onto the BBO crystal (Fig. 4.5). This lens is not present

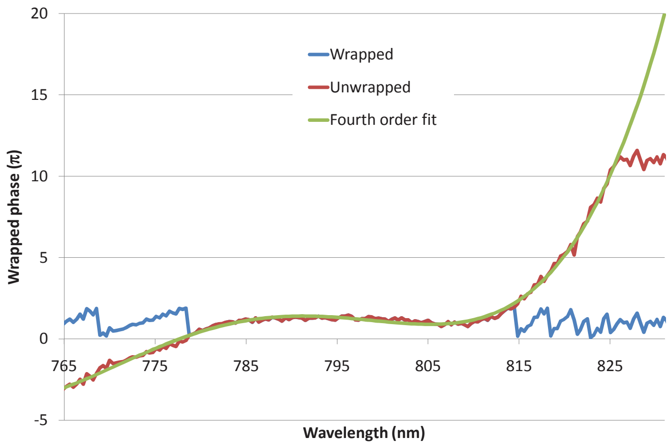

Figure 4.9: Best phase mask using a spectrometer to measure the SHG signal. The wrapped phase (blue line, left axis) can be unwrapped to obtain the red line (right axis). A fourth order polynomial fits the unwrapped phase neatly (green line).

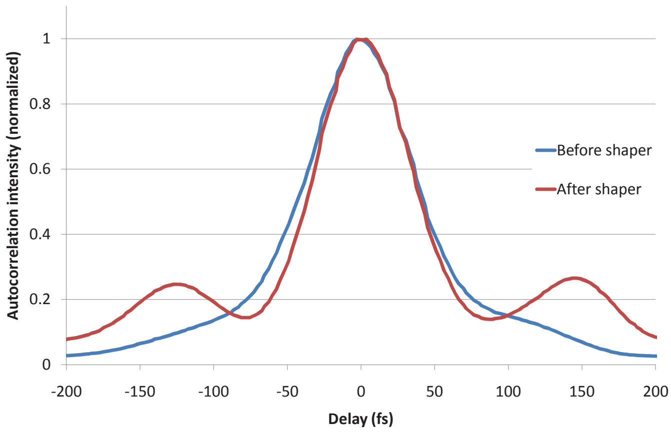

Figure 4.10: Autocorrelation traces for the input pulse (blue) and the optimised pulse after the shaper (red). Higher order distortions can be seen in the optimised pulse.

in the beam path to the autocorrelator. The phase of the measured pulse will thus include the phase distortions caused by the lens, but will not have undergone the chromatic aberrations of the lens when measured by the autocorrelator. While this can by itself explain the observed side lobes in the autocorrelation, another interesting phenomenon has to be discussed.

In the incoming spectrum (Fig. 4.11) fringes can clearly be seen, corresponding to a pulse train in time. With a fringe spacing of 3.0 ± 0.2 nm, this indicates a pulse separation of 0.71 ± 0.05 ps. While the autocorrelator has a temporal range of only 500 fs, some qualitative statements can be made if it is forced beyond this range. By doing this, small pulses are indeed found at sev-

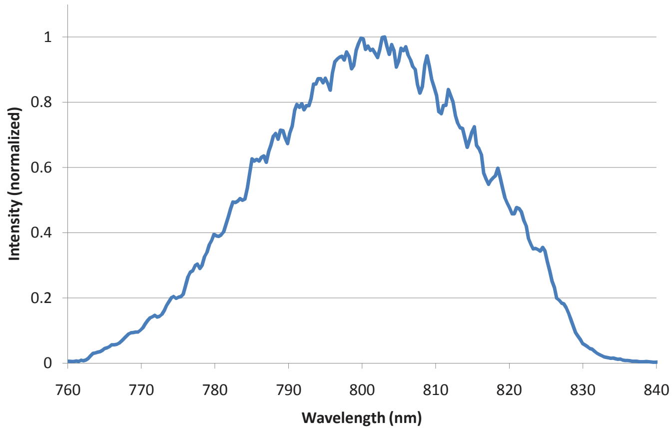

Figure 4.11: Spectrum showing fringes with a spacing of 3.0±0.2 nm. The corresponding pulse train with 711 fs spacing was also found in the autocorrelation.

eral hundred femtoseconds before and/or after the main pulse. Whether the pulses cause the modulation of the spectrum, or the spectrum creates pulses is a moot point. What we do know, is that the SHG optimisation is influenced by the presence of the pulse train.

The same oscillation present in the spectrum is found in the optimised phase mask of Fig. 4.9. The period of oscillation is established to be 2.7±0.3 nm, corresponding to the fringe spacing in the spectrum; the amplitude of these oscillations is 0.25π. Apparently, the peak SHG intensity is found if the pulse train is compensated for by phase shaping.

Given these aberrations, it is not to be expected that the peak SHG intensity corresponds to a single, narrow pulse. Nor should we expect to measure this with an autocorrelation. Instead, the transform limited pulse for this pulse train may very well include side lobes in an autocorrelation. Since the SHG optimisation converged to identical solutions on multiple occasions and with different parameterisations, it is save to conclude that a good enough optimum has been reached. Of course, the global optimum would require an unwrapped phase, which is impossible to find with these parameterisations.

5

In situ FROG

Creating complicated pulse shapes in order to control processes is one thing, but to then learn something about the process from this optimisation is another. The first step is to determine what the full time-frequency characteristics of the optimised pulse are. Ideally, such pulse characterisation is done at the sample position. This avoids creating distortions as seen in Section 4.7, where the beam path to the experiment and the beam path to the characterisation setup have significantly different behaviours. However, conventional setups for pulse characterisation such as frequency-resolved optical gating (FROG)[20] require multiple beam paths or well-known reference pulses. Obtaining a good characterisation of a femtosecond pulse at the precise location of an experiment is a challenge.

Herein lies a unique opportunity for a two dimensional SLM. By dividing the spatial dimension of the SLM into multiple parts, several distinct beams can be created from one incoming laser beam. To each of these, a separate spectral phase shape can be applied. The newly created beams are inherently parallel and can therefore easily be combined in a sample. This allows for simple pump-probe experiments, shaped single beam CARS,[8] or other multibeam

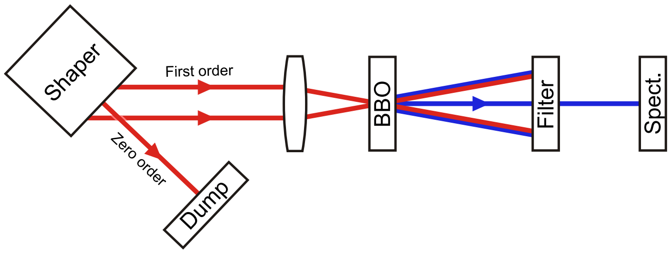

Figure 5.1: FROG setup with an SLM in place of a delay line. Two beams coming off the SLM are overlapped in a BBO crystal. The cross term emerges separate from the main SHG signals and can be measured with a spectrometer. A time delay is added to one of the beams with a linear phase mask across the spectral dimension, applied to the top half of the SLM.

schemes while needing only one incoming beam. What is more, we can also create second harmonic generation FROG (SHG-FROG) traces of our pulses as they pass through a sample.

5.1 Setup

Two beams are generated from the main beam by applying a linear phase in the spatial dimension of the SLM. By diverting the two beams we prevent interference from the zero order diffraction of the background. The background is then dumped somewhere along the beampath. Some spatial filtering is applied to the two beams with a vertical slit to clip the unwanted wings of the beam.

In a similar setup as used in the SHG optimisation (Section 4.5) and detailed here in Fig. 5.1, the two beams are focussed onto a BBO crystal of 25 μm thickness. Three distinct SHG signals now emerge from the crystal. Both beams will generate their own second harmonic signals, but a third pulse is generated by the combination of the two incoming beams. In other words, this signal is formed by the interaction of the two separate beams and can be used to obtain the time-frequency characteristics of the two.

Since the incident angles of the two beams are different, the three SHG signals will be separated spatially and focussed onto a spectrometer to measure the resulting spectrum. With an iris the desired beam can be selected.

Applying a linear phase in the spectral domain corresponds to a delay in time. Consequently, it is possible to replace the delay line of conventional multi-shot FROG setup with a simple phase mask on the SLM. When the SHG spectrum is measured for a range of time delays between the two beams, we obtain a complete FROG trace. We now have made a FROG setup that uses only one beam path, has no moving parts, and will give the spectrum and phase characteristics of a pulse at the sample position.

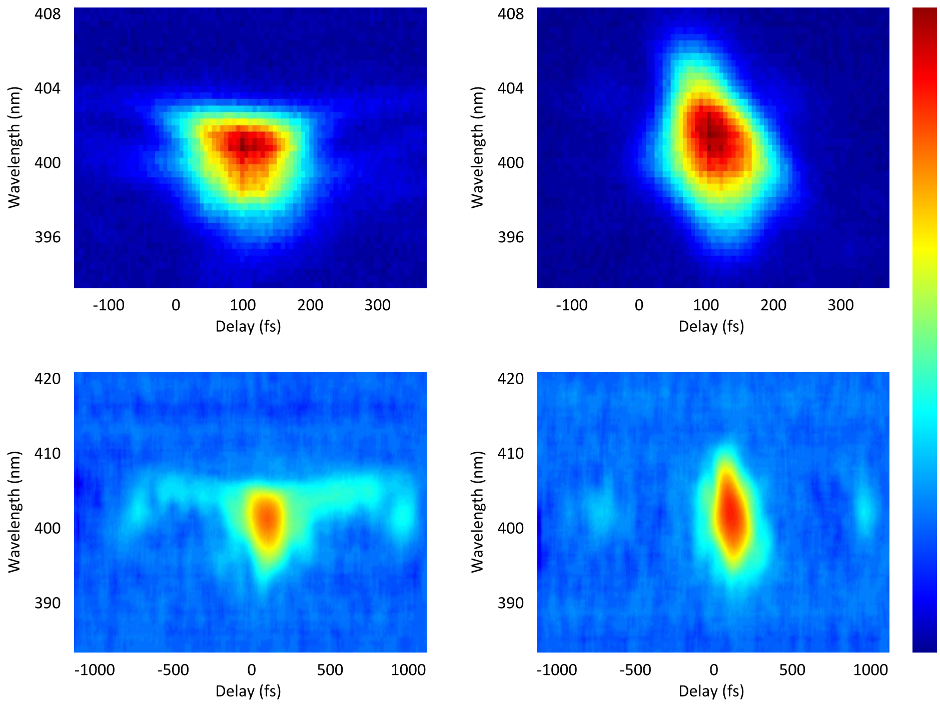

Figure 5.2: FROG traces of the raw pulse coming from the shaper (left) and the optimised pulse (right). The intensity has been normalised to the total intensity in both figures. The two figures at the bottom cover a wider temporal range and display the logarithm of the intensity for better contrast. At 700 fs satellite pulses can be seen, as described in Section 4.7.

FROG versus X-FROG

With this setup two different experiments can be imagined. In the classical FROG scheme, the to-be-measured pulse is characterised by gating it with a time delayed version of itself. This is the most trivial scheme that we can do with the in situ FROG setup.

However, if we use the knowledge gained in the SHG optimisation, then cross correlation FROG (X-FROG) becomes a possibility. One of the two beams can be made transform limited by shaping it with the phase resulting from the SHG optimisation (see Section 4.7). This gives a well characterised beam that can easily be adapted to any changes in the optical setup and with which X-FROG becomes feasible.

5.2 FROG measurements

Figure 5.2 gives the FROG traces for the non-optimised and optimised pulses, using the phase mask found in Section 4.7. If the contrast is increased by taking the logarithm of the intensity, a satellite pulse at 0.76±0.04 ps become visible. This pulse was already predicted based on the fringe spacing in the spectrum of Fig. 4.11 on page 45. The time delay compares well with the previously calculated value of 0.71±0.05 ps. In the optimised pulse most of the energy has been taken from the satellite pulse and put into the main peak. This FROG trace also indicates that no further improvement can be obtained through SHG optimisations, as the unwanted pulse train has already been reduced to the noise level.

While there is excellent control over the delay through the application of linear phase masks on the SLM, the effective spectral bandwidth of the FROG traces in Fig. 5.2 is insufficient. Only a part of the spectrum is efficiently converted by the BBO crystal into a second harmonic signal that sticks out above the noise level. The wings of the amplifier's spectrum are thus not represented in the measured SHG spectra.

This can be easily tested by applying a phase mask to the edges of the SLM (e.g. a linear phase delay) and observing the change in the FROG trace. The FROG trace was only changed when spectral components from 795 nm up to 820 nm were manipulated. The combination of decreasing intensity in the wings of the spectrum and they reduced phase matching efficiency for the same spectral components reduces the SHG signals to below the noise level.

The effective bandwidth of the FROG setup is thus limited by the signal to noise ratio in the wings of the spectrum. For proper retrieval of the FROG traces, the information contained within these wings is also required. Several factors influence this signal to noise ratio.

First of all, the bandwidth is defined by the phase matching bandwidth of the non-linear crystal and the numerical aperture of the focussing optics. By changing the orientation of the crystal, the central phase matched wavelength could be tuned. No combination of lenses or concave mirrors and BBO crystals was available that allowed for a sufficiently broad phase matching bandwidth.

The easiest way of improving the signal to noise ratio is probably to replace the fibre spectrometer with a more sensitive alternative. This would also present an opportunity to improve the spectral resolution of the FROG traces.

Finally, the beam parameters such as angle and separation of the two beams in the FROG setup can be optimised. In our experiments the two beams emerged parallel from the shaper with a vertical separation of 200 pixels on the SLM. By changing the linear phase masks on the spatial dimension of the SLM, the angle between the two beams can be arbitrarily chosen. while larger angles give rise to a different phase matching bandwidth in the crystal, this comes at the cost of increased geometric smear. Changing the vertical separation of the two beams

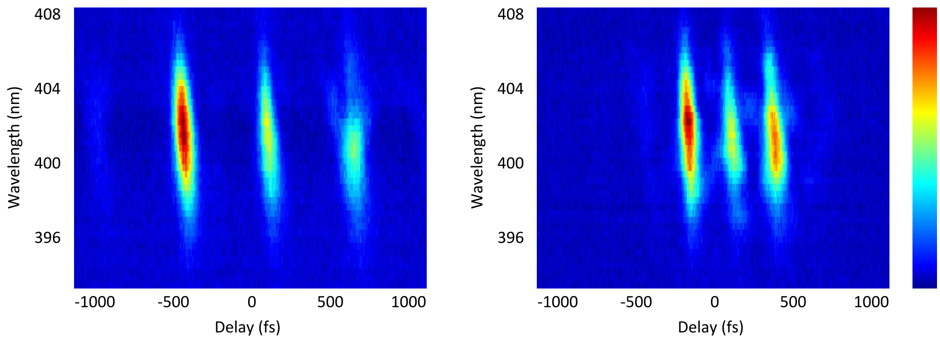

Figure 5.3: Pulse trains characterised with X-FROG. A 128 pixel (left) and 256 pixel (right) period sinusoidal phase masks were used to create these pulse trains.

can also improve the spectral bandwidth of the FROG trace, but may make it harder to separate the cross term from the other two SHG signals.

There is some asymmetry observable in the FROG traces of Fig. 5.2. Conventional SHG-FROG traces should be symmetric as they are essentialy an autocorrelation of the pulse. The fluctuating phase mask on the SLM, breathing of the amplified pulses, and geometric smear in the FROG setup all contribute to this asymmetry.

Spectrally shaped pulses

Despite the aforementioned limitations, shaped pulses could be characterised using the in situ FROG setup. First, simple sinusoidal phase masks have been applied to the SLM, resulting in pulse trains with controllable temporal spacing. Two of these pulse trains are shown in Fig. 5.3. With modulation periods of 128 and 256 pixels, pulse spacing is experimentally found to be 489±5 fs and 235±5 fs. The expected temporal spacing is given by Eq. 5.1 and is 471 fs and 235 fs, respectively.

$$\Delta\tau = \frac{\lambda^2}{c\Delta\lambda} \tag{5.1}$$

The FROG traces in Fig. 5.3 show the time-frequency characteristics of two different pulse trains obtained through X-FROG. The ability to do X-FROG with this setup is a great advantage. The optimised, transform limited pulse can be used to accurately characterise more complicated pulse shapes. Optimisation of the reference pulse can even be performed with the same experimental setup when the SHG signal from only one beam is measured and optimised.

As a final showcase experiment a complicated pulse was made with the SLM and simultaneously characterised using the in situ FROG setup. The results are shown in Fig. 5.4 wherein

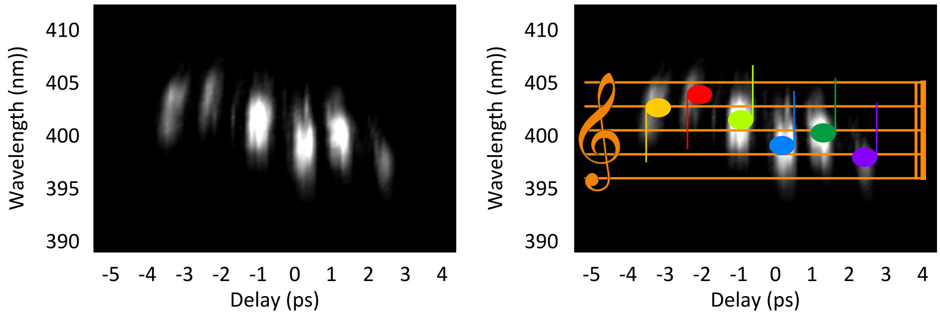

Figure 5.4: A complicated pulse is shaped with the SLM and simultaneously characterised by the same SLM. Shown are the opening notes of *The Entertainer* by Scott Joplin.

the opening notes of *The Entertainer* by Scott Joplin are played with a femtosecond pulse laser. Again, the spectral bandwidth was limited, so only 650 out of 1920 pixels of the spectral dimension of the SLM were used to generate this pulse.

The pulse shown in Fig. 5.4 also demonstrates the large temporal bandwidth obtainable with the SLM: a range of 10 ps is easily obtained. Even greater delays can probably be obtained, though for very large delays the accuracy of the FROG trace will decrease as parasitic diffraction off the phase mask begins to play an increasingly important role.

6 Learning loop

As the systems we wish to research become more intricate, we can rely less on a priori knowledge of the mechanisms and processes involved. To selectively excite molecules, for example, requires full knowledge of the molecular Hamiltonian as well as a precise analytical description of laboratory conditions. A new technique was thus developed wherein feedback from the experiment is used to obtain optimal input.[21] From these experimentally obtained inputs one can then attempt to decipher the underlying mechanics of the system under scrutiny.

Figure 6.1 gives an implementation of a closed loop experiment in the field of femtosecond lasers. An initial guess is made for a suitable pulse shape and applied using a spatial light modulator. This pulse is fed into the experiment — be it a second harmonic optimisation, pump-probe experiment, or coherent anti-Stokes Raman scattering microscopy — and a response is measured. Finally, the experimental result is sent to an algorithm which chooses new pulse shapes such that the measured signal is optimised. As the new pulse shapes are fed into the shaper, the circle is completed and we can speak of a closed-loop optimisation.

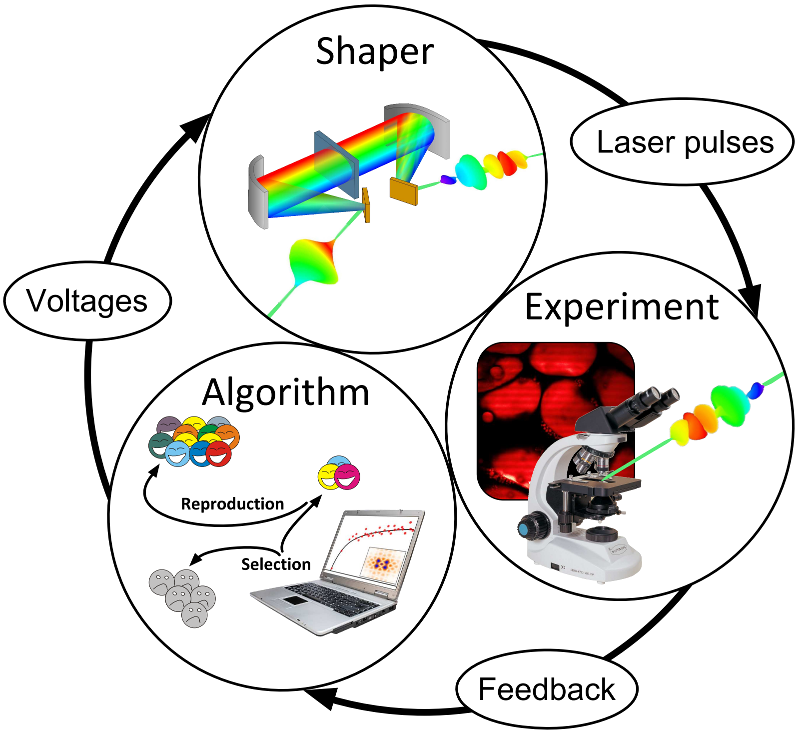

Figure 6.1: The concept behind the learning loop. Voltages from the evolutionary algorithm are turned into pulse shapes with the spatial light modulator. The experimental outcome then serves as basis for the algorithm to suggest improved pulse shapes.

6.1 Evolutionary algorithms

The pulse shaper setup has already been discussed in great detail in Chapter 4 and an application of spectral shaping was shown with the optimisation of Second Harmonic Generation (SHG). In this chapter we will look at the algorithm that steers this optimisation process.

A modern development in many-parameter optimisations is the field of evolutionary algorithms. While Barricelli ran the first evolutionary simulations in 1954,[22] it was not until the 1980's that evolutionary algorithms were used in real-world applications. Recent advances in computing power have fueled further research in evolutionary algorithms.

In all types of evolutionary algorithms, complex search landscapes are traversed stochastically using processes derived from biological evolution. Solutions are represented as a genetic code and in each iteration a population of solutions is evaluated. The best individuals are then selected and a new generation is formed after mutation and recombination of the genetic information.

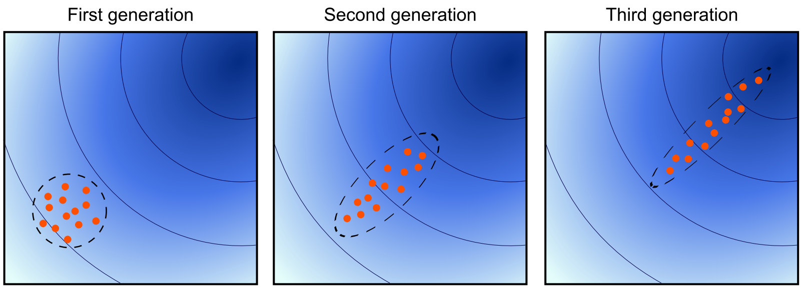

Figure 6.2: Concept behind the covariance matrix adaptation. As the generations develop, the search direction is adapted to the landscape.

While the basis of selection — survival of the fittest [23] — in biological systems is simultaneously the motivation for evolution, in the case of mathematical problems an artificial fitness function has to be defined as the objective of the optimisation. It is on the basis of this fitness function that the selection is made. Those individuals that are selected are then referred to as the *parents* with the next generation consisting of their *offspring*. In general, the fitness of the offspring will exceed that of the previous generation and the algorithm moves towards an optimum.

6.2 Covariance Matrix Adaptation

A very modern, robust, and fast implementation of an evolutionary algorithm is the covariance matrix adaptation (CMA) evolutionary strategy, developed by Nikolaus Hansen and coworkers.[24,25] While the name sounds frightful and the supporting mathematics is non-trivial, the key idea is relatively simple and easily explained in a low-dimensional search space.

The different parameters in an optimisation — such as the pixels of an SLM — are usually not independent. Instead of changing one parameter at a time, finding the optimal solution may require the optimisation of several parameters in unison. Those parameters are then covariant as they *vary co*llectively. If we represent this in a matrix and adapt unto an optimum, we have obtained the covariance matrix adaptation.

In two dimensions we can imagine this as shown in Fig. 6.2. Around an initial point a random set of points is sampled and the fitness values collected. Based on these points, a new solution is obtained, around which a new generation of individuals is randomly chosen. Now, the key element of the CMA strategy is that the direction of the search is also improved. By taking into account the path of the search, the slope of the search landscape can be determined. New points will then be sampled along this path, thereby greatly speeding up the optimisation routine. As the generations progress in Fig. 6.2, the distribution of individuals around the mean is oriented in the direction of the slope of the landscape.

All closed loop optimisations in this project were done using a CMA evolutionary strategy. The code is the same as used by Fanciulli et al.[26], with a more detailed description given in the manuscript of Savolainen.[3]

In the common (μ, λ) notation with the number of parents μ and number of offspring λ, the evolutionary strategy is implemented as a $(20, 40)$ system. In other words, from 40 individuals the best 20 are interbred to form the new generation. No elitism is present, so none of the parents survive the next generation and the optimisation routine can temporarily worsen its result as it moves out of local optima. The values for μ and λ were based on earlier work by Fanciulli et al.[26] wherein the population size (μ, λ) and initial mutation rate σ_0 were optimised to obtain the best and fastest convergence. It must be noted that these values are optimal for a given parameterisation and experiment. Changing the application of the algorithm could well influence the optimal population size or other settings of the algorithm.

6.3 Parameterisations

The CMA evolutionary algorithm spews out a list of numbers in exchange for fitness values of the previous generation. It is then up to the experimentalist to convert those numbers into a phase mask for the SLM. Although there is no direct link between these parameterisations and the algorithm itself, it is expected — and observed in the laboratory — that the choice of parameterisation will have an effect on the behaviour of the algorithm.

One possible choice for the parameterisation of the phase is a simple polynomial function. A sixth or seventh order polynomial, for example, usually approximates the phase distortions in the shaper setup to a good degree. With only eight parameters, this would thus seem an obvious, ideal, non-complex choice. However, it is practically impossible to get the algorithm to converge with this parameterisation. Due to the nature of the polynomial function, the parameterisation is highly sensitive to small changes in its parameters. Coupled with ever-present noise in the experiments, a more robust algorithm is required than is available. More complicated parameterisations will take longer to converge because of their higher dimensionality, but at least we are assured that they will converge to a local or global optimum.

We will distinguish two classes of parameterisations: direct and indirect parameterisations. The indirect parameterisations make use of basis functions to describe the phase. The Taylor expansion and Fourier transform are typical examples thereof. A more complicated version uses the Von Neumann representation[27*] and its applicability in the laboratory is currently being investigated. Indirect parameterisations were not used in this work, but they have the possible advantage of greatly reducing the dimension of the optimisation task if they can maintain robustness.

The direct parameterisations are those where the numbers from the algorithm are linearly scaled to obtain a phase pattern. Different forms of this class of parameterisations have been used in the experiments with which this manuscript concerns itself.

6.3.1 Free optimisation

The simplest parameterisation of a phase mask is of course a free optimisation wherein all pixels are optimised individually. The obvious drawback is the immense number of parameters of the optimisation problem this results in. In practice, a free optimisation can be used if the dimension is reduced to a more manageable degree. The algorithm works well with around 200 parameters given the inherent noise in the experiments. For a shaper with 1920 pixels in one dimension (we are assuming a one dimensional problem in this case), this would mean a binning of 10 pixels.

6.3.2 Wrapping

A key issue in finding a suitable parameterisation and adapting the algorithm to it, lies in the wrapping of the phase. On the SLM, the phase mask is displayed modulo 2π. The algorithm, however, has completely different, arbitrary restrictions as to the range of values that it can produce. A choice thus has to be made as to if, where, and how the phase is wrapped.

In the case of indirect parameterisations, the answer is clear: the phase is wrapped as the last step before applying it to the SLM. Such parameterisations naturally allow for an unlimited range of the phase and one can not establish a direct relation between the numbers of the algorithm and the wrapping of the phase. It is in the case of the direct parameterisations, which we use, that properly implementing the wrapping operators requires some thought.

The optimal phase routinely covers a range of 20π or more. In the case of a free optimisation, this would stipulate a very large search space. However, an initial guess of the phase must be restricted to a smaller range to ensure convergence; the algorithm will not be able to find

[*] The Von Neumann representation is named after the Hungarian mathematician John von Neumann. Von Neumann was one of the builders of the first computer: the ENIAC. He also constructed its successor at the Institute for Advanced Study in Princeton. It was with this machine that the Norwegian-Italian mathematician Nils Aall Barricelli used the first evolutionary algorithms.[22,28]

a search path if the initial step size is too big, as is the case for a random distribution in a 20π range.

It turns out that the solutions that are found will never span the entire 20π range. In fact, only about 2π will be used, regardless of the available range of values. It is significantly more difficult for the algorithm to find an unwrapped phase than to accommodate phase wrapping within the parameterisation. A local optimum is thus found. It must be noted that this is dependent on the behaviour at the boundaries as is discussed in Section 6.3.4.

For a good direct parameterisation, we must thus accommodate wrapping of the phase. In the case of the free optimisation, this is not the case. Wrapping occurs minimally over a range of 10 pixels in case of 10 pixel binning. We would thus like to improve the wrapping within the parameterisation.

6.3.3 Moving nodes

The free optimisation gives good results in the laboratory, so it forms a good starting point for an improved parameterisation. What we need for good wrapping, is to have two nodes — points where the phase is defined by the algorithm; in between nodes the phase is interpolated — on adjacent pixels of the SLM between which the phase jumps with 2π. We can do this by giving each node two parameters: one that describes its amplitude, and one that describes its position on the SLM. This allows two nodes to move to adjacent pixels if wrapping needs to occur.

Of course, this method doubles the number of parameters, complicating the search space significantly. However, we can make do with lower resolution (binning more pixels) as the nodes can move to better accommodate features in the solution. In this way, the dimension of the problem can be kept at a reasonable level.

The code used for the moving nodes parameterisation is given and explained in Appendix C.

6.3.4 Boundary behaviour

We have established that wrapping of the phase will occur, regardless of the available search space. It thus seems most reasonable to restrict the algorithm to a 2π range. This is precisely the implementation of Fanciulli et al.[26] The variations in the genetic codes are restricted to ensure that the solutions x_n from generation n are always in the interval $[lb, ub] = [0, 2\pi]$. If

x_{n+1} exceeds the lower or upper boundary, a new x'_{n+1} is chosen such that:

$$x'_{n+1} = \begin{cases} \mathscr{U}(lb, x_n) & \text{if} \quad x_{n+1} < lb \\ \mathscr{U}(x_n, ub) & \text{if} \quad x_{n+1} > ub \end{cases} \tag{6.1}$$

where $\mathscr{U}(a, b)$ denotes sampling uniformly from $[a, b]$.[3]

On second thought, this is not an ideal implementation. Going from 2π to 0 is now a large step for the algorithm, while the two values are equivalent in terms of phase. Ideally, the boundary conditions would thus be implemented as follows:

$$x'_{n+1} = \begin{cases} x_{n+1} + 2\pi & \text{if} \quad x_{n+1} < lb \\ x_{n+1} - 2\pi & \text{if} \quad x_{n+1} > ub \end{cases} \tag{6.2}$$

Unfortunately, this implementation is not trivial in CMA. The algorithm will remember the stepsizes taken in order to calculate the optimal search direction. Wrapping x_n at the boundaries would also require wrapping the associated stepsize. Implementing this in the algorithm was unfortunately outside the scope of this project. For future use of direct parameterisations, however, this could be a useful addition.

Alternatively, an attempt can be made to obtain an unwrapped phase, even with a direct parameterisation. Instead of rotating boundaries as in Eq. 6.2, additional parameters could be added to the parameterisation that add multiples of 2π at appropriate points in the phase. This would allow the algorithm to jump between different local optima that have different wrapping points, in search of the global unwrapped optimum. An actual implementation of this does not yet exist. Probably, two optimisations have to be run in parallel: one to find the optimal phase mask, and one to find the optimal wrapping.

One change had to be made to the boundary conditions to reduce the problems caused by having a large stepsize between 0 and 2π. The wrapping points are found early in the optimisation, both for the free optimisation, as well as with moving nodes. This means that later in the optimisation one might find the wrapping point to be off by a few pixels. If the phase is restricted to $[lb, ub] = [0, 2\pi]$, the algorithm will be unable to resolve this issue. By extending the boundaries slightly below zero and beyond 2π, this can be alleviated. While not as good a solution as the rotating boundaries, this does allow the algorithm to further improve its wrapping. In the end, the boundary conditions were completely removed, without detrimental effects to the efficiency of the algorithm. The nature of the direct parameterisations is such that there is a large hit in fitness for large deviations from the mean value. Solutions were never found outside $[-0.5\pi, 2.5\pi]$.

6.4 Simulated SHG optimisation

As an optimisation run in the lab can take up to three hours, it is easiest to test different parameterisations in simulations. This does require simple search landscapes where the analytical solution exists of the fitness function.

As a testbed, the SHG optimisation experiment discussed in Section 4.5 is used. The global optimum for this problem is well defined and earlier research has proven this problem to be a useful reference system for comparing evolutionary algorithms.[26] Our results will also be compared to the earlier findings by Fanciulli et al.

Two parameterisations were compared and the evolutionary algorithm was also adapted to the parameterisation in each case. Not much is yet known of the influence of the parameterisation on the efficiency of the algorithm. The work by Fanciulli et al. has shown that for a given parameterisation, there is a large dependence of the efficiency of the optimisation on the settings of the algorithm. Three control knobs are defined: the population and offspring sizes (μ, λ), and the initial stepsize σ_0. In the simulations presented here, an additional variable has been added in the form of the dimension *dim* of the search space.

Fanciulli et al. also compared two different evolutionary algorithms: the covariance matrix adaptation (CMA) and derandomized adaptation (DR2). In our work, only the CMA algorithm is used, as this turned out to be the most reliable algorithm.[26]

The following parameter values were combined in the simulations, wherein the number of pixels on the simulated phase mask was 1024.

$$\mu \in \{1, 5, 10, 20\},$$

$$\lambda \in \{5, 10, 20, 40, 80\},$$

$$\sigma_0 \in \{0.01, 0.05, 0.1, 0.2\},$$

$$dim \in \{32, 64, 96, 128, 160, 192, 224, 256\}$$

In each simulation, the number of function evaluations was kept constant at 20.000. This equates to keeping the number of measurements in a laboratory experiment constant. Large population sizes (high λ) thus went through fewer generations than small population sizes. Of course, only combinations where $\mu \leq \lambda$ have been simulated. Moreover, combinations of high

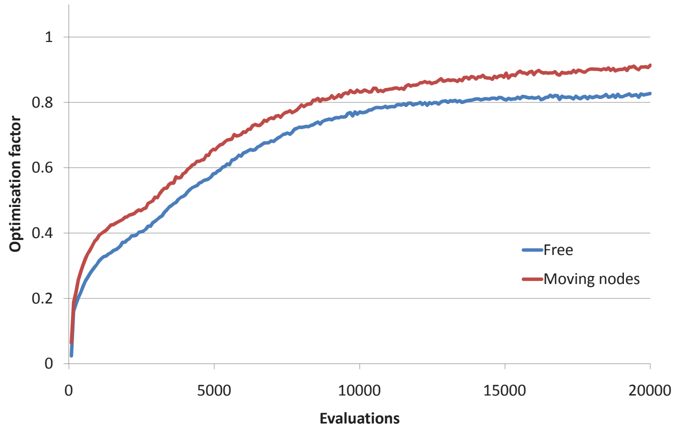

Figure 6.3: Simulated SHG optimalisation with 1 equal to a theoretical flat phase. The moving node parameterisation ($dim = 96$) shows a 10% improvement over the free optimisation ($dim = 128$).

dimensionality (dim) and many generations (low λ) could not be computed due to memory problems as the CMA algorithm stores covariance matrices for each dimension and each generation. With a few thousand generations (e.g. 20.000 evaluations divided over 5 individuals) and 256 nodes per individual, this results in a million covariance matrices to be stored at once.

To simulate laboratory conditions, 5% noise is added, as in the experiments by Fanciulli and co-workers. In contrast with that work, however, our simulations used a positively correlated fitness, with a fitness of 1 corresponding to an ideal flat phase. In the simulations by Fanciulli et al. the fitness is defined as the inverse of this value and their optimisation is actually a minimisation problem. For the evolutionary algorithm this is not of any importance.

6.4.1 Simulation results

The best optimisations after 20.000 evaluations are shown in Fig. 6.3 for the free parameterisation and the moving nodes parameterisation. What is plotted in the graphs is the mean fitness of all individuals in each generation, averaged over 10 optimisation runs. While the general behaviour during the optimisation is similar for both parameterisations, the moving nodes show a 10% improvement over the free parameterisation.

In Fig. 6.3, unity represents the theoretical flat phase. This is not achievable due to the finite dimensionality of the optimisation, the inherent imperfection of wrapping, and of course the artificially introduced noise. However, it is clear that the advanced phase wrapping of the moving nodes parameterisation does offer a significant improvement in SHG optimisation.

We can also plot the fitness value after 20.000 evaluations as a function of the four “control

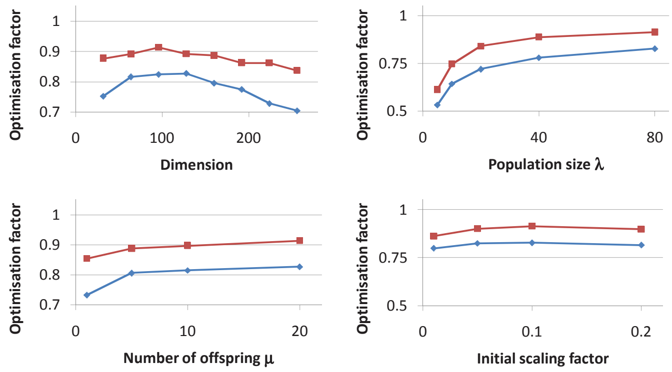

Figure 6.4: Simulated SHG optimalisation with a free optimisation (blue) and moving nodes (red). The four graphs show the dependence on the four "control knobs" of the algorithm.

knobs" of the algorithm: dimensionality, (μ, λ), and the initial scaling factor. Fig. 6.4 compares the two different parameterisations as a function of these four variables.

From Fig. 6.4, it is clear that the optimal number of points in a parameterisation is not equal to the number of pixels of the SLM. This is due to the finite number of evaluations in each optimisation. Low dimensional search spaces are optimised faster, but end at lower optimisation factors. High dimensional search spaces, on the other hand, are slow in improving but will ultimately find a better phase mask. Given the time constraint of 20.000 evaluations, a dimensionality of 96–128 is optimal for both parameterisations.

This optimal dimensionality should be compared not only to the size of the phase mask (1024 pixels), but also to the complexity of the phase distortions. In particular, the number of wrapping points in the final phase mask may have a large influence on the optimal number of control points in the parameterisation. If there are many wrapping points, then a high number of points will be required to follow this complicated phase mask. If very little wrapping occurs, then a lower dimensionality will greatly improve the speed of the optimisation without compromising the final fitness.

Overall, Fig. 6.4 shows that the behaviour of the algorithm does not greatly depend on the parameterisation used. Of course, this is partly because the two chosen parameterisations are both direct parameterisations that have largely the same characteristics. Nevertheless, it is clear that the moving nodes parameterisation always shows an improvement over the free parameterisation used by Fanciulli et al.

The paper by Fanciulli et al.[26] can not be directly compared to these results as there is a non-linear scaling between the fitness values of these two projects. However, we can compare our results with the conclusions drawn by the authors.

For the CMA algorithm, Fanciulli et al. find an optimal population size of $(\mu, \lambda) = (20, 40)$. Our simulations suggest that a larger population size leads to better results. The optimisation efficiency is only to some extent dependent on μ, but the best results are clearly found for $\lambda = 80$.

Key here is the different number of evaluations used in this work and that of Fanciulli et al. They stopped after 5.000 evaluations, while our simulations ran for 20.000 evaluations. If we analyse the optimisation at different numbers of evaluations, then we see the optimal λ shift to lower values. Below 12.000 evaluations, we also find the best optimisation for $\lambda = 40$. If the number of evaluations is further restricted to 4.000, then the best optimisation is found with $\lambda = 20$. It appears that given a low number of evaluations, it is more effective to roughly map out the search landscape by going through as many generations as possible, than it is to accurately search locally for optima by having many randomly distributed individuals in each generation.

6.5 Experimental SHG optimisation

Using a tried and proven[3] pulse shaping setup (PUSH lab) with a one dimensional pulse shaper from CRI, several actual SHG optimisation experiments were performed. The shaper is a 640 pixel LC SLM that can do both phase and amplitude shaping. In this case, only 360 were used for pure phase shaping.

The two different parameterisations were tested for $dim \in \{64, 128, 192\}$, with the other variables of the algorithm fixed at $(\mu, \lambda) = (20, 40)$ and $\sigma_0 = 0.1$. The optimisations were run for 250 generations or, equivalently, 10.000 fitness evaluations. As each optimisation took 2.5 hours, every experiment was repeated only twice.

The results of these experiments are shown in Fig. 6.5. From these curves we can immediately conclude two things. First of all, the influence of laboratory noise is much greater than accounted for in the simulations. Secondly, the moving node parameterisation gives on average a 30% improvement over the free parameterisation. This is significantly more than predicted by the simulations.

The subtle dependence of optimisation efficiency on the dimension of the parameterisation (as shown in Fig. 6.4) is lost in the noise in laboratory experiments. However, it must be noted that a direct comparison between the simulations and experiments is not completely fair. The simulated phase mask consisted of 1024 pixels, while only 360 pixels of the shaper in the PUSH lab were used for the SHG optimisations. The ratio of the optimisation dimen-

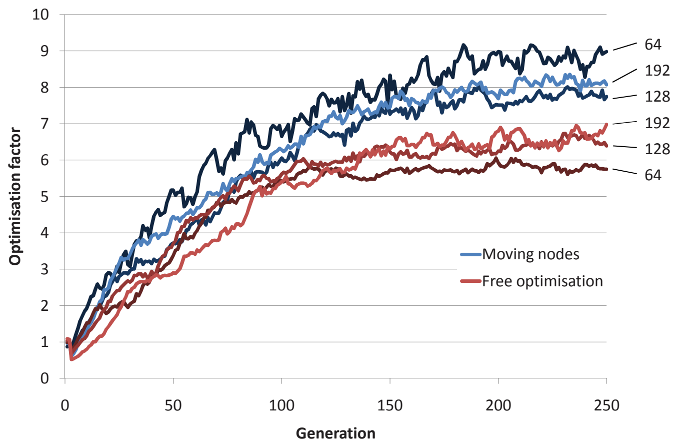

Figure 6.5: SHG optimisations measured in PUSH lab for two different parameterisations and different dimensionality.

sion to the size of the SLM is thus off by a factor of three. However, a simple linear scaling between the experiments and simulations will not suffice, as the complexity of the induced phase distortions is independent of the size of the SLM. The parameterisation must have sufficient points to accommodate the overall shape of these phase distortions, as well as all the wrapping points.

7

Conclusions

A two dimensional spatial light modulator was bought in the form of the Holoeye HEO1080P Phase Only Modulator. In both the spatial and spectral domain novel calibration and characterisation techniques have been developed and demonstrated. The spectral shaping capabilities have been tested through a second harmonic generation in a closed loop experiment, driven by a CMA evolutionary strategy.

A superpixel-by-superpixel characterisation of the SLM is performed, showing positional dependence of the maximum phase retardance of the SLM. Use of the SLM is further complicated by strong temporal fluctuations in the shaping effectiveness. In later experiments, this is combined with the instability of the laser, and other noise sources. The need for robust optimisation algorithms and good normalisation of experiments is clear. The SHG optimisation performed in this work shows that, despite the limitations of the SLM, good and consistent control over the pulse shape can be obtained.

The influence of different parameterisations of phase profiles on the efficiency of evolutionary algorithms has been investigated. Using SHG optimisations as a testbed, both simulations and laboratory experiments have shown that the optimisation efficiency can be improved up to 30% by replacing the common free optimisation with a novel moving node approach.

For direct parameterisations, the efficiency of the CMA evolutionary strategy is mostly invariant under change of the population dynamics parameters (μ, λ) and the initial stepsize σ_0. Simulations do show some dependence on these parameters, but this is lost in the noise inherent to laboratory work. Future research will extend this analysis to include other parameterisations, including indirect parameterisations using more complex basis functions. It is expected that the indirect parameterisations will require significant optimisation of the different variables of the evolutionary strategy.

Finally, a new type of FROG setup is developed as proof-of-principle of spatio-temporal shaping. Offering in situ pulse characterisation with just one beam and no moving parts, this FROG method allows quick and easy characterisation of complicated pulses at the precise experimental conditions with no further alignment necessary.

The possibility of both FROG and X-FROG has been demonstrated, and it has been shown that complicated pulse shapes can be simultaneously generated and characterised using a single two dimensional SLM.

Acknowledgements

Seven years of study have made me a cynic. When I began in the fall of 2001, I had a romantic view of the world in which scientists discovered great truths about the fundamental principles of the universe and were driven purely by the beauty of science. I have since come to realise that this is not how science is done. Publications and impact factors are the motive forces behind all research, and students work only as hard as a minimal passing grade requires rather than being enamoured by the Physics they are taught.

I am now a full member of this scientific community as I too get excited at the sight of my Hirsch index going up. It is really my romantic view that is misguided, for when we consider the great natural philosophers of old it is easy to forget that they too were subject to politics and intrigue in securing funding, and ultimately ranked by the grandeur of their publications. In several hundred years, who will be remembered as the Newton, Aristotle, or Einstein of today, and who will be forgotten as so many "lesser" scientists have been? No matter how insignificant their contribution or small their fame, every scientist forms an integral part of the scientific community as they add — in their own way — to our understanding of the universe. I merely ask that they spell-check their paper's title before publication.

Spending a year with the Optical Sciences group has, however, restored my love for science. Jennifer Herek has created an excellent environment for young minds to develop and find their own way in the world of ultrafast optics. The hands-off approach to supervising my project suited us both, and I am grateful for the opportunity given to and the trust placed in me as I went my own way. My brief status reports — few and far between — seemed enough to give us both the assurance that the project was progressing properly.

Ali, Di, and Janne deserve a special word of thanks. They have been a great help in facilitating the laboratory work and helping me with my first steps in a proper optical laboratory.

The invaluable assistance of Frans and Jeroen in this can of course also not be left without mention.

All my Optical Sciences colleagues deserve heartfelt thanks for providing a wonderful place to work on my Master's thesis. Pancake brunches, group retreats, and daily shopping trips, but also many insightful discussions, quick tips and troubleshooting, a comfortable open door policy, and a keen eye for design and presentation skills; I can not imagine a better environment in which to conclude my studies.

This thesis is the culmination of seven years at the University of Twente. This time has not been spent solely in class-rooms and laboratories. Some thought must therefore also be given to all those whose friendship, company, and support have been a continuous source of joy during these years. This thought is given with ease and pleasure, and I can only wonder where I grew more as a person: within the lecture halls of the university, or outside on the open water.

Finally, I owe my family great gratitude for supporting me in many ways. Even when two continents separate us, you are always in my heart and it is comforting to know that I am in yours.

A SLM manufacturers

This list contains contact details of the SLM manufacturers that were contacted during this project.

Holoeye Photonics
http://www.holoeye.com

Stefan Osten - R&D Germany
Technical contact
Telephone: +49 (0)30 6392 3661
E-mail: stefan.osten@holoeye.com

Hamamatsu

http://www.hamamatsu.com

Arie van Gool - Netherlands office
Sales contact

Telephone:	+31 (0)36 538 2123
E-mail:	avangool@hamamatsu.de

Ludwig Schleinkofer - Director R&D Marketing
Technical contact Europe

Telephone:	+49 (0)81 5237 5100
E-mail:	lschleinkofer@hamamatsu.de

Yuji Kobayashi - R&D Japan
Technical contact Japan

Telephone:	+81 53 434 3311
E-mail:	yujikoba@etd.hpk.co.jp

Boulder Nonlinear Systems

http://www.bnonlinear.com

Pieter Kramer - Laser 2000 Benelux
Sales and technical contact

Telephone:	+31 (0)297 266 191
E-mail:	pkramer@laser2000.nl

B Fringe shift calibration

On several occasions, calibrations were done using shifting interference patterns. This chapter explains the methods used.

Using the setup in Fig. B.1, an interference pattern is made from two overlapping beams. This interference pattern is recorded on a CCD in the form of a simple webcam, and the 2D image is reduced to a one dimensional fringe pattern (Fig. B.2). If a phase delay is applied to one of the two beams, the interference pattern shifts. Given a good gamma curve (Section 3.3), the fringes will move linearly for an increasing phase delay, shown in Fig. B.3. For each phase delay — that is, for each vertical line in Fig. B.3 — the local minima of the interference pattern are estimated. By fitting second order polynomials around these points, the precise location of each minimum is found. In Fig. B.4, these minima are

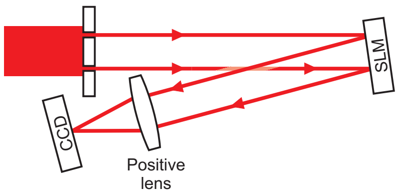

Figure B.1: Interference of two parallel beams. The lens overlaps the two beams on the CCD where an interference pattern is observed. To one beam a variable phase delay is applied.

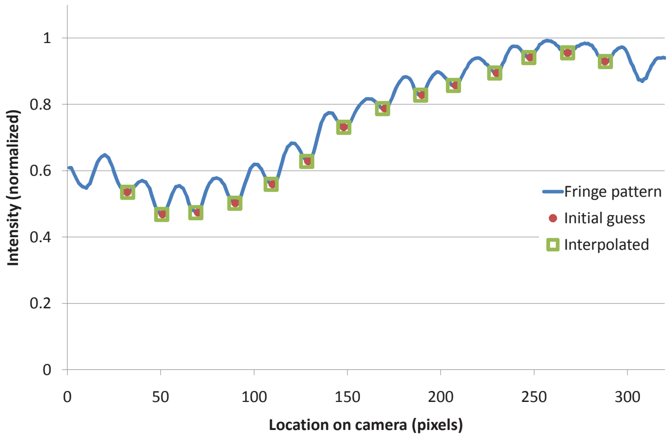

Figure B.2: Fringe pattern for one phase delay (blue line) with the minima indicated. An initial estimate is made of the minima locations (red circles). A second order polynomial is fit at these minima to obtain an improved, interpolated location of the minima (green squares).

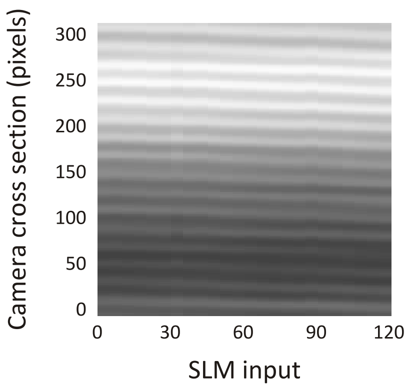

Figure B.3: Shifting interference pattern as a function of phase delay in one arm of the interferometric setup of Fig. B.1.

plotted as a function of the SLM input applied to one of the beams.

A robust algorithm finds adjacent points in Fig. B.4, even if datapoints are missing. Through these sets of points linear fits are made. The spacing between the lines then gives the shift associated with a 2π phase shift. The slope of the lines (in pixels per input value) allows us to determine the relation between the SLM input (the colour value sent to the SLM in the range 0–255) and the resulting phase delay.

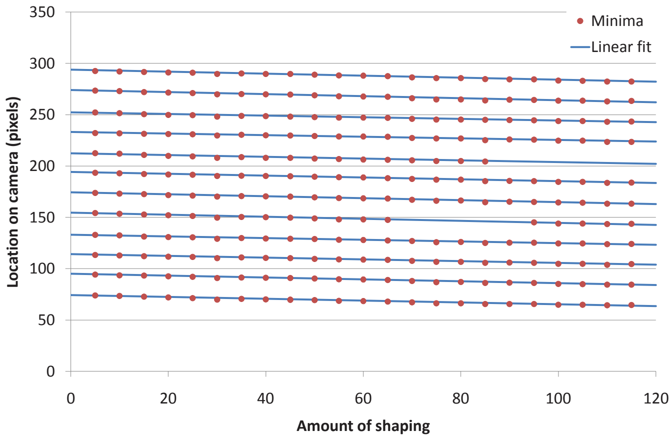

Figure B.4: Local minima of Fig. B.3. A linear fit is made through adjacent points. The algorithm can overcome missing points by extrapolating from previous points. Datapoints have been intentionally removed for this example.

C Moving nodes parameterisation

The moving nodes parameterisation (see Section 6.3.3) uses two genes to describe each point of a phase mask. One number gives the phase, while another defines the position of the node on the phase mask. The MATLAB code for this parameterisation is given on the next page.

The shaper is divided into N segments for a total of N+2 nodes. Two nodes are fixed at the first and last pixel of the SLM. There is some overlap between the segments to allow fine-tuning of the wrapping positions. Because the segments overlap there is the possibility of having two nodes at the same position. This is solved by shifting the overlapping nodes by one pixel. Finally, the nodes are sorted by position, and a linear interpolation finds the complete phase mask for the SLM.

```
function phase_mask = movingnodes(genes);

%Split up genes in phase and position
phase = genes(1:(end/2+1))';
position = genes((end/2+2):end)';

%Divide available space (1024px) over number of nodes
N = length(genes)/2-1;
xb = floor(linspace(2,1023,N+1));
xs = [xb(3:end)-xb(1:end-2), xb(end)-xb(end-1)] - 1 - eps;
xb = xb(1:end-1);

%Calculate position of each node
xb = xb + floor(xs.*position);

%Deal with overlapping nodes
q = find(xb(2:end)==xb(1:end-1))+1;
xb(q) = xb(q)+1;
xb = [1, xb, 1024];

%Sort nodes for interpolation
[xb,b]= sort(xb);
phase = phase(b);

%Interpolate linearly between nodes
y1 = 1*interp1(xb,phase,1:1024,'linear');

%Return final phase
phase_mask = y1 * 2*pi;
```

Bibliography

[1] Rabitz, H., de Vivie-Riedle, R., Motzkus, M., & Kompa, K. **2000**. Whither the future of controlling quantum phenomena? *Science*, *288*(5467), 824–828.

[2] Patrice, T. (Ed.) **2003**. *Comprehensive Series in Photochemistry & Photobiology: Photodynamic Therapy*, vol. 2. The Royal Society of Chemistry.

[3] Savolainen, J. **2008**. *Coherent Control of Biomolecules*. Ph.D. thesis, University of Twente.

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

[5] Vellekoop, I. M. **2008**. *Controlling the propagation of light in disordered scattering media*. Ph.D. thesis, University of Twente.

[6] Feurer, T., Vaughan, J. C., Koehl, R. M., & Nelson, K. A. **2002**. Multidimensional control of femtosecond pulses by use of a programmable liquid-crystal matrix. *Optics Letters*, *27*(8), 652–654.

[7] Vaughan, J. C. **2005**. *Two-Dimensional Ultrafast Pulse Shaping and its Application to Coherent Control and Spectroscopy*. Ph.D. thesis, Massachusetts Institute of Technology.

[8] Gundogdu, K., Stone, K. W., Turner, D. B., & Nelson, K. A. **2007**. Multidimensional coherent spectroscopy made easy. *Chemical Physics*, *341*(1–3), 89–94.

[9] Frumker, E., & Silberberg, Y. **2007**. Phase and amplitude pulse shaping with two-dimensional phase-only spatial light modulators. *Journal of the Optical Society of America B: Optical Physics*, *24*(12), 2940–2947.

[10] Frumker, E., & Silberberg, Y. **2007**. Femtosecond pulse shaping using a two-dimensional liquid-crystal spatial light modulator. *Optics Letters*, *32*(11), 1384–1386.

[11] Herek, J. L., Offerhaus, H. L., Bäck, T., Otto, C., & Mosk, A. P. **2008**. *A Smart CARS Microscope*. STW Perspective Smart Optics Systems program project proposal.

[12] Postma, S., van Rhijn, A. C. W., Korterik, J. P., Gross, P., Herek, J. L., & Offerhaus, H. L. **2008**. Application of spectral phase shaping to high resolution CARS spectroscopy. *Optics Express*, *16*(11), 7985–7996.

[13] Chimento, P. F. **2007**. *Orbital angular momentum mode analysis with a micro mirror array*. Report, University of Twente.

[14] Dufresne, E. R., Spalding, G. C., Dearing, M. T., Sheets, S. A., & Grier, D. G. **2001**. Computer-generated holographic optical tweezer arrays. *Review of Scientific Instruments*, *72*, 1810.

[15] Jesacher, A., Maurer, C., Furhapter, S., Schwaighofer, A., Bernet, S., & Ritsch-Marte, M. **2008**. Optical tweezers of programmable shape with transverse scattering forces. *Optics Communications*, *281*(8), 2207–2212.

[16] Faulkner, H. M. L., & Rodenburg, J. M. **2004**. Movable aperture lensless transmission microscopy: A novel phase retrieval algorithm. *Physical Review Letters*, *93*(2), 023903.

[17] Pasienski, M., & DeMarco, B. **2008**. A high-accuracy algorithm for designing arbitrary holographic atom traps. *Optics Express*, *16*(3), 2176–2190.

[18] Franken, P. A., Hill, A. E., Peters, C. W., & Weinreich, G. **1961**. Generation of optical harmonics. *Physical Review Letters*, *7*(4), 118–119.

[19] Bloembergen, N. **2000**. Nonlinear optics: past, present, and future. *IEEE Journal on Selected Topics in Quantum Electronics*, *6*(6), 876–880.

[20] Trebino, R. **2000**. *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses*. Norwell: Kluwer Academic Publishers.

[21] Judson, R. S., & Rabitz, H. **1992**. Teaching lasers to control molecules. *Physical Review Letters*, *68*(10), 1500–1503.

[22] Barricelli, N. A. **1954**. Esempi numerici di processi di evoluzione. *Methodos*, 45–68.

[23] Darwin, C. R. **1872**. *On the Origin of Species*. London: John Murray, 6^{th} ed.

[24] Hansen, N., & Ostermeier, A. **1996**. Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. *Proceedings of IEEE International Conference on Evolutionary Computation*, 312–317.

[25] Hansen, N., Muller, S. D., & Koumoutsakos, P. **2003**. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation*, *11*(1), 1–18.

[26] Fanciulli, R., Willmes, L., Savolainen, J., van der Walle, P., Bäck, T., & Herek, J. L. **2007**. Evolution strategies for laser pulse compression. *Lecture Notes in Computer Science*, *4926*, 219–230.

[27] Fechner, S., Dimler, F., Brixner, T., Gerber, G., & Tannor, D. J. **2007**. The von Neumann picture: a new representation for ultrashort laser pulses. *Optics Express*, *15*(23), 15387–15401.

[28] Dyson, G. **2007**. Barricelli's universe. *Make*, *06*, 190–192.