Absolute Heterodyne Interferometer for Strongly Aspherical Mirrors

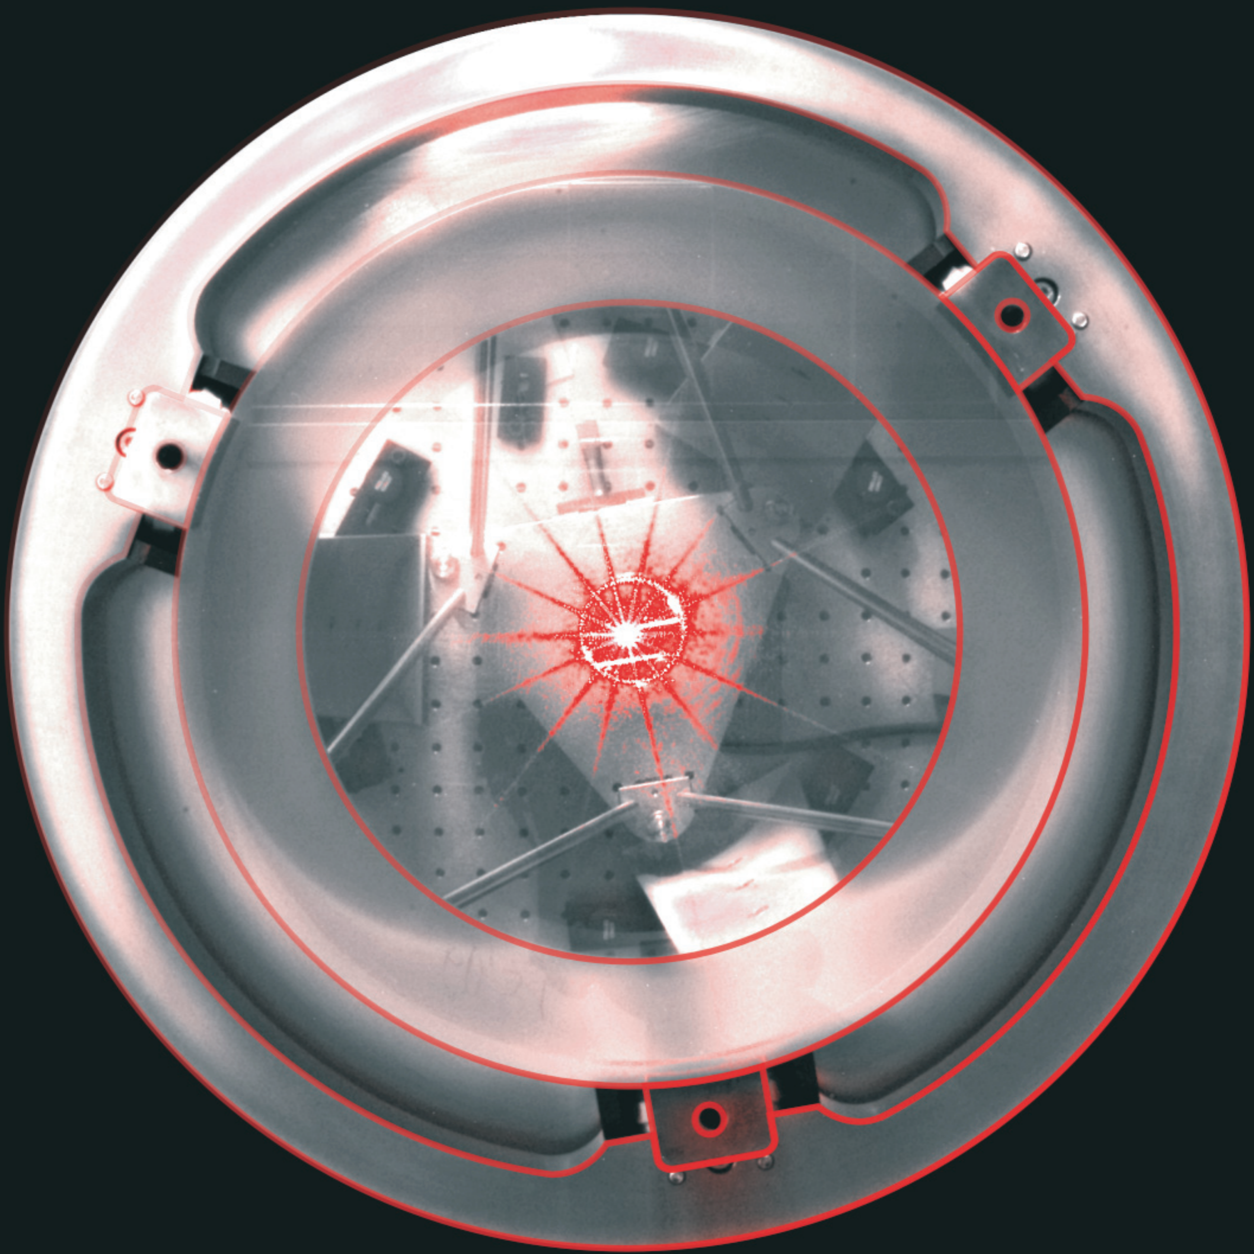

Max Lukas Krieg

Absolute Heterodyne Interferometer for Strongly Aspherical Mirrors

Absolute Heterodyne Interferometer for Strongly Aspherical Mirrors

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. J.T. Fokkema,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen
op dinsdag 30 november 2004 om 15.30 uur

door

Max Lukas KRIEG,

Master of Science in Physics, University of Auckland (New Zealand)
geboren te Jegensdorf, Zwitzerland

Dit proefschrift is goedgekeurd door de promotor,
Prof. dr. ir. J.J.M. Braat.

Samenstelling promotiecommissie:

Rector Magnificus	Voorzitter
Prof. dr. ir. J.J.M. Braat	Technische Universiteit Delft, promotor
Prof. dr. ir. P. Urbach	Technische Universiteit Delft
Prof. dr. ir. P.M.J. van den Hof	Technische Universiteit Delft
Prof. dr. ir. P. French	Technische Universiteit Delft
Prof. W. Osten	Universität Stuttgart
Dr. A. den Boef	ASML, Veldhoven
Dr. G. Seitz	Carl Zeiss, Oberkochen, Duitsland
Prof. dr. ir. L. J. van Vliet	Technische Universiteit Delft, reservelid

ISBN 90-901-8783-9

This work was supported by the Dutch Technology Foundation (STW), ASML (Veldhoven, The Netherlands), Carl Zeiss (Oberkochen, Germany), and the Dutch Organization for Applied Scientific Research (TNO–TPD).

Printed in The Netherlands.

Contents

Preface

This thesis discusses the design, construction and testing of a novel interferometer to measure, with extreme accuracy, the low spatial frequency shape of extreme ultraviolet lithography (EUVL) mirror substrates. These mirrors will in turn be used to manufacture integrated circuits with unprecedented feature sizes smaller than 50nm.

The work continues on from the initial research of Klaver[1] at the TU Delft from 1995 to 2000, where he laid the foundation for the development of an interferometer that can claim to yield accurate absolute measurement results, in contrast with the vast majority of interferometers that only produce relative measurement results. While great care was taken to maintain general applicability to most types of potential EUVL mirror substrates that may require measurement, several practical concessions were made in the design of the prototype discussed here to enable proof-of-principle experiments to be carried out on one particular EUVL mirror candidate.

To a large part, the conclusions reached by Klaver will be used as starting points for the practical implementation of the interferometer. However, a number of significant changes in the principle of the interferometer, especially the light-source, will require a renewed treatment of subjects already covered in the preceding work. Also, many of the more general points will be developed to specifics which can then be applied directly to the design and construction of the device.

The structure of this thesis is as follows: Chapter 1 will contain a brief introduction to the field of extreme ultraviolet lithography, serving as motivating factor for the research undertaken. Key challenges will be mentioned qualitatively, and the approach to the problem will be sketched. Chapter 2 deals with the concept of the interferometer, reviewing several earlier concepts and conclusions, and presenting the specific measurement approach culminating in the design of the prototype interferometer. Chapter 3 is devoted to the theoretical challenge of interpreting the raw measurement data to yield a useful measurement of the mirror substrate’s shape. Largely mathematical in nature, it develops- and compares the performance of- several inverse propagation algorithms (IPAs) to fulfill this end. Chapter 4 treats the design and construction of the multiple wavelength light source required by the interferometer. Chapter 5 deals with several types of sensors which could be used in the interferometer. The various methods required to acquire and subsequently process measurements with these sensors will be discussed. Chapter 6 focuses on the design and construction of the interferometer itself, including the frame which connects the mirror substrate under test with the other interferometer components, and includes a discussion of the optical fibers used in the interferometer. Chapter 7 gives the details and results of measurements on a particular mirror substrate using our instrument. These results are discussed in Chapter 8, together with other, more general considerations and suggestions for improvements and future work. Finally, the appendix contains specifications of materials and equipment used.

[1] R. G. Klaver, "Novel interferometer to measure the figure of strongly aspherical mirrors." Delft: Delft University of Technology, 2001.

1 Introduction

The interferometer discussed in this thesis is intended as a measurement tool for aspheric, concave mirrors to be used in extreme ultra-violet lithography (EUVL). To achieve this goal, it will need to operate at the limits of accuracy for an instrument of its class, outperforming the current state of the art.

1.1 The need for an Ångström accuracy interferometer

The research presented here is being supported by a coalition of semiconductor industry- and optical fabrication representatives, as well as governmental bodies. The reason for this broad backing is the fact that the accuracy with which projection optics can currently be measured forms a bottle-neck for the continued development of semiconductor devices.

The semiconductor industry is currently following a self-fulfilling prophecy known as "Moore's Law". The law, originally stated in 1965[1] as merely an extrapolation of observed trends, predicts that the number of components per micro-chip will double every year (which was later adjusted to every two years). Ever since its stipulation, manufacturers have strived to follow or even beat this law in an effort to remain competitive.

To facilitate this increasing feature density on silicon wafers, the size of the patterns to be written has had to decrease, until the limits of the processes used were reached. This limit is largely optical in nature, since the resolution of optical lithography processes is closely linked with the writing wavelength:

$$R = k_1 \frac{\lambda}{NA} \tag{1.1}$$

Where *NA* stands for the Numerical Aperture, and k_1, referred to as the "process factor", is an empirical coefficient which depends on a number of technology-specific factors.

Once the numerical aperture and process factor have been optimized, the only way to maintain the trend of decreasing feature size is to decrease the wavelength of the light used in the lithography process. This has already occurred a number of times since 1975, moving from 436nm to 365nm by 1985, to 248nm by 1995, and to 193nm in 2002. EUV lithography represents the jump to 13nm.

While alternative methods to classical lithography are actively being pursued[2], prototype EUVL systems have recently become operational[3] (Figure 1-1), and continue to be developed.

> *"[…] a breakthrough lithography technology currently under development, will become the volume production standard. Known as Extreme Ultraviolet (EUV) lithography, this technology uses reflected rather than directly transmitted light which allows the patterning of lines smaller than 50 nm. […] Intel anticipates building processors using EUV technology in the second half of the decade."*
>
> -Intel Fall report[4], 2002

As mentioned in the quote above, the move to EUVL also means a move to reflective projection optics, as opposed to the refractive, lens-type projection optics used to date, the reason being that most materials become highly absorptive at the EUVL wavelength. These projection optics require large NA aspheric mirrors with unprecedented figure accuracies.

Mirror asphericity is desirable because it enables the design of optical systems with considerably fewer aberrations, compared to a system using a similar number of spherical optics instead. The reason for larger numerical apertures should be clear from (1.1), while the stringent figure accuracy requirements are a consequence of the very short wavelength used. In order to perform well in an optical sense – that is, to produce focal spots or images which are diffraction limited – the root mean square (rms) figure aberrations must be well below $\lambda/14$ (Maréchal condition), while in practise $\lambda/50$ is often strived for. For EUV optics, this translates to sub-nm accuracy requirements.

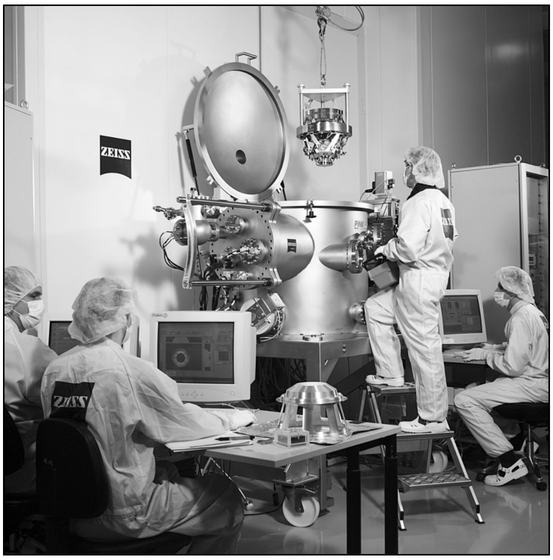

Figure 1-1 - Final assembly of EUVL illumination system, courtesy of Carl Zeiss SMT.

A number of other areas of development share the need for large aspheric reflectors with sub-nm figure accuracies. These include microscopy in the "water window"[5] (2.3nm-4.4nm), X-ray astronomy (5nm-31nm), spectroscopy and plasma diagnostics[6].

There are already a number of methods available which can measure the mid- and high- spatial frequencies of EUVL mirrors[7] ($1\mu m^{-1}$-$1mm^{-1}$ – phase shifting interferometric microscopy and $1\mu m^{-1}$-$50\mu m^{-1}$ – atomic force microscopy, respectively), but a number of challenges have hampered efforts to meet the low spatial frequency requirements (from 1/[clear aperture] to $1mm^{-1}$).

1.2 The state of the art

Besides the interferometer described here, there are four other types of metrology systems under development which have reached, or could potentially reach the accuracy requirements for EUVL projection optics. After a brief review, including an up-to date evaluation of their performance, we will state the place of our interferometer in the context of these methods.

1.2.1 At-wavelength phase-shifting point-diffraction interferometry.

The phase-shifting point-diffraction interferometer (PSPDI) by Medecki[8], has recently been extended to perform interferometry at the actual EUV wavelength of 13.4nm by Goldberg and Naulleau[9-14].

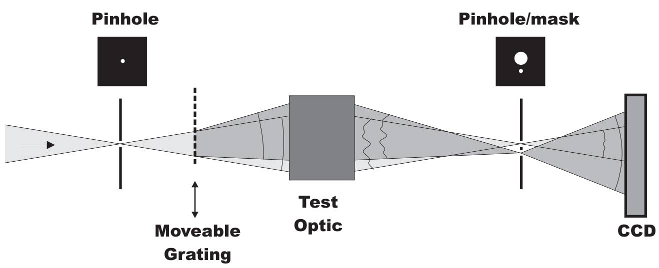

Figure 1-2 – Schematic diagram of the working principle of a PSPDI.

There are several configurations of this interferometer, one of which is shown in Figure 1-2. The method uses a pinhole and grating to generate two spherical, angularly sheared wavefronts, which are aberrated by the test optic. One of these is subsequently spatially filtered to a spherical wavefront by another pinhole, while the other is simply transmitted through the mask. The aberrated and spherical wavefronts then interfere at the CCD detector, where the resulting interference pattern is measured. By moving the grating, a phase-shift can be introduced between the measurement and reference beams, allowing the use of phase-shifting interferometry (PSI) to accurately determine the optical phase difference between the two beams.

There are a number of obvious limitations and sources of error in this interferometer, which we will discuss shortly. Nonetheless, this is currently the most accurate method for measuring EUVL reflection optics available, with repeatabilities of 0.006nm and accuracies of 0.02nm being reported from comparisons with other methods[9, 15]. This accuracy is mainly due to the short metrology wavelength being used.

Unfortunately, using a shorter wavelength also means that the surface can not deviate significantly from the reference wavefront without inducing excessive fringe-densities at the detector. To circumvent this short-coming, severe aspheres are measured in small “patches”, which are subsequently “stitched” together to give the complete surface form. Given a sufficiently generous overlap between patches, this method is very effective.

Clearly, the presence of the grating will introduce some wavefront aberrations, although these are somewhat averaged out by the process of stepping the grating during phase-shifting. There is also a trade-off between the efficiency of the system and the reference wavefront sphericity, as smaller pinholes (creating more spherical wavefronts) mean less transmission.

This method is unsuitable for optical shop testing, since it requires a coherent EUV lightsource. Currently, the only realistic sources are free electron lasers (FELs) and synchrotron radiation, both of which require large particle accelerators.

1.2.2 Fizeau interferometry.

Carl Zeiss has reported surprisingly good results with a Fizeau-type interferometer, despite the large number of optical components contained therein[16]. Repeatability is stated as 0.07nm, and the accuracy is claimed to be below 0.15nm for large NA EUV optics[17].

The interferometer, shown schematically in Figure 1-3, is of the Fizeau type with compensation optics, where the Fizeau plate is tilted to introduce a spatial carrier frequency in the interference pattern, allowing the local phase to be retrieved from a single image with methods similar to Takeda’s Fourier transform technique[18]. It is operated at visible wavelengths.

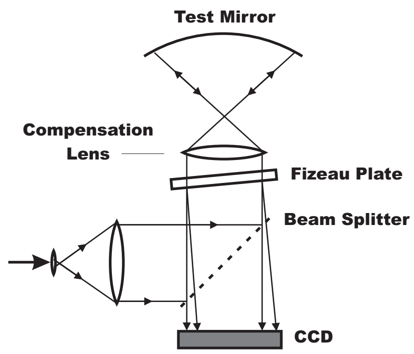

Figure 1-3 – Carl Zeiss Fizeau interferometer with multi-fringe DMI.

The large number of optical components in the beam path raises several issues. In order to give reliable results, the aberrations introduced by these optics must be below the accuracy requirements, or at least known with similar accuracy, so that they may be subtracted from the final measurements. Great care has been taken to calibrate this interferometer with reference sphere mirrors and rotation of internal optics to average non-rotationally symmetric aberrations. Even so, such optics are prohibitively expensive to produce for a regular optical shop instrument, requiring regular and extensive calibration procedures due to ageing.

1.2.3 At-wavelength Hartmann sensor.

Mercère has recently reported on a Hartmann sensor which is capable of measuring the wavefront from EUVL projection optics at operational wavelengths with 0.11nm accuracy[19].

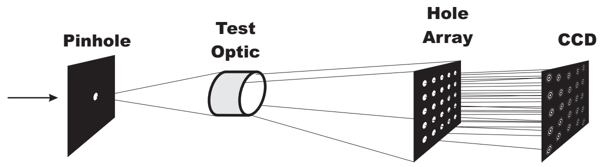

Figure 1-4 – Hartmann sensor

The principle of a Hartmann sensor is illustrated in Figure 1-4. The local gradient of a wavefront is measured at the pinhole array position by measuring the displacement of the spots projected onto the CCD.

Once again, this is a method which profits from the short metrology wavelength. Although, in contrast to the PSPDI mentioned above, the approach does not require a coherent source, it does require a high intensity source at around 13.4nm, because of the low detector efficiency at these wavelengths and the severe losses at the pinhole array. The requirements on the lightsource are therefore what make this method unsuitable for optical shop testing.

1.2.4 Sommargren interferometer

The Sommargren interferometer[20-23] is, like our interferometer, a fiber based instrument, operating in the visible. Various configurations exist, some using two fibers, but we will focus on the one-fiber configuration which has recently been reported to achieve an accuracy of 0.25nm[22], shown in Figure 1-5. Improvements to bring this figure down to a projected 0.089nm are underway.

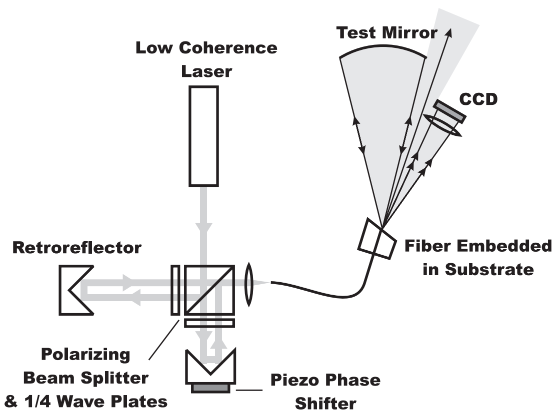

Figure 1-5 – Schematic diagram of one arrangement of the Sommargren interferometer.

A low coherence length laser (<2mm) is used to couple two relatively delayed beams into a fiber, which then illuminates both the test mirror and the CCD detector. The end of the fiber is embedded in a substrate which has been super-polished at a slight angle, to direct the light reflected back from the test mirror to the detector. By adjusting the relative delay between the two arms of the light source to twice the fiber-mirror separation, interference occurs despite the short coherence length of the laser used. The piezo phase stepper can then be used to perform PSI to determine the optical phase difference between the reference and measurement arms with great accuracy. The resulting phase map then needs to be converted into a mirror-figure with an inverse propagation algorithm.

The development of this interferometer has been going on in parallel with the work reported in this thesis, and shares several salient features.

Potential sources of error include the lack of stability and traceability of the wavelength of the low-coherence light source, the surface properties of the fiber substrate, residual and parasitic fringes, and alignment errors. Since the NA of the fiber is used to illuminate the mirror and detector at the same time, the NA of the measurement arm is considerably less than the NA of the fiber, meaning that larger optics will have to be measured in parts. Furthermore, the modulation depth of this arrangement cannot exceed 0.5, because of the mutually incoherent parts of the two beams.

1.3 Conclusion

It is clear that there are several methods already available to measure the figure of EUVL optics with the required accuracy. With the exception of the Sommargren interferometer, none of the methods outlined are suitable for in-process control of such optics in a standard optical shop. The purpose of our interferometer should therefore be the sub-nm accuracy measurement of large NA aspheres for EUV or X-ray applications, in an optical shop environment, without the constant need for re-calibration, both in a reasonable time and at a reasonable cost.

1.4 References

[1] G. E. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, 1965.

[2] Carl Zeiss Website, "Nano Lithography," vol. 2004, 2004. (http://www.smt.zeiss.com/C1256E4600305472/Contents-Frame/8F27C3099EA4FC53C1256E540047800E)

[3] "Carl Zeiss SMT closes gap in EUV lithography," in *Electro Optics Magazine*, vol. 33.

[4] Intel, "Fall 2002 Update - Expanding Moore's Law," Intel 2002.

[5] M. Berglund, "A Compact Soft X-ray Microscope Based on a Laser-Plasma Source." Stockholm: Royal Institute of Technology, 1999.

[6] N. Kaiser, S. Yulin, T. Feigl, H. Bernitzki, et al., "EUV and soft X-ray multilayer optics," in *Advances in Optical Thin Films*, vol. 5250, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2003, pp. 109-118.

[7] R. G. Klaver, "Novel interferometer to measure the figure of strongly aspherical mirrors." Delft: Delft University of Technology, 2001.

[8] H. Medecki, E. Tejnil, K. A. Goldberg, and J. Bokor, "Phase-shifting point diffraction interferometer," *Optics Letters*, vol. 21, pp. 1526-1528, 1996.

[9] P. Naulleau, K. A. Goldberg, E. H. Anderson, P. Batson, et al., "At-wavelength characterization of the extreme ultraviolet Engineering Test Stand Set-2 optic," *Journal of Vacuum Science & Technology B*, vol. 19, pp. 2396-2400, 2001.

[10] K. A. Goldberg, P. Naulleau, P. Batson, P. Denham, et al., "Extreme ultraviolet alignment and testing of a four-mirror ring field extreme ultraviolet optical system," *Journal of Vacuum Science & Technology B*, vol. 18, pp. 2911-2915, 2000.

[11] K. A. Goldberg, P. Naulleau, and J. Bokor, "Extreme ultraviolet interferometric measurements of diffraction-limited optics," *Journal of Vacuum Science & Technology B*, vol. 17, pp. 2982-2986, 1999.

[12] P. P. Naulleau, K. A. Goldberg, S. H. Lee, C. Chang, et al., "Extreme-ultraviolet phase-shifting point-diffraction interferometer: a wave-front metrology tool with subangstrom reference-wave accuracy," *Applied Optics*, vol. 38, pp. 7252-7263, 1999.

[13] K. A. Goldberg, P. Naulleau, S. Lee, C. Bresloff, et al., "High-accuracy interferometry of extreme ultraviolet lithographic optical systems," *Journal of Vacuum Science & Technology B*, vol. 16, pp. 3435-3439, 1998.

[14] K. A. Goldberg, P. P. Naulleau, P. E. Denham, S. B. Rekawa, et al., "Preparations for extreme ultraviolet interferometry of the 0.3 numerical aperture Micro Exposure Tool optic," *Journal of Vacuum Science & Technology B*, vol. 21, pp. 2706-2710, 2003.

[15] K. A. Goldberg, P. Naulleau, J. Bokor, H. N. Chapman, et al., "Testing extreme ultraviolet optics with visible-light and extreme ultraviolet interferometry," *Journal of Vacuum Science & Technology B*, vol. 20, pp. 2834-2839, 2002.

[16] B. Dorband and G. Seitz, "Interferometric testing of optical surfaces at its current limit," *Optik*, vol. 112, pp. 392-398, 2001.

[17] H. Handschuh, J. Froschke, M. Julich, M. Mayer, et al., "Extreme ultraviolet lithography at Carl Zeiss: Manufacturing and metrology of aspheric surfaces with angstrom accuracy," *Journal of Vacuum Science & Technology B*, vol. 17, pp. 2975-2977, 1999.

[18] M. Takeda, H. Ina, and S. Kobayashi, "Fourier-Transform Method of Fringe-Pattern Analysis for Computer-Based Topography and Interferometry," *Journal of the Optical Society of America*, vol. 72, pp. 156-160, 1982.

[19] P. Mercere, P. Zeitoun, M. Idir, S. Le Pape, et al., "Hartmann wave-front measurement at 13.4 nm with lambda(EUV)/120 accuracy," *Optics Letters*, vol. 28, pp. 1534-1536, 2003.

[20] D. A. Tichenor, A. K. Ray-Chaudhuri, S. H. Lee, H. N. Chapman, et al., "Initial results from the EUV Engineering Test Stand," in *Soft X-Ray and Euv Imaging Systems II*, vol. 4506, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2001, pp. 9-18.

[21] J. S. Taylor, G. E. Sommargren, D. W. Sweeney, and R. M. Hudyma, "The fabrication and testing of optics for EUV projection lithography," in *Emerging Lithographic Technologies II*, vol. 3331, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 1998, pp. 580-590.

[22] G. E. Sommargren, D. W. Phillion, M. A. Johnson, N. Q. Nguyen, et al., "100-picometer interferometry for EUVL," in *Emerging Lithographic Technologies VI, Pts 1 and 2*, vol. 4688, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2002, pp. 316-328.

[23] G. E. Sommargren, "Diffraction methods raise interferometer accuracy," *Laser Focus World*, vol. 32, pp. 61-&, 1996.

2 Concept

As mentioned in the previous chapter, we aim to present an interferometer capable of sub-nm accuracy measurements of large NA aspheres for EUV or X-ray applications, in an optical shop environment, without the constant need for re-calibration, both in a reasonable time and at a reasonable cost. In order to fulfil these requirements, a number of challenges must be overcome.

2.1 Challenges

The prohibitive cost of a EUV source[1], coherent or not, strongly indicates the use of longer wavelengths for the purposes of optical shop metrology. However, the use of longer wavelengths places a bigger demand on the phase-measurement accuracy. In our case, this accuracy must reach λ/10'000 in order to be effective.

Since this is the first time that such an interferometer has been realized, it is essential that a certain amount of flexibility in the design of the light-source is maintained to allow for a comparison between the results obtained with various modes of operation. Such flexibility invariably comes at the expense of efficiency and financial costs. In contrast, the design of the interferometer frame was kept as simple as possible, optimized for the measurement of only one particular test mirror. In this way, the required positioning stability could be attained at a reasonable cost.

The positioning stability of interferometer components is critical to the accurate measurement of the EUVL substrates[2]. This stability has to be maintained not only during measurement, where the presence of vibrations and drift will introduce systematic errors, but also between measurements, to avoid the need for frequent re-calibration. The stability requirements extend, beyond the mere positioning stability of components, to stability requirements of the atmosphere in which the measurement is performed, as refractive index changes in the atmosphere will easily introduce excessive errors.

In trying to eliminate the potential errors and calibration procedures associated with reference- and intermediate optics in the interferometer beam path, we must find a way to build an interferometer with no such components, while maintaining the flexibility to measure a variety of substrates.

To ensure traceability of the measurements to international standards, the measurement wavelength must be kept stable in the long term. The need for wavelength stability also arises from a number of other design aspects of the interferometer.

Despite the use of longer wavelengths, the large asphericity of some EUVL mirrors means that the resulting fringe pattern would be too dense to be resolved by standard CCD type sensors. A way to overcome this under-sampling of the fringes has to be found[3].

Finally, since the measurement data is not trivially related to the shape of the mirror, an algorithm must be found which retrieves the mirror shape from the measured data with the required accuracy.

The following section will summarize the approach chosen to attempt to fulfil these goals.

2.2 Approach

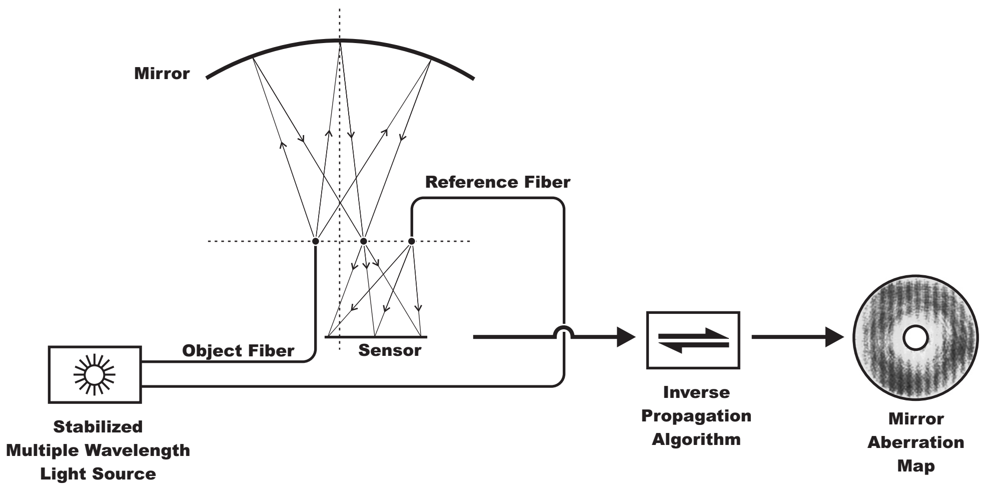

Figure 2-1 – Schematic diagram of interferometer principle

Figure 2-1 shows a schematic diagram of the various components of our interferometer. These will be discussed in detail during subsequent chapters, and have been partially treated in a series of papers[4-7].

The primary light source used here is a stabilized HeNe laser, capable of maintaining the wavelength accuracy required for our purposes. The accuracy of our measurements will be based on the accuracy of this laser.

By passing this light through a single-mode optical fiber, we obtain a highly spherical wavefront, which we will use as a reference shape against which the EUVL substrates will be measured[8]. The advantages of this approach are that we immediately have a large range of curvatures at our disposal, by simply adjusting the fiber position, and that this wavefront reference essentially does not require repeated calibration (see chapter 6).

The object fiber tip is placed near the centre of curvature of the reflector, displaced a few hundred microns in the horizontal direction and pointed towards the area of interest on the mirror. The reflection off the mirror will come to a focus in the same horizontal plane as the fiber before falling on the sensor.

Light from the reference fiber is brought to overlap with the reflection from the mirror on the sensor, where the two beams interfere.

In order to guarantee an accurate and stable interference pattern, the whole set-up must be stable both in the short- and long term. In the short term, vibrations could cause excess noise, while long term drifts could cause systematic errors to appear in our measurements.

With these considerations in mind, a new set of ultra-stable optical mounts was designed and manufactured for the light source. The interferometer itself, housing the mirror under test, the fibers and the sensor, was designed entirely from Invar – a material with a particularly low thermal expansion coefficient, using a well established hexapod design which gives excellent immunity against deformations and vibrations.

To reduce the influence of atmospheric changes and inhomogeneities on our measurement, both the light-source and interferometer can be placed in a Helium atmosphere. Helium has a refractive index seven times closer to unity than that of regular air, hence reducing the influence of density, pressure and temperature changes by the same factor[2].

To meet our accuracy requirements, the optical phase difference between the two beams has to be measured with $2\pi/10'000$ accuracy. Such accuracies can only be obtained by using either phase stepping- or heterodyne interferometry[9].

Our lightsource allows for a step-wise or continuous phase shift between the reference and object fibers[7]. A sequence of images captured from the sensor between discrete phase steps can be analysed with phase-stepping interferometry (PSI) techniques to yield the local optical phase difference at every pixel. Heterodyne techniques can be applied to measure the phase at the sensor for a continuous phase shift, equivalent to a slight wavelength offset between the two fibers – provided the sensor is capable of heterodyne detection.

Such a detector is indeed currently under development, and an alternative novel sensor has been obtained to demonstrate the principle in the meantime.

Due to a combination of the large asphericities of the mirror under test and the separation of the fibers, a very dense fringe pattern is produced at the sensor. The resolution of available sensors is not high enough to distinguish some of these fringes, making it impossible to retrieve an optical path difference (OPD) map using standard methods.

We overcome this problem by using multiple wavelength interferometry[10, 11]. The same setup is used to measure the mirror with a slightly different but stable and well-known wavelength produced by our lightsource. The difference between the two measured interferograms allows us to resolve the phase ambiguities caused by the under-sampling of the fringes.

Once an accurate OPD map is obtained, we use a specially developed inverse propagation algorithm to deduce the shape of the mirror from our data. Our task is impeded by the presence of severe diffraction effects, rendering purely geometrical optics insufficiently accurate for this purpose. A combination of geometrical optics and rigorous diffraction methods are therefore employed to calculate the shape of our mirror with the required accuracy.

2.3 Conclusion

Our approach resolves many of the problems inherent in the methods mentioned in the previous chapter, especially with regards to cost effective optical shop measurements. The utilization of a precise and traceable wavelength, together with a stable reference wavefront makes our interferometer useful as an absolute measurement device.

2.4 References

[1] M. Berglund, "A Compact Soft X-ray Microscope Based on a Laser-Plasma Source." Stockholm: Royal Institute of Technology, 1999.

[2] B. Dorband and G. Seitz, "Interferometric testing of optical surfaces at its current limit," *Optik*, vol. 112, pp. 392-398, 2001.

[3] G. E. Sommargren, D. W. Phillion, M. A. Johnson, N. Q. Nguyen, et al., "100-picometer interferometry for EUVL," in *Emerging Lithographic Technologies VI, Pts 1 and 2*, vol. 4688, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2002, pp. 316-328.

[4] R. G. Klaver and J. J. M. Braat, "Novel interferometer to measure the figure of aspherical mirrors as used in EUV lithography," in *Emerging Lithographic Technologies IV*, vol. 3997, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2000, pp. 784-793.

[5] R. G. Klaver, H. van Brug, and J. J. M. Braat, "Interferometer to measure the form figure of aspherical mirrors as used in EUV lithography," in *Laser Metrology and Inspection*, vol. 3823, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 1999, pp. 123-132.

[6] M. L. Krieg, R. G. Klaver, and J. J. M. Braat, "Absolute optical path difference measurement with angstrom accuracy over ranges of millimetres," in *Optical Measurement Systems for Industrial Inspection II: Application in Industrial Design*, vol. 4398, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2001, pp. 116-126.

[7] M. L. Krieg, G. Parikesit, and J. J. M. Braat, "Three-wavelength laser light source for absolute, sub-Angstrom, two point source interferometer," in *Optical Measurement Systems for Industrial Inspection III*, vol. 5144, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2003, pp. 227-233.

[8] G. E. Sommargren, "Diffraction methods raise interferometer accuracy," *Laser Focus World*, vol. 32, pp. 61-&, 1996.

[9] D. Malacara, *Optical shop testing*, 2nd ed: Wiley-Interscience, 1992.

[10] R. Dandliker, K. Hug, J. Politch, and E. Zimmermann, "High-accuracy distance measurements with multiple-wavelength interferometry," *Optical Engineering*, vol. 34, pp. 2407-12, 1995.

[11] R. Dandliker, Y. Salvade, and E. Zimmermann, "Distance measurement by multiple-wavelength interferometry," *Journal of Optics*, vol. 29, pp. 105-14, 1998.

3 Inverse Propagation Algorithm

In designing an inverse propagation algorithm to retrieve the shape of the reflector under test from our interferometer data, we aim to fulfil two main requirements: sufficient accuracy and a reasonable execution time.

Unfortunately, these are conflicting requirements, as increased accuracy invariably means increased execution time. We therefore look for an algorithm which makes use of as many symmetries and simplifications applicable to our instrument, without unduly affecting the algorithm's accuracy. A brief treatment of this problem has recently been published[1], but we will take the opportunity to develop these ideas in more detail here.

The error budget set up by Klaver[2] to attain a final measurement accuracy of 0.1nm rms requires an inversion algorithm accuracy equivalent to ~λ/10'000. At this level, the consideration of diffraction effects is not only unavoidable, but must be performed with particular accuracy. Assumptions and approximations which are routinely made to simplify diffraction calculations must now be scrutinised to ensure that they do not violate our requirements. As we shall see in the course of this chapter, a hybrid method combining ray tracing methods with diffraction calculations is the most promising candidate to meet our requirements at a minimal computational cost. We will therefore begin with a detailed description of the raytracing method used, before justifying its use with rigorous diffraction theory. Following chapter 3.2 which identifies the conditions under which the raytracing approach fails to give an accurate description of our measurements, we show a number of ways to improve our algorithm to properly deal with diffraction effects. These methods are vindicated by comparison against a brute-force forward propagation algorithm which provides the required accuracy.

3.1 Raytracing approach

In discussing this approach, we will use the notation shown in Figure 3-1:

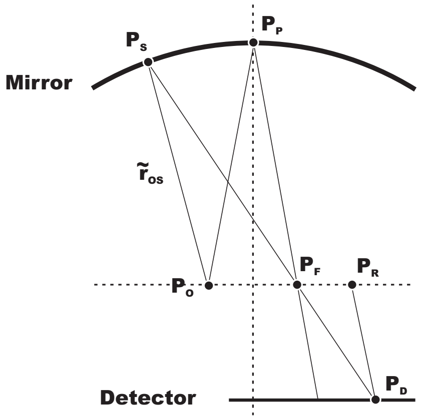

P_O – Location of point-source
P_P – Particular point on mirror
P_F – Nominal focal point
P_D – Point of interest on detector
P_S – Point of reflection on mirror surface
P_R – Position of reference fiber.
$\tilde{\mathbf{n}}(P_S)$ – Normal to mirror at P_S (not shown)
$\tilde{\mathbf{r}}_{OS}$ – Vector from P_O to P_S

Figure 3-1 – Notation convention for raytracing approach.

In addition to the vector $\tilde{\mathbf{r}}_{OS}$, vectors connecting an arbitrary pair of points, P_J to P_K, will be written as $\tilde{\mathbf{r}}_{JK}$. Our treatment of the forward problem necessarily assumes an exact knowledge of the locations of all the interferometer components.

As we will show shortly, the process of inversion turns out to be simpler and computationally less expensive than the forward problem of simulating our measured quantity, the optical path difference (OPD), at the detector. Nonetheless, we will begin with a review of the forward problem before describing the inversion approach.

3.1.1 Forward raytracing problem

The initial task for obtaining the OPD map at our detector consists of finding the points of reflection on the mirror surface (P_S) which will send rays from the fiber-tip to the various pixel locations (P_D). We begin by choosing a set of NxM points on the mirror surface (dependent on our desired resolution) which generously cover the complete mirror aperture. The normals to the surface of the mirror at these points are evaluated, and the directions of the reflected rays are calculated from the law of reflection. The intersection of these rays with the detection plane is then found analytically. The resulting ray intersections with the detection plane will not coincide with our pixel locations, and so we interpolate between the initial ray positions on the mirror to find rays which should intersect more closely to our pixels. The process of tracing this new set of rays to the detection plane is then repeated. To ensure that the reflected rays intersect with the detection plane within 0.1nm of the actual location of the pixel centres (P_D), three to four iterations of this interpolation and raytracing process are required. The whole procedure takes a matter of seconds for a grid of 25x25 points on a modern computer.

For every pixel location (P_D), we now have the corresponding point of reflection on the mirror (P_S), from which we can calculate the total optical path length (OPL) of the ray from the object fiber tip (P_O) to the pixel (r_{OD}+ r_{DS}). The OPD is found by subtracting the path from the reference fiber (P_O) to the pixel:

$$\begin{aligned} OPD\left(P_D\right) &= r_{OS} + r_{SD} - r_{RD} \\ &= OPL\left(P_D\right) - r_{RD} \end{aligned} \tag{3.1}$$

In this way we can get an OPD map for our detector. This OPD is not the true optical path difference between the object- and reference-beams however, since there is an optical delay before the reference fiber tips. While the geometric OPD given by (3.1) has a large offset, the delay before the reference fiber tip is adjusted to reduce the offset of the true OPD to zero.

By way of example, we will generate the OPD and OPL maps for an aspheric mirror with radius of curvature (ROC) ~340mm, aperture diameter of 160mm and a p-v deviation from the best-fit sphere of approximately 4.6μm. (See Figure 3-2 for the mirror shape and Figure 3-3 for the resulting simulated OPD and OPL maps). The positions of the various interferometer components are those given in Figure 6-8b.

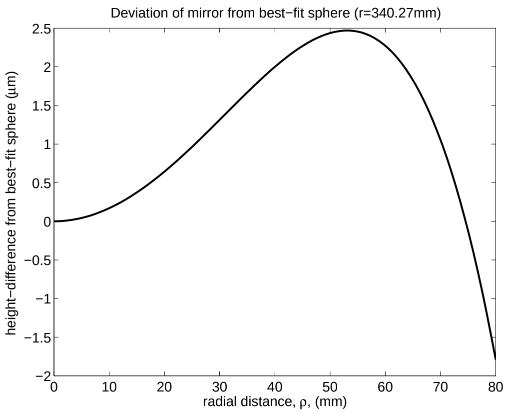

Figure 3-2 – Deviation from best-fit sphere of the rotationally symmetric example mirror.

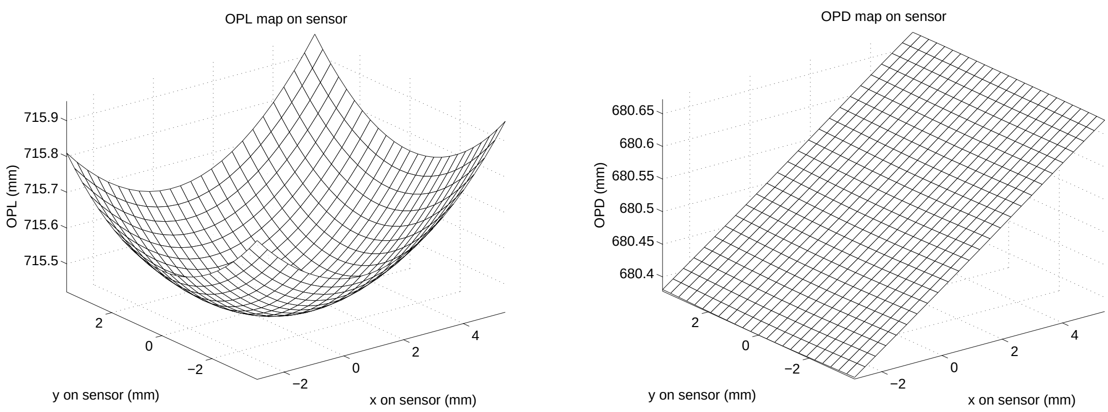

Figure 3-3 – Optical path length (OPL) from object fiber to detector and optical path difference (OPD) at detector, as given by (3.1).

3.1.2 Inverse raytracing problem

Although we do not have access to the absolute OPD from our interferometer, we have a very sensitive measurement of the relative OPD. We can estimate the offset of the relative OPD from the absolute OPD to within a few tens of microns for a well-constructed interferometer by using measurements of the geometry of the set-up.

We will first show how we can retrieve the shape of the mirror, given the exact positions of the key components of the interferometer and the true value of the OPD. The consequences of errors in our estimates of these quantities, and the steps required to correct for them will be discussed later.

We begin by calculating the OPL from our measurement of the OPD at the detector, using (3.1). From the OPL between the object fiber tip and the pixel position, we can conclude that the point of reflection (P_S) lies somewhere on a prolate spheroid with P_O and P_D as foci (Figure 3-4).

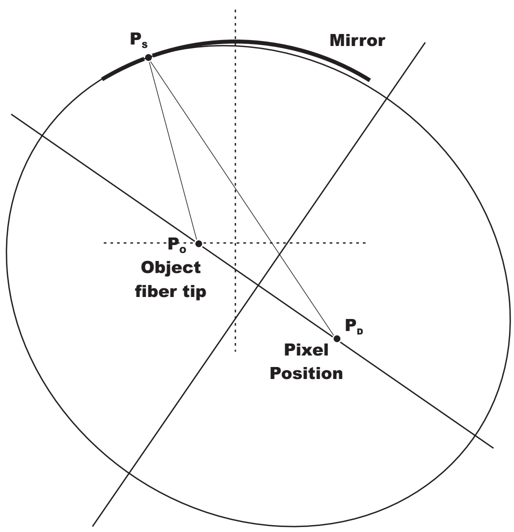

Figure 3-4 – Prolate spheroid traced out by fixed OPL around fiber tip and pixel location.

From a single-point measurement it is therefore impossible to uniquely determine the point of reflection. However, we may use the OPL of neighbouring pixels to estimate the normal to the wavefront at the pixel, and hence the direction of the ray $\tilde{\mathbf{r}}_{SD}$.

Making the assumption of a locally plane wave (an assumption later justified by the resulting inversion accuracy), we can write the OPL function in the immediate vicinity of a particular pixel located at (x_D, y_D, z_D,) as:

$$OPL \approx \frac{\lambda}{2\pi}\left(k_x\left[x_D+\varepsilon_x\right]+k_y\left[y_D+\varepsilon_y\right]+k_z\left[z_D+\varepsilon_z\right]+c\right) \tag{3.2}$$

with:

$$\left|\tilde{\mathbf{k}}\left(x_D,y_D,z_D\right)\right|=\sqrt{k_x^2+k_y^2+k_z^2}=\frac{2\pi}{\lambda} \tag{3.3}$$

The direction of the wave-vector, $\tilde{\mathbf{k}}(x_D, y_D, z_D)$, is identical to the direction of the ray $\tilde{\mathbf{r}}_{SD}$. Thanks to (3.3), it is sufficient for us to find two out of the three wave-vector components – the third resulting from the normalization constraint.

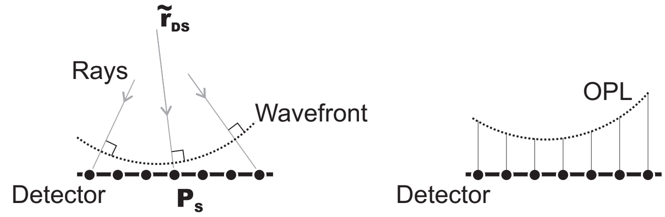

Figure 3-5 – diagram illustrating rays, wavefront and OPLs at detector

The partial derivatives of the OPL function in the x- and y- direction are therefore sufficient to retrieve the direction of the local ray:

$$\tilde{\mathbf{r}}'_{SD} \equiv \frac{\tilde{\mathbf{r}}_{SD}}{\left|\tilde{\mathbf{r}}_{SD}\right|} \approx \left(\frac{\partial OPL}{\partial x}, \frac{\partial OPL}{\partial y}, \sqrt{1 - \left(\tfrac{\partial OPL}{\partial x}\right)^2 - \left(\tfrac{\partial OPL}{\partial y}\right)^2} \right) \tag{3.4}$$

Where $\tilde{\mathbf{r}}'_{SD}$ is the unit vector in the direction of the calculated wavefront normal.

We now have enough information for a unique solution for the point of reflection, given by the intersection of a line from P_D and direction $\tilde{\mathbf{r}}_{DS}$, with the prolate spheroid. Mathematically, this is equivalent to solving a quadratic equation:

$$\tilde{\mathbf{r}}_{DS} = \tilde{\mathbf{r}}'_{DS} \frac{r_{DO} d_A - OPL\sqrt{d_A^2 + d_B^2}}{2\left(d_A^2 + \dfrac{OPL^2 d_B^2}{OPL^2 - r_{DO}^2} \right)} \tag{3.5}$$

$$d_A = \frac{\left\langle \tilde{\mathbf{r}}'_{DS}, \tilde{\mathbf{r}}_{DO} \right\rangle}{r_{DO}},\ d_B = \tilde{\mathbf{r}}'_{DS} - \tilde{\mathbf{r}}_{DO} \frac{d_A}{r_{DO}} \tag{3.6}$$

The mirror shape retrieved with this method exhibits a remarkable stability with respect to errors in the estimated ray direction, as a consequence of the identical local gradients of the prolate spheroid and the mirror surface at the point of reflection. Even the curvatures of the two surfaces are matched very closely, so that the estimate of the mirror shape is correct to better than first order for an error in the ray direction.

Nonetheless, care must be taken when calculating the wavefront normals from our sampled OPD function. The fact that our forward simulations automatically generate the correct ray directions, allows us to evaluate and compare the accuracy of schemes to perform this task. The best results were obtained by fitting a bivariate quadratic (including cross-terms) to six points on our OPL function surrounding the point of interest, in an arrangement similar to that shown in 25.3.27 of Abramowitz & Stegun[3] and using the resulting expression to determine the local gradient. This method has a minute amount of noise suppression, having six fitting parameters compared to the seven data points used in the fit.

Performing this inversion on ideal data shows some marginal residual error (rms = 1.3pm), dominated by edge effects of the wavefront normal retrieval algorithm (Figure 3-6). We will subsequently subtract this residual error from retrieved surface shapes of non-ideal mirrors, giving a perfect retrieval for the ideal mirror shape by definition.

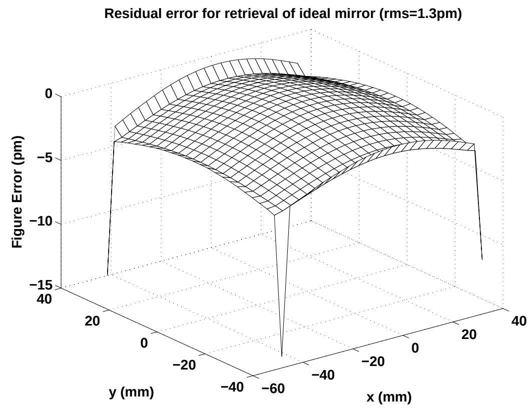

Figure 3-6 – Residual error when directly inverting ideal, simulated data (rms=1.3pm)

To ensure that the accuracy of this procedure is not dependent on this particular mirror shape, or its circular symmetry, we repeat the inversion for data generated with the previous mirror profile to which a Gaussian “bump” of several nm has been added (see Figure 3-7)

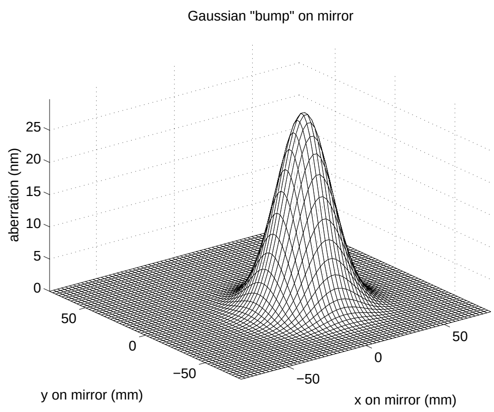

Figure 3-7 – Gaussian “bump” added to ideal mirror surface profile.

The error figure between the actual and retrieved shapes in Figure 3-8 shows that the algorithm retrieves the correct mirror shape with an rms error of merely 2.0pm. This is more than sufficient for our purposes.

We can therefore state that we can successfully retrieve the surface profiles of mirrors, given the exact OPD distribution and interferometer parameters, under the assumption that raytracing is valid.

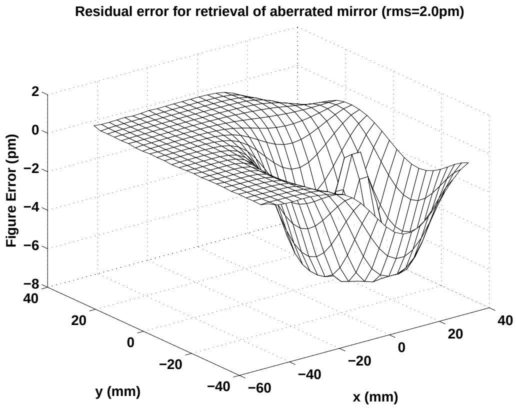

Figure 3-8 – Residual error when inverting data from a mirror aberrated by a Gaussian “bump” (rms=2.0pm) c.f. Figure 3-7.

As already mentioned, this inversion requires much less computing time than the forward problem – the inversion executing in well under a second for a 25x25 grid of data. When investigating the effect of parameter errors, we will therefore change the parameters for the inversion rather than the parameters for the forward problem.

3.1.3 Parameter errors

We will now consider the effect of errors in our assumptions about the positions of the various interferometer components. In this section, the coordinate origin will be placed at the nominal centre of the detector.

Object fiber position P_O:	(x_O,y_O,z_O)	Detector centre P_{DC}:	(x_{DC},y_{DC},z_{DC})
Reference fiber position P_R:	(x_R,y_R,z_R)	Detector orientation ΔO_O:	$(\Delta\phi_X, \Delta\phi_Y, \Delta\phi_Z)$
OPD offset:	(OPD_0)	Detector size S_O:	(Sx_D, Sy_D)

Table 3-1 – Interferometer parameters considered as potential error sources.

An error in any of the parameters listed in Table 3-1 will cause a figure error to be introduced into our retrieved mirror shape. A global translation of the fibers and the sensor is equivalent to the opposite translation of the mirror. The same is true for a global rotation about a fixed point. Such errors can therefore be grouped under the term “alignment errors”, and should be reported separately from the purely “figure” errors. The error introduced by a 1μm shift of the mirror in the x-direction is shown in Figure 3-9 as an example.

Due to the almost spherical shape of the mirror, a rotation about a point close to the mirror surface (tilt) would be virtually identical to a horizontal shift plus an offset. It is important to note that “tilt” cannot be considered a purely linear term, as seen in Figure 3-9, while piston (translation in the z-direction) can. Another type of error which is considered permissible beyond the 0.1nm accuracy otherwise desired for the figure measurement is defocus, defined as a figure error proportional to the square of the radial co-ordinate in the horizontal plane. Alignment errors and defocus are considered acceptable within certain limits specified by the manufacturers of such mirrors[2] and we will therefore report two sets of figure errors for every parameter variation: one without subtraction of alignment and defocus errors, and one with this subtraction.

In addition to the parameters listed in Table 3-1, the direction of the fibers and the positions of the individual pixels can affect the accuracy of the inverse propagation algorithm. The wavefront from the fiber is only spherical around the central part of the intensity distribution (see chapter 6.3), and care must therefore be taken to ensure that the fiber-output is directed at the centre of the area of the reflector under test. As long as the area under test falls within the region of acceptable wavefront sphericity, the pointing direction of the fiber has no systematic effects on the retrieved shape of the reflector, due to the spherical symmetry of the wavefront. The relative positions of the individual pixels can be determined with independent calibration methods (see section 5.2.2), and we assume that such a calibration has already been carried out. The parameters listed in Table 3-1 are then sufficient to represent any remaining alignment errors for the sensor.

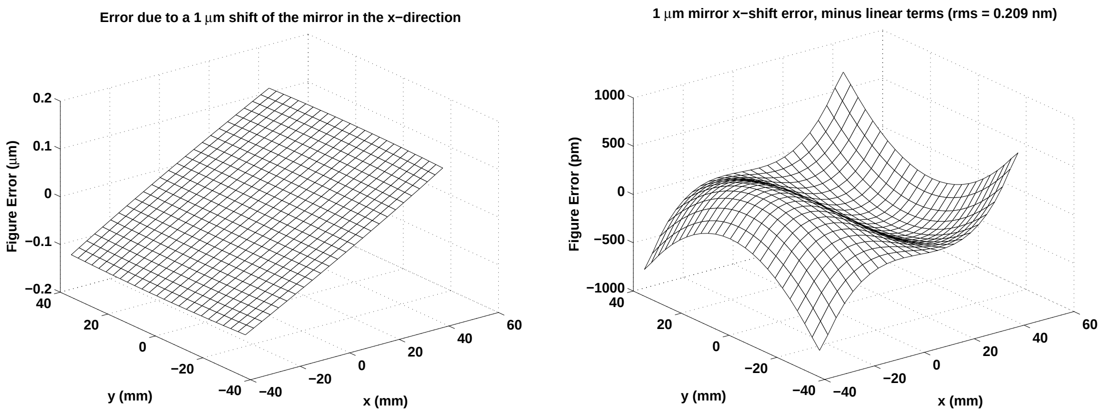

Figure 3-9 – Errors in the surface height measurement for a displacement of the mirror in the x-direction. Note the nonlinearity of this error term. The equivalent error profile for a purely spherical mirror would differ from the above by only a few pm.

Figure 3-10 to Figure 3-24 show the various errors in the retrieved mirror shape for 1μm errors in the various component positions, 1mrad errors in the orientation, and 1ppm (parts per million) errors for the CCD size. It should be clear that the most severe figure errors are introduced by incorrect estimates in the horizontal positions of interferometer components, where they cause what appears to be shear-errors in the direction of displacement. The obtainable positioning accuracy of the fibers in the xy plane can be considered to be in the 1-5μm range, while the z-positions could be determined with even better accuracy. The CCD horizontal position can also be determined to within about 1μm by noting the location of visible mirror features (such as apertures and/or markings) projected on the CCD. This is still not sufficient to guarantee a figure error of 0.1nm however, as can be seen from Table 3-2.

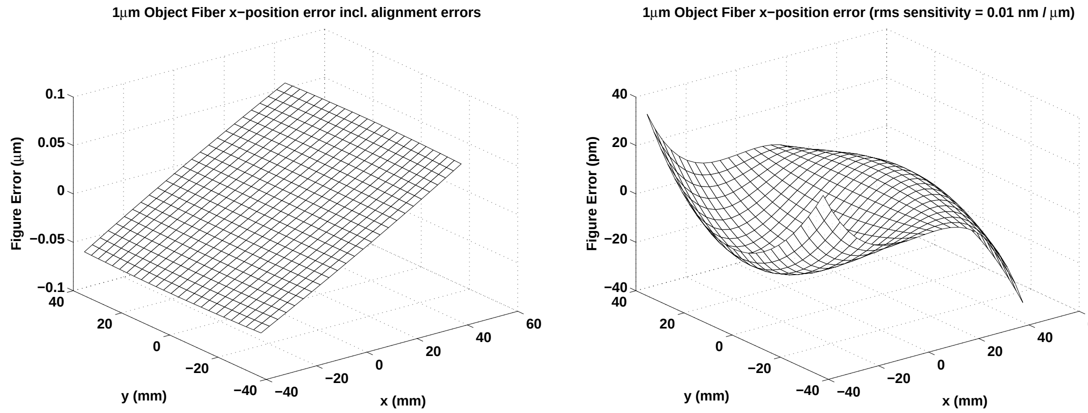

Figure 3-10 – Figure error due to an erroneous estimate of the x-coordinate of the object fiber

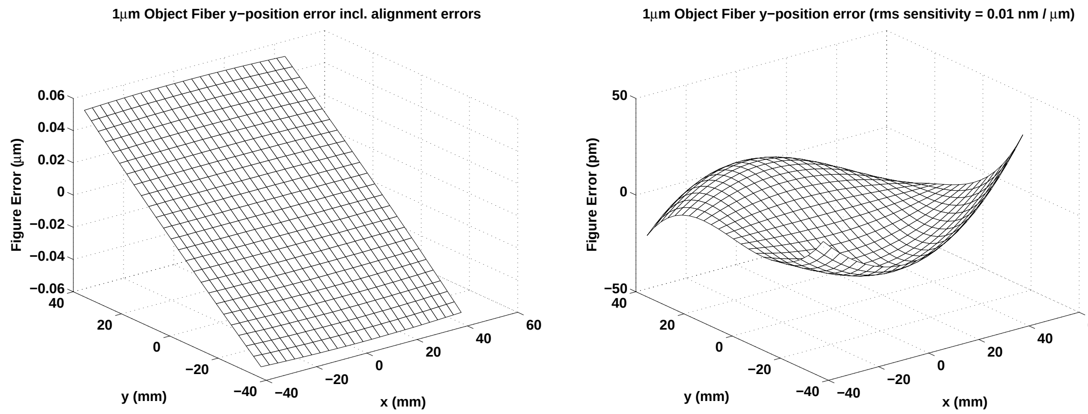

Figure 3-11 – Figure error due to an erroneous estimate of the y-coordinate of the object fiber

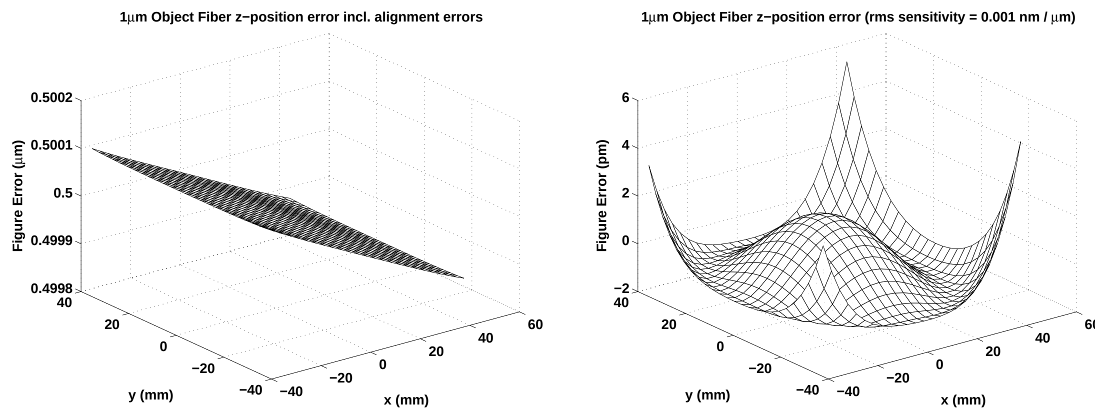

Figure 3-12 – Figure error due to an erroneous estimate of the z-coordinate of the object fiber.

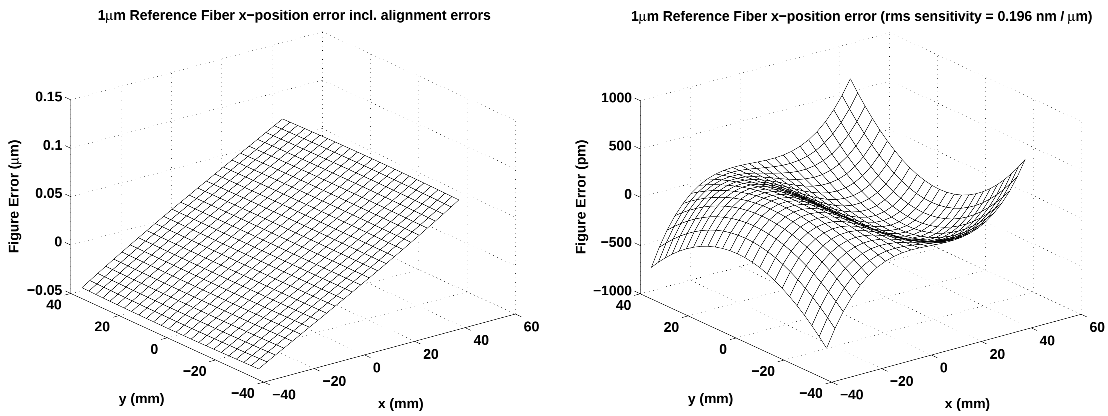

Figure 3-13 – Figure error due to an erroneous estimate of the x-coordinate of the reference fiber.

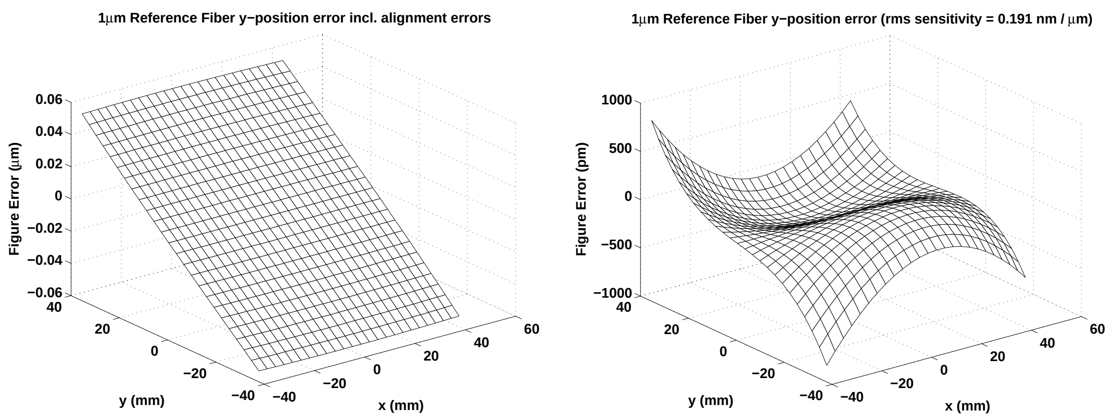

Figure 3-14 – Figure error due to an erroneous estimate of the y-coordinate of the reference fiber

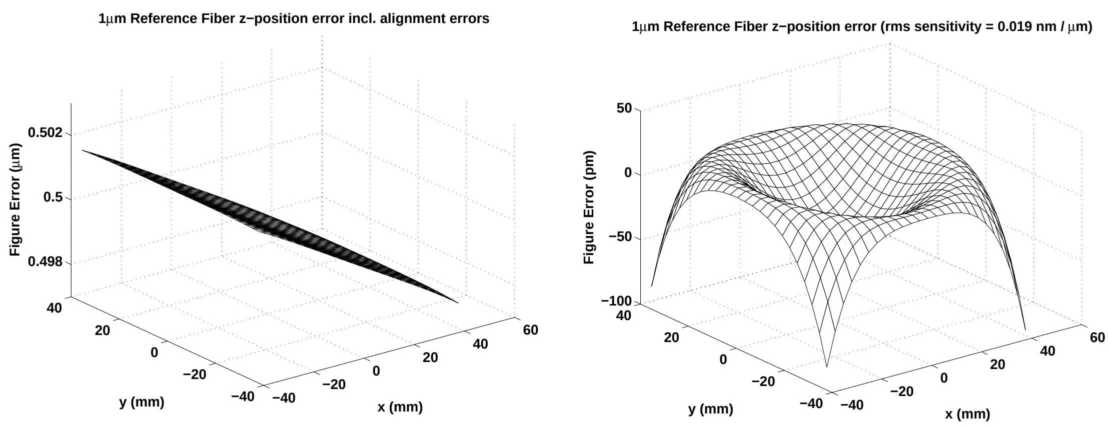

Figure 3-15 – Figure error due to an erroneous estimate of the z-coordinate of the reference fiber

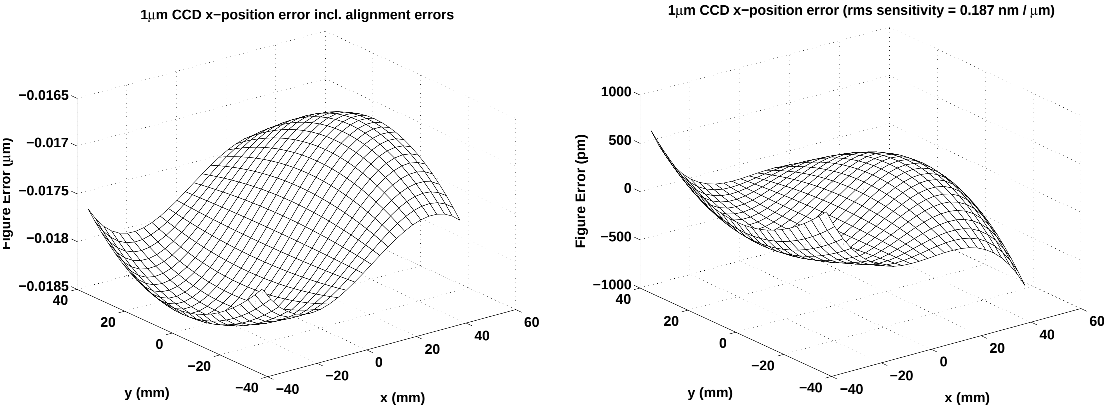

Figure 3-16 – Figure error due to an erroneous estimate of the x-coordinate of the CCD centre

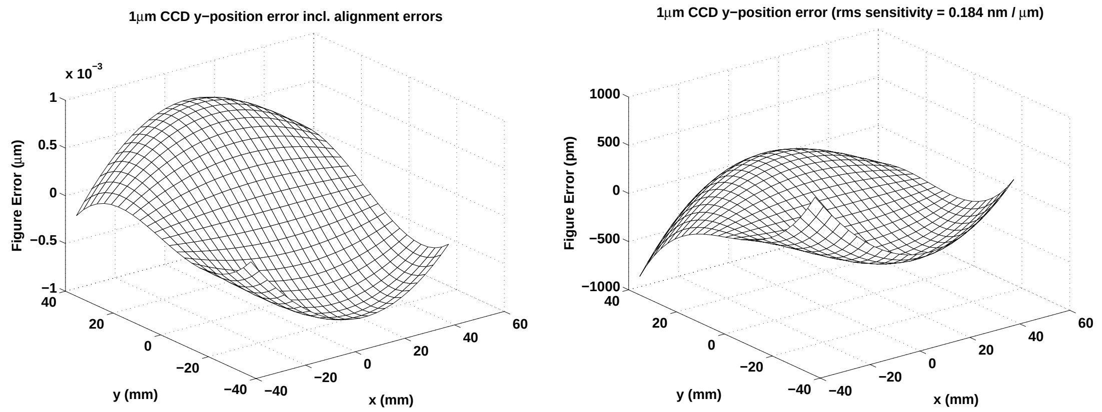

Figure 3-17 – Figure error due to an erroneous estimate of the y-coordinate of the CCD centre

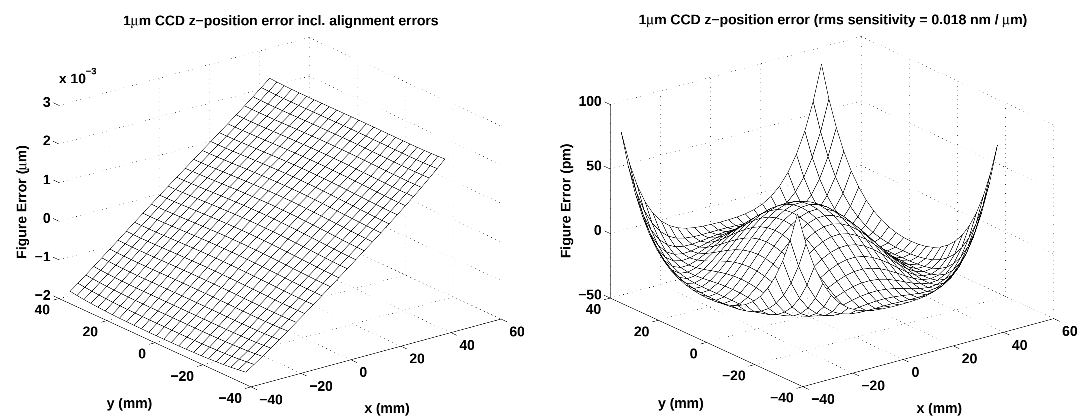

Figure 3-18 – Figure error due to an erroneous estimate of the z-coordinate of the CCD centre

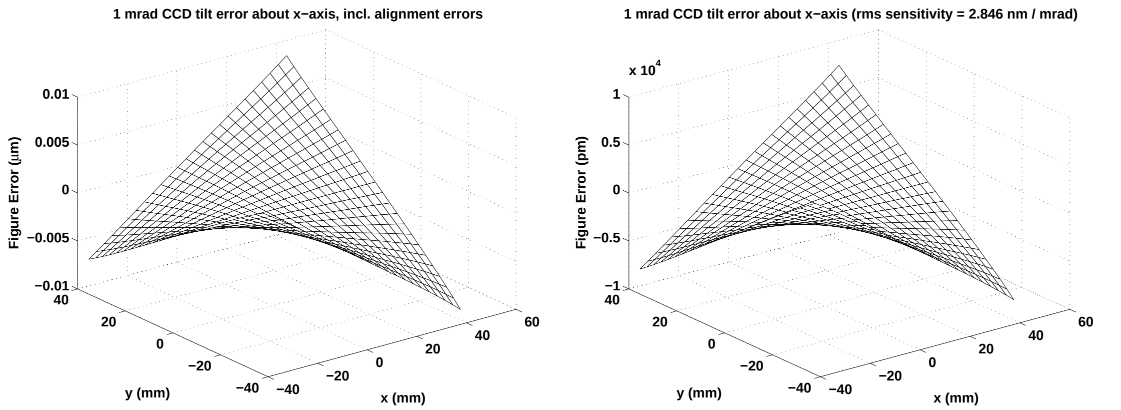

Figure 3-19 – Figure error due to an erroneous estimate of the CCD orientation about the x-axis.

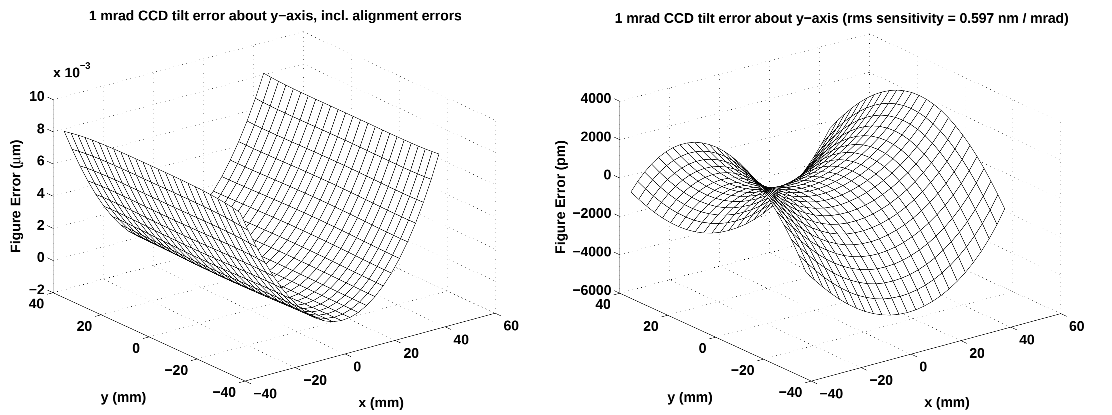

Figure 3-20 – Figure error due to an erroneous estimate of the CCD orientation about the y-axis.

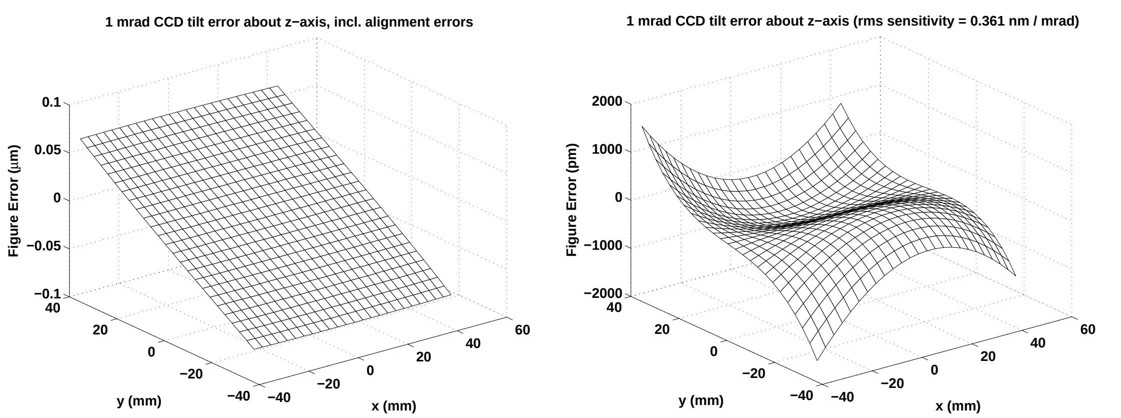

Figure 3-21 – Figure error due to an erroneous estimate of the CCD orientation about the z-axis.

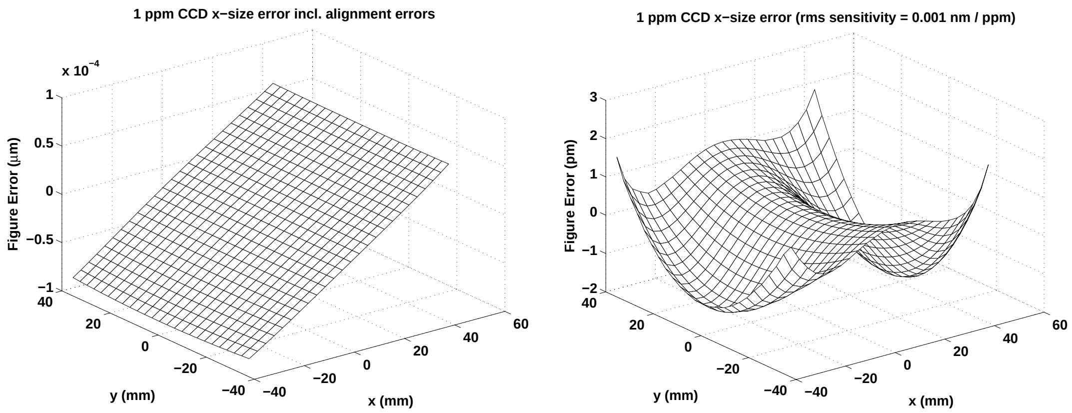

Figure 3-22 – Figure error due to an erroneous estimate of the x-size CCD

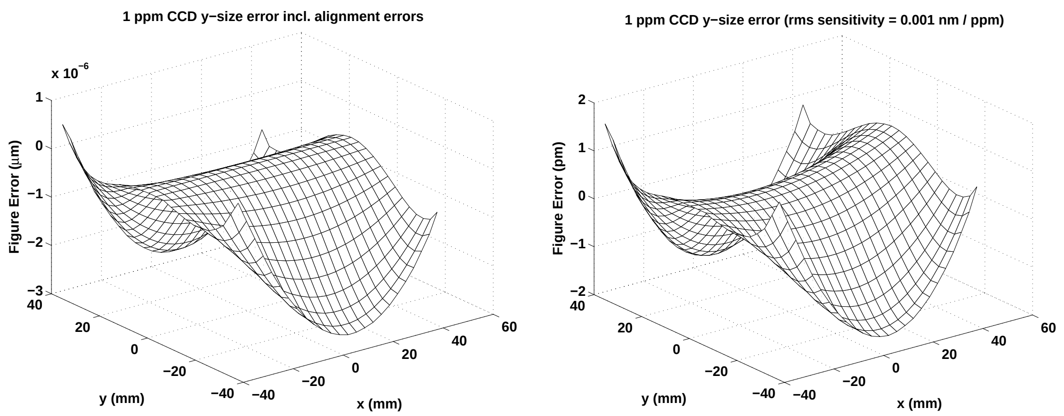

Figure 3-23 – Figure error due to an erroneous estimate of the y-size CCD

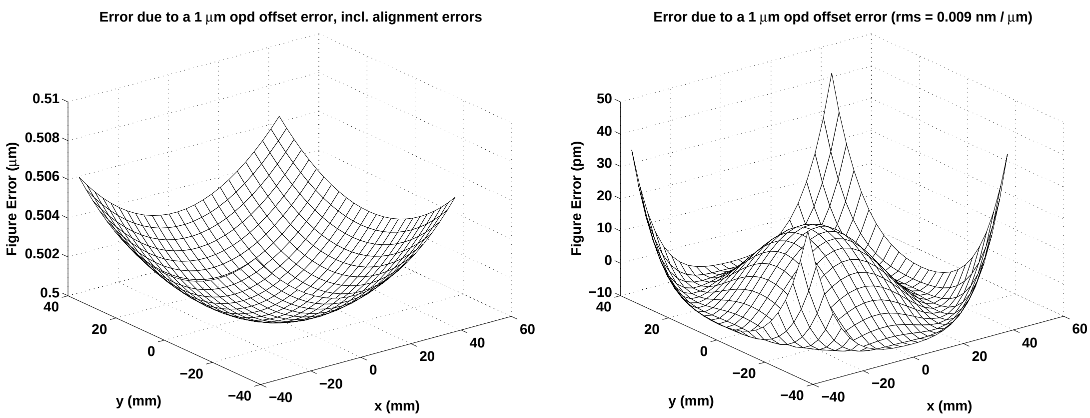

Figure 3-24 – Figure error due to an erroneous estimate of the OPD offset.

The above error figures allow us to make an error-budget for the positioning of the interferometer components. Table 3-2 below shows the positioning accuracies which, individually, cause figure errors below 0.1nm. Similar errors are marked by the symbols: ∗, +, and † which correspond to the x-,y- and z-fiber positioning errors respectively.

$^*x_{of}$	10.0 μm
$^+y_{of}$	10.0 μm
$^\dagger z_{of}$	100.0 μm
$^*x_{rf}$	0.5 μm
$^+y_{rf}$	0.5 μm
$^\dagger z_{rf}$	5.0 μm
†OPD	10.0 μm

$^+x_{CCD}$	0.5 μm
$^*y_{CCD}$	0.5 μm
$^\dagger z_{CCD}$	5.0 μm
$\Delta\phi_x$	30.0 μrad
$\Delta\phi_y$	200.0 μrad
$^+\Delta\phi_z$	200.0 μrad
Sx_{CCD}	0.1 %
Sy_{CCD}	0.1 %

Table 3-2 – Interferometer parameter accuracies, each producing <0.1nm figure errors.

To decrease the influence of these errors, we can try to optimize the interferometer parameters such, that they give a best fit of the resulting retrieved mirror shape to the ideal mirror shape. Only one member of a set of similar parameters needs to be fitted, resulting in a total of 7 possible fitting parameters. This procedure is likely to result in an overly optimistic estimate of the error figure for our mirror, since any actual figure errors present on our mirror of the type shown in the preceding series of figures will be significantly attenuated by such a fitting procedure. Chapter 7 shows how such parameter fitting can be achieved.

Ideally, a parameter optimization of this type would first be performed for a particularly well characterized surface, such as a spherical reflector, and the resulting calibrated interferometer parameters adopted for the retrieval of future mirror shapes. Alternatively, extensive metrology of the interferometer can be performed a-priori, to determine the positions and orientation of the various interferometer components.

3.1.4 Summary

Provided that the assumptions and approximations underpinning the raytracing approach are valid, this inversion technique gives an estimate of the figure of the reflector under test within ~2pm rms with our interferometer. The influence of positioning errors should be considered an effect of the interferometer type rather than of the inversion approach, but they will most likely dominate the low-frequency error landscape in our case.

Examples of reflectors for which the raytracing description is sufficient would be ones where the smooth reflective surface extends beyond the illuminated area, and contains no obstructions. For reflectors where the area of interest is close to the edges or contains obstructions, diffraction plays a significant role; distorting the optical phase from that calculated by geometrical optics, and hence affecting our OPD measurement.

3.2 Effect of diffraction

To obtain an idea of the extent to which diffraction can be expected to affect the above method, we turn to the paper by Sherman and Chew[4], where the problem of focused fields encountering an aperture is treated analytically using the Debye integral. This integral is based on the assumption that the angular spectrum of the field goes to zero at the angles corresponding to the edge of the aperture. Furthermore, the treatment is restricted to circular apertures on converging spherical waves and hence is not directly applicable to our interferometer. Despite these two major draw-backs, we will use this analytical result for a qualitative assessment of the effects of diffraction because of the high calculation speed afforded by this approach.

By letting the object fiber-tip coincide with the centre of curvature along the axis of a spherical-cap mirror of similar dimensions to ours, we may use the analytical results stated by Sherman and Chew. See Figure 3-25 for a diagram and the notation used.

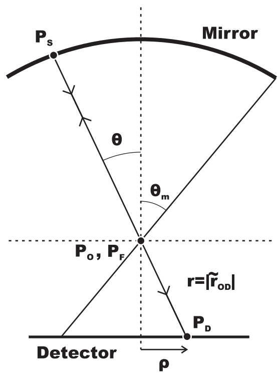

Figure 3-25 – Notation used for the treatment of diffraction by Sherman and Chew.

It is important to note that these results are cast in a form that explicitly includes the geometrical optics contribution. We see that the diffraction contribution depends on the intensity of our incident wavefront only at the aperture boundary. This is consistent with the notion of a boundary diffracted wave, and allows us to attribute the effect of diffraction exclusively to the field at the aperture boundaries. That is to say that figure errors in the interior of the reflector should not alter the contribution due to diffraction, provided they introduce no caustics at the detector. The fact that this is also true for the more general cases will be shown in section 3.5.

Geometrically illuminated region:

$$
\begin{aligned}
D(x,y,z) = & -A(\theta)\frac{e^{ikr}}{r} \\
& +\frac{1}{2}\sqrt{\frac{ik\pi}{2r}}e^{ikr-ik\rho\sin\theta_m}\times\frac{A(\theta_m)H_0^{(1)}(k\rho\sin\theta_m)\sin\theta_m}{\cos\left(\frac{1}{2}\theta-\frac{1}{2}\theta_m\right)} \\
& \times\operatorname{erfc}\left\{\sqrt{(2\text{ikr})}\sin\left(\tfrac{1}{2}\theta_m-\tfrac{1}{2}\theta\right)\right\} \\
& +\frac{1}{2}\sqrt{\frac{ik\pi}{2r}}e^{ikr+ik\rho\sin\theta_m}\times\frac{A(\theta_m)H_0^{(2)}(k\rho\sin\theta_m)\sin\theta_m}{\cos\left(\frac{1}{2}\theta+\frac{1}{2}\theta_m\right)} \\
& \times\operatorname{erfc}\left\{\sqrt{(2\text{ikr})}\sin\left(\tfrac{1}{2}\theta_m+\tfrac{1}{2}\theta\right)\right\}
\end{aligned}
\tag{3.7}
$$

Geometrical Shadow:

$$
\begin{aligned}
D(x,y,z) = & -\frac{1}{2}\sqrt{\frac{ik\pi}{2r}}e^{ikr-ik\rho\sin\theta_m}\times\frac{A(\theta_m)H_0^{(1)}(k\rho\sin\theta_m)\sin\theta_m}{\cos\left(\frac{1}{2}\theta-\frac{1}{2}\theta_m\right)} \\
& \times\operatorname{erfc}\left\{\sqrt{(2\text{ikr})}\sin\left(\tfrac{1}{2}\theta-\tfrac{1}{2}\theta_m\right)\right\} \\
& +\frac{1}{2}\sqrt{\frac{ik\pi}{2r}}e^{ikr+ik\rho\sin\theta_m}\times\frac{A(\theta_m)H_0^{(2)}(k\rho\sin\theta_m)\sin\theta_m}{\cos\left(\frac{1}{2}\theta+\frac{1}{2}\theta_m\right)} \\
& \times\operatorname{erfc}\left\{\sqrt{(2\text{ikr})}\sin\left(\tfrac{1}{2}\theta+\tfrac{1}{2}\theta_m\right)\right\}
\end{aligned}
\tag{3.8}
$$

Where $A(\theta)$ is the angular amplitude distribution (in this example taken as unity) and $H_0^{(1)}$ and $H_0^{(2)}$ are the Hankel functions of the first and second kind. θ_m is the half-opening angle of the aperture.

Figure 3-26 shows a cross-section of the field amplitude profile at a detector 35mm from focus, for a spherical mirror (R_m = 340mm) with a central obstruction (radius 25mm) illuminated with a uniform amplitude distribution.

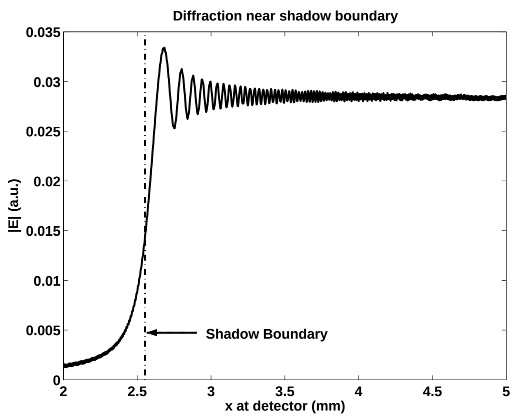

Figure 3-26 – Amplitude of diffracted field near shadow boundary

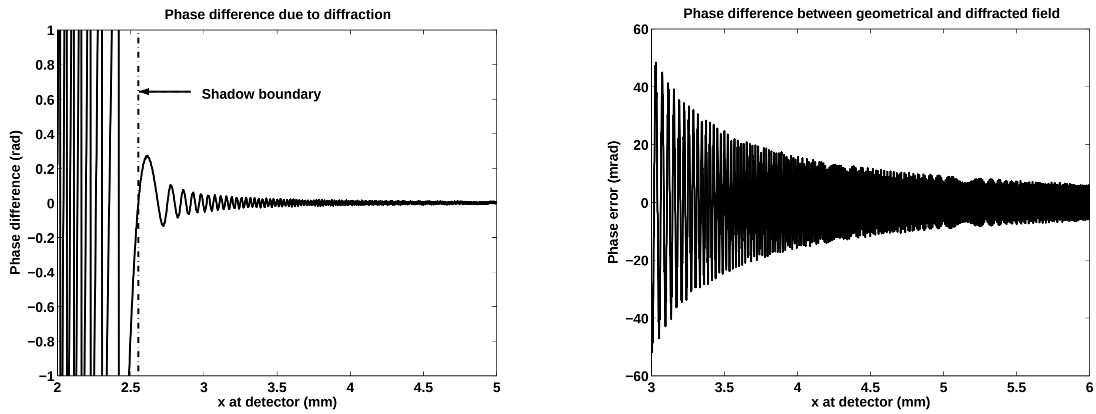

Figure 3-27 – Phase-difference with geometrical field near shadow boundary, and magnified view further away.

Figure 3-27 shows the difference in optical phase between the geometrical optics field and the diffracted field. The oscillations near the shadow boundary exceed 0.2rad, which roughly translates into a 20nm figure error introduced when retrieving the mirror shape using the raytracing approach. We see that the oscillations do not drop to zero very quickly, but increase in frequency.

The physical extent of our pixels will cause an averaging of the phase over the pixel area, so that the measured phase is actually the convolution of the actual phase with the pixel shape, sampled at the pixel locations, (see Figure 3-28). The phase fluctuations are soon under-sampled, but also attenuated by the measurement process. For our example geometry, we can say that diffraction introduces excessive figure errors over a rim 1.8mm wide, surrounding the central obstruction at the detection plane. This translates to a rim 17mm wide on the mirror itself. We will refer to this region as the diffraction rim.

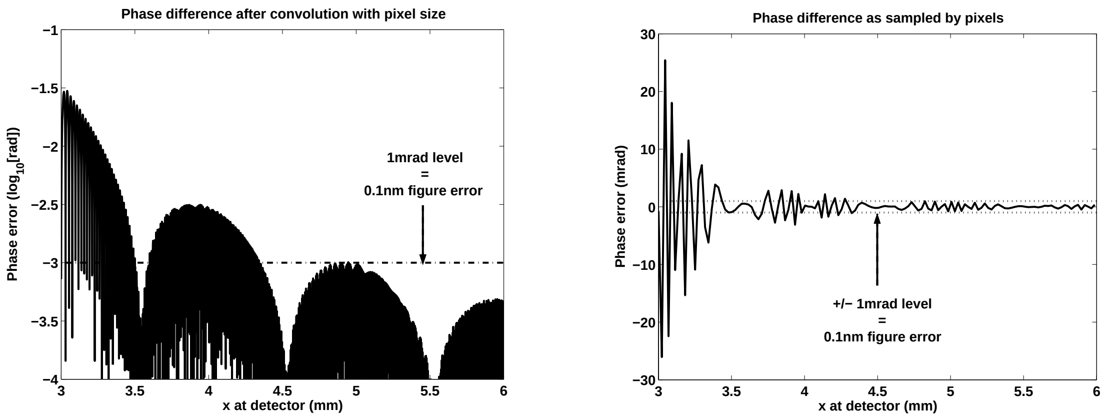

Figure 3-28 – Phase error remaining after low-pass filtering by 11μm pixels.

In the diffraction rim, the OPD measurement deviates from that predicted by geometrical optics by more than 0.1nm, as a series of oscillations in the direction of the normal to the boundary of the obstruction.

3.3 Tempered Raytracing approach

Due to the systematic and oscillatory nature of the diffraction effects, we may still be able to use our raytracing inverse propagation algorithm if we can find a way to reduce or remove these effects. Such approaches can be divided into two classes:

1. Physical reduction of diffraction effects:
   a. Damping of diffraction oscillations through vignetting.
   b. Averaging out diffraction effects with a moving aperture.

2. A-posteriori reduction of diffraction effects:
   a. Filtering out of the non-geometrical field.
   b. Parameter-fitting to the measured field.
   c. Subtracting an estimate of the non-geometrical field.

We will now briefly discuss these various approaches.

3.3.1 Physical approach

Damping

The damping approach requires us to "vignette" the boundaries of the mirror. This could be achieved either with an external mask with a tapered transmission profile near the edges, or by applying an increasingly absorbing coating to the edges of the mirror, letting the reflectivity fall off smoothly over several hundred wavelengths (see Figure 3-29). While we may state that a wider vignetting rim will cause a smaller diffraction rim, we have to assume that the figure of the mirror is altered over the entire area where this vignetting takes place. The optimum width of such a border still needs to be determined by rigorous calculations, and would depend on the physical process used to generate the vignetting mask.

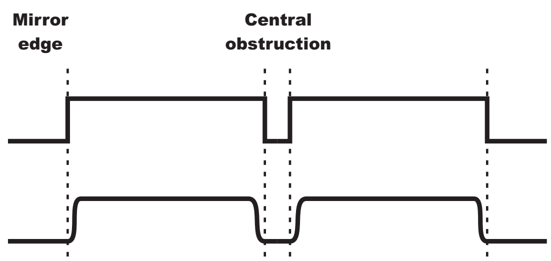

Figure 3-29 – Cross-section of mirror reflectivity before and after vignetting.

Moving aperture

The fact that the diffraction oscillations appear to have a zero mean value can be exploited by placing a smaller aperture in front of the mirror, and either varying its position, orientation or both between repeated measurements, thus averaging out the diffraction effects. We will illustrate this process by averaging 25 phase profiles as in Figure 3-27, which have been shifted with respect to each other randomly in a range of 0-0.3mm. This is equivalent to a random shift in the position of the artificially added aperture at the mirror of up to 3mm. The resulting reduction of the diffraction rim to 0.8mm at the detector (equivalent to 7.8mm at the Mirror) can be seen in Figure 3-30.

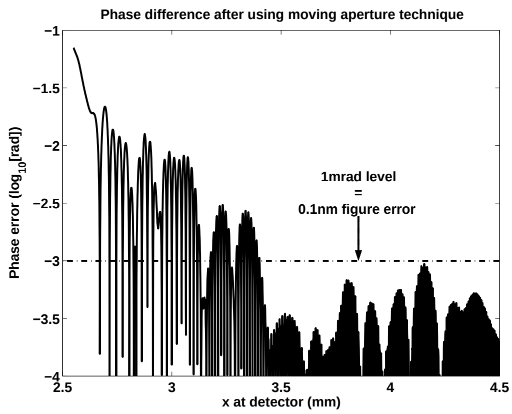

Figure 3-30 – Reduction in phase error after using moving aperture technique (c.f. Figure 3-28)

3.3.2 A-posteriori methods

Filtering

While the Zernike polynomials are generally considered an excellent basis for the representation of mirror figure errors, they will not represent the diffraction patterns well at lower orders.

The filtering approach would therefore consist of fitting, in a regularized least-squares sense (using the truncated singular value decomposition – TSVD, for example), the lower order Zernike polynomials to the OPL function on the region of interest (which excludes regions in the geometrical shadow). The resulting decomposition would have a “smoothed” appearance, which excludes patterns with high spatial frequencies, such as those generated by diffraction.

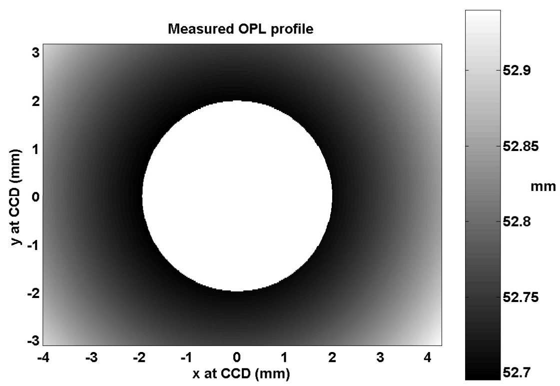

Figure 3-31 – Measured OPL profile (with arbitrary offset)

To illustrate this procedure, we will filter some data obtained from our interferometer. It should be kept in mind that this data still contains several other features and noise – notably wavefront aberrations originating from the cover-glass of the CCD. The raw OPL data, defined at a set of positions (x_D,y_D), is shown in Figure 3-31. Note the masked region in the middle, which generously covers the geometrical shadow region produced by a central hole of the mirror under test. The mask was applied manually after the measurements were taken.

We now choose, somewhat arbitrarily, a "centre" for this OPL profile (x_c,y_c), coinciding with the centre of the masked region. While the Zernike decomposition is quite stable with respect to offsets from a centre of curvature, better results can be expected if there is some degree of rotational symmetry about this centre.

In order to maximize the sensitivity of our decomposition without exceeding the domain of the Zernike polynomials, we now assign "ρ" and "θ" values to every point of the OPL profile, scaled in such a way that the maximum value of "ρ" is 1:

$$\rho = \frac{\sqrt{(x_D - x_C)^2 + (y_D - y_C)^2}}{\max\left[\sqrt{(x_D - x_C)^2 + (y_D - y_C)^2}\right]} \tag{3.9}$$

$$\theta = \text{atan}\left[(y_D - y_C), (x_D - x_C)\right] \tag{3.10}$$

We then generate the values for the first 39 Zernike polynomials on this domain (more, or fewer polynomials may be chosen), and store them as an Nx39 array of numbers, where N is the number of un-masked points on our OPL profile (166302 for the OPL data shown in Figure 3-31 above). We will denote this "image basis" matrix with "**Z**", while the Nx1 vector of OPL values will be denoted by "**p**". We are now looking for a 39x1 vector, "**a**", which minimizes the expression:

$$\left\| \mathbf{aZ} - \mathbf{p} \right\|^2 \tag{3.11}$$

In other words, we wish to find the Zernike coefficients "**a**" which give the best representation of our OPL values, "**p**", in a least-squares sense. The Moore-Penrose inverse could now be applied to this problem to find a solution for "**a**".

However, although we have only 39 Zernike polynomials to represent several thousand OPL data points, it is not guaranteed that there is a unique solution to this problem, nor is it certain that this solution will be "well-behaved"[5]. In order to "regularize" this inverse, we make use of the truncated singular value decomposition (TSVD), which strikes a balance between minimizing the norm in equation (3.11) and minimizing the sum of squares of the Zernike coefficients:

$$\left\| \mathbf{a} \right\|^2 \tag{3.12}$$

In the case when the Zernike polynomials are not linearly independent over our particular domain, this will prevent the solution from diverging. We will refer the reader to the excellent treatment of the TSVD by Tan[5] for further details of this method, and proceed by using the Matlab® function "pinv" to calculate this inverse for us:

$$\mathbf{Y} = \mathrm{pinv}(\mathbf{Z}) \tag{3.13}$$

Now, "**Y**" is a 39xN matrix, which allows us to estimate the "**a**" which minimizes (3.11) (See Figure 3-32):

$$\mathbf{a} = \mathbf{pY} \tag{3.14}$$

So that we now have:

$$\mathbf{p}_{LPF} = \mathbf{aZ} \tag{3.15}$$

$$\mathbf{p}_{HPF} = \mathbf{p} - \mathbf{aZ} \tag{3.16}$$

With:

$$\mathbf{p} = \mathbf{p}_{LPF} + \mathbf{p}_{HPF} \tag{3.17}$$

Here, $\mathbf{p}_{\mathrm{LPF}}$ is the Zernike decomposition (denoted "LPF" due to the low-pass filtering effect this decomposition has), and $\mathbf{p}_{\mathrm{HPF}}$ is the remaining, "high-pass filtered" component of the original OPL profile. Both of these are shown in Figure 3-33, where the diffraction patterns are clearly visible on $\mathbf{p}_{\mathrm{HPF}}$, along with detector related artefacts.

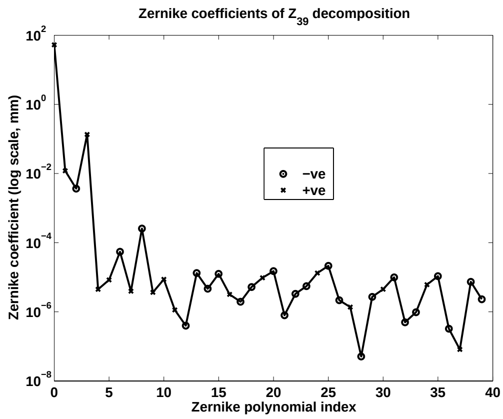

Figure 3-32 – The Zernike coefficients obtained after decomposing the OPL data in Figure 3-31.

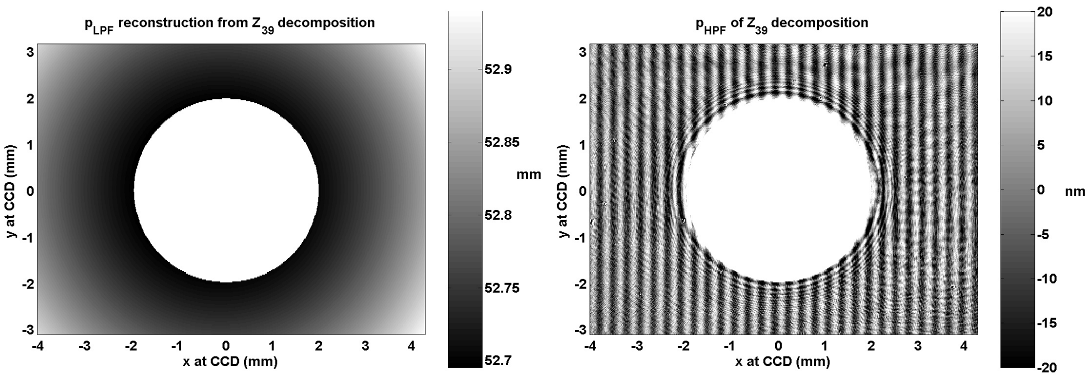

Figure 3-33 – Zernike-39 reconstruction, $\mathbf{p}_{\mathrm{LPF}}$ (left) and difference with actual OPL data (right). The prominent vertical stripes are an artefact possibly due to the CCD cover-glass, or pixel positioning errors (see section 5.2.2).

To improve the visibility of the diffraction fringes, we have removed the vertical stripes (possibly an artefact of the CCD structure. See chapter 5.2.2), and smoothed the image with a uniform 3x3 pixel kernel. The result is shown in Figure 3-34.

A certain amount of “leakage” of the diffraction pattern into the low-order Zernike decomposition is inevitable, and cannot easily be quantified without prior knowledge of the actual diffraction pattern. In addition, a lot of high spatial frequency features which may be due to actual figure errors rather than diffraction would also be filtered out by this approach. While an intuitively simple and computationally inexpensive approach, we do not consider this to be an accurate or reliable means of retrieving the reflector surface shape from our measurements.

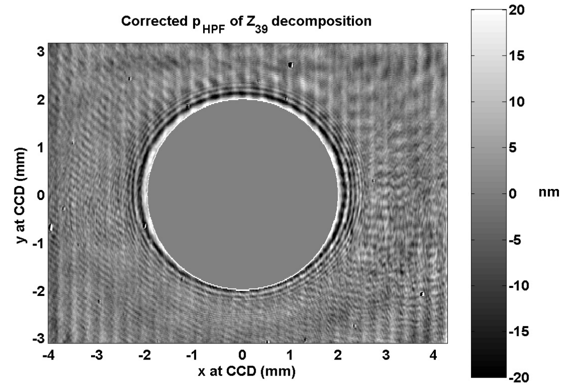

Figure 3-34 – Smoothed and corrected $\mathbf{p}_{HPF}$, accentuating the diffraction rings around the geometrical shadow region.

Parameter fitting
By parameterizing the shape of our surface (through a set of Zernike coefficients, or amplitude coefficients of other functions able to adequately represent the surface shape), as well as parameterizing the shape of the aperture (see section 3.6), we could perform an optimization of these parameters, provided we have an accurate forward-model of diffraction for our situation. An initial set of parameters, representing the nominal surface and aperture shapes, would be used to calculate the expected OPL function. The difference-squared misfit of this function with our measurement of the OPL could then be used as a cost function for a non-linear optimization routine over the parameter space for the surface and aperture shapes. Provided such an optimization converges, the resulting parameters would represent a surface and aperture shape which generates an OPL profile very similar to the one observed.

While much more rigorous than the previous approach, several problems make this approach impractical. For one, an accurate forward model of diffraction for our situation is required to simulate the theoretical OPL functions. As we will see in chapters 3.4 and 3.5, such models are indeed available, but take a significant amount of computation time. To adequately represent our surface and the shape of any apertures, a large number of parameters would need to be fitted, and consequently a much larger number of iterations of the forward problem would be required for the optimization to converge. We consider the time taken for this approach to be prohibitive.

Subtraction

As we will show shortly, the electromagnetic field at our detector is the sum of the field as predicted by geometrical optics, plus a “boundary diffracted wave” (BDW) field arising from the edges of our apertures. If a good estimate of the BDW field is available (taking into account the spatial filtering effect of our pixels), this can simply be subtracted from our measurements, leaving only the geometrical optics field, before proceeding with the raytracing approach for inversion.

Calculating the BDW field also requires an accurate knowledge of the shape of any diffracting boundaries. To this end, an algorithm has been developed to retrieve the shape of the diffracting boundary from our measurement data and a reasonable initial guess at the boundary shape.

This is so far the most promising method for dealing with diffraction. The following three sections will detail the required algorithms to perform the calculation of the electromagnetic field at the detector and retrieve the boundary shape from the measured data.

3.4 Numerical evaluation of diffraction integral.

Our model for diffraction will have to be more flexible than that of the previous section, which was restricted to axially symmetric systems. It will have to fulfil a number of requirements that rule out most asymptotic or approximate methods. The algorithm will have to be able to evaluate the diffraction integral over an arbitrary three dimensional surface shape, not restricted to a plane or a spherical cap. This is due to the combination of large NAs and severe asphericities of our reflectors. The accuracy should be adjustable, to check that the algorithm converges to within our desired precision. Also, any assumptions or approximations made must influence the final phase by less than the 0.1mrad equivalent to our 0.1nm accuracy.

For these reasons alone, a direct numerical integration of the Rayleigh-Sommerfeld integral seems the most logical choice. Once results of such an approach are available, we are at liberty to use these to evaluate the applicability of approximate methods.

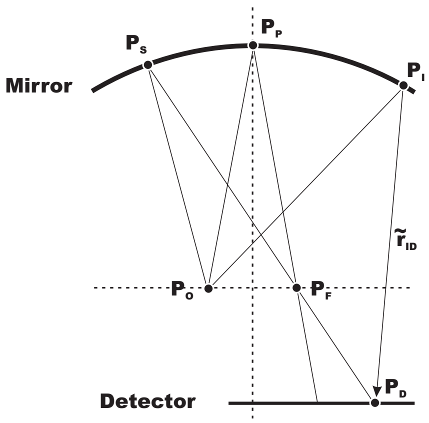

P_O – Location of point-source
P_P – Particular point.
P_F – Nominal focal point
P_D – Point of interest on detector
P_S – Point of stationary phase
P_I – Arbitrary point, for integration.
$\tilde{n}$ (P_I)– Normal to mirror at P_I
$\tilde{r}_{ID}$ – Vector from P_I to P_D

Figure 3-35 – Notation convention for brute-force approach.

By making use of symmetries inherent in our situation, we can choose an adaptive type of quadrature which is scaleable, to allow us to trade off accuracy against computation time. For the purpose of this approach, we will use the notation shown in Figure 3-35. The point of stationary phase (P_S) is approximated by the point of reflection of the geometric ray ending at P_D.

As indicated, we will assume that the mirror is illuminated by a point-source. However, we allow for the possibility of a non-uniform amplitude distribution of this light source, Am($\tilde{r}_{OI}$), to model realistic fiber output. The electric field at P_D is then given by:

$$\begin{aligned} U\left(P_D\right) &= \frac{1}{ik\lambda} \iint_{Mirror\ Surface} \frac{Am\left(\tilde{r}_{OI}\right)}{r_{OI}} e^{ikr_{OI}} \frac{1}{r_{ID}} \left(1 - \frac{1}{ikr_{ID}}\right) e^{ikr_{ID}} \cos\left(\tilde{n}(P_{ID}), \tilde{r}_{ID}\right) dA \\ &= \frac{1}{ik\lambda} \iint_{Mirror\ Surface} B\left(P_I\right) e^{iC\left(P_I\right)} dA \end{aligned} \tag{3.18}$$

We have separated out the real, slowly varying amplitude factor $B(P_I)$ and the quickly oscillating exponential, with argument $C(P_I)$. The contour-plot of $C(P_I)$ in Figure 3-36 for a typical mirror shows the apparent symmetry about the point of stationary phase. This symmetry is not perfect, especially if the mirror is aspheric. Nonetheless, by choosing a polar co-ordinate system in the plane normal to the mirror axis, centered on the point of stationary phase, we are able to choose a non-uniform spacing for the mantissa of the radial co-ordinate ρ, and a uniform spacing in the angular co-ordinate θ, to take advantage of symmetry properties.

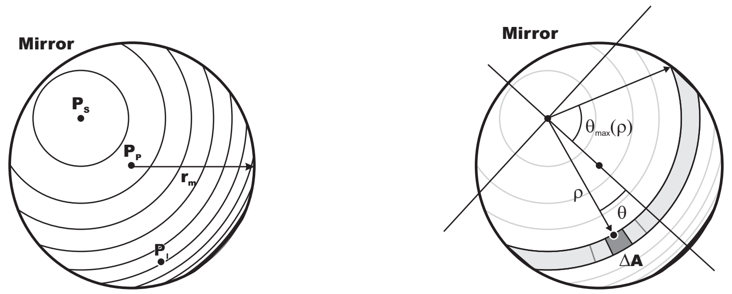

Figure 3-36 – Contours of $C(P_I)$ over mirror surface, and polar co-ordinate system centered at P_S used for numerical integration.

We re-write (3.18) to reflect this change in co-ordinate system as follows:

$$U\left(P_D\right)=\frac{1}{ik\lambda}\iint\limits_{Mirror\ Surface} B\left(\rho,\theta\right)e^{iC\left(\rho,\theta\right)}dA \tag{3.19}$$

The radial mantissa is chosen with increasing density, to account for the growing number of oscillations (Figure 3-37). As with the method of stationary phase, we see that the primary contribution to our integral will come from the stationary point and from an area around the boundary of the mirror, near the stationary point, where the contours are truncated. The region anterior to the stationary point with respect to the mirror axis contributes less to the integral due to the quick averaging that occurs by the large number of oscillations.

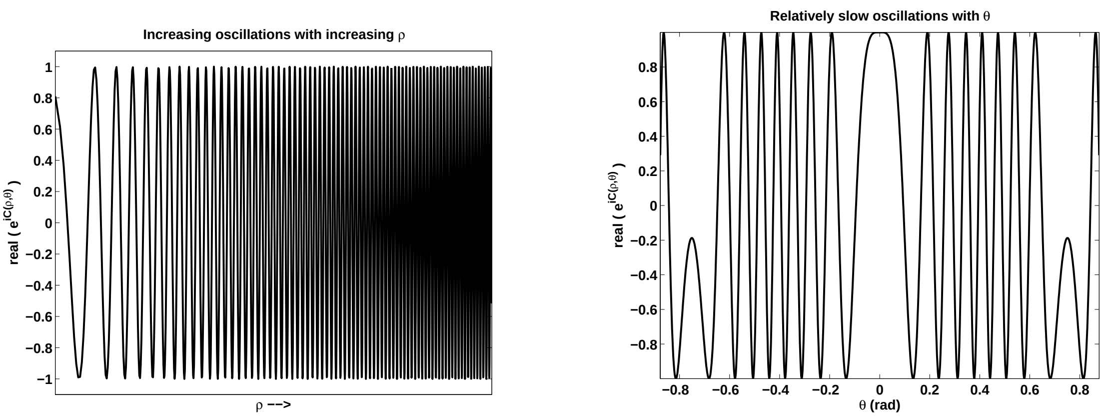

Figure 3-37 – Behavior of integrand with ρ and θ.

The results of such a brute-force calculation are shown in Figure 3-39 and Figure 3-40, where they are also compared to the results of a boundary-diffracted wave (BDW) approach[6]. The accuracy of the numerical result which samples each oscillation with 100 points is better than 1%.

It should be noted here that because our mirror is a three-dimensional surface there are a number of implications for the evaluation of the integral. First, the obliquity factor must be evaluated with respect to the local surface normal. Secondly, since (3.19) is an integral over 2 variables in the XY plane, the values of B and C are actually those of the points on the mirror which project to the corresponding point on the XY plane, and the area element, *dA* is the three-dimensional surface element on the mirror surface rather than the area element in the XY-plane.

This method is obviously very time-consuming – a single point of Figure 3-39 taking more than 10 minutes to compute on a modern PC. Fortunately, a comparison with the boundary diffracted wave (BDW) approach outlined in the next section has already resulted in a considerable saving of computational time while maintaining our desired accuracy. This new approach takes just over one second to calculate each point.

3.5 Boundary Diffracted Wave approach

The BDW approach will not only serve to perform the necessary diffraction calculations in a reasonable time, but it will clearly separate the influence of the geometrical optics field (from raytracing) from the influence of the diffracted field, which arises solely at aperture edges.

Born & Wolf[6] (Ch. 8.9) show that the problem of solving the full diffraction integral for a point source behind a planar aperture can be split into two separate problems: Finding the geometrical optics field, and adding the so-called boundary-diffracted wave field. The notation and resulting expression for the BDW field, $U^{(d)}\left(P_c\right)$, can be seen in Figure 3-38 and equation (3.20).

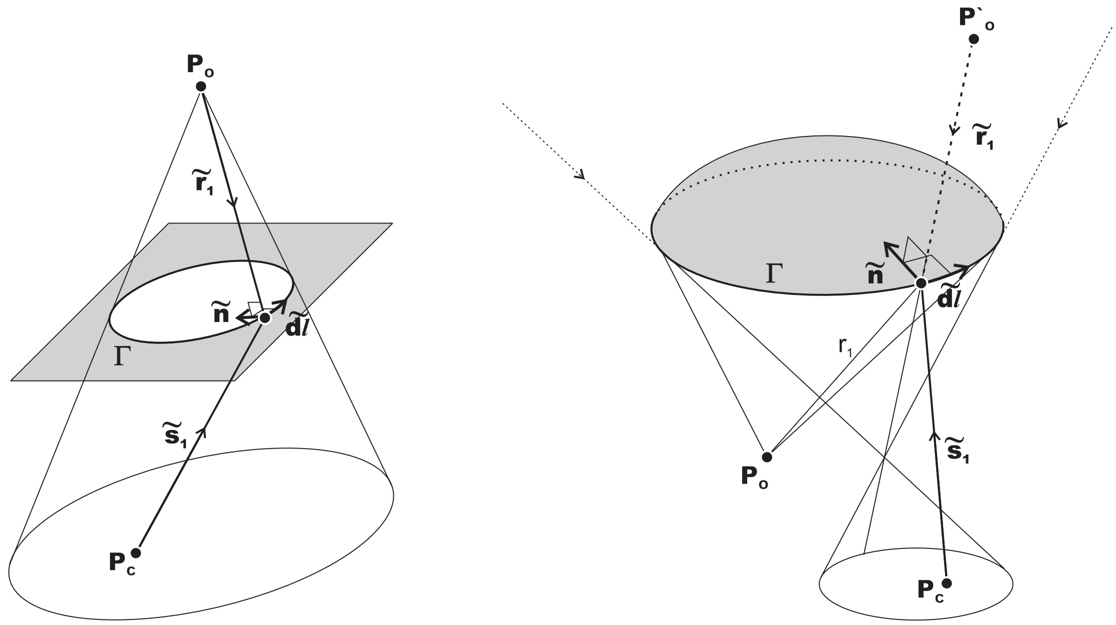

Figure 3-38 – Notation from Born & Wolf (left) and equivalent notation for our situation (right).

$$U^{(d)}\left(P_c\right)=\frac{1}{4\pi}\oint_{\Gamma} Am\left(\tilde{\mathbf{r}}_1\right)\frac{e^{ik\left(r_1+s_1\right)}}{r_1 s_1}\times\frac{\cos\left(\tilde{\mathbf{n}},\tilde{\mathbf{s}}_1\right)}{\left[1+\cos\left(\tilde{\mathbf{s}}_1,\tilde{\mathbf{r}}_1\right)\right]}\times\sin\left(\tilde{\mathbf{r}}_1,\mathbf{d}\tilde{l}\right)dl \tag{3.20}$$

The transferral of this theory to the case of a point source in front of a concave mirror of limited spatial extent is non-trivial, due to the presence of caustics in the region of the geometrical focus of this situation, and the somewhat arbitrary choice of effective point-source location for each integration path element. While our approach seems intuitively correct, the fact that the reflected wave may no longer be spherical makes the rigorous treatment of Born & Wolf inapplicable to this situation.

The phase of the geometrical optics field is calculated from the optical path lengths (OPLs), while the field intensity is calculated from the initial intensity distribution from the fiber multiplied by the ratio of the local ray-densities at the detector (*D(P$_c$)*) and just after the fiber (*D(P$_s$)*).

To calculate the local ray-density accurately, two extra rays are traced very close to each ray of interest over an equal distance before reflection, their end-points forming a triangular surface element with an area inversely proportional to the local ray density. The ray of interest is then traced on to the detection plane, and the two accompanying rays are traced through the system to the same length, before forming another triangle from their end-points with an area inversely proportional to the ray density at the detector:

$$U^{(g)}\left(P_c\right) = Am\left(P_s\right)\sqrt{\frac{D\left(P_c\right)}{D\left(P_s\right)}}e^{ik\mathrm{OPL}\left(P_c\right)} \tag{3.21}$$

The total field is then given by the sum of the geometrical and boundary-diffracted fields:

$$U\left(P_c\right) = U^{(g)}\left(P_c\right) + U^{(d)}\left(P_c\right) \tag{3.22}$$

Each separate problem is significantly easier to solve than the full Sommerfeld diffraction integral, even when using stationary phase methods.

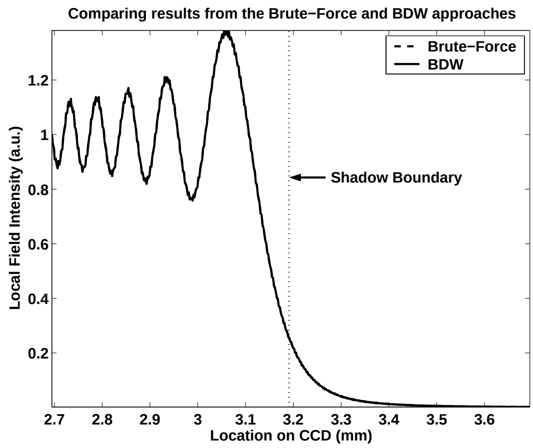

Figure 3-39 – Field intensity at detector calculated with brute-force and BDW models.

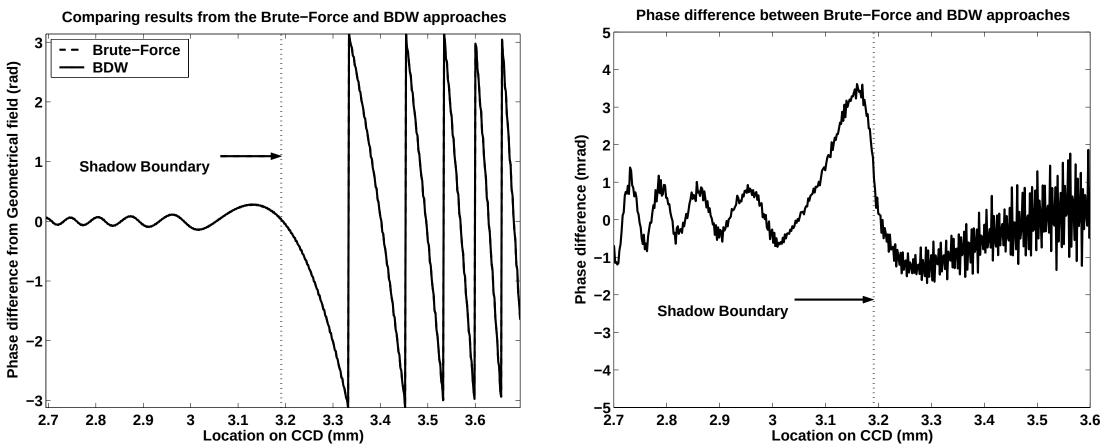

Figure 3-40 – Phase deviation from geometrical optics, calculated with brute-force and BDW methods (left) and difference between phase calculated with the two methods (right).

Since our approach is not entirely rigorous, we need to establish whether it still yields sufficiently accurate results. To do this, we will calculate the field at our detector for a realistic set of parameters, and compare results obtained using this new method with the results obtained using our rigorous, numerical brute-force integration of the Sommerfeld integral. The results of this comparison can be seen in Figure 3-39 and Figure 3-40. The difference between the two methods is limited only by the numerical accuracy for the brute-force approach, while the computational speed of the BDW approach is significantly faster (500s for the BDW approach, vs. 63hrs for the brute-force approach for Figure 3-39)

The geometrical phase in the shadow-region was calculated by simply extending the nominal mirror shape. The phase-difference between the geometrical and diffracted fields in the shadow region is therefore somewhat arbitrary.

This is a practical method for calculating the diffraction pattern to be subtracted from the measured phase profile to yield the purely geometric component, which in turn can be inverted with the raytracing based inversion technique.

3.6 Boundary retrieval method

While we are now able to quickly and accurately calculate the diffraction field for a given mirror form and boundary shape, we need a-priory knowledge of both in order to subtract the correct BDW field. This section will deal with a means of retrieving the shape of the boundary of the mirror from the measured data.

The method relies on the fact that the diffraction patterns are of a characteristic form, very different from both random noise and the mirror form itself:

$$\text{OPL} = \text{Geometrical Phase} + \text{BDW Phase} + \text{Noise} \tag{3.23}$$

If we can now generate a complete set of basis functions which represent the BDW phase contribution from all reasonable boundary shapes, we should be able to de-compose the OPL function into these “boundary diffracted wave basis functions” (BDWBFs), and hence deduce the shape of the aperture which gave rise to the observed OPL, very much in the same way as the filtering approach outlined in section 3.3.2 deduced the best fit Zernike coefficients to the OPL.

Although we may expect the geometrical phase component of the OPL to be naturally orthogonal to the BDWBFs (due to the oscillatory nature of the latter), their large magnitude (~mm rms) compared to the magnitude of the BDWBFs (~nm rms) magnifies even small leaks of components of the geometrical phase along the BDWBFs. For that reason, we will first subtract the low-order Zernike components from both the OPL and our BDWBFs, to reduce this “leakage”. The noise is implicitly assumed to be largely orthogonal to any systematic basis functions.

It remains for us to construct such a complete set of BDWB functions which will represent all reasonable diffracting boundary shapes. We may assume good a-priory knowledge of this boundary shape, both from the mirror manufacturer’s specifications and by visual inspection of the interferograms. The vast majority of boundary shapes can be represented parametrically:

$$\left[x(\theta), y(\theta), z(\theta)\right] = \left[\rho(\theta)\cos(\theta), \rho(\theta)\sin(\theta), z(x(\theta), y(\theta))\right], \tag{3.24}$$

for some “centre” of the boundary (x_{BC},y_{BC}), where $z(x,y)$ is the mirror height profile. The boundary is therefore completely defined by the function $\rho(\theta)$. Let us therefore fix such a centre, and write our default guess at the aperture shape as $\rho_{def}(\theta)$.

We can now represent all reasonable aperture shapes (including position offsets etc.) with the addition of a number of finite Fourier components to this function:

$$\rho(\theta) = \rho_{def}(\theta) + \sum_{j=1}^{M} b_j F_j(\theta), \tag{3.25}$$

with:

$$F_j(\theta)=\begin{cases}\sin\left(\tfrac{1}{2}j\theta\right) & j \text{ even}\\ \cos\left(\tfrac{1}{2}\left[j-1\right]\theta\right) & j \text{ odd}\end{cases} \tag{3.26}$$

For the case of an almost circular aperture, Figure 3-41 shows how the various Fourier components affect the shape of the aperture.

The aperture is now defined in terms of the default function and the vector of Mx1 Fourier coefficients "**b**". Using our Zernike decomposition of the OPL together with the raytracing IPA to obtain a default mirror shape, $z_{def}(x,y)$ we can calculate the default BDW phase map as outlined in the previous section. The process which maps a particular mirror form and boundary shape to a BDW phase map on the detector (x_D,y_D) will be denoted by: $BDWP(x_D,y_D|z(x,y),\rho(\theta))$.

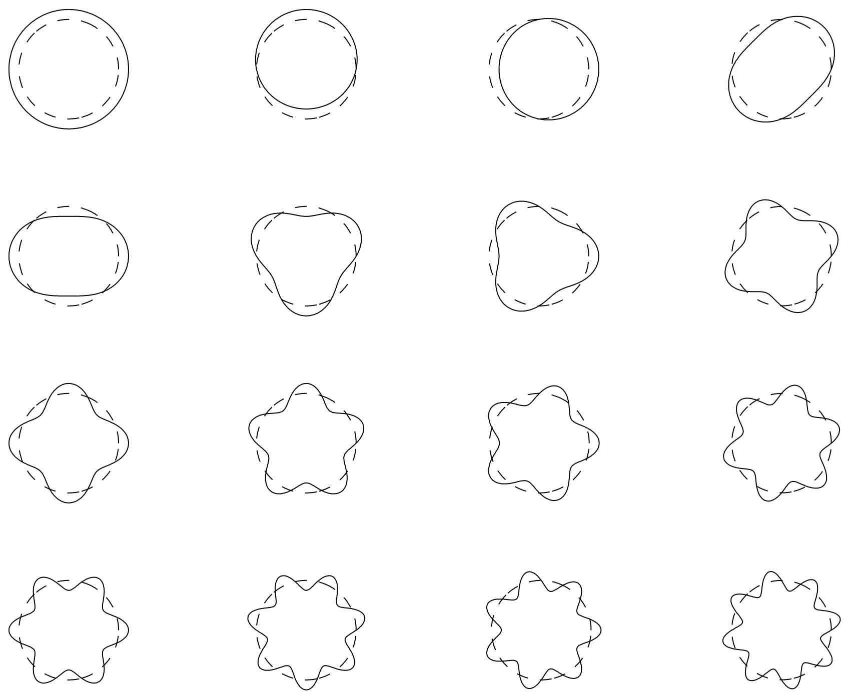

Figure 3-41 – Default estimate of boundary shape (dashed) aberrated by the first 16 Fourier functions.

It can be observed that this BDWP function is sufficiently well behaved to be, to first order, linear under addition of small Fourier components as outlined above. We may therefore write:

$$\begin{aligned}BDWP\left(x_D,y_D\,|\,z(x,y),\rho(\theta)\right)&\approx BDWP\left(x_D,y_D\,|\,z(x,y),\rho_{def}(\theta)\right)\\&\quad+\sum_{j=1}^{M}b_j BDWBF_j\left(x_D,y_D\,|\,z(x,y),\rho_{def}(\theta)\right)\end{aligned} \tag{3.27}$$

Where our basis functions (the BDWBFs discussed earlier) are now given by:

$$\varepsilon \times BDWBF_j\left(x_D, y_D \mid z(x,y), \rho_{def}(\theta)\right) \equiv BDWP\left(x_D, y_D \mid z(x,y), \rho_{def}(\theta) + \varepsilon F_j(\theta)\right) - BDWP\left(x_D, y_D \mid z(x,y), \rho_{def}(\theta)\right) \quad (3.28)$$

An example of the BDWBFs corresponding to a circular initial guess at the boundary shape can be seen in Figure 3-43.

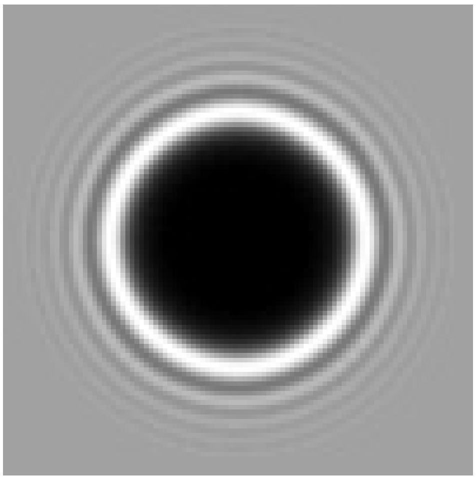

Figure 3-42 – Example of a BDW pattern at the detector for a circular boundary shape (arbitrary scale).

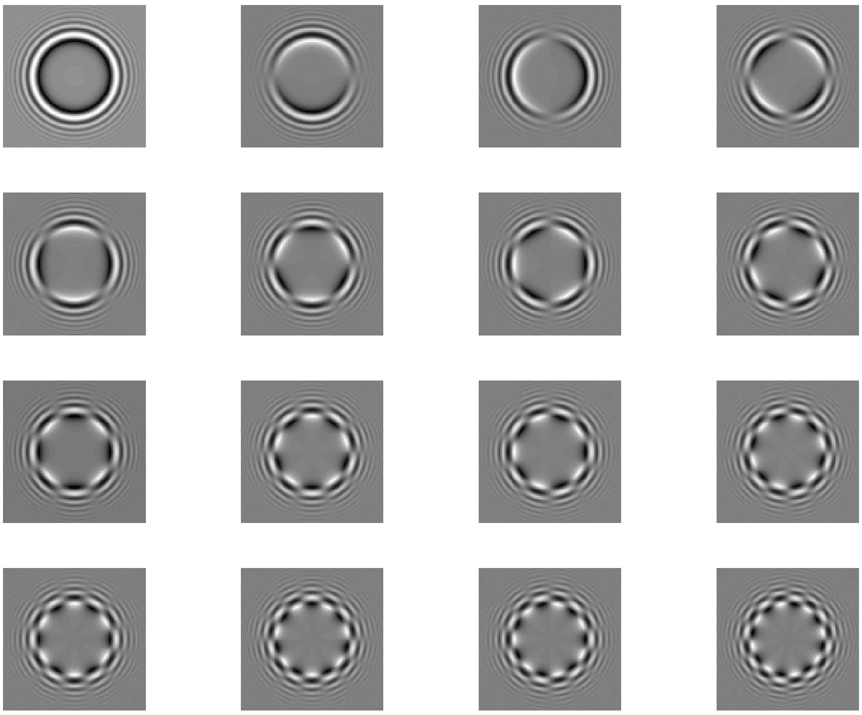

Figure 3-43 – First 16 BDW basis functions for a circular default boundary shape, corresponding to the 16 aberrated boundary shapes of Figure 3-41.

Returning to the notation of section 3.3.2, we can now express our desire to find the coefficients "**b**" which best match our BDW phase once again in matrix notation, by searching for the "**b**" which minimizes:

$$\left\| \mathbf{p} - \mathbf{aZ} - \mathbf{d} - \mathbf{bB} \right\|^2 \tag{3.29}$$

Where "**d**" is the Nx1 BDW phase pattern due to the default boundary shape, $BDWP(x_D,y_D|z(x,y),\rho_{def}(\theta))$, and "**B**" is the NxM matrix of all the BDWBFs as defined above. Following the arguments in section 3.3.2 leads to the regularized solution for "**b**":

$$\mathbf{b} = \left[\mathbf{p} - \mathbf{aZ} - \mathbf{d}\right] \mathrm{pinv}(\mathbf{B}) \tag{3.30}$$

From this solution and equation (3.25), we can obtain our improved estimate of the boundary shape.

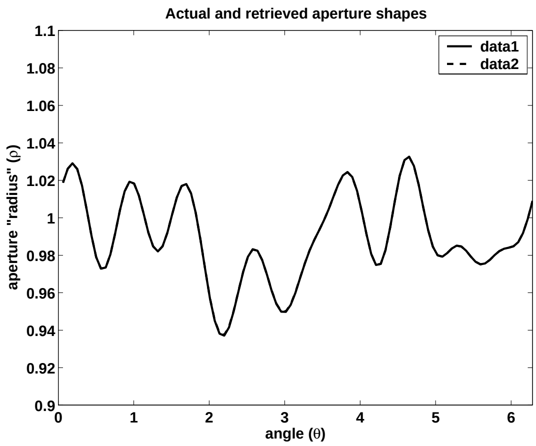

Figure 3-44 – Example of a simulated boundary shape (solid) and the retrieved boundary shape (dashed, indistinguishable) from a circular initial guess of the boundary shape ($\rho_{def}(\theta)$ = 1).

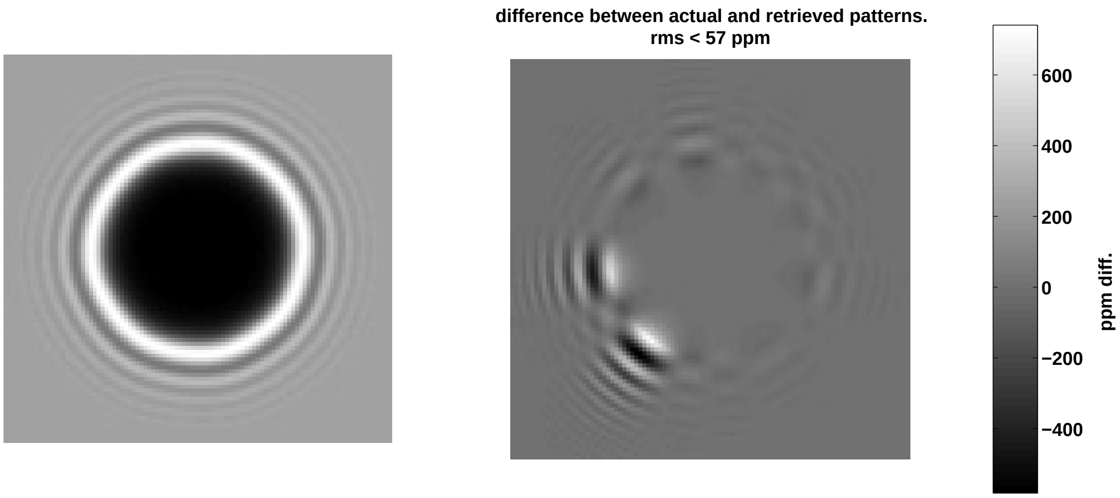

Figure 3-45 – BDW pattern (left) produced by simulated boundary shape in Figure 3-44, and the deviation from the BDW pattern from the retrieved boundary shape (right).

The results of a simulation of such a procedure are shown in Figure 3-44 and Figure 3-45. This then allows us to calculate the BDW phase using the methods in the previous section, which can be subtracted from our OPL map to give an estimate of the purely “geometrical” OPL map. This finally allows us to retrieve an estimate of the mirror shape using the raytracing IPA of section 3.1.2.

We may iterate this procedure, as outlined in Figure 3-47, using our new mirror shape- and boundary shape- estimates as starting points to improve the accuracy of our estimates.

While the above examples are of an internal boundary (a “hole” in the reflector), external boundaries can be treated with the same approach. Furthermore, it is possible to retrieve the boundary shape from measurements of the diffraction pattern’s intensity rather than of its phase, using the methods discussed in this section.

3.7 Direct back-propagation

One last potential method for determining the shape of the reflector from our measurements of the OPL is the propagation of our measured wavefront back to the nominal mirror position. We cannot use this method to determine the mirror shape directly. Instead we will be deriving a retardation to the ideal field expected at the nominal surface, as described in the previous section, which can then be used to deduce the surface shape which would be equivalent to this retardation. (See Figure 3-46)

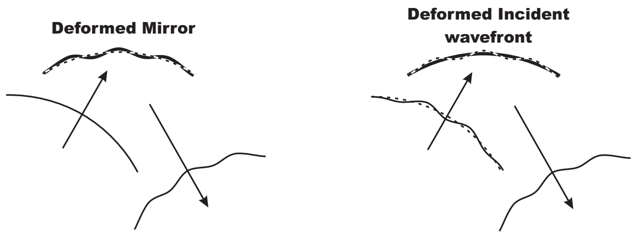

Figure 3-46 – Equivalence of mirror form aberrations and deformations in the incident wavefront.

While being the most intuitively appealing, this method unfortunately suffers from several problems. In terms of computation time, back-propagation would still be prohibitive in the absence of a more efficient forward algorithm (the BDW approach is not applicable here, due to the absence of sharp boundaries in the diffracted field at the detector). Furthermore, while the method would give the correct mirror shape if the exact diffraction pattern at the detector were available to us, the under-sampling and spatial filtering inherent in our measurement process will doubtlessly introduce a new class of distortions in the retrieved mirror shape.

3.8 Conclusion

We can now combine the results from the previous chapters into an over-all strategy for an accurate retrieval of the measured surface shape. This procedure is illustrated in the flowchart below:

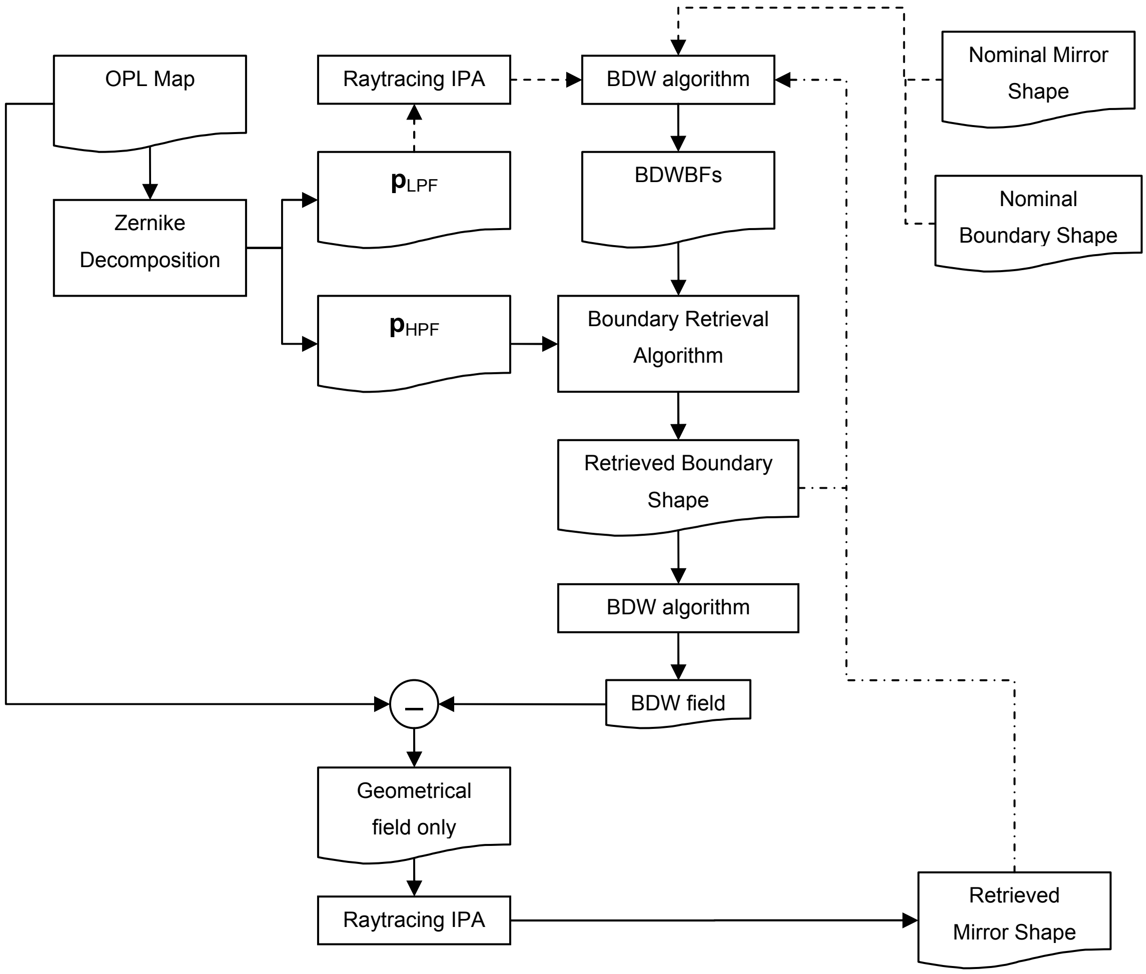

Figure 3-47 – Flowchart summarizing the total inverse propagation algorithm, yielding the mirror shape from the OPL measurements.

We have presented several potential approaches to retrieving the shape of reflectors under test in our novel interferometer. For reflectors where diffraction effects are negligible, a geometrical optics approach yields fast and accurate results. The geometrical optics approach in the presence of obstructions has been determined to be valid except within a so-called “diffraction rim” around the projected obstructions. Various ways to reduce the detrimental effects of diffraction were discussed. Motivated by one of these approaches, a novel scheme for accurate diffraction calculations was developed. Based on this scheme, a “boundary retrieval” algorithm was developed to accurately retrieve the shape of any obstructions present. A combination of the above methods yields a fast hybrid scheme capable of dealing with the diffraction effects present to retrieve the shape of reflectors under test from measurement data from our novel interferometer.

3.9 References

[1] M. L. Krieg and J. J. M. Braat, "Inverse propagation algorithm for angstrom accuracy interferometer," in *Interferometry XII: Techniques and Analysis*, vol. 5531, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, K. Creath and J. Schmit, Eds., 2004.

[2] R. G. Klaver, "Novel interferometer to measure the figure of strongly aspherical mirrors." Delft: Delft University of Technology, 2001.

[3] M. Abramowitz and I. A. Stegun, *Handbook of mathematical functions; with formulas, graphs, and mathematical tables*, 7th pr. ed: Dover, 1970.

[4] G. C. Sherman and W. C. Chew, "Aperture and Far-Field Distributions Expressed by the Debye Integral-Representation of Focused Fields," *Journal of the Optical Society of America*, vol. 72, pp. 1076-1083, 1982.

[5] S. M. Tan and C. Fox, *Inverse Problems*. Auckland: University of Auckalnd, 2001.

[6] M. Born and E. Wolf, *Principles of optics; electromagnetic theory of propagation, interference and diffraction of light*, 9th ed: Pergamon, 1999.

4 Light Source

The multiple wavelength light-source used in our interferometer provides up to three stable wavelengths, used to increase the unambiguity range of the interferometer beyond the λ limit of a single wavelength interferometer. Several challenges are involved in the design of such a light source. The wavelengths will have to be stabilized to specific desired wavelengths, the OPD drifts within the light source will need to be minimized and cross-talk between the reference and object beams will need to be minimized, all while maintaining as much flexibility as possible with respect to potential sensor types. Given a fast-sampling type detector, or a 6+ bucket lock-in pixel CCD as described in chapter 5, the light source described here should even allow measurement at all three wavelengths simultaneously.

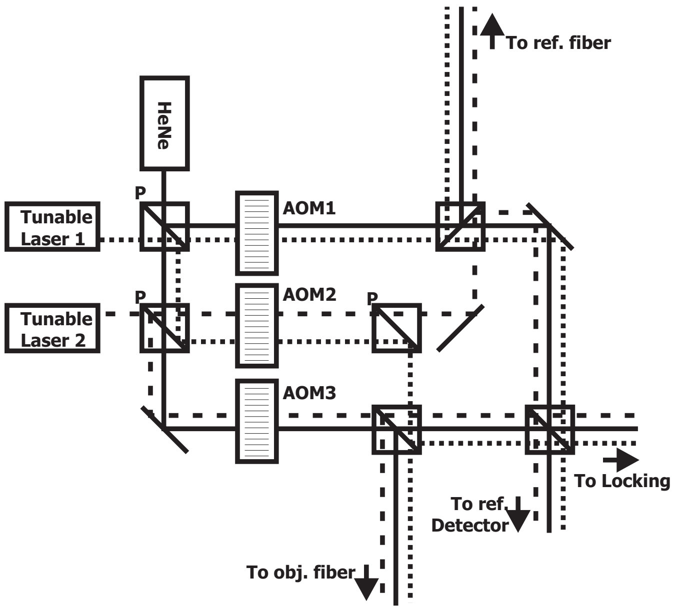

Figure 4-1 – Schematic diagram of the 3λ lightsource. Beamsplitters marked with a “P” are polarizing.

The light source for our interferometer should provide light which, on arrival at the sensor, carries enough information about the shape of the mirror to determine its shape to within our accuracy requirements. As with classical interferometry, we will end up measuring optical path differences (OPDs) between two optical paths, using the wavelength(s) of our light source as length standards. It is therefore necessary to be able to guarantee a certain level of stability for the wavelengths used. Section 4.4 examines this topic in detail.

Several methods of OPD measurement are available to us, each with their own advantages: Heterodyne interferometry, phase-shifting interferometry and superheterodyne interferometry. While our light source is capable of producing light compatible with all three methods, the available sensors restrict us to the first two methods. Figure 4-1 shows the full, 3-λ set up capable of the various

OPD detection modes, described in a recent paper[1]. Depending on which detection mode is chosen, along with the number of wavelengths used, this set-up can be simplified considerably. These set-ups will be shown and discussed in the subsequent sections.

4.1 Simultaneous heterodyne detection

Heterodyne detection is a well known[2,3] and particularly accurate method for determining the OPD between two optical paths. As opposed to standard interferometry, where the OPD information is encoded on the intensity of an interference pattern, heterodyne interferometry encodes the OPD on the phase of an oscillating interference pattern. Heterodyne interferometry is inherently more accurate than standard interferometry because it is less sensitive to intensity noise.

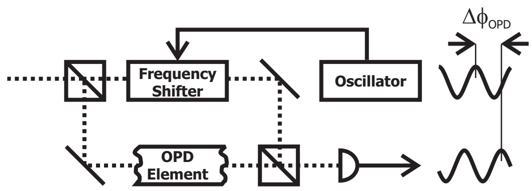

Figure 4-2 – Schematic diagram of a heterodyne interferometer

Figure 4-2 shows a schematic diagram explaining the principle of heterodyne interferometry. The optical phase difference between the object and reference beam is encoded as the phase of the amplitude modulated output signal, which has a frequency equal to the frequency difference introduced by the frequency-shifter. Mathematically, we can express this as follows: Let U_o and U_r be the electric field due to the object and reference beams respectively. ν_{opt} is the optical frequency of the reference object beam, ν_{shift} is the frequency shift introduced in the reference beam and $\Delta\phi_{opd}$ is the optical phase difference introduced by the object under test in the object arm of the interferometer:

$$
\begin{aligned}
U_o &= A_1 e^{i2\pi\nu_{opt}t+\Delta\phi_{OPD}} \\
U_r &= A_2 e^{i2\pi[\nu_{opt}+\nu_{shift}]t}
\end{aligned}
\tag{4.1}
$$

$$
\begin{aligned}
I_{\text{interf}} &= \left(U_o + U_r\right)\left(U_o + U_r\right)^* \\
&= A_1^2 + A_2^2 + A_1 A_2 \cos(2\pi\nu_{shift} - \Delta\phi_{OPD})
\end{aligned}
\tag{4.2}
$$

The intensity at the output of the interferometer, I_{interf}, shows a distinct frequency component at the difference frequency, ν_{shift}, This intensity modulation has a phase equal to the optical phase difference due to the test object. The goal of a heterodyne interferometer is to measure this phase difference as accurately as possible.

4.1.1 Wavelength multiplexing

Of course, the optical phase difference observed is due to a physical path-length difference between the reference and object beams:

$$
\Delta\phi_{OPD} = 2\pi \frac{L_{OPD}}{\lambda}\ , \tag{4.3}
$$

which is where the dependence on the wavelength comes in. Different wavelengths will give a different proportionality constant between the path-length difference and the optical phase difference, which can be exploited, as shown in section 4.3, to lead to an increased unambiguity range. For now, let us simply accept that the use of several wavelengths is desirable, and proceed with the problem of implementing a heterodyne interferometer measuring at several wavelengths simultaneously.

For a simultaneous measurement at all three wavelengths, it is necessary to send all three wavelengths through the same object-path, so that all beams measure the same parts of the object. This however means that the signals from the three different wavelengths also need to be separated out again afterwards, so that the measured phase differences can be uniquely attributed to their respective wavelengths. If the wavelengths are sufficiently different from one another, they may be easily separated out by spectral means, such as filters, gratings or prisms. The wavelengths required for our purposes (see section 4.3.1) are unfortunately too close to one another for spectral separation. Also, these means of separation would require additional optical elements in the interferometer, hence introducing unwanted OPD errors. We therefore need an alternative means of de-multiplexing the signals due to the three wavelengths.

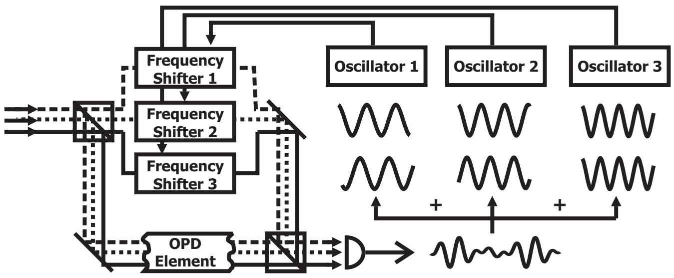

Figure 4-3 – Diagram of a three-wavelength heterodyne interferometer. The three wavelengths are actually co-propagating in the object arm, and have only been separated for visual reasons.

By providing a different frequency shift for each wavelength (Figure 4-3), we obtain an output signal which is intensity modulated by a superposition of the three shift-frequency signals. Cross-terms arising from interference between the different wavelengths can be neglected, as they are in the GHz range, and are therefore filtered out by the bandwidth-limited response of the photodiode.

The signals can be said to be shift-frequency multiplexed. The three frequencies can then be de-multiplexed by electronic means which is usually done automatically during the phase-detection step, with phase detection methods such as mixing or lock-in detection for example. Counters form an exception here, and the three signals would have to be separated using band-pass filters first, before being passed to the counters for phase measurement.

The shift frequencies must also be compatible with the sensors' detection capabilities, which unfortunately vary widely between sensors. For a CSEM-type lock-in pixel sensor as described in chapter 5.4, the shift frequencies should ideally be in the MHz range, whereas a fast-sampling type sensor is limited to a few kHz or less.

The reasons for the complicated nature of the set-up shown in Figure 4-1 should now be apparent. By tracing the path of the three different wavelengths, it can be seen how the three acousto-optic modulators are used as frequency shifters to produce our three shifting frequencies. If we denote the three frequencies driving the AOMs as ν_{AOM1}, ν_{AOM2} and ν_{AOM3}, the shift frequencies will be given by:

$$\begin{aligned} \nu_{shift1} &= \nu_{AOM3} - \nu_{AOM1} \\ \nu_{shift2} &= \nu_{AOM2} - \nu_{AOM1} \\ \nu_{shift3} &= \nu_{AOM3} - \nu_{AOM2} \end{aligned} \tag{4.4}$$

These can be varied between 0-5MHz with μHz resolution using the digital synthesizer modules driving the AOMs.

4.1.2 Cross-talk

An important feature of the arrangement shown in Figure 4-1, is the complete absence of cross-talk (often referred to as “heterodyne nonlinearity” in the literature[4]) for the HeNe laser, our primary metrology wavelength. The acousto-optic modulators have some intrinsic birefringence, and will therefore affect the state of polarization of the incoming beams slightly. By carefully tracing the beams of either one of the two tuneable lasers in Figure 4-1, it should become obvious that the polarizing beam-splitter after AOM2 will therefore not act exactly as shown, instead causing a small amount of mixing between light destined for the reference and object arms. It can also be seen that the reference and object beams of the HeNe laser will experience no such mixing before entering the fibers, even in the face of strongly birefringent AOMs or imperfect beam-splitters. For the tuneable laser wavelengths, this cross-talk will limit the accuracy of the phase-measurement by introducing a nonlinear phase error which, to first order, is given by[4, 5]:

$$\Delta\phi_{nonlin.} \approx \varepsilon \sin\left(\Delta\phi_{OPD}\right) \tag{4.5}$$

Where $\varepsilon \ll 1$ is the fractional amount of amplitude cross-talk. Even so, this error can be calibrated out, provided the degree of mixing stays constant. Methods discussed in Chapter 5.2.2 automatically reduce the influence of this type of error.

4.1.3 Two-wavelength set up.

If only two wavelengths should be required, the set-up can be simplified considerably, as shown in Figure 4-4. In addition to eliminating cross-talk for both wavelengths, this set up has the added advantage of permitting the use of a very much simplified frequency locking scheme, as described in Section 4.4.4.

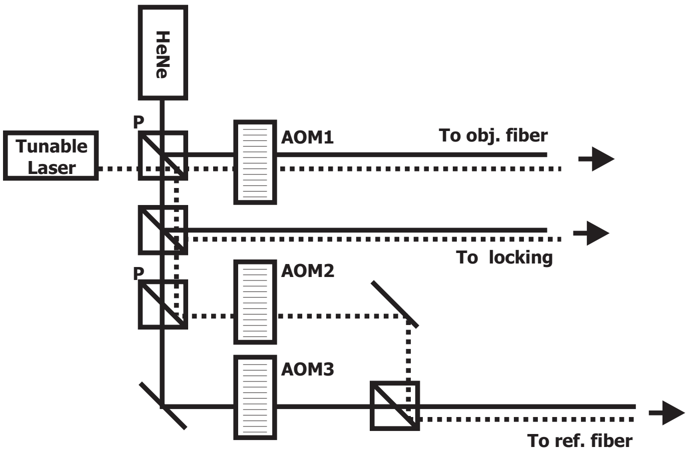

Figure 4-4 – Two-wavelength light-source permitting simultaneous 2λ heterodyne interferometry.

If, instead of measuring the two wavelengths simultaneously, it should be sufficient to measure the two wavelengths sequentially, the set-up can be simplified even more, to that shown in Figure 4-5. For the mirror measurements reported in this thesis, the configuration in Figure 4-4 was used. A photograph of the actual set up can be found at the end of this chapter in Figure 4-31.

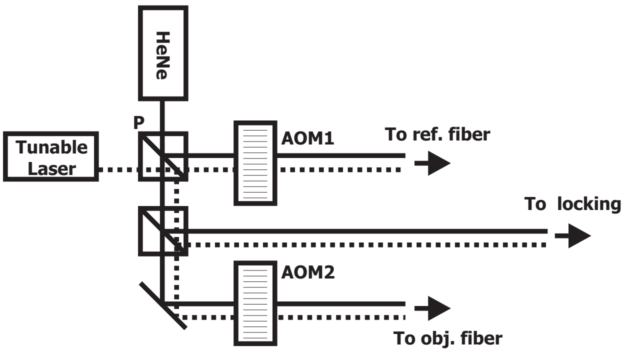

Figure 4-5 – Simplified light source for sequential 2λ interferometry

4.2 Phase shifting interferometry

The same set-up as shown in Figure 4-5 above can also be used to perform phase shifting interferometry (PSI). PSI has a long and rich history[2, 6], and has been the method of choice for 2D metrology purposes for many years[7].

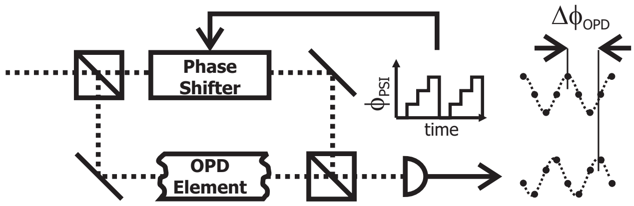

Figure 4-6 – Schematic diagram of a Phase-shifting interferometer

Figure 4-6 illustrates the working of PSI schematically. The similarities with heterodyne interferometry should immediately be obvious. In fact, PSI may be regarded as a discretized version of heterodyne interferometry. For a fixed phase-increment between measurements of size 2π/n, we can write the intensity at the photodiode at phase-step number 'k' as:

$$I_k = A_1^2 + A_2^2 + 2A_1A_2\cos\left(\frac{2\pi k}{n} + \Delta\phi_{OPD}\right) \tag{4.6}$$

After which the phase can be calculated from 'n' successive intensity measurements using the discrete Fourier transform (see 14.8.2 of ref.[8]):

$$\Delta\phi_{OPD} = \tan^{-1}\left[\frac{\sum_{k=0}^{n-1} I_k \sin\left(\frac{-2\pi k}{n}\right)}{\sum_{k=0}^{n-1} I_k \cos\left(\frac{-2\pi k}{n}\right)}\right] \tag{4.7}$$

We need not restrict ourselves to uniform- or even known- phase steps, as described in[7]. However, the most robust methods of phase-retrieval do use phase steps of 2π/n, and hence we will confine ourselves to this class of phase-shifts.

The main advantage of PSI over heterodyne interferometry is that it can use a standard CCD to perform 2D OPD measurements. The CCD records a static interferogram after every phase shift, and the phase can be calculated at every pixel over a number of frames as in the single-point case. Up until recently, it was impossible to perform 2D heterodyne measurements with a comparable pixel-resolution.

4.2.1 AOM phase stepping

Phase-stepping has traditionally been performed using mechanical phase-shifters, such as piezo-mounted retro-reflectors or gratings, which suffered from significant phase-shifting inaccuracies due to piezo nonlinearity and hysteresis[2]. Our set up offers the alternative method of acousto-optic phase shifting, which can be regarded as analogous to the shifted-grating method. Here, all AOMs are driven at the same frequency, but the phase of one of the AOMs is changed by a fixed increment before recording each interferogram. This approach has been previously reported (Bass[2], chapter 30.7), and is more accurate than mechanical phase-shifting, since neither hysteresis nor nonlinearity affect the phase-shifting accuracy, nor does the phase-shifter require calibration. The digital signal synthesizers allow the phase of the driving signal to be adjusted with 32bit accuracy – equivalent to an accuracy of 1.5 nrad.

While some of the drawbacks of piezo phase shifters have been significantly reduced by integrating capacitive position sensors in their design, the AOM option is still more attractive here, because it allows the operation of the light source in both heterodyne- and PSI-mode without any modifications to the setup.

To perform a PSI measurement, a standard video frame-rate CCD can be used and the vertical synchronization signal can be used to trigger a phase-shift of a fixed magnitude at the signal synthesizer before every frame. A frame-grabber is then used to capture each successive frame, until enough frames are captured to perform the phase measurement. The frames are then post-processed on a PC as described in[7] to obtain the OPD at each pixel.

The disadvantage over heterodyne detection is that it is slower, and hence phase-drifts during the measurement process have a greater influence on the measurements. These effects can be reduced to some extent (as described in chapter 5.2, where we treat PSI in more detail), but the intrinsic advantage of heterodyne interferometry remains. Furthermore, PSI does not allow for simultaneous measurement of several wavelengths. These will have to be performed sequentially.

4.3 Phase ambiguity

Despite the fact that our interferometer is designed to minimize the range of OPDs over the sensor (see chapter 6), a certain range of OPDs is unavoidable, since the object- and reference-beams are necessarily angularly separated, and the mirror is severely aspherical.

Klaver[9] estimated the range of OPDs that would result from an optimum positioning of the interferometer components, and concluded that an OPD range of 3mm was to be expected for the severest aspheres. This means that the OPDs between adjacent pixels of the sensor could frequently exceed one wavelength of light. For heterodyne- and PSI-interferometers, this would pose a virtually insurmountable challenge, due to the resulting phase ambiguity. These interferometers can only measure the optical *phase* difference between two beams, translating to an OPD measurement with modulo λ. It would therefore be impossible to uniquely determine the number of integer wavelengths by which the OPD at two adjacent pixels differs. (See Figure 4-7 for an analogy)

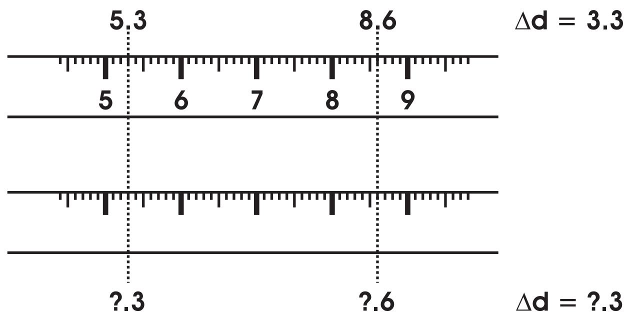

Figure 4-7 – Analogy to illustrate concept of phase ambiguity

At the very least, our light source should allow us to decode the *absolute* optical path difference between pixels and preferably even the total absolute path length of the two beams. While the latter calls for an unambiguity range of 3mm, the former constraint is less severe, requiring an unambiguity range just larger than the maximum OPD between any two adjacent pixels.

The phase ambiguity can be resolved in a number of ways, each placing different requirements on the light-source and the detector. One such method, called frequency modulated continuous wave (FMCW) interferometry, was proposed by Klaver[9] in his thesis as a viable candidate. This method essentially amounts to an extremely sensitive time-of-flight measurement of a frequency-chirped beam. We have exhaustively investigated this option[10], and concluded that it was not suitable for our application for two main reasons. Firstly, the method was unable to measure distances with the required accuracy to overcome the one-λ phase ambiguity, and secondly, the requirements on the sensor and the subsequent data processing would have been unrealistic.

As with heterodyne interferometry, the interference pattern of a FMCW interferometer is dynamic, but this time the OPD information is encoded on the instantaneous *frequency* of the oscillating interference pattern. A potential sensor would therefore have to be able to measure both the phase and frequency of an amplitude modulated signal at each pixel. This could probably only be achieved by means of a fast-sampling type camera as explained in chapter 5.3, and the data-processing would be considerably more intensive.

A method which can overcome the phase ambiguity without changing the requirements of the sensor is multiple-wavelength interferometry.

4.3.1 Multiple wavelength interferometry

Multiple-wavelength interferometry is well established[6, 11], and essentially consists of performing a standard interferometric phase measurement with two or more different wavelengths. With an appropriate choice of wavelengths, the obtained phase measurements can be used to determine the OPD unambiguously, or at least with a much larger unambiguity range.

Let us assume that a phase measurement has already been performed at our primary metrology wavelength, λ_1, yielding a phase, ϕ_1. By introducing a second laser with wavelength λ_2, we can perform another phase measurement. The difference between the phases measured with λ_1 and λ_2 give a difference phase, $\Phi_{1,2}$:

$$\Phi_{1,2} = \phi_1 - \phi_2 , \tag{4.8}$$

which repeats every $\Lambda_{1,2}$. This $\Lambda_{1,2}$, referred to as the synthetic wavelength, is given by:

$$\Lambda_{1,2} = \frac{\lambda_1 \lambda_2}{\left| \lambda_1 - \lambda_2 \right|} \tag{4.9}$$

To extend the unambiguity range to 3mm and beyond, we will choose a synthetic wavelength of approximately 6mm, and hence $\lambda_2 \approx \lambda_1 \pm \lambda_1/10'000$. To allow an absolute OPD determination however, the accuracy of a measurement with this synthetic wavelength must reach an accuracy better than the unambiguity range of the previous system - i.e. 632.8nm for the case of only two wavelengths. This requirement would demand a synthetic phase ($\Phi_{1,2}$) measurement accuracy better than $2\pi/10'000$, placing a requirement on the frequency stability of the sources, as well as the detection system. For sources with similar wavelengths, the relationship between synthetic phase stability and wavelength stability, when measuring a distance, *d*, is:

$$\varepsilon\left(\Phi_{1,2}\right) = \frac{2\pi d}{\lambda} \sqrt{\left(\frac{\varepsilon\left(\lambda_1\right)}{\lambda_1}\right)^2 + \left(\frac{\varepsilon\left(\lambda_2\right)}{\lambda_2}\right)^2} \tag{4.10}$$

Where $\varepsilon(\Phi_{1,2})$ and $\varepsilon(\lambda)$ are the maximum allowed deviations from the mean for the synthetic phase and for the wavelength respectively. In our case we require a fractional wavelength stability of 10^{-8} of the source. While this requirement can

be guaranteed for the primary source (see appendix A.1), the stability of our λ_2 source is unlikely to meet this criterion. It is therefore necessary to introduce an additional wavelength, λ_3, yielding another synthetic wavelength together with λ_1: $\Lambda_{1,3}$.
To distribute the demands on the frequency stability of the two additional sources evenly, the synthetic wavelength $\Lambda_{1,3}$ should be approximately 60μm, giving $\lambda_3 \approx \lambda_1 \pm \lambda_1/100$. Now, to go from our 6mm synthetic wavelength to the 60μm synthetic wavelength, a mere phase measurement accuracy of $2\pi/100$ is required, translating to a fractional wavelength stability of 10^{-6}. The same accuracy is also required of the synthetic phase measurement $\Phi_{1,3}$ to bridge the gap between the 60μm and 632.99nm wavelengths. The details of the required wavelengths and their stabilities, are summarized in Table 4-1.

	Wavelength	Stability	Fractional stability	Frequency ν	Frequency stability
λ_1	632.99141nm	±0.00001nm	10^{-8}	474.393THz	5MHz
λ_2	λ_1±0.06330nm	±0.00063nm	10^{-6}	$\nu_1 \pm$ 47GHz	474MHz
λ_3	λ_1±6.32991nm	±0.00063nm	10^{-6}	$\nu_1 \pm$ 4.7THz	474MHz

Table 4-1 Wavelengths and stabilities required for multi-λ light source.

Given light sources meeting these specifications, we proceed as follows: The synthetic phase $\Phi_{1,2}$ is measured, allowing us to determine 6mm OPDs to an accuracy better than 60μm. This data will allow us to add the appropriate number of 2π phase jumps to the measurement of $\Phi_{1,3}$, giving us a measurement accuracy better than 600nm over the total 6mm range, which in turn can be used to add the correct number of 2π phase shifts to the ϕ_1 phase measurements. Provided the phase-measurement accuracy for our primary wavelength is $2\pi/1000$, we now have a resolution of 0.06nm over a range of 6mm. If we merely need to overcome the phase ambiguity between adjacent pixels, wavelengths λ_1 and λ_3 will be sufficient.

To generate these wavelengths, we have an actively stabilized HeNe laser from Melles Griot, a tuneable external cavity diode laser by EOSI (now New Focus), and a similar laser from Newport, as described in appendices A.1-A.3. While the HeNe laser already fulfils our requirements, the tuneable lasers first need to be locked and stabilized to the desired wavelengths.

4.4 Wavelength stabilization

The ultimate accuracy of our interferometer depends on the accuracy of the primary metrology wavelength used. This wavelength is provided by an actively stabilized HeNe laser by Melles Griot. Its cavity length is thermally adjusted to equalize the two Zeeman split modes of the cavity, of which only one is used as output. Since the two modes are only equal in amplitude when centred around the middle of the HeNe gain spectrum (see Figure 4-8), the laser is stabilized with respect to the centre of this spectrum. Measurements by the Dutch metrology institute (NMI) confirm that the laser is stable to better than 2MHz over 8 hours (see appendix A.1).

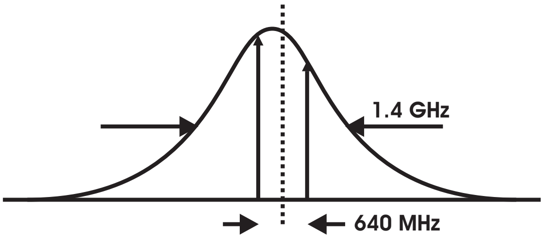

Figure 4-8 – HeNe gain profile and two Zeeman-split modes, illustrating the working principle of the Melles Griot HeNe locking scheme (not to scale).

The choice of this particular wavelength was a result of the availability and low cost of a sufficiently stabilized source, together with the large amount of optical components designed for this wavelength. While we were free to choose our primary metrology wavelength using these criteria, the wavelengths of our other two lasers are consequently fixed by the specifications in Table 4-1.

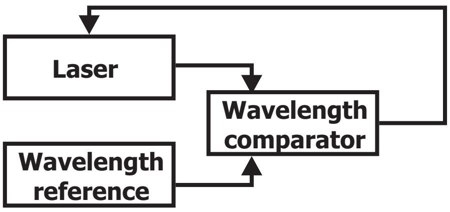

Figure 4-9 – Generic feedback locking scheme.

Tuneable lasers are able to lase at the unusual wavelengths required, as well as providing the possibility to dynamically adjust the wavelength by a generic feedback type locking scheme, as illustrated in Figure 4-9.

The two most common types of wavelength comparators for stabilization purposes either use a molecular gas absorption spectrum or the mirror spacing of a Fabry-Perot cavity as wavelength reference. While gas absorption locking has the advantage of offering an absolute wavelength reference, there may not be any absorption at our wavelengths of interest. With Fabry-Perot locking on the other hand, any wavelength can be locked to, but this wavelength may not remain stable for a sufficiently long time. Both options were evaluated, and the results will be discussed in the subsequent subsections.

4.4.1 Molecular gas absorption

The method by which the Melles Griot HeNe laser wavelength is locked can be considered as a type of gas absorption locking, since the HeNe gain spectrum is related to its absorption spectrum. The gain profile of the tuneable lasers is far too broad and temperature dependent to be useful as an absolute wavelength reference, nor are other lasers available with gain profiles which just so happen to be centred on our desired wavelengths. Instead we will have to rely on the absorption profile of another substance to serve as wavelength reference.

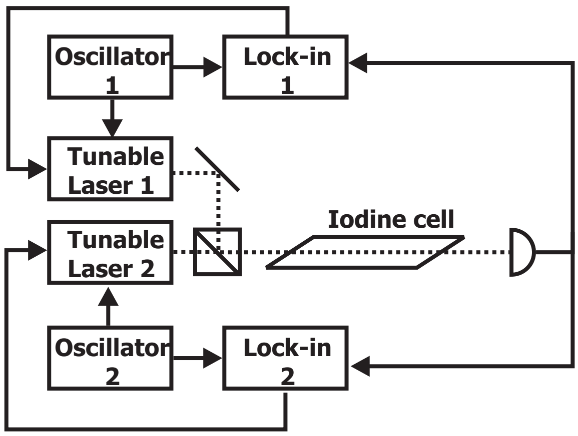

Figure 4-10 – gas absorption locking scheme for two lasers using an iodine gas cell.

Given such a substance, the laser would be locked to the appropriate absorption peak by the well-known[2, 12] scheme illustrated in Figure 4-10. By lightly modulating the wavelength of our laser, we produce a synchronous amplitude modulation signal, whose amplitude is roughly proportional to the gradient of the absorption profile at the current wavelength. When at the peak of the absorption profile, there will be no amplitude modulation at the frequency of the wavelength modulation. Instead, the amplitude modulation will occur at twice this frequency, and its harmonics, as can be seen in Figure 4-11 and the following treatment.

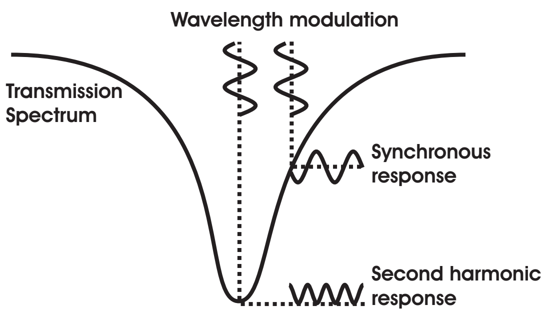

Figure 4-11 – Diagram illustrating the concept of modulation locking.

Given a transmission profile of a substance as a function of wavelength, $T(\lambda)$, let us examine the effect of a small oscillation in the wavelength being transmitted through it:

$$
\begin{aligned}
I_{trans} &= I_0 T\left(\lambda_0 + \varepsilon \sin\left(\omega t\right)\right) \\
&= I_0 \sum_{n=0}^{\infty} \frac{\varepsilon^n}{n!} \frac{\partial^n T}{\partial \lambda^n}\bigg|_{\lambda_0} \sin^n\left(\omega t\right) \\
&\approx I_0 \left[T\left(\lambda_0\right) + \left(\varepsilon \frac{\partial T}{\partial \lambda}\bigg|_{\lambda_0} \right) \sin\left(\omega t\right) - \left(\frac{\varepsilon^2}{4} \frac{\partial^2 T}{\partial \lambda^2}\bigg|_{\lambda_0} \right) \cos\left(2\omega t\right) + \frac{\varepsilon^2}{4} \frac{\partial^2 T}{\partial \lambda^2}\bigg|_{\lambda_0} \right]
\end{aligned}
\tag{4.11}
$$

To second order, the response exhibits a DC component, a component at the modulation frequency proportional to the gradient of $T(\lambda)$, and one at twice the modulation frequency, proportional to the second derivative of $T(\lambda)$.

We can therefore use the amplitude of the amplitude modulation component at the wavelength modulation frequency as a feedback signal. This is best achieved by using a lock-in amplifier which, besides being highly frequency selective, has the added advantage of amplifying our weakly modulated signal.

The gas selected as a potential candidate for exhibiting absorption peaks near our desired wavelengths was molecular iodine – a well characterized substance, commonly used as wavelength standard[13-18], with a rich absorption spectrum near our region of interest[19].

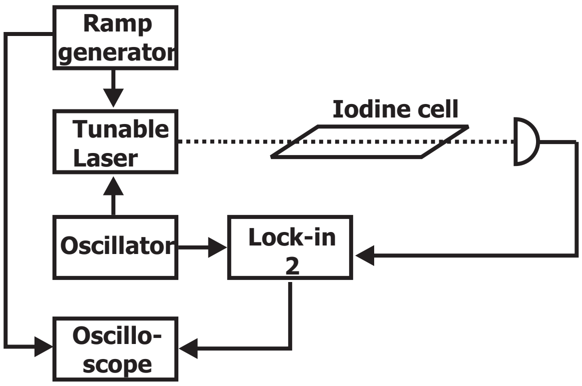

Figure 4-12 – Set up to confirm the presence of iodine absorption lines.

The presence of iodine's main absorption peak group (transition 11-5, R(127) - consisting of 14 peaks around 632.9913nm, and transition 6-3, P(33), consisting of 10 peaks around 632.9920nm, all merged into one peak by Doppler broadening), was confirmed by measuring the lock-in response of the set up shown in Figure 4-12 while slowly sweeping the wavelength of the tuneable laser past the expected location of the peak. The ramp generator acted on the high sensitivity input (piezo) of the tuneable laser, while the oscillator acted on the low sensitivity input (current). The results can be seen in Figure 4-13.

The width of the detected signal is primarily due to Doppler broadening, which is 367.6 MHz at room temperature. The irregularities are due to a combination of laser amplitude variations during tuning, and the fine structure of the peak group.

Unfortunately no reliable absorption peaks were found using this set up around our wavelengths of interest. Such absorption peaks should exist (transitions: 7-3, R(120); 7-3, P(114); 7-3, P(115); 8-4, R(60); 8-4, R(61); 8-4, P(54); 8-4, P(55), and transitions: 7-4, R(93); 7-4, R(94); 7-4, P(87); 7-4, P(88)), but would be significantly weaker than the main absorption peak group around 632.9913nm. While the sensitivity of the measurement could be improved by several orders of magnitude by using an intra-cavity configuration for the Iodine cell, this was not considered practical, requiring a re-design of the commercial tuneable lasers. The alternative option of Fabry-Perot locking was investigated instead.

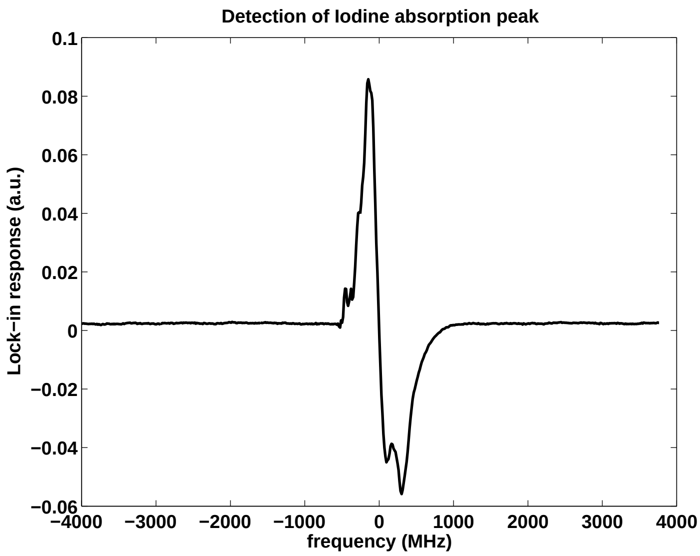

Figure 4-13 – Lock-in response from set up in Figure 4-12, demonstrating the detection of Iodine absorption lines. Frequency given as de-tuning from primary Iodine absorption peak

4.4.2 Fabry Perot locking

A Fabry-Perot etalon essentially consists of a two-mirror cavity, the transmissivity of which is a very sensitive function of the frequency of incoming light and the mirror separation.

Peak transmission occurs periodically with a period called the “Free Spectral Range” (FSR) (see Figure 4-14), dependent on the mirror separation, d, according to:

$$FSR = \frac{c}{2d} \tag{4.12}$$

The transmission peaks therefore occur at the following frequencies:

$$f_m = m\frac{c}{2d}, \qquad m = 1,2,3,.. \tag{4.13}$$

The full width at half maximum (FWHM) of these transmission peaks, is usually expressed as a fraction of the FSR and the inverse of this ratio is called the

"Finesse" ($\mathscr{F}$) of the cavity. In a loss less cavity with two identical mirrors, the finesse depends only on the reflectivity of these mirrors according to:

$$F = \frac{\pi\sqrt{R}}{1-R} \tag{4.14}$$

Because of this sharply peaked transmission function, a Fabry Perot cavity can serve as a frequency reference. It is important that the mirror spacing remains constant if such a cavity is to be used for this purpose.

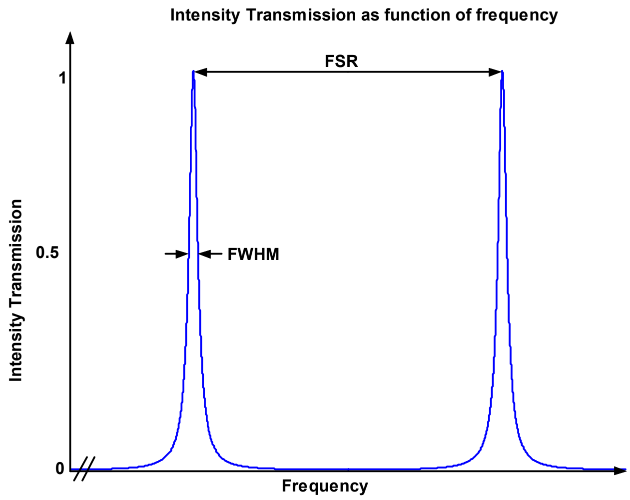

Figure 4-14 – Transmission spectrum of a Fabry Perot Etalon.

For our requirements, a cavity with a FSR of 47GHz would be suitable, giving an approximate cavity-spacing of 3.19mm according to (4.12). We could then tune the cavity spacing such that there is a transmission peak at our stabilized HeNe wavelength (m≈10'000 in (4.13)), which would place our 2nd desired wavelength on an adjacent transmission peak, and our 3rd wavelength on the 100th transmission peak from the HeNe peak.

We have purchased such a cavity – a refurbished Burleigh Fabry-Perot (see appendix A.8), and fitted it with a pair of parallel mirrors. The transmission of our tuneable laser, as a function of frequency offset is shown in Figure 4-15. While the mirror spacing and hence also the FSR, are continuously adjustable, the finesse of the cavity is limited by the mirror reflectivity. Our cavity exhibits a finesse better than 60, equivalent to a FWHM of 780 MHz. Locking to within a small fraction of this FWHM is routinely achieved in the laboratory[12], and this cavity can be expected to fulfil our requirements.

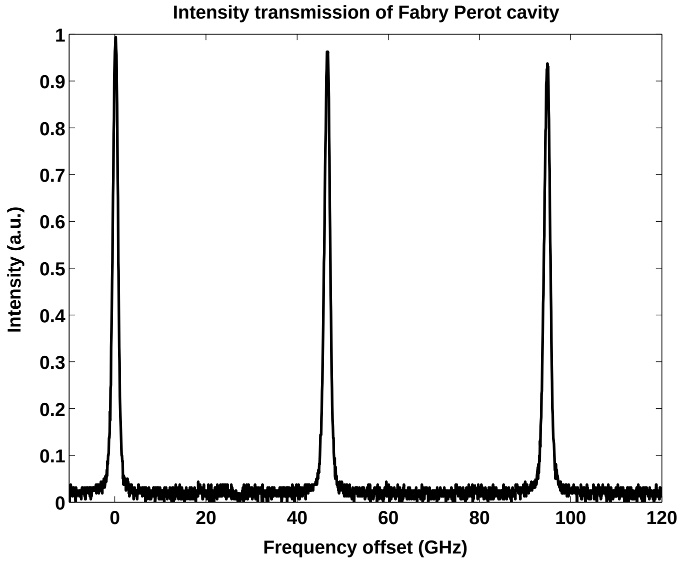

Figure 4-15 – Measured transmission characteristics of Burleigh Fabry Perot cavity.

Locking could proceed as illustrated schematically in Figure 4-16. The cavity spacing would be weakly modulated, causing the transmitted intensity of the laser to be modulated as well. Similar to gas absorption locking (equation (4.11)), the amplitude of the resulting intensity modulation at the frequency of the cavity modulation would be roughly equal to the derivative of the transmission spectrum (see Figure 4-17), and could therefore serve as an error signal to be fed back to the laser to stabilize it.

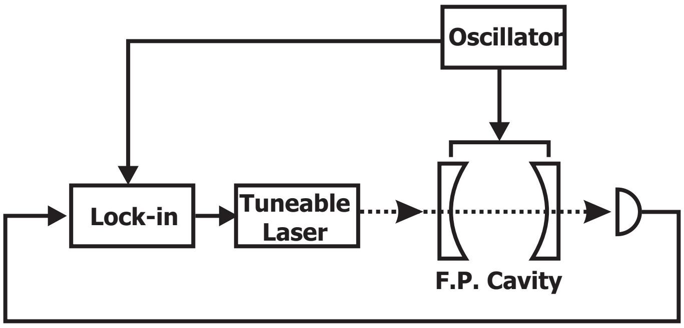

Figure 4-16 – Scheme to lock a tuneable laser to a F.P. cavity. The flat cavity mirrors are shown curved only for aesthetic reasons.

Near the locking region, the error-signal (V_e) is approximately proportional to the detuning of the laser's actual optical frequency (ν_a) from the desired lock frequency (ν_0). The laser's actual frequency, in turn, is made to differ from its 'natural' frequency (ν_n- without a feedback signal) by a constant times the feedback signal:

$$V_e = A(\nu_a - \nu_0) \quad (4.15)$$

$$\nu_a = \nu_n + BV_e \quad (4.16)$$

Figure 4-17 shows such an error signal measured with our cavity.

By appropriately amplifying or attenuating this signal (varying '*A*'), we can ensure an arbitrarily tight lock, limited only by the level of noise in the signal. In our case, we are limited to an accuracy of ~2 MHz – well within our requirements:

$$\begin{aligned} \nu_a &= \nu_0 + \frac{\nu_0 - \nu_n}{BA - 1} \\ &\approx \nu_0 + \frac{\nu_0 - \nu_n}{BA} \end{aligned} \tag{4.17}$$

Unfortunately, the transmissivity of a Fabry-Perot cavity is not only highly sensitive to the wavelength of the incoming radiation, but also to the mirror spacing. Differentiating (4.13) with respect to the mirror spacing yields:

$$\frac{\partial f}{\partial d} = -m\frac{c}{2d^2} \tag{4.18}$$

So, the mirror spacing would need to kept stable to better than 3nm to fulfil the 475MHz frequency stability requirement in Table 4-1.

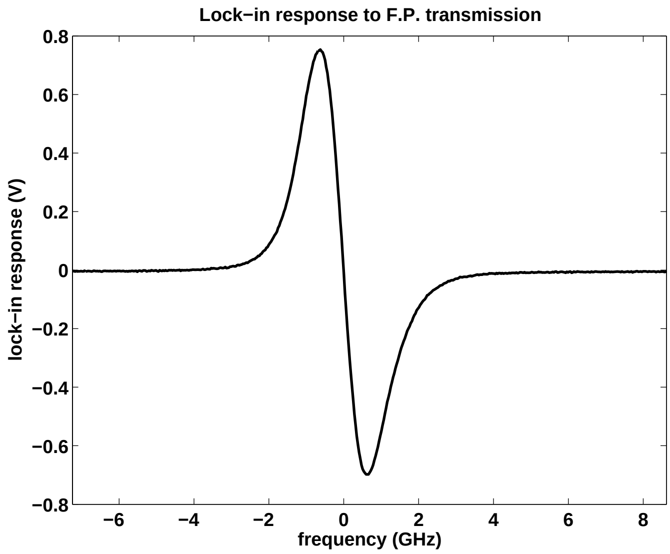

Figure 4-17 – Error signal measured from Fabry-Perot locking set-up.

One way to overcome this problem is to actively stabilize the cavity to our reference standard – the HeNe laser. Instead of feeding the error signal arising from a transmitted HeNe beam back to the laser, it would be fed back to control the mirror spacing of the Fabry-Perot, as shown in Figure 4-18.

The challenge here is separation of the different signals passing through the cavity. It is essential that the control signal for the cavity is only due to the HeNe, while the tuneable lasers should only be controlled by their respective transmission signals. The following sections show a method for achieving this for three and two wavelengths respectively.

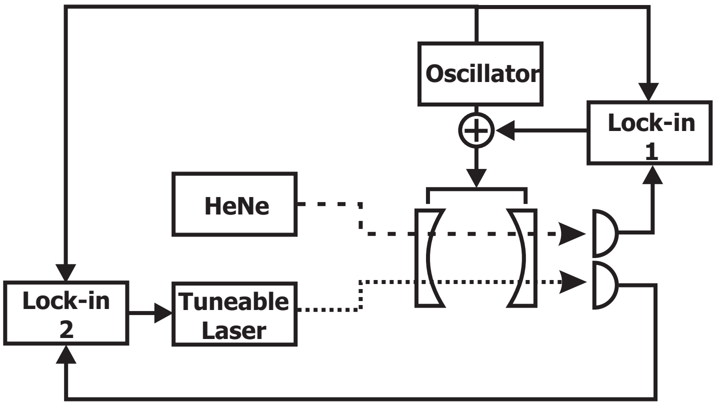

Figure 4-18 – Simplified scheme to lock a tuneable laser to a reference laser, via a Fabry Perot cavity. The flat cavity mirrors shown as curved for aesthetic reasons only.

4.4.3 Three-wavelength locking scheme

As discussed in section 4.1.1, it is not feasible to separate the three wavelengths using spectral means. However, given the light produced by the set-up shown in Figure 4-1, each wavelength can uniquely be identified by its intensity modulation frequency. Depending on the AOM settings, the beams of the three lasers might be modulated at 15, 20 and 25kHz, for example. By sending this bundle of three beams through the set-up shown in Figure 4-19 each beam is attenuated according to the instantaneous transmission characteristics of the Fabry-Perot cavity at their respective wavelengths.

Since the cavity modulation is limited to about 100Hz, the process of cavity modulation can be seen as quasi-static with respect to the intensity modulation of the laser beams themselves. We can therefore first separate the beams out by de-modulation at their respective frequencies, and then use the resulting signals to generate three feedback signals using de-modulation at the cavity modulation frequency.

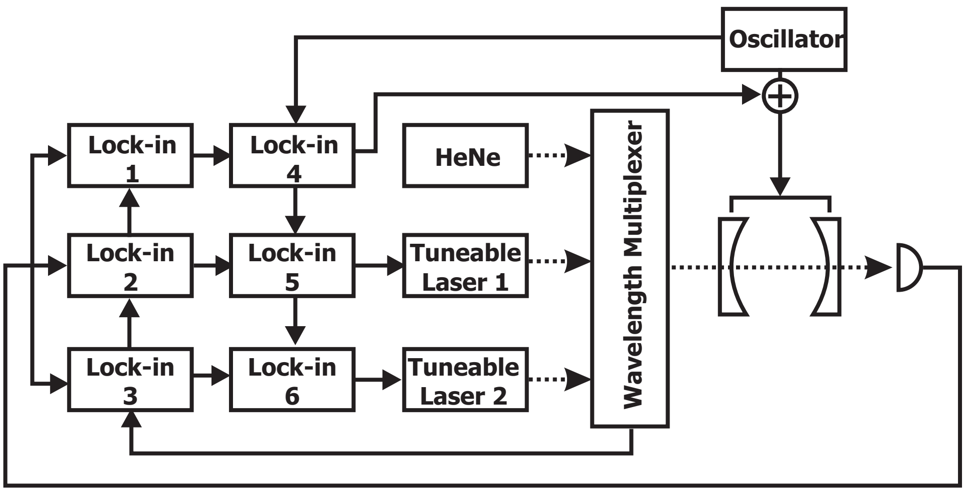

Figure 4-19 – Full 3λ locking scheme, using shift-frequency multiplexing (c.f. Figure 4-1 & Figure 4-3).

Due to the limited number of lock-in amplifiers available at the time, only a limited 2λ demonstrator version of such a set-up was constructed, according to Figure 4-20. The switches allow us to either lock the cavity to the HeNe, the tuneable laser to the cavity, or the cavity to the tuneable laser.

The presence of suitable error signals was confirmed by measuring at “A” while leaving the loop at “A” open, and sweeping the cavity past both the HeNe and Tuneable laser transmission peaks. The peaks were intentionally separated to demonstrate the absence of cross-talk between the signals (Figure 4-20).

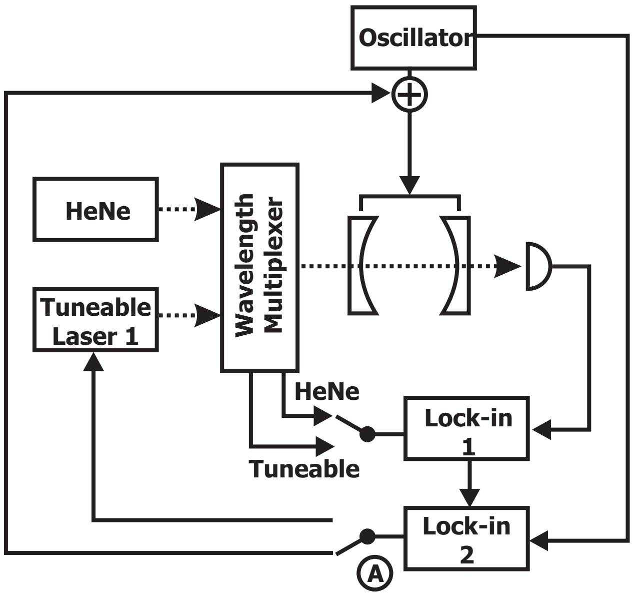

Figure 4-20 – Two-λ demonstrator scheme for shift-frequency multiplexed wavelength locking

The lock was confirmed by first manually tuning the cavity close to resonance before monitoring the signal at “A” and closing the loop. Figure 4-22 shows the expected drop in error-signal as soon as the loop is closed.

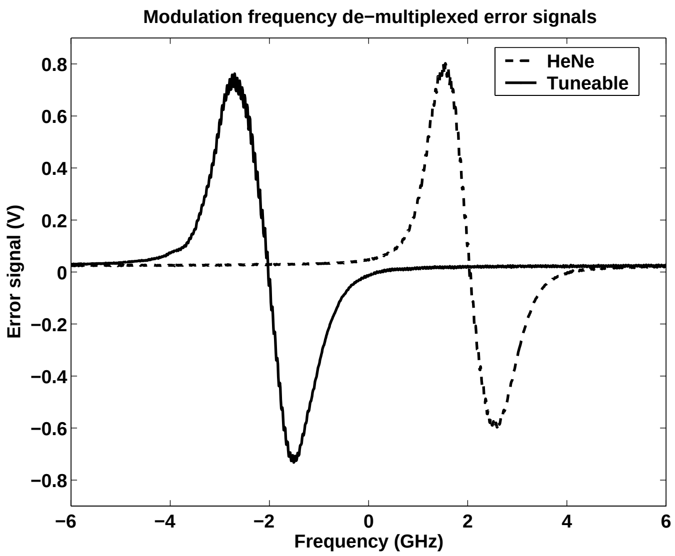

Figure 4-21 – Error signals from open feedback loop at location “A” of Figure 4-20

It should be noted that the feedback signal remains around 0.05V - 1/40 of the p-p amplitude of the error signal in Figure 4-21. From Figure 4-21, we can determine the value of *A* in (4.15) to be ~1.6 GHz/V, and hence deduce that the lock is maintained within 80MHz of our desired frequency – well within the requirements in Table 4-1.

To test the lock's immunity to cross-talk, the cavity was locked to the HeNe laser, while the tuneable laser was slowly swept past its cavity resonance. The tuneable laser was then locked to the cavity while the cavity separation was slowly swept past the HeNe resonance. No loss of lock occurred in either case.

This demonstrator set-up shows the feasibility of using this locking method, once the availability of appropriate sensors warrants its use (see the beginning of this chapter). In the meantime, it is sufficient to lock only one extra laser at a time, since the available sensors are only capable of measuring the different wavelengths sequentially. In fact, only two lasers are required to measure our current mirror, since its asphericity is sufficiently small to introduce an OPD range well below 3mm.

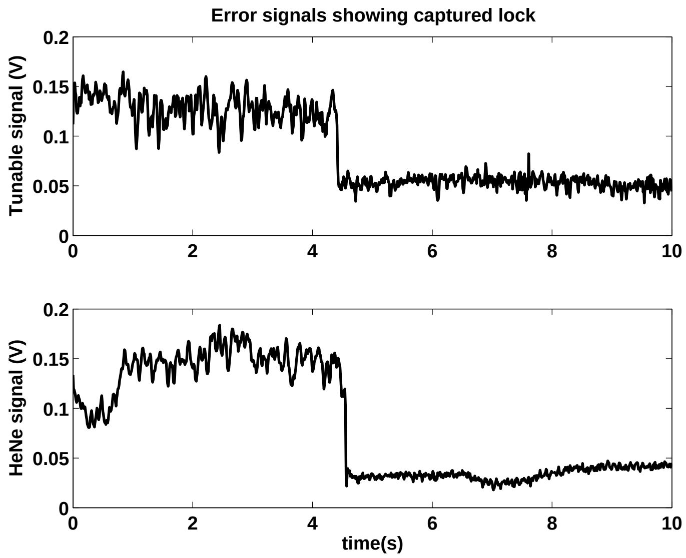

Figure 4-22 – Feedback loop error signals observed before and after active lock.

4.4.4 Two-wavelength locking scheme

When only two wavelengths are involved (one being the stabilized HeNe wavelength here), instead of having to resort to a complicated de-modulation scheme of the type in section 4.4.3, the two can be separated by making use of polarization effects. A beam, as provided by the set-up in Figure 4-4 or Figure 4-5, can be sent through our Fabry-Perot cavity and the two wavelengths can then be separated out again using a polarizing beam splitter, as shown in Figure 4-23.

Such a set-up was constructed using our Fabry-Perot, HeNe laser and EOSI tuneable laser. The Fabry-Perot was modulated with 100Hz, with an amplitude of approximately 0.4nm (equivalent to 60 MHz of optical frequency). The resulting signals at the photodiodes were first high-pass filtered at 75Hz and fed to two Femto lock-in modules (See appendix A.6 for specifications), which used the 100Hz driving signal as reference signal. The 75Hz high-pass filter is necessary to prevent a DC overload of the Lock-in modules during an active lock. After separately attenuating the de-modulated error signals from the HeNe and tuneable wavelengths, they were fed to the cavity control unit and tuneable laser piezo-input respectively. The existence of a lock was confirmed by observing the output from the two photodiodes. During lock, a 200Hz signal with a large DC offset is seen, corresponding to modulation around the peak transmission of the Fabry-Perot cavity.

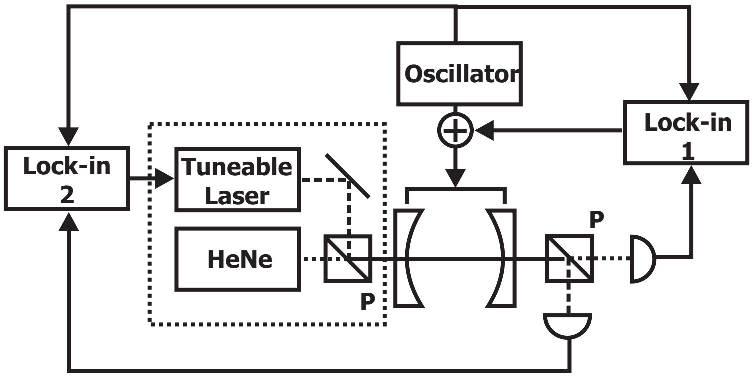

Figure 4-23 – Two-λ locking scheme using polarization multiplexing. The dotted box is equivalent to the light source before the "locking" output of Figure 4-4 or Figure 4-5.

To quantify the quality of the lock, we tuned the tuneable laser to the HeNe wavelength and measured a beat-frequency as a function of time (using a fast NewFocus photodiode), with- and without an active lock. The results can be seen in Figure 4-24. It is clear that even though the tuneable laser already has very good intrinsic frequency stability, the active lock improves this by a factor of 20, placing it well within the specifications in Table 4-1.

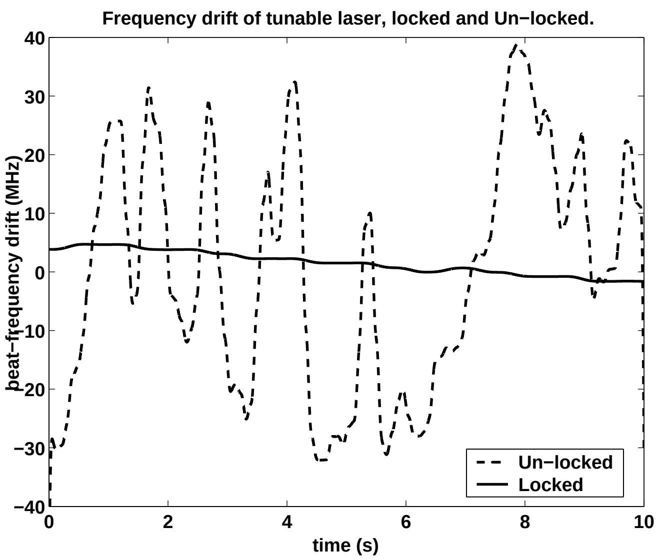

Figure 4-24 – Beat-frequency of HeNe and tuneable laser, as function of time, with- and without active locking

We now have a light source with the required wavelength properties to measure our mirror. Besides the issue of wavelength stability, we also have to consider the effect of OPD drifts within the light source, as excessive drifts can have a detrimental effect on our phase measurements (See chapter 5.2.2). In order to minimize these drifts, we have to take mechanical, thermal and atmospheric stability into consideration. These are the topics covered in the following sections.

4.5 OPD drifts

Any changes in the optical path lengths inside the interferometer must be regarded as detrimental, as they will give rise to phase drift between the reference- and object-beams.

The three major contributors to unwanted OPD changes inside the light source are mechanical vibrations, thermal expansion, and atmospheric turbulence. We have tried to reduce each of these contributions by adapting the design of the light-source and its components. We will now deal with these in turn.

4.5.1 Mechanical considerations

The first steps towards reducing the amount of mechanical vibrations is the choice of a Newport honeycombed optical table together with four vibration isolators (See appendix A.7 for specifications). This system is used to reduce the effect of building vibrations on the relative positions of optical components on the table. Short of complex suspension systems or expensive active damping solutions, these components offer the best means of isolating our optical table from external vibrations.

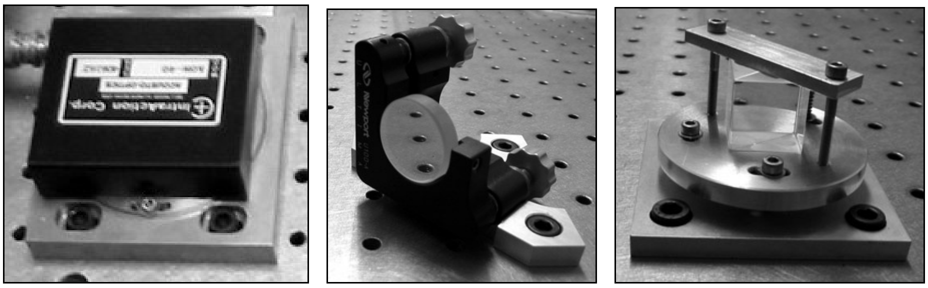

Figure 4-25 – Custom designed optical mounts for AOMs, mirror-holders and beam-splitters

Even so, the optical components fixed on the table are able to pick up acoustic noise present in the room. While sources of such noise will be reduced to a minimum during measurements by switching off unnecessary equipment, and placing as much equipment as possible in a sound-damped cabinet, a certain amount of exposure to sources of vibration will remain. These include the cooling fan of the tuneable laser driver, the presence of the experimenter, any vibrations that pass through the damping system of the table, and the mechanical vibrations of the Fabry-Perot cavity during the wavelength lock.

To reduce the effect of these vibration sources, we have designed mounts for our optical components which should outperform commercially available mounts. The main method for improvement envisaged was a reduction in the beam-height, from 100 mm to just 27mm. Furthermore, the number of parts per component was kept to a minimum, while maintaining the required degrees of freedom to adjust the optics in the set-up. Once adjusted, the components can then be fixed in position by the application of counter forces. The resulting designs for beam cube-, AOM- and mirror-mounts are shown in Figure 4-25 (for technical drawings, see appendix A.5).

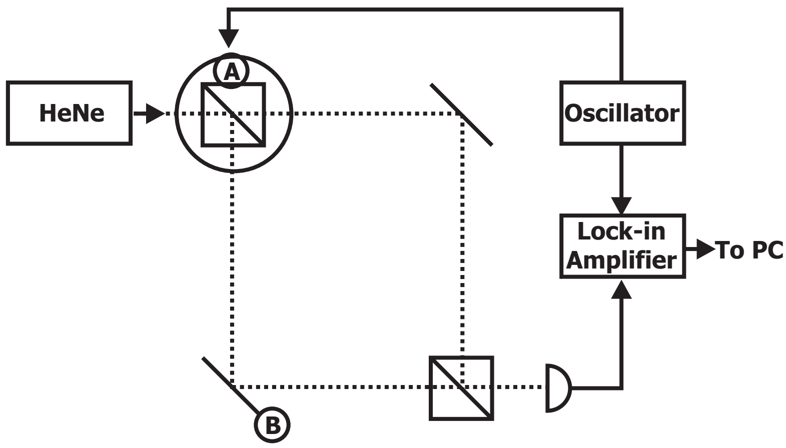

Figure 4-26 – Test-setup to measure the immunity to vibration of optical mounts. Positions ‘A’ and ‘B’ indicate the placement of the piezo-mounted mass to test the beam-splitter mount and mirror-mount respectively. The same set-up was constructed twice, once for each set of components

The stability of the beam cube mounts and mirror mounts was quantitatively compared with that of their Newport equivalents by synchronous detection of OPD oscillations under active excitation. The set-up used was a Mach-Zehnder configuration (as illustrated in Figure 4-26), tuned to the middle of a fringe. A piezo-mounted mass of 100mg was attached to the mounts near the optical components, and the piezo was driven by a fixed amplitude signal at several different frequencies. The amplitude of the resulting OPD variations at the driving frequency was measured with a lock-in detector (using the amplitude of a full-fringe swing as reference), giving the component’s response for that frequency. The range of frequencies tested was from 50Hz to 1 kHz, corresponding to the lower acoustic range of the vibrational spectrum.

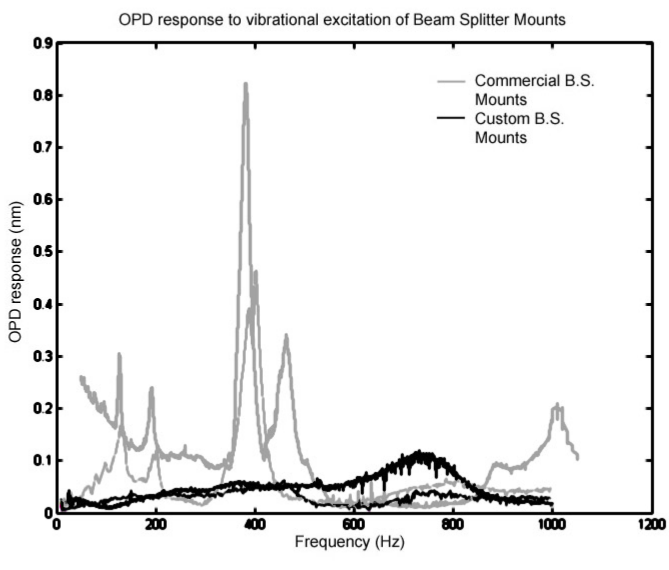

Figure 4-27 – Comparison of vibrational sensitivity between commercial and custom designed beam-splitter mounts

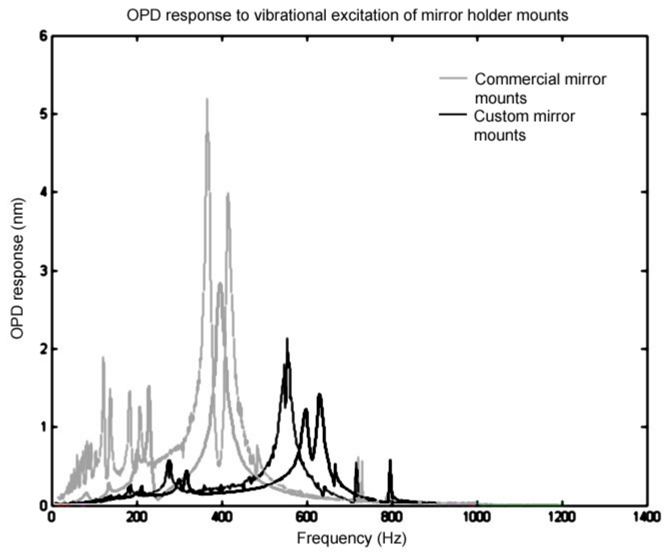

Figure 4-28 – Comparison of vibrational sensitivity between commercial and custom designed mirror mounts

A comparison of the response from our custom designed components with that from the commercial mounts (Figure 4-27 & Figure 4-28) shows a clear improvement.

Despite this reduction in vibration sensitivity, the set-up still exhibits low frequency phase drifts. This is most likely due to a combination of temperature variations and air turbulence. With an appropriate combination of materials, the temperature sensitivity of the interferometer can be reduced.

4.5.2 Thermal considerations

The five most important materials used in building the interferometer are: steel (optical table), invar (interferometer frame – see Chapter 6), silica (fibers), BK7 (beam-splitters) and aluminium (custom mounts). Each of these materials has a different thermal expansion coefficient, and the thermal dependence of the reference- and object-beam OPLs is defined by different combinations of these materials. The optical path *difference* between the reference and object beam must be kept below 3mm for our interferometer, but we have some liberty as to the total optical path *length*. By choosing to minimize the total OPL, we reduce the influence of several phase disturbing effects, such as turbulence, temperature drift, angular drifts, etc. With a non-minimal OPL however, it is theoretically possible to completely eliminate the effect of uniform temperature drifts, by equalizing the combined, effective thermal expansion coefficient of the reference- and object-arms.

Although most of the customized components were manufactured from aluminium, aluminium does not contribute to the effective thermal expansion coefficient of the two arms. This is because all aluminium components are tightly screwed onto the stainless steel faceplate of the optical table. The fact that aluminium and stainless steel have significantly different expansion coefficients still gives rise to problematic temperature effects, such as warping, bending and creeping of the aluminium components secured to the optical table, but the contribution of these effects to the OPD cannot be estimated with any confidence. It would be advisable to use only one type of material in the construction of commercial versions of this interferometer, to eliminate such warping problems. Referring to Figure 4-29 we can see why the flint-glass inside the AOMs is also neglected – each beam passes through the same amount of this glass, and it therefore does not contribute to the OPD.

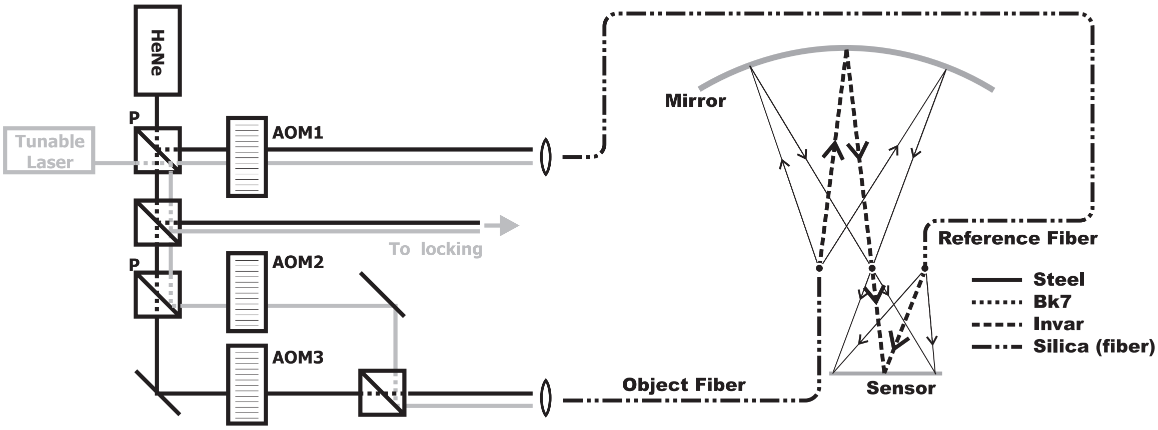

Figure 4-29 – Setup for thermal drift calculations. Materials influencing the optical path lengths differently are indicated.

The contribution of the invar to the OPD is fixed, as the path-length inside the interferometer frame is fixed. The path-length through BK7 will also be considered to be fixed, since the arrangement of beam-splitters can not be changed, but we would theoretically be at liberty to add a “dummy” section of BK7 glass to one of the arms, if this should somehow prove to be advantageous. The only two variables we can now play with are the OPD through the silica-core fibers, and the OPD defined by steel. We can find the optimum combination of these two OPDs by requiring that the total OPD of the chief-ray through the interferometer be zero, and that the OPD change with temperature should be zero.

The dependence of the OPL through a particular material on temperature can be approximated by:

$$\begin{aligned} OPL &= n \times l \\ &\approx n_0(1+\xi\Delta T) \times l_0(1+\alpha\Delta T) \\ &\approx n_0 l_0\left(1+\left[\alpha+\xi\right]\Delta T\right) \end{aligned} \tag{4.19}$$

$$\Rightarrow \partial OPL / \partial T \approx n_0 l_0 \left[\alpha + \xi\right] \tag{4.20}$$

Were l_0 and n_0 are the lengths and refractive indices at the nominal temperature respectively, α is the linear expansion coefficient, and ξ is the fractional change of refractive index with temperature, $(1/n)(dn/dT)$. The dependence of the latter two coefficients on temperature is neglected. For the paths through air (the length of which is defined by stainless steel or invar), n is assumed here to be unity, and ξ is assumed to be negligible.

material/property	α ($\times 10^{-6}K^{-1}$)	ξ ($\times 10^{-6}K^{-1}$)	n
Stainless steel	17.6	-	-
Aluminium	23.1	-	-
Invar	1.7	-	-
Silica fibers	0.55	8.24	1.4571
BK7	7.1	1.78	1.5151

Table 4-2 Selected material properties at 20°C and for λ=632.8nm.

The physical path length difference between object and reference beams for a particular material is trivially defined by:

$$\Delta l_{material} = l_{material_{object}} - l_{material_{reference}} , \tag{4.21}$$

The instantaneous optical path difference (OPD) between the object and reference paths is simply given by:

$$OPD = \Delta l_{invar} + \Delta l_{steel} + n_{silica}\Delta l_{silica} + n_{BK7}\Delta l_{BK7} \tag{4.22}$$

The change in OPD with temperature can now be found using:

$$\partial OPD / \partial T = \sum_{materials} \left[\alpha_{material} + \xi_{material}\right] n_{material} \Delta l_{material} \ . \tag{4.23}$$

	l_{obj} (mm)	l_{ref} (mm)	OPL (mm)	OPD (mm)	$\partial OPD / \partial T$ (mm/K)
Flint	60.0	60.0	180.0	0.0	0.0
BK7	75.4	10.0	129.4	99.1	871.1x10^{-6}
Air/Steel	643.4	82.4	725.8	561.0	9873.1 x10^{-6}
Air/Invar	680.0	0.0	680.0	680.0	1155.9 x10^{-6}
Silica	250.0	1169.7	1419.7	-1340.1	-11900.1 x10^{-6}
Total			**3134.9**	**0**	**0**

Table 4-3 – Path lengths of thermally compensated set-up in Figure 4-29.

These last two equations can now be set to zero, and we can solve for Δl_{steel} and Δl_{silica} using the material constants in Table 4-2, Δl_{invar} = 0.680m – twice the test-mirror's ROC – and Δl_{BK7} = -0.0654m – the total thickness of the last three beam-splitters. From these calculations, we get: Δl_{steel} = 0.561m and Δl_{silica} = -0.920m. In other words, the object path through air, defined by stainless steel, should be 0.561m longer than the equivalent reference path, and the object fiber should be 0.920m shorter than the reference fiber, to get zero net thermal expansion difference between the object and reference arms.

This stainless-steel path-length difference can be achieved naturally, by choosing the longer arm of the light-source as the object arm (an assumption we have already made, by including the three prism-lengths in the object-, rather than the reference-arm in Figure 4-29). However, we need at least 25cm of object fiber to connect the interferometer frame with our light source, so the reference fiber would have to be at least 1.021m long.

Many things could go wrong over such a length – especially since we have so far neglected the fact that the temperature distribution over the various components will not be uniform. If we wish to make use of the “zero net thermal expansion” configuration, we must ensure that the temperature changes occur as uniformly throughout the interferometer as possible. Future designs of this interferometer should therefore place the fibers, optical elements and interferometer frame in good thermal contact with the rest of the set-up.

	l_{obj} *(mm)*	l_{ref} *(mm)*	OPL *(mm)*	OPD *(mm)*	$\partial OPD / \partial T$ *(mm/K)*
Flint	60.0	60.0	180.0	0.0	0.0
BK7	10.0	75.4	129.4	-99.1	-871.1×10^{-6}
Air/Steel	82.4	347.2	429.6	-264.7	-4658.7×10^{-6}
Air/Invar	680.0	0.0	680.0	680.0	1155.9×10^{-6}
Silica	250.0	467.0	1044.7	-316.2	-2807.9×10^{-6}
Total			**2463.7**	**0**	$\mathbf{7181.7 \times 10^{-6}}$

Table 4-4 – Path lengths of shortest-path set-up in Figure 4-30.

Because good thermal contact was not made one of the design criteria for this prototype, the various components are likely to experience very different, and unpredictable temperature changes. For such a case, it is wiser to choose the shortest possible total optical path length (while maintaining a OPD of zero, of course), since the temperature fluctuations induce a fractional change in the optical path length.

Such a configuration which minimizes the total optical path length while maintaining a zero OPD is shown in Figure 4-30, with the corresponding distances given in Table 4-4. The object- and reference-arms are swapped, with the fiber-couplers as close as practically possible to the output-ports of the light source. The length difference between the object and reference fibers needed to set the OPD to zero is -0.217m, with the shorter length of fiber again chosen as being 25cm long for practical reasons.

This is the configuration chosen for our measurements. In future versions of the interferometer, where a more uniform distribution of temperatures can be guaranteed, it may be advisable to use the configuration shown in Figure 4-29 instead.

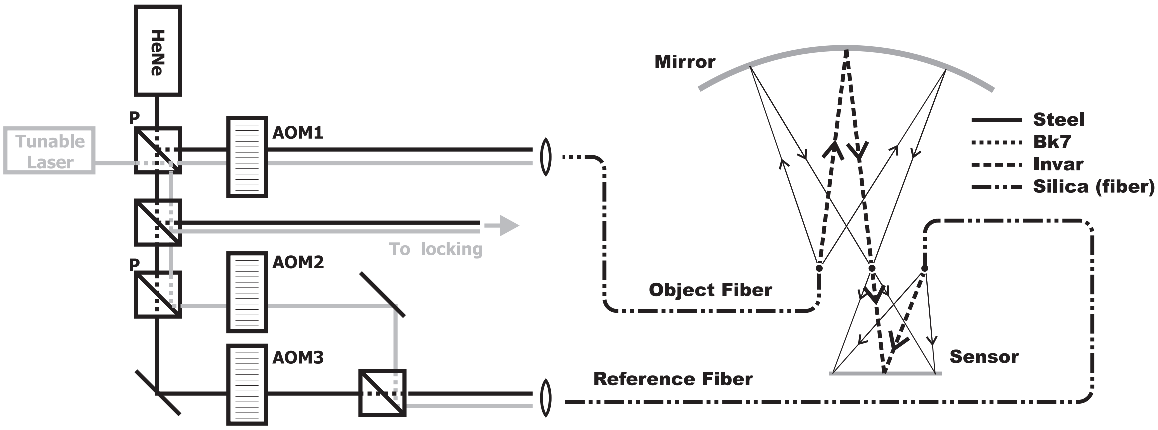

Figure 4-30 – Alternative arrangement of the interferometer, minimizing the total OPL.

4.5.3 Atmospheric considerations.

The last remaining factor contributing to OPD drift in the light source is the inhomogeneous nature of the atmosphere. Other authors have looked at OPD errors in high-accuracy interferometry due to atmospheric disturbances in a laboratory context[4, 20, 21], but it may be useful to re-iterate some of their conclusions, and add some comments to their findings.

Although the refractive index of air differs from unity by only ~256ppm, its dependence on a number of environmental parameters can introduce short- to medium-term fractional OPD drifts of 3.4×10^{-8} – equivalent to roughly 40nm, for the in-air path lengths of the set-up in Figure 4-29, where the OPD in air is particularly long. Borboff[4] traces such drifts to temperature inhomogenities caused by convection, and shows that placing the set-up in an enclosure can reduce these drifts to 1.1×10^{-8}.

Slower, long-term (~10hrs) fractional OPD drifts of up to 1.6×10^{-6} (equivalent to 2μm for our set-up) remain, and have been shown by Estler[20] to be strongly correlated to pressure changes.

The dependence of the refractive index of air on temperature and pressure variations is easily demonstrated by starting with the Claussius Mossotti equation[22]:

$$n = \sqrt{1 + \frac{3}{\frac{M}{\rho R_m} - 1}} \approx 1 + \left(\frac{3}{2}\frac{R_m}{M}\right)\rho \qquad (4.24)$$

Where M is the molecular mass of the compound, R_m is its molar refractive index, and ρ its density. The approximation holds for most gases at optical frequencies where n~1.

Using the ideal gas law as a simplified model, we can relate the density of the gas to its temperature and pressure:

$$\rho = \frac{P}{RT} \tag{4.25}$$

R being the Gas constant (8.314$JK^{-1}mol^{-1}$), P the pressure and T the temperature. Consequently, (4.24) can be re-written:

$$n \approx 1 + \left(\frac{3}{2} \frac{R_m}{MR} \right) \frac{P}{T} \tag{4.26}$$

The resulting sensitivity of the refractive index of air to temperature and pressure is shown in Table 4-5, along with its dependence on humidity and CO_2 concentration as reported by Estler[20]. Wile Estler dismisses the influence of CO_2 concentration variations, the presence of an experimenter might render this assumption invalid.

Parameter	Nominal Value	Realistic changes (/8hrs)	Fractional OPD sensitivity	Realistic OPD drifts (our set-up)
Pressure	101.3 kPa	0.6 kPa	2.68×10^{-6} kPa^{-1}	2000nm
Temperature	20.0 °C	0.4 °K	-1×10^{-6} K^{-1}	40nm
Humidity	40 %	2 %	-0.01×10^{-6} $\%^{-1}$	20nm
CO_2	0.03 %	5.6 %	149×10^{-6} $\%^{-1}$	1000nm

Table 4-5 – Sensitivity of refractive index of air to environmental parameters

An average adult exhales approximately 6L of air every minute, the CO_2 concentration of which is ~187 times larger than the usual 0.03% of the ambient atmosphere[23]. Referring to Table 4-5, this change could therefore cause local fractional OPD drifts of up to 8.4×10^{-6} – equivalent to ~10μm for our set-up.

Humidity drifts appear to have the least influence on the refractive index of air, considering that a realistic humidity drift of 2% only results in a fractional OPD drift of 2×10^{-8}.

Regardless of the actual magnitude of the OPD drifts due to atmospheric conditions, they can be reduced by a factor of around 7.3 by replacing the air in the set-up with Helium, as reported by Dörband and Seitz[24], where the use of Helium flooding increased the attainable repeatability of measurements from 0.058nm to 0.015nm. Since Helium has a refractive index 7.3 times closer to unity than air, the temperature and pressure effects outlined above will also be reduced by this factor. The fact that the Helium would be provided from a standard source, would mean that the issues of humidity and CO_2 concentration also fall away.

In order to replace the air in the light source with helium, the surface of the optical table beneath the light source was first coated in adhesive plastic foil, and a chamber was built around it (see Figure 4-31). The chamber was then sealed with vacuum sealant to minimize leaks. An exit valve can be opened to flush the chamber with Helium. Using an oxygen meter, we can determine when the gas exiting the valve has a sufficiently low oxygen concentration. At this point, the valve will be closed with a balloon, serving as an indicator of a slight over-pressure. This over-pressure is necessary to maintain the helium atmosphere inside the container, even in the presence of small leaks.

Even without the addition of helium to the chamber, the influence of atmospheric turbulence is reduced, since it is isolated form the effects of turbulence originating outside the chamber.

4.6 Conclusion

In designing this light source, three main requirements were taken into consideration: versatility, wavelength stability and optical path length stability. Its versatility is evident in that it can be used for sequential as well as simultaneous multiple wavelength heterodyne interferometry, phase shifting interferometry as well as superheterodyne interferometry. Furthermore, it can be easily scaled down for less severe requirements.

The required wavelength stability can be met for all three wavelengths thanks to a novel locking scheme. To the best of our knowledge, this is the first time such a locking scheme has been implemented. A simpler version of the locking scheme is perfectly adequate for the measurements performed in this thesis.

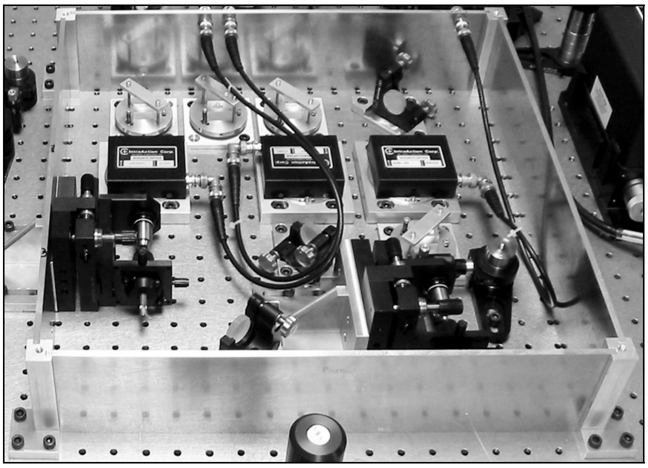

Figure 4-31 – Light source equivalent to the schematic set-up in Figure 4-4.

Several avenues for reducing the optical path length drift have been investigated and implemented. In doing so, a novel set of optical mounts has been designed. The use of helium gas to replace air in the set-up is also beneficial for the interferometer stage discussed in chapter 6. In future, the optical path length drifts can further be reduced by choosing an appropriate combination of materials for the construction of the light source.

4.7 References

[1] M. L. Krieg, G. Parikesit, and J. J. M. Braat, "Three-wavelength laser light source for absolute, sub-Angstrom, two point source interferometer," in *Optical Measurement Systems for Industrial Inspection III*, vol. 5144, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2003, pp. 227-233.

[2] M. Bass and Optical Society of America., *Handbook of optics*, 2nd ed. New York: McGraw-Hill, 1995.

[3] R. N. Shagam and J. C. Wyant, "Heterodyne Interferometer for Measuring Surface Profiles from Interferograms," *Journal of the Optical Society of America*, vol. 67, pp. 1365-1365, 1977.

[4] N. Bobroff, "Residual Errors in Laser Interferometry from Air Turbulence and Nonlinearity," *Applied Optics*, vol. 26, pp. 2676-2682, 1987.

[5] C. M. Wu and R. D. Deslattes, "Analytical modeling of the periodic nonlinearity in heterodyne interferometry," *Applied Optics*, vol. 37, pp. 6696-6700, 1998.

[6] Y. Y. Cheng and J. C. Wyant, "2-Wavelength Phase-Shifting Interferometry," *Applied Optics*, vol. 23, pp. 4539-4543, 1984.

[7] B. V. Dorrio and J. L. Fernandez, "Phase-evaluation methods in whole-field optical measurement techniques," *Measurement Science & Technology*, vol. 10, pp. R33-R55, 1999.

[8] D. Malacara, *Optical shop testing*, 2nd ed: Wiley-Interscience, 1992.

[9] R. G. Klaver, "Novel interferometer to measure the figure of strongly aspherical mirrors." Delft: Delft University of Technology, 2001.

[10] M. L. Krieg, B. L. Swinkels, and J. J. M. Braat, "Characterization of the frequency modulated continuous wave subsystem of an angstrom accuracy absolute interferometer," in *Interferometry XI: Applications*, vol. 4778, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2002, pp. 131-141.

[11] R. Dandliker, Y. Salvade, and E. Zimmermann, "Distance measurement by multiple-wavelength interferometry," *Journal of Optics*, vol. 29, pp. 105-14, 1998.

[12] T. Ikegami, S. Sudo, and Y. Sakai, *Frequency stabilization of semiconductor laser diodes*. Boston: Artech House, 1995.

[13] S.-H. Lu, R.-H. Hsu, C.-P. Lai, and C.-J. Chen, "Iodine Stabilized He-Ne Laser at 633nm with Simple Line-Locking Device and the Intercomparison Results," *Precision Electromagnetic Measurements*, 1996.

[14] H. P. Layer, "A portable iodine stabilized helium-neon laser," *IEEE-Transactions-on-Instrumentation-and-Measurement.*, vol. IM-29(4), pp. 358-61, 1980.

[15] A. J. Wallard, "A practical approach to the design and construction of iodine stabilized lasers," Nat.-Phys.-Lab., -Teddington,-UK March 1979: 1979.

[16] H. P. Layer, "Recent developments in iodine stabilized helium-neon lasers," presented at CPEM-Digest-1980.-Conference-on-Precision-Electromagnetic-Measurements, Braunschweig, West Germany, 1980.

[17] W. G. Fastie, "An Iodine Absorption Line as a Primary Standard of Length," *Journal of the Optical Society of America*, vol. 47, pp. 120-120, 1957.
[18] C. S. Edwards, G. P. Barwood, P. Gill, and W. R. C. Rowley, "A 633 nm iodine-stabilized diode-laser frequency standard," *Metrologia*, vol. 36, pp. 41-45, 1999.
[19] A. Zarka, J. M. Chartier, J. Aman, and E. Jaatinen, "Intracavity iodine cell spectroscopy with an extended-cavity laser diode around 633 nm," *IEEE Transactions on Instrumentation and Measurement*, vol. 46, pp. 145-148, 1997.
[20] W. T. Estler, "High-Accuracy Displacement Interferometry in Air," *Applied Optics*, vol. 24, pp. 808-815, 1985.
[21] H. Handschuh, J. Froschke, M. Julich, M. Mayer, M. Weiser, and G. Seitz, "Extreme ultraviolet lithography at Carl Zeiss: Manufacturing and metrology of aspheric surfaces with angstrom accuracy," *Journal of Vacuum Science & Technology B*, vol. 17, pp. 2975-2977, 1999.
[22] P. W. Atkins, *Physical chemistry*, 4th ed. New York: Freeman, 1990.
[23] C. A. Guenter, M. H. Welch, and J. C. Hogg, *Clinical aspects of respiratory physiology*. Philadelphia: Lippincott, 1978.
[24] B. Dorband and G. Seitz, "Interferometric testing of optical surfaces at its current limit," *Optik*, vol. 112, pp. 392-398, 2001.

5 Sensor

In the early stages of this project, it was concluded that the type of sensor required to perform measurements with sufficient accuracy was not commercially available, and a separate PhD project should be devoted to its design, testing and further development. At the time of writing, no prototype sensor was yet available and alternative sensors had to be used to perform the measurements presented in this thesis. The optical system has been designed in such a way as to permit the implementation of a number of different detection schemes, each suited to one or more of the sensors considered. Two of these sensors, a standard CCD and a phase-measuring pixel detector, will be treated in detail and subjected to performance tests to estimate to which extent they realistically limit the measurement accuracy. A third potential sensor of the “fast sampling” type, will also be discussed briefly.

5.1 Requirements

Ultimately, the sensor must be able to measure the relative optical phase of the object and reference beams from our interferometer with sufficient accuracy to permit the inverse propagation algorithm to calculate the shape of the mirror under test to within the required accuracy of 0.1nm. Since the three sensors treated in this chapter are each based on a very different measurement strategy, it is impossible to specify a comprehensive list of quantitative requirements without first discussing the measurement strategy. While the subsequent chapters will do just that, we can first make a number of qualitative comments about requirements common to all three types of sensors.

The optical path difference (OPD) between the reference and optical beam is expected to vary very rapidly over the area of the sensor, resulting in a fringe density of ~50mm^{-1}. While the fringe density usually places a restriction on the minimum pixel density in order to resolve the fringes, the fact that we employ multiple wavelength interferometry means that in our case, this condition is relaxed to a restriction on the pixel size rather than the density.

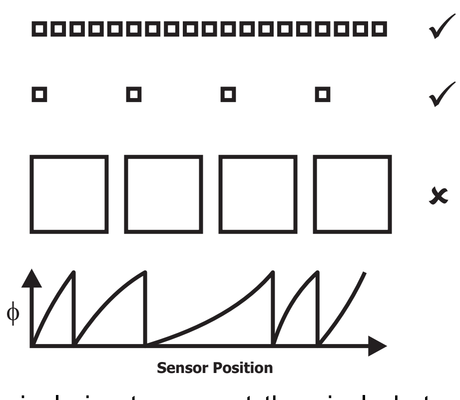

Figure 5-1 – Restriction on pixel size to prevent the pixel photosensitive area from spanning more than one fringe. The top two situations (marked with ticks) are acceptable, while the third situation (marked with a cross) is not.

In other words, it is acceptable to have more than one fringe between adjacent pixels, but we may not have more than one fringe over a pixel's photosensitive area (see Figure 5-1).

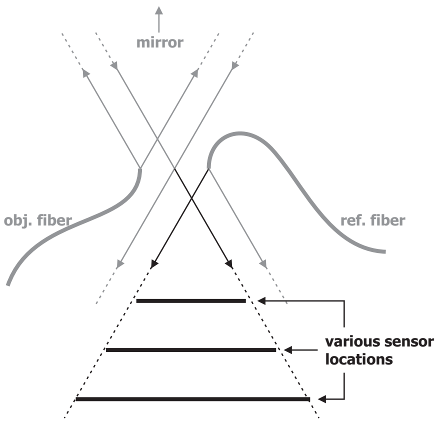

Figure 5-2 – Dependence of sensor size on sensor position

Unlike conventional applications where miniaturization is often the goal, a bigger total sensor area is advantageous for our interferometer. In our case, placing the sensor further away from the fiber tips typically improves the reference wavefront sphericity as well as reducing the interferometer's sensitivity to positioning errors of the pixels on the sensor and relaxing the minimum pixel size requirement. However, we need to maintain the same field of view regardless of the sensor's position. Therefore, the further from the fiber tips we wish to place our sensor, the bigger the total sensor area needs to be (see Figure 5-2).

Since our light source only has a finite output power and the reflection off our mirror substrate is a mere 4%, the sensor needs to be sufficiently sensitive to allow the shot noise limited signal-to-noise ratio (SNR) to meet or exceed our requirements. Although averaging can be used to reduce the influence of such noise sources, our measurement time should be limited to ~1s to prevent drifts in the positions of interferometer elements from affecting our accuracy, and hence sufficient amount of charge should be collected by the sensor in this time.

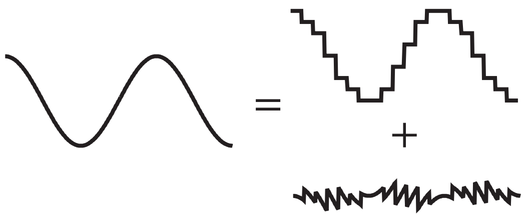

Figure 5-3 – Decomposition of an analogue signal into its digital form and the remaining quantization noise

Even in the case of noise-free detection, an upper limit is placed on our phase measurement accuracy by the bit-depth of the data from the sensor. Errors caused by this limited bit-depth are said to be due to quantization noise (see

Figure 5-3). However, since the total noise is a combination of quantization noise and shot noise (as well as other types of noise like thermal- and flicker noise), we will experience no significant gain in accuracy by reducing the quantization noise far below the shot noise, or vice-versa. Most sensors are designed in such a way that when they are used near their saturation point, the electronic noise lies just below the quantization noise.

Since there are drifts in the optical path length within our light source as well as in the positions of our interferometer elements, the quicker a phase measurement is performed, the more immune it is to such drifts. This is also the primary reason for considering heterodyne methods over phase shifting methods for the final interferometer.

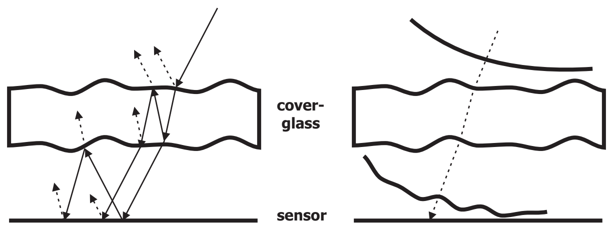

Figure 5-4 – Diagram illustrating the effects of multiple reflections and wave- front distortions due to a cover-glass

Another common source of error is the presence of a protective “cover-glass” on most commercial sensors. Despite the fact that the light from the interferometer passes through this glass almost common-path, the distortions of the reference and object wave fronts differ appreciably, so that a net optical path difference (OPD) error is introduced. Multiple reflections at the interfaces also serve to introduce systematic errors in our measurement (see Figure 5-4). Our ideal sensor would have no such cover glass.

It is clear that several of these requirements are conflicting (such as the sensor area- and pixel positioning accuracy requirements or the measurement time- and sensitivity requirements). Klaver[1] has already looked at the resulting trade-offs in detail, but recent developments have relaxed several of the requirements assumed in this treatment. Instead of generating a new and generally applicable parameter space of suitable sensor characteristics, we will look at the performance of several available sensors for our interferometer; allowing us to identify the characteristics which limit their accuracy.

5.2 CCD

Although high speed CCDs are now available with frame rates exceeding 1kHz, we have chosen a simple video frame rate CCD to make our proof-of-concept measurements. Should other, more promising, detectors prove to be unfeasible, such high frame rate cameras could potentially further improve the measurement accuracy of this approach.

Depending on the type of shutter mode and frame rate of the camera, we can use the light source in either continuous heterodyne- or discrete phase stepping mode. For high frame rates and full-frame shutter modes, continuous heterodyne operation will be the most suitable, since the abrupt phase-stepping of the Acousto-optic modulators (AOM) introduces noticeable transients when the phase is stepped at frequencies above 1kHz. In the absence of a full-frame shutter, the effective time at which a sample is recorded varies from pixel to pixel, and it is therefore better to use phase-stepping.

The CCD used in our proof-of-concept measurements is a Monochrome Sony XC-77RR-CE interline transfer CCD. It has a sensing area of 8.8x6.6mm, containing 756 x 580 pixels. The specifications of the camera state a cell size of 11x11μm, but no information is given about the size of the actual photosensitive area per pixel. Since smaller pixel sizes are favourable for our application, we have assumed a worst-case scenario of cell size = pixel size.

5.2.1 Data acquisition

For this camera, phase stepping interferometry (PSI) was chosen over heterodyne detection due to the absence of a full-frame shutter mode. Figure 5-5 shows the test-setup used to evaluate the camera. A phase-step of $\pi/2$ was performed before each recorded frame, by triggering a phase-step in the AOM driving signal of the light source with a synchronization signal from the CCD's driving electronics. We then record several frames at 25 fps (frames per second) using a National Instruments frame-grabber card with a bit-depth of 8 bits.

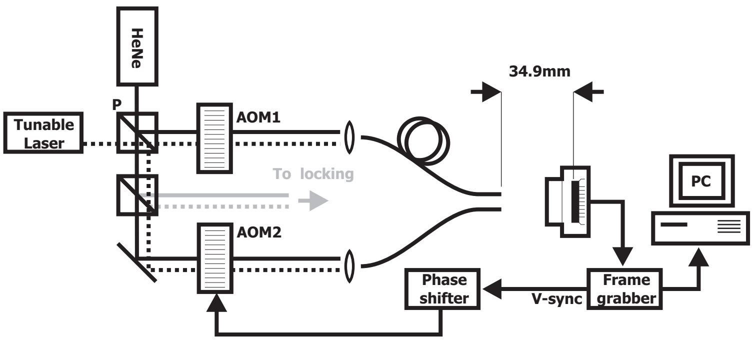

Figure 5-5 – Test setup to evaluate the CCD

The data collected consists of a series of intensity-images (40 in our case), each with a phase shift of $\pi/2$ with respect to the previous frame. Figure 5-6 shows a sequence of four such images. These sequences of intensity-images now need to be turned into phase-maps. The well known Wyant 4-bucket algorithm[2] given by equation (5.1), shows how four successive intensity values at a point in the interferogram ($I_1,\ldots,I_4$) can combined to give the phase at that point:

$$\phi = \arctan\left(\frac{I_4 - I_2}{I_1 - I_3}\right) \tag{5.1}$$

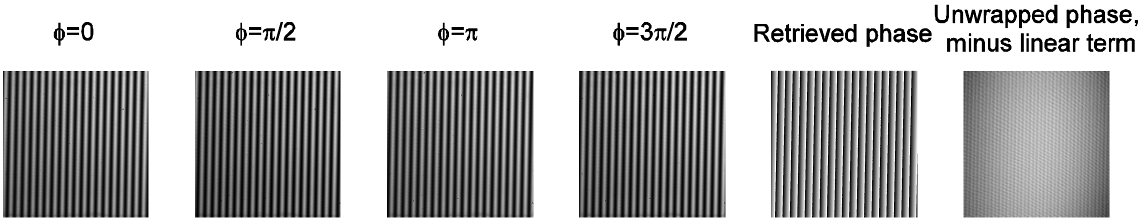

Figure 5-6 – Sequence of four consecutive intensity images having undergone incremental phase-shifts. Also shown is the resulting retrieved phase.

Besides this algorithm, there is a myriad of other algorithms to perform the same task. They vary in terms of the number of intensity images considered (Schwider-Hariharan - 5 images, Zhao & Surrel – 6, Hibino – 9, etc.), the size of phase-steps required between images (Carre – arbitrary phase-step, Larkin & Oreb - $\pi/3$, etc.), and the way in which the intensities are combined before the arctangent operation is performed. These algorithms have been extensively reviewed[3], and numerically, as well as experimentally, compared for robustness against various systematic and stochastic sources of error[4]. Essentially, it can be concluded that algorithms using more images are more accurate, even when compared to the average of a sequence of phase values obtained with lower-image-number algorithms spanning the same total number of images. Based on the relevant reviews in the literature[3-5] and our own numerical simulations, we conclude that the windowed discrete Fourier transform method (WDFT), shown in equation (5.2), is the most suitable for our purposes. In fact, the Wyant algorithm is simply a special case of the WDFT algorithm, using N=4, and a uniform windowing function.

$$\phi = \arctan\left(\frac{\sum_{j=0}^{N-1} w_j I_j \sin\left(-\frac{2\pi}{N} j\right)}{\sum_{j=0}^{N-1} w_j I_j \cos\left(-\frac{2\pi}{N} j\right)}\right), \quad N \in \{4,8,12,\ldots\} \tag{5.2}$$

with:

$$w_j \geq 0 \quad \forall j \tag{5.3}$$

Here, w_j represent the discrete values of a “windowing function”, such as the Hanning, Blackman-Harriss, or Hamming windows (Figure 5-7). We have chosen to use the Hanning window. The results of applying this algorithm to our data can also be seen in Figure 5-6.

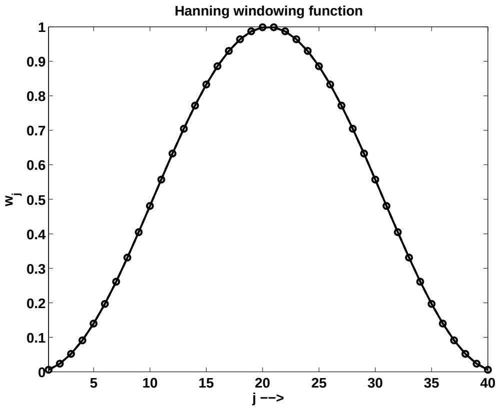

Figure 5-7 – Hanning windowing function for WDFT method.

Although the choice of algorithm is strongly dependent on the expected sources of error in the measurement, we have presented the choice of algorithm before the treatment of the error sources to enhance readability.

5.2.2 Error sources

An accurate phase measurement is impeded by a number of sources of error, both systematic and stochastic in nature. The stochastic errors include:

- Intensity fluctuations during the measurement process.
- Shot-, flicker- and thermal-noise.
- Zero-mean phase fluctuations (e.g. due to vibrations).
- Quantization noise (although technically a systematic source of error, its behaviour is largely stochastic)

Our WDFT method already averages these effects over the measurement period. Due to the large number of available pixels, a further reduction of these errors can be expected from the averaging of neighbouring pixels. Regions where the modulation depth is below the noise limit can easily be identified and discarded before further processing.

The systematic sources of error present a greater challenge:

- Camera nonlinearity. The resulting phase error is a function of both mean intensity and modulation depth.
- Systematic phase-drift during the measurement (i.e. with nonzero mean). The resulting phase error is a function of the actual phase value only
- Cross-talk between the reference- and object-beams before entering the fibers. The resulting phase error is function of the actual phase and the modulation depth.
- Wavefront distortion and multiple reflections from cover-glass. The resulting phase error is a function of pixel position and interferometer geometry.
- Pixel positioning- and dimension- errors. The resulting phase error is a function of pixel position.

- Electrical inter-pixel cross-talk. The resulting phase error is similar to a weak spatial filtering of the phase map.

Intensity response
The robustness against a nonlinear intensity response is a particular strength of the WDFT algorithm. A nonlinear intensity response of the camera is equivalent to a deformation of the sinusoidal fringe profile. This in turn is equivalent to the introduction of higher harmonics, which are automatically filtered out by the algorithm, especially in the presence of small stochastic phase variations. Although the camera could be calibrated, and a correction applied to the data before the phase-retrieval algorithm is used, the excellent linearity and noise characteristics of the Sony camera, together with the robustness of the algorithm make this unnecessary. (See also the discussion in 14.10.2 in ref[2]).

Phase Drift
Phase drift does introduce a significant error in the measurement, and while most phase retrieval algorithms are immune to a linear phase drift, nonlinear drifts are not cancelled. Figure 5-8 shows a typical phase drift during a 40-frame measurement and the resulting phase-error from our WDFT algorithm.

The fact that this phase error is only a function of the actual phase, offers us a unique way of reducing its influence after having performed the measurements, which will be explained below. As far as we are aware, this method has not previously been reported, and takes advantage of the large fringe density in our interferograms – a situation otherwise often avoided in metrology.

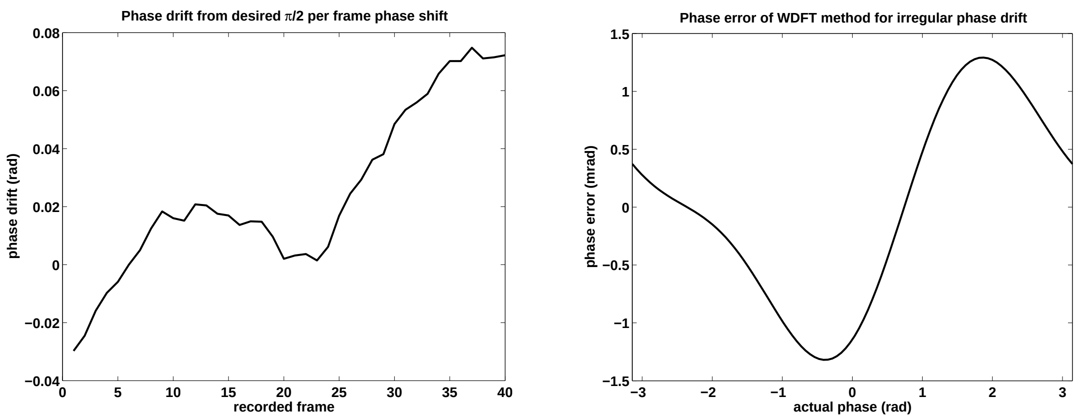

Figure 5-8 – Typical phase drift during 40-frame acquisition, and corresponding phase error

Having a large number of fringes over a large number of pixels, especially if no "stationary" points (of low fringe density) are present, means that the statistical distribution of phase values over the CCD should be uniform. That is, there should be no intrinsic preference for particular phase values. This is in agreement with simulations of interference patterns similar to ones to be expected in our set-up. Introducing phase-drift causes a significant and easily observable change in the statistical distribution of measured phases.

This can be understood with the help of the diagram in Figure 5-9. Due to the nonlinearity of the actual- to measured-phase mapping, certain regions of actual phase will be mapped to much smaller regions of measured phase, meaning that such measured phases will occur more often in the interferogram, and vice versa.

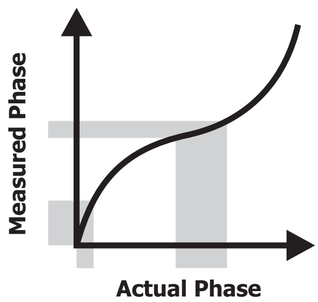

Figure 5-9 – Redistribution of phase values causing a change in probability density function

The expression relating the mapping function to the statistical distributions can easily be derived to be:

$$p_{meas}(\phi_m) = \frac{1}{\left(\frac{\partial\phi_m}{\partial\phi_a}\right)} p_{actual}(\phi_a) \tag{5.4}$$

Where '*p*' stands for the probability density function. If we now assume a uniform distribution of the actual phases and only a very small phase-error in our measurement, so that $\phi_m = [1+\varepsilon(\phi_a)]\phi_a$, we can write:

$$p_{measured}(\phi_m) = \frac{1}{1+\dfrac{\partial\varepsilon}{\partial\phi_a}} \approx 1 - \frac{\partial\varepsilon}{\partial\phi_m} \tag{5.5}$$

$$\Rightarrow \varepsilon(\phi) \approx \int_{-\pi}^{\phi} \left[1 - p_{measured}(\phi_a)\right] d\phi_a \tag{5.6}$$

Although we cannot uniquely determine the drift from the statistical distribution of measured phases, we can find the actual- to measured-phase mapping (by finding ε), and use its inverse to correct our observed phases.

We therefore proceed as follows: Our series of intensity images is turned into a phase-map by using the WDFT algorithm, and a histogram of the phase distribution is calculated for 128 phase regions. After low-pass filtering the distribution to discard stochastic noise, we use the resulting "smoothed" histogram to calculate the forward-mapping function according to equation (5.6). The mapping is inverted using an interpolated look-up table, and applied to the phase-map.

The application of this method is shown in Figure 5-10. The test setup of Figure 5-5 was used to generate an interferogram with a low fringe-count for easy unwrapping. A distinct reduction of artefacts – vertical streaks in the same direction as the interference fringes – can be seen. It is important to realize that although the results look similar, this is not a spatial filtering technique. The observed artefacts are removed purely by re-mapping the phases. The method can be used to reduce phase errors due to unknown phase drifts during a PSI measurement. It should not be used for interferograms with areas of low fringe-density, since such regions will upset the uniform phase distribution.

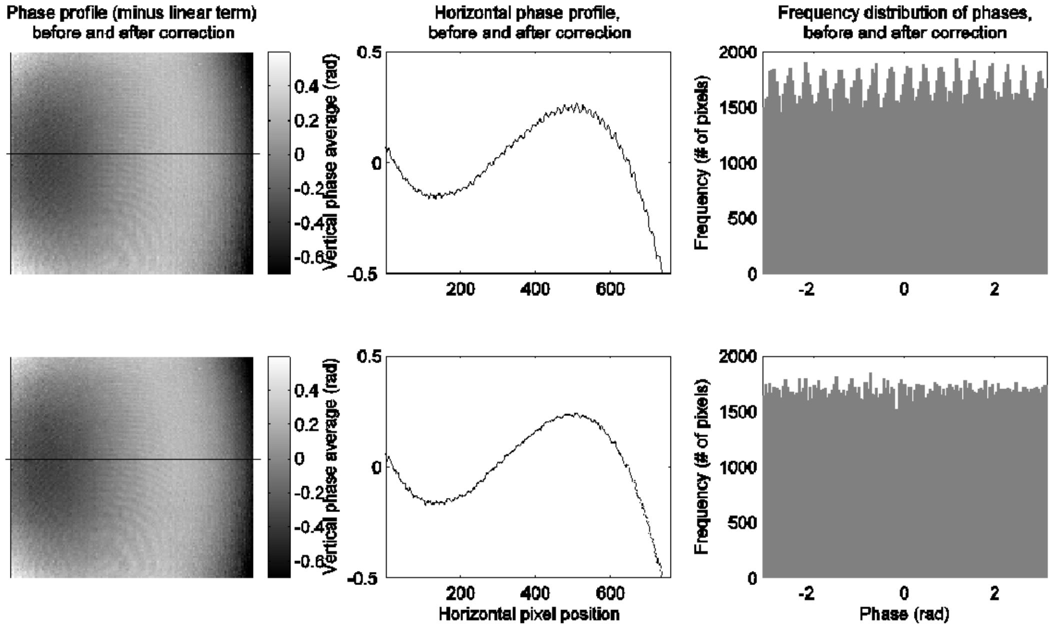

Figure 5-10 – Difference between corrected (bottom) and uncorrected (top) interferograms. The artefacts remaining in the corrected image could be due to multiple reflections in the CCD cover glass.

Cross-talk

Another error in our phase measurement is that of cross-talk between the reference- and object-beams before entering the fibers. For regions of fairly uniform modulation depth, this error is also corrected by the above procedure, since it behaves like a phase-dependent error (See chapter 4). When the modulation-depth varies significantly over the CCD, the error becomes modulation-depth dependent and as such, more difficult to remove. Great care has been taken in designing the light source to avoid any cross-talk, especially at the primary metrology wavelength, where there is no crosstalk whatsoever. The secondary and tertiary wavelengths have been measured to exhibit cross-talk below 10^{-3}. We believe it is inadvisable to attempt modulation-depth dependent correction of this error, as we cannot confidently distinguish between errors due to cross-talk, and errors due to the actual mirror figure in this case.

Coverglass related errors.

The presence of a cover-glass introduces two types of errors. Firstly, wavefront distortion is introduced even by a perfectly plane-parallel cover-glass, which manifests itself as a low spatial frequency OPD error. In the case of a textured cover-glass surface, this wavefront distortion is of similar spectral content as the cover-glass surface features. Secondly, multiple reflections between the front- and back-surface of the cover-glass, as well as reflections involving the CCD itself, will interfere with the purely refracted beams, causing a small phase error, which in turn also translated to an error in the measured OPD. High spatial frequencies are characteristic of multiple reflection induced OPD errors. Novak[6] states that, compared to the reflections occurring at the front- and back-surfaces of the cover-glass, the reflections involving the CCD surface may be neglected.

While the form of the OPD errors due to wavefront distortion is not very sensitive to the position and separation of the fibers, the multiple-reflection induced errors are very sensitive to these factors. The converse is true for the size of these errors. The OPD errors induced by wavefront distortion increase with increasing fiber separation, while the size of the multiple reflection induced OPD errors does not depend on the fiber separations at all, instead depending almost exclusively on the reflectivities of the various surfaces involved.

To illustrate this, several types of cover-glass shapes were simulated using our raytracing approach, and the resulting OPD errors were calculated. The approach involved the tracing of four rays to each detection point (See Figure 5-11): One purely refracted ray and one multiply reflected ray from each the “reference” and “object” point sources. The object point-source was placed at the focus of the object beam as an approximation. The multiply relfected rays underwent two reflections at the back- and front-face of the cover glass, n=1.54 with an estimated reflectivity of ~0.5% due to antireflection coatings. The starting directions of the rays were adaptively chosen to intersect the CCD surface at the exact pixel positions. Figure 5-11 – Figure 5-14 show the various cover-glass shapes and the resulting OPD errors.

The un-wrapped total effective OPL for each point-source was calculated as follows:

$$OPL \approx OPL_t + \frac{\lambda}{2\pi} R \sin\left(\frac{2\pi}{\lambda}\left[OPL_l - OPL_t\right]\right) \tag{5.7}$$

Where OPL_t and OPL_l stand for the “transmitted” and “reflected” OPLs respectively, and R is the intensity reflection coefficient of one surface.

The first term may be attributed to the wavefront distortion, while the second term is due to the field arising from the multiple reflections. To illustrate the difference between OPD errors from these two effects, the terms are shown separately in the figures. The approximation above is valid for reflectivities of the surfaces below ~5%.

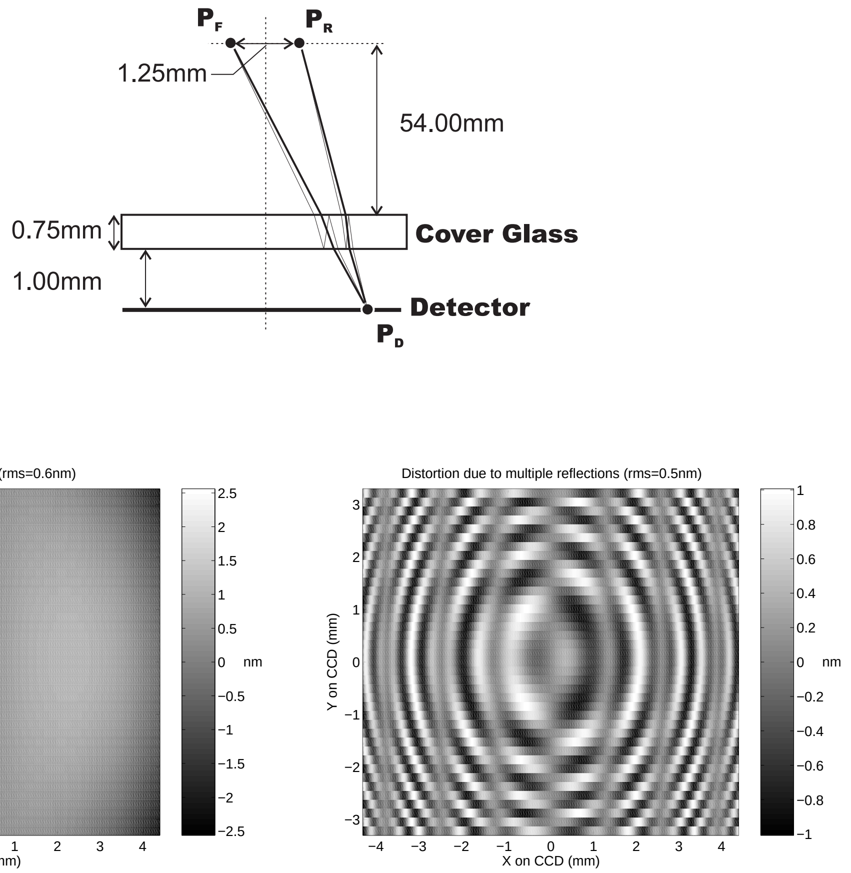

Figure 5-11 – OPD errors introduced by a plane-parallel cover glass (top). Errors due to wavefront distortion (bottom left) have been separated from errors due to multiple reflections (bottom right). Piston and tilt have been removed.

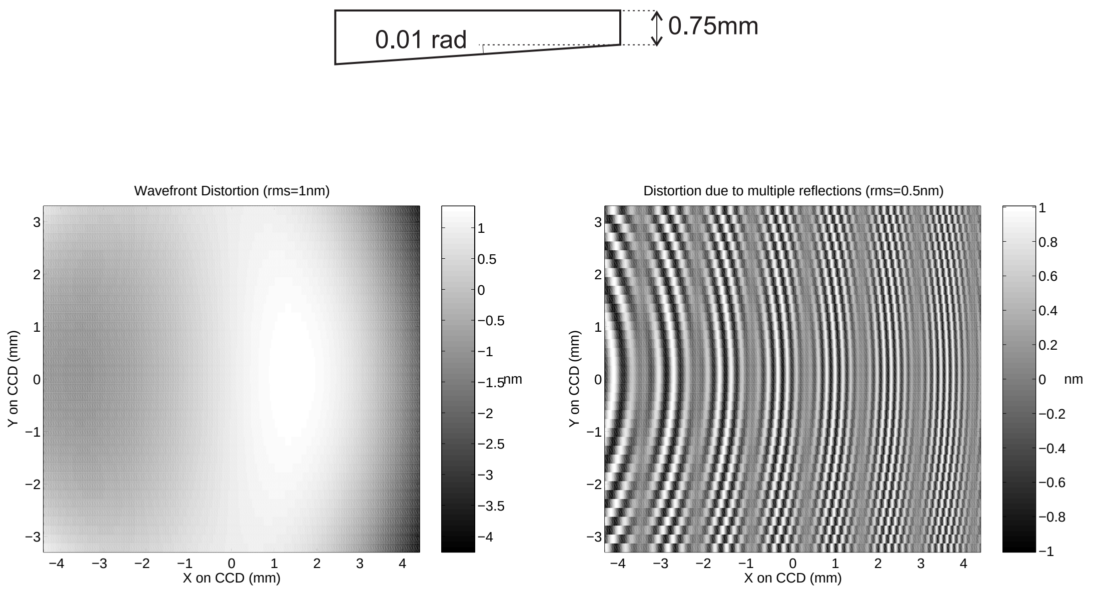

Figure 5-12 – OPD errors introduced by a wedged cover glass (top). Errors due to wavefront distortion (bottom left) have been separated from errors due to multiple reflections (bottom right). Piston, tilt and wavefront errors due to a plane glass have been removed.

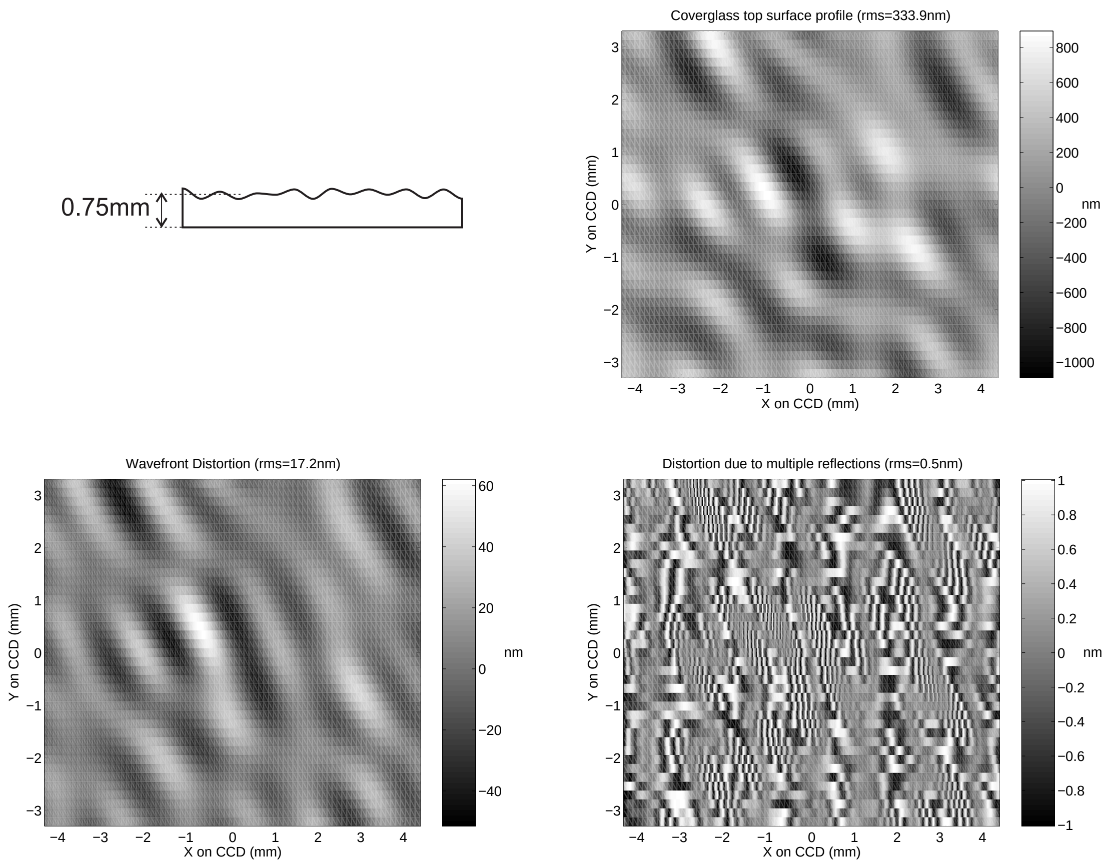

Figure 5-13 – OPD errors introduced by a cover glass with a 334nm rms random surface profile (top left and right). Piston, tilt and wavefront errors due to a plane glass have been removed. Note the difference between the surface profile and the resulting error.

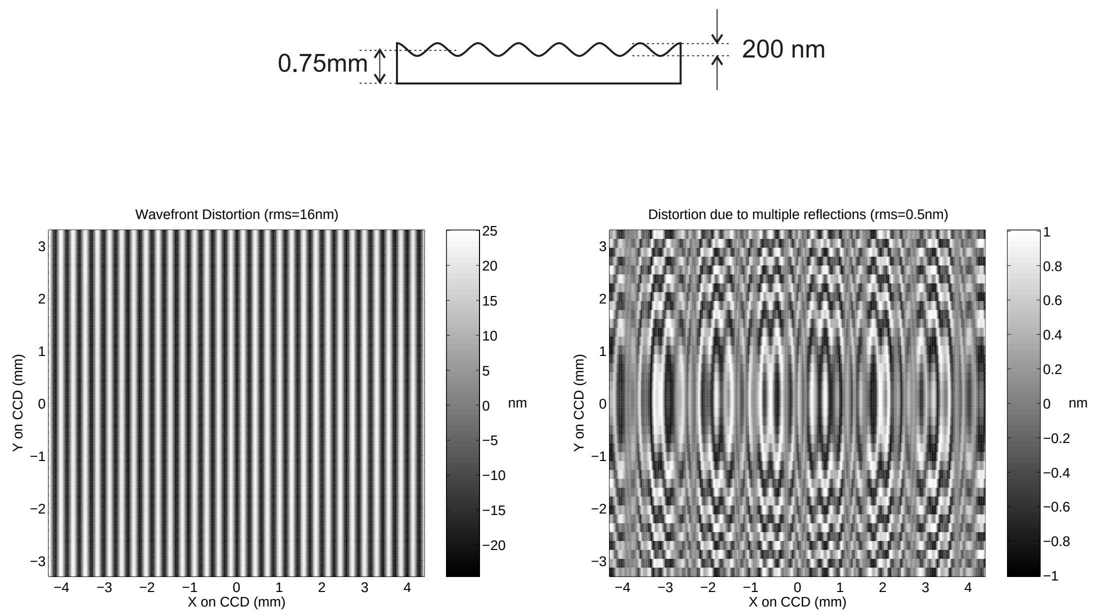

Figure 5-14 – OPD errors introduced by a cover glass with a sinusoidal ripple on the front surface, with a period of 0.28mm and an amplitude of 100nm (top). Piston, tilt and wavefront errors due to a plane glass have been removed.

The effect of interference between the transmitted and multiply reflected beams on the intensity distribution of one point source alone is also well approximated by the following formula:

$$I \approx I_t \left\{ 1 + R \cos\left(\frac{2\pi}{\lambda} \left[OPL_l - OPL_t \right] \right) \right\} \tag{5.8}$$

The intensity distribution on the CCD from a single point is therefore an indication of the type of effects to be expected on the OPD due to multiple reflections.

Figure 5-16 shows artefacts potentially due to the presence of a CCD cover-glass. A low fringe density image obtained with our testing set-up was un-wrapped, and the theoretical OPD for this configuration (Figure 5-15) was subtracted. The diagonal streaks observed are similar in structure to fringes which appear on intensity images with only one beam illuminating the CCD, pointing towards multiple reflections as possible cause.

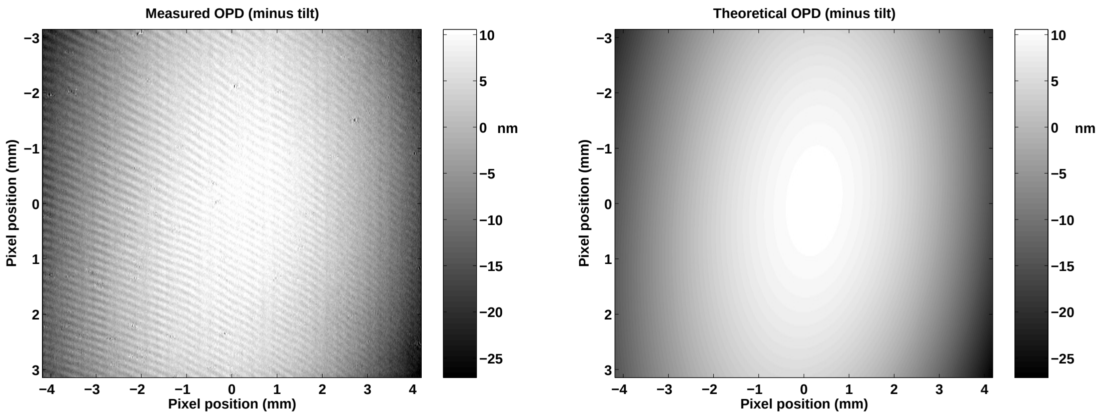

Figure 5-15 – OPD measurement from test-setup shown in Figure 5-5, together with the theoretical OPD profile, for a small fiber separation of ~0.13mm

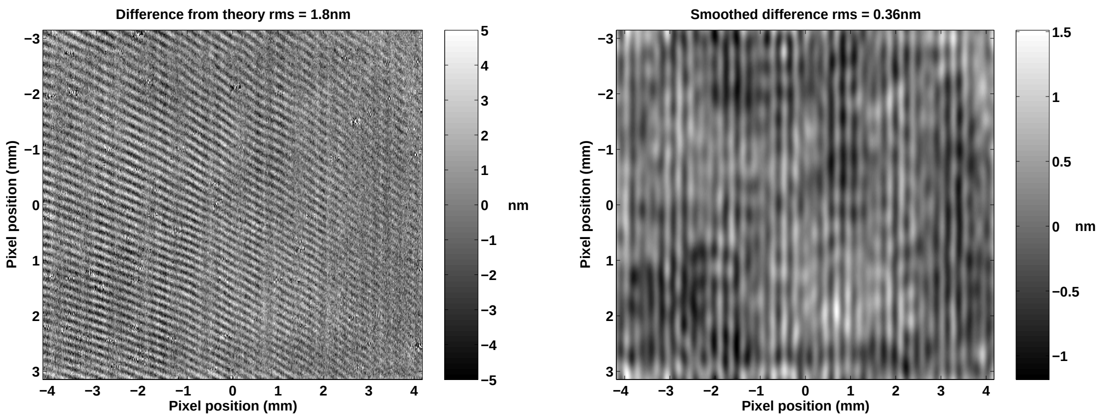

Figure 5-16 – Difference image of the two images in Figure 5-15, showing artefacts potentially due to the presence of a cover glass in the un-filtered image (left) and in a smoothed version of the same image (right)

Use of spatial filtering to average out the diagonal streaks in Figure 5-16 reveals yet another error structure on our interferogram. We will see in the next section how this structure becomes more dominant as the fringe density is increased, and will form a limiting factor of our measurement accuracy.

We have tried to remove the cover glass from a Sony CCD described above, but the removal process proved difficult, even for the optical workshop of TPD. During the removal process, the CCD was damaged, causing missing lines and excessive noise in the measured data. However, we believe that it should be possible to obtain CCDs from their manufacturers without a cover-glass. Future work will be carried out with such sensors.

As the fiber separation (and hence also the fringe-density) is increased towards the separation used in our final setup, we find that the spatial frequency of the diagonal-stripe features quickly exceeds the sampling density, and the feature ceases to be visible. This is consistent with the behaviour of multiple reflection related patterns arising from a significant wedge in the cover glass. The vertical stripe pattern on the other hand maintains the same spatial frequency and becomes more pronounced (i.e. grows in amplitude). We can therefore not attribute these two patterns to the same cause. While the latter behaviour is consistent with wavefront-aberration induced errors, the next subsection will show data which suggests a different cause.

Pixel positioning- and size-errors

Another potential cause for the vertical lines seen in Figure 5-16, and again in our measurement results of chapter 7, are pixel positioning and size-errors.

Such errors can occur during the lithography process of a CCD if the mask alignment or even the masks themselves are imperfect. In such a case, the pixel position errors and pixel size errors can be expected to be related. The vertical stripes mentioned earlier are characteristic of a simultaneous and systematic variation of pixel size- and position. We will now treat the effects of these two types of errors separately before estimating the magnitude of the errors required to produce the observed effect.

Due to the finite size of an individual pixel, problems arise when the fringe period approaches the dimensions of a pixel. The pixel effectively averages the fringe-pattern over its area[2], which affects the observed modulation depth. The ratio of the measured modulation depth to the actual modulation depth is called the Modulation Transfer Function (MTF) of the CCD, and depends on the pixel size (d_{pixel}) and the fringe separation ($d_{fringes}$):

$$MTF = \mathrm{sinc}\left(\frac{\mathrm{d}_{\mathrm{pixel}}}{\mathrm{d}_{\mathrm{fringes}}}\right) \tag{5.9}$$

When the pixel size is exactly one fringe period, the MTF is zero. This gives us the restriction of a maximum pixel size as mentioned in section 5.1. The modulation depth will vary over the CCD due to two factors – a varying amplitude ratio between the reference and object beams, and a variation of the

MTF. The former factor will cause the modulation depth to vary slowly and smoothly over the area of the CCD, due to the Gaussian intensity profile of the two beams. Any other variation in modulation depth could be due to a variation in pixel size.

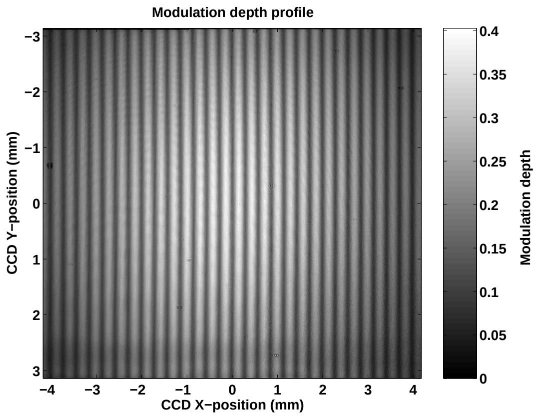

Figure 5-17 – Modulation depth profile showing oscillations in the x-direction.

Figure 5-17 shows such a variation of modulation depth in an interferogram taken with our test setup, using parameters similar to our final setup. The period of the pattern is 26.2 pixels. The fiber-to-CCD distance was 34.6mm, and the fiber-to-fiber distance was 1.25mm, giving a fringe spacing of 17.5μm, compared to the reported pixel-dimensions of 11x11μm. The MTF is therfore 0.47 – meaning that the modulation depth cannot exceed this value. The oscillations in the x-direction are an indication that the sizes of the pixels vary along the x-direction. We needn't correct the resulting variation in the CCD's MTF, but an accompanying effect on the pixel spacing could pose a serious problem.

Up until now, we have assumed that the points at which the OPD function is being sampled are regularly spaced in a plane, according to the manufacturer's specifications of the CCD area and number of pixels. Figure 5-17 should give us cause to draw this assumption into question. We can quantify the effect of an in-plane pixel position error as follows:

$$\begin{aligned} OPD_{measured}(x_i) &= OPD(x_i + \varepsilon_i) \\ &= OPD(x_i) + \varepsilon_i \left.\frac{\partial OPD}{\partial x}\right|_{x_i} + \frac{\varepsilon_i^2}{2!} \left.\frac{\partial^2 OPD}{\partial x^2}\right|_{x_i} + \ldots \end{aligned} \tag{5.10}$$

Where x_i is our assumed position of the pixel, and ε_i is the error in this assumption. For both the test setup in Figure 5-5 and our actual measurement configuration, the OPD function is predominantly linear, with the first derivative dominating over the second- and successive derivatives by several orders of magnitude. To first order, our OPD function is given by:

$$OPD(x_i) \approx OPD(0) + \frac{s}{d} x_i \tag{5.11}$$

Where *s* is the separation of the two fiber-tips in the test set-up, or the distance between the focus of the reflected object beam and the reference fiber-tip in the full setup, and *d* is the distance between the fiber-tips and the CCD. Therefore, our measured OPD, in the presence of a small positioning error, is:

$$OPD_{measured}(x_i) \approx OPD(x_i) + \varepsilon_i \frac{s}{d} \tag{5.12}$$

In other words, the OPD error will be directly proportional to the positioning error. For the test setup used to generate Figure 5-17, we have *s*=1.25mm and *d*=34.9mm. Hence the positioning error will appear directly on our OPD measurement, attenuated by a factor of 27.92.

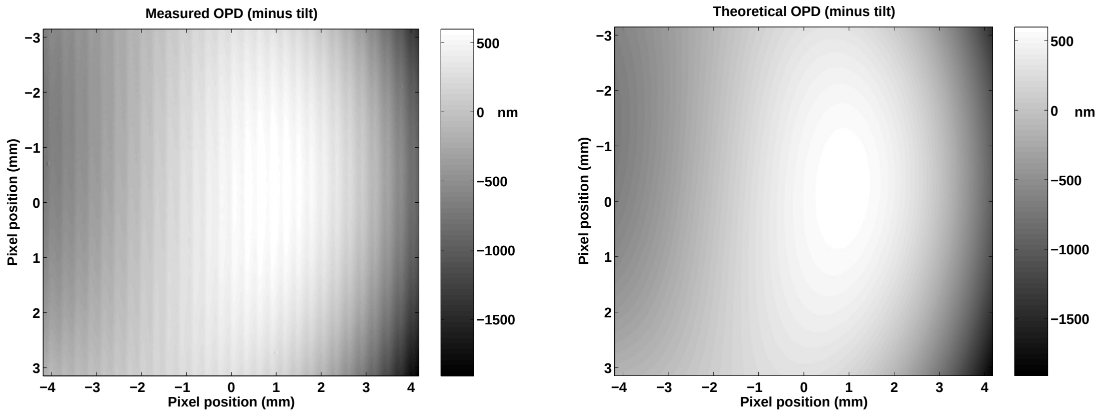

Figure 5-18 –Measured and theoretical OPD from our test-setup, with a high fringe density similar to the one to be expected from our final measurement set-up.

Figure 5-18 shows such an OPD measurement – from the same measurement which also gave rise to the modulation depth profile of Figure 5-17. The difference between our OPD measurement and the theoretical OPD (Figure 5-19) shows the same 26.2-pixel-period sinusoidal “horizontal wobble” as the modulation depth profile – giving us reason to believe that this error is indeed due to non-uniformities of the features on the CCD.

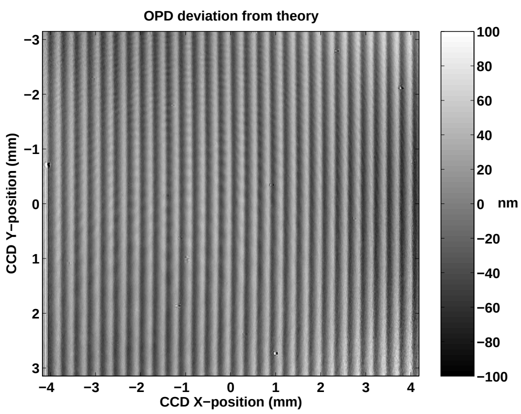

Figure 5-19 – Diagram showing measured OPD minus the theoretical values (rms = 29nm). The deviations show the same structure as the modulation depth variations of Figure 5-17.

Using the amplitude of the oscillations in the cross-section of this measurement error as an indication, together with equation (5.12), we arrive at a maximum absolute deviation from the ideal pixel positions of +/- 0.5μm (See Figure 5-20). The pixel-to-pixel positioning error is significantly less of course.

Such positioning errors seem excessive, despite the age of the CCD. At the time of manufacture, <1μm processes were already being used. Nonetheless, the experimental evidence strongly points towards a systematic pixel position- and size-error.

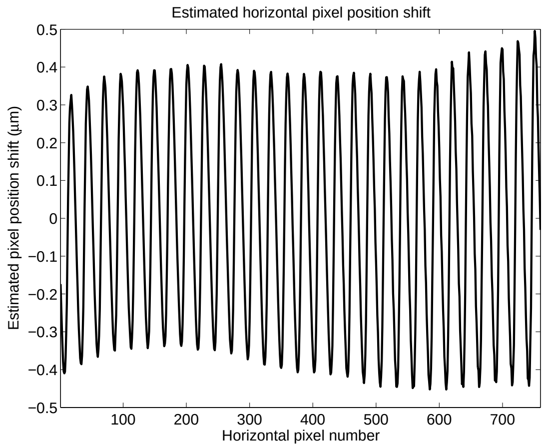

Figure 5-20 – Cross-section of pixel-positioning errors derived from deviation of the measured OPD from theory.

It is observed that the amplitude of these errors for different fiber-to-fiber distances decreases with reduced fringe density. This is in accordance with equation (5.12), which indicates a reduced sensitivity to CCD positioning errors at lower OPD slopes. Wavefront aberration errors, discussed in the previous subsection (see Figure 5-14), can also explain this behaviour, but fail to account for the observed variation in modulation depth.

We can attempt to correct for in-plane CCD geometry errors by estimating the correct effective location of the pixels to first order, using our error measurements from the test-setup together with equation (5.12). Once we have an estimate of the correct pixel locations, we can use interpolation techniques to get the data at our desired points. Such a calibration procedure has recently also been carried out by Sommargren[7] using a setup virtually identical to ours. Unfortunately, the presence of a coverglass makes such a calibration impossible.

Considering that <0.2μm processes are now available, we may expect a significant improvement in the accuracy of PSI methods using CCDs. Another way to reduce the influence of these errors is to measure further from the fiber tips. In order to maintain the same field of view, the CCD would need to be proportionally larger.

5.2.3 Multiple wavelength data

With several measurements shown above, as with the final measurement set-up, the fringe density is so high that standard phase-unwrapping techniques cannot be applied. This is the reason for introducing a second measurement wavelength with which the same measurement is performed. Subtracting the two phase profiles obtained this way, gives rise to a “synthetic phase” profile (see chapter 4), which pertains to a “synthetic wavelength”, given by (4.9), much longer than either measurement wavelength. The synthetic phase profile has fewer phase-jumps than either of the measured phase profiles, and can be unwrapped with the algorithm mentioned in the following subsection. Ideally, this unwrapped synthetic phase profile could be turned into a rough absolute OPD measurement with an accuracy better than 600nm, thus allowing us to calculate the correct number of 2π phase-shifts for every pixel of the phase-profile from our primary measurement wavelength. This then gives the absolute OPD with the measurement accuracy for our primary measurement wavelength.

The interferometer and light source are indeed designed to allow this approach, but the limited accuracy of the sensor requires us to proceed differently. As mentioned previously, the OPD function is predominantly linear – the deviation from a best-fit plane being only a few wavelengths over the CCD area. Consequently, we can perform the following steps:

1. The measured phase profiles are turned into the synthetic phase profile.
2. The synthetic profile is un-wrapped.
3. The best-fitting plane to this phase profile is determined.
4. This best-fit plane is scaled by a factor of $\Lambda_{1,2}/\lambda_1$ to give an estimate of the best-fitting plane to the λ_1 phase profile.
5. The scaled best-fit plane is subtracted, modulo 2π, from the λ_1 phase profile, giving a profile with a reduced fringe density.
6. This profile is un-wrapped, and the best fit plane previously subtracted is added again.

The resulting absolute phase profile differs from the originally measured λ_1 phase profile only by an integer number of 2π steps at each pixel and can be turned into an OPD profile by scaling it by $\lambda_1/2\pi$. This approach has been used to un-wrap all of the measurements shown in this section.

By fitting a plane to the synthetic phase profile, we average out all stochastic-, and several systematic errors over the area of the CCD – in essence this is a noise reduction technique. For very severe aberrations, we would be at liberty to fit a more complicated shape instead of a plane, and proceed as above.

5.2.4 Unwrapping routine

Due to the presence of “bad” pixels in all but the high-end cameras, our 2D unwrapping routine should exhibit good immunity to noise. Furthermore, the un-wrapping routine should be able to deal with “masked regions” – a fringe pattern where certain regions are excluded from the unwrapping process, either because of low modulation depth or because of diffraction effects (e.g. just inside our central aperture). We will now describe a fast and simple hybrid method exhibiting these properties.

To deal with possible noise, we will first temporarily "smooth" our interference pattern. As we shall see later, this smoothing effect does not carry over to our final answer.

Instead of naively smoothing the phase map per se, which would also smooth the critical discontinuities at 2π phase jumps, we will first decompose our phase map into its sine and cosine components, before smoothing each one and re-assembling our phase map via the arctangent function. The result is a smoothing effect which leaves our discontinuities intact (see Figure 5-21 & Figure 5-22). The degree of smoothing required depends on the spatial frequency of our features of interest and the amount of noise. The main purpose of this smoothing is merely to prevent the false identification of 2π phase jumps. We are not trying to reduce the noise in our final measurement.

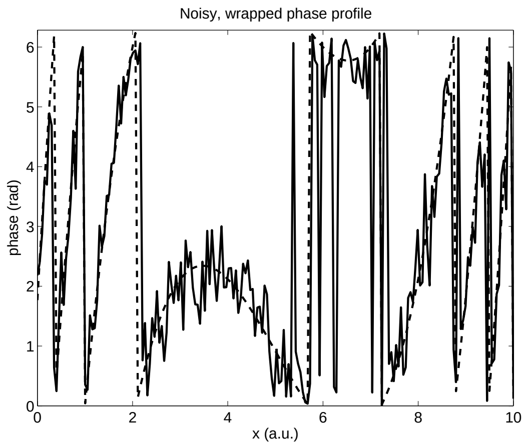

Figure 5-21 – Exaggerated example of a wrapped, noisy phase profile (the noise-free phase profile is represented by the dashed line)

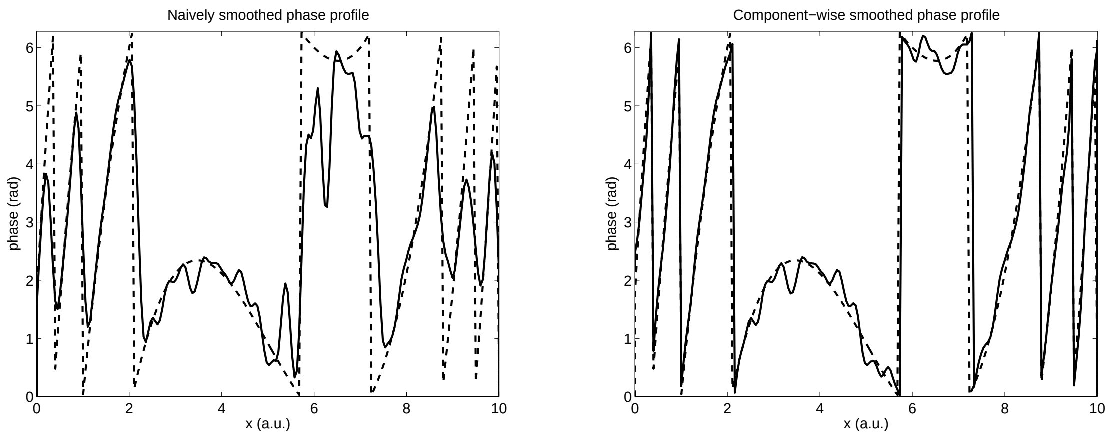

Figure 5-22 – Comparison between naïve and component-wise smoothing, using the same smoothing kernel. (Dashed line – noise free phase profile)

We now take the difference between our smoothed and original phase maps (modulo 2π), to be added again after unwrapping the smoothed phase profile.

The unwrapping used here consists of a sequence of four 1D unwrapping steps: from left to right, from top to bottom, from right to left and from bottom to top. Each of these unwrapping steps deals with the presence of masked regions by extending the masked regions in the direction of unwrapping, as if casting a shadow (see Figure 5-23). It is important to note that there are some choices of masks which will not permit a complete unwrapping of the un-masked regions. For our purposes however, this approach is sufficient.

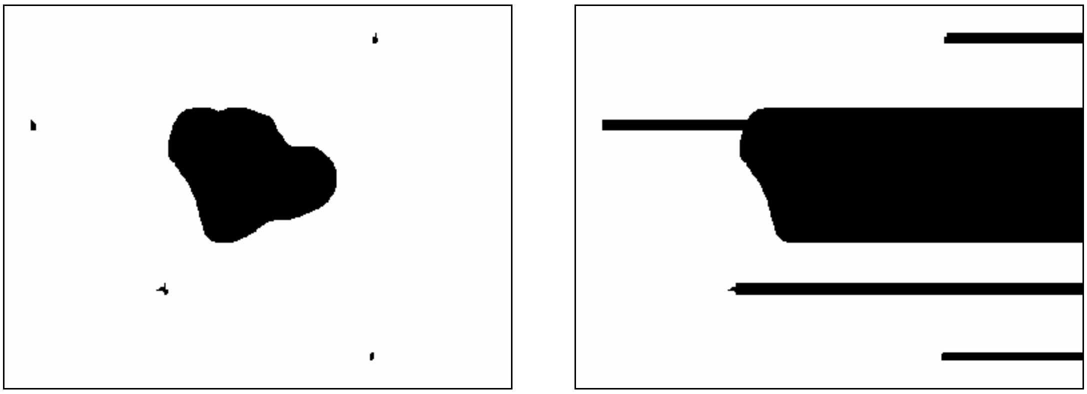

Figure 5-23 – Example of a masked region (left) and the extended mask used for left-to-right unwrapping (right). The black regions are not involved in the un-wrapping procedure.

5.2.5 Summary

Figure 5-24 summarizes the methods outlined in this section to interpret measurements taken with a CCD-type sensor to retrieve the absolute OPD function produced by our interferometer. The steps show how the intensity measurements are converted to phase measurements, how these measurements are corrected for systematic errors and how the under-sampling of the phase function can be overcome by using the phase data from the other wavelength of our multiple wavelength light source.

Performing these steps on data from the test-setup, allows us to estimate the accuracy with which OPD measurements can be made using our CCD. With similar parameters as for the actual measurement set-up, the OPD function for two fiber-tips was measured, and the theoretical OPD subtracted to obtain an error-image.

The stochastic noise in the error image can be reduced significantly by spatial averaging. Since the CCD provides many more data-points than required for our spatial resolution of the test mirror, such averaging is permissible. An increased bit-depth (>8 bits) could further reduce the stochastic noise observed. The remaining error image has an RMS value of 29nm – due primarily to CCD geometry errors. These errors can be corrected using calibration procedures in the absence of a cover-glass. Unfortunately, the presence of the cover-glass makes such a calibration impossible, and will itself also contribute a small amount of error directly.

Newer CCDs have both higher bit-depths and better feature positioning accuracies than our CCD. If one of sufficient size can be obtained without the cover-glass, it is likely that our measurements could be performed to the required accuracy using the methods outlined in this chapter.

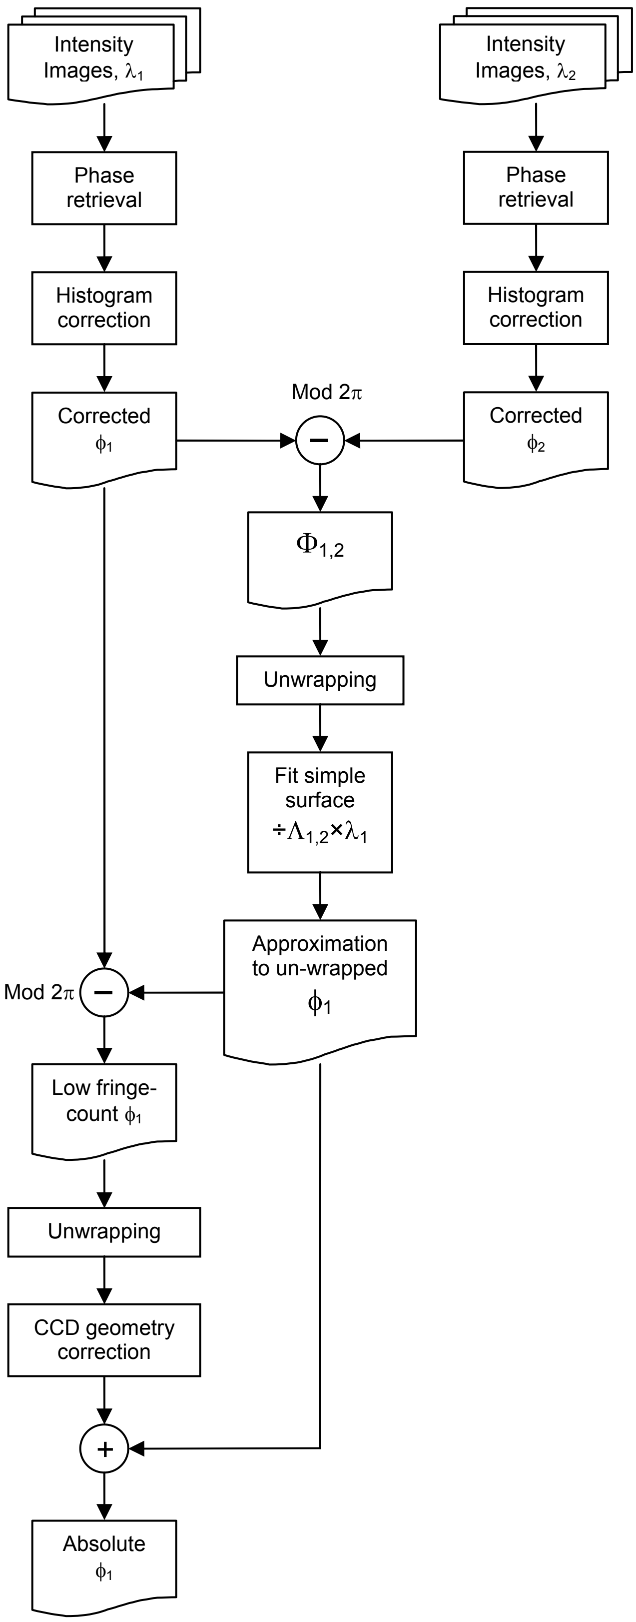

Figure 5-24 – Flowchart showing the procedure to convert a series of PSI intensity images into an absolute phase profile, with a number of error-reducing steps

5.3 Fast-sampling type sensor

A fast-sampling type sensor is essentially the same as a CCD sensor, but with a much faster frame-rate than the video frame rate of 25fps. The primary advantage of a fast-sampling type sensor is the possibility of measuring the interferogram of both wavelengths simultaneously, together with the ability to use heterodyne interferometry rather than phase-shifting interferometry. For the latter, the sensor must also have a full-frame shutter – that is, all pixels must sample the light intensity during the same period of time over the whole CCD.

Such sensors are available with frame-rates exceeding 100kHz, but this high frame rate usually comes at the cost of resolution and bit-depth. The minimum size restriction, the maximum pixel dimensions and requirements on positioning accuracies all remain the same as those for the standard CCD, and the problems of a cover-glass persist also.

Given such a sensor, with a frame rate *F*, and the ability to store *N* consecutive images, we could choose the following heterodyne frequencies for our two wavelengths:

$$\begin{aligned} f_1 &= F\left(\frac{1}{2}-\frac{1}{N}\right) \\ f_2 &= F\left(\frac{1}{2}-\frac{3}{N}\right) \\ N &\in 8,10,12... \end{aligned} \tag{5.13}$$

The first frequency is chosen to be the highest resolvable frequency by this method, since less noise is generally present at higher frequencies. The second heterodyne frequency is chosen in such a way as to find a balance between a reduction in cross-talk with the first heterodyne frequency (requiring a large frequency difference), and the desire to have the frequency as high as possible out of noise considerations (requiring a small frequency difference). Of course, both frequencies should have an integer number of oscillations during the measurement time. If significant amounts of phase-noise are found, these frequencies should be spaced more widely, to prevent cross-talk.

A sequence of measured images can then be analysed by using the WDFT method for each pixel along the “time” dimension, just as in the phase-shifting case, with the difference that we would pick out the phase of two frequency components, one for each wavelength, rather than just one.

Since no such sensor is available, we will show merely a simulation of how a sequence of intensity values captured with such a camera would be turned into the two separate phase values for the two wavelengths. Figure 5-25 shows a simulation of 36 successively acquired intensity values for one pixel and the corresponding discrete Fourier transform. The argument of the complex values corresponding to our heterodyne frequencies will give us the phase of that signal. The statistical noise in such a case is averaged by the process, and reduced by a factor proportional to $\sqrt{N}$, but phase-drift and other systematic errors on the CCD, such as non-uniform shutter-times, geometry errors etc. will doubtlessly increase this error further.

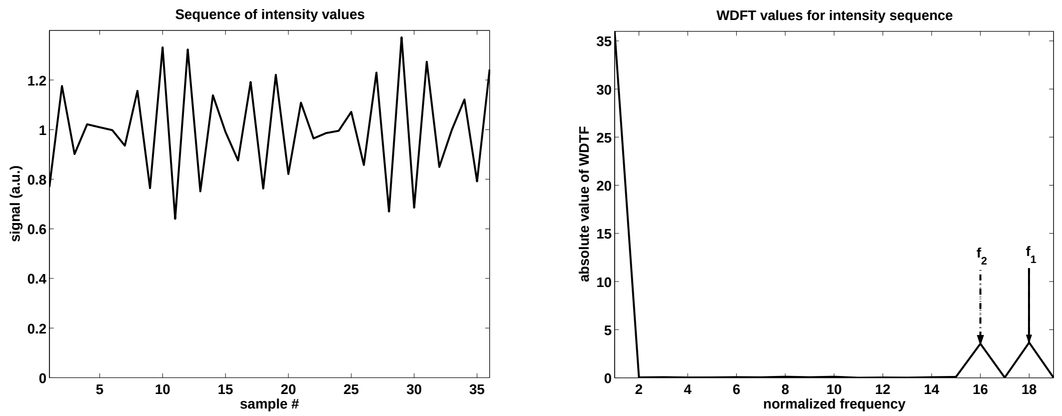

Figure 5-25 – Simulated sequence of intensity values and corresponding WDFT values for a fast sampling type sensor.

Once the phases have been obtained for every pixel of the sensor, we would proceed as with the phase signals obtained by a CCD to retrieve the total OPD.

While such a sensor offers decreased sensitivity to the low-frequency noise present in the setup, due to the higher measurement frequency and the simultaneous nature of the two-wavelength measurement, there are also considerable disadvantages to its use, not the least of which is the prohibitive cost of such a sensor. Another disadvantage is the large amount of data that must be transferred in order to obtain the phase values for each pixel. This makes this option also much more demanding in terms of post-processing hardware. Furthermore, read noise has been known to increase with higher clock frequencies, thus decreasing our SNR.

As such sensors become more wide-spread, they may well become a viable option for performing our phase measurements, taking full advantage of the abilities of our light-source.

5.4 Phase-measuring pixel sensor

Although originally developed for time-of-flight measurements, the phase-measuring pixel sensors developed by Lange and Seitz[8, 9] offer 2D, high-bandwidth, heterodyne detection without the drawbacks of large amounts of data transport or processing. The elegantly simple working principle of this class of sensors is based on synchronous detection.

The importance of such a device is highlighted by the following quote by Dorrio & Fernandez[3]:

"…The fact that the [heterodyne] measurement is carried out point-by-point, has resulted in the fact that this method is practically nonexistent in industry, and therefore is relegated to being a precise laboratory evaluation method"

Each phase-measuring pixel has, in addition to a photosensitive area, *N* charge-storage "bins" around it. A signal, synchronized with the heterodyne frequency, switches the charge produced successively from bins 1 to *N* during one period of the heterodyne signal. This process is repeated for many periods, until one of the bins reaches its capacity. At this point, the charges are read out, and can be processed by an "*N*-step" algorithm, like equation (5.2), into a phase value for each pixel (Figure 5-26 shows this process for the case of four bins). For five or more bins, it is even possible to simultaneously measure two phases from different, well-chosen, heterodyne frequencies.

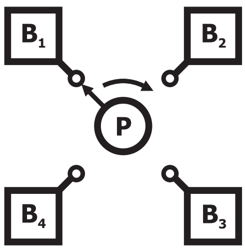

Figure 5-26 – Diagram showing how charge produced at a phase-measuring pixel is successively shunted to different charge storage bins.

Such sensors have recently come into production. Up until now, sensors with 5 or more bins have only been built as linear arrays, but a 2D array of 4-bin sensors has been built by CSEM, Switzerland[8].

This sensor is not yet commercially available, but a 2-bin predecessor, working on a slightly modified principle, designed for time-of-flight measurements, has been purchased to evaluate its potential to act as sensor in our interferometer.

5.4.1 Working principle

The Swiss Ranger 2 from CSEM, Switzerland, is such a phase-measuring pixel camera, with which the phase of amplitude modulated light with frequencies between 80kHz and 20MHz can be measured at every pixel by means of a two-bin sampling scheme outlined in the thesis by Lange [8].

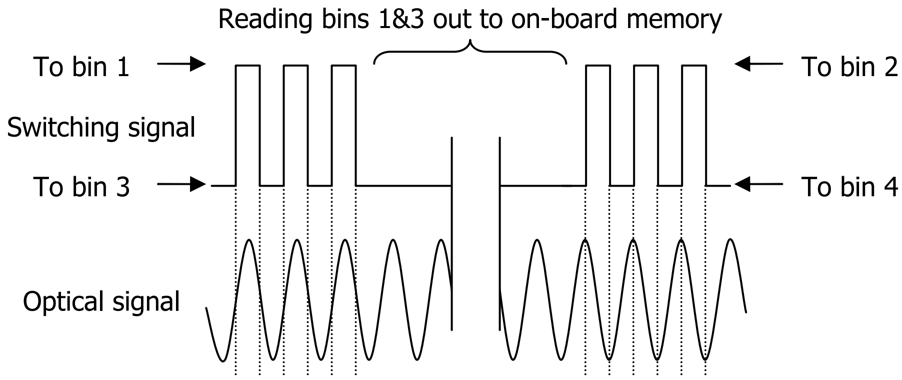

Figure 5-27 – Schematic diagram of the two-bucket sampling scheme of the SR2 camera.

Every pixel has only two charge storage bins to which the charge produced by the photo-sensitive area of the pixel is alternatively shunted, synchronous to a switching signal with an identical frequency to the modulation signal to be measured (see Figure 5-27). This process is carried out for a few milliseconds, and the charges stored in the bins are read out to an on-board memory. The process is then repeated again, but with a switching signal shifted by 90° with respect to the previous switching signal. After reading out the bins a second time, four samples per pixel are available, corresponding to 0°, 90°, 180° and 270° relative to the original switching signal.

In the absence of any phase shifts during the integration time, the phase of the amplitude modulation can be calculated according to the well-known formula:

$$\phi = \tan^{-1}\left(\frac{b_4 - b_2}{b_1 - b_3}\right) \tag{5.14}$$

Where b_1, b_2, b_3 and b_4 stand for the charges stored in bins 1-4.

The test setup in Figure 5-28 was used to evaluate the camera. While this setup allows the use of two wavelengths simultaneously, all measurements, except for the ones in Figure 5-36, use only one wavelength at a time.

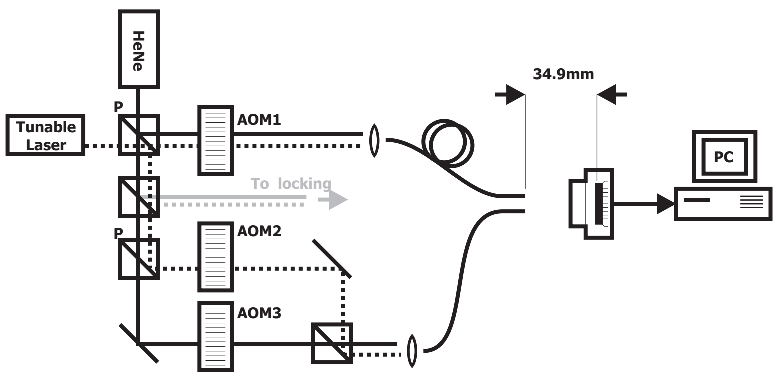

Figure 5-28 – Test setup to evaluate CSEM phase measuring pixel sensor.

A typical set of test measurement data is shown in Figure 5-29, together with the resulting raw- and unwrapped phase.

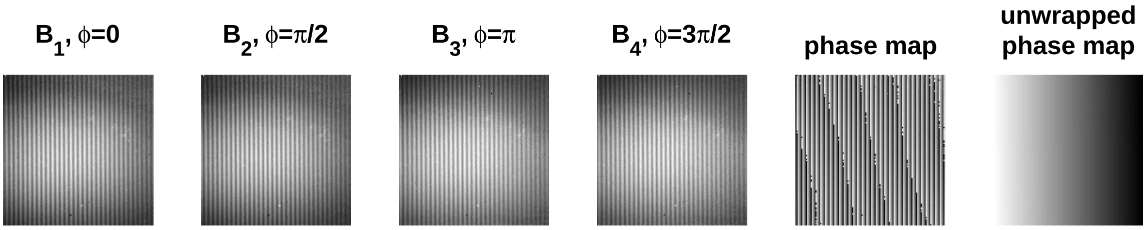

Figure 5-29 – Four images from one CSEM sensor measurement, resulting phase, and unwrapped phase.

The four images are returned as one data set via a USB interface before a new reading can be taken.

5.4.2 Error sources

Test measurements carried out under identical conditions with the Sony CCD using PSI, and with the CSEM sensor using heterodyne interferometry, show that the CCD sensor significantly outperforms the CSEM sensor, despite the inherent advantages of heterodyne interferometry.

Phase Drifts

Because the 0 and 180 samples are measured at a different time to the 90 and 270 samples, phase drifts during the measurement introduce considerable errors. The expressions for the charges in bins 1-4, for an arbitrary optical intensity function *o(t)* are:

$$\begin{aligned} b_1 &= \sum_{j=0}^{N-1} \int_{j/f}^{(j+\frac{1}{2})/f} o(t)dt & b_3 &= \sum_{j=0}^{N-1} \int_{(j+\frac{1}{2})/f}^{(j+1\frac{1}{2})/f} o(t)dt \\ b_2 &= \sum_{j=N+M}^{2N+M-1} \int_{(j+\frac{1}{4})/f}^{(j+\frac{3}{4})/f} o(t)dt & b_4 &= \sum_{j=N+M}^{2N+M-1} \int_{(j+\frac{3}{4})/f}^{(j+1\frac{1}{4})/f} o(t)dt \end{aligned} \tag{5.15}$$

Where *f* is the frequency of the switching signal, and we assume an integration time of *N* periods for each bin pair, with a pause of *M* periods for the reading out of bins 1&3. It can easily be verified that substituting $o(t) = \sin\left(2\pi f \times t + \phi\right)$ into (5.15), and applying (5.14) to the result gives the correct phase. For a linear phase drift with time, we can write:

$$o(t) = \sin\left(2\pi\left[f + \varepsilon\right] \times t + \phi - \pi\left[2N + M + \tfrac{1}{4}\right]\varepsilon / f\right) \tag{5.16}$$

Where the last term is there to set the “average” phase due to the phase drift to zero. Substituting this into (5.15) and applying (5.14) to the result allows us to calculate the resulting phase error for some realistic camera parameters. The default parameters will be taken as 80kHz for the modulation frequency, 4ms integration time, and 0.1 rad/s phase drift. The time taken to read out the bins is always around 2.735ms, but rounded to the nearest integer and ¼ periods of the switching signal.

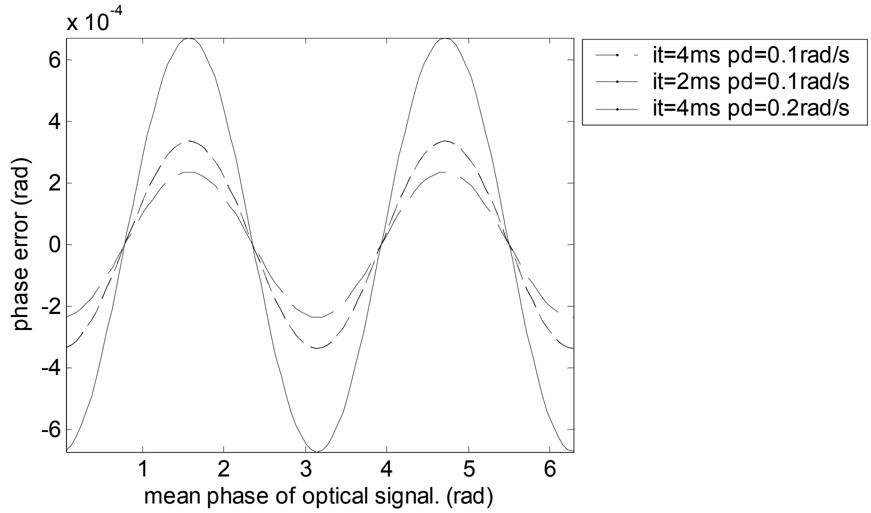

Figure 5-30 – Phase error in SR2 camera for different combinations of integration time (it) and phase drift (pd), at 80kHz optical modulation.

From Figure 5-30 it should be clear that the over-all shape of the phase-error curve does not change with any of our parameters, but the amplitude is strongly affected by both the integration time and the phase drift. The modulation frequency has virtually no influence on the phase error.

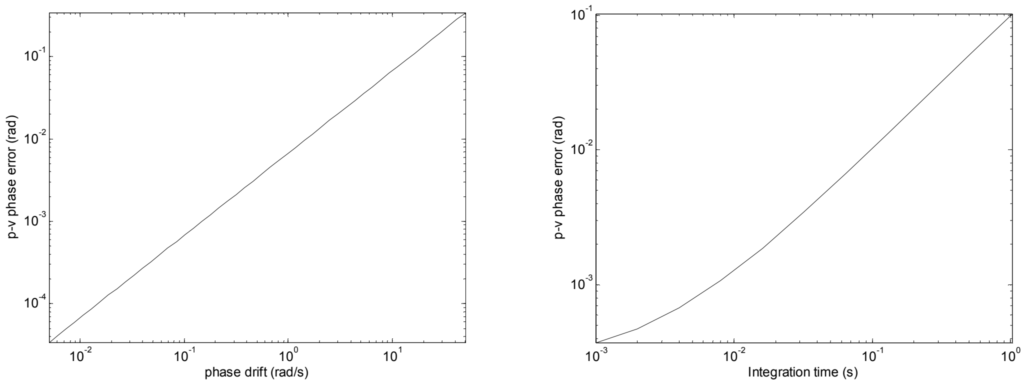

Figure 5-31 – p-v phase error due as function of phase drift and integration time respectively, for 80kHz modulation of SR2 camera.

For the default parameters, we can therefore expect a peak-to-valley (p-v) phase error of 0.6 mrad. Longer integration times, and of course larger phase drifts, cause bigger p-v phase errors as is shown in Figure 5-31.

A proper 4-bin (or more) sensor would exhibit a much lower sensitivity to phase drifts, because all bins are effectively collected at the same time. In that case, only the phase-drift during one heterodyne period contributes to the error, while the 2-bucket scheme is influenced by the drift over the total measurement time.

Nonetheless, since this is a phase-dependent error, the histogram correction outlined in section 5.2.2 can be used to reduce the effect of this error.

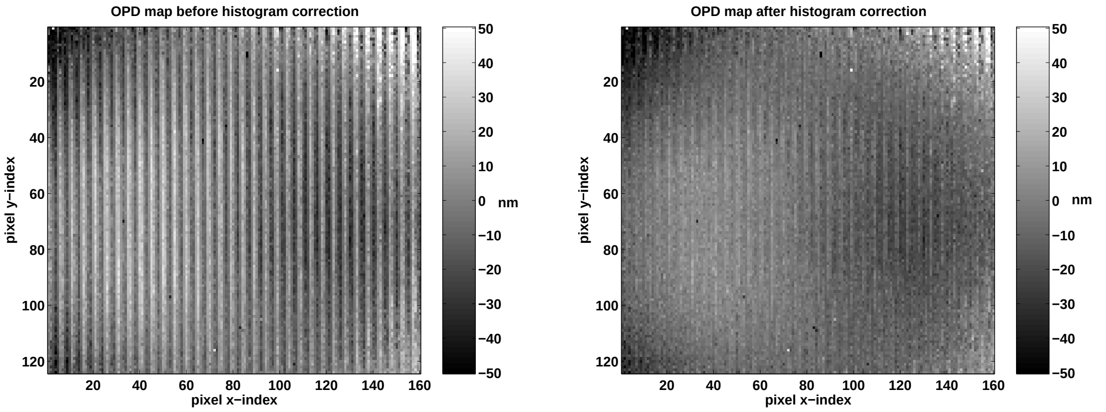

Figure 5-32 – Comparison between an un-corrected (left) and a corrected interferogram (right - tilt has been removed in both cases)

Figure 5-32 shows the OPD map (minus tilt) obtained from the data in Figure 5-29, and serves to illustrate the extent with which artefacts on an interferogram can be removed with this method. It is important to stress once again that this is not a spatial filtering technique, and will work for more complicated fringe patterns.

Despite a significant reduction in the phase-drift induced errors, they remain the dominant source of measurement inaccuracy. By using a 2-tap, rather than a 4-tap approach, the measurement accuracy of this sensor is sacrificed for the sake of ease of manufacture.

Pixel positioning- and size-errors
Another major drawback of the CSEM sensor is the fact that it has larger pixels than the CCD described in section 5.2. The manufacturer's specifications for the CSEM sensor are given in Table 5-1.

Number of pixels	160 x 128
Pixel pitch	39.2μm x 54.8μm
Fill factor	16.3%
Chip size	8.3mm x 7.6mm

Table 5-1 – Manufacturer's specifications for the CSEM phase-measuring pixel sensor

From the pixel pitch and fill-factor, we can estimate the pixel size to be 15.7μm x 21.9μm if scaled proportionally in both orthogonal directions. As calculated previously, the fringe period for our measurement set-up will be approximately 17.7μm – only marginally bigger than the horizontal pixel size. This would cause a reduction of the measured modulation depth to 12.5% of the actual value due to the MTF of the sensor (see (5.9)), causing a significant reduction of signal-to-noise ratio. Ironically, since the bit-depth of the sensor is 10-bits rather than our CCD's 8-bits, the theoretical SNR of both sensors would then be about the same (the MTF of the CCD is about 4 times better than that of the CSEM camera, but the quantization noise in the CCD is 4 times worse than that in the CSEM sensor).

Several test measurements with the set-up shown in Figure 5-28 give us reason to believe that the pixel size is unfortunately bigger than the calculated values above. It was observed that the measured modulation depth already drops to zero with a fringe-spacing of 30μm (see Figure 5-33) – meaning that this is likely to be the actual effective horizontal pixel size. This is unfortunate, since we do not have a lot of leeway to adjust our fringe density. Placing the sensor further away is an option, but this then also reduces our field of view.

The noise-floor of the measurements was too high to be able to distinguish effects due to positioning errors.

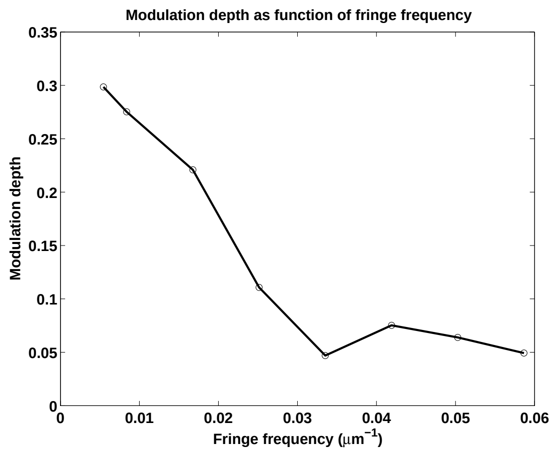

Figure 5-33 – Modulation depth as function of fringe density, tracing out the expected |sinc|-profile of the MTF. Zero is at a fringe spacing of ~30μm.

Electronic noise
Further bad news from the test measurements is the level of noise observed. While the quantization noise of most sensors is chosen to lie just above the electronic noise, this seems not to be the case for the CSEM sensor. Our measurements indicate a 0.3% level of noise just below saturation, compared to the 0.1% to be expected from the bit-depth of the sensor. This may be due to the fact that the camera is being operated at frequencies of 80kHz-100kHz, rather than the design frequency of 20MHz. A detailed analysis of the electronic noise sources of this CCD falls outside the scope of this thesis, but a more detailed treatment can be found in the thesis of Lange[8].

Wavefront distortion and multiple reflections
We were fortunate enough to be able to receive a sensor which was taken off the assembly line before the cover-glass was attached. The issues of multiple reflections and wavefront distortions therefore do not come into play here.

Intensity response
Because this is a fairly novel sensor, we have decided to test the intensity response of the detector with the help of the set-up shown in Figure 5-34. In an effort to obtain a uniform intensity distribution, and LED was used to avoid speckles, and a ground-glass plate further diffused the resulting illumination. Measurements for a typical pixel are shown in Figure 5-35.

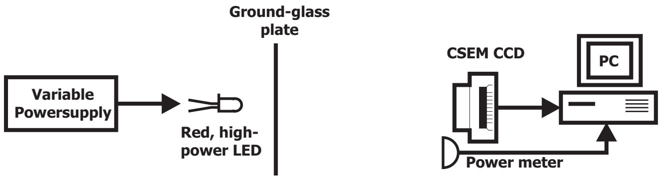

Figure 5-34 – Intensity response calibration setup for CSEM CCD.

While the response itself is fairly linear until saturation is approached, the difference in response between the two sets of charge storage bins of a particular pixel is of concern. We have therefore measured the ratio of the response for the two bins of every pixel, and applied a correction factor to the data from the second and fourth bin of every pixel in an effort to reduce this error.

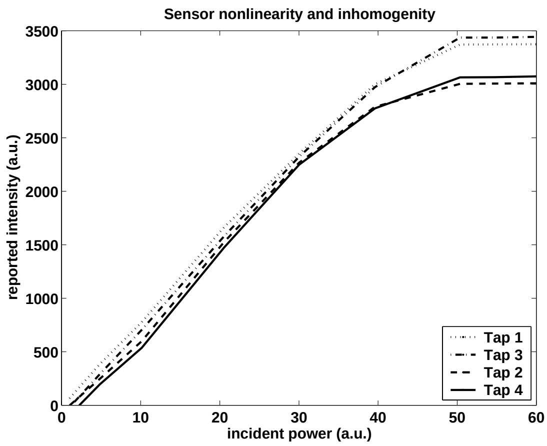

Figure 5-35 – Intensity response of a typical pixel, showing nonlinearity and different sensitivities between taps.

Also, the saturation level should be approximately the same for all pixels. The fact that this is not the case is evidence of irregularities in the CCD structure.

5.4.3 Beat frequency demultiplexing.

A unique feature of this camera is its ability to selectively measure the phase of one particular wavelength when more than one wavelength from the multiple wavelength source described in chapter 4 falls on the sensor at the same time.

This separation, or de-multiplexing, of the wavelengths is achieved by matching the sensor's de-modulation frequency to the heterodyne frequency of that particular wavelength.

The use of two (or more) wavelengths simultaneously has the disadvantage of reducing the modulation depth of each signal, because the light from the other wavelength(s) form a "background" intensity level on the modulated signal. Under ideal conditions, with N wavelengths, the modulation depth will be limited to 1/N for all wavelengths.

Figure 5-36 shows such a measurement using two wavelengths simultaneously in the test-setup shown in Figure 5-28. The best-fit plane to the un-wrapped phase profile for ϕ_1 was determined with the approach outlined in section 5.2.3. The heterodyne frequencies of the two wavelengths were 100kHz and 80kHz respectively. As far as we are aware, this is the first time that a 2D heterodyne synthetic wavelength measurement has been performed with such a sensor.

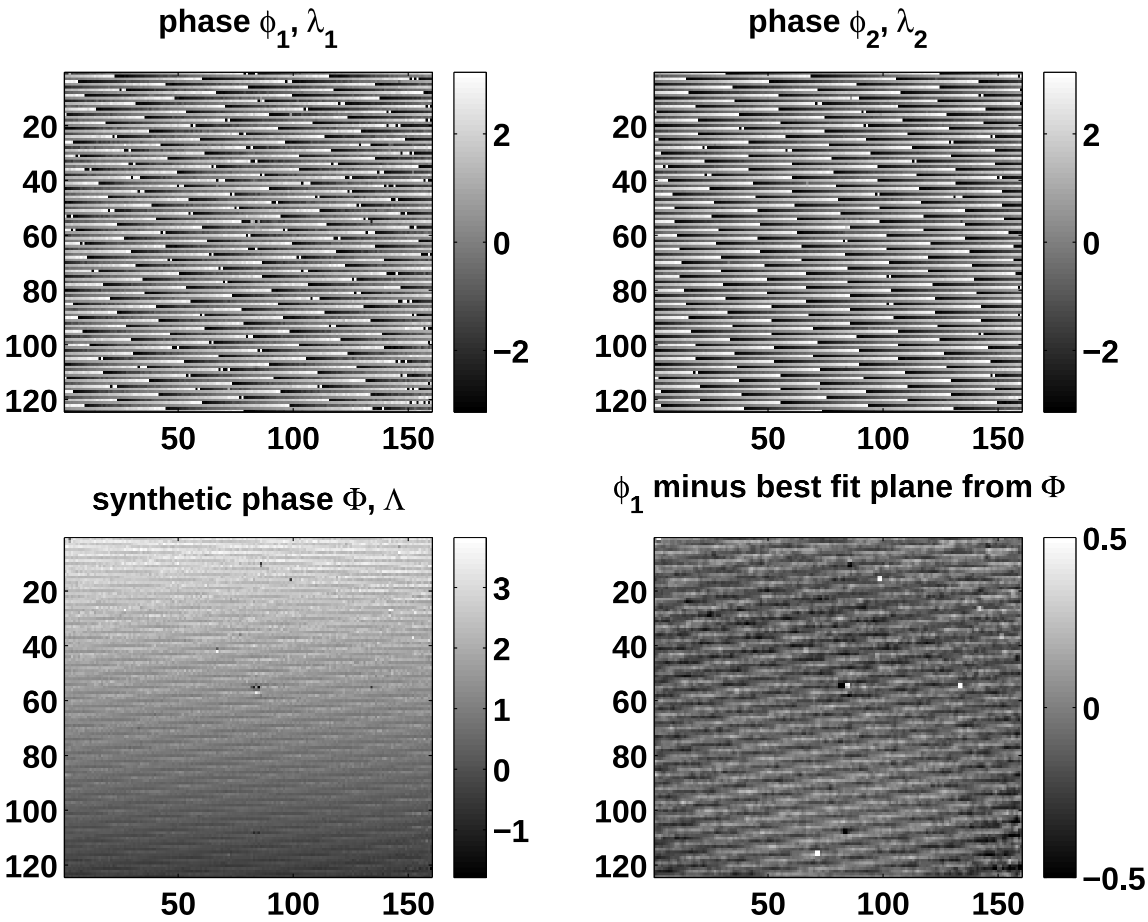

Figure 5-36 – Measurement of separate phases for λ_1 and λ_2, while both wavelengths are illuminating the sensor simultaneously. Resulting synthetic phase and its application to remove phase ambiguity from ϕ_1. Phases are measured in radians.

5.5 Conclusion

Although a sensor meeting our initial requirements was unavailable at the time of writing, significant progress has been made, with a view to achieving the desired measurement accuracy with commercially available sensors.

The best results were achieved with a standard video frame-rate CCD using the interferometer in phase shifting mode. An rms accuracy of 29nm can be achieved this way – limited primarily by CCD geometry errors. The accuracy of the method after correction for the CCD geometry errors is estimated at 6nm, with a further reduction to 1nm with spatial averaging of neighbouring pixels. The presence of a cover glass makes the correction of CCD geometry errors impractical. Currently available CCD sensors without a cover glass should be able to meet our requirements.

A novel phase-measuring pixel sensor was also tested as a candidate for heterodyne measurements, with initially discouraging results. To the best of our knowledge, the principle of 2D heterodyne multiple wavelength interferometry has been demonstrated for the first time with such a sensor, but this approach has not yielded the necessary measurement accuracy. Due to the sensor's 2-tap approach, it is particularly sensitive to phase-shift errors and errors due to phase-drift. The observed nonlinear response causes the errors to be both amplitude and phase dependent, preventing a-posteriori correction of the data. The large difference in responses between adjacent pixels is also cause for concern. Improvements on such sensors in future, especially the use of a 4-tap approach and improved linearity, may well make them the most promising candidates for this form of interferometry due to the flexibility and efficiency inherent in these sensors.

5.6 References

[1] R. G. Klaver, "Novel interferometer to measure the figure of strongly aspherical mirrors." Delft: Delft University of Technology, 2001.

[2] D. Malacara, *Optical shop testing*, 2nd ed: Wiley-Interscience, 1992.

[3] B. V. Dorrio and J. L. Fernandez, "Phase-evaluation methods in whole-field optical measurement techniques," *Measurement Science & Technology*, vol. 10, pp. R33-R55, 1999.

[4] K. Creath, "Choosing a phase-measurement algorithm for measurement of coated LIGO optics," in *Laser Interferometry X: Techniques and Analysis and Applications, Pts a and B*, vol. 4101, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2000, pp. 47-56.

[5] K. A. Goldberg and J. Bokor, "Fourier-transform method of phase-shift determination," *Applied Optics*, vol. 40, pp. 2886-2894, 2001.

[6] E. Novak, A. Chiayu, and J. C. Wyant, "Errors caused by nearly parallel optical elements in a laser Fizeau interferometer utilizing strictly coherent imaging," presented at Optical Manufacturing and Testing II. 27-29 July 1997 San Diego, CA, USA [SPIE], 1997.

[7] G. E. Sommargren, D. W. Phillion, M. A. Johnson, N. Q. Nguyen, A. Barty, F. J. Snell, D. R. Dillon, and L. S. Bradsher, "100-picometer interferometry for EUVL," in *Emerging Lithographic Technologies VI, Pts 1 and 2*, vol. 4688, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2002, pp. 316-328.

[8] R. Lange, "3D Time-of-Floght Distance Measurement with Custom Solid-State Image Sensor in CMOS/CCD-Technology," in *Department of Electrical Engineering and Computer Science*. Siegen: University of Siegen, 2000.

[9] R. Lange and P. Seitz, "Solid-state time-of-flight range camera," *IEEE Journal of Quantum Electronics*, vol. 37, pp. 390-397, 2001.

6 Interferometer

The design of the interferometer used to obtain the results in this thesis was aimed at one specific test mirror available. While future designs should focus on greater flexibility, in this design, stability was made the top priority. We will begin by giving details of the test-mirror used, before proceeding to outline the specifications and structural design of the interferometer (technical drawings can be found in appendix A.9). This chapter will also include a section on the sphericity of the wavefronts generated by the optical fibers.

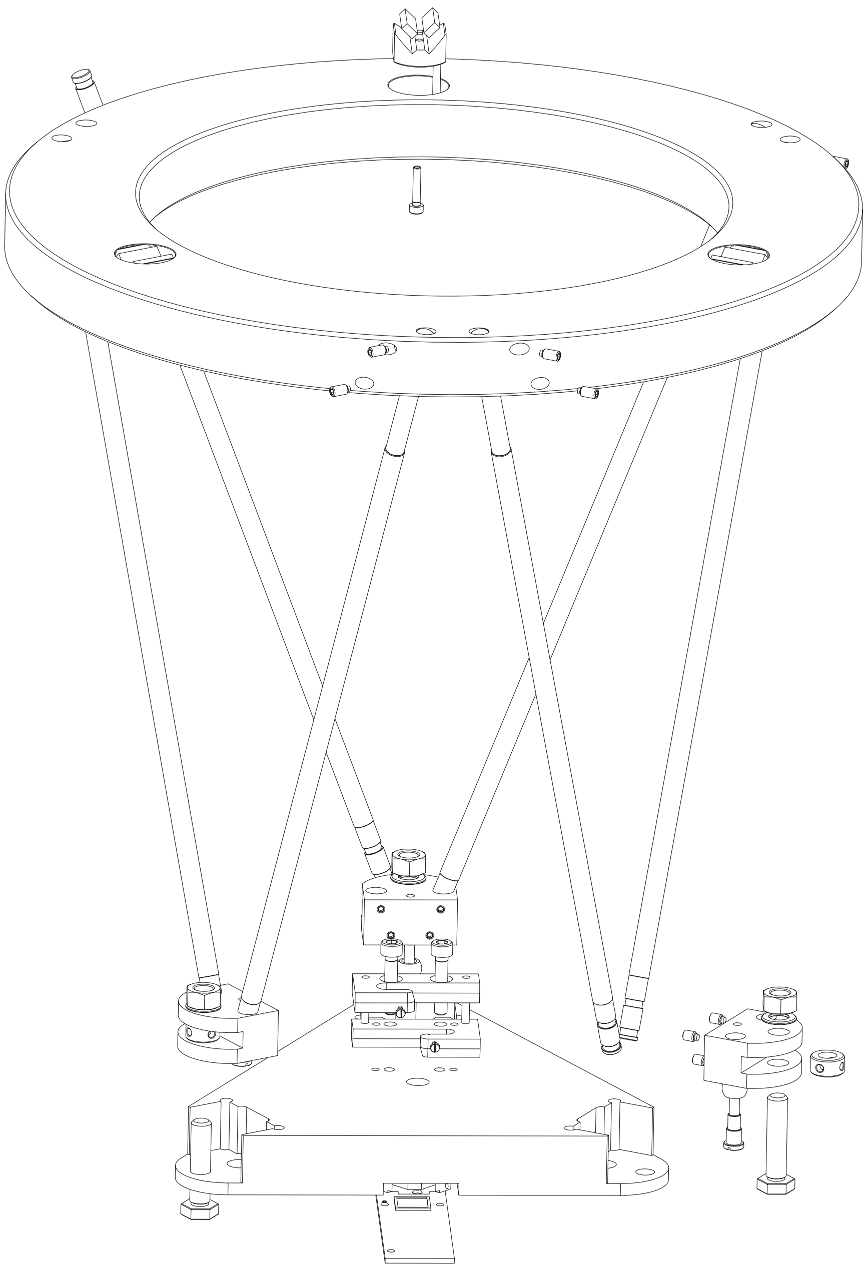

Figure 6-1 – Exploded view of interferometer frame.

6.1 Test Object

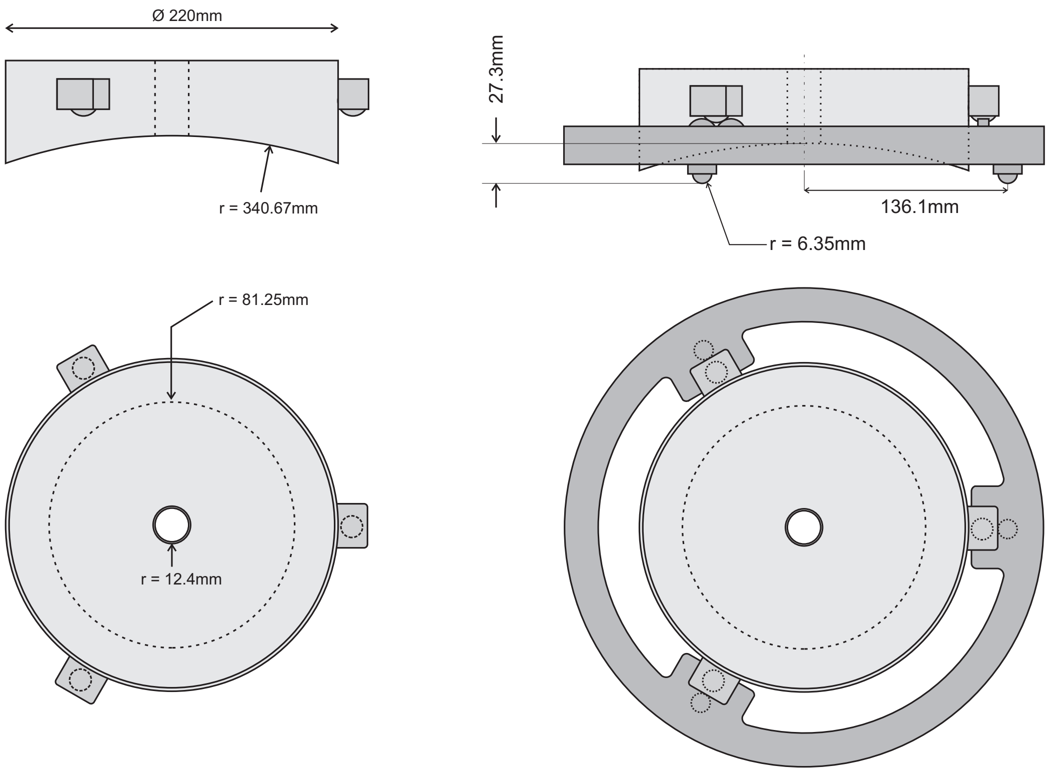

Figure 6-2 – Test mirror and mounting structure

The object under test is an aspherical mirror substrate, made entirely of Zerodur (See Figure 6-2). The aperture is circular, centred on the rotational axis of symmetry. The radius of the reflective surface of interest is 81.25mm, and has a hole at the centre with a diameter of 24.8mm. The nominal radius of curvature is 340.67mm (equivalent to a nominal focal length of 170.34mm; f=R/2), giving a "standard" NA for the mirror of 0.48. The rim beyond the region of interest is not polished, having the appearance of ground/frosted glass. Ideally, the substrate would be rotationally symmetric, and would have a P-V vertical deviation from the best-fit sphere of 4.6μm as shown in Figure 6-3.

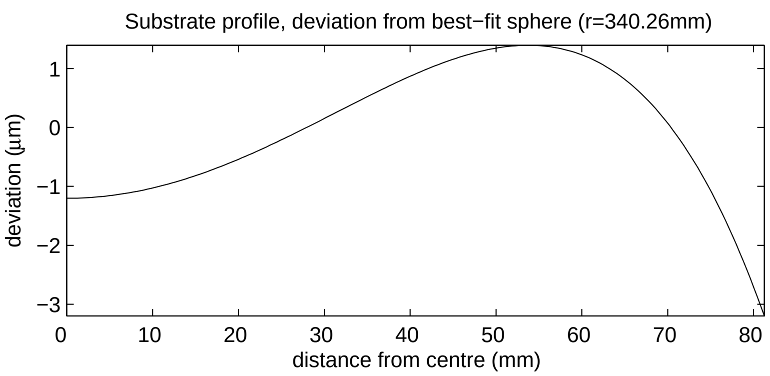

Figure 6-3 – Vertical deviation of mirror from best-fit sphere.

The substrate was delivered with a mounting structure to facilitate accurate placement in measurement systems. The mounting structure essentially consists of a metal ring, with three disc-grooves to receive the positioning spheres on the mirror substrate, designed to minimize frictional placement errors. These mounting structures are proprietary, and we will have to restrict ourselves to the essential dimensions of the combination of mounting structure and substrate.

The three positioning spheres provided on the mounting structure, to be used to interface with the measurement setup, are equally spaced by 120˚ at 136.1mm from the centre. The intention is to leave the mounting structure in place while the mirror substrate is being removed and replaced.

6.2 Interferometer frame

The classic hexapod structure of the interferometer frame was suggested by designers at TNO/TPD, to guarantee stiffness and stability without inducing excess stresses. The design and manufacture of the instrument was carried out in close collaboration with the drafting bureau (VeCaTek fijnmechanica) and the manufacturer (Instrumek B.V., Schiedam), resulting in the blueprint shown in Figure 6-1. The top ring serves as an interface with the mounting structure of the test mirror, while the triangular base plate accommodates both the fibers and the sensor.

In accordance with the interferometer concept outlined in chapter 2, the various interferometer components should be positioned as shown in Figure 2-1. The exact distances depend on a number of factors, and are shown in

Figure 6-8.

6.2.1 Fiber positioning

The height of the mirror apex above the fiber-tips should be equal to the nominal radius of curvature of the mirror, i.e. 340.67mm, to match the wavefront curvature optimally to the mirror curvature, hence minimizing our fringe-density.

The fiber positions are determined by balancing the desire for a low fringe-density, and the need for the reflected beam to pass between the two fibers unobstructed. This means that we need to find the minimum fiber separation which allows the reflected beam to pass between them without being obstructed. The focus of the reflected beam will occur the same distance from the mirror axis, but on the opposite side, to the object fiber. We will start with a very much simplified model of the beam shape – that of a cone with the focus as apex, and the principal ray as axis – and derive an initial estimate of the required fiber separation from this. We will later use raytracing to check this estimate and if necessary, adjust it.

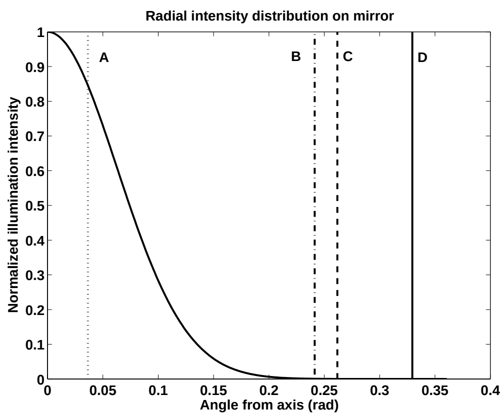

Figure 6-4 – Angular intensity distribution on mirror. A = edge of central hole, B = edge of region of interest, C = half opening angle of unobstructed cone, D = Edge of substrate.

Figure 6-4 shows the intensity distribution due to our fiber at the mirror, along with the various mirror features. At the edge of our region of interest, the illumination intensity is merely 0.065% of the peak intensity, dropping to 10^{-4}% at the edge of the substrate. The cone of unobstructed light should contain the region of interest, but extending the opening angle of this cone much further would require a bigger fiber separation than necessary. As a compromise, we have chosen the half opening angle of the cone which should remain unobstructed as 0.26 rad (approximately 15˚), compared to the 0.24 rad subtended by the edge of the region on interest. At the angle of 0.26 rad, the intensity has dropped to 10^{-4} of the peak intensity, and the ground-glass nature of the rim beyond the region of interest will, on average, reduce the amount of light scattered between our fibers much further.

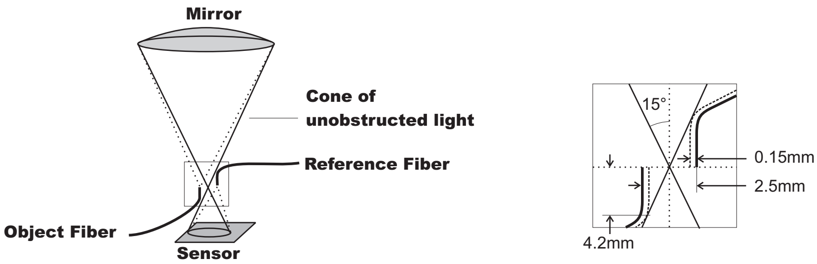

Figure 6-5 – Derivation of optimal fiber separation

Nonetheless, there is an obvious danger of part of the random speckle pattern from this rim passing between our fibers and interfering with the measurement beams. No precautions against this have yet been taken, but the placement of an additional (optically black) aperture to obscure the rim beyond the region of interest would solve this problem. Alternatively, the rim could be polished, but at such an angle that the reflected light does not pass between our fibers.

Figure 6-5 shows this cone of unobstructed light, together with the positioning of the fibers. Of several methods to fix fibers to metal components, the most promising was the use of an elastic band to press the fiber into a v-groove cut into the metal component (See Figure 6-6 and Figure 6-12).

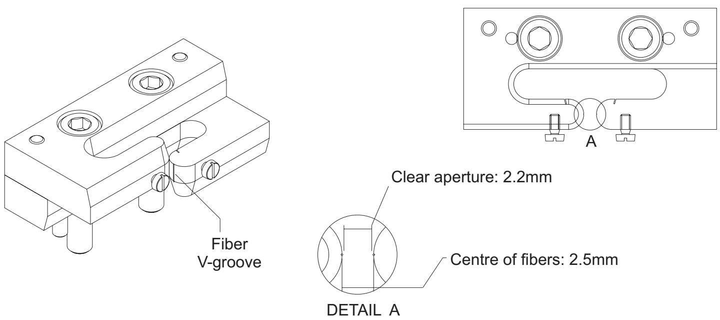

Figure 6-6 – Fiber holder design to mount object- and reference fibers.

The minimal height required to bend the fibers clear of the cone of unobstructed light using this method of affixing fibers was determined experimentally to be ~4.2mm. Taking the thickness of the elastic band and the fibers into consideration (~0.15mm), the minimal fiber separation was determined to be ~2.5mm.

From this arrangement, we see that the clear aperture volume near the focus is actually the combination of the aforementioned cone and a cylinder with radius

2.2mm (Figure 6-5). Provided no significant amount of light falls outside this clear aperture volume, our beam may be considered to be unobstructed. Using the raytracing programme mentioned in chapter 3, a set of rays covering the entire mirror area of interest was traced a few cm past the focus, confirming that no rays fell outside this clear aperture volume.

Provided that no improvements are made which allow the mounting of fibers over a shorter vertical distance, the minimal fiber distance is given by:

$$d_{fib} = 0.3 + 4.2 \times \mathrm{NA} \tag{6.1}$$

Where the NA is the standard numerical aperture of the reflective optic to be measured, not to be confused with the NA of our system, since our fiber is placed near the centre of curvature, rather than at the focal point of the reflector. The maximum NA which can reliably be measured in one go is limited to below 0.26 by the fiber wavefront sphericity (see section 6.3.1). Using fibers with smaller or tapered cores, this NA could possibly be extended further.

6.2.2 Sensor positioning

Ideally, we would want to position the sensor at such a distance from the focus, as to just fill the sensor area with the projection of the region of interest of the mirror. However, we are limited to a certain minimum distance, based on the fringe-density as compared to the pixel dimensions (see chapter 5.2.2). At the time the interferometer was being designed and manufactured, the primary sensor being considered was the phase-measuring sensor from CSEM (chapter 5.4). The sensor was not yet available for testing, and the design had to be based solely on the specification of the sensor.

Based on the specified pixel pitch (39.2 x 54.8 μm) and fill factor (16.3%), the pixel size was estimated at 15.8 x 22.1 μm. Using (5.11), we can see that the minimum focus-sensor distance (where the pixel size is equal to the fringe density) is 31.3 mm. At this distance, the modulation depth would be zero, and no measurements would be possible. In order to retain a 10% modulation depth (see (5.9)) this distance has to be increased to 34.9mm.

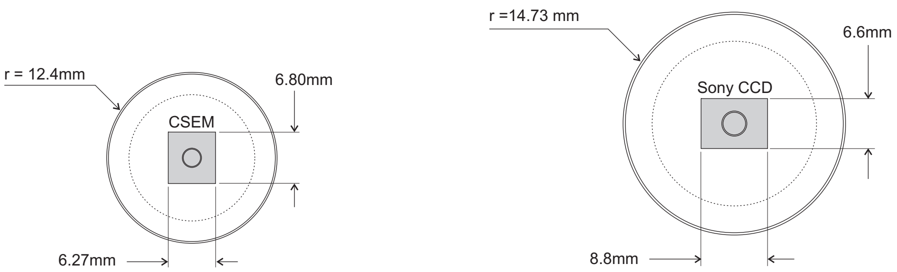

Figure 6-7 – Sensors and projected mirror area at fiber-sensor distances of 34.9mm (CSEM sensor) and 45.5mm (Sony CCD) respectively.

We can now no longer measure the entire mirror area of interest. Instead, we will be measuring only the area shown in Figure 6-7. Unfortunately, the CSEM sensor did not perform as well as was expected, and better results were obtained with a standard CCD in PSI mode. The minimum fiber-sensor distance of this CCD was limited to 45.5mm by the sensor housing. Figure 6-7 also shows the area of the mirror which can be measured with this arrangement. It should be noted that since the CCD is not directly fixed to the interferometer frame, the location of the CCD can be adjusted to some degree, unlike the CSEM sensor, which would have been fixed in such a way as to be centred on the projection of the mirror for stability purposes.

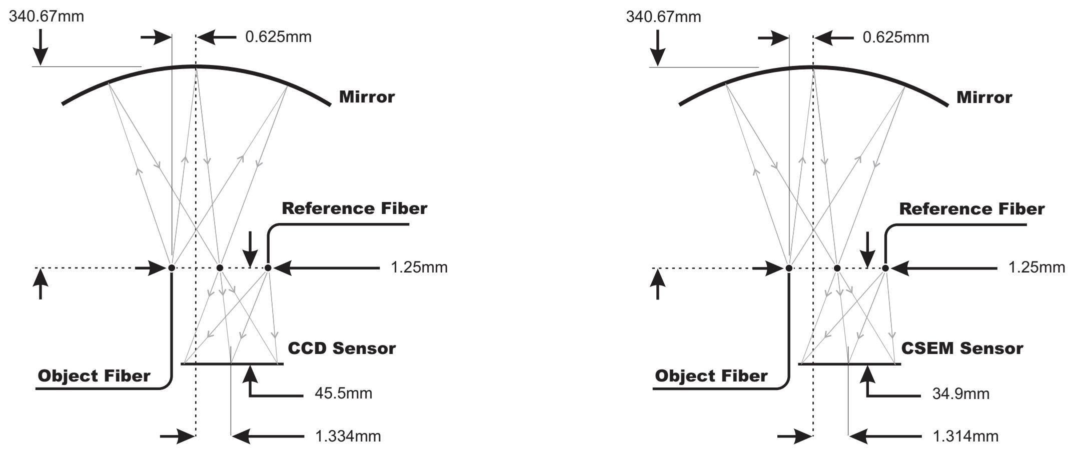

Figure 6-8 – Positioning of interferometer components

The relative positions of the various interferometer components are now completely determined.

6.2.3 Temperature stability.

In order to minimize the influence of temperature variations on the positioning accuracy of the various interferometer components, Invar has been chosen as the material for the entire interferometer frame. The extremely low thermal expansion coefficient of $1.7\times10^{-6}K^{-1}$ makes this material more than an order of magnitude less susceptible to temperature fluctuations than stainless steel.

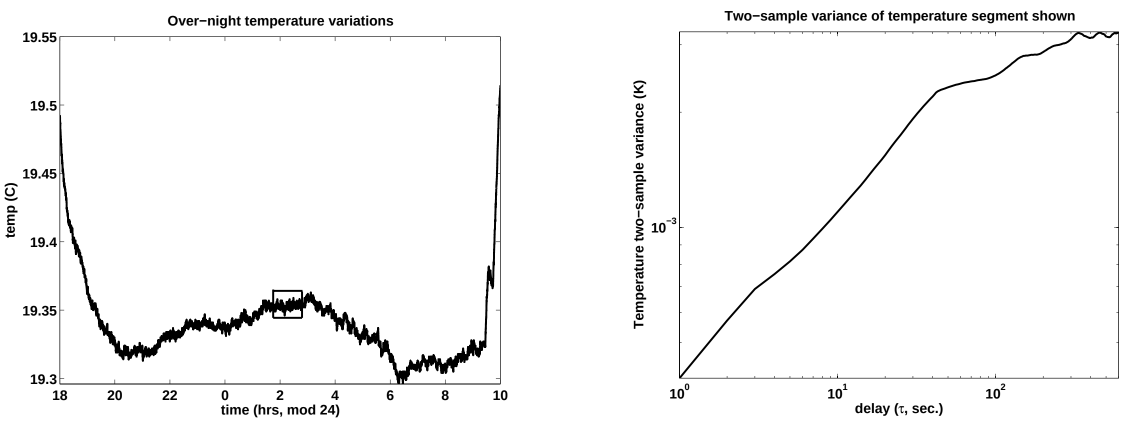

Figure 6-9 – Over-night temperature variations of optical table, and two-sample variance of indicated segment.

Figure 6-9 shows an over-night temperature measurement of the surface of our optical table, using a mK accuracy temperature probe. No active temperature control was used. The marked section is analysed in terms of the two-sample variance – giving the rms temperature difference between two samples taken a time τ apart. This analysis is used to establish that the temperature variation during a 40 frame PSI measurement (lasting 3.3s) is approximately 0.71mK. We assume that the day-to-day repeatability of the temperature is better than 0.05K, using active temperature control if necessary. These values are used to obtain the positioning stabilities in Table 6-1.

The effect of position errors on the retrieved mirror surface is discussed in more detail in chapter 3, and the results are also incorporated in Table 6-1. The short term positioning stability is sufficient to guarantee that the interferogram does not drift by a significant percentage of a fringe (in comparison to OPD drifts caused by drifts within the light source), and the figure errors due to temperature drifts during one measurement are well below the required 0.1nm. Long term drifts are considerably below the accuracy with which the corresponding quantities can be measured, and also do not induce significant figure errors compared to our desired 0.1nm accuracy level. The long term mirror-height drift is also consistent with the requirement of a maximum 50nm error in the measurement of the absolute radius of curvature.

It should be noted that these estimates are not complete, as they may depend on the geometry of the interferometer setup, but they give a convincing indication that positioning drifts due to temperature variations are unlikely to contribute to our figure errors.

	Short-term (3.3s)		Long term (24h)	
	Drift	Figure error	Drift	Figure Error
Mirror Height	0.411 nm	0.008 pm	28.9 nm	0.58 pm
Fiber separation	0.003 nm	0.001 pm	0.2 nm	0.06 pm
Fiber x-position	0.001 nm	< 0.001 pm	0.1 nm	0.02 pm
CCD x-position	0.002 nm	< 0.001 pm	0.1 nm	0.01 pm
CCD z-position	0.055 nm	0.001 pm	3.9 nm	0.08 pm

Table 6-1 – Position variations and resulting figure errors due to temperature fluctuations.

6.2.4 Placement accuracy.

A comprehensive analysis of placement errors is outside the scope of this thesis. Nonetheless, we have attempted to optimize placement accuracy by using a V-groove positioning system for both the fibers and the mirror support structure.

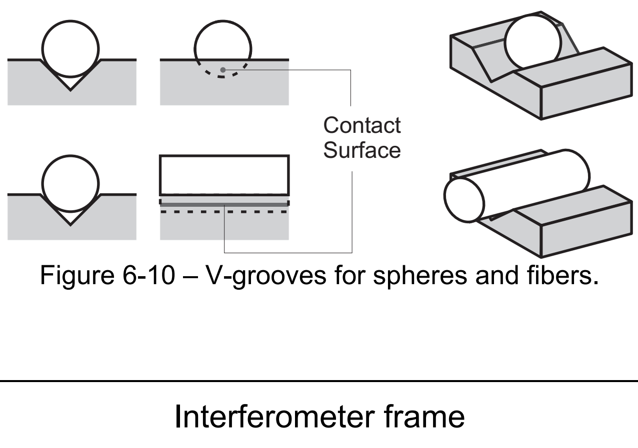

Figure 6-10 – V-grooves for spheres and fibers.

To mount the mirror support, the interferometer frame contains three V-grooves pointing towards the geometrical centre of the mounting structure to receive the positioning spheres (Figure 6-1).

This well-known design singly restricts the position of the mounting structure in all degrees of freedom. Any uniform expansion or contraction of either the interferometer frame or the mounting structure about their centre does not affect the mirror position, and does not induce stresses in either the mount or the frame. In order to reduce deformation errors, the V-grooves are clad with ceramic strips to distribute the pressure from the otherwise extremely small contact area between the groove walls and the positioning spheres. Deformation errors may still occur on the hardened positioning spheres of the mounting structure.

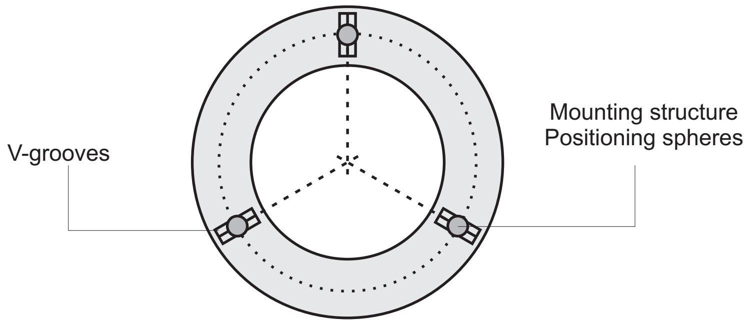

Figure 6-11 – Kinematic V-groove mount for mirror support structure

The fibers are first lightly fastened to the v-grooves with the elastic band, and then pressed down gently with the help of a gauge block, until the fiber tip is level with the fiber-holder surface. The elastic band is then tightened, and the gauge block is removed. With this procedure, sub micron vertical positioning accuracy with respect to the fiber-holder can be achieved.

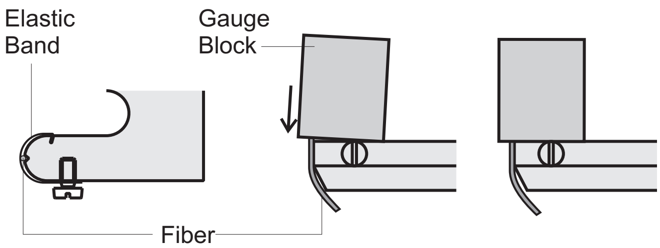

Figure 6-12 – Fiber mounting procedure. Drawings not to scale.

The interferometer consists of a number of sub-assemblies to facilitate manufacture. These include the v-grooves for the support structure and the fiber holders. To ensure accurate placement of these sub-assemblies, tight tolerancing was used in combination with guiding pins which define the locations of connected components with respect to each other.

6.2.5 Adjustability.

Once the interferometer is assembled, only three degrees of freedom of the base plate remain adjustable.

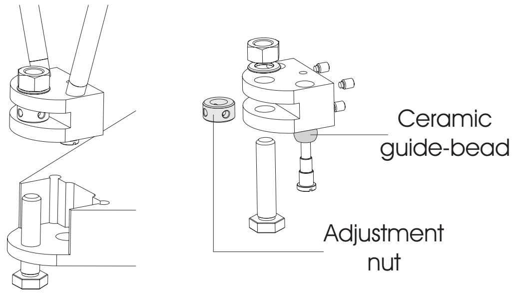

Figure 6-13 – Adjustable corners of base-plate (to scale).

The base plate is connected to the rods via corner pieces which allow the height of the base plate corners to be adjusted with an adjustment nut. The nut and corresponding bolt have a very fine thread of 0.5mm, and are also made entirely from invar. To guarantee that the in-plane position of the base plate is stable during adjustment, a ceramic guide-bead fixed to the corner piece fits snugly into a guide-hole in the base-pate. Once the corners have been optimally adjusted, they can be locked into place with the bolts on top of the corner pieces.

Since the fibers and CCD are fixed respective to the base-plate, it is easier to understand the effect of adjusting the corner height in terms of a rotation of the mirror about one of the base-plate edges. Using adjustments over the 2mm range of travel of these adjustment screws, we have a 0.3 mm out-of-plane adjustment range, with a sensitivity of ~0.5μm and a 1μm in-plane adjustment range with a sensitivity of ~1.5nm.

This adjustment range is in proportion to the different sensitivities of the figure measurement to errors in the in-plane and out-of-plane positions of the interferometer components.

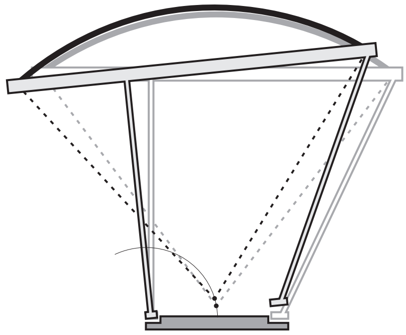

Figure 6-14 – Effect of height-adjustment of corner blocks (exaggerated motion).

6.2.6 Summary

The hexapod interferometer frame presented here is designed specifically for our test-object and optimized for stability. The positions of the various interferometer components depend on a number of factors, and will need to be adjustable in future versions of this interferometer, designed to measure a range of different mirrors. While care must be taken to ensure thermal equilibrium during a measurement, no active temperature stabilization is required since the low expansion coefficient of the material used in the construction of the interferometer keeps the figure errors introduced by temperature variations well below the required levels. The use of V-grooves makes accurate and repeatable positioning possible without introducing undue stresses. Finally, the limited amount of adjustability present in this design is representative of the sensitivity of the instrument to positioning errors.

6.3 Optical Fiber

The optical fibers play a key role in this interferometer. We have previously assumed that they produce spherical waves with an approximately Gaussian intensity distribution. In this section, we will justify this assumption, and show what limitations apply.

The optical fiber used in this interferometer is a Newport F-SPV polarization-preserving single-mode fiber. appendix A.4 lists the fiber's parameters as specified by the manufacturer.

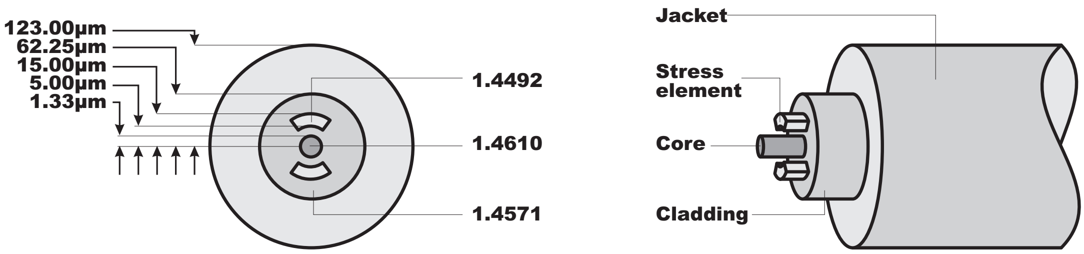

Figure 6-15 – Composition of F-SPV polarization maintaining fiber (not to scale).

Like most polarization preserving fibers, this fiber consists of a cylindrical core, surrounded by a cladding layer with stress-elements, which in turn is covered by a protective plastic jacket (See Figure 6-15). The core has a higher refractive index than the cladding ($n_{core} > n_{cladd}$).

In the following section, we will try to estimate the degree of sphericity which we can expect from the wavefronts of such fibers. To guarantee applicability to our situation, we will calculate the field from the fiber at a distance equal to the fiber-CCD distance (34.9mm) and the fiber-mirror distance (340.67) in our interferometer, without making the far-field assumption.

6.3.1 Modeled fiber output

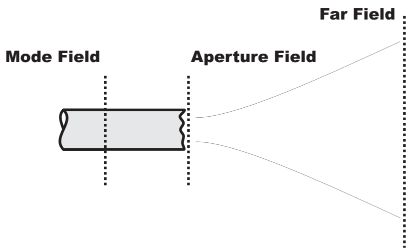

Figure 6-16 – The three field regions used to describe fiber output.

In an attempt at deriving a good model for the fiber output, we will need to consider three distinct field regions[1]: the mode field, the aperture field and the far field region, see Figure 6-16. The mode field is the field inside the fiber, far from either end. An expression for the mode field can be found by solving Maxwell's equations with the boundary conditions appropriate to a cylindrical waveguide, and will take into consideration the birefringent nature of the fiber. The aperture field is the field just outside the fiber end-face, and will be affected by the fiber surface properties. Finally, the far-field distribution is found by propagating the aperture field the appropriate distance through free space.

Mode field

Neglecting the stress elements for the moment, solutions for cylindrical waveguides are readily found and have been exhaustively discussed in the literature[2, 3]. For weakly-guiding fibers such as our own, the solutions are of the following form:

$$E_x(r) = \mathrm{e}^{i(\omega t - \beta z)} \begin{cases} AJ_l(hr) & r < a \\ BK_l(qr) & r \geq a \end{cases} \tag{6.2}$$

With:

$$r = \sqrt{x^2 + y^2} \tag{6.3}$$

The constant A is a normalization constant, but the other constants depend on the fiber specifications as follows:

$$B = A\frac{J_l(ha)}{K_l(qa)} \quad (continuity) \tag{6.4}$$

$$q^2 = \left(\frac{2\pi}{\lambda}\right)^2 \left(n_1^2 - n_1^2\right) - h^2 \tag{6.5}$$

$$h\frac{J_{l+1}(ha)}{J_l(ha)} = q\frac{K_{l+1}(qa)}{K_l(qa)} \tag{6.6}$$

$$\beta^2 = n_1^2\left(\frac{2\pi}{\lambda}\right)^2 - h^2 \tag{6.7}$$

(6.5) and (6.6) can be used together to solve for h and q, which in turn give β. The solutions to this set of simultaneous equations is not trivial, and numerical methods need to be used. For single mode fibers operated at wavelengths longer than the so-called “cut-off wavelength”, only one solution for β exists which gives a propagating mode:

$$E_x(r) = \mathrm{e}^{i(\omega t - \beta z)} \begin{cases} AJ_0(hr) & r < a \\ BK_0(qr) & r \geq a \end{cases} \tag{6.8}$$

h, q and β still need to be evaluated numerically, and depend on the refractive index difference, core-diameter and operational wavelength.

Taking our fiber parameters as given in appendix A.4, we can observe the slight wavelength dependence of the mode-field (see Figure 6-17)

The effect of the stress elements is two-fold. For one, the core (and to a lesser extent, the cladding), becomes birefringent. For our fiber, the birefringence (n_x-n_y) is 3.2x10^{-4}. The other effect is a slight ellipticity of the core.

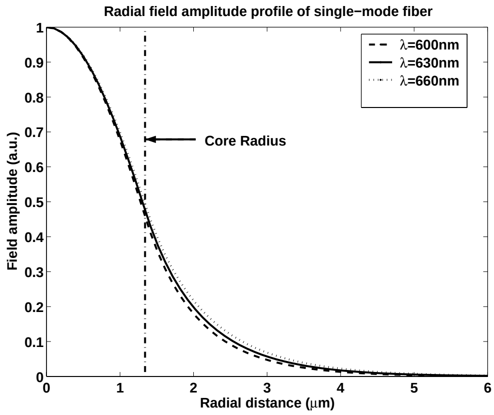

Figure 6-17 – Radial field distribution for several wavelengths in our fiber.

According to Yariv[2], the birefringence has little effect other than to destroy the degeneracy of the two orthogonally polarized modes supported in the fiber. Consequently, two orthogonally polarized modes, with virtually identical mode profiles but experiencing different refractive indices, are supported by the fiber. Since we launch linearly polarized light into the fiber along one of the optical axes, we can restrict ourselves to just one of these modes.

The ellipticity can be treated with perturbation analysis[4, 5] to show that the resulting solutions are identical to the circularly symmetrical ones save for a co-ordinate transformation mapping a circle to the appropriate ellipse. The ellipticity of our fiber is approximately 1.2:1, and we can therefore simply replace (6.3) with:

$$r = \sqrt{x^2 + \left(y/1.2\right)^2} \tag{6.9}$$

It should be noted that this fundamental mode has the potential to couple into so-called cladding modes, which propagate within the air-cladding or jacket-cladding interface. Such coupling can occur through imperfections in the fiber and sharp bends. However, these modes are generally lossy and do not propagate far in the fiber. We have neglected these here, and rely on the experimental results to justify this assumption.

Aperture field

Although there is surprisingly little literature exclusively on the subject of finding the aperture field from the mode field, several papers concerned with finding the mode field from the far-field distribution have been published, making the assumption that the aperture field can be derived simply by multiplying the local mode field amplitude (given by equation (6.8) in our case) with the Fresnel transmission coefficients for a plane-wave perpendicular to an interface, using the local refractive index at the fiber end face[1, 6, 7].

$$E_{ap}(x,y) = E_{mo}(x,y) \times \frac{2n(x,y)}{n(x,y)+1} \tag{6.10}$$

Equation (6.10) shows how the aperture-field amplitude (E_{ap}) can be calculated from the mode-field amplitude (E_{mo}) and the local refractive index (n) using this approach.

A more accurate solution might be found by taking the finite extent of the plane wave into account by de-composing it into its angular spectrum before applying the Fresnel transmission coefficients. Young[1] notes that the difference between these two methods is negligible. Another option, beyond the scope of this thesis, would be to rigorously solve Maxwell's equations with the appropriate boundary conditions present at the air-fiber interface and using the mode field as incident field. We will be satisfied with the first approach, confident that the difference between the results will be small, especially considering the fact that the fiber surface shape has not actually been measured, and some assumptions about this will need to be made anyway.

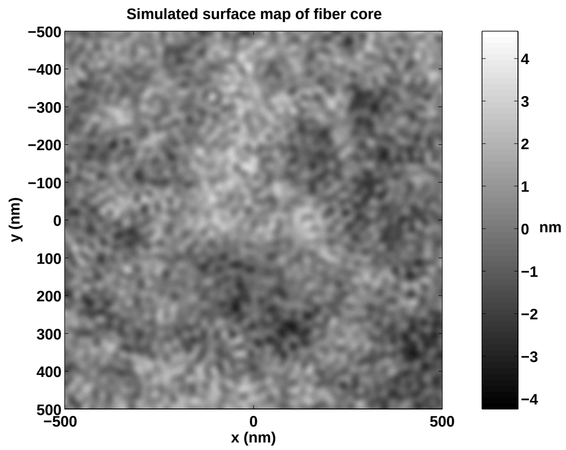

Figure 6-18 – Simulated surface profile of fiber core.

Atomic force microscopy (AFM) Measurements on cleaved single mode fibers[8] have shown that the act of cleaving causes the core to recede by a few nm at the end-face. Furthermore, the cleaved surface generally exhibits an rms roughness of 0.5nm with a 1/f spatial frequency distribution, stopping at 0.08nm^{-1}. We simulate this by filtering white noise appropriately and normalizing to give σ=0.5nm. The resulting height profiles (see Figure 6-18) compare very well with the measurements reported by Poumellec[8]. We then add a 3nm deep cylindrical dip at the centre, having a radius of 1.35μm and with a smooth transition over 0.5μm, as described in Poumellec's paper.

In accordance with the paper by Young[1], we will model the transition of the mode field to the aperture field as follows: The amplitude of the aperture field is given by multiplying the mode field distribution with the Fresnel transmission coefficient using the local refractive index, as shown in equation (6.10).

The phase of the aperture field is found by calculating the retardation experienced at different locations of the fiber end-face, due to the varying height profile:

$$\phi(x,y) = \frac{2\pi}{\lambda}\left[1 - n(x,y)\right] \times h(x,y) \tag{6.11}$$

Here, $\phi(x,y)$ is the aperture phase profile, $n(x,y)$ is the fiber's refractive index profile, and $h(x,y)$ is the local height profile. Typical aperture field values are shown in Figure 6-19

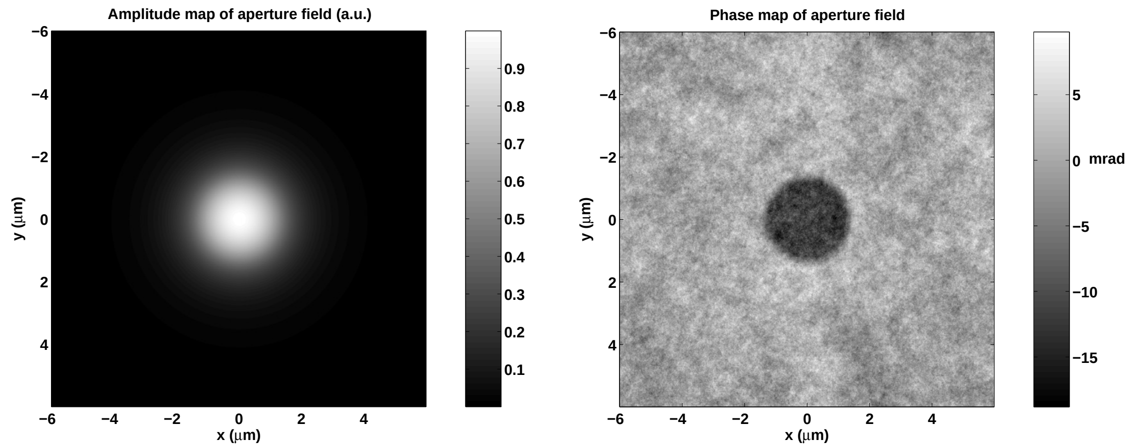

Figure 6-19 – Amplitude and phase of aperture field.

Far field

While the far field of cylindrically symmetric aperture fields can be evaluated via the Hankel transform[7], we have to resort to an evaluation of the two-dimensional diffraction integral. Owing to the stochastic nature of the phase profile to be propagated, analytical solutions do not exist, and a numerical evaluation will have to be performed instead. Although the results from vector- and scalar diffraction theory differ appreciably, especially at large NAs, Young[7] shows that this difference is mainly in the field amplitude, while the phase function may be considered unaffected for our purposes.

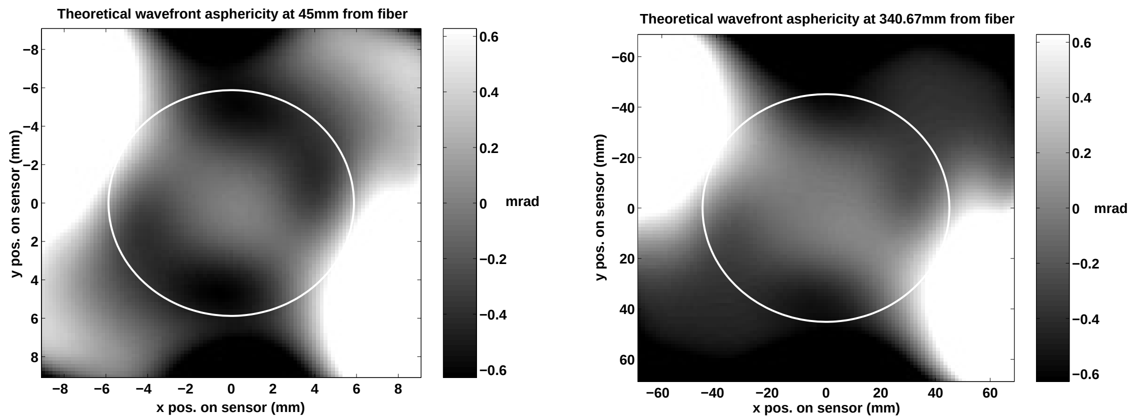

Figure 6-20 – Optical phase difference between spherical wave and simulated fiber output for 45mm and 340mm. Phases outside our tolerance appear clipped.

We therefore evaluate the Sommerfeld diffraction integral numerically, using the aperture field derived above, to evaluate the far field at a plane z=34.9mm and z=340.67mm from the fiber tip. These brute-force calculations are considerably simplified by the fact that the aperture field is confined to such a small area, allowing us to choose a 1000x1000 element grid for a 12x12μm area, to yield a grid resolution of 12nm – equivalent to the highest spatial frequency components in the fiber's height profile. Numerical simulations show that increasing the grid resolution does not alter the results by more than 0.3μrad. The resulting field distribution and the associated integration kernel can therefore be stored entirely in the working memory of a standard PC, and the double integral evaluated very quickly as a discrete sum. Figure 6-20 shows the resulting far-field distribution and the phase difference from a spherical phase profile. It should be noted that the centre of curvature of the spherical wave has not been chosen to coincide with the fiber end-face, and has instead been taken as the best-fit location. This focal shift is analogous to the focal shift in Gaussian beams. The images of the phase difference are clipped at +/-2π/10'000, equivalent to a wavefront error of 0.06nm, to show regions of unacceptable wavefront errors. In addition, the largest inscribing circle of the region of acceptable phase deviation is displayed on the image.

We can conclude that we may expect the measured wavefronts to be spherical to within 0.06nm for a half opening angle of 0.13 rad at both 45mm and 340.67mm from the fiber tip. This means that the largest standard NA of a reflective optic which can reliably be measured in one go using this fiber is 0.26 (See the closing remarks of section 6.2.1). To measure optics with larger NAs, such as our example optic, the measurement either needs to be performed in several parts, or fibers with smaller core diameters/apertures at the end-face need to be used.

6.3.2 Experimental measurement of fiber output.

Using an experimental setup as shown in Figure 5-5, we can estimate the wavefront sphericity from our fibers. Similar experiments have been carried out [6, 9, 10], which are in agreement with our findings, and place the wavefront sphericity at the theoretically predicted 0.06nm[9].

In these experiments, the separation between the fiber tips was approximately 0.44mm, and the fiber-to-sensor distance was 163 mm, producing fringe-patterns as shown in Figure 6-21. Using phase shifting interferometry as explained in 5.2.4, we retrieved the optical path difference (OPD) map also shown in Figure 6-21.

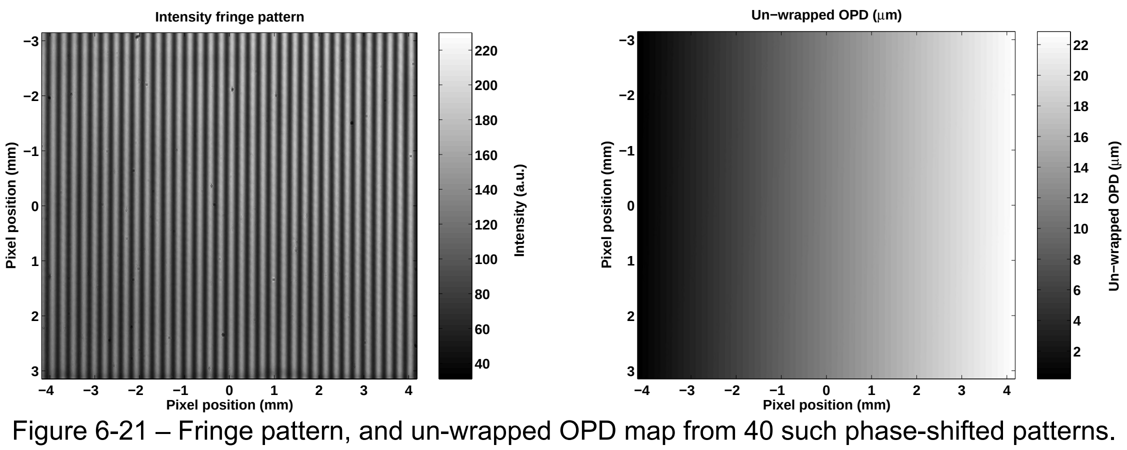

Figure 6-21 – Fringe pattern, and un-wrapped OPD map from 40 such phase-shifted patterns.

We then performed a parameter optimization on the fiber-tip locations to minimize the misfit between our data and the theoretically predicted OPD. The remaining misfit can be seen in Figure 6-23, along with a low-pass filtered version. The noise seen before low-pass filtering is made up of stochastic noise and patterns introduced by multiple reflections.

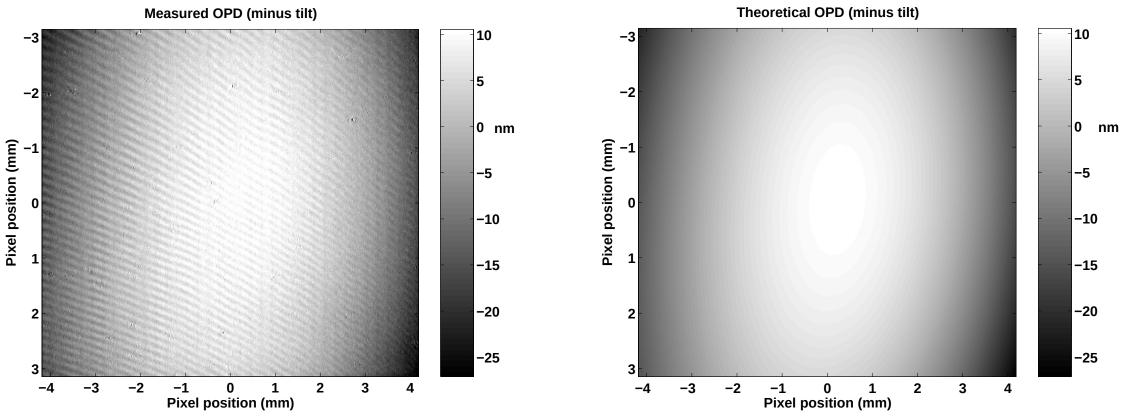

Figure 6-22 – OPD minus tilt. Measured (left) and model (right)

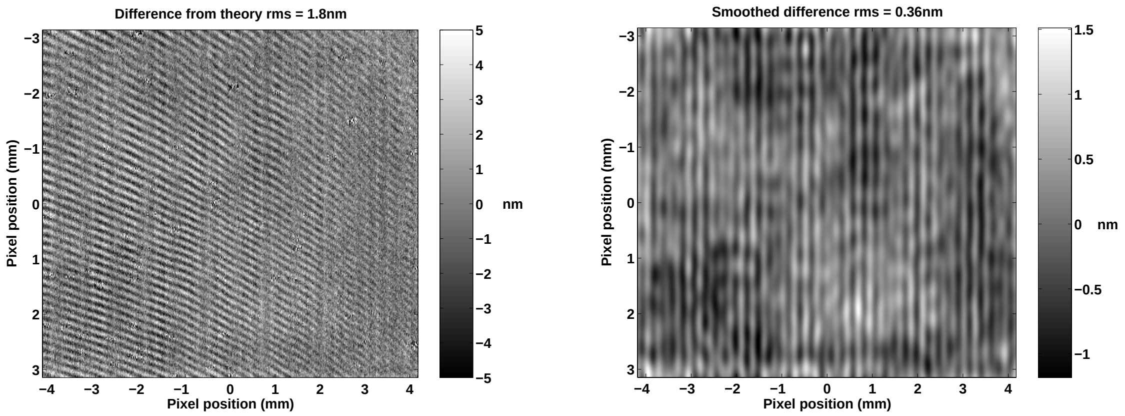

Figure 6-23 – Misfit between measurement and model (left), and low-pass filtered version to get rid of parasitic reflection artifacts (right)

After filtering, the error image is dominated by residual fringes, resulting from nonlinear phase drifts during the measurement. The reduction of the rms noise from 1.8nm to 0.36nm after filtering (convolution with a 5x5 pixel kernel), is consistent with a stochastic noise process, and we may conclude that our measurement is noise limited.

The theoretical and actual OPD functions can therefore be said to agree to within 0.36nm rms. It should be noted that the systematic pixel positioning errors discussed in 5.2.2 play a negligible role in these measurements, because of the extremely low fringe density compared with the actual mirror form measurements.

The OPD is basically the difference between the laterally sheared wavefronts from the two fibers:

$$OPD = \frac{\lambda}{2\pi}\left[\hat{\phi}_1\left(x,y,z\right)-\hat{\phi}_2\left(x-\delta x,y,z\right)\right] \tag{6.12}$$

Where we have assumed, without loss of generality, that the two fibers are separated only along the x-axis. $\hat{\phi}_1$ and $\hat{\phi}_2$ are the absolute phase functions of the two fibers, and δx is the distance by which the two are separated.

If the two wavefronts are identical, the OPD by is given by their shearing function, which can be approximated by the first derivative:

$$OPD = \frac{\lambda}{2\pi}\left[\hat{\phi}\left(x,y,z\right)-\hat{\phi}\left(x-\delta x,y,z\right)\right] \approx \frac{\lambda \delta x}{2\pi}\frac{\partial \hat{\phi}}{\partial x} \tag{6.13}$$

Our measurement of the wavefront sphericity is therefore less sensitive to low-spatial frequency aberrations than to higher-frequency ones. Additional measurements, where the fiber-tip locations remain approximately constant while the fiber direction is varied, could increase our sensitivity to such low-spatial frequency aberrations, but have not yet been carried out.

We can state that the measured OPD function is consistent with the assumption that the fiber wavefronts are spherical to within our measurement accuracy, and the theoretically predicted wavefronts are spherical to within the required accuracy over a half-opening angle of 0.13 rad for typical fiber to CCD distances.

6.4 Conclusion

The interferometer frame presented in this chapter is sufficient to guarantee the measurement accuracies required of our measurements. The simplicity and extreme stability of this design has come at the cost of flexibility – having been optimized for one particular mirror substrate only. The generic design of a hexapod structure still lends itself ideally for extension to an adjustable interferometer frame however, by using commercially available hexapods, such as ones produced by Physik Instrumente (PI) for example (See Figure 6-24). These hexapods can reach absolute positioning accuracies of ~1μm, and sensitivities of 50nm, allowing the use of off-axis substrates, various focal lengths and the ability to measure larger NA substrates in sections.

Figure 6-24 – Adjustable hexapod structure (M-850) by Physik Instrumente (PI).

A certain degree of adjustability could also be introduced to the fiber holders, to facilitate the measurement of smaller NA optics by decreasing the fiber separation.

This chapter has also shown the wavefronts produced by the optical fibers should theoretically be spherical within our requirement of 0.06nm over a half opening angle of 0.13 rad, sufficient for optics with a standard NA of 0.26. This prediction is in agreement with our experimental observations.

6.5 References

[1] M. Young, "Mode-field diameter of single-mode optical fiber by far-field scanning," *Applied Optics*, vol. 37, pp. 5605-5619, 1998.

[2] A. Yariv, *Optical electronics in modern communications*, 5th ed. New York: Oxford University Press, 1997.

[3] D. Marcuse, *Theory of Dielectric Optical Waveguides*, 2nd ed: Academic Press, 1991.

[4] A. W. Snyder and W. R. Young, "Modes of Optical-Waveguides," *Journal of the Optical Society of America*, vol. 68, pp. 297-309, 1978.

[5] J. Noda, K. Okamoto, and Y. Sasaki, "Polarization-Maintaining Fibers and Their Applications," *Journal of Lightwave Technology*, vol. 4, pp. 1071-1089, 1986.

[6] H. G. Rhee and S. W. Kim, "Absolute distance measurement by two-point-diffraction interferometry," *Applied Optics*, vol. 41, pp. 5921-5928, 2002.

[7] M. Young and R. C. Wittmann, "Vector Theory of Diffraction by a Single-Mode Fiber - Application to Mode-Field Diameter Measurements," *Optics Letters*, vol. 18, pp. 1715-1717, 1993.

[8] B. Poumellec, P. Guenot, R. Nadjo, B. Keita, and M. Nicolardot, "Information obtained from the surface profile of a cut single-mode fiber," *Journal of Lightwave Technology*, vol. 17, pp. 1357-1365, 1999.

[9] G. E. Sommargren, "Diffraction methods raise interferometer accuracy," *Laser Focus World*, vol. 32, pp. 61-&, 1996.

[10] G. E. Sommargren, D. W. Phillion, M. A. Johnson, N. Q. Nguyen, A. Barty, F. J. Snell, D. R. Dillon, and L. S. Bradsher, "100-picometer interferometry for EUVL," in *Emerging Lithographic Technologies VI, Pts 1 and 2*, vol. 4688, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 2002, pp. 316-328.

7 Measurements

The measurements presented in this chapter were made on the EUVL mirror substrate discussed in chapter 6. We have used the Sony CCD in phase-shifting interferometry (PSI) mode as mentioned in chapter 5 to obtain the OPD measurements for two wavelengths. The two wavelengths (632.99nm and 637.32nm) were produced with the stabilized multiple wavelength source as described in chapter 4.

7.1 Results

Figure 7-1 shows two typical PSI fringe-patterns used to derive the optical phase difference between the reference- and object-beams.

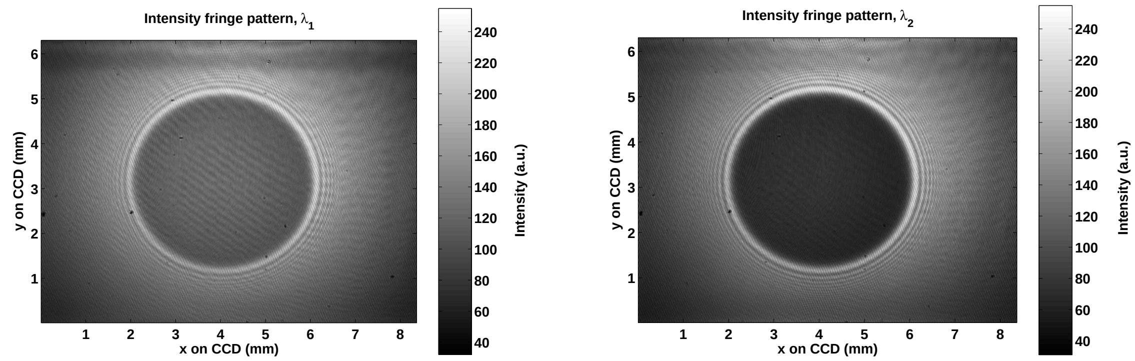

Figure 7-1 – Intensity fringe patterns obtained from interferometer with test substrate described in chapter 6, using λ_1 and λ_2.

The resulting phase profiles from λ_1 and λ_2 are shown in Figure 7-2, and the synthetic phase profile is shown in Figure 7-3. Histogram correction has already been carried out. After following the procedure outlined in chapter 5 of how to use the synthetic phase profile to unwrap the under-sampled phase map for λ_1, we obtain the optical path difference (OPD) map shown in Figure 7-5. The central obstruction is masked to avoid problems during the unwrapping phase.

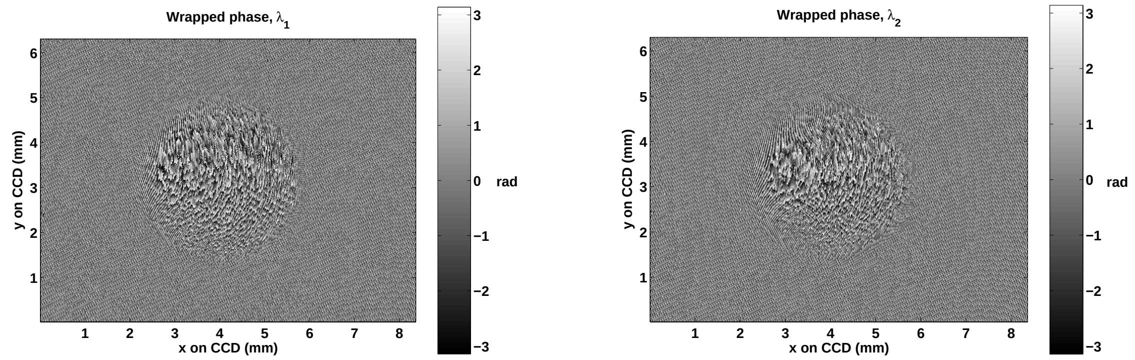

Figure 7-2 – Phase profiles retrieved from a sequence of 36 phase-stepped intensity fringe patterns, using the WDFT algorithm described in chapter 5.

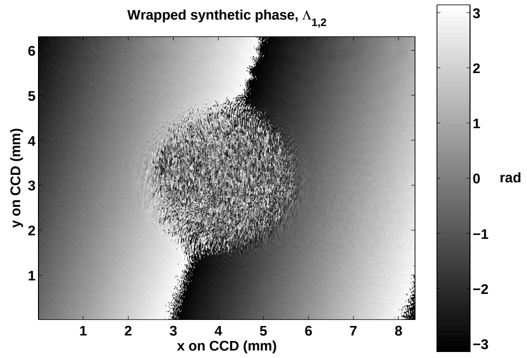

Figure 7-3 – Wrapped synthetic phase, obtained by subtracting the two phase maps shown in Figure 7-2, modulo 2π.

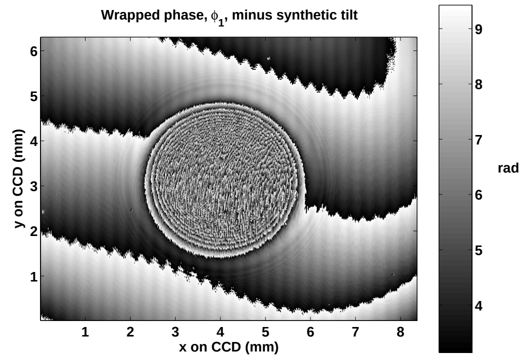

Figure 7-4 – Phase map for λ_1 (Figure 7-2) after subtracting (modulo 2π) linear phase terms derived from un-wrapped synthetic phase profile (Figure 7-3)

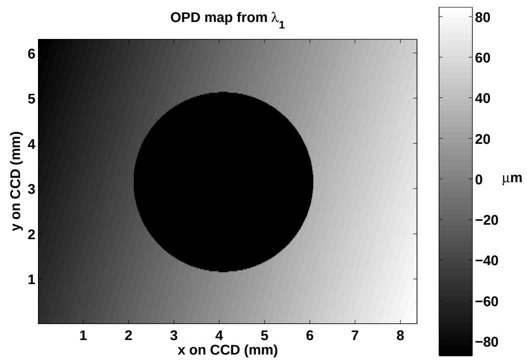

Figure 7-5 – OPD map for λ_1, obtained by un-wrapping the phase-map of Figure 7-4 and re-adding the linear terms previously subtracted.

We now fit this OPD map with the first 38 Zernike polynomials as outlined in chapter 3.3.2.

Figure 7-6 displays our OPD map alongside the resulting Zernike-38 fit (with tilt removed for clarity). Figure 7-7 displays the remainder of the OPD map after the first 38 Zernike components are removed, to reveal some diffraction artefacts around the obstruction and systematic artefacts due to CCD pixel positioning errors. After applying the CCD calibration as discussed in chapter 5.2.2, we obtain the second image of Figure 7-7. The rms values of these residual errors are 13.6nm and 11.6nm respectively, and it should be apparent from the figure that the CCD calibration procedure was unable to remove the majority of the

CCD artefacts. This is not surprising, since the calibration procedure assumes the absence of a cover-glass.

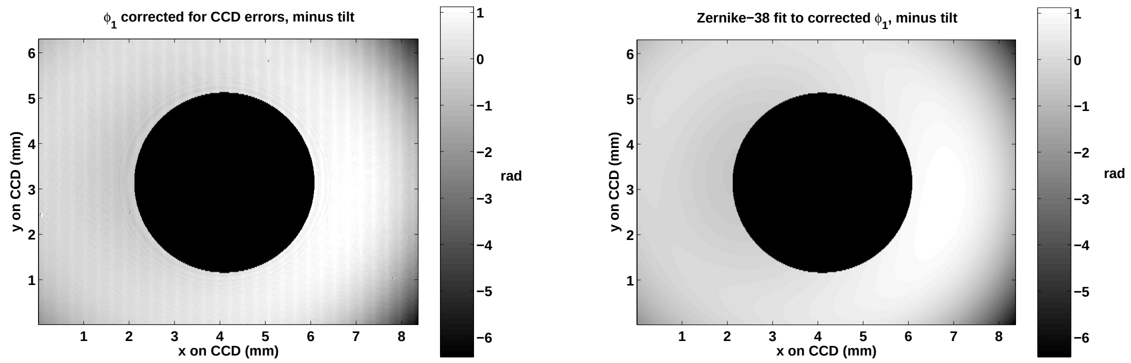

Figure 7-6 – Phase map from Figure 7-5 and the Zernike-38 fit. Both with tilt removed.

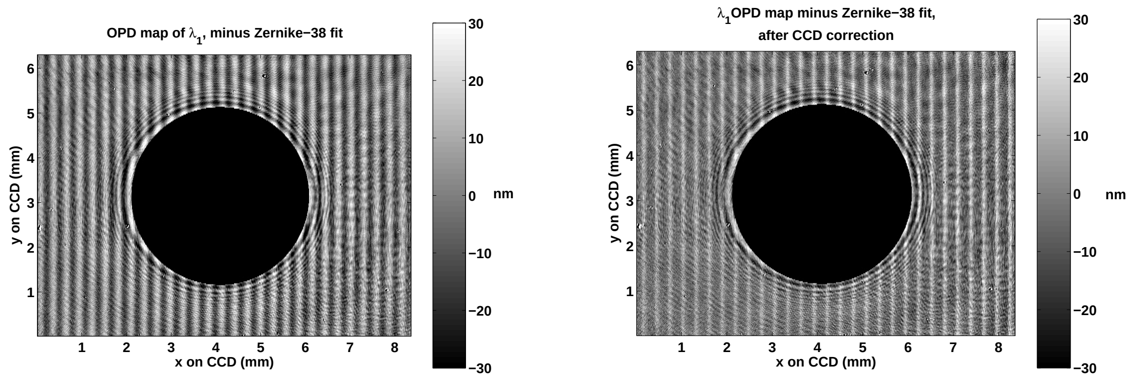

Figure 7-7 – difference between Figure 7-5 and the Zernike-38 fit. Also shown is the same difference image after applying the CCD correction discussed in chapter 5.

7.2 Analysis

For an accurate inversion, it is necessary to first know the position values for our interferometer components. Since our interferometer has neither been calibrated nor extensively measured, we must expect the nominal position values to differ significantly from the actual ones. We will be using two approaches to estimate the actual position values. We begin by using some simple features of the measured interferogram together with a number of assumptions to get an initial estimate of the parameters. This estimate is then refined by launching an optimization routine which aims to minimize the misfit between the measured and the ideal mirror shape.

7.2.1 Parameter estimation using interferogram features.

A number of position values can be estimated from features of our OPD map. They are:

1. The vertical distance from the CCD to the focal point – by measuring the size of the central obstruction projected onto the CCD, compared to the size of the obstruction on the mirror.
2. The distance (and direction) from the reference fiber to the focal point – by measuring the slope of the OPD function
3. The in-plane location of the object fiber, with respect to the CCD – by locating the centroid of the central obstruction projected onto the CCD.

The separation of the object- and reference-fiber was measured using a microscope mounted on a milling-machine translation table at 2.50mm +/- 0.01mm. In order to obtain an estimate for the remaining component positions, two additional assumptions are made:

1. The fibers and focus lie in the same horizontal plane (implicitly, we also assume that the object fiber is in the horizontal plane of the mirror's centre of best-fit curvature.
2. The object fiber, focal point and reference fiber lie along a straight line.
3. The CCD surface is perpendicular to the plane defined by the object fiber, reference fiber and mirror apex.

These assumptions are not strictly correct, but if the tolerancing specified by the interferometer frame manufacturers is reliable, the above assumptions will hold to within a few tens of μm. It is important to keep in mind the fact that once the relative positions of the interferometer components are known, either by direct measurement, or by calibration, the procedures described in this section are no longer necessary.

Vertical CCD – focal point distance

The rays from the boundary of the central obstruction passing through the focus to the CCD trace out a circle (approximately – an ellipse is a more accurate description due to the off-axis nature of the chief ray), the radius of which depends on the vertical distance of the CCD from the focus as follows:

$$R_{proj.} \approx R_{on\ mirror} \left| \frac{z_{s,f}}{z_{CCD,f}} \right| \tag{7.1}$$

Our task of measuring the radius of the projected circle is complicated by two factors. One: we do not know the centre of this circle and two: due to diffraction there is no sharp shadow boundary for this projection

Ironically, the latter complication is easily circumvented, and actually helps us in resolving the first complication. From our treatment of diffraction due to the central obstruction, we know that the shadow boundary coincides with the first zero-crossing of the absolute phase difference between the diffracted field due to the obstruction and the unobstructed (geometrical) field (see Figure 3-27). By subtracting the lower spatial frequencies of our un-wrapped interferogram from the same, we artificially re-create such an interference experiment (See Figure 7-8).

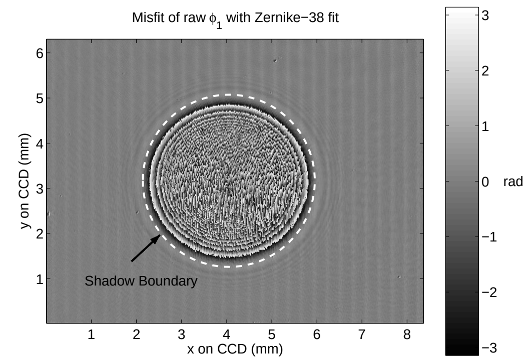

Figure 7-8 - Phase map for λ_1 (Figure 7-2) after subtracting (modulo 2π) the 38 Zernike polynomials obtained by fitting outside the geometrical shadow.

A plot of the phase as a function of the radial distance from the centre of the projected obstruction should strongly resemble Figure 3-27, and the radius of the projected obstruction can be found by locating the first zero-crossing of the absolute phase difference.

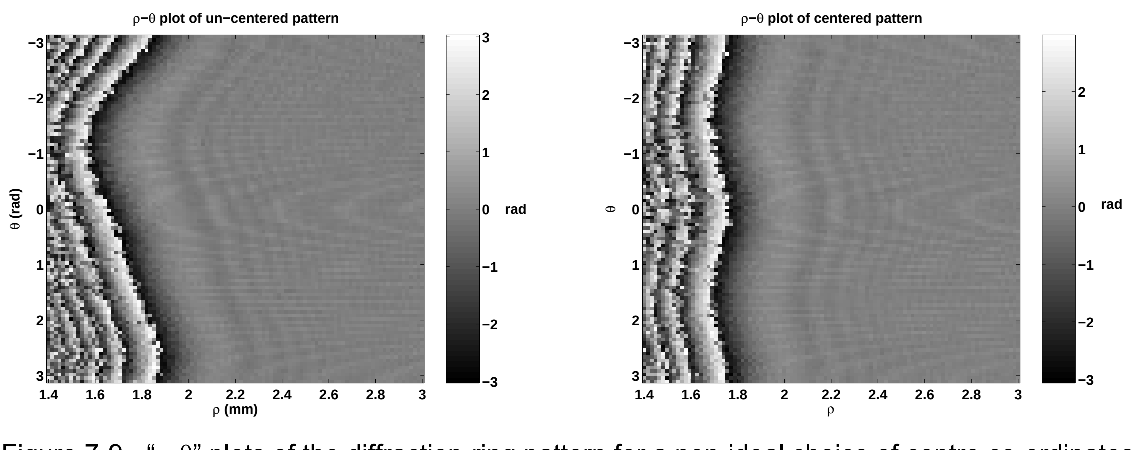

Figure 7-9 - "ρ-θ" plots of the diffraction-ring pattern for a non-ideal choice of centre co-ordinates (left) and an optimum choice of centre co-ordinates (right)

A "ρ-θ" plot of the phase as a function of radial distance and angle from a point near- but not at- the centre of the projected obstruction (Figure 7-9) should help the reader to understand the approach we will use to determine the centre of the projected obstruction.

After averaging along the "θ" co-ordinate (see Figure 7-10), an error in our guess of the location of the centre of the obstruction will cause a smearing-out of the radial diffraction pattern, hence decreasing its modulation depth. We will therefore measure the modulation depth of these radial phase functions, and find the location for which the resulting modulation depth is largest.

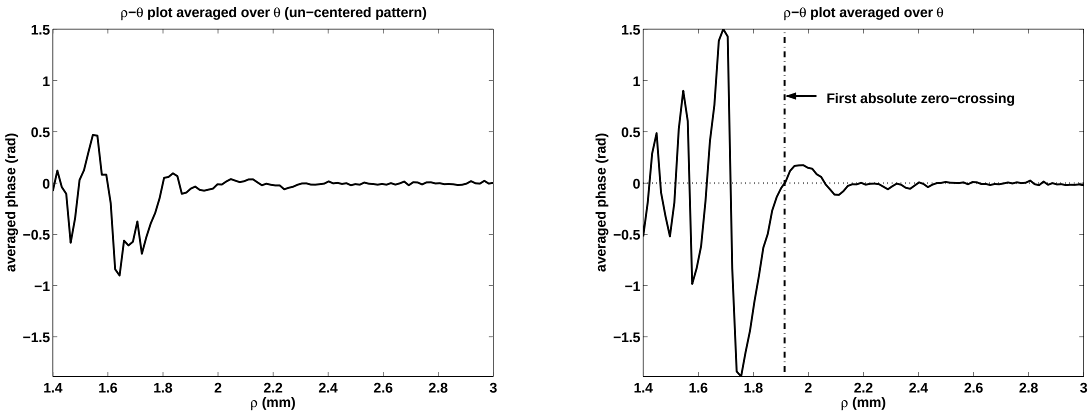

Figure 7-10 – The plots in Figure 7-9 were averaged along the "θ" co-ordinate to obtain plots which ideally should resemble the theoretical diffraction patterns predicted in chapter 3.

Figure 7-9 also shows the ρ-θ plot resulting from our best estimate of the centre of the obstruction obtained this way. The remaining deviations are due to the slight ellipticity expected of our projection, and perhaps an anamorphic scaling error for our pixel locations.

The zero-crossing marking the shadow-boundary is clearly visible in the radial phase plot, giving us an estimate of the radius of the projected obstruction of 1.937mm. Inserting the known radius of the obstruction on the mirror (12.4mm) into (7.1) above, together with the assumed focus-mirror distance of 340.67mm, gives us an estimate for the focus-CCD distance of 53.208mm.

Reference fiber – focal point distance

The slope of the OPD function for our interferogram is:

$$\left(\frac{\partial OPD}{\partial x}, \frac{\partial OPD}{\partial y} \right) = \left(0.0168, -0.0051 \right)$$

The slope of the OPD function is a function of the fiber-focus separation and the vertical distance of the CCD from the fibers (see (5.11)). Consequently, the location of the reference fiber relative to the focal point is

(Δx, Δy)=(0.8837mm, -0.2697mm)

This indicates an in-plane rotation of the CCD relative to the axis defined by the object fiber – reference fiber separation by -16.97˚.

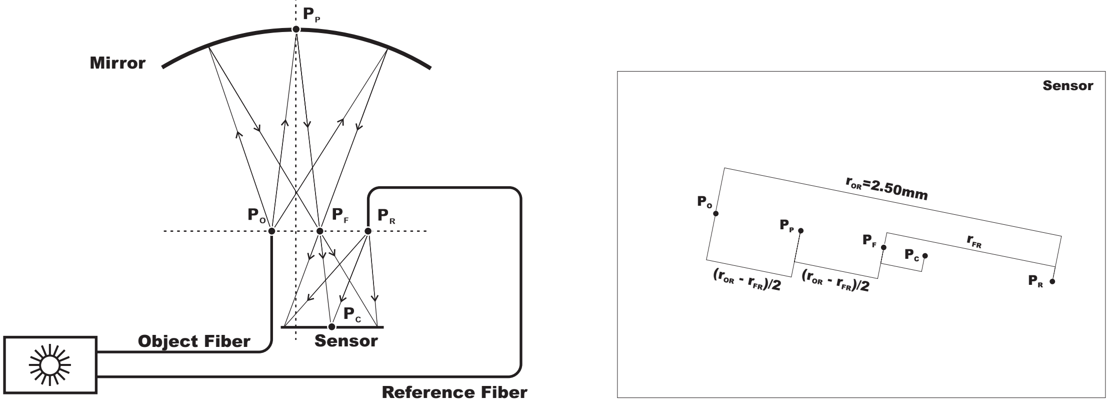

Figure 7-11 – Relative positions of relevant interferometer components. Front-view (left) and top-view (right). Neither diagram is to scale.

Position of CCD relative to fibers and focus

Knowing the distance of the focal point to the reference fiber, together with our measurement of the reference- to object-fiber distance, also gives us the horizontal distance between the mirror apex and the focus (see Figure 7-11). The centre of the projected central obstruction can therefore serve as a reference point on the CCD. The position of this point on the CCD, relative to the apex of the mirror, is given by:

$$\left(x_p, y_p, z_p\right) \approx \left(x_f, y_f, z_f\right)\frac{z_p}{z_f} \tag{7.2}$$

Inversion results

Using the Zernike-38 fit to our measured data with our estimate of the parameters (shown in Table 7-1) to perform an inversion, yields the retrieved mirror form shown in Figure 7-12. The Error figure with the nominal mirror shape is shown in Figure 7-13. With piston, tilt and defocus removed, it has an rms value of 4.2nm.

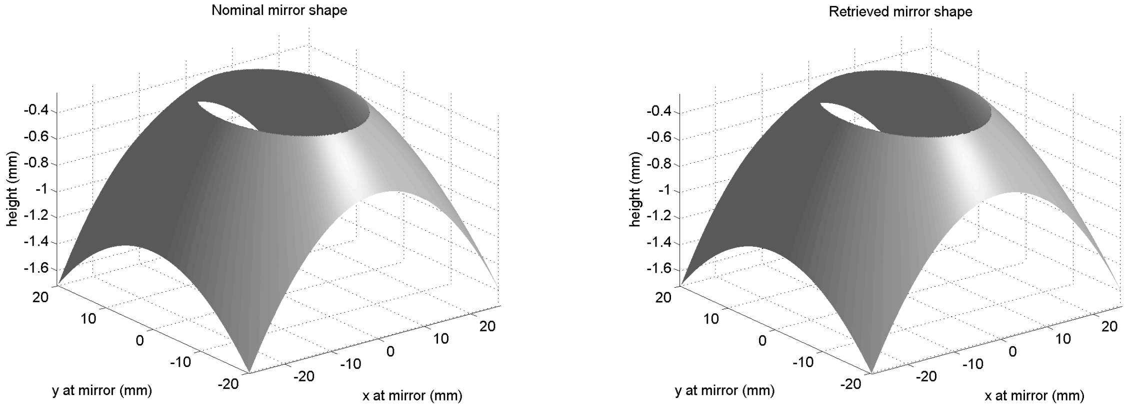

Figure 7-12 – Nominal and retrieved mirror surface shape, using estimated parameters.

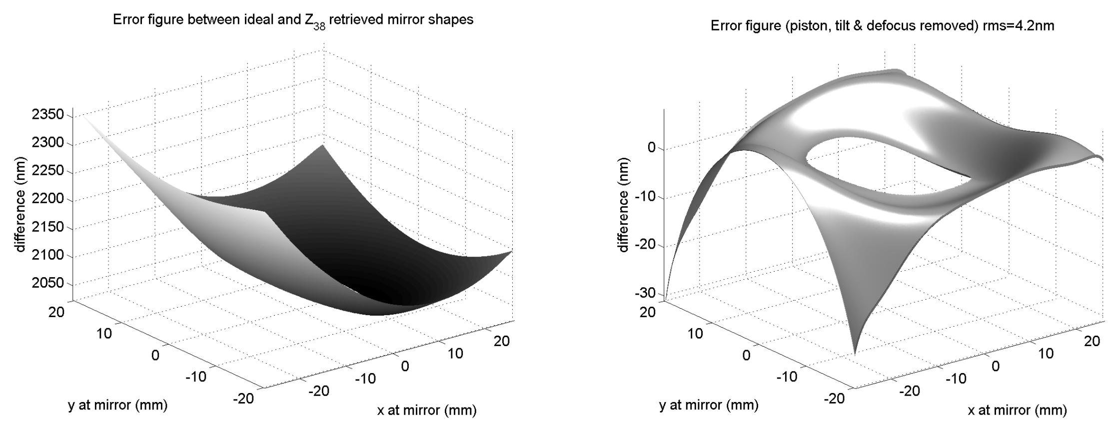

Figure 7-13 – Difference between nominal and retrieved surface shapes, using initial parameter estimates. Piston, tilt & defocus are removed in the second image, resulting in a 4.2 nm rms difference.

Use of Zernike-38 fit

As outlined in chapter 3.3.2, the use of the Zernike-38 fit to our data has a number of advantages over the use of the raw data. Apart from naturally reducing the effects of diffraction, a number of CCD-dependent errors are also filtered out.

Due to the expected magnitude of the cover-glass and CCD related errors introduced, the use of the Zernike-38 fit seems justified as a means of removing diffraction effects. The marginal improvement gained by using the complex hybrid inverse propagation algorithm (IPA) of chapter 3.8 will be insignificant in comparison to the remaining systematic errors. Nonetheless, once improved sensors are available, the hybrid IPA will be able to deal with diffraction effects more rigorously.

To confirm that the use of the Z_{38} data does not introduce any significant errors with low spatial frequencies, we also show the results of the inversion with the raw data and the parameters estimated above (Figure 7-14 & Figure 7-15). The inversion with the raw data has an rms deviation from the nominal surface shape of 10.1nm.

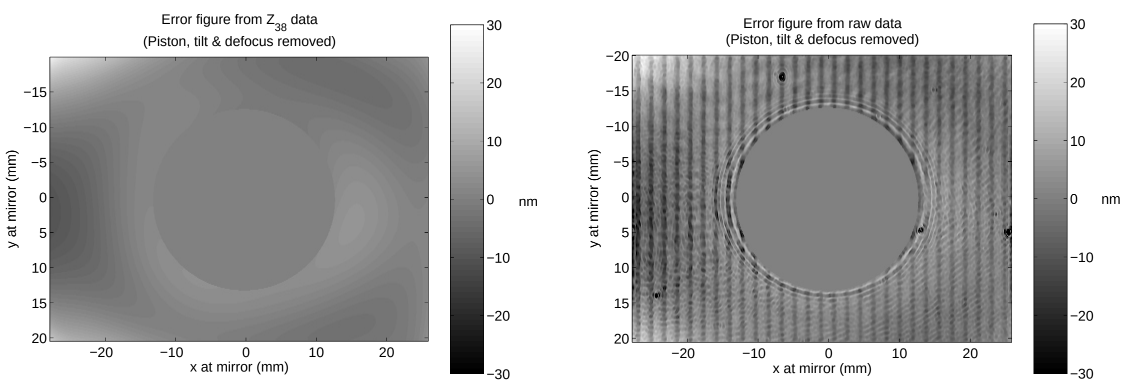

Figure 7-14 – Retrieved error figure from Z_{38} fit to raw data (left) and from raw data (right).

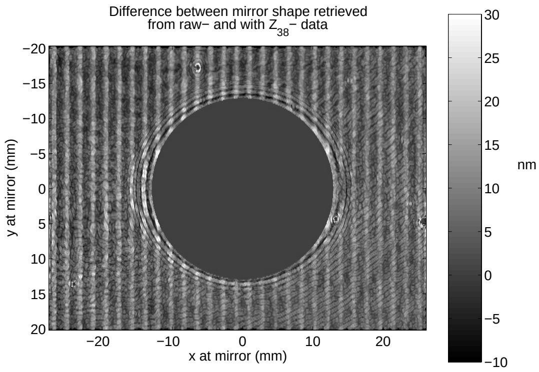

Figure 7-15 – Difference between error figures in Figure 7-14. Note the absence of low spatial frequency components.

7.2.2 Parameter estimation using optimization.

We now apply an optimization procedure outlined in chapter 3.3.2 to retrieve the set of interferometer parameters which minimize the misfit between the measured- and the nominal data (given in Table 7-1). With this set of parameters, we get significantly better agreement before subtraction of piston, tilt and de-focus errors compared with the initially estimated parameters (see Figure 7-16 cf. Figure 7-13).

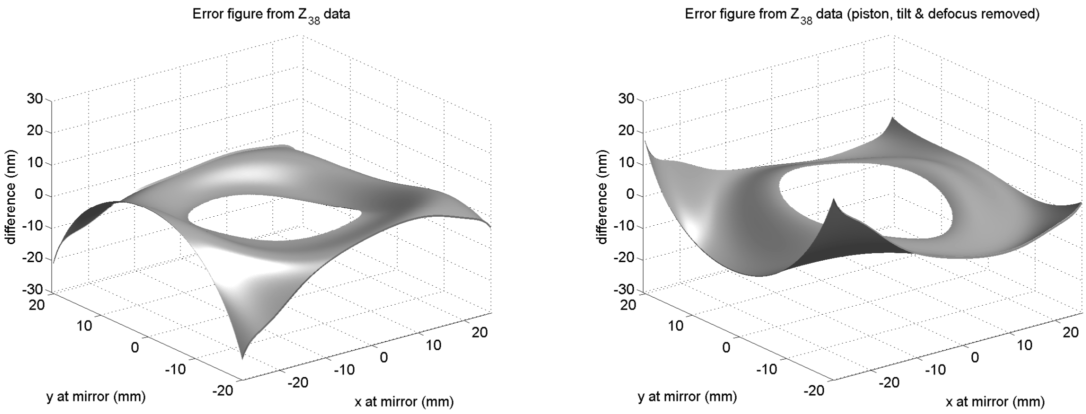

Figure 7-16 – Difference between nominal and retrieved surface shapes, using optimized parameter estimates. Piston, tilt & defocus are removed in the second image, resulting in a rms difference of 3.9 nm.

Parameter	Nominal value (mm)	Estimated value using interferogram features	Estimated value using optimization
Object fiber (x,y,z)	(-0.625, 0, -340.67)	(-0.75, 0.21, -340.67)	(-0.75, 0.21, -340.67)
Reference fiber (x,y,z)	(1.875, 0, -340.67)	(1.65, -0.46, -340.67)	(1.65, -0.46, -340.76)
CCD centre (x,y,z)	(1.334, 0, -386.17)	(0.87, -0.24, -393.83)	(0.91, -0.29, -393.83)
OPD at centre	(681.32)	(681.32)	(681.23)

Table 7-1 – comparison between nominal and estimated positions of interferometer components.

7.3 Conclusion

Due to the excessive noise introduced by the inadequate sensor used to make these measurements, it is difficult to objectively assess the resulting instrument accuracy. It was already clear from chapter 3.1.3 that knowledge of the correct interferometer parameters plays a significant role in the accuracy of the final measurement. While we have addressed this issue “a-posteriori”, it should be possible to obtain the positions of the interferometer components by means of a combination of independent metrology and self-calibration, using as yet to be determined procedures. Once these have been obtained, they can directly be used for the retrieval of surface shapes, without the need for the complicated estimation steps outlined above.

The fact that the Z_{38} fit to the raw data, together with our initial, rough estimate of the interferometer parameters has already resulted in an rms error of merely 4.2 nm (after subtraction of piston, tilt and de-focus), is proof of the principle of this novel interferometer, and an encouraging indication that the target accuracy of 0.1nm could be attained with improvements in both the quality of the sensor and the determination of the interferometer parameters.

8 Discussion

In this thesis, we have demonstrated the use of a novel interferometer for the accurate measurement of the shape of extreme ultraviolet lithography (EUVL) aspheric mirror substrates at spatial frequencies up to $1mm^{-1}$.

Although the accuracy of 4nm reported here for our preliminary measurements falls short of the desired accuracy of 0.1nm, several improvements not yet implemented are likely to increase the accuracy of the instrument to desired levels. The most significant source of error at this time is believed to be the sensor. A sensor designed specifically to meet the requirements of our interferometer was still under development at the time of writing. Consequently, our proof-of-concept measurements were obtained with a significantly inadequate sensor instead. All other subsystems meet the requirements for the desired measurement accuracy.

In developing the interferometer presented here, advances have been made in a number of fields. These developments and suggestions for future work will now be discussed in the same order as they appear in the body of the thesis:

The original concept of the interferometer has been revised to use multiple wavelength interferometry instead of the originally proposed frequency modulation continuous-wave (FMCW) techniques, to gain a significant simplification of the detection subsystem and the added flexibility of using either phase shifting or heterodyne methods to perform our measurements.

The novel nature of the interferometer required the development of a unique inverse propagation algorithm to retrieve the shape of the surface under test from the measurement data. This has been achieved by using a combination of analytic raytracing methods and numerical diffraction methods, based on the idea of boundary diffracted waves, to obtain a good balance between computational speeds and accuracy. A rigorous method for diffraction calculations was developed and used to confirm the accuracy of the hybrid methods used in the inverse propagation algorithm. While the current model is workable, improvements in terms of speed and efficiency are certainly possible.

The light source constructed is capable of providing a stable set of wavelengths which can be used to perform full-field multiple wavelength heterodyne interferometry at three wavelengths simultaneously. In the absence of a suitable sensor, the light source can also be used to perform sequential multiple wavelength interferometry using phase shifting methods. To ensure a minimum of drift and vibration susceptibility, several optical mounting structures were re-designed to be used at lower beam-heights. The custom designed mounts have been shown to outperform commercially available mounts. To achieve the required flexibility and stability for the wavelengths used, tuneable lasers were stabilized relative to a reference laser, using an actively stabilized Fabry-Perot cavity. Direct beating between the stabilized lasers and the reference laser show that the obtained stability far exceeds our requirements. Drifts in the light source could be further reduced by a more careful choice of materials and further miniaturization of the components.

Although a sensor specifically designed to meet the requirements of the interferometer was still under development at the time of writing, significant progress was made with commercially available sensors. A phase-measuring active pixel sensor was used to demonstrate full-field heterodyne interferometry, as well as the beat frequency de-multiplexing of multiple wavelengths, allowing interferometry at several wavelengths simultaneously. The use of a standard CCD sensor in conjunction with phase shifting interferometry was also discussed, and calibration techniques were applied to reduce the influence of a number of systematic errors inherent in this method. This approach proved to be the most accurate of the methods compared using the available sensors. Even in the absence of the desired custom made active pixel sensor, the use of state-of-the-art CCD sensors together with improved calibration techniques may be sufficient to perform measurements to the desired accuracy with standard phase-shifting techniques in the future.

An interferometer frame which allows the stable placement of the fibers, sensor and the surface under test at the required positions without obstructing the beam path has been designed and constructed entirely from invar to ensure optimum immunity against temperature fluctuations. The frame is also designed for accurate re-placement of the fibers and the substrate under test to allow in-process monitoring of the substrate. Future versions of the interferometer frame will need to have a far greater adjustment range to permit the measurement of a variety of substrates. A comprehensive metrology strategy should also be devised which will allow the determination of the relative positions of the interferometer components with greater accuracy.

Measurements on the sphericity of the wavefronts produced by the optical fibers were in agreement with the theoretical models, which predict a deviation from sphericity below +/- 0.06nm over an opening angle of 0.13 rad. By tapering and polishing the fiber ends, a larger effective NA may be achieved, both in terms of the intensity distribution and the wavefront sphericity. This may become necessary in future, for large NA optics.

Preliminary measurements of a test substrate, performed using phase shifting techniques and a standard CCD, yielded promising results. The measurement resulted in a retrieved a mirror shape within 4.2nm rms of the nominal mirror shape, and an even smaller deviation from the mirror shape as measured by conventional interferometry techniques. The three main factors thought to limit the instrument’s accuracy are the non-uniformity of the features on the CCD, the presence of a cover-glass on the sensor and insufficient a-priory knowledge of the relative positions of the interferometer components. All of these factors can be overcome or reduced, leading to measurements with the required accuracy in the foreseeable future. The result will be an ultra-precise metrology instrument suitable for use in optical workshops – both in terms of cost and ease of operation.

Summary

For the past thirty years, microchips have doubled in complexity every two years. This increasing complexity required that the size of the structures written on silicon halve at the same rate. A fundamentally limiting factor to the size of microchip structures is the wavelength of the lithographic projection processes used in their manufacture. Consequently, the wavelengths used to produce microchips have shrunk from 436nm, at the boundary of the visible spectrum, to 193nm, in the ultraviolet, between 1975 and 2002.

A next generation process aims to use light with a wavelength of 13nm, in what is called extreme ultraviolet lithography (EUVL). Unlike previous processes, which could use lens systems to project the required patterns onto the microchips, EUVL requires the use of mirror projection systems. The mirrors used are highly aspheric and must be manufactured with unprecedented accuracies, of the order of 0.1nm. There are currently no affordable and easy to use systems to measure such mirrors with the required accuracy in optical workshops during manufacture.

While the primary push for the development of such a measurement tool has come from the semiconductor industry, there are a number of other technology branches which could benefit from increased reflector accuracies. These include astronomy at ultra-short wavelengths, plasma physics, and biological microscopy.

This thesis describes the construction and use of a novel interferometer to accurately measure the shape of such EUVL mirror substrates. Since tools to measure the surface roughness of these mirrors are readily available, we are concentrating on the measurement of spatial frequencies below $1mm^{-1}$ for the entire mirror surface.

The main advantages of our interferometer over competing instruments are its independence from reference optics, which could introduce significant errors in other types of interferometers, its ability to measure the whole surface of most EUVL optics in one go and the possibility of using a more accurate type of interferometry - heterodyne interferometry - instead of the usual phase stepping. In contrast with a number of methods already available, this instrument is suitable for use in optical workshops, both in terms of cost and ease of use.

Although the accuracy of 4nm reported here for our preliminary measurements falls short of the desired accuracy of 0.1nm, several improvements not yet implemented are likely to improve the accuracy of the instrument to desired levels. The most significant source of error at this time is believed to be the sensor. A sensor designed specifically to meet the requirements of our interferometer was still under development at the time of writing.

In constructing the interferometer presented here, advances have been made in a number of fields:

The novel nature of the interferometer required the development of a unique mathematical tool - an inverse propagation algorithm - to retrieve the shape of the

surface under test from the measurement data. This has been achieved using a combination of analytic raytracing and numerical diffraction methods based on the idea of boundary diffracted waves, to obtain a good balance between computational speed and accuracy. A rigorous method for diffraction calculations was also developed and used to confirm the accuracy of the fast hybrid method finally used in the inverse propagation algorithm.

The light source constructed is capable of providing a stable set of wavelengths which can be used to perform full-field multiple wavelength heterodyne interferometry at three wavelengths simultaneously. In the absence of a suitable sensor, the light source can also be used to perform sequential multiple wavelength interferometry using phase shifting methods. To ensure a minimum susceptibility to drift and vibrations, several optical mounting structures were re-designed from scratch. The custom designed mounts have been shown to outperform commercially available mounts.

Two different types of sensors were tested and compared. A commercially available CCD sensor already allowed us to make measurements coming close to the desired accuracy by using calibration techniques to reduce the influence of a number of systematic error sources. A recently developed sensor, with phase-measuring active pixels, was used to demonstrate new approaches to interferometry: full-field heterodyne interferometry, as well as the beat frequency de-multiplexing of multiple wavelengths, allowing interferometry at several wavelengths simultaneously.

An interferometer frame, which allows the stable placement of the mirror and various other components, has been designed and constructed entirely from invar to ensure optimum immunity against temperature fluctuations.

Theoretical models were used to show that substrates with numerical apertures as large as 0.26 may be measured with the desired accuracy, limited only by the optical fibers used. Measurements carried out on the optical fibers showed that the sphericity of the wave fronts produced was in agreement with our theoretical models. By tapering and polishing the fiber ends, it may be possible to measure optics with even larger numerical apertures.

Preliminary measurements of a test substrate, performed using phase shifting techniques and a standard CCD, yielded promising results. The measurement resulted in a retrieved mirror shape within 4.2nm rms of the nominal mirror shape, and an even smaller deviation from the mirror shape as measured by conventional interferometry techniques. The three main factors thought to limit the instrument's accuracy are the non-uniformity of the features on the CCD, the presence of a cover-glass on the sensor and insufficient a-priori knowledge of the relative positions of the interferometer components. All of these factors can be overcome or reduced, leading to measurements with the required accuracy in the foreseeable future. The result will be an ultra-precise metrology instrument suitable for use in optical workshops.

Samenvatting

Gedurende de afgelopen dertig jaar is de complexiteit van microchips elke twee jaar verdubbeld. Deze toenemende complexiteit eist dat de structuren die op het silicium worden geschreven overeenkomstig kleiner worden. Een fundamentele beperking voor de afmetingen van microchipstructuren wordt gevormd door de golflengte van het lithografische proces waarmee de chips worden vervaardigd. Om deze reden zijn de gebruikte golflengtes voor het maken van microchips tussen 1975 en 2002 afgenomen van 436nm naar 193nm.

Bij een toekomstig fabricageproces – extreem ultraviolet lithografie (EUVL) – wordt gestreefd naar het schrijven met golflengtes van 13nm. In plaats van lenzen, zoals tot nu toe gebruikt in de lithografie, vereist EUVL toepassing van spiegel-projectiesystemen. De daarin toegepaste spiegels zijn sterk asferisch, en moeten met een uitzonderlijke nauwkeurigheid van rond de 0,1nm worden vervaardigd. Er zijn tot dusverre geen eenvoudige en betaalbare meetsystemen die de vorm van zulke spiegels met de vereiste nauwkeurigheid tijdens het fabricageproces in een optische werkplaats kunnen meten.

Terwijl de ontwikkeling van dergelijke meetsystemen voornamelijk van belang is voor de halfgeleiderindustrie, zijn er ook andere sectoren die van een hogere meetnauwkeurigheid kunnen profiteren, waaronder de plasmafysica, biologische microscopie en astronomie bij ultrakorte golflengtes.

Dit proefschrift beschrijft de constructie en het gebruik van een nieuw type interferometer voor het nauwkeurig meten van de vorm van spiegelsubstraten zoals gebruikt in EUVL. Aangezien er al meetsystemen bestaan voor het meten van de oppervlakteruwheid van spiegelsubstraten, concentreren wij ons op het meten van de lage ruimtelijke frequenties tot 1mm^{-1} over het gehele spiegeloppervlak.

Onze interferometer kent meerdere voordelen ten opzichte van concurrerende instrumenten. Hij is onafhankelijk van referentie-optiek die grote fouten in andere interferometers kan introduceren, hij meet het hele oppervlak van de meeste substraten in een keer en kan gebruik maken van een nauwkeurigere interferometrische techniek – heterodyne interferometrie – in plaats van de gewone fase-stap technieken. In tegenstelling tot andere meetinstrumenten is onze interferometer geschikt voor gebruik in een optische werkplaats, zowel qua kosten als op het punt van gebruiksgemak.

Hoewel bij onze inleidende metingen tot nu toe maar een nauwkeurigheid van 4nm is bereikt, in plaats van de gewenste 0,1nm, zijn er meerdere verbeteringen mogelijk die alsnog tot de gewenste nauwkeurigheid kunnen leiden. De sensor introduceert op dit moment vermoedelijk de grootste fouten in onze metingen. Een sensor die voldoet aan de eisen van onze interferometer was tijdens het schrijven van dit proefschrift nog steeds in ontwikkeling in een parallel-project.

Tijdens de constructie van onze interferometer is er op een aantal gebieden vooruitgang geboekt:

De aard van de interferometer vereiste de ontwikkeling van een uniek rekenmodel – een invers propagatie-algoritme – om de vorm van de spiegel uit de meetgegevens terug te vinden. Dit model bestaat uit een combinatie van analytische bundelpropagatie en numerieke diffractiemethoden gebaseerd op het concept van rand-buigingsgolven, om een goede balans tussen rekentijd en nauwkeurigheid te bereiken. Verder werd ook een nauwkeuriger, maar langzamer rekenmodel ontwikkeld om de betrouwbaarheid van onze methode te valideren.

De door ons gebouwde lichtbron maakt het mogelijk om heterodyne interferometrie met drie golflengtes tegelijkertijd uit te voeren. Als daarvoor geen bijbehorende sensor beschikbaar is, kan de bron nog steeds worden gebruikt om voor elke golflengte een reeks fase-stap metingen uit te voeren. Ter vermindering van de gevoeligheid van de opstelling voor vibraties en fluctuaties, zijn er meerdere optische houders ontworpen. De nieuwe houders zijn aanzienlijk stabieler dan commerciële houders.

Twee verschillende types sensoren zijn met elkaar vergeleken. Een standaard CCD was al voldoende om metingen uit te voeren die bijna de gewenste nauwkeurigheid behaalden, waarbij ijkmethodes zijn toegepast om een aantal systematische fouten te reduceren. Een nieuw type sensor, met fase metende pixels, stond heterodyne metingen toe over de gehele sensor, evenals het meten met meerdere golflengtes tegelijkertijd, door het de-multiplexen van de golflengten op basis van hun modulatiefrequenties.

Om een optimale stabiliteit ten opzichte van temperatuurfluctuaties te bereiken, is een behuizing voor de spiegel en de andere componenten van de interferometer ontworpen die volledig is gemaakt van invar.

Rekenmodellen hebben aangetoond dat spiegelsubstraten met numerieke aperturen tot 0,26 gemeten kunnen worden, waarbij de gebruikte optische vezels de beperking vormen. Metingen aan de optische vezels hebben laten zien dat de bolvormigheid van de gegenereerde golffronten in overeenstemming was met het rekenmodel. Door het verkleinen en polijsten van de vezeluiteinden zal het wellicht mogelijk zijn bij nog grotere numerieke aperturen metingen uit te voeren.

Voorlopige metingen van een testsubstraat, uitgevoerd met een standaard CCD en fase-stap techniek, waren veelbelovend. De metingen toonden een verschil van 4,2nm ten opzichte van de nominale spiegel vorm, en zelfs een kleiner verschil ten opzichte van de vorm zoals gemeten met conventionele methodes. De drie belangrijkste oorzaken voor de beperkte nauwkeurigheid van het instrument zijn pixel-positioneerfouten, de aanwezigheid van een dekglas op de sensor en onvoldoende kennis van de relatieve posities van de verschillende onderdelen van de interferometer. Al deze factoren kunnen worden gereduceerd of geëlimineerd, waardoor de gewenste nauwkeurigheid kan worden bereikt. Het resultaat zal een ultra-nauwkeurig meetinstrument zijn, geschikt voor gebruik in een optische werkplaats.

A. Equipment specifications

A.1 HeNe Laser

The HeNe laser used here was an actively stabilized HeNe laser from MellesGriot, model 05-STP-903.

Calibration by the Dutch national institute of standards measured the wavelength of the laser at:

$$(632.991\ 410\ 1 \pm 0.000\ 001\ 0)\ \text{nm}$$

We will briefly list the technical specifications along with some reported performance data:

Model:	05-STP-903	Output mode:	TEM_{00}
Power:	1 mW	Polarization:	1000:1
Power stability:	~1% rms	Noise:	<0.05% rms
Wavelength:	632.991410 nm	Freq. stability (1min/1h/8h):	0.3/0.8/1.2 MHz
Beam diameter:	0.5 mm	Temp. dependence:	0.5 MHz / K
Divergence:	1.8 mrad		

Table A - 1: Specifications of HeNe laser used.

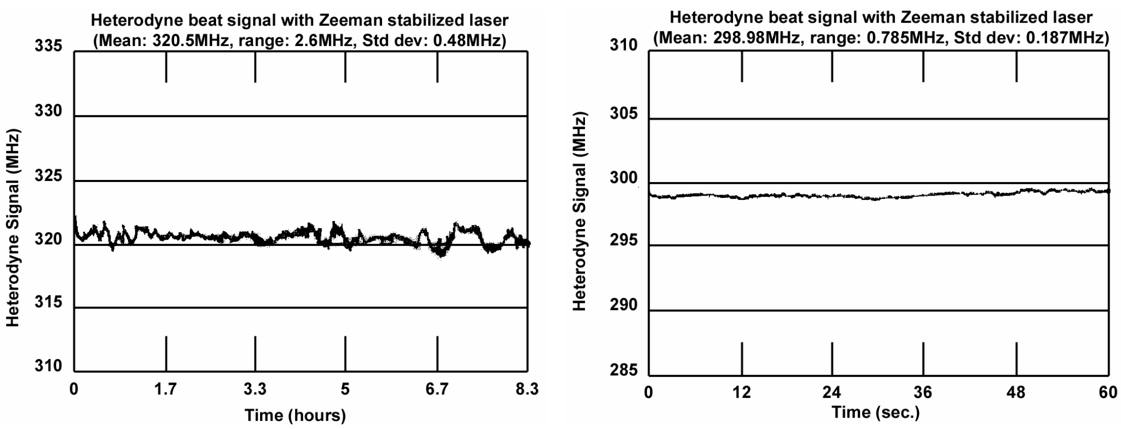

Figure A - 1: Typical frequency stability measurements, relative to a Zeeman stabilized laser, as provided by Melles Griot

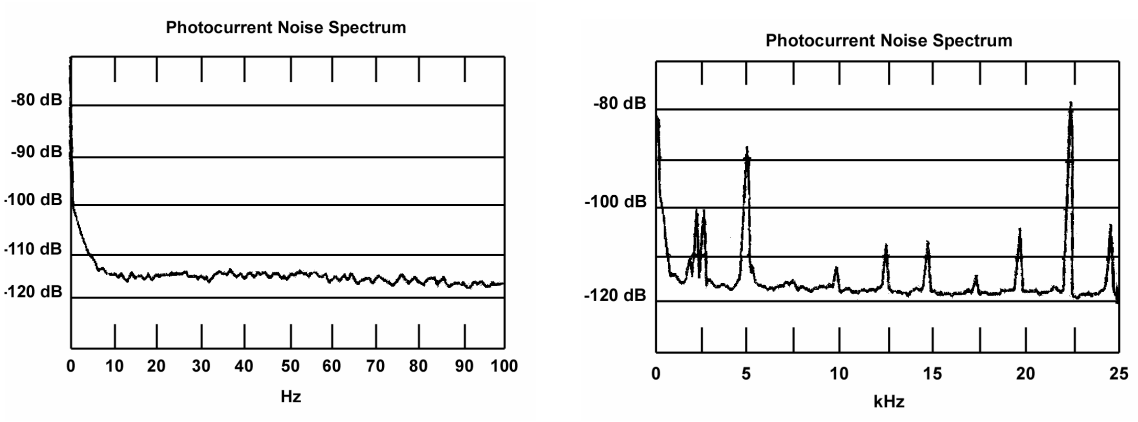

Figure A - 2: Typical photocurrent noise spectrum as provided by MellesGriot. The harmonics of 2.5kHz are probably due to the high voltage amplifier.

A.2 Tunable Laser Newport/EOSI 2010

The primary tunable laser used for most of our experiments was a Newport (formerly EOSI) model 2010 external caviy diode laser based on a Littman-Metcalf design.

Model:	2010M	Power:	5 mW
Wavelength:	635±4 nm	Power stability:	~1% rms
Fine Tuning Range (piezo):	100GHz	Beam shape:	2.5x0.7 mm
Polarization:	1000:1	Divergence:	<1 mrad
Max. modulation rate:	500Hz		

Table A - 2: Specifications of the Newport 2010 laser used.

A.3 Tunable Laser New Focus 6210

Another tuneable laser used for several of our experiments (such as the three-wavelength locking scheme) was a 6210 tuneable external cavity diode laser from New Focus.

Model:	2010M	Power:	4 mW
Wavelength:	633±5 nm	Power stability:	~1% rms
Fine Tuning Range (piezo):	70GHz	Beam shape:	2.5x0.7 mm
Polarization:	1000:1	Divergence:	<1 mrad
Max. modulation rate:	2000Hz		

Table A - 3: Specifications of the New Focus 6210 laser used.

A.4 Polarization maintaining single-mode fiber

A key component in our interferometer is the optical fiber used as a point-source to generate our reference spherical wave. The fibers used here are Newport F-SPV single-mode polarization preserving fibers, manufactured by Fibercore as model HB600.

Model:	F-SPV	Cut-off wavelength:	550±50nm
Index Profile:	Step	Stress-elements:	Bow-tie
N.A:	0.14-0.18	Birefringence:	0.31×10^{-3}
Mode-field diameter:	3.2μm	Ellipticity (typ.)	1.2:1
Cladding diameter:	125±1μm	Core ref. index:	1.4610
Coating dameter:	245±12μm	Cladd. Ref. index:	1.4571

Table A - 4: Specifications of the F-SPV fiber used.

A.5 Custom components.

Stability concerns required the design of new, ultra-low mounting structures for beam-splitting cubes, AOM modules and mirror-mounts. Only the most essential degrees of freedom were allowed to remain. The following drawings were used in the construction of these mounts.

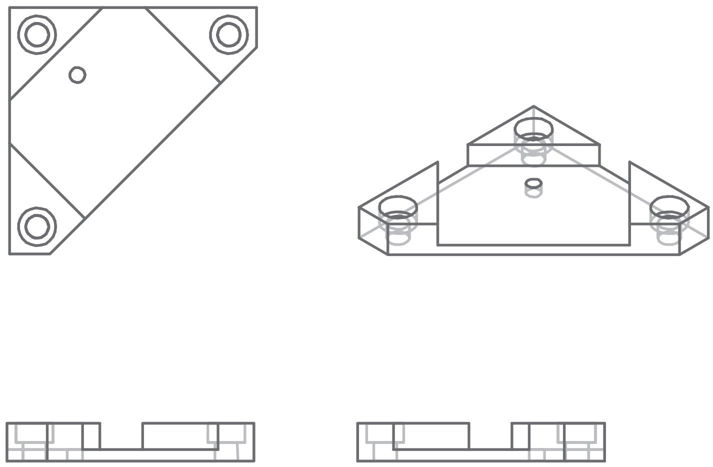

Figure A - 3: Mirror-mount holders for imperial optical table.

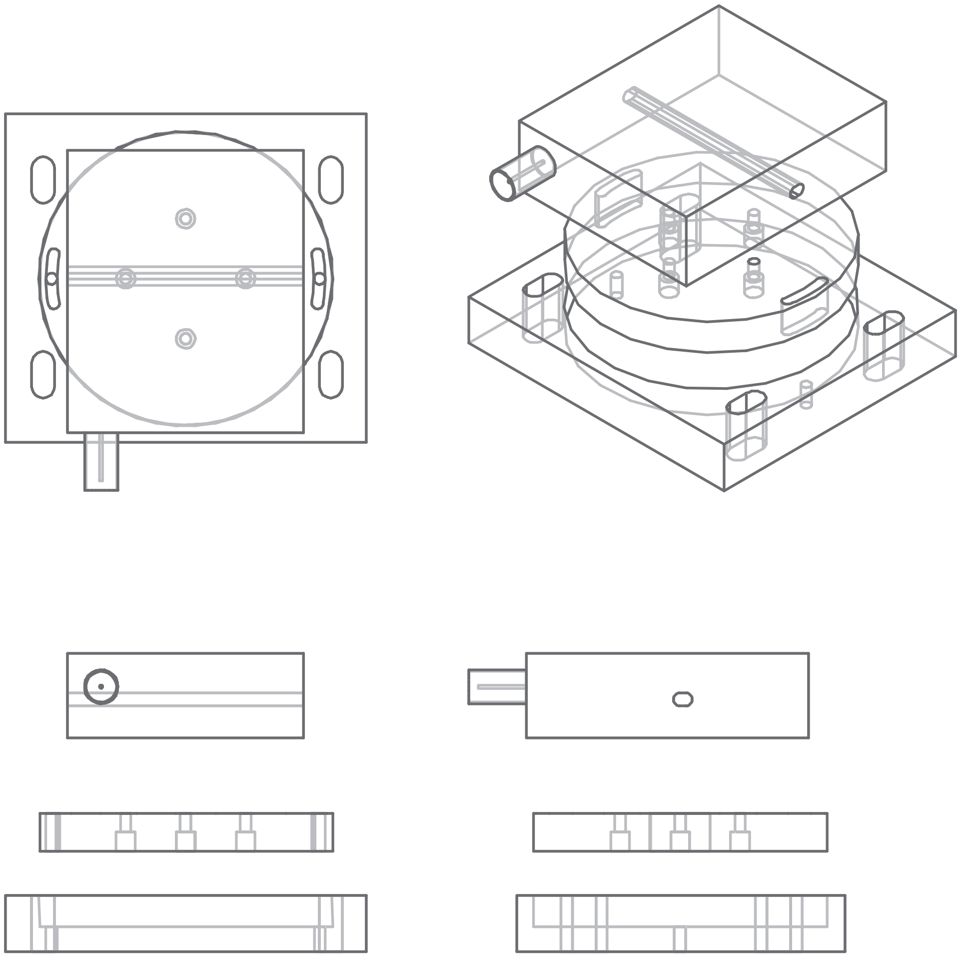

Figure A - 4: Adjustable AOM mounts (incl. AOM module) for imperial optical table.

The beam-splitting cube holders shown above can pivot and rotate about the centre of the steel bearing shown. To fix the top disc in place, three screws are fed through the slits in the disc, and screwed into the threaded holes of the plate below. One of these screws can also contain a spring to provide a flexible counter-force, making the mounts continuously adjustable (see Figure A - 6).

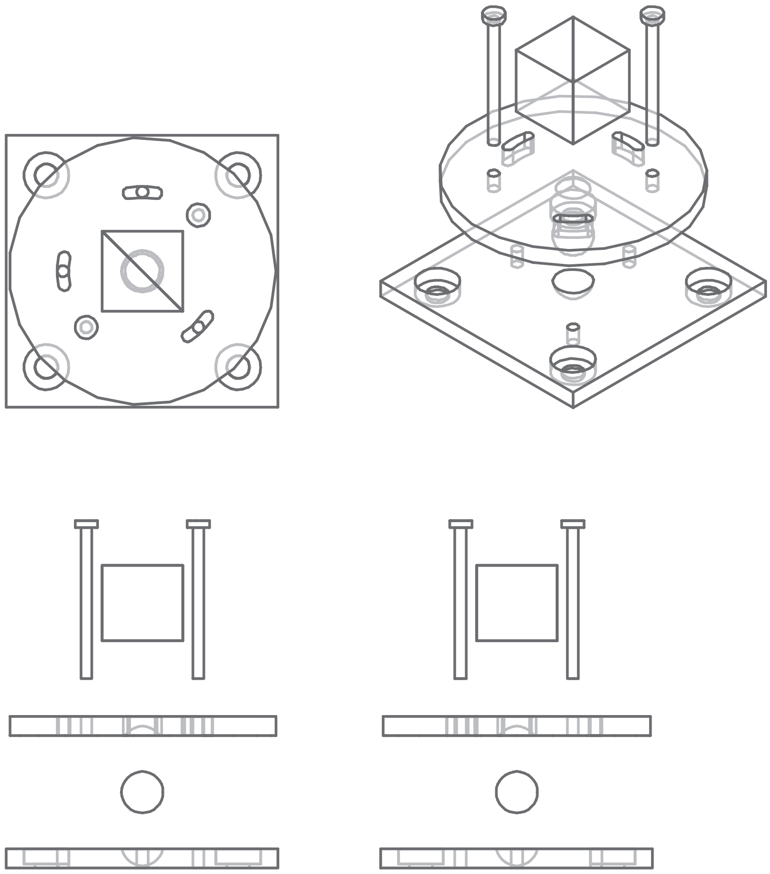

Figure A - 5: Adjustable beam-splitting cube mounts (incl. cube) for imperial optical table.

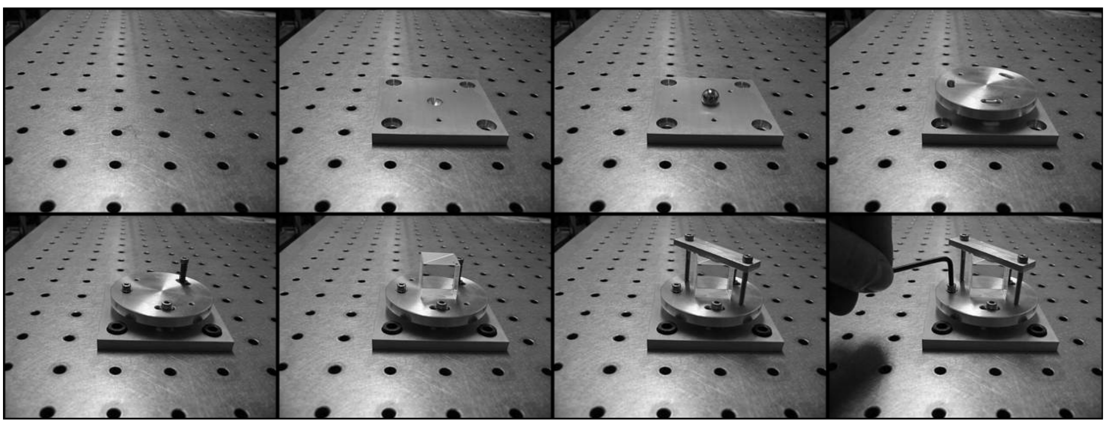

Figure A - 6: Assembly of an adjustable beam-splitting cube mount.

A.6 Lock-in modules.

The lock-in modules used for the wavelength stabilization of the tuneable lasers are LIA-BV-150-H single-board lock-in amplifiers by Femto Messtechnik GMBH. These compact lock-in modules offered all the necessary flexibility needed for our demodulation purposes without the additional overhead of stand-alone lock-in amplifiers. The specifications can be found in Table A - 5.

Model:	LIA-BV-150-H	Reference acquisition time:	<2sec.
Input range:	3μV-1V	Phase shifter resolution:	<1.4° at f<60 kHz
Input noise:	12nV/√Hz	Phase shifter drift:	<100 ppm/K
Input impedance:	1MΩ//4pF	Time constant range:	300μs – 1s
Input gain drift:	100ppm/K	Demodulator dyn. reserve:	15,35 or 55 dB
CMRR:	110dB @ 1kHz	Ouput impedance	50Ω
Pre-filters (high, 6dB/Oct):	2Hz-10kHz	Output DC-stability:	5,50 or 500 ppm/K
Pre-filters (low, 6dB/Oct):	100Hz-1MHz		

Table A - 5: Specifications of the LIA-BV-150-H Lock-in modules used.

A.7 Table and support

The Optical table used was a Newport RPR series, with a set of XL-B vibration isolation legs. The specifications of these components are given in the tables below.

Model:	RPR-48-12	Thickness	106 kg/m^2
Dimensions:	2.4mx1.2mx0.3m	Max. dyn. defl. coeff.:	<2x10^{-3}
Surface tickness:	4.8 mm	Max. Rel. Motion Value*:	<3x10^{-7}mm
Surface flatness:	0.1 mm over 0.36 m^2	Deflection under load$^+$:	<1.3x10^{-3}mm
Surface material:	ferromagnetic steel (w. damping layer)		
Mounting holes:	M6 thread on 1 in. grid (non-standard)		

Table A - 6: Specifications of the optical table used. * - maximum relative motion value derived and confirmed by measurement for a typical table on isolators in a typical laboratory environment with a vibration PSD 10^{-10} g^2/Hz. $^+$ - Measured with a 114kg load at the centre of the table.

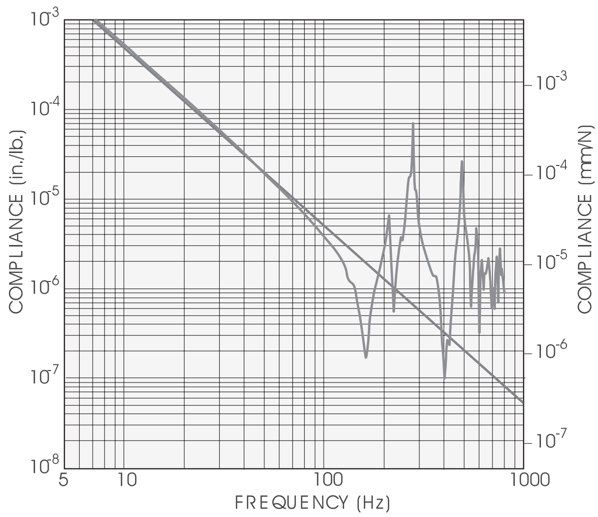

Figure A - 7: Compliance curve for optical table as reported by manufacturer.

Model:	Four Newport XL-B legs	Acceptable load:	700-2700kg
Type:	Pneumatic	Settling time:	~1.5s
Leg height:	0.405m	Max. air pressure:	6.5 kg/cm^2
Active air volume:	0.45m^3		

Table A - 7: Specifications of the vibration isolation legs used.

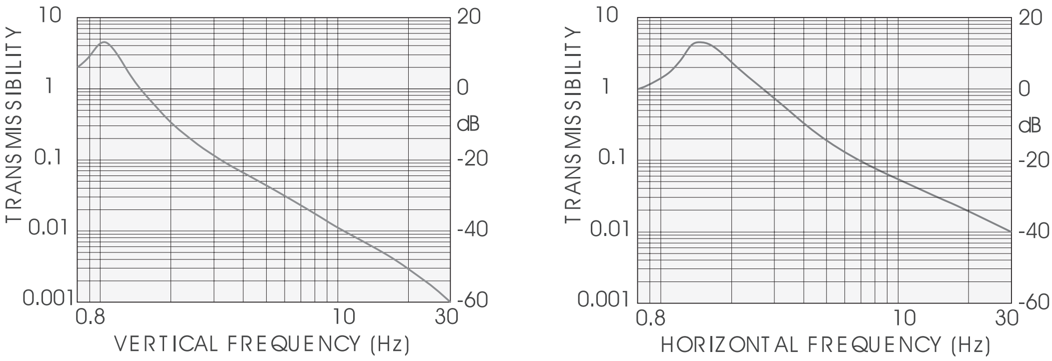

Figure A - 8: Transmissibility curves of isolation legs for vertical and horizontal vibrations, as reported by manufacturer.

A.8 Fabry-Perot cavity.

The Fabry-Perot cavity used for our wavelength stabilization scheme was a refurbished Burleigh RC-150 (See Fig. A – 9). The specifications are listed in the table below.

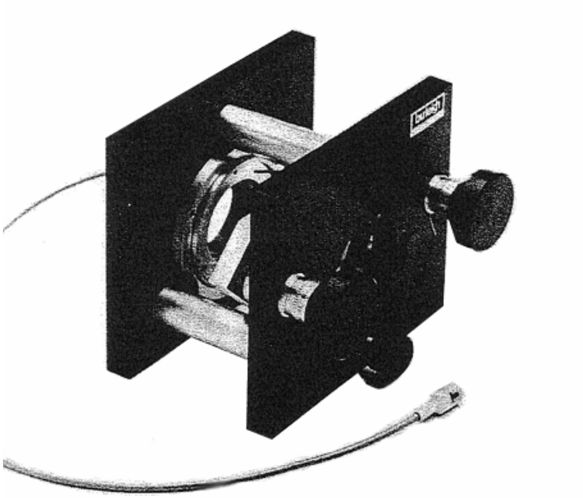

Figure A - 9: Compliance curve for optical table as reported by manufacturer.

Model:	RC-150	PZT range:	1.75μm
Materials used:	Al & super invar	Scan linearity:	<0.1%
Aperture:	50mm	Mirror type:	Flat, 10' wedge
Mirror separation:	0-10mm	Mirror flatness:	λ/100
Adjustments:	Coarse sliding 3-point vernier 3-point PZT	Coatings:	Multi-layer dielectric (front) AR coating (back wedge)

Table A - 8: Specifications of the RC-150 Fabry-Perot.

A.9 Interferometer frame

Some drawings used for the construction of the interferometer frame (not to scale):

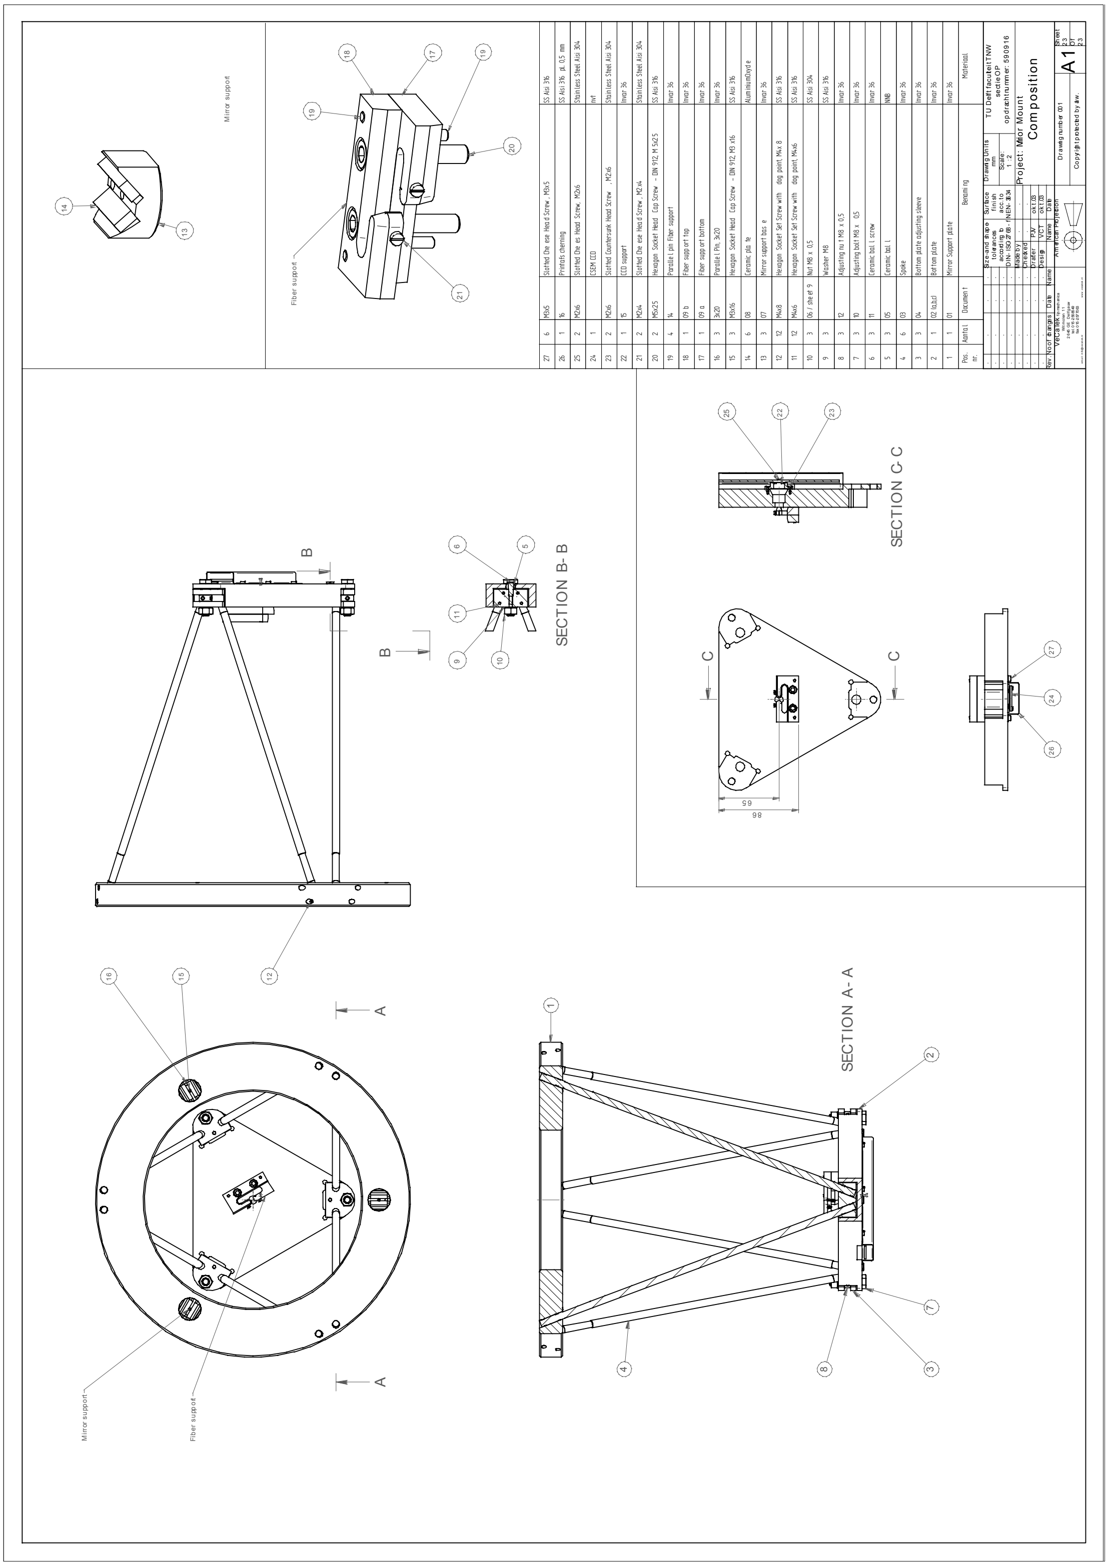

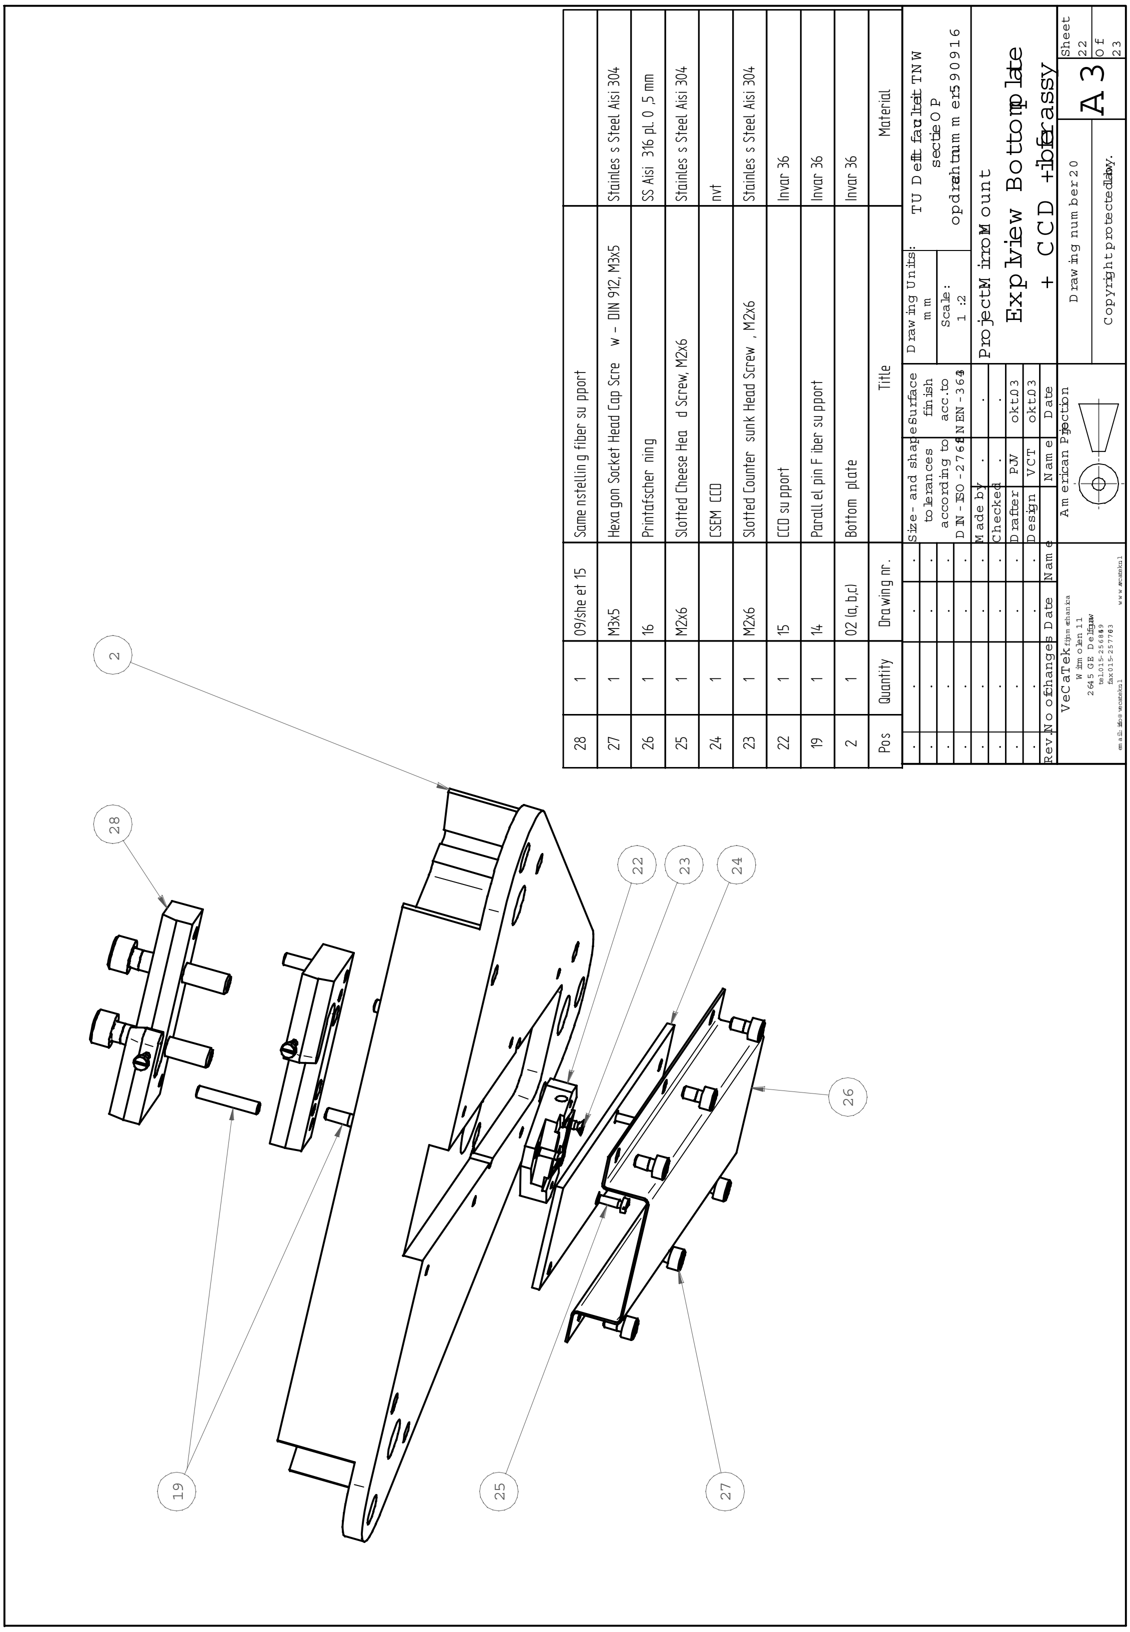

28 1 09/sheet 15 Same nstelling fiber support
27 1 M3x5 Hexagon Socket Head Cap Screw - DIN 912, M3x5 Stainless Steel Aisi 304
26 1 16 Printafscherming SS Aisi 316 pl. 0,5 mm
25 1 M2x6 Slotted Cheese Head Screw, M2x6 Stainless Steel Aisi 304
24 1 CSEM CCD nvt
23 1 M2x6 Slotted Counter sunk Head Screw , M2x6 Stainless Steel Aisi 304
22 1 15 CCD support Invar 36
19 1 14 Parallel pin Fiber support Invar 36
2 1 02 (a,b,c) Bottom plate Invar 36
Pos Quantity Drawing nr. Title Material
Drawing Units: mm
Scale: 1:2
TU Delft faculteit TNW sectie OP opdrachtnummer 590916
Project Mirror Mount
Explview Bottomplate + CCD + fiber assy
A3
Sheet 22 of 23
Drawing number 20
Copyright protected by.
Drafter PJV okt.03
Design VCT okt.03
American Projection
VeCaTek
2
28
22
23
24
26
19
25
27

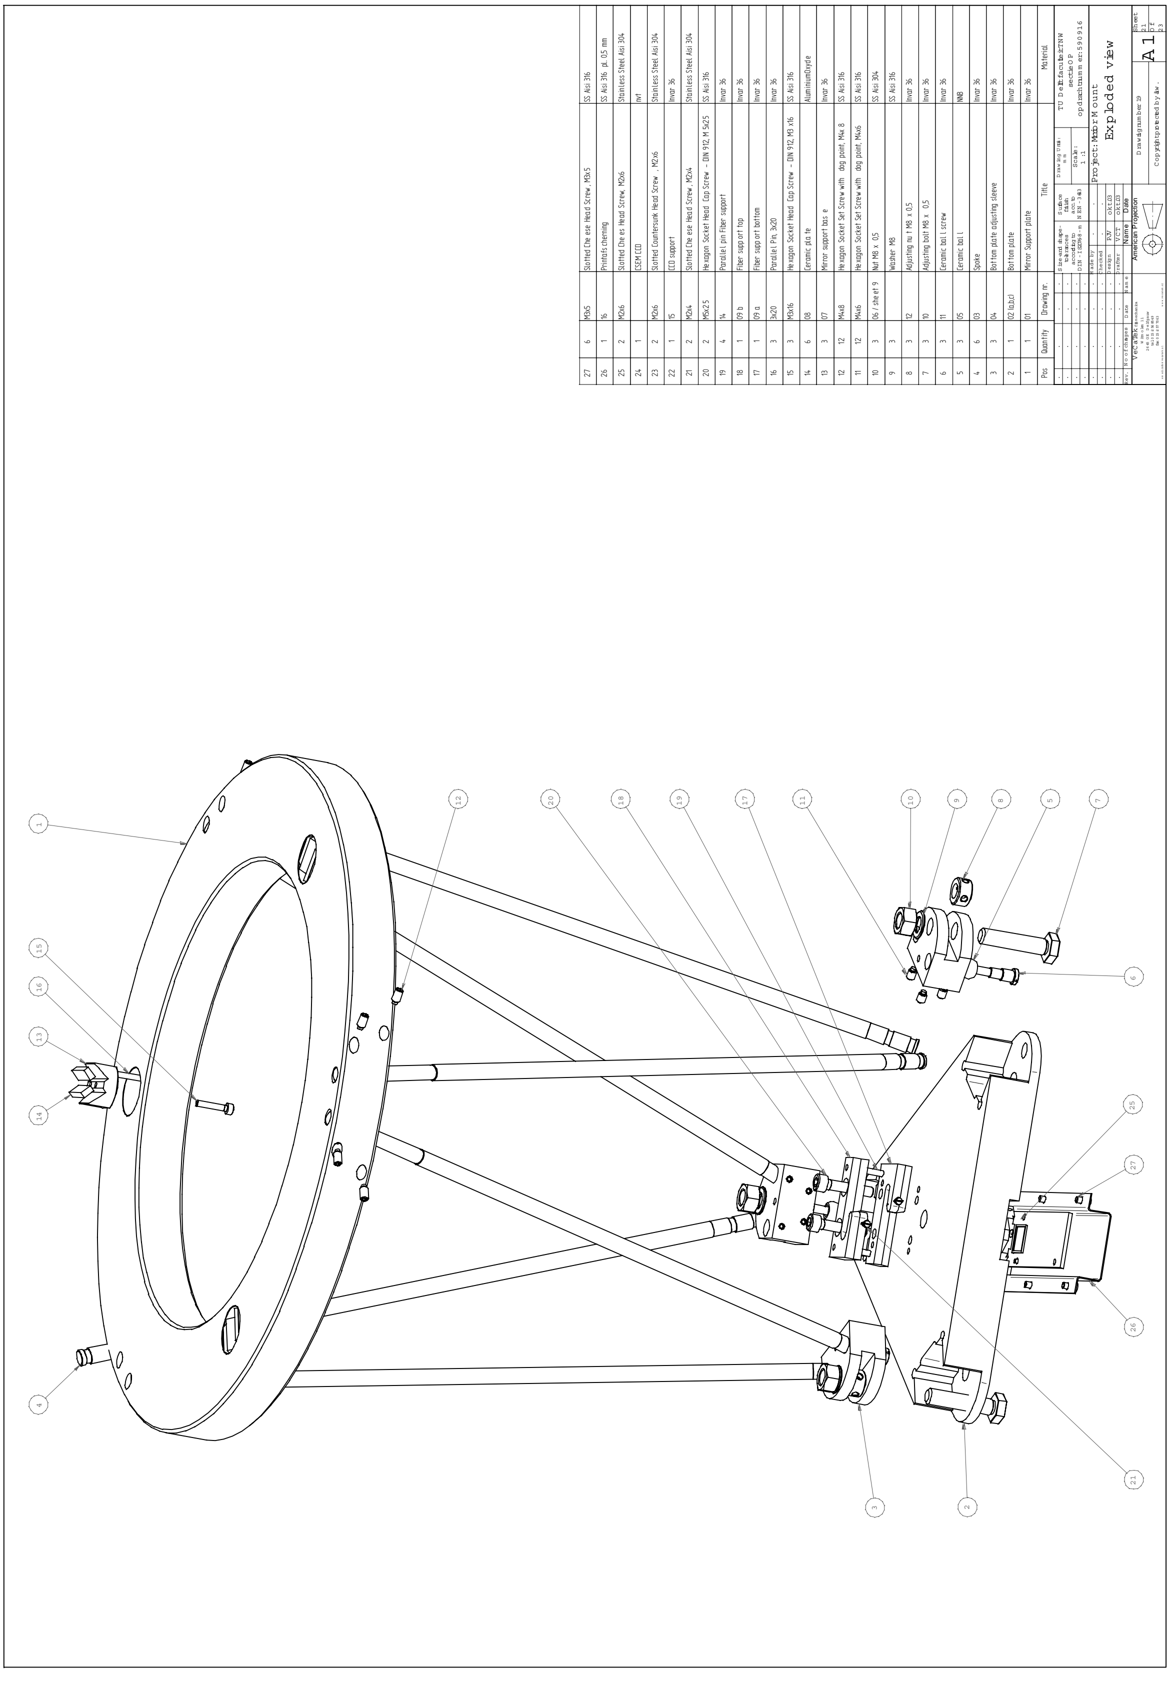

Pos	Quantity	Drawing nr.	Title	Material
27	6	M3x5	Slotted Cheese Head Screw, M3x5	SS Aisi 316
26	1	16	Printafscherming	SS Aisi 316 pl. 0,5 mm
25	2	M2x6	Slotted Chees Head Screw, M2x6	Stainless Steel Aisi 304
24	1		CSEM CCD	nvt
23	2	M2x6	Slotted Countersunk Head Screw, M2x6	Stainless Steel Aisi 304
22	1	15	CCD support	Invar 36
21	2	M2x4	Slotted Cheese Head Screw, M2x4	Stainless Steel Aisi 304
20	2	M5x25	Hexagon Socket Head Cap Screw - DIN 912, M 5x25	SS Aisi 316
19	4	14	Parallel pin Fiber support	Invar 36
18	1	09 b	Fiber support top	Invar 36
17	1	09 a	Fiber support bottom	Invar 36
16	3	3x20	Parallel Pin, 3x20	Invar 36
15	3	M3x16	Hexagon Socket Head Cap Screw - DIN 912, M3 x16	SS Aisi 316
14	6	08	Ceramic plate	AluminiumOxyde
13	3	07	Mirror support base	Invar 36
12	12	M4x8	Hexagon Socket Set Screw with dog point, M4x 8	SS Aisi 316
11	12	M4x6	Hexagon Socket Set Screw with dog point, M4x6	SS Aisi 316
10	3	06 / sheet 9	Nut M8 x 0,5	SS Aisi 304
9	3		Washer M8	SS Aisi 316
8	3	12	Adjusting nut M8 x 0,5	Invar 36
7	3	10	Adjusting bolt M8 x 0,5	Invar 36
6	3	11	Ceramic ball screw	Invar 36
5	3	05	Ceramic ball	NNB
4	6	03	Spoke	Invar 36
3	3	04	Bottom plate adjusting sleeve	Invar 36
2	1	02 (a,b,c)	Bottom plate	Invar 36
1	1	01	Mirror Support plate	Invar 36

Acknowledgements

This project would never have been possible without the help of the dozens of people who have poured hours – in some cases, days – of their valuable time into it.

First and foremost, I would like to thank my supervisor, Joseph Braat, for his support throughout my time here in Delft. His guidance was the perfect mix of encouragement, gentle reminding of deadlines, constructive criticism and trust as far as I am concerned. It is primarily thanks to him that I will look back on my PhD with pride and satisfaction.

René Klaver, whose extensive opus on the theoretical foundations of the envisaged instrument has provided me with invaluable reference material time and time again, needs to be acknowledged as the founding father of the concept of this absolute interferometer. His enthusiasm for the project, despite being bogged down with the odious task of writing it all up at the time we first met, has infected me and lives on to this day.

Roland Horsten and Rob Pols are the dynamic duo, who have turned all the crazy ideas on paper into tangible, working and beautiful pieces of equipment. They never ran out of patience and always took the time to explain the practical aspects of my ideas. I cannot emphasize enough the depth to which I am indebted to them. In addition, people such as Aad v.d. Lingen and Marnix Schnitger provided outstanding support for all things experimental.

The Optics department could not function without the organizational backbone formed by Yvonne van Aalst and Lucia Heijenga. From organizing conference visits to lending me money for the candy machine, they were always there for me.

The countless colleagues and ex-colleagues of the optics department have always played the dual role of professional and personal colleagues. To mention them all by name would doubtlessly unnecessarily strain the patience of the reader, but I cannot avoid highlighting among them Arthur, Oscar and Iciar. The latter two both shared my office, although consecutively, and have kept me sane. All three are excellent people and have become very dear friends.

My parents were not given the opportunity to participate in my life in Delft from close by, having remained in beautiful New Zealand to look after my brother and sister instead. Nonetheless, they have unerringly supported me with phone-calls, e-mails and letters. I am grateful to them for being such wonderful parents and role-models.

Finally Janneke, the love of my life. Together we moved across the world to discover what life was like on the northern hemisphere, and I believe it is fair to say that we succeeded. You have given me love, comfort, company, support and friendship. I could not have done this without you. Thank you for your patience and understanding during the tough times – I promise that this will be my last Ph.D.

Biography

Max Lukas Krieg was born in Jegensdorf, Switzerland, on March 7, 1978. In 1990, he moved together with his parents from Niederurnen, Switzerland, to Warkworth, New Zealand, where he began his secondary education. After graduating as Dux of his school in 1995, he began his tertiary education at Auckland University with a B.Sc., majoring in Physics. In 1998 he received his Bachelor of Science before commencing an M.Sc. on the generation and detection of coherent THz radiation. He graduated with first class honours in March 2000, and moved to The Netherlands with his now fiancée, Janneke Bastiaanssen. He started his research as a Ph.D. student at Delft University of technology in May 2000, on the measurement of the figure of aspherical mirrors, as used in EUV lithography, under the supervision of professor J.J.M. Braat.

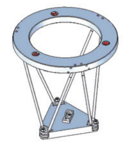
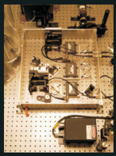
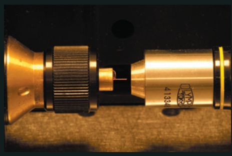
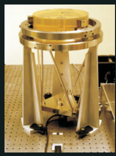
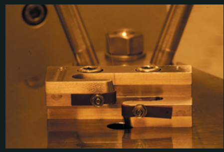
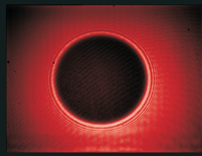
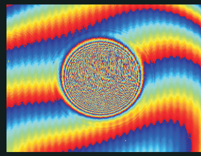

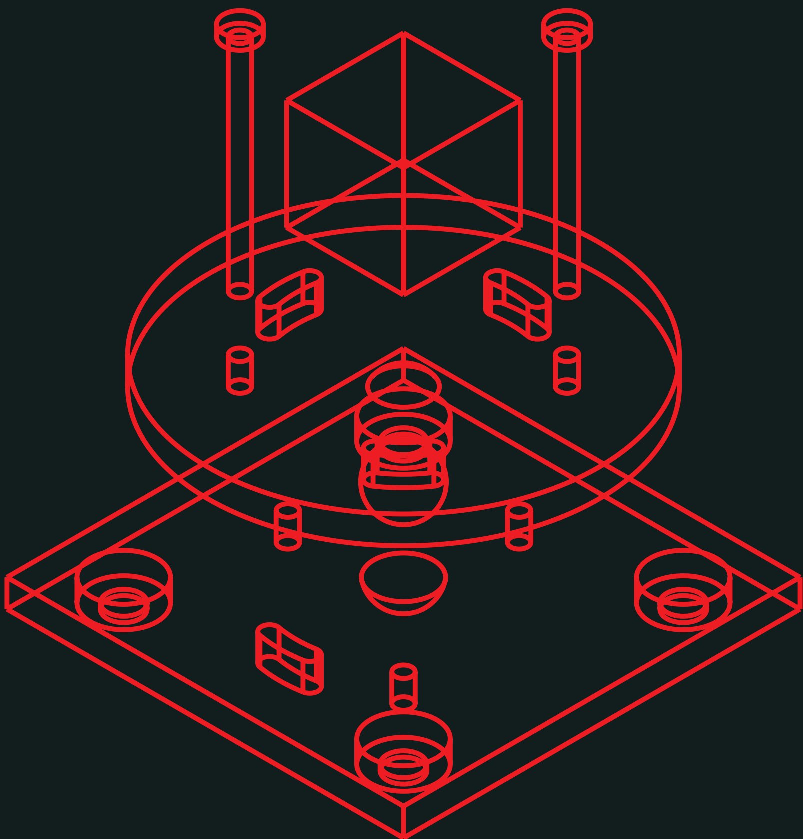

ISBN 90-901-8783-9