Robust Order N Wavelet Filterbanks to Perform 2-D Numerical Integration Directly from Partial Difference or Gradient Measurements

by

Peter John Hampton

Bachelor of Engineering, University of Victoria, 2002
Master of Applied Science, University of Victoria, 2006

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Electrical and Computer Engineering

Robust Order N Wavelet Filterbanks to Perform 2-D Numerical Integration Directly from Partial Difference or Gradient Measurements

by

Peter John Hampton

Bachelor of Engineering, University of Victoria, 2002
Master of Applied Science, University of Victoria, 2006

Supervisory Committee

Dr. Panajotis Agathoklis, Co-Supervisor
(Department of Electrical and Computer Engineering)

Dr. Colin Bradley, Co-Supervisor
(Department of Mechanical Engineering)

Dr. Wu-Sheng Lu, Departmental Member
(Department of Electrical and Computer Engineering)

Dr. Martin Jun, Outside Member
(Department of Mechanical Engineering)

Supervisory Committee

Dr. Panajotis Agathoklis, Co-Supervisor
(Department of Electrical and Computer Engineering)

Dr. Colin Bradley, Co-Supervisor
(Department of Mechanical Engineering)

Dr. Wu-Sheng Lu, Departmental Member
(Department of Electrical and Computer Engineering)

Dr. Martin Jun, Outside Member
(Department of Mechanical Engineering)

Abstract

In this dissertation, a new method for the numerical integration of two-dimensional partial differences is presented. The approach is based on obtaining an estimate of the 2-D Haar wavelet decomposition of the integrated differences by filtering and down-sampling the partial difference measurement data as an intermediate step. Then, this decomposition estimate is synthesized into an estimate of the integrated differences.

The filterbanks required for estimating this decomposition are derived directly from the 2-D Haar Wavelet Analysis Filterbank. The order of operations of this process is manipulated in a novel way so that gradient or partial difference data can be used as input to the filterbank instead of the image data. The original data can then be obtained from this decomposition estimate using unmodified 2-D Haar Wavelet Synthesis Filterbanks. This use of the wavelet decomposition brings a reduction in computation complexity to less than 10 operations per pixel of the result.

This dissertation shows that the data used for this algorithm may be calculated partial differences or discretely sampled gradient data measurements. This data set may have any-sized convex area of support as long as it is on a Cartesian grid. The method is stable as a component of a closed loop system as shown by simulations of a recently developed woofer-tweeter adaptive optics control system.

Table of Contents

List of Tables

List of Figures

Acknowledgements

I would like to thank Dr. Colin Bradley for providing me with the opportunity to develop my research in his cutting edge adaptive optics research laboratory. I thank Dr. Pan Agathoklis for his help in developing the formulas and processes found in this document. I thank both of my supervisors for their patience and advice as we advanced this work from the conceptual stages to published presentations.

I thank Dr. Rodolphe Conan for the use of the wave-front simulation software package that he has developed. I would also like to thank Glen Herriot and Dr. Jean-Pierre Véran of HIA for their instruction on adaptive optics concepts. I also thank my external examiner, Dr. Lisa Poyneer, for her thorough constructive criticism.

I would like to acknowledge financial support from NSERC, CFI and BCKDF.

I would like to thank a few of my past professors. They include Dr. Wu-Sheng Lu, for being the first to instruct me on the use of the Haar wavelet filters, which led to this discovery; Dr. Michael Adams, for imparting his knowledge of efficient implementations of wavelet filterbanks, which led to the low computation cost of the algorithms; and Mr. Cameron Inglis, for teaching me calculus at Aldergrove Community Secondary School, which gave me the confidence to pursue higher education.

To my parents, Bob and Rita Hampton, my sisters, Leigh-Anne Povey and Laura Hampton and my best friends, Patrick McMorris and Andrew Caines, this work wouldn't be complete without your support through the greatest challenge of my life.

To my girlfriend, Sue Lang, your love has given me the resolve to reach this milestone so we may continue to those of our future.

Abbreviations

AO	Adaptive Optics
DM	Deformable Mirror
DWT	Discrete Wavelet Transform
ELT	Extremely Large Telescope
HWAF	Haar Wavelet Analysis Filterbank
HWSF	Haar Wavelet Synthesis Filterbank
MGCG	Multi-Grid Conjugate-Gradient
MS	Mean Square
PSF	Point-Spread Function
RMS	Root Mean Square
SH-WFS	Shack-Hartmann Wave Front Sensor
SNR	Signal to Noise Ratio
SVD	Singular Value Decomposition
TMT	Thirty Meter Telescope
WFS	Wave Front Sensor
W-T	Woofer-Tweeter

Dedication

I dedicate this work, in loving memory, to my grandmother

Anne Hampton

Who peacefully passed from this Earth during the defense of this dissertation

And in loving memory, to my grandfather

Harry Stocker

You will always be an inspiration to my inventive spirit

Chapter 1

Introduction

1.1 Motivation for numerical integration of partial differences

There are at least three sensor types: (i) direct measurements with respect to a common reference, (ii) partial difference (i.e., a relative measurement with respect to a neighbouring grid point) and (iii) gradient, (i.e., a slope measurement at a grid point).

All three sensor types are capable of determining the rate of change either directly or through the application of simple mathematics. However, it is a challenging problem (in more than one dimension) to obtain the equivalent of a direct measurement with respect to a common reference with the second and third sensor types. Integration methods have been developed in the literature and they allow partial difference or gradient sensors to provide an estimate of direct measurements with respect to a common reference, rather than with respect to a neighbouring point.

A short list of uses for gradient sensors include measuring: (i) wave fronts (Widiker, Harris and Duncan 2006), (ii) magnetic fields (Miller 1965), (iii) flow-induced temperature gradient (van Oudheusden 1989), or (iv) determining cracks in stress experiments on optical materials (Tippur, Krishnaswamy and Rosakis 1991). For these sensors, integration of the gradient measurements is the only means to obtain an estimate of the discretely sampled field.

Material in this chapter has been published as:
Hampton, P. J., P. Agathoklis, and C. Bradley, "Wave Front Reconstruction over a Circular Aperture using Gradient data Extrapolated via the Mirror Equations." *Applied Optics* 48, (July 2009), 4018-4030.

1.1.1 Adaptive optics

In an adaptive optics (AO) system, wave front sensors (WFS) are used to detect the aberrations created in a plane wave caused by the atmosphere. Deformable mirrors (DMs) (Cornelissen, Bierden and Bifano 2006) are used to correct these aberrations. The sensors used in AO applications are:

(i) The Shack-Hartman WFS (SH-WFS) inspired by (Hartmann 1907) measures the angle of arrival of the wave front at a sample point. They have been developed over the years to now use off-the-shelf components (Widiker, Harris and Duncan 2006);

(ii) The pyramid WFS (Ragazzoni 1996) approach, which calculates the partial differences of a wave front from four pupil plane images; and

(iii) The curvature WFS (Roddier and Roddier 1993), which obtains images on either side of the focal plane to obtain a measure of the combination of the curvature (Laplacian operator) and the wave front gradient on expanding radial lines at the edge of the pupil.

The third approach requires the use of Poisson solvers to obtain the wave front, not numerical integration. The SH-WFS and pyramid WFS are gradient sensors. Integration can be used for these sensors.

Alternate uses of AO are in ophthalmology (Fernández, Iglesias and Artal 2001) and laser communication (Tyson 1996). The discussions on AO in this dissertation pertain to the astronomical application.

The method of wave front reconstruction from gradient measurements has been considered extensively in the literature. Several techniques have been proposed to obtain

such reconstructions and are described in Chapter 2. Discussion of models of the specific components and control algorithms are left to Chapter 5.

1.1.2 Tomography

One current research area for advanced AO systems is the use of tomography (Tokovinin and Viard 2001) to reconstruct the volume of turbulence above the telescope from several SH-WFSs. An intermediate step in this process is to integrate the gradient measurements of the SH-WFSs into wave front phase (D. Gavel 2004).

1.1.3 Gradient camera

Future applications in digital photography will require an image to be reconstructed from gradient measurement data. For example, this occurs with the camera proposed in (Tumblin, Agrawal and Raskar 2005), where static gradients are measured (instead of static intensities) and these gradients are used to reconstruct the image. The intent of this system is to reduce saturation of the data by recording the local difference between two pixels rather than the intensity at each pixel. The method proposed in (Tumblin, Agrawal and Raskar 2005) for image reconstruction is based on using a Poisson solver to convert the gradient data to an estimate of the intensity image. One approach to implementing a Poisson solver is to begin with the Dirichlet boundary conditions that require measurement of the image on the boundary. The method chosen in (Tumblin, Agrawal and Raskar 2005) instead solves the boundary condition intensities by adding the given boundary gradient measurement to the estimate of the neighbouring point with a starting condition of 0 for the first boundary pixel. This is 1-D integration. Then it iteratively solves the interior points of the image.

1.1.4 Computer vision

The 3-D computer vision system described in (Kovesi 2005) directly measures the surface normal, at each point, of the object. This data may be used directly or it may be converted to 2-D gradients. This allows for 2-D integration methods to obtain the depth of the 3-D object at all points of the 2-D grid.

1.1.5 Special effects and photo editing

It has been shown in the literature that there are benefits to merging pictures in the gradient domain because the errors are more difficult to detect visually (Peréz, Gangnet and Blake 2003). This paper demonstrates several special effects, such as graffiti removal/creation, face swapping and wrinkle removal from faces.

Another approach is to use the gradient space to compress the dynamic range of an image (Fattal, Lischinski and Werman 2002). This method is used to show detail that would otherwise be lost in bright or dark regions of an image.

The algorithms are performed on gradient data and iterative Poisson methods of Section 2.3 are used to return to the image domain. These processes may benefit from the approach in Chapter 3.

1.1.6 Scientific modelling

The numerical integration of partial differences is a mathematical concept and its applications are as limitless as differentiable functions. The Poisson equation is used in modelling complex phenomena that occur in nuclear power plants, such as neutron transport, fluid flow and structural behaviour (Tanaka 2006). Scientific applications, including these, often use more than two dimensions. Therefore it is an area of future work to develop integration algorithms for higher numbers of dimensions.

1.2 Contribution of thesis

This dissertation presents a new method for the numerical integration of two-dimensional partial differences. The approach is based on obtaining an estimate of the 2-D Haar wavelet decomposition of the integrated differences by filtering and down-sampling the partial difference measurement data as an intermediate step. This decomposition estimate is then synthesized into an estimate of the integrated differences.

The filterbanks required for this process are derived directly from the 2-D Haar Wavelet Analysis Filterbank (HWAF). The order of operations of this process is manipulated in a novel way so that either gradient or partial difference data can be used as input to the filterbank instead of direct measurements. The integrated data can then be obtained from this decomposition estimate using unmodified 2-D Haar Wavelet Synthesis Filterbanks (HWSF). This use of the wavelet decomposition brings a reduction in computational complexity to less than 10 operations per pixel of the result.

This dissertation shows that the data used for this algorithm may be calculated partial differences or discretely sampled gradient data measurements. This data set may have any sized convex area of support as long as it is on a Cartesian grid.

When there is no measurement noise, the only errors are from shapes that are invisible to the sensor (i.e., the constant offset and a shape resembling a checkerboard at the highest pixel resolution). When noise is present, the noise is significantly attenuated. The method is stable as a component of a closed loop system when measurement noise and sensor modelling errors are present, as shown by simulations of a recently developed Woofer-Tweeter adaptive optics control system.

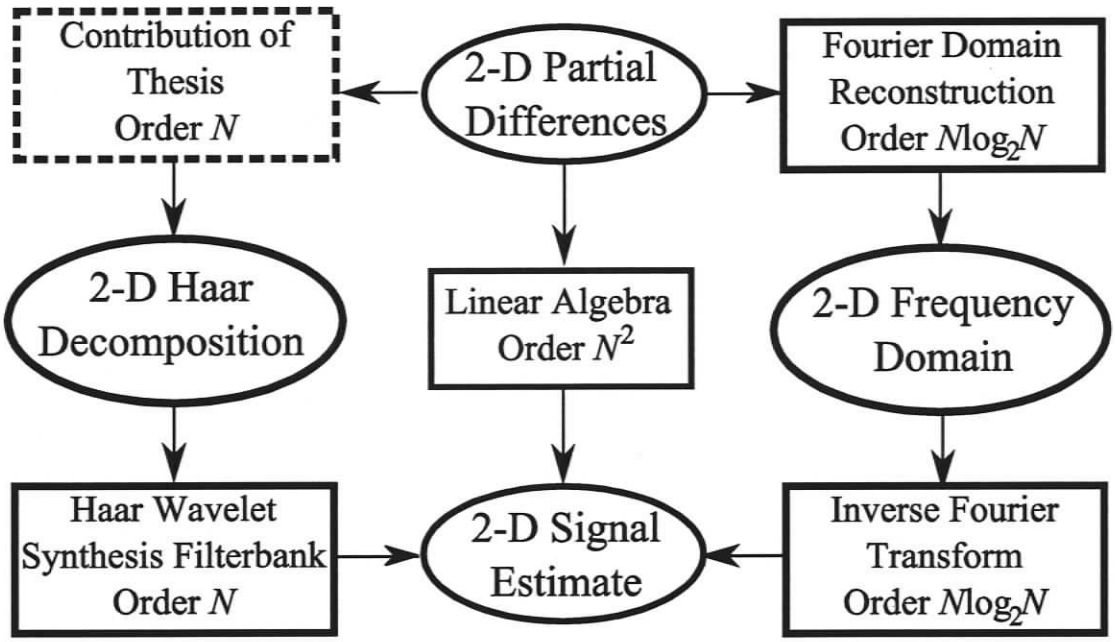

Figure 1.1: Primary contribution of thesis. This figure shows some solution paths to determining a 2-D signal from its respective 2-D partial differences. The 'Order' is in reference to how quickly the required computations increase as the data size, N, increases. The new path is fastest.

1.3 Dissertation outline

This dissertation is organized as follows:

Chapter 1 is an overview of the applications of the integration of partial differences. Applications range from the qualitative (special effects and cameras) to the quantitative (tomography, adaptive optics and other scientific applications). The contribution to numerical integration of this dissertation is discussed.

Chapter 2 shows the continuous gradient model and discrete partial difference operators. It describes integration methods and Poisson equation solvers developed in the literature. The iterative Poisson solver converges on the estimate of the integral of first order gradients using the second order Laplacian operator. This is known to converge to the least squares optimal integration. Other methods directly integrate the first-order gradients directly.

Chapter 3 develops the wavelet based integration method. This method was originally published as (Hampton, Agathoklis and Bradley, A new wave-front reconstruction method for adaptive optics systems using wavelets 2008). It is shown that the method provided is O(N) and attenuates the noise in the partial difference data. This chapter includes the integration algorithm and its proof, which is based primarily on the reordering of linear filter operations. An efficient implementation of the filterbank structure is also provided.

Chapter 4 presents the approach when only circular data sets are available as published in (Hampton, Agathoklis and Bradley, Wave Front Reconstruction over a Circular Aperture using Gradient data Extrapolated via the Mirror Equations 2009). It is shown that the gradient data may be extrapolated using the mirror equations or by forcing the gradient to be zero along the diagonal lines. These allow the integration method developed in Chapter 3 to be applied to any sized convex data set.

Chapter 5 uses the material of Chapters 3 and 4 and demonstrates its robustness on an Adaptive Optics (AO) simulation. A proposed technique of controlling multiple deformable mirrors is shown that takes advantage of the multi-resolution nature of wavelet-based integration. This chapter shows that there can be performance gains when the speed of this algorithm is used to reduce the number of delays in an AO system.

Chapter 6 summarizes the future work and conclusions. The primary areas of future work are to increase the number of dimensions and the quality of the integrations. Improvements and implications for AO systems are also provided.

Appendix A was originally published as (Hampton et al. 2008). It is provided here to support the assertions of the dynamics of the control system made in Chapter 5 and as a resource for measuring the dynamics of a completed AO system.

Appendix B is the proof of concept of an AO Woofer-Tweeter (W-T) control system for the Thirty Meter Telescope (TMT). This was also presented and published as (Hampton et al. 2006). It provides an explanation of the use of the Singular Value Decomposition (SVD) (Watkins 2002) for the application of separating error signals in a W-T AO system and a possible control system architecture.

Chapter 2

Gradient Operators and Integration Methods

2.1 Abstract

The gradient models are defined in the continuous and discrete domains in this chapter. This is followed by a literature review of methods to integrate gradient data. The equations of the integration methods are provided when the process is composed of only a few steps. For more complicated approaches, the reader is directed to the respective publications that explain the techniques.

2.2 Gradient of continuous functions

Consider a continuous function represented by $\Phi(x,y)$. The continuous gradient of this function is represented as shown in (2.1).

$$\begin{aligned}\nabla\Phi &= \frac{\partial\Phi}{\partial x}\hat{u}_x + \frac{\partial\Phi}{\partial y}\hat{u}_y \\ &= \dot{\Phi}_x\hat{u}_x + \dot{\Phi}_y\hat{u}_y\end{aligned} \tag{2.1}$$

where x and y are the horizontal and vertical directions and $\hat{u}$ represents a unit vector.

A gradient field describes the spatial rate of change of the continuous function, Φ, and has some important properties that have to be satisfied if gradient measurement data

---

Material in this chapter has been published as:

Hampton, P. J., P. Agathoklis, and C. Bradley. "Wave Front Reconstruction over a Circular Aperture using Gradient data Extrapolated via the Mirror Equations." *Applied Optics* 48, (July 2009): 4018–4030.

is extrapolated beyond the measured region. A critical property of such a gradient field is that the integral of the gradient, along a path, is equal to the difference in intensity between the end point and the starting point of Φ. This extends to the rule that the closed path integral of gradients (i.e. the curl) must equal zero, since the start and end points are equal. There are effectively infinite possibilities to extend a gradient field so that these conditions are satisfied.

Extrapolation is necessary for this algorithm when the provided gradient data set does not fill $2^M - 1 \times 2^M - 1$, such as in AO where telescopes provide data on a circle. Solutions of the extrapolation process are discussed in Chapter 4.

2.2.1 Fried discrete model of partial differences

The discretization of (2.1) has been modelled by Fried (1977) in the following manner. Let φ be an element of matrix Φ representing the discretely sampled continuous function and $\dot{\varphi}_x$ and $\dot{\varphi}_y$ be the measured gradient elements in the x and y directions respectively. The Fried geometry used in modelling the gradient is shown in Figure 2.1. The arrows indicate the directions of the gradients and the gradient measurement point is where they cross. This shows that the measurement point for $\dot{\varphi}_x$ and $\dot{\varphi}_y$ is in the center of the four discrete pixels. It can be seen from this figure that the grid points of φ are not the same as those for $\dot{\varphi}_x$ and $\dot{\varphi}_y$.

It follows from this discrete model (Fried 1977) that the average change in the x direction is (2.2) and the average change in the y direction is (2.3).

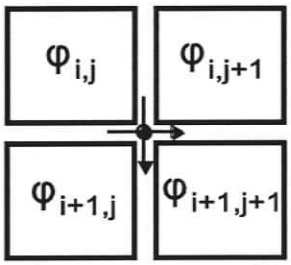

Figure 2.1: Fried alignment of gradient measurements. Phase pixels represented by squares. Measurement is modelled as the discrete gradient of the phase pixels in each of the two Cartesian directions. Measurement is at the crossing of the two arrows.

$$\dot{\varphi}_{x,p,q} = 0.5\left(-\varphi_{i,j} + \varphi_{i,j+1} - \varphi_{i+1,j} + \varphi_{i+1,j+1}\right) \tag{2.2}$$

$$\dot{\varphi}_{y,p,q} = 0.5\left(-\varphi_{i,j} - \varphi_{i,j+1} + \varphi_{i+1,j} + \varphi_{i+1,j+1}\right) \tag{2.3}$$

where the elements of matrices Φ, $\dot{\Phi}_x$ and $\dot{\Phi}_y$ are φ, $\dot{\varphi}_x$ and $\dot{\varphi}_y$ respectively. Assuming a grid spacing of 1 unit, the indices representing the two grids, (i,j) for Φ and (p,q) for the gradients, are related by

$$i = p - 0.5 \tag{2.4}$$

$$j = q - 0.5 \tag{2.5}$$

indicating that the grid of the wave-front elements, φ, is offset by half the pitch from the measurements of the x and y direction slopes, $\dot{\varphi}_x$ and $\dot{\varphi}_y$, as shown in Figure 2.1. The positive vertical direction is downward and the positive horizontal direction is to the right.

An alternate notation (used in Chapter 3) is to denote the $2^M \times 2^M$ matrix representing the discretely sampled continuous function as ${}^M\Phi$, where M is a positive integer. The discrete gradient measurements are then

$${}_F\tilde{x}_{p,q} = 0.5\left(-\varphi_{i,j} + \varphi_{i,j+1} - \varphi_{i+1,j} + \varphi_{i+1,j+1}\right) = \dot{\varphi}_{x,p,q} \tag{2.6}$$

$${}_F\tilde{y}_{p,q} = 0.5\left(-\varphi_{i,j} - \varphi_{i,j+1} + \varphi_{i+1,j} + \varphi_{i+1,j+1}\right) = \dot{\varphi}_{y,p,q} \tag{2.7}$$

The horizontal and vertical gradients of an image ${}^{M}\Phi$, using the Fried geometry, can then be represented as matrices ${}^{M}_{F}\tilde{X}$ and ${}^{M}_{F}\tilde{Y}$ respectively and will be of size $\left(2^{M}-1\right)\times\left(2^{M}-1\right)$.

2.2.2 Hudgin discrete model of partial differences

An alternate alignment was developed by Hudgin (1977) for the case where the sensor is aligned between two pixels instead of four.

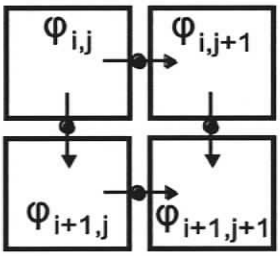

Figure 2.2: Hudgin alignment of gradient measurements. Phase pixels represented by squares. Measurement is modelled as the discrete gradient of the phase pixels in each of the two Cartesian directions. Measurement is at the dot in the center of the arrows.

Using the Hudgin geometry (Hudgin 1977) in Figure 2.2, the difference between two points is sensed and the gradient data is:

$${}_{H}\tilde{x}_{i,q} = -\varphi_{i,j} + \varphi_{i,j+1} \tag{2.8}$$

$${}_{H}\tilde{y}_{p,j} = -\varphi_{i,j} + \varphi_{i+1,j} \tag{2.9}$$

Thus, the horizontal and vertical gradients of an image ${}^{M}\Phi$ of size $2^{M}\times 2^{M}$, using the Hudgin geometry can be represented as matrices ${}^{M}_{H}\tilde{X}$ and ${}^{M}_{H}\tilde{Y}$ respectively and will be of size $2^{M}\times\left(2^{M}-1\right)$ and $\left(2^{M}-1\right)\times 2^{M}$, respectively. This non-square structure is the result of filtering ${}^{M}\Phi$ in only one direction for each of the gradient directions.

2.3 Poisson solver

The reconstruction of Φ based on the gradient data is known to be convertible to a Poisson equation (Herrmann 1980), which equates combinations of the second order partial derivatives. A method used by (Tumblin, Agrawal and Raskar 2005) and (Vogel and Yang 2006) is an iterative solution of equating the Laplacian, ∇^2, of Φ with the divergence, $\nabla\bullet$, of the gradient of Φ (Fattal, Lischinski and Werman 2002).

$$\nabla^2\Phi = \nabla\bullet\left(\frac{\partial\Phi}{\partial x}\bar{u}_x + \frac{\partial\Phi}{\partial y}\bar{u}_y\right) \tag{2.10}$$

It is presented by Herrmann (1980) that when noise is present on the gradient measurements, the least-square solution is obtained when equality is achieved in (2.11).

$$\frac{\partial^2\Phi_{estimate}}{\partial x^2} + \frac{\partial^2\Phi_{estimate}}{\partial y^2} = \frac{\partial\left(\dot{\Phi}_x + \eta_x\right)}{\partial x} + \frac{\partial\left(\dot{\Phi}_y + \eta_y\right)}{\partial y} \tag{2.11}$$

where η is the noise. Assuming a Fried alignment, this is discretized similarly to (Vogel and Yang 2006) by using 2-D filters. The values of the filter coefficients are chosen here to be correct according to the definition of the discrete gradient model of (Fried 1977). The Laplacian operator for Hudgin alignment is shown in (Ellerbroek 2002) and (Vogel and Yang 2006).

$$\begin{bmatrix} 1 & 0 & 1 \\ 0 & -4 & 0 \\ 1 & 0 & 1 \end{bmatrix} \otimes \Phi_{estimate} = \begin{bmatrix} -1 & 1 \\ -1 & 1 \end{bmatrix} \otimes \tilde{X}_F + \begin{bmatrix} -1 & -1 \\ 1 & 1 \end{bmatrix} \otimes \tilde{Y}_F \tag{2.12}$$

An iterative solution then becomes

$$\Phi_{estimate}[k+1] = \Phi_{estimate}[k] + g\left(\begin{bmatrix} 1 & 0 & 1 \\ 0 & -4 & 0 \\ 1 & 0 & 1 \end{bmatrix} \otimes \Phi_{estimate}[k] - \begin{bmatrix} -1 & 1 \\ -1 & 1 \end{bmatrix} \otimes \tilde{X}_F - \begin{bmatrix} -1 & -1 \\ 1 & 1 \end{bmatrix} \otimes \tilde{Y}_F\right) \tag{2.13}$$

where g is a gain term. This is stable for $0 < g \leq 0.25$ as determined by trial and error. For a 64×64 data set, the above converges to visually acceptable results in hundreds of iterations. It converges to numerically accurate results after thousands of iterations with $g = 0.25$, which is the same as using the Jacobi method described in (Watkins 2002). It was shown in (Hockney 1965) that the number of iterations can be reduced to approximately 64 for this sized data. Therefore, it is still a great benefit to this approach if an initial approximation can be obtained quickly prior to this slow convergence to the least-squares optimal result.

2.4 1-D numerical integration

There is a very simple method of integrating Hudgin aligned difference data. Consider the following 4-element example.

$$\vec{\varphi} = \begin{bmatrix} \varphi_0 & \varphi_1 & \varphi_2 & \varphi_3 \end{bmatrix} \tag{2.14}$$

$$\begin{aligned} {}_H\vec{x} &= \begin{bmatrix} {}_H x_0 & {}_H x_1 & {}_H x_2 \end{bmatrix} \\ &= \begin{bmatrix} \varphi_1 - \varphi_0 & \varphi_2 - \varphi_1 & \varphi_3 - \varphi_2 \end{bmatrix} \end{aligned} \tag{2.15}$$

If ${}_H\vec{x}$ is known and $\vec{\varphi}$ is unknown then we have 3 knowns and 4 unknowns. It is well known that the mean value is lost when the only data is the difference measurements. Therefore one may arbitrarily set $\varphi_0 = 0$. This results in

$$\tilde{\varphi}_K = \sum_{k=0}^{K-1} {}_H x_k \tag{2.16}$$

$$\begin{aligned} \vec{\varphi}_{estimate} &= \begin{bmatrix} \tilde{\varphi}_0 & \tilde{\varphi}_1 & \tilde{\varphi}_2 & \tilde{\varphi}_3 \end{bmatrix} \\ &= \begin{bmatrix} 0 & {}_H x_0 & {}_H x_0 + {}_H x_1 & {}_H x_0 + {}_H x_1 + {}_H x_2 \end{bmatrix} \end{aligned} \tag{2.17}$$

In the absence of noise this results in perfect reconstructions except for the error of the initial point, even when applied to 2-D data sets. However, noise is greatly amplified

by this process since any estimated value of the integration includes the running sum of the noise, η.

$$\tilde{\varphi}_K = \varphi_K - \varphi_0 + \sum_{k=0}^{K-1} \eta_k \tag{2.18}$$

The result is the standard deviation of the noise on a particular element of the integral is amplified by the square root of its index, *K*. This approach is unacceptable for integration of realistic measurements (Kovesi 2005). Chapter 3 presents a method for 2-D integration that attenuates noise instead.

2.5 Pseudo-inversion

A common approach is to represent the Cartesian gradient measurements as vectors, $\vec{\dot{\varphi}}_x$ and $\vec{\dot{\varphi}}_y$, and use linear algebra in the following equation (Ellerbroek 2002):

$$\begin{bmatrix} \vec{\dot{\varphi}}_x \\ \vec{\dot{\varphi}}_y \end{bmatrix} = G\vec{\varphi} + \vec{\eta} \tag{2.19}$$

where $\vec{\varphi}$ is the 2-D wave front phase represented as a 1-D vector, $\vec{\eta}$ is the measurement noise and *G* is a matrix which models the discretized gradient. The commonly used discretizations are Fried (Fried 1977) and Hudgin (Hudgin 1977) models described in Section 2.2.

There have been several approaches proposed in the literature for solving for $\vec{\varphi}$ in (2.19). The solution using linear algebra is based on converting the 2-D structure into 1-D as in (2.19) and determining the pseudo-inverse matrix shown in (2.20) for any *G* (Ellerbroek 2002), (Fried 1977), (Hudgin 1977).

$$\vec{\varphi}_{estimate} = \left(G^T G\right)^{\dagger} G^T \begin{bmatrix} \vec{\dot{\varphi}}_x \\ \vec{\dot{\varphi}}_y \end{bmatrix} \tag{2.20}$$

Such an approach is $O(N^2)$ where N is the number of phase screen pixels for the matrix-vector multiply. This also requires a one-time cost of $O(N^3)$ flops at least due to the matrix-matrix multiplies to compute the pseudo-inverse as shown in (2.20).

Rather than use gradient operator models, an interaction matrix may be measured directly between the WFS data, $\vec{d}$, and the DM actuator commands, $\vec{a}$. The measured interaction matrix is denoted here as matrix $D_{d,a}$ instead of the gradient matrix, G.

$$\vec{d} = D_{d,a}\vec{a} \tag{2.21}$$

The Singular Value Decomposition (SVD) described in (Watkins 2002) is used to perform the pseudo-inversion of the interaction matrix since it allows suppression of invisible or near invisible modes.

$$U\Sigma V^T = D_{d,a} \tag{2.22}$$

where U and V are matrices that each contain a set of orthogonal modes. Matrix U is WFS modes and V is actuator modes in this case. Matrix Σ has non-zero entries on the diagonal referred to as "singular values." These non-zero entries are positive and arranged in decreasing order. It is important not to confuse this matrix with the summation operator.

The conversion of error in WFS data space to DM actuator space is the pseudo-inverse of $D_{d,a}$ and is denoted as $M_{a,d}$.

$$\begin{aligned} M_{a,d} &= D_{d,a}^{\dagger} \\ &= V\Sigma^{\dagger}U^T \end{aligned} \tag{2.23}$$

The inverse of V or U is the respective transpose (Watkins 2002). The magnitude of the singular values compared to the first singular value is a measure of visibility of the respective vectors in matrix V and U. Small singular values are nearly invisible and can

be dominated by system noise. These values are set to zero in order to ignore poorly sensed modes. Then, the pseudo-inverse of Σ is obtained by the element-by-element reciprocal of all non-zero values. This is the method used for comparison to classic AO in Chapter 5. The relationships between the three matrices are further explained in Appendix A.

It is preferable in AO to incorporate the statistics of the atmosphere and the noise in a reconstruction. One approach to this is to obtain a minimum variance reconstructor. This method is described in (Ellerbroek 2002) for the case of multiple WFS input. A simplification for single WFS systems is

$$\vec{\phi}_{estimate} = \left(G_{\psi}^{T} C_{\eta\eta}^{-1} G_{\psi} + C_{\psi\psi}^{-1}\right)^{-1} G_{\psi}^{T} C_{\eta\eta}^{-1} \vec{s} \tag{2.24}$$

where $C_{\psi\psi}$ and $C_{\eta\eta}$ are the covariance matrices of the turbulent phase and noise, respectively. In order to maintain the values of $C_{\psi\psi}$, the system must either be in open loop or must approximate the open loop slopes through the use of an interaction matrix.

2.6 Multi-grid conjugate gradient

Another approach based on the work of (Ellerbroek 2002) is an iterative method called the Multi-Grid Conjugate-Gradient (MGCG) method (Gilles, Ellerbroek and Vogel 2002). This method has been shown to give an estimate of the wave front in $28N\log_2 N$ operations where N is the number of data points that describe the shape of the wave front on a square. Since then, an implementation was presented in (Gilles 2003) that costs $520N$ operations for the required two iterations. Then (Vogel and Yang 2006) reduced the computations to $50N$ for the necessary two iterations. This approach preconditions the

matrices of (2.24) by multiplying each side by the inverse of some matrix C defined in a manner that improves the convergence speed.

$$C^{-1}\left(G_{\psi}^{T}C_{\eta\eta}^{-1}G_{\psi}+C_{\psi\psi}^{-1}\right)\vec{\phi}_{estimate}=C^{-1}G_{\psi}^{T}C_{\eta\eta}^{-1}\vec{s} \tag{2.25}$$

The method of (Gilles, Ellerbroek and Vogel 2002) and their later works take advantage of the sparseness in the associated matrices and the known statistics of turbulent distortion caused by the atmosphere. The approach is multi-grid, meaning that the resolution of the phase estimate begins coarse and iteratively gains resolution as the process executes. This method includes curvature regularization via one or two iterations of the Poisson solver of Section 2.3.

Another method of improving the convergence speed of MGCG algorithms is the use of Haar wavelet transforms to improve the condition number of the matrices involved (Tanaka 2006).

2.7 Shapelets

One method developed for the field of computer vision correlates the computed gradients of Gaussian shapes of differing sizes (referred to as "shapelets") and finds the correlation in the Fourier domain between these shapelets and the gradient measurements to reconstruct the image (Kovesi 2005). This method uses the surface normal of a 3-D object measured as angles. It is stated in (Kovesi 2005) that when partial difference measurements are available instead, it is preferable to use the method of (Frankot and Chellappa 1988). This family of approaches is explained in Section 2.8.

2.8 Fourier transforms

Several algorithms take advantage of fast Fourier transforms to reduce the computational complexity down to $O(N\log N)$. Another approach for AO is to determine the solution for Φ in (2.19) through integrating the gradient data in the Fourier domain and then computing the inverse Fourier transform to obtain Φ (Freischlad and Koliopoulos 1986), (Frankot and Chellappa 1988), and (Poyneer, Gavel and Brase 2002). The computation cost is dominated by three Fourier transforms that cost $7.5N\log_2 N$ in total and some multiple of N flops to combine them (Poyneer, Macintosh and Véran 2007). When using the notation of this dissertation, the foundation of these methods are (2.26) to (2.28) with the noise term ignored.

$$F\{X\} = G_{\omega x} F\{\Phi\} \tag{2.26}$$

$$F\{Y\} = G_{\omega y} F\{\Phi\} \tag{2.27}$$

$$\Phi_{estimate} = F^{-1}\left\{\frac{G_{\omega x}^{*} F\{X\} + G_{\omega y}^{*} F\{Y\}}{|G_{\omega x}|^2 + |G_{\omega y}|^2}\right\} \tag{2.28}$$

where F and F^{-1} are the Fourier and inverse Fourier transforms respectively. Matrices $G_{\omega x}$ and $G_{\omega y}$ are the Fourier transforms of the 2-D filters that perform the gradient operation (e.g., the Fried or Hudgin models). The points must be ignored where (2.28) results in a divide by zero.

The use of the Fourier cosine series has been proposed by (Talmi and Ribak 2006) who claimed that it converges on a wave front estimate in only $19N$.

2.9 Wavelet transforms

Wavelet filterbanks have been used to decompose matrix G in the same way a digital picture may be decomposed. This is in an effort to selectively increase the sparseness, reduce the computational complexity of matrix-vector multiplies to $O(N^{3/2})$ and reduce the computation of the pseudo-inverse to $O(N)$ (Dowla, et al. 2000).

Using the notation of (2.19), this problem is modified using a wavelet filter (2.29)

$$W\vec{s} = WGW^T W\vec{\phi} \tag{2.29}$$

where W may be any 1-D wavelet filter. The inverse of W is defined as the transpose of W. WGW^T is the 2-D wavelet decomposition of G. The solution is

$$\vec{\phi}_{estimate} = W^T \left[WGW^T \right]^{\dagger} W\vec{s} \tag{2.30}$$

To avoid confusing the method of (Dowla, et al. 2000) with the wavelet method developed by (Hampton, Agathoklis and Bradley 2008) in Chapter 3, (2.31) shows the essence of the latter method.

$$W_{Haar}\left({}^{M}\Phi\right)W_{Haar}^{T} = T\left\{\left[{}^{M}X \quad {}^{M}Y\right]\right\} \tag{2.31}$$

There is some transform, T{ }, of the partial differences of ${}^{M}\Phi$ that equals the Haar wavelet transform of ${}^{M}\Phi$. This transform is developed in (Hampton, Agathoklis and Bradley 2008) and is reprinted in Chapter 3. The phase estimate is then the inverse 2-D Haar wavelet transform of (2.31).

One major difference is that the signals remain two dimensional. Another is that the method of (Hampton, Agathoklis and Bradley 2008) is lossless when there is no noise with the exception of data invisible to the sensor. The method of (Dowla, et al. 2000) is

defined as a "lossy" method since it zeros out small values in the decomposition. The amount of loss is a trade-off for computation speed.

2.10 Conclusions

The Poisson solver operates on the divergence of gradient data, rather than the gradient itself. This changes the problem from integrating first order derivatives to equating combinations of the second-order derivative. Either approach produces an estimate of the integrated gradients.

The approach of 1-D integration only require $1N$ or $2N$ flops, but is dismissed in the literature as causing an excessive amount of error. The 1-D approach is the only known approach that is computationally faster than the wavelet-based 2-D integration process presented in Chapter 3, which can provide a solution in $5.4N$ or $10.2N$ flops. MGCG is also a fast approach at $50N$ flops and has the advantage of being applicable to a wider range of problems than integration. Approaches that use the Fourier transform, including Shaplets, are $O(N\log_2 N)$. Wavelet transforms have been used to precondition interaction matrices to reduce computations to $O(N^{3/2})$ from the $O(N^2)$ of the pseudo-inverse techniques.

The Poisson solver is known to converge iteratively on a least-squares optimal result. Any integration method may be used as the initial state of a Poisson solver for further reduction in error. Using a fast method as an initial estimate would decrease the number of iterations required for convergence with the Poisson solver.

Chapter 3

Reconstruction of the Haar Wavelet Decomposition

3.1 Abstract

A new technique for integrating 2-D partial differences (or discretely sampled gradient measurements) on a square data set is presented. This technique is based on obtaining the Haar wavelet image decomposition of the integrated data by appropriate filtering and down-sampling of the gradient measurement data. The use of the wavelet decomposition leads to an algorithm with complexity of $O(N)$, where N is the number of data points in the reconstruction and, at the same time, allows "denoising" of the data using the wavelet coefficients. The proposed technique is illustrated with two examples and its reconstruction accuracy, computational speed and denoising performance are discussed. Results indicate that the proposed technique is a computationally efficient and accurate technique for integrating partial differences.

Material in this chapter has been published as:

3.2 Wavelet filter preliminaries

3.2.1 Haar wavelet filter definition

The orthogonal Haar filters given by (Vetterli and Kovacevic 1995):

$$H_L(z) = \left(1 + z^{-1}\right)/\sqrt{2} \tag{3.1}$$

$$H_H(z) = \left(1 - z^{-1}\right)/\sqrt{2} \tag{3.2}$$

can be used to obtain decompositions of data by filtering in either the vertical or horizontal directions. To distinguish between these directions, z_h will indicate horizontal filtering of rows, whereas z_v will indicate vertical filtering of columns.

The relationship between the gradient measurements obtained using Fried and Hudgin geometries of (2.8) to (2.9) can be represented by the Haar filters and are shown in Figure 3.1 and equations (3.3) and (3.4).

$${}^{M}\widetilde{X} = {}^{M}_{F}\widetilde{X} = \frac{{}^{M}_{H}\widetilde{X}H_L(z_v)}{\sqrt{2}} = {}^{M}\Phi H_H(z_h)H_L(z_v) \tag{3.3}$$

$${}^{M}\widetilde{Y} = {}^{M}_{F}\widetilde{Y} = \frac{{}^{M}_{H}\widetilde{Y}H_L(z_h)}{\sqrt{2}} = {}^{M}\Phi H_L(z_h)H_H(z_v) \tag{3.4}$$

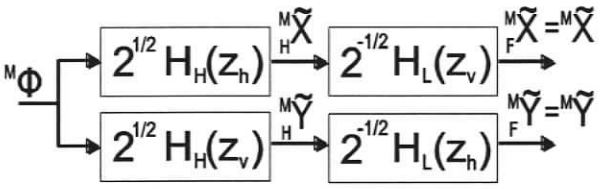

Figure 3.1: Relationship between Hudgin and Fried aligned gradient models and the Haar filters.

As shown in Figure 3.1, ${}^{M}\Phi$ is filtered by the high pass Haar filter (3.2) in z_h to obtain the x-gradients and in z_v to obtain the y-gradients in Hudgin alignement. This is also multiplied by $\sqrt{2}$ to cancel the $\sqrt{2}$ in the denominator. The Hugin aligned gradients are divided by $\sqrt{2}$ and filtered by the low pass Haar filter (3.1) to convert to Fried.

3.2.2 Filter factorization through a down-sampler

Consider a signal $S(z)$ that has every even-numbered element discarded by down-sampling. This signal is then filtered by $H(z)$. The result of this process is indistinguishable from signal $S(z)H(z^2)$ followed by down-sampling. This is one aspect of the noble identities discussed in several works including (Park and Kim 2000) and (Li 2003).

$$\downarrow_2 \{S(z)\} H(z) = \downarrow_2 \{S(z) H(z^2)\} \tag{3.5}$$

This concept permits filter steps in the HWAF to be moved to any earlier step in the process as long as the exponents in the z-domain are doubled each time it is moved before a down-sampling-by-2 operator. For example, $H_H(z^2) = (1 - z^{-2})/\sqrt{2}$. This can also be factored as in (3.6) showing that every time $H_H(z)$ is moved before a down-sampler, the resulting filter can be factored to obtain $H_H(z)$ again. Therefore, the process can be recursive for this specific filter. If the factorization of (3.6) is not executed after each down-sampler, (3.7) shows the resulting filter after passing through K factor 2 down-samplers. These factorizations are used extensively to prove the following algorithms in Section 3.6.

$$H_H(z^2) = \sqrt{2} H_H(z) H_L(z) \tag{3.6}$$

$$H_H(z^{2^K}) = 2^{K/2} H_H(z) \prod_{k=0}^{K-1} H_L(z^{2^k}) \tag{3.7}$$

3.2.3 Haar wavelet filterbank structure

The structure of 1 stage of the 2-D HWAF is shown in Figure 3.2.

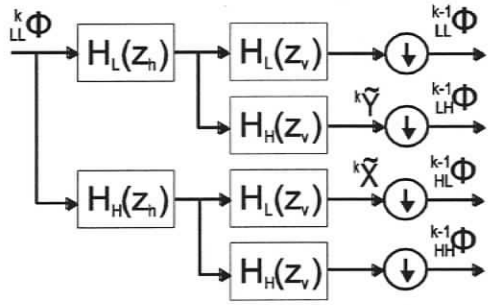

Figure 3.2: One stage of the HWAF. Stages cascade by connecting the LL output to the LL input of the next stage.

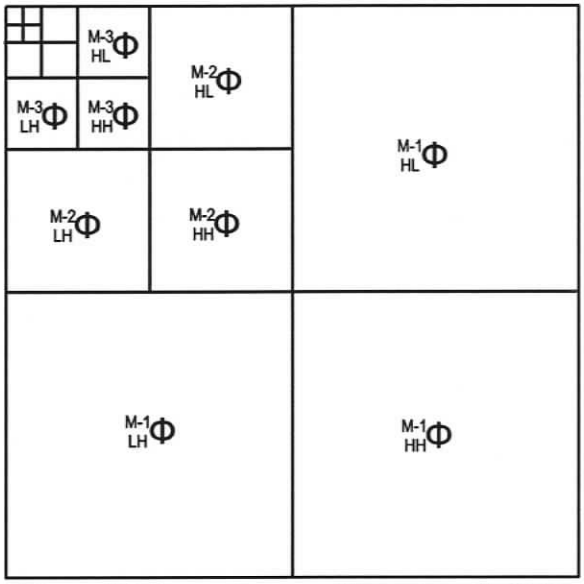

Figure 3.3: Image decomposition using a HWAF.

Let the symbol Φ represent a 2-D image. The superscript in the top left of ${}^{k}\Phi$ indicates the stage and the size of the 2-D image, i.e., ${}^{k}\Phi$ is of size $2^k \times 2^k$ for $k = 1 \leq k \leq M$. The subscripts LL, LH, HL and HH introduced in Figure 3.2 indicate the order that the Haar filters were applied to the signal. By cascading wavelet filter blocks such as the one in Figure 3.2 to each LL output, an M-stage 2-D filterbank is obtained where the output of the last stage is a block of 2×2 pixels. This M-stage HWAF leads to the decomposition of Figure 3.3 which is the same structure as in (Park and Kim 2000).

Using the equivalent analysis filters for M-stage 2-D filterbanks, the non-recursive structure of the Haar wavelet decomposition of the image can be obtained as

$$ {}^{M-m}_{LL}\Phi = \downarrow_{2^m} \left\{ {}^{M}\Phi \prod_{k=0}^{m-1} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-1} H_L\left(z_v^{2^k}\right) \right\} \tag{3.8} $$

$$ {}^{M-m}_{LH}\Phi = \downarrow_{2^m} \left\{ {}^{M}\Phi H_H\left(z_v^{2^{m-1}}\right) \prod_{k=0}^{m-1} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-2} H_L\left(z_v^{2^k}\right) \right\} \tag{3.9} $$

$$ {}^{M-m}_{HL}\Phi = \downarrow_{2^m} \left\{ {}^{M}\Phi H_H\left(z_h^{2^{m-1}}\right) \prod_{k=0}^{m-2} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-1} H_L\left(z_v^{2^k}\right) \right\} \tag{3.10} $$

$$ {}^{M-m}_{HH}\Phi = \downarrow_{2^m} \left\{ {}^{M}\Phi H_H\left(z_h^{2^{m-1}}\right) H_H\left(z_v^{2^{m-1}}\right) \prod_{k=0}^{m-2} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-2} H_L\left(z_v^{2^k}\right) \right\} \tag{3.11} $$

where $\downarrow_k \{\ \}$ is the down-sampling function in both horizontal and vertical directions by the factor k. The first valid data point of filtering is where the data fully supported the filter. Down-sampling by two retains data points on odd columns and rows.

In this chapter, the structure of the HWAF given by (3.8) to (3.11) will be modified by rearranging the order of operations to allow for image gradient data to be used as the input for the decomposition. The original image can then be reconstructed by applying a standard HWSF to the image decomposition.

The original image can then be reconstructed by applying a standard HWSF to the image decomposition, i.e., the decomposition provided by the HWAF is perfectly reversible by the HWSF, as in Figure 3.4.

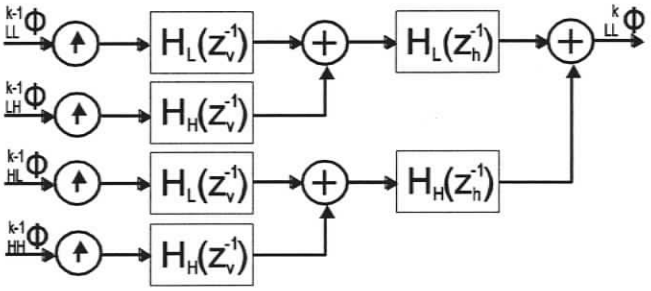

Figure 3.4: Structure of HWSF. Stages cascade by connecting the LL output to the LL input of the next stage. Upsampling doubles resolution and inserts zeros for unknown data points.

3.3 Haar decomposition from gradient data

The approach for integrating partial differences presented in this section is based on generating the decomposition given in Figure 3.5, from ${}^{M}\widetilde{X}$ and ${}^{M}\widetilde{Y}$, (3.3) and (3.4), and then obtaining the original data using standard HWSFs applied to this decomposition.

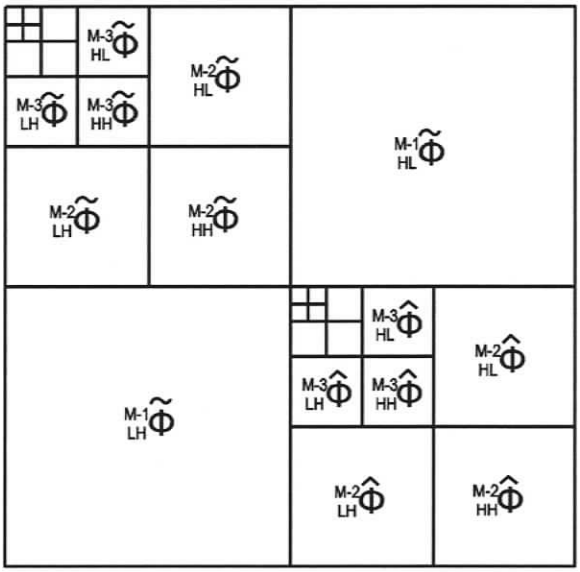

Figure 3.5: Image decomposition obtained from Fried aligned gradient data.

The motivation for developing the decomposition of Figure 3.5 is based on the similarity between the Haar wavelet and the representation of the gradient measurements using the Fried model. Comparing Figure 3.1 to Figure 3.2, it is clear that gradient data that matches the Fried model perfectly provides the filtering necessary to obtain ${}^{M-1}_{LH}\Phi$

and ${}^{M-1}_{HL}\Phi$. All that is required is the down-sampling of the measurements to obtain the upper-right and lower-left part of the decomposition shown in Figure 3.3. Obtaining the diagonal sections of Figure 3.5 using the gradient data requires the development of filterbanks, which are no more than a manipulation of the order of operations of the linear filters in the HWAF and the HWSF. The algebraic proof is provided in Section 3.6. This manipulation can be considered as converting the Haar decomposition from a denoising/compression tool into an important intermediate step when solving 2-D discrete partial differential equations.

3.4 Integration algorithm

The estimation of the 2-D numerical integral is obtained in two steps, the *Analysis* and the *Synthesis* steps. The filterbank for the first three quadrants of Figure 3.5 is given directly as Figure 3.6. An efficiently implemented lattice structure of this filterbank is shown as Figure 3.7.

Analysis Step: Given ${}^{M}\widetilde{X}$ and ${}^{M}\widetilde{Y}$, in (3.3) and (3.4), generate the decomposition of Figure 3.5:

HL and LH quadrants (upper-right and lower left-quadrants)

$$ {}^{M-1}_{HL}\widetilde{\Phi} = \downarrow_2 \left\{ {}^{M}\Phi H_H(z_h) H_L(z_v) \right\} = \downarrow_2 \left\{ {}^{M}\widetilde{X} \right\} \tag{3.12} $$

$$ {}^{M-1}_{LH}\widetilde{\Phi} = \downarrow_2 \left\{ {}^{M}\Phi H_L(z_h) H_H(z_v) \right\} = \downarrow_2 \left\{ {}^{M}\widetilde{Y} \right\} \tag{3.13} $$

LL quadrant (upper-left quadrant of Figure 3.5)

For $k = M$ to 2

$$ {}^{k-1}\widetilde{Y} = \sqrt{2} \downarrow_2 \left\{ {}^{k}\widetilde{Y} H_L(z_h^2) H_L^2(z_v) \right\} \tag{3.14} $$

$$ {}^{k-1}\widetilde{X} = \sqrt{2} \downarrow_2 \left\{ {}^{k}\widetilde{X} H_L^2(z_h) H_L(z_v^2) \right\} \tag{3.15} $$

$$ {}^{k-2}_{LH}\widetilde{\Phi} = \downarrow_2 \left\{ {}^{k-1}\widetilde{Y} \right\} \tag{3.16} $$

$$ {}^{k-2}_{HL}\widetilde{\Phi} = \downarrow_2 \left\{ {}^{k-1}\widetilde{X} \right\} \tag{3.17} $$

$$ \begin{aligned} {}^{k-2}_{HH}\widetilde{\Phi} &= \sqrt{2} \downarrow_4 \left\{ {}^{k}\widetilde{Y} H_H(z_h^2) H_L^2(z_v) \right\} \\ &= \sqrt{2} \downarrow_4 \left\{ {}^{k}\widetilde{X} H_L^2(z_h) H_H(z_v^2) \right\} \end{aligned} \tag{3.18} $$

When noise is present, the average of the two lines of (3.18) is used as in Figure 3.7.

end

$$ {}^{0}_{LL}\widetilde{\Phi} = 2^{M} \, mean(\Phi) \tag{3.19} $$

i.e., ${}^{0}_{LL}\widetilde{\Phi}$ is proportional to the mean value of the image.

HH quadrant (lower-right quadrant of Figure 3.5):

For k = M to 2

$$ {}^{k-1}\hat{X} = \begin{cases} \sqrt{2}\downarrow_2 \left\{ {}^{M}\tilde{X}H_L\left(z_h^2\right)H_H^2\left(z_v\right)\right\} & for \quad k = M \\ \sqrt{2}\downarrow_2 \left\{ {}^{k}\hat{X}H_L\left(z_h^2\right)H_L^2\left(z_v\right)\right\} & for \quad k < M \end{cases} \tag{3.20} $$

$$ {}^{k-1}\hat{Y} = \begin{cases} \sqrt{2}\downarrow_2 \left\{ {}^{M}\tilde{Y}H_H^2\left(z_h\right)H_L\left(z_v^2\right)\right\} & for \quad k = M \\ \sqrt{2}\downarrow_2 \left\{ {}^{k}\hat{Y}H_L^2\left(z_v\right)H_L\left(z_v^2\right)\right\} & for \quad k < M \end{cases} \tag{3.21} $$

$$ {}^{k-2}_{LH}\hat{\Phi} = \downarrow_2 \left\{ {}^{k-1}\hat{X}\right\} \tag{3.22} $$

$$ {}^{k-2}_{HL}\hat{\Phi} = \downarrow_2 \left\{ {}^{k-1}\hat{Y}\right\} \tag{3.23} $$

$$ \begin{aligned} {}^{k-2}_{HH}\hat{\Phi} &= \begin{cases} \sqrt{2}\downarrow_4 \left\{ {}^{M}\tilde{Y}H_H^2\left(z_h\right)H_H\left(z_v^2\right)\right\} & for \quad k = M \\ \sqrt{2}\downarrow_4 \left\{ {}^{k}\hat{Y}H_L^2\left(z_h\right)H_H\left(z_v^2\right)\right\} & for \quad k < M \end{cases} \\ &= \begin{cases} \sqrt{2}\downarrow_4 \left\{ {}^{M}\tilde{X}H_H\left(z_h^2\right)H_H^2\left(z_v\right)\right\} & for \quad k = M \\ \sqrt{2}\downarrow_4 \left\{ {}^{k}\hat{X}H_H\left(z_h^2\right)H_L^2\left(z_v\right)\right\} & for \quad k < M \end{cases} \end{aligned} \tag{3.24} $$

When noise is present, the average of the results using X or Y in (3.24) is used.

end

Synthesis Step: Given the image decomposition of Figure 3.5 obtained in the Analysis step:

Apply an M-1-stage HWSF to HH quadrant to obtain ${}^{M-1}_{HH}\hat{\Phi}$.

Apply an M-stage HWSF to all 4 quadrants, including ${}^{M-1}_{HH}\hat{\Phi}$, to obtain estimate of 2-D numerical integration.

Algorithm end.

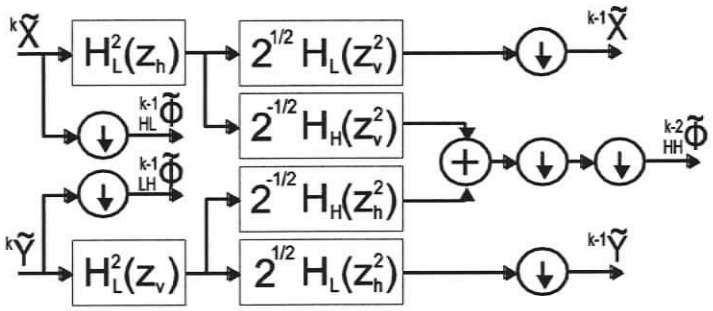

Figure 3.6: Gradient analysis filterbank structure for HL, LH and LL quadrants.

The above equations have been obtained by combining (3.3) or (3.4) with (3.9) through (3.11) to obtain the wavelet decomposition given in Figure 3.5 from the gradient signals ${}^{M}\widetilde{X}$ and ${}^{M}\widetilde{Y}$, instead of the image, ${}^{M}\Phi$ as input (see Section 3.6).

Two entries in the decomposition given in Figure 3.5 cannot be obtained from the gradient data. The first is ${}^{0}_{LL}\widetilde{\Phi}$ which is the mean of the image and could be found with an additional sensor to measure the mean as suggested in (3.19). In AO, the mean is called piston and it may be acceptable that this value would be set to 0. In other applications, such as the one considered in Section 3.7.2, ${}^{0}_{LL}\widetilde{\Phi}$ can be estimated from the knowledge of the signal dynamic range.

The other entry is ${}^{0}_{LL}\hat{\Phi}$. The HH quadrant represented by ${}^{M-1}_{HH}\Phi$ in Figure 3.3 is entirely composed of 2×2 waffle, just as ${}^{M-1}_{LH}\Phi$ and ${}^{M-1}_{HL}\Phi$ are 2×2 tilts and ${}^{M-1}_{LL}\Phi$ is 2×2 pistons. Any waffle-type shape is composed of linear combinations of these 2×2 waffle modes. However, ${}^{M-1}_{HH}\Phi$ cannot be solved directly because it contains one invisible waffle shape. This mode is proportional to a checkerboard pattern given as $(-1)^{i+j}$, where *i* and *j* represent the row and column of the image. It was determined that the invisible component of ${}^{M-1}_{HH}\Phi$ can be isolated by applying an M-1 stage HWAF to ${}^{M-1}_{HH}\Phi$. This allows the other visible components to be solved directly. The magnitude of

invisible waffle is directly proportional to ${}_{LL}^{0}\hat{\Phi}$. As suggested by the subscripts, it is directly proportional to the mean value of all the 2×2 waffle modes that compose ${}_{HH}^{M-1}\Phi$. Such a high-frequency component is known to be unusable in AO and is expected to not have a significant magnitude in digital photography. This is supported by the small reconstruction error found in Section 3.7. Therefore, ${}_{LL}^{0}\hat{\Phi}$ can be set to 0 and ignored without significant error.

Further, it is known from atmospheric turbulence models that the high spatial frequency shapes are reduced by a factor of $f^{11/3}$ in power as spatial frequency increases (Noll 1976), which makes the signal to noise ratio (SNR) poor at high frequencies. It may be desirable to remove the localized waffle from the corrected modes (D. T. Gavel 2003). The Haar decomposition automatically isolates compactly supported waffle modes into the HH quadrant. This allows for removal of these modes simply by not computing the respective portion of the decomposition. This is only shown for the open loop case in the example section and will be shown for the closed loop case in Chapter 5.

3.5 Computational complexity

The computational complexity of the proposed algorithm can be established based on the complexity of the two-dimensional separable discrete wavelet transform (DWT) (Vetterli and Kovacevic 1995). Both the Analysis and Synthesis steps of the proposed algorithm involve convolutions with separable wavelet filters followed by down-sampling in both dimensions and thus, based on (Vetterli and Kovacevic 1995), the computational complexity will be of order $O(N)$, where $N = 2^M \times 2^M$. Thus, the proposed algorithm is scaling linearly with the number of sensors.

Each 2-D down-sampling operator in Figure 3.7 reduces the amount of data by a factor of 4. There are 8 multiplies and additions after the first set of down-samplers and 4 (ignoring multiplies by -1) after the second set of down samplers. To obtain the HH quadrant, another filterbank of similar structure is used with the same computation cost. Each subsequent stage operates on one quarter of the data of the previous stage, leading to the following complexity for the Analysis step:

$$\begin{aligned} C_{Analysis}(N) &\le 2\left(8\frac{N}{4}+4\frac{N}{16}\right)\sum_{m=0}^{M-2}\frac{1}{4^m} \\ &\le 4.5N\left(\frac{4}{3}\right)=6N \end{aligned} \tag{3.25}$$

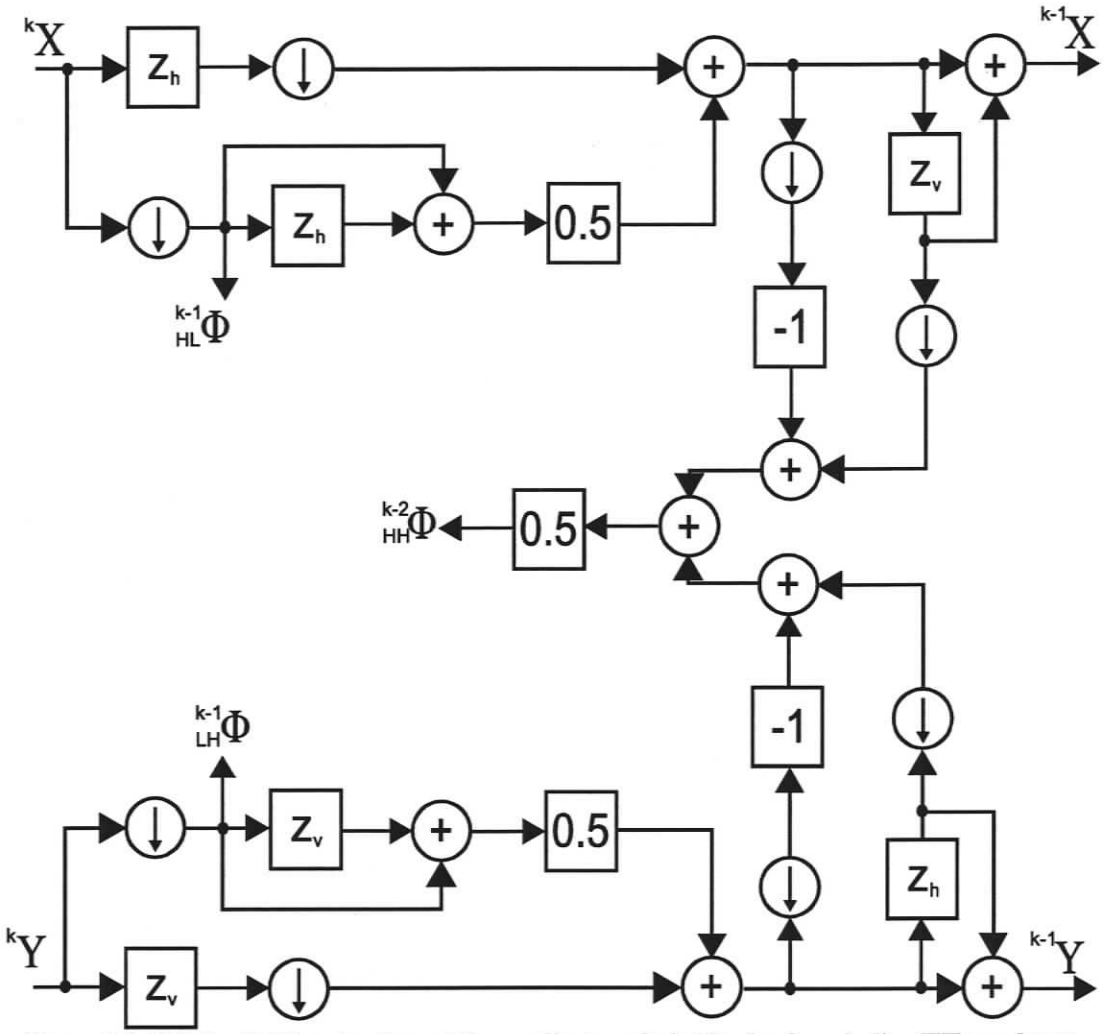

Figure 3.7: Efficient lattice structure of the gradient analysis filterbank excluding HH quadrant.

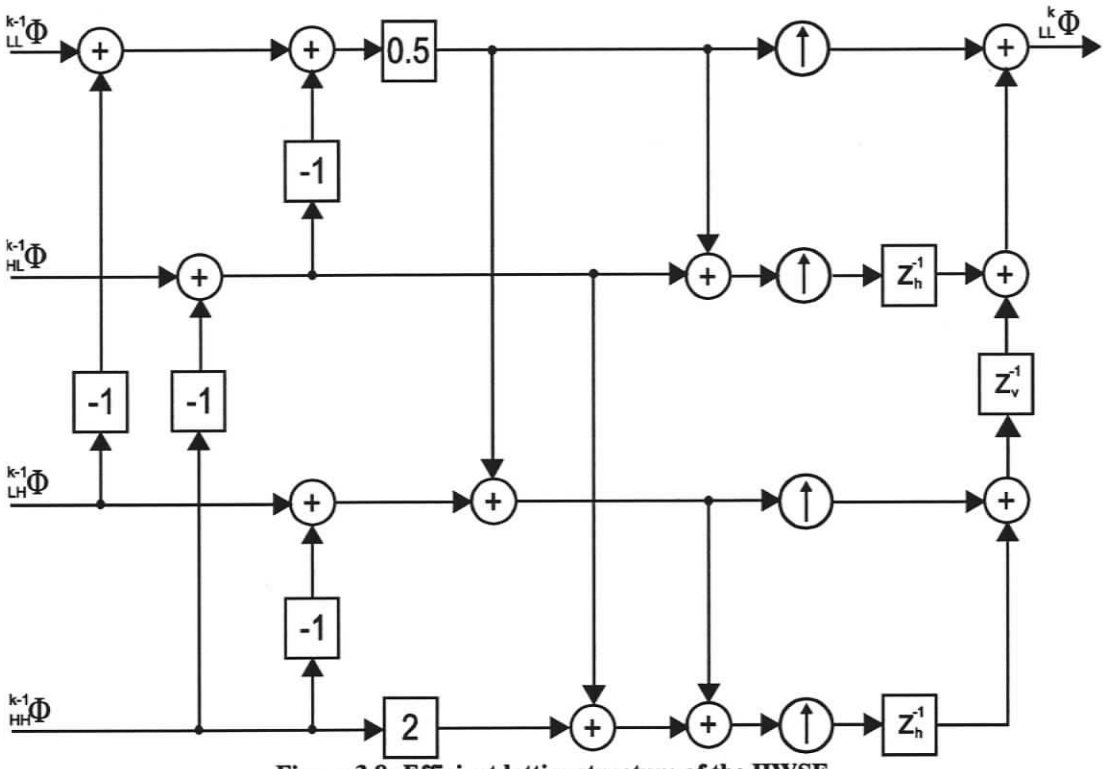

Figure 3.8: Efficient lattice structure of the HWSF.

In the synthesis step, an efficient lattice structure of the HWSF is provided as Figure 3.8. Prior to the up-sampling operations, there are 10 flops to produce 4 pixels of the reconstruction. Therefore, there are 2.5 flops per pixel of one stage of the HWSF. Multiplying this by 4/3 as in (3.25) results in a computation cost of $3.33N$ for M stages of the HWSF. The HH quadrant also requires M-1 stages of the HWSF, so it comes at a cost of $0.84N$ operations. This leads to a $C_{Synthesis}(N)$ of $4.17N$ and a total complexity for the reconstruction of

$$C_{Total}(N) = C_{Analysis}(N) + C_{Synthesis}(N) = 10.17N \tag{3.26}$$

Ignoring the HH quadrant has the benefit of ignoring half the computations of the analysis stage and almost half the computations of the synthesis stage, reducing the lower limit of total computations to $5.33N$. This is less than expected from the above discussion

because the operations of the last stage of the HWSF does not need to compute anything sourced from ${}^{M-1}_{HH}\Phi$.

The performance of the proposed algorithm will be illustrated and evaluated in Section 3.7.

3.6 Algorithm proof

The proof of how the gradient data can be transformed into the Haar decomposition is based entirely on algebraic manipulation of the standard HWAF and HWSF. The approach of the following algebra is to move the high pass filtering step that occurs at the end of the standard HWAF to the beginning via the noble identities. The model of a gradient sensor includes these highpass filters, so once the highpass filter is algebraically moved to the first step of the filterbank, it is trivial to move such filters out of the filterbank and into the sensor model. The main steps of developing (3.12) to (3.24) from (3.8) to (3.11) are given in this section.

HL and LH quadrants: (3.12) and (3.13) follow directly from (3.9) and (3.10) for $m = 1$ and the definitions of ${}^{M}\widetilde{X}$ and ${}^{M}\widetilde{Y}$ given by (3.3) and (3.4).

LL quadrant: To derive (3.16), first (3.9) is factored to obtain an expression for ${}^{M-m}_{LH}\widetilde{\Phi}$ which involves the available gradient data ${}^{M}\widetilde{Y}$ instead of the image data ${}^{M}\Phi$ as in (3.27). The use of square brackets helps to indicate how the algebraic expansion in the second line relates to the first line. These brackets are dropped for the reordering in the third line and finally ${}^{M}\Phi H_L(z_h)H_H(z_V)$ is directly replaced by ${}^{M}\widetilde{Y}$ through the use of (3.4). The ${}^{M-m}_{HL}\widetilde{\Phi}$ and ${}^{M-m}_{HH}\widetilde{\Phi}$ terms are factored in similar ways.

$$
\begin{aligned}
{}^{M-m}_{LH}\tilde{\Phi} &= \downarrow_{2^m} \left\{ {}^{M}\Phi \left[H_H\left(z_v^{2^{m-1}}\right) \right] \left[\prod_{k=0}^{m-1} H_L\left(z_h^{2^k}\right) \right] \prod_{k=0}^{m-2} H_L\left(z_v^{2^k}\right) \right\} \\
&= \downarrow_{2^m} \left\{ {}^{M}\Phi \left[\sqrt{2}^{m-1} H_H(z_v) \prod_{k=0}^{m-2} H_L\left(z_v^{2^k}\right) \right] \left[H_L(z_h) \prod_{k=1}^{m-1} H_L\left(z_h^{2^k}\right) \right] \prod_{k=0}^{m-2} H_L\left(z_v^{2^k}\right) \right\} \\
&= \sqrt{2}^{m-1} \downarrow_{2^m} \left\{ {}^{M}\Phi H_L(z_h) H_H(z_V) \prod_{k=1}^{m-1} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-2} H_L^2\left(z_v^{2^k}\right) \right\} \\
&= \sqrt{2}^{m-1} \downarrow_{2^m} \left\{ {}^{M}\tilde{Y} \prod_{k=1}^{m-1} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-2} H_L^2\left(z_v^{2^k}\right) \right\}
\end{aligned}
\tag{3.27}
$$

The first line of (3.27) is the non-recursive form shown in (3.9). The identity of (3.7) is used to obtain the second line. The third line is simple reordering of terms so that the use of (3.4) to obtain ${}^{M}\tilde{Y}$ becomes obvious. From the above non-recursive equation and the definition of ${}^{k-1}\tilde{Y}$ in (3.14), the recursive equation (3.16) can be obtained using:

$$
\begin{aligned}
{}^{M-m}_{LH}\tilde{\Phi} &= \downarrow_{2^m} \left\{ {}^{M}\tilde{Y} \sqrt{2}^{m-1} \prod_{k=1}^{m-1} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-2} H_L^2\left(z_v^{2^k}\right) \right\} \\
&= \downarrow_{2^{m-1}} \left\{ \downarrow_2 \left\{ {}^{M}\tilde{Y} \sqrt{2}^{m-1} H_L\left(z_h^2\right) H_L^2(z_v) \prod_{k=2}^{m-1} H_L\left(z_h^{2^k}\right) \prod_{k=1}^{m-2} H_L^2\left(z_v^{2^k}\right) \right\} \right\} \\
&= \downarrow_{2^{m-1}} \left\{ \downarrow_2 \left\{ {}^{M}\tilde{Y} \sqrt{2} H_L\left(z_h^2\right) H_L^2(z_v) \right\} \sqrt{2}^{m-2} \prod_{k=1}^{m-2} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-3} H_L^2\left(z_v^{2^k}\right) \right\} \\
&= \downarrow_{2^{m-1}} \left\{ {}^{M-1}\tilde{Y}(z) \sqrt{2}^{m-2} \prod_{k=1}^{m-2} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-3} H_L^2\left(z_v^{2^k}\right) \right\} \\
&\vdots \\
&= \downarrow_2 \left\{ {}^{M-m+1}\tilde{Y}(z) \right\}
\end{aligned}
\tag{3.28}
$$

Starting from (3.27) in the first line of (3.28) the lowpass filters are factored out of the product operators in the second line. In the third line, the order of down-sampling is adjusted through the use of the noble identity of (3.5), which modifies the initial and final values of the product indices. In the fourth line, ${}^{M-1}\tilde{Y}$ is substituted using (3.14) which leads to an equation that is the same as the first line of (3.28) with $m \rightarrow m-1$ and $M \rightarrow$

$M-1$. Repeating this circular process leads to the last line of (3.28) which after the substitution $m = M\text{-}k\text{+}2$, becomes (3.16).

This shows that for the computation of ${}^{M-2}_{LH}\widetilde{\Phi}$, ${}^{M-1}\widetilde{Y}$ is computed first by filtering ${}^{M}\widetilde{Y}$ and then this data is down-sampled to produce ${}^{M-2}_{LH}\widetilde{\Phi}$. In the next iteration, the computation of ${}^{M-3}_{LH}\widetilde{\Phi}$ can start with ${}^{M-1}\widetilde{Y}$ which is available from the previous step, rather than begin with ${}^{M}\widetilde{Y}$.

The derivation of (3.17) can be carried out in a similar way. From (3.10) and the definition of ${}^{M}\widetilde{X}$ it follows that

$$\begin{aligned}
{}^{M-m}_{HL}\tilde{\Phi} &= \downarrow_{2^m}\left\{{}^{M}\Phi H_H\left(z_h^{2^{m-1}}\right)\prod_{k=0}^{m-2}H_L\left(z_h^{2^k}\right)\prod_{k=0}^{m-1}H_L\left(z_v^{2^k}\right)\right\} \\
&= \downarrow_{2^m}\left\{{}^{M}\Phi\sqrt{2}^{m-1}H_H\left(z_h\right)\prod_{k=0}^{m-2}H_L\left(z_h^{2^k}\right)\prod_{k=0}^{m-2}H_L\left(z_h^{2^k}\right)H_L\left(z_v\right)\prod_{k=1}^{m-1}H_L\left(z_v^{2^k}\right)\right\} \\
&= \sqrt{2}^{m-1}\downarrow_{2^m}\left\{{}^{M}\tilde{X}\prod_{k=0}^{m-2}H_L^2\left(z_h^{2^k}\right)\prod_{k=1}^{m-1}H_L\left(z_v^{2^k}\right)\right\}
\end{aligned} \tag{3.29}$$

which, using the definition of ${}^{k-1}\widetilde{X}$ in (3.15), leads to (3.17) in a similar way as with (3.16) from (3.28).

Then (3.18) can be obtained by developing a non-recursive form of (3.11) which involves the available gradient data ${}^{M}\widetilde{X}$ instead of the image data ${}^{M}\Phi$:

$$\begin{aligned}
{}^{M-m}_{HH}\tilde{\Phi} &= \downarrow_{2^m}\left\{{}^{M}\Phi H_H\left(z_h^{2^{m-1}}\right)H_H\left(z_v^{2^{m-1}}\right)\prod_{k=0}^{m-2}H_L\left(z_h^{2^k}\right)\prod_{k=0}^{m-2}H_L\left(z_v^{2^k}\right)\right\} \\
&= \downarrow_{2^m}\left\{{}^{M}\Phi\sqrt{2}^{m-1}H_H\left(z_h\right)\prod_{k=0}^{m-2}H_L\left(z_h^{2^k}\right)\prod_{k=0}^{m-2}H_L\left(z_h^{2^k}\right)H_H\left(z_v^{2^{m-1}}\right)H_L\left(z_v\right)\prod_{k=1}^{m-2}H_L\left(z_v^{2^k}\right)\right\} \\
&= \sqrt{2}^{m-1}\downarrow_{2^m}\left\{{}^{M}\tilde{X}\prod_{k=0}^{m-2}H_L^2\left(z_h^{2^k}\right)H_H\left(z_v^{2^{m-1}}\right)\prod_{k=1}^{m-2}H_L\left(z_v^{2^k}\right)\right\}
\end{aligned} \tag{3.30}$$

and use the approach in (3.28) to obtain the recursive form of (3.18).

HH quadrant : To obtain the equations for the HH quadrant, ${}^{M-1}_{HH}\Phi$ is decomposed with an M-1-stage HWAF to the form shown in Figure 3.5. This implies that ${}^{M-m}_{LL}\hat{\Phi}$ is obtained by applying (3.8) to ${}^{M-1}_{HH}\Phi$. Then ${}^{M-1}_{HH}\Phi$ is substituted in this expression using (3.11) and after changing the order of down-sampling one obtains the following equation:

$$
\begin{aligned}
{}^{M-m}_{LL}\hat{\Phi} &= \downarrow_{2^{m-1}} \left\{ {}^{M-1}_{HH}\Phi \prod_{k=0}^{m-2} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-2} H_L\left(z_v^{2^k}\right) \right\} \\
&= \downarrow_{2^{m-1}} \left\{ \downarrow_2 \left\{ {}^{M}\Phi H_H(z_h) H_H(z_v) \right\} \prod_{k=0}^{m-2} H_L\left(z_h^{2^k}\right) \prod_{k=0}^{m-2} H_L\left(z_v^{2^k}\right) \right\} \\
&= \downarrow_{2^m} \left\{ {}^{M}\Phi H_H(z_h) H_H(z_v) \prod_{k=1}^{m-1} H_L\left(z_h^{2^k}\right) \prod_{k=1}^{m-1} H_L\left(z_v^{2^k}\right) \right\} \\
&= \downarrow_2 \left\{ {}^{M}_{H}\tilde{X} \frac{H_H(z_v)}{\sqrt{2}} \right\} \; for\; m = 1 \\
&= \downarrow_2 \left\{ {}^{M}_{H}\tilde{Y} \frac{H_H(z_h)}{\sqrt{2}} \right\} \; for\; m = 1
\end{aligned}
\tag{3.31}
$$

Equation (3.31) is not solvable with Fried aligned gradient data. However, it can be factored using Hudgin aligned data and directly solved with $m = 1$. Equations (3.32) to (3.34) are obtained following similar operations as in (3.31) when the HH quadrant must be solved with Fried data.

$$
\begin{aligned}
{}^{M-m}_{LH}\hat{\Phi} &= \downarrow_{2^m} \left\{ {}^{M}\Phi H_H(z_h) H_H(z_v) H_H\left(z_v^{2^{m-1}}\right) \prod_{k=1}^{m-1} H_L\left(z_h^{2^k}\right) \prod_{k=1}^{m-2} H_L\left(z_v^{2^k}\right) \right\} \\
&= \downarrow_{2^m} \left\{ {}^{M}_{F}\tilde{X} H_H^2(z_v) \prod_{k=1}^{m-1} H_L\left(z_h^{2^k}\right) \prod_{k=1}^{m-2} H_L^2\left(z_v^{2^k}\right) \right\}
\end{aligned}
\tag{3.32}
$$

$$
\begin{aligned}
{}^{M-m}_{HL}\hat{\Phi} &= \downarrow_{2^m} \left\{ {}^{M}\Phi H_H(z_h) H_H(z_v) H_H\left(z_h^{2^{m-1}}\right) \prod_{k=1}^{m-2} H_L\left(z_h^{2^k}\right) \prod_{k=1}^{m-1} H_L\left(z_v^{2^k}\right) \right\} \\
&= \downarrow_{2^m} \left\{ {}^{M}_{F}\tilde{Y} H_H^2(z_h) \prod_{k=1}^{m-2} H_L^2\left(z_h^{2^k}\right) \prod_{k=1}^{m-1} H_L\left(z_v^{2^k}\right) \right\}
\end{aligned}
\tag{3.33}
$$

$$
\begin{aligned}
{}^{M-m}_{HH}\hat{\Phi} &= \downarrow_{2^m}\left\{ {}^{M}\Phi H_H(z_h) H_H(z_v) H_H\left(z_h^{2^{m-1}}\right) H_H\left(z_v^{2^{m-1}}\right) \prod_{k=1}^{m-2} H_L\left(z_h^{2^k}\right) \prod_{k=1}^{m-2} H_L\left(z_v^{2^k}\right)\right\} \\
&= \downarrow_{2^m}\left\{ {}^{M}_{F}\tilde{X} H_H\left(z_h^{2^{m-1}}\right) H_H^2(z_v)\sqrt{2}^{m-1} \prod_{k=1}^{m-2} H_L\left(z_h^{2^k}\right) \prod_{k=1}^{m-2} H_L^2\left(z_v^{2^k}\right)\right\} \qquad (3.34) \\
&= \downarrow_{2^m}\left\{ {}^{M}_{F}\tilde{Y} H_H^2(z_h) H_H\left(z_v^{2^{m-1}}\right)\sqrt{2}^{m-1} \prod_{k=1}^{m-2} H_L^2\left(z_h^{2^k}\right) \prod_{k=1}^{m-2} H_L\left(z_v^{2^k}\right)\right\}
\end{aligned}
$$

The recursive form of (3.20) to (3.24) follows from (3.32) to (3.34) in a similar way as obtaining (3.17) from (3.28).

3.7 Examples

In this section, the proposed technique is illustrated and its performance is evaluated using two experiments. In the first one, a square wave front is reconstructed by this numerical integration method based on measurements using the Fried model as the SH-WFS. The second is an example of reconstruction of a photograph from Hudgin aligned numerical measurements of its gradient data. These are to demonstrate performance when the sensor matches these models perfectly. A discussion on modelling error follows this section.

3.7.1 Atmospheric phase screen reconstruction

To evaluate the performance of the proposed algorithm for wave-front reconstruction, atmospheric phase screens are generated using a Kolmogorov turbulence model with r_0 of 0.25 m, wave length of 500 nm, pupil diameter of 32 m and outer scale of 22 m. These are the same parameters as in (Vogel and Yang 2006) with the exception of the outer scale which was not given in (Vogel and Yang 2006) and was chosen based on measurements at the Paranal ESO observatory in Chile (Martin and Conan 2000). The phase screen generation software was developed according to the discussion in

(Nakajima 1988). The gradient data (${}^{M}\widetilde{X}$ and ${}^{M}\widetilde{Y}$) are obtained by using a simulated sensor that uses the Fried aligned operators of (2.6) and (2.7). The phase screen is 128×128 so that the Fried sensor model will provide two 127×127 sets of data.

The phase screen is reconstructed using the proposed technique based on the simulated sensor data contaminated, where appropriate, with Gaussian white noise. The input SNR is obtained using ${}^{M}\widetilde{X}$ and ${}^{M}\widetilde{Y}$ as the signal, similarly to (Vogel and Yang 2006). To evaluate the capacity of the algorithm to remove high-frequency waffle, two reconstructions are produced for each test. The first is based on using the full decomposition of Figure 3.5 while in the second the HH quadrant is suppressed (i.e., equating all entries of ${}^{M-1}_{HH}\hat{\Phi}$ to zero). The quality of the reconstruction is evaluated using the root mean square (rms) normalized residual error of the pixels of a single phase screen calculated by

$$E_R = \frac{\left\|\Phi_{\text{Reconstructed}} - \Phi_{original}\right\|}{\left\|\Phi_{original}\right\|} \tag{3.35}$$

as well as the 66 Zernike modes (corresponding up to radial order 10) of a largest inner circle inside the 128×128 phase screen (Noll 1976).

In the first part of this experiment, the phase screen was reconstructed assuming no measurement noise. The full wave-front reconstruction using the proposed algorithm is shown in Figure 3.9(a). There is no visible difference between the full reconstruction and the reconstruction with neglected HH (and thus it is not included here). The error for the full reconstruction is shown in Figure 3.9(c) and is composed only of a pure checkerboard pattern corresponding to the ${}^{0}_{LL}\hat{\Phi}$ entry in the decomposition of Figure 3.5. This entry is invisible to the sensor and the algorithm does not attempt to reconstruct

this shape, just as it does not reconstruct the mean value. E_R for the full reconstruction is 0.000503 rms where the amplitude of the phase screen is normalized to 1 rms. For AO examples, ${}^{0}_{LL}\widetilde{\Phi}$ has been assumed to be zero. The assumption of ${}^{0}_{LL}\widetilde{\Phi}$ equal to zero is accurate since the simulated wave-front has a mean of zero. In AO, the mean does not contribute to the error of wave-front correction so its actual value is irrelevant.

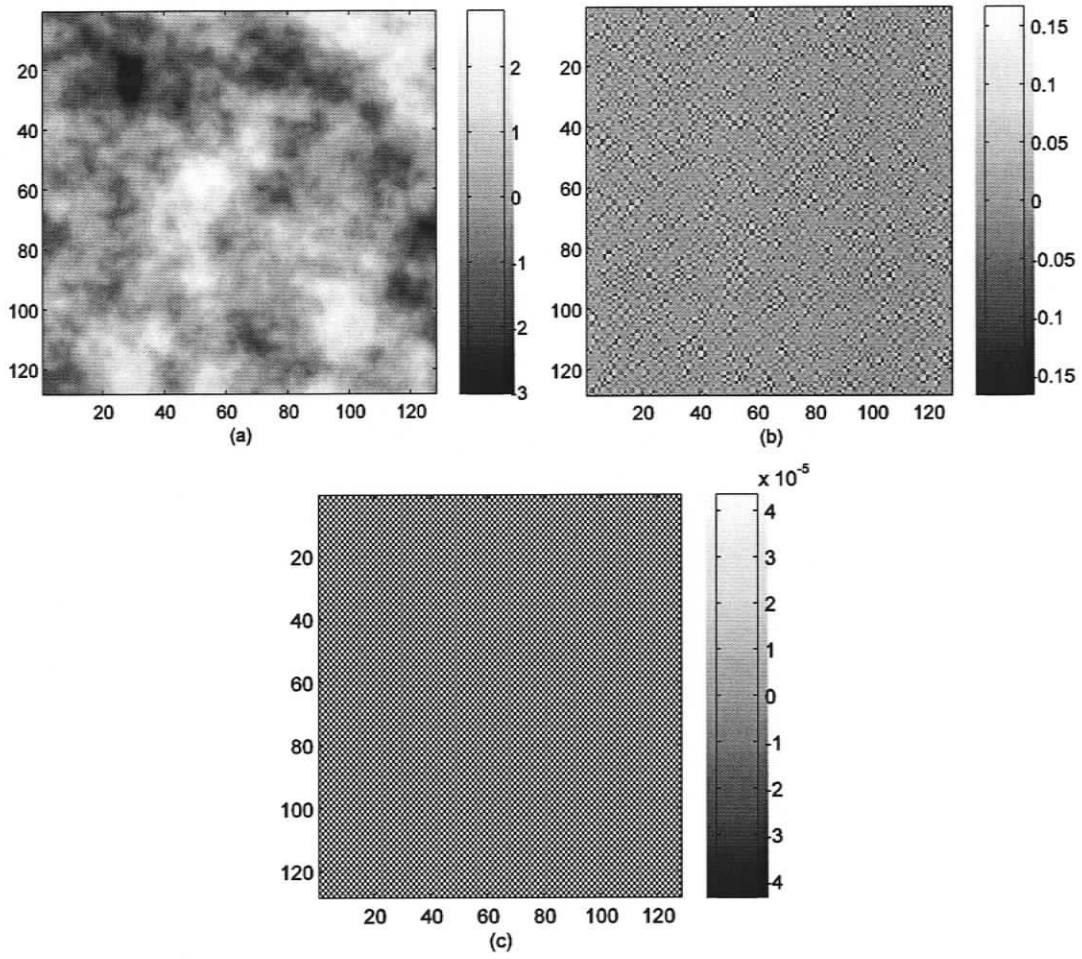

Figure 3.9: (a) Reconstructed wave-front; (b) residual error when HH quadrant is suppressed; (c) residual error from full reconstruction.

The error for the reconstruction neglecting HH is given in Figure 3.9(b) and the corresponding E_R is 0.0383 rms. The Zernike coefficients for the original phase screen are shown in Figure 3.10(a) and the Zernike coefficients of the reconstruction error for

the two reconstructions are shown in Figure 3.10(b). The low reconstruction error and the low relative magnitude of the Zernike coefficients of the reconstruction error indicate excellent reconstruction. Further, when there is no measurement noise, the full reconstruction gives better results.

In the second part of this experiment, Gaussian white measurement noise with SNR = 20 (i.e., 26 dB) is added to the gradient data and the phase screen is reconstructed from the noisy gradient data as in (Vogel and Yang 2006). The E_R for the full reconstruction is 0.0232 rms and the reconstruction neglecting HH is 0.0495 rms. Further, the Zernike coefficients of the reconstruction error for the full reconstruction and the reconstruction neglecting HH are shown in Figure 3.10(c). For comparison, the Zernike coefficients of the reconstruction error obtained from directly projecting the gradient data to Zernike modes are shown as the dark line. For the 66 modes shown, the rms of the Zernike coefficients of the reconstruction error is only ~1.5 times larger than the rms of the Zernike coefficients of the measurement noise.

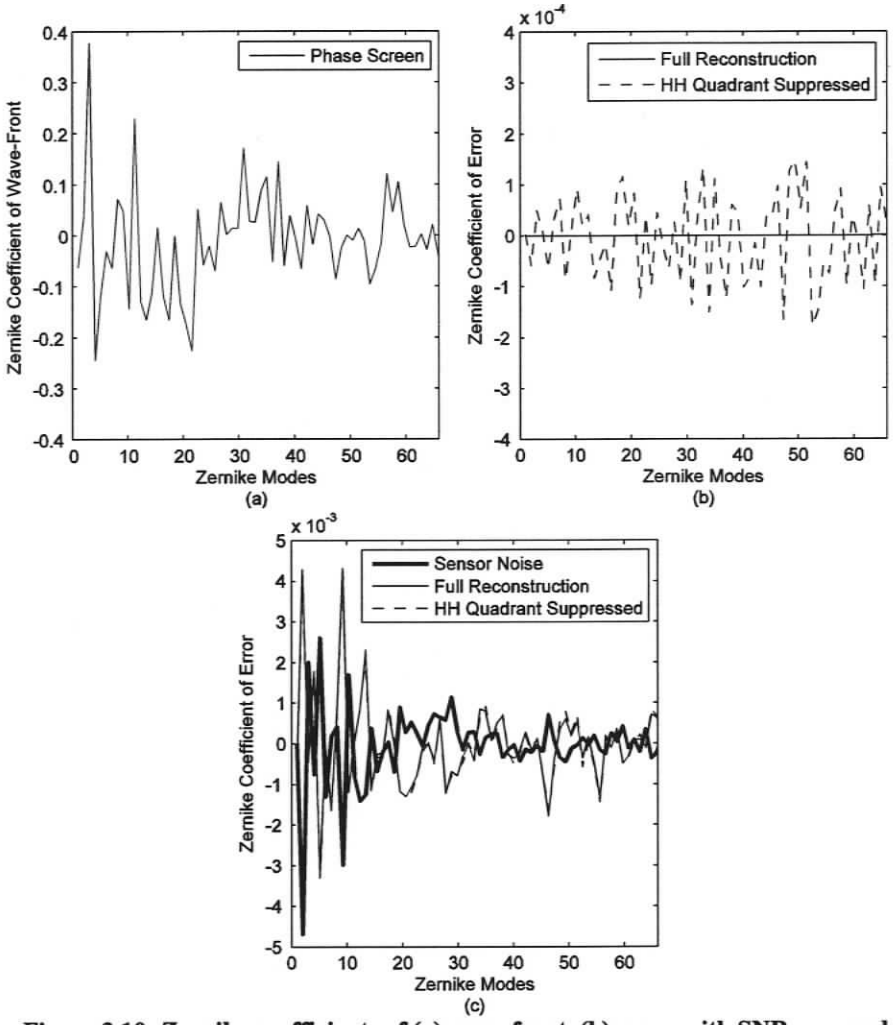

Figure 3.10: Zernike coefficients of (a) wave front, (b) error with SNR = ∞, and (c) error with SNR = 20.

To further evaluate the performance of the proposed algorithm in the presence of measurement noise, a series of reconstructions using noisy gradient data with several values of SNR were performed for both the full reconstruction and the reconstruction neglecting HH. The results with respect to the normalized residual error are presented in Figure 3.11.

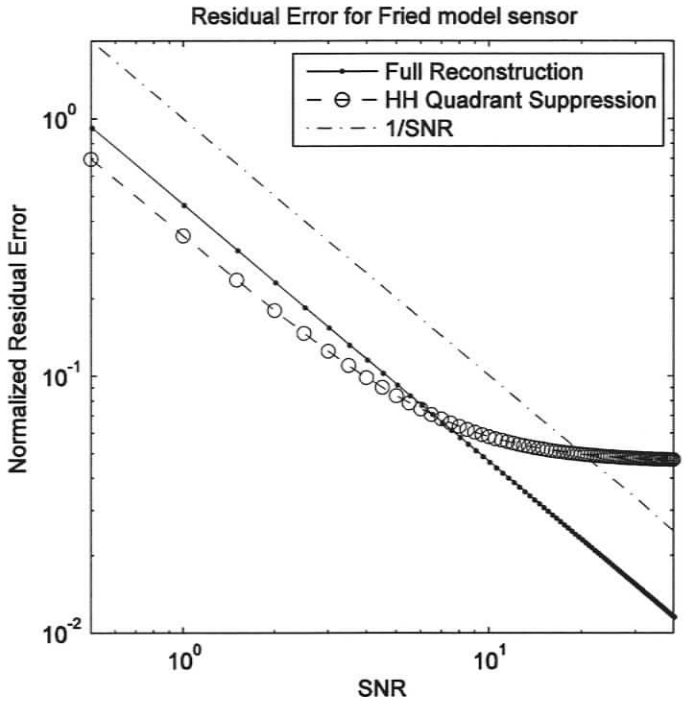

Figure 3.11: Reconstruction error compared to the inverse of the SNR.

To evaluate the computational complexity of the proposed method and the effect of neglecting the HH quadrant on the computational speed, wave fronts of various resolutions were generated and the computation times are given as Figure 3.12. This shows that the computation time is linear with the number of pixels for large numbers of pixels. It also shows that the algorithm requires at least a 128×128 reconstruction for the computation time to be significant compared to MatLab's overhead because the time per pixel decreases significantly until that size. These computations were made using MatLab with an implementation that is not as efficient as the one discussed in Section 3.5 and thus the computation times are not indicative of computation time in a real time implementation.

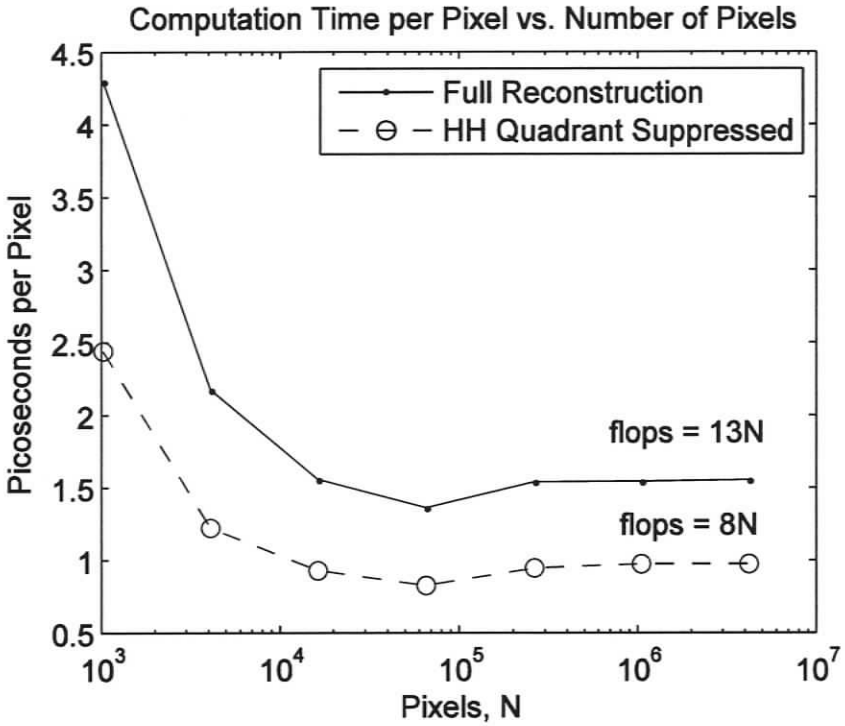

Figure 3.12: Computation time per pixel for a MatLab implementation.

3.7.2 Photograph reconstruction

A 2048×2048 image is being considered and the Hudgin aligned gradient data, ${}^{M}\widetilde{X}_{H}$ and ${}^{M}\widetilde{Y}_{H}$, are obtained using (2.8) and (2.9). The Hudgin model measures the difference between two neighbour pixels in the horizontal and vertical directions. The proposed algorithm can solve for the original data by first converting this data to Fried aligned data as in (3.3) and (3.4).

Atmospheric turbulence follows a known statistical model. In general, photographs do not. This experiment shows the algorithm's performance when solving discrete partial differential equations for unknown statistics of the data.

The image is reconstructed from noiseless gradient data using the proposed technique. The full reconstruction of the 8-bit image is perfect after rounding to integer values.

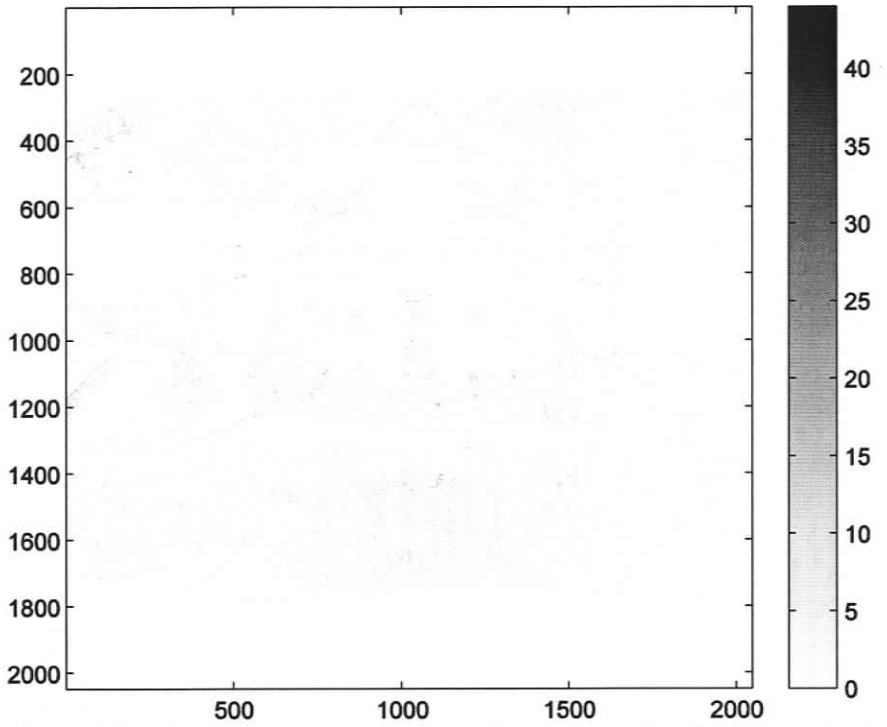

Figure 3.13: Absolute reconstruction error from suppressing the HH quadrant.

Figure 3.13 shows the error when the data from the HH quadrant, i.e., ${}^{M-1}_{HH}\hat{\Phi}$, is not computed. The dynamic range of this error is about 88 (compared to 255 dynamic range of image). However, the rms of the reconstruction error when ignoring the HH quadrant is only -31 dB (0.028 rms).

Figure 3.14: Reconstruction from gradient data corrupted to SNR = 2 (or 6 dB).

To illustrate the effect of noise, the reconstruction of Figure 3.14 is carried out using noisy gradient data. The reconstructed image using the proposed technique is shown for the case when the SNR of the noisy gradient data is 6dB (i.e., SNR = 2). The reconstruction error in Figure 3.14 is -20.3 dB (0.095 rms). This is a visually acceptable reconstruction despite significant input noise.

Measurement noise strongly affects the quality of the estimate of data in the HH quadrant. Figure 3.15 shows the SNR of the gradient data used as input and the SNR of the reconstruction error for both cases: full reconstruction, and reconstruction while neglecting the HH quadrant. The full reconstruction error is shown to be -14.3 dB smaller than the inverse of the SNR. This means that the ratio of the reconstructed image to the reconstructed error is 5.2 times greater than the input signal to noise ratio. This result is independent of any denoising techniques that could be applied.

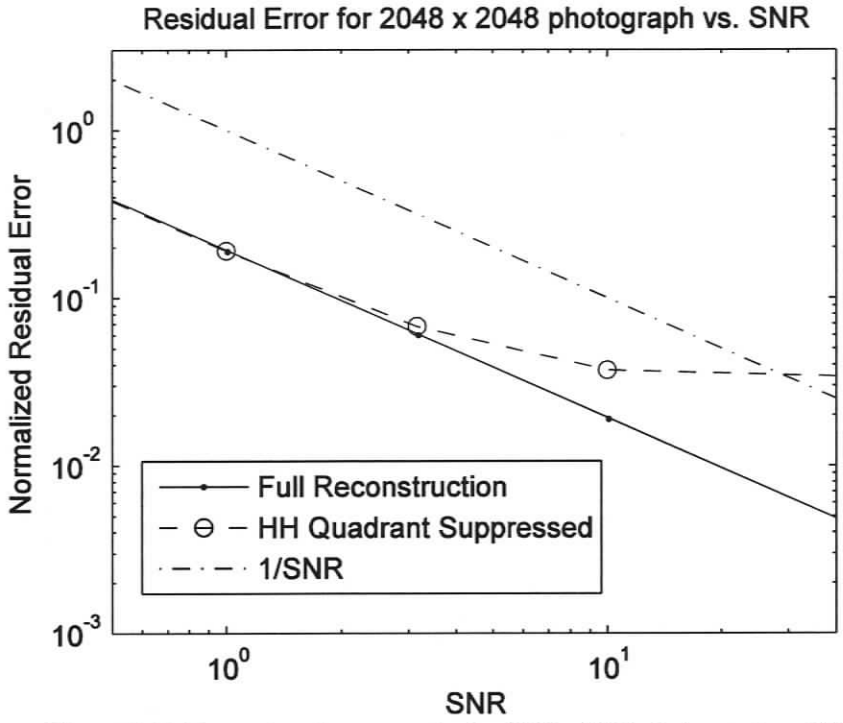

Figure 3.15: Reconstruction error for the 2048 x 2048 photograph vs. SNR.

It can be seen from Figure 3.15 that when the input SNR is reduced to about 3, the $\hat{\Phi}_{HH}^{M-1}$ estimate becomes corrupted and the full reconstruction does not give lower reconstruction error than the reconstruction neglecting the HH quadrant.

Remark: In this example, the mean value is assumed not known and was set to zero for reconstruction. The reconstruction has a dynamic range of -75.959 to 179.041. Since the original image is an 8-bit image and coincidentally used the full 0 to 255 dynamic range, this implies that the mean should be chosen to be 75.959. However, the reported error ignores mean value error.

3.8 Discussion

These examples indicate that the proposed reconstruction technique is a fast and efficient method to reconstruct simulated wave-fronts and/or images from gradient data when the application model matches the Hudgin or Fried models. By applying the method to

photographs and to atmospheric turbulence models, it shows that the reconstruction is not dependant on assumed statistical properties of the data. For the 128×128 case, the full reconstruction improved the SNR by a factor of approximately 2 as opposed to approximately 5 for the 2048×2048 case. This suggests that noise rejection may improve as the data sets get larger and may help mitigate the effect of dividing the number of collected photons between more sensor points.

An important consideration is to determine how well the algorithm performs when the model does not match the sensor. The Fried model is only an approximation of how a real SH-WFS would operate. Naturally, reconstruction errors are introduced when the Fried sensor model does not match the SH-WFS.

The proposed algorithm was also operated on an alternate model of the SH-WFS. This model used by (Vogel and Yang 2006) takes a 4×4 block of phase screen pixels and determines the average slope in the horizontal and vertical directions. Due to mathematical cancelling from numerical integration of the difference of pixels, our interpretation of this sensor is given as (3.36) and (3.37).

$$x_{p,q} = \frac{1}{3}\sum_i\sum_k \begin{bmatrix} -\phi_{i,k} & 0 & 0 & \phi_{i,k+3} \\ -\phi_{i+1,k} & 0 & 0 & \phi_{i+1,k+3} \\ -\phi_{i+2,k} & 0 & 0 & \phi_{i+2,k+3} \\ -\phi_{i+3,k} & 0 & 0 & \phi_{i+3,k+3} \end{bmatrix} \tag{3.36}$$

$$y_{p,q} = \frac{1}{3}\sum_i\sum_k \begin{bmatrix} -\phi_{i,k} & -\phi_{i,k+1} & -\phi_{i,k+2} & -\phi_{i,k+3} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \phi_{i+3,k} & \phi_{i+3,k+1} & \phi_{i+3,k+2} & \phi_{i+3,k+3} \end{bmatrix} \tag{3.37}$$

where p and q are the row and column of the measurement point on the sensor grid and $i = 4p-1$ and $k = 4p-1$ are the row and column representing the top corner of the

respective 4×4 cell of the phase screen. The indices i and k were chosen so that a 512×512 phase screen would generate 127×127 sensor grid points to sense 128×128 blocks in Fried alignment.

This shows that alternate valid SH-WFS models may be very different from the assumed Fried model. The main issue making this model different from the Fried model is that the Fried model would consist of 2 sets of 16×16 nonzero weighting factors. This alternate model has weighting from only 8 pixels for each direction and thus this model only uses 12.5% of the area of support of the Fried model. The gradient data generated by both methods were calculated for the same phase screen and it was determined that the difference between the two data sets was 0.416 rms, normalized to the ideal Fried measurements. This modelling error is similar to SNR = 2.4, though it is dependent on the phase screen itself and is not a random signal.

Gradient measurements were generated by (3.36) and (3.37) for a 127×127 sensor grid. Even with these large modelling errors that are not independent of the signal, the full reconstruction generated with the proposed algorithm had 0.136 rms reconstruction error. The reconstruction neglecting the HH quadrant resulted in a reduced error of 0.081 rms. For comparison, these relative errors are consistent with the results shown in Fig. 10 for a SNR value of about 3.

To improve the reconstruction using the proposed algorithm in the case of model mismatch, a simple lowpass filter was applied to the gradient data using the alternate model. This lowpass filtering reduced the difference between the gradient data from the alternate model and the Fried model to 0.227 rms (from 0.416 rms in the previous experiment). Using the filtered data, the full reconstruction resulted in an error of 0.055

rms and the reconstruction neglecting HH in an error of 0.060 rms. For comparison, in (Vogel and Yang 2006), using the MGCG the reconstruction error was improved from 0.20 rms to less than 0.01 rms by using curvature regularization. Similarly, the performance of the proposed algorithm in the presence of the gradient model mismatch is expected to improve by reducing the effects of model mismatch. Since even a simple lowpass filter improved the results, it is expected that implementing filters minimizing the curvature, such as in MGCG, would further improve results. Curvature regularization must be explicitly implemented in order to be a least-squares solution.

Another area of possible improvement of the proposed algorithm is the use of wavelets to denoise a wave-front using the HH part as shown by (Donoho 1995) instead of simply eliminating the HH part. As was shown in Section 3.3, an intermediate step for solving the wave front with the proposed algorithm is the Haar wavelet decomposition. Therefore, denoising techniques could be performed at this stage before proceeding to the final estimate to alleviate the effects of measurement noise in the input. An alternate approach to denoising suggested by (Kolaczyk 1997) is to threshold the noise based on the Poisson nature of photon noise.

Finally, AO systems operate on gradient data that have circular support. It has been shown in the literature (Freischlad and Koliopoulos 1986), (Poyneer, Gavel and Brase 2002) and (Talmi and Ribak 2006) that gradient data with circular support can be extrapolated to a square for use with the square Fourier transform and Fourier series based reconstruction methods. The solution for this problem is shown in Chapter 4.

3.9 Conclusions

A new technique for numerical integration from gradient measurements with square support is presented. This technique is based on obtaining the Haar wavelet image decomposition of the integral of the gradients by appropriate filtering and down-sampling of the gradient measurement data. The filters required for filtering are obtained by modifying the standard 2-D Haar filterbanks to use gradient data instead of image data as input. The use of the wavelet decomposition leads to an algorithm shown to be of complexity of $O(N)$, where N is the number of data points in the square reconstruction. It is asserted that this can be implemented in $10.17N$ for full reconstruction and as few as $5.33N$ when disregarding the high-frequency components. The particular implementation for these experiments was $13N$ for full reconstruction and $8N$ for the HH quadrant suppressed reconstruction. An added benefit of using this decomposition as an intermediate step is that wavelet denoising techniques could be applied without the usual computation cost of obtaining the decomposition of the wave-front. The proposed technique is illustrated by two examples and a discussion on sensor model errors. The reconstruction has been shown to be accurate despite relatively large noise signals (e.g., −20.3 dB error from 6 dB SNR for the 4-megapixel photograph). Results indicate that the proposed technique is a computationally efficient and accurate technique for wave-front and/or image reconstruction from Hudgin and Fried sensor models.

Chapter 4

Gradient Extrapolation from Circular Apertures via the Mirror Equations

4.1 Abstract

In this chapter, methods for extrapolating gradient data outside a circular aperture from measurements obtained within a circular aperture are presented. The proposed methods are required to be computationally efficient and to avoid the excitation of additional waffle modes in Fried alignment. It is shown that by using an octagon as an intermediate step from the circle to the square in the extrapolation process, the computations and/or residual reconstruction error can be reduced. The resulting computational cost is as low as $O(N^{1/2})$ where N is the number of measurement points. The performance of the extrapolation methods are studied in connection with the wavelet reconstruction method of the previous chapter. Experiments indicate that as expected, a significant reconstruction error occurs if no extrapolation is used. Further, the proposed extrapolation techniques lead to a reconstruction with data that is marginally different from a pupil-masked reconstruction using data from a square aperture.

Material in this chapter is published as:
Hampton, P. J., P. Agathoklis, and C. Bradley. "Wave Front Reconstruction over a Circular Aperture using Gradient data Extrapolated via the Mirror Equations." *Applied Optics* 48, 2009, 4018-4030

4.2 Introduction

The method developed in Chapter 3 is based on obtaining the wavelet decomposition of numerical integration directly from the gradient measurements. This method has computational complexity of O(N) and it allows an effective elimination of waffle modes, which are often dominated by noise contamination (D. T. Gavel 2003). However, this technique, as well as techniques based on the Fourier transform and cosine series are based on gradient data being available on a square, which is not the case in astronomical applications where there is a circular pupil and the gradient data is available over this circular aperture.

For reconstruction using techniques that require gradient measurements on a square, the measurements available on a circular aperture have to be extrapolated to a square. The extrapolated data measurements have to satisfy the zero-curl condition, but beyond this, the actual values of the extrapolated phase are inconsequential because this extrapolated region will be ignored after reconstruction. Therefore, many solutions to the extrapolation problem exist. The objective is to find those that require the least computation and do not magnify measurement noise and/or waffle modes. Waffle is the highest frequency that can be represented on a grid. It resembles a checkerboard pattern. These modes are well known to be of low magnitude in atmospheric turbulence (Noll 1976) and it is desirable to remove this and local waffle shapes from the reconstruction (D. T. Gavel 2003).

Some of the possible techniques proposed so far for FFT is to assume that the extrapolation is periodic (Freischlad and Koliopoulos 1986). Another approach is to assume that the extrapolated region is flat (Poyneer, Gavel and Brase 2002). An

advantage of this approach is that the gradient data in the extrapolated region is zero. The only significant computation cost is to determine the correct transition from the pupil edge to the flat region. One approach uses a polar coordinate system with a logarithmic radial axis to extrapolate the central circle of data that is often blocked by the secondary mirror and/or the hole in the center of the primary mirror (Talmi and Ribak 2006).

It was reported by (Poyneer, Gavel and Brase 2002) that waffle modes could be excited when extrapolating in Fried alignment (Fried 1977). One solution presented is to convert the measured data to a modified-Hudgin geometry as an intermediate step. One cause of the waffle excitation in Fried alignment in (Poyneer, Gavel and Brase 2002) is that the extrapolation approach forces the gradients in the extrapolated region to be zero. This strict condition coupled with the invisibility of waffle may excite local waffle shapes during the extrapolation process.

In this chapter, new approaches for extrapolating gradient data from the circular aperture to a square are presented so that the extrapolated gradient data can be used for wave front reconstruction in AO. These extrapolation processes are based on the desired properties of the reconstructed phase (i.e., continuous phase, waffle not excited), rather than based on properties of the gradient data that leave the waffle as ambiguous until reconstruction occurs. The first one is based on extrapolating the gradient measurements so that the extrapolated gradients along the diagonals are zero (i.e., the reconstructed phase will be flat on the diagonals). This is the "extension" method that (Poyneer, Gavel and Brase 2002) applied to the two diagonal grids of the Fried geometry. This allows a simple calculation of the extrapolated data adjacent to the boundary and a simple extension beyond it using only data copying. The other approach proposed here is based

on considering the circular boundary as a mirror surface and extending the data using the mirror equations. A curved mirror can be used to extrapolate the data from the circle to an octagon and then flat mirrors can be used to extend the data from the octagon to a square. The resulting extrapolated data will be one of the many possible solutions satisfying the zero-curl condition.

The proposed mirror-based approach has several advantages. When using reflections, the extrapolation is based on data further in the interior of the pupil rather than only from the edge. When a gradient sensor does not satisfy the zero curl condition, the gradient data is reflected without being forced to have a zero-curl. Another advantage of the reflection approach is that part of the extrapolation can be performed using flat mirrors. In such a case, the gradient data can be simply copied outside the octagonal or rectangular data set with the proper signs, and thus greatly reduce the computational cost.

4.3 Extrapolation of gradient measurement

Given gradient measurements on a circular aperture, the goal of this chapter is to provide a computationally fast technique to extrapolate the gradient data to a square aperture so that they can be used for phase reconstruction. This can be done using the same extrapolation method from the circular aperture to the square. An alternate approach presented here is based on extrapolating from a circle to an octagon and from the octagon to the square. The motivation for this is that, as shown later, the extrapolation from the octagon to any square is computationally negligible. The individual steps of such an extrapolation from circle to octagon to a square are illustrated in Figure 4.1 and the detailed description of the equations is discussed in Sections 4.3.1 and 4.3.2.

The obtained gradients of Figure 4.1(e) can be processed by any reconstruction algorithm that requires a square set of gradient data. The reconstructed image shown in Figure 4.2 is then masked to the original circular aperture shown in Figure 4.3. The photograph used here is to demonstrate the structure of the reflections, which is harder to visualize using atmospheric turbulence phase screens.

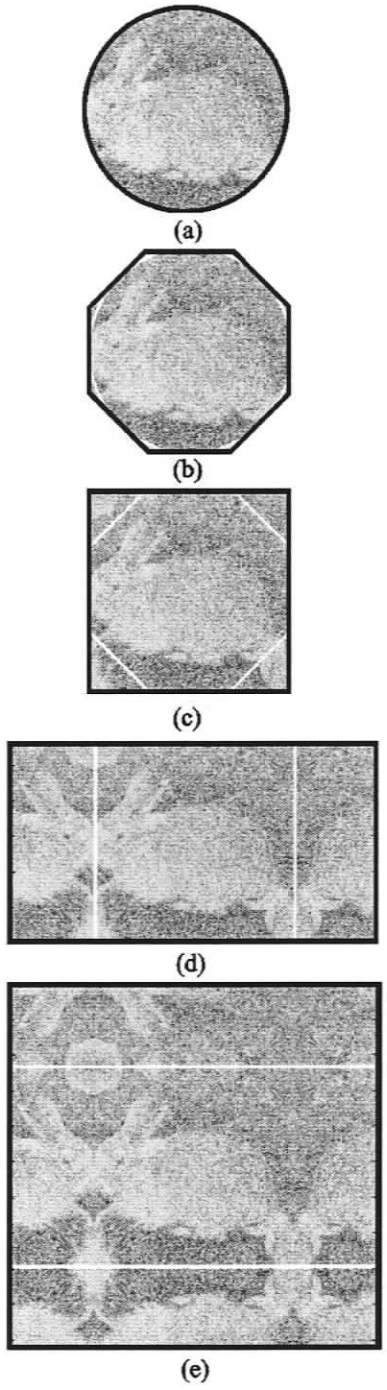

Figure 4.1: Extrapolation process of converting a circular gradient data set to any larger sized square. Solid white lines indicate simulated mirror locations. Solid black lines indicate the new data set edge. The figures display the logarithm of the magnitude of the gradients, which tend to look like pencil sketches for image gradient data. Reflections are to be processed on each of the two data sets individually.

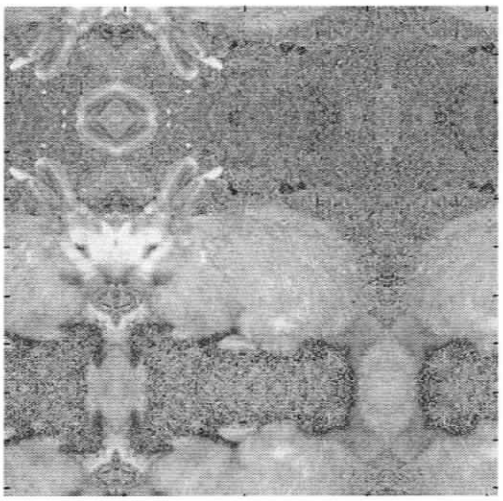

Figure 4.2: Reconstruction of data represented by Figure 4.1(e). Result resembles a kaleidoscope.

Figure 4.3: Pupil-masked representation of image shown in Figure 4.1.

4.3.1 Diagonally flat extrapolation

As mentioned earlier, there are many possible approaches to the extrapolation of gradient measurement data outside the circular aperture so that the zero-curl condition is satisfied. The one presented in this section is the "extension" method of (Poyneer, Gavel and Brase 2002) as applied to Fried alignment, where the gradients along the diagonal in the extrapolated region are zero, i.e., the pixels along the diagonal in the extrapolated region have the same intensity. The diagonal direction is the same in the upper-left and lower-right quadrants and upper-right and lower-left quadrants respectively. The case of extrapolating gradient data in the upper left quadrant is discussed in more detail here.

Consider gradient data modelled using the Fried model as presented in Section 2.2.1. It has been shown in (Poyneer, Gavel and Brase 2002) that these gradients will satisfy the zero-curl condition in a closed path consisting of a square rotated by 45 degrees as shown in Figure 4.4. For the upper left quadrant, $\dot{\varphi}_{p,q}$ is the gradient for the point to be extrapolated using the other three gradient measurements.

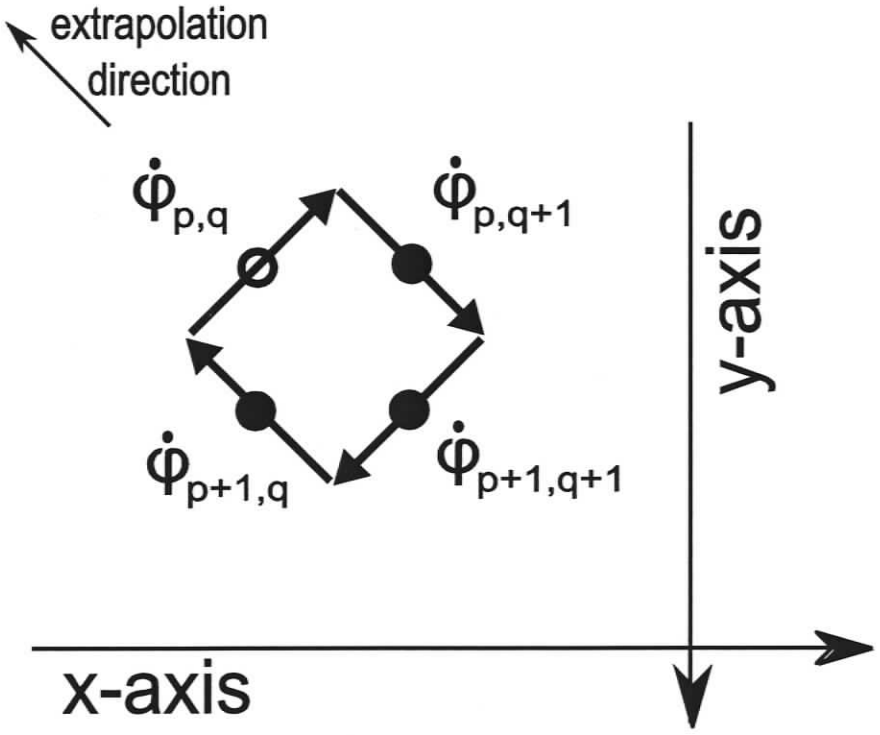

Figure 4.4: Smallest closed path of measured gradients in Fried alignment. Example of process in quadrant 2. Variable name convention holds for all quadrants.

For the zero-curl condition the x and y components of the gradients, $\dot{\varphi}$, need to be rotated by 45 degrees in the direction of the integration path and thus the scaling factors in the x and y directions can be ignored. Therefore, the equation for the closed path when the curl equals zero is equivalently represented by (4.1).

$$\begin{aligned}&\dot{\varphi}_{x,p,q} - \dot{\varphi}_{y,p,q} + \dot{\varphi}_{x,p,q+1} + \dot{\varphi}_{y,p,q+1} = \\ &\dot{\varphi}_{x,p+1,q+1} - \dot{\varphi}_{y,p+1,q+1} + \dot{\varphi}_{x,p+1,q} + \dot{\varphi}_{y,p+1,q}\end{aligned} \tag{4.1}$$

One possible choice is to require that the extrapolated gradient, $\dot{\varphi}$, is zero in the diagonal direction perpendicular to the path shown in Figure 4.4. This implies the condition:

$$\dot{\varphi}_{y,p,q} = -\dot{\varphi}_{x,p,q} \tag{4.2}$$

for quadrant 2, which when substituted in (4.1) leads to one equation for one unknown. The same approach can be used for points in the other three quadrants. The only difference is that condition (4.2) becomes:

$$\dot{\varphi}_{y,p,q+1} = \dot{\varphi}_{x,p,q+1} \tag{4.3}$$

$$\dot{\varphi}_{y,p+1,q} = \dot{\varphi}_{x,p+1,q} \tag{4.4}$$

$$\dot{\varphi}_{y,p+1,q+1} = -\dot{\varphi}_{x,p+1,q+1} \tag{4.5}$$

for quadrants 1, 3 and 4 respectively.

Once the extrapolation is extended beyond the edge of the circular aperture, then in the upper left quadrant, see Fig 5, it is $\dot{\phi}_{x,p,q}$ and $\dot{\phi}_{y,p,q}$ that must be extrapolated and $\dot{\phi}_{x,p+1,q+1}$ and $\dot{\phi}_{y,p+1,q+1}$ are the previously extrapolated data. Further, since $\dot{\phi}_{p+1,q}$ and $\dot{\phi}_{p,q+1}$ are outside the aperture, they are zero in the diagonal direction and thus contribute zero along the diagonal integration path. It follows then that $\dot{\phi}_{x,p,q} = \dot{\phi}_{x,p+1,q+1}$ and $\dot{\phi}_{y,p,q} = \dot{\phi}_{y,p+1,q+1}$ is a choice that satisfies the zero-curl condition. Clearly, extrapolation beyond the edge of the aperture requires merely data copying and no further calculations. A similar method is used for all other quadrants. This extrapolation technique is used only up to an area equal to or smaller than the smallest rectangle and further extrapolation is conducted with the reflection-based extrapolation technique in the following section.

4.3.2 Extrapolation by reflection

An alternative approach to obtaining a valid gradient field outside the circular aperture is based on using the mirror equations of optics. This is explained in the rest of this section.

4.3.2.1 Reflection using a curved mirror

Lemma 1: Given gradient data $\nabla \Phi(r_{in},\theta)$ measured inside a circular aperture of radius R, then the gradient data can be extrapolated by reflecting the measured data (the inside region) to the unmeasured region (the outside region) using (4.6):

$$\nabla\Phi(r_{out},\theta)=\frac{R^2}{r_{out}^2}\left(-\frac{\partial\Phi(r_{in},\theta)}{\partial r_{in}}\hat{u}_r+\frac{\partial\Phi(r_{in},\theta)}{r_{in}\partial\theta}\hat{u}_\theta\right) \tag{4.6}$$

where R is the radius of the curved mirror, $r_{in}\in r\le R$, $r_{out}\in r>R$. This extrapolation, which is based on the standard curved mirror equation with focal length equal to R, leads to a gradient field that satisfied the zero-curl condition.

Proof: From the curved mirror equation in optics, the relationship between the inner and outer radius is obtained in (4.7) and is shown in the diagram of Figure 4.5.

$$r_{in}=\frac{R^2}{r_{out}} \tag{4.7}$$

The reflection using a curved mirror is described using the simple equality of (4.8)

$$\Phi_{in}(r_{in},\theta)=\Phi_{out}(r_{out},\theta) \tag{4.8}$$

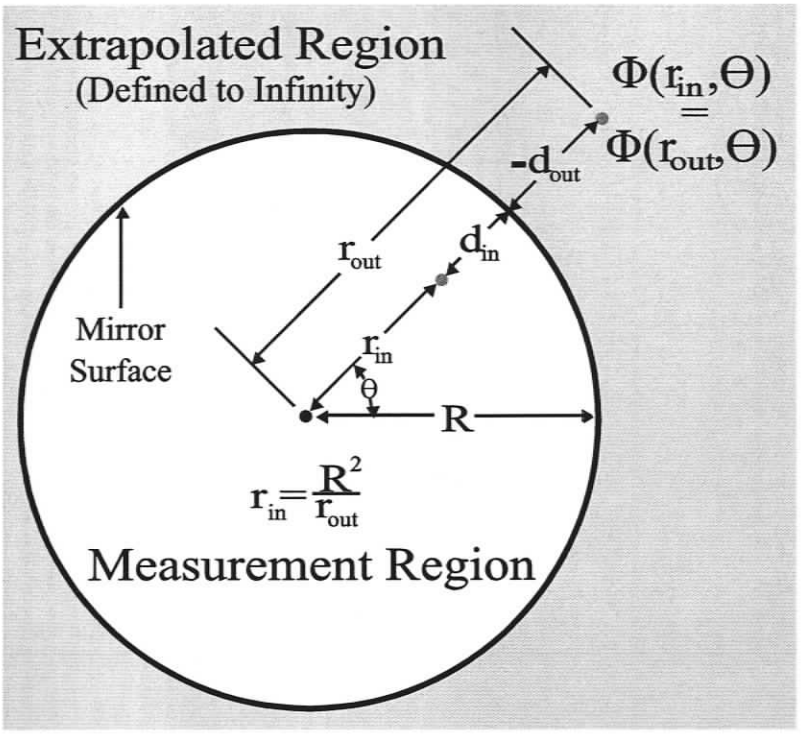

Figure 4.5: Fundamental relationships of curved mirror reflection.

and the polar gradient of $\Phi_{out}(r_{out},\theta)$ can be obtained using its partial derivative

$$\nabla\Phi(r_{out},\theta)=\frac{\partial\Phi(r_{out},\theta)}{\partial r_{out}}\hat{u}_r+\frac{\partial\Phi(r_{out},\theta)}{r_{out}\partial\theta}\hat{u}_\theta \tag{4.9}$$

By substituting (4.8) in (4.9), (4.10) is obtained:

$$\nabla\Phi(r_{out},\theta)=\frac{\partial\Phi(r_{in},\theta)}{\partial r_{out}}\hat{u}_r+\frac{\partial\Phi(r_{in},\theta)}{r_{out}\partial\theta}\hat{u}_\theta \tag{4.10}$$

which leads to (4.11):

$$\nabla\Phi(r_{out},\theta)=\left(\frac{\partial r_{in}}{\partial r_{out}}\right)\frac{\partial\Phi(r_{in},\theta)}{\partial r_{in}}\hat{u}_r+\left(\frac{r_{in}}{r_{out}}\right)\frac{\partial\Phi(r_{in},\theta)}{r_{in}\partial\theta}\hat{u}_\theta \tag{4.11}$$

This is a scaled version of the gradients inside the circular aperture. Substituting (4.7) and the derivative of (4.7) with respect to r_{out} into (4.11) followed by factoring, leads directly to (4.6). This relationship is determined as a gradient of a scalar function Φ

and therefore automatically satisfies the condition that the curl equals zero for any closed path.

Equation (4.6) defines the gradient field outside the circular aperture using the gradient data measured inside the circular aperture. In the next section it will be shown how this relationship can be used to extrapolate the gradient measurements when the measurements are carried out using a rectangular grid.

4.3.2.2 Curved mirror reflection on a rectangular grid.

Lemma 2: Given gradient data $\nabla \Phi(r_{in}, \theta)$ measured on a rectangular grid inside a circular aperture of radius R, then the data can be extrapolated by reflecting the measured data (the inside region) to the unmeasured region (the outside region) by using (4.12):

$$\begin{bmatrix} \frac{\partial \Phi_{out}}{\partial x} \\ \frac{\partial \Phi_{out}}{\partial y} \end{bmatrix} = B^{-1} G B A \vec{d} \tag{4.12}$$

where $\vec{d}$ is a vector of the nearest 4 gradient measurements to the location inside the circular aperture as shown in Figure 4.6.

$$\vec{d} = \begin{bmatrix} \dot{\varphi}_{x,p,q} & \dot{\varphi}_{x,p,q+1} & \dot{\varphi}_{x,p+1,q} & \dot{\varphi}_{x,p+1,q+1} & \dot{\varphi}_{y,p,q} & \dot{\varphi}_{y,p,q+1} & \dot{\varphi}_{y,p+1,q} & \dot{\varphi}_{y,p+1,q+1} \end{bmatrix}^T \tag{4.13}$$

Matrix A is a 2x8 interpolation matrix given by:

$$A = \begin{bmatrix} \vec{a} & null \\ null & \vec{a} \end{bmatrix} \tag{4.14}$$

where "null" is a 1x4 zero vector and

$$\vec{a} = \begin{bmatrix} (1-\Delta x)(1-\Delta y) & \Delta x(1-\Delta y) & (1-\Delta x)\Delta y & \Delta x \Delta y \end{bmatrix} \tag{4.15}$$

Matrix B represents the rotation matrix from Cartesian to polar co-ordinates given by

$$B = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \tag{4.16}$$

and matrix G is the gain factor applied to the inner data to scale it appropriately for use in the region outside the circular aperture:

$$G = \frac{R^2}{r_{out}^2}\begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} \tag{4.17}$$

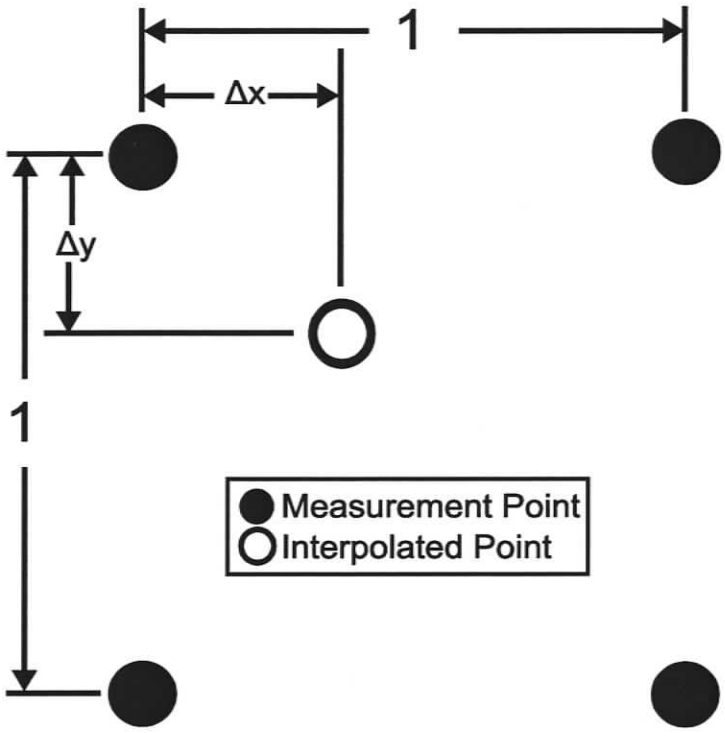

Figure 4.6: Example of desired measurement point between 4 measurement grid positions. Indices for measurement points increase towards the right for *p* and downwards for *q* in (4.14).

Equation (4.12) follows when (4.6) is used to obtain the gradient values on the rectangular grid outside the circular aperture. The resulting values for r_{in}, for a given r_{out}, on the grid will, in general, not be on the inside rectangular grid and therefore the corresponding inside gradient values will be obtained using an interpolation given by:

$$\vec{d}_{in} = A\vec{d} \tag{4.18}$$

where A and d are defined in equations (4.13) and (4.14) respectively. An example of the relative locations of the inside measurement points '•', outside extrapolated points 'x' and the inside interpolated points 'o' is given in Figure 4.7. The circular aperture in Figure 4.7 is the largest circle inscribed in the square.

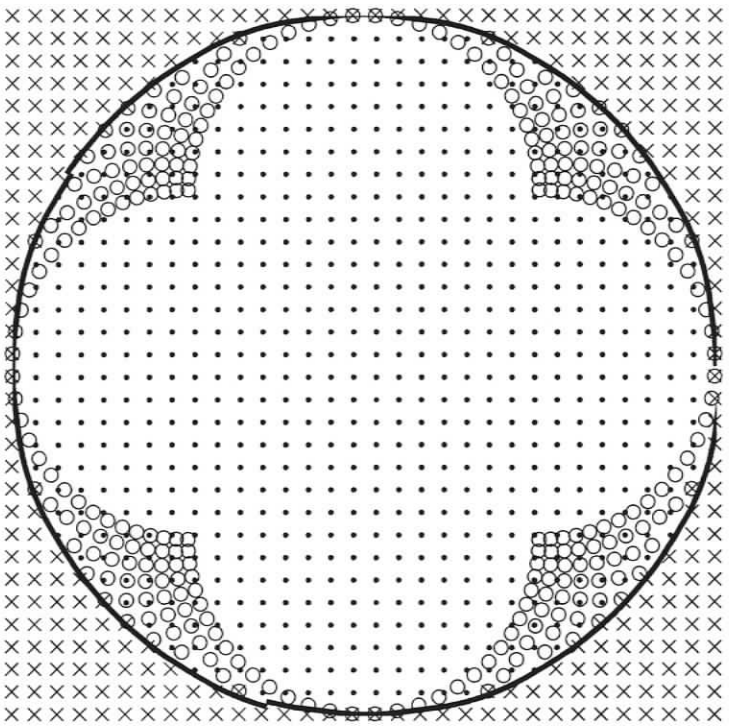

Figure 4.7: Relationship between extrapolated grid points, 'x', and the corresponding required measurement positions, 'o'. Measurement data, '·', do not coincide with 'o' so the process requires interpolation of the actual measurement points by matrix *A*. The circle is the edge of the pupil.

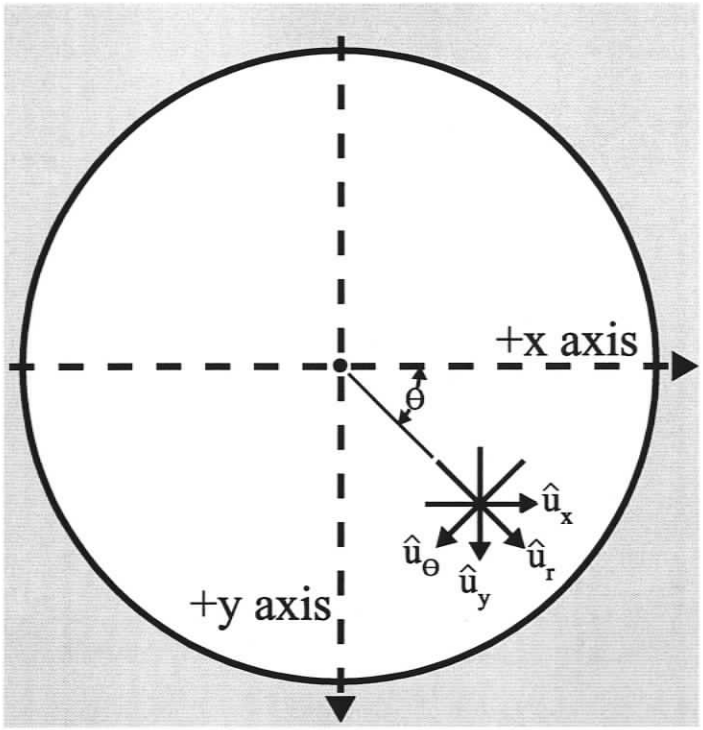

Figure 4.8: Example of relationship between co-ordinate systems.

After the gradient data has been interpolated, it must be rotated from Cartesian co-ordinates to polar co-ordinates. A single point example for the relationship between these co-ordinate systems is given in Figure 4.8 and the corresponding rotation matrix is given by B in (4.16).

This leads to

$$\begin{bmatrix} \dfrac{\partial\Phi(r,\theta)}{\partial r} \\ \dfrac{\partial\Phi(r,\theta)}{r\partial\theta} \end{bmatrix} = B \begin{bmatrix} \dfrac{\partial\Phi(r,\theta)}{\partial x} \\ \dfrac{\partial\Phi(r,\theta)}{\partial y} \end{bmatrix} \tag{4.19}$$

Once in polar co-ordinates, lemma 1 indicates that the gradient must be scaled by the factor of $(R/r_{out})^2$ and the sign of the radial partial derivative changed. This is carried out using matrix G in (4.16). Finally, the reflected gradient data is converted back to Cartesian co-ordinates using the inverse of matrix B.

For implementation, $B^{-1}GB$ can be implemented as a single 2×2 matrix for each extrapolation point that is computed offline. Matrix A would also be computed offline and implemented in a way to avoid all multiplications by 0. This reduces the computation cost to 12 multiplications and 9 additions per extrapolated grid point (i.e., 21 flops per extrapolation).

Remark: Extrapolation using curved mirrors can not be used to compute the gradient for data points immediately adjacent to the curved boundary. Since the technique uses reflections, the gradient perpendicular to the mirror will be near zero and the gradient tangential to the mirror will be nearly the same as the tangential gradient values near the mirror edge. A computationally simple method is, then, to interpolate these data points. This provides a result that is both close to the true gradient and provides lowpass filtering to alleviate any present measurement noise. This approach is computationally simpler than the one in (Poyneer, Gavel and Brase 2002), which is based on finding systems of linear equations to force the curl for a closed path to become zero.

4.3.2.3 Flat mirror reflection

Reflection of the gradient measurement data using a flat mirror is a very simple operation and it can easily be shown that the reflected gradient data will satisfy the zero curl condition. When the mirror surface is chosen to be parallel to the principal axes of the rectangular measurement grid, reflection involves only a positive or negative copying of the gradient data and does not involve any computations. If the mirror surface is chosen in any of the diagonal directions, reflection of the gradient data implies a rotation by 90^{o}. Rotating Cartesian gradient data by 90^{o} is simply a swap of the horizontal and vertical directions with the appropriate signs depending on the diagonal direction. Clearly flat

mirror reflection requires less computation than the curved mirror reflection discussed in the previous section. From the computational point of view, therefore, it is convenient to use the curved mirror extrapolation (or the diagonally flat from Section 4.3.1) to fill up the points from the circular aperture to an octagonal aperture and then continue with flat mirror reflections. For an octagonal aperture, the flat mirror surfaces are defined such that reflected outer pixels are positioned on the extension of the same rectangular grid as the inner pixels shown in Figure 4.9.

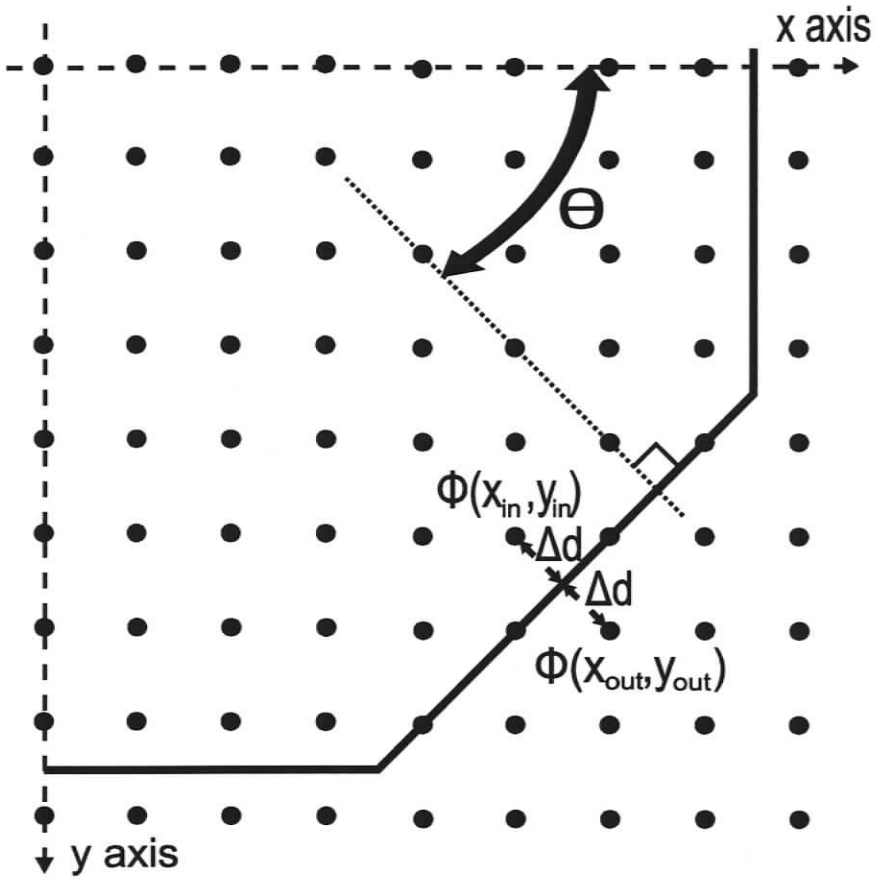

Figure 4.9: Lower-right quadrant of the octagon. Octagon is symmetric about the x and y axis.

The simple relationship between the inner gradient measurements and the outer gradient extrapolations can be given by (4.20)-(4.22)

$$\begin{bmatrix} \dfrac{\partial\Phi(x_{out},y_{out})}{\partial x} \\ \dfrac{\partial\Phi(x_{out},y_{out})}{\partial y} \end{bmatrix} = \begin{bmatrix} -\cos 2\theta & -\sin 2\theta \\ -\sin 2\theta & \cos 2\theta \end{bmatrix} \begin{bmatrix} \dfrac{\partial\Phi(x_{in},y_{in})}{\partial x} \\ \dfrac{\partial\Phi(x_{in},y_{in})}{\partial y} \end{bmatrix} \tag{4.20}$$

$$x_{in} = x_{out} - 2\Delta d\cos\theta \tag{4.21}$$

$$y_{in} = y_{out} - 2\Delta d\sin\theta \tag{4.22}$$

where θ is the angle from the positive x axis to the normal of the flat mirror, Figure 4.9. In this case $\theta=0.25\pi k$ where k is integer $k=0,..,7$. For k even, this is reduced to a mirror surface that is parallel to the Cartesian axes and, therefore, the extrapolation of the gradient measurements does not involve any interpolation. For k odd, the matrix in (4.20) has entries that are ±1 for the sine terms and 0 for the cosine terms that model the rotation due to the diagonal mirror.

The case $\Delta d = 0$ is a special case that occurs for points on the simulated diagonal mirror surface. For such points, the extrapolated gradient cannot be computed using reflection and is computed using the diagonally flat extrapolation approach discussed in Section 4.3.1.

Remark: It is interesting to note that although gradients in AO systems are considered to be measured using a circular aperture, the difference between the circular and an octagonal pupil is only a few points for low and medium resolutions. As an example, Figure 4.10 shows that using a WFS with 15×15 resolution, there are only 2 points per quadrant that are fully within the octagon but outside the circular aperture.

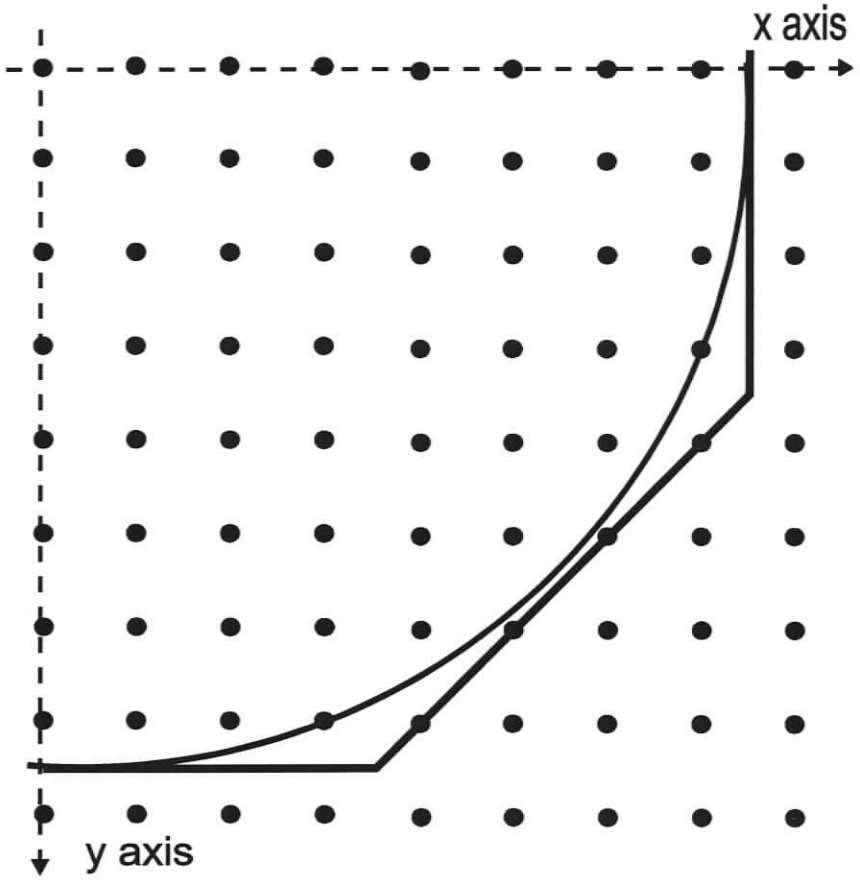

Figure 4.10: Circular and octagonal apertures on a square 15 x 15 grid. For this resolution, there are only 8 total points (2 per quadrant) that are fully within the octagon but outside the circular aperture. Points on the pupil edge are considered outside the pupil.

4.4 Extrapolation process

With the tools developed in the previous sections, the process of converting a circular aperture to a square one can be summarized in the block diagram of Figure 4.11.

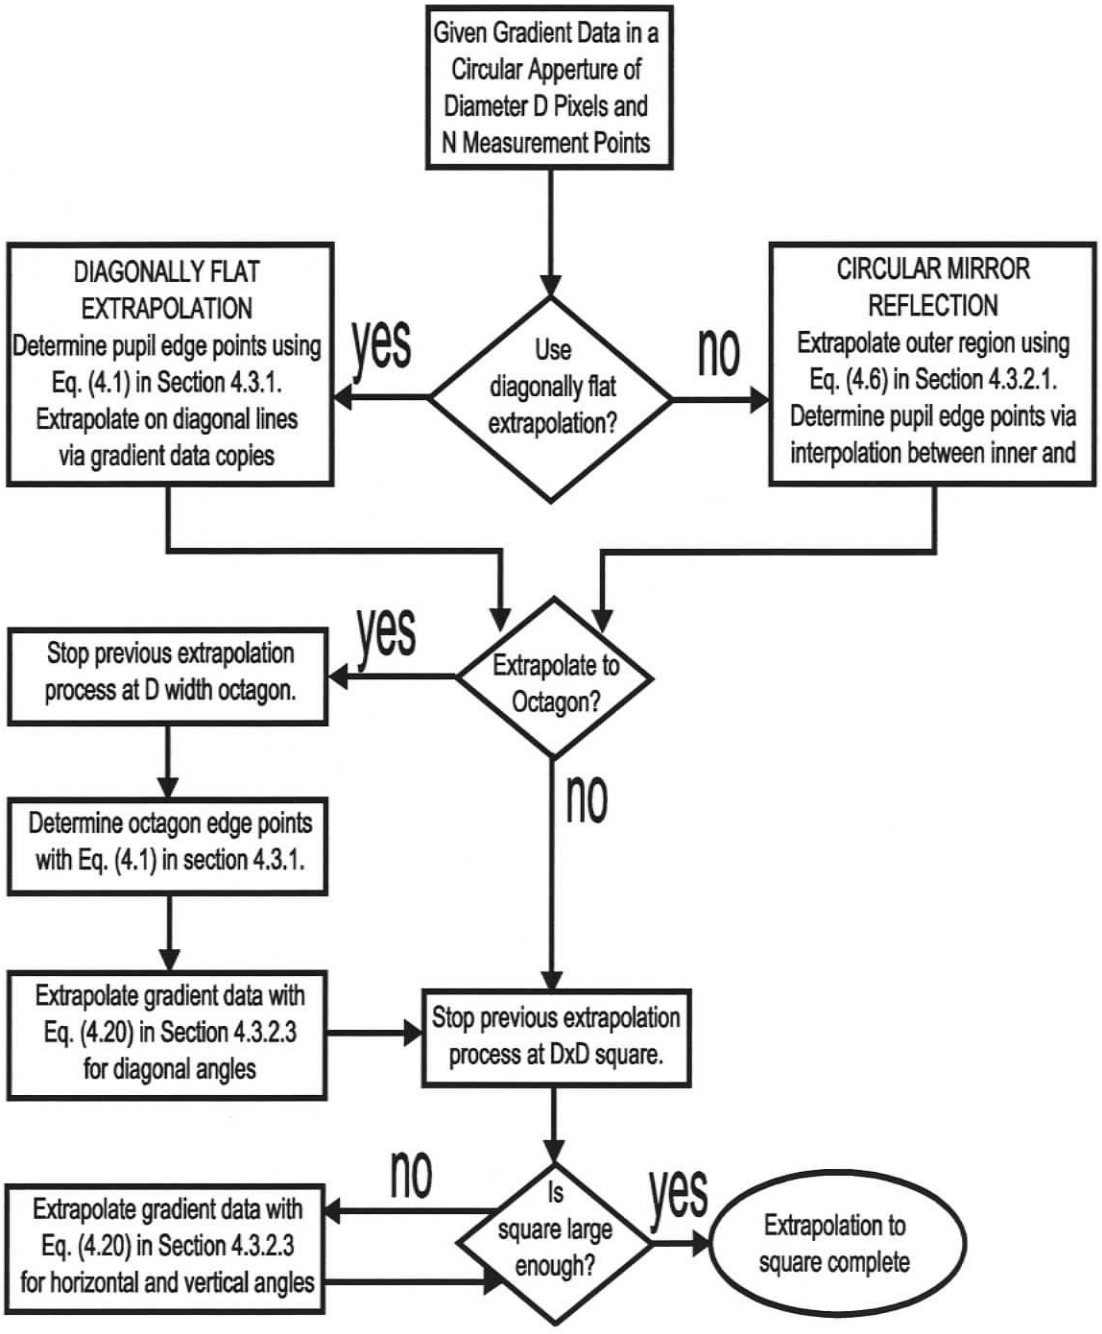

Figure 4.11: Decision flowchart for presented extrapolation options.

4.4.1 Computational complexity

The computational requirements for the extrapolation can be estimated as follows. The diagonally flat extrapolation requires computation only for the points on the boundary of the circular aperture. The computation required to evaluate both x and y gradients using (4.1) is 6 flops per data location. Every other extrapolated point using the diagonally flat method beyond the points on the boundary is merely a data copy. It is obvious that the number of data points on the circular boundary is the circumference measured in pixels. So, if there are N pixels within the circular aperture, then the number of pixels is approximately $D=2(N/\pi)^{1/2}$ on the diameter and πD on the boundary. Therefore, the computation cost of the diagonally flat extrapolation is approximately $21N^{1/2}$.

The extrapolation using the circular mirror equations in (4.12) has a cost of 21 flops per extrapolation point. Thus, the total number of operations depends on whether the extrapolation ends at the octagon, the smallest square or a square of any size. The area of the smallest square that contains the circular pupil is D^2 and the area of the smallest octagon that contains the circular pupil is $(8^{1/2} - 2)D^2$. The difference between the area of the smallest square and the circle is about $0.273N$ pixels and the area between the octagon and the circle is $0.055N$. Therefore, the cost of using the circular mirror approach is $5.74N$ to extrapolate to the smallest square and only $1.16N$ if this process is stopped at the octagon.

4.5 Experiments

The performance of the proposed extrapolation techniques is evaluated by a series of experiments using a simulated phase screen that represents aberrations caused by the Earth's atmosphere. This phase screen was generated using the same turbulence

parameters as the ones used in (Vogel and Yang 2006) and Chapter 3, which was published as (Hampton, Agathoklis and Bradley 2008).

Atmospheric phase screens are generated using a Kolmogorov turbulence model with r_0 of 0.25 m, wave length of 500 nm, pupil diameter of 32 m and outer scale of 22 m. These are the same parameters as in (Vogel and Yang 2006) and (Hampton, Agathoklis and Bradley 2008) with the exception of the outer scale which was not given in (Vogel and Yang 2006) and was chosen by (Hampton, Agathoklis and Bradley 2008) based on measurements at the Paranal ESO observatory in Chile (Martin and Conan 2000). The phase screen generation software was developed according to the discussion in (Nakajima 1988). The gradient data are obtained from the generated phase screen using a simulated sensor that uses the Fried aligned sensor model of (2.2) and (2.3). The gradient measurements are contaminated with additive zero mean Gaussian white noise as in (Gilles, Ellerbroek and Vogel 2002), (Vogel and Yang 2006), and (Hampton, Agathoklis and Bradley 2008) with various SNR values. The generated phase screen is 1024×1024 and the resulting square gradient measurement grid is 1023×1023. Phase screens smaller than 1024×1024 resolution are obtained via simple 2×2 averaging filters and down-sampling by 2 operations as in Figure 4.12.

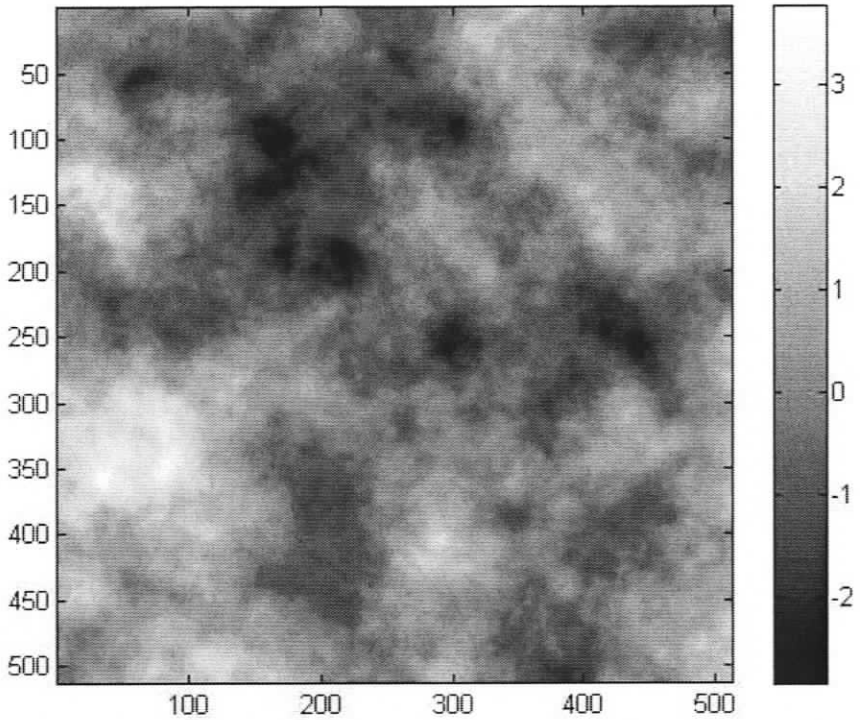

Figure 4.12: Atmospheric turbulence phase screen down-sampled to 512 x 512. Amplitude normalized to 1 unit rms. $r_0 = 0.25$m, $L_0 = 22$ m, 32 m width, 500 nm wavelength, 25 m/s wind speed.

Two simulation experiments will be discussed in the rest of this section. The first is using a 64×64 phase screen while the second is dealing with phase screens of sizes varying from 8×8 to 1024×1024.

4.5.1 Modern adaptive optics resolution

For the first set of simulation experiments the resolution of 64×64 pixels is chosen to conform to the current highest resolution deformable mirrors available, such as the Boston Micromachines 64×64 MEMS (Cornelissen, Bierden and Bifano 2006). A 64×64 phase screen is generated and the noisy gradient measurement data is obtained from the phase data contained in a circular aperture of 63 pixels diameter. The missing gradient measurements are extrapolated from the noisy gradients measurements inside the circular aperture using the techniques proposed in the previous section and the

extrapolated 63×63 gradient measurements are used for phase reconstruction. For the gradient extrapolation, four different options have been evaluated. They are given by all possible combinations of the flow chart in Figure 4.11. For comparison of the error when the phase screen is reconstructed without any extrapolation, in Figure 4.13 to Figure 4.16 the "No Extrapolation" case is included

The results from the first set of simulation experiments using the wavelet-based reconstruction method by (Hampton, Agathoklis and Bradley 2008) are shown in Figure 4.13. Figure 4.13 shows the normalized residual rms error for these experiments as a function of the gradient measurement SNR. It also compares the quality of the reconstruction using extrapolation to the ideal case, i.e., the quality of reconstruction using gradient measurements over the full square. Each plot in this section is obtained as the mean of 10 separate reconstructions that are each contaminated by a different noise profile. It can be seen from Figure 4.13 that using the diagonally flat extrapolation leads to nearly the same reconstruction quality as having gradient measurements over the full square. The curved mirror-based extrapolations always have an error above a minimum amount, which is indicated by the horizontal asymptotes as the SNR becomes high, i.e., there are reconstruction errors even when there is no measurement noise. This is due to interpolation. The diagonally flat extrapolations, on the other hand, do not have a measured horizontal asymptote, which implies that the reconstruction would be lossless when there is no noise.

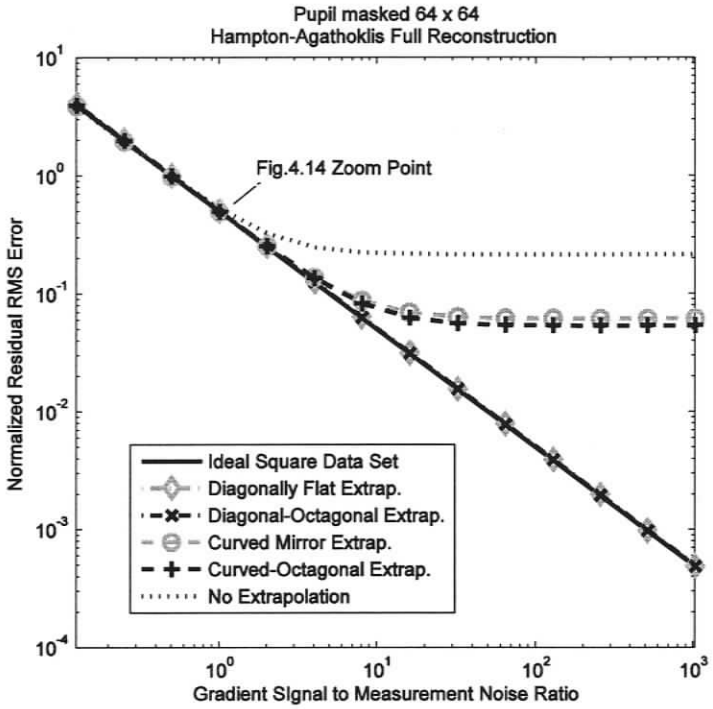

Figure 4.13: Residual error from using the proposed extrapolation methods with the Hampton-Agathoklis reconstruction method. For each pair of similar lines, the black line is better at high gradient SNR and the grey line is inferior. Figure 4.14 shows performance at low gradient SNR.

Figure 4.14 shows a close-up of the low SNR region in Figure 4.13, an area where all the methods converge to similar performance. This plot shows that the curved mirror extrapolation outperforms the others in terms of reconstruction error. It even surpasses the ideal square aperture case, likely due to the noise suppressing properties of the interpolation (lowpass) filters. A computationally faster method is the curved-octagonal extrapolation, which only uses the curved mirror equations to extrapolate gradients to the smallest octagon that surrounds the circle. The performance of this method is very close to that of the ideal square aperture case. This plot also shows that there is a performance improvement in reflecting the diagonal extrapolation with a simulated octagonal mirror. There is no additional computation cost to this decision since the diagonally flat process properly defines gradient on the octagonal mirror edge automatically.

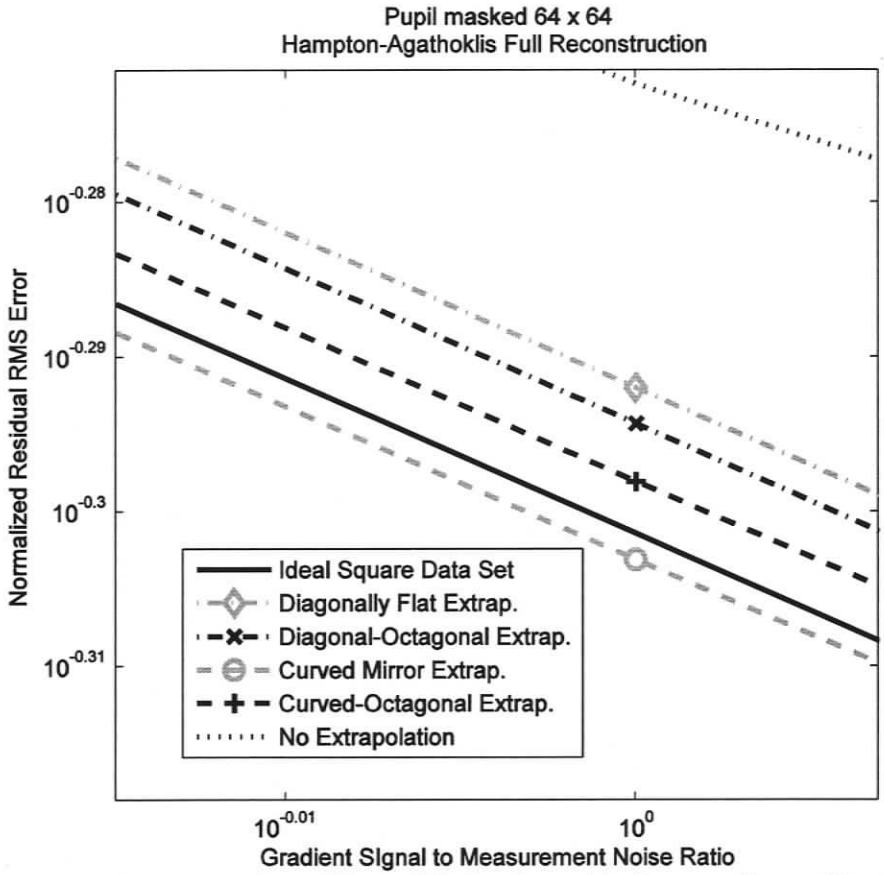

Figure 4.14: Zoomed representation of Figure 4.13 showing residual error at low gradient SNR. The curved mirror extrapolation was the worst method at high gradient SNR and is shown here to provide better reconstruction for low gradient SNR.

The results presented in Figure 4.13 and Figure 4.14 show that there are performance trade-offs between the various extrapolation approaches that are dependent on the SNR of the gradient measurements. The circular mirror reflections become better in the low SNR region for the Hampton-Agathoklis reconstruction method and the circular-octagonal approach is indistinguishable from the ideal square aperture case in this region.

One interesting question is the effect of the proposed extrapolation and reconstruction approach to the waffle modes. It is well known (D. T. Gavel 2003) that it is desirable to suppress the waffle modes associated with an AO correction. One of the

advantages of the wavelet-based reconstruction approach given in (Hampton, Agathoklis and Bradley 2008) is that in the wavelet decomposition of the phase obtained directly from the gradient sensor measurements, the waffle modes are organized into a well-defined quadrant. All waffle modes can easily be removed, therefore, from the reconstruction by ignoring this section of the decomposition. This provides the additional benefit of reducing the computation cost of the reconstruction by about 40%. To evaluate the impact of removing the waffle modes, the same approach as that presented in Figure 4.13 was used with the wavelet reconstruction of (Hampton, Agathoklis and Bradley 2008), with the waffle modes ignored. Results are shown in Figure 4.15 and Figure 4.16.

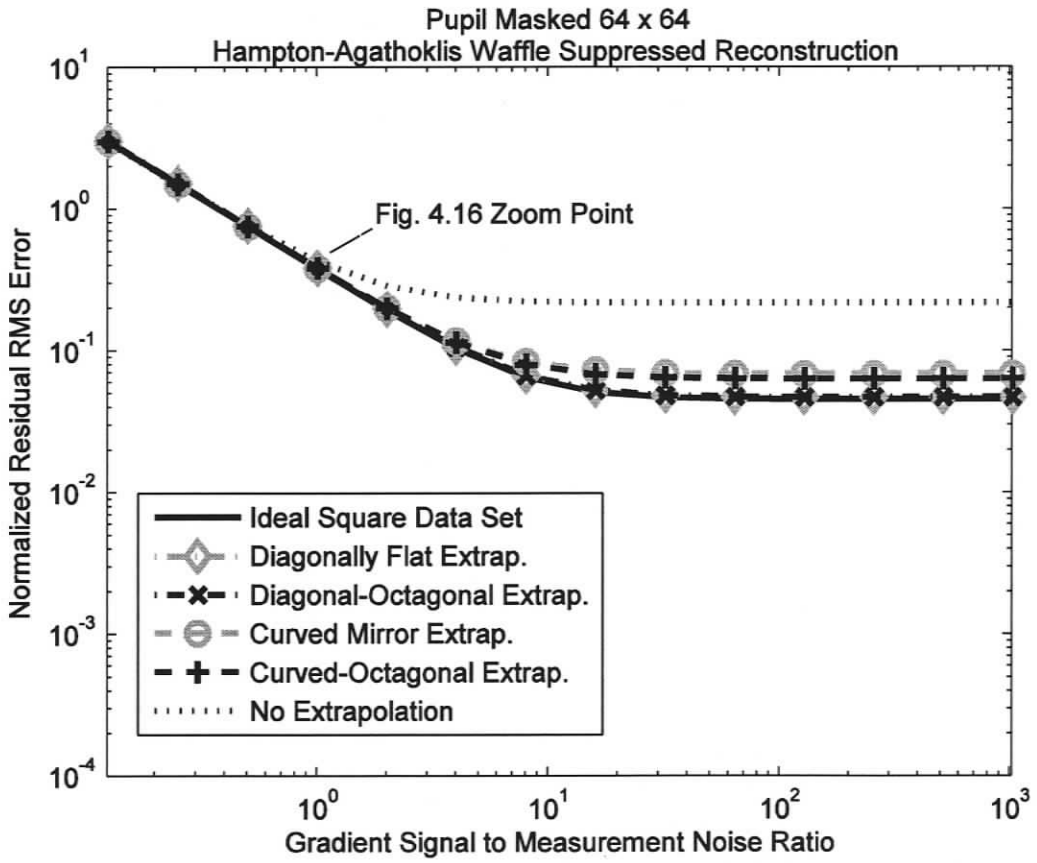

Figure 4.15: Residual error from using the proposed extrapolation methods and the Hampton-Agathoklis reconstruction method with all waffle modes suppressed. For each pair of similar lines, the black line is better at high gradient SNR and the grey line is inferior. Figure 4.16 shows performance at low gradient SNR.

Figure 4.15 indicates that the reconstruction based on gradient data extrapolated using a curved mirror is about the same as that in Figure 4.13. For the other extrapolation techniques, the removal of all waffle modes leads to a minimal error (of about 5% in this case) which cannot be overcome. Further, it can be seen that the diagonally flat-based extrapolations are very close to the ideal case of having measurements in a full square. This indicates that when the waffle modes are neglected, the diagonally flat extrapolation is as good as having gradient measurements for the full square. Figure 4.16 shows a close-up view of Figure 4.15 for low gradient SNR. It can be seen that for low SNR, the curved mirror extrapolations provide the best reconstruction.

The above results indicate that the proposed approach for wave front reconstruction on a circular aperture is promising in correcting aberrations using deformable mirrors (DM) with resolutions of 64×64 actuators such as the one in (Cornelissen, Bierden and Bifano 2006). The model of the interaction between a SH-WFS, DM, control computer and intervening electronics may be developed by the methods presented by (Hampton et al. 2008) and evaluation of the performance of the DM can be determined by the methods presented by (Wallace et al. 2006).

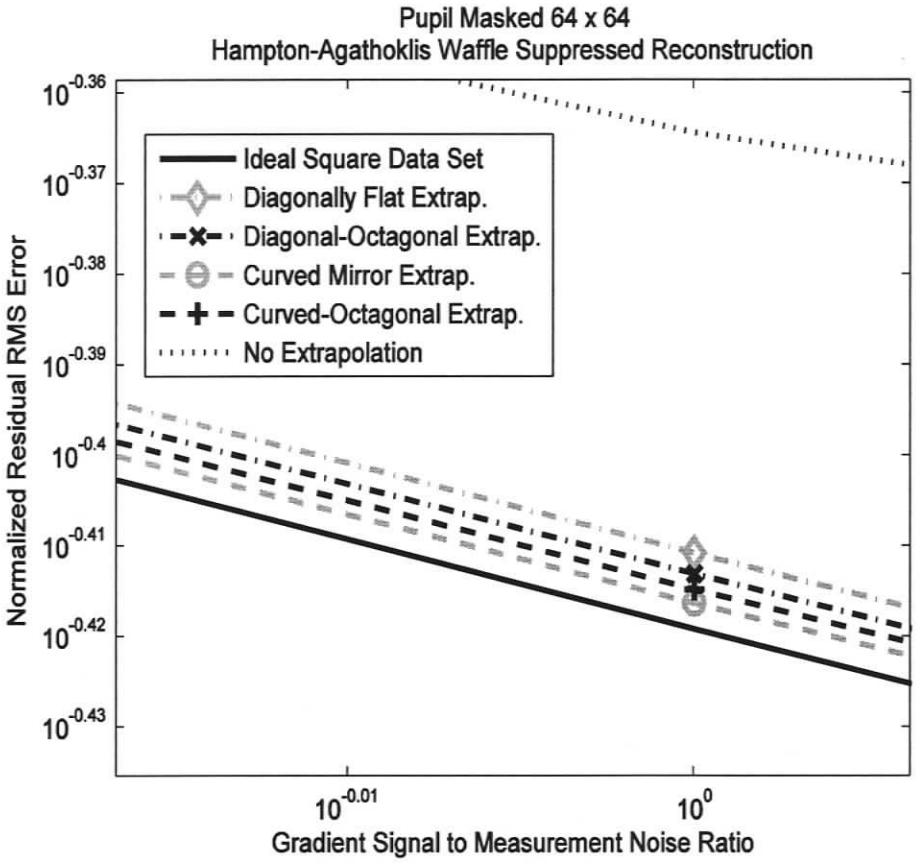

Figure 4.16: Zoomed representation of Figure 4.15 showing residual error at low gradient SNR when waffle modes are suppressed.

The next set of simulations compares the performance of the previously discussed extrapolation-reconstruction methods using noisy gradient measurements in a circular aperture. The results are presented in Figure 4.17. They indicate that the full reconstruction method performs better for gradient SNR greater than 10 while suppressing waffle modes performs better when the gradient SNR is 10 or less. These results are for gradient data defined as in the Fried model.

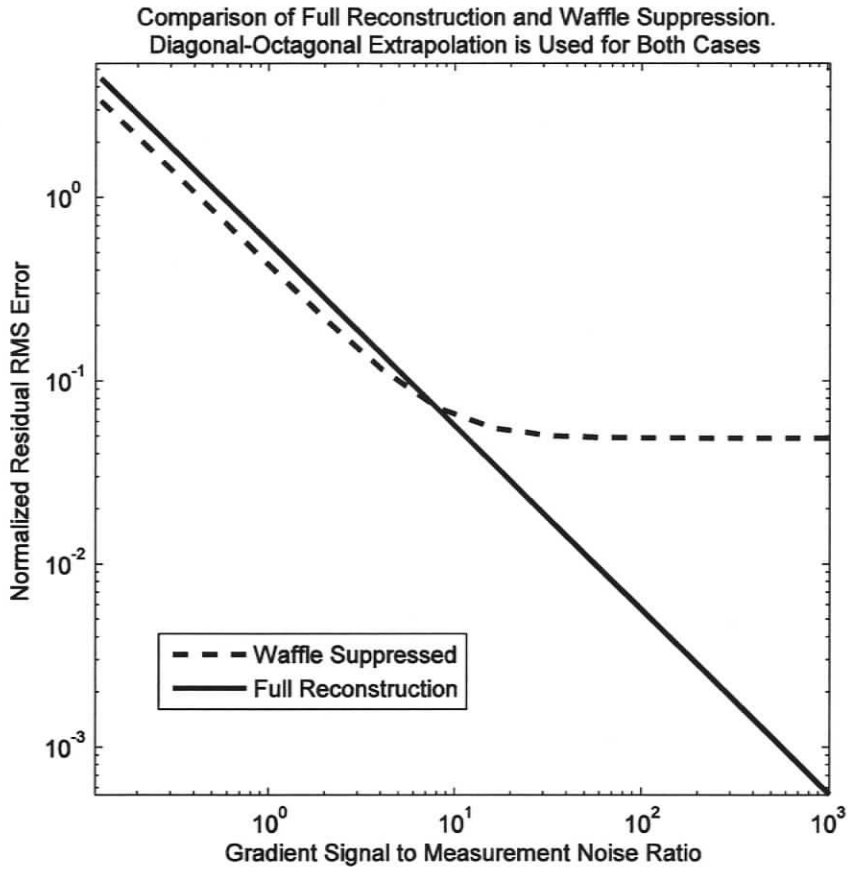

Figure 4.17: Comparison between the full reconstruction and waffle suppression. Both are combined with the diagonal-octagonal extrapolation method.

This performance of the Hampton-Agathoklis method indicates that this method could be a good choice for open loop AO control, such as the VOLT experiment by (Andersen, et al. 2008). Further, the Hampton-Agathoklis method could also be used for a fast estimation of the phase screen as an intermediate step applied to tomographic reconstruction techniques such as the ones presented by (Tokovinin and Viard 2001) or (D. Gavel 2004). Increased accuracy of the process is expected when used as the initial condition of the Poisson solver.

Remark: It was determined that the gradient extrapolation methods do not perform well for Fourier transform-based reconstruction algorithms. The reason for this is that the

methods presented in this work do not explicitly satisfy the strict requirement of the Fourier transform methods that the curl must be zero across the edge when the data is tiled. This is because it is not a criterion of the wavelet based approach. For a method that handles this requirement for Fourier transforms, the works of (Poyneer, Gavel and Brase 2002), and (Poyneer and Véran 2005) are recommended.

4.5.2 Reconstruction for varying resolutions

The second set of experiments deals with phase screens that are $2^m \times 2^m$ where m is integers from 3 to 10 in order to validate that the proposed approach can be used with a wide range of data sizes. In this set of experiments noisy gradient measurements over circular apertures with a diameter up to 1023 pixels were generated for a constant pupil size of 32 m, as discussed at the beginning of this section. These measurements were extrapolated using the diagonally flat extrapolation method and reconstructed using the Hampton-Agathoklis full reconstruction method. The results presented in Figure 4.18, demonstrate that the diagonally flat reconstruction process does not have a lower limit for the residual error when the gradient SNR is in the range of SNR = 0.1 to SNR = 1000. This indicates that for current AO systems with pixel diameters from 7 to 63 as well as future AO systems with higher resolutions up to 1023 pixel diameter, there is no limit on the quality of reconstruction due to extrapolation errors.

Figure 4.18 allows comparison of the accuracy of reconstruction among resolutions for a given gradient SNR. It is unlikely that the gradient SNR would remain equal as the resolution changes in a sensor application since the noise level may remain unchanged but the gradient signal strength would vary. A method shown in (Poyneer, Gavel and Brase 2002) is to determine the amount that the mean square (MS) noise is amplified

when converted to MS error of the reconstruction. A simplified version with noiseless reconstruction error of zero is (4.23).

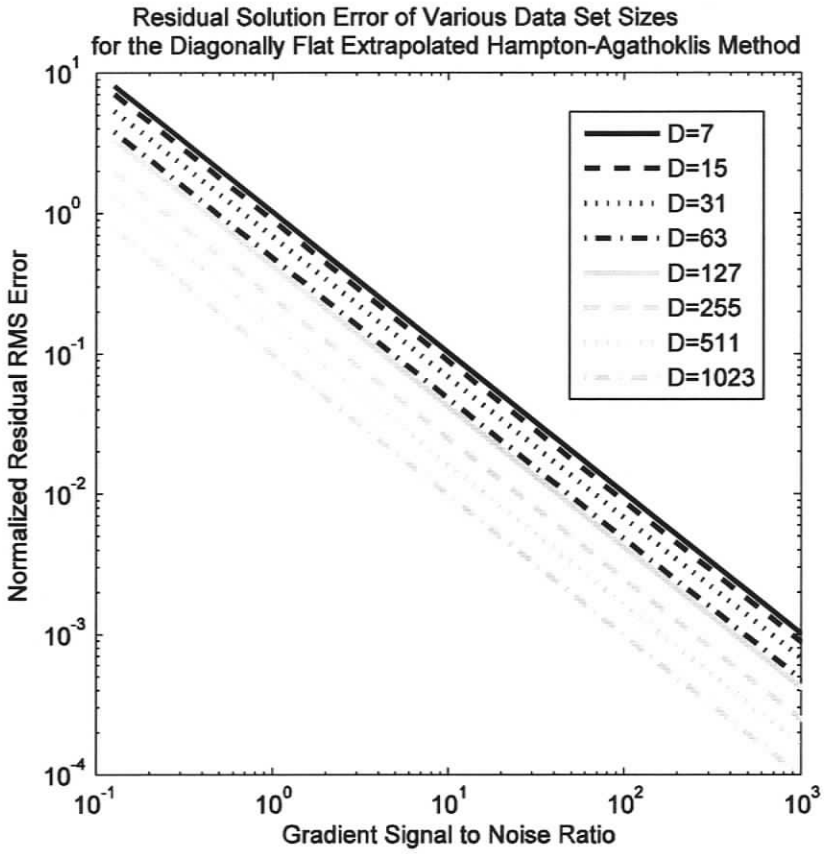

Figure 4.18: The residual reconstruction error of various sizes of data sets. Extrapolation method is diagonally flat. Reconstruction method is the Hampton-Agathoklis method. The unit of diameter, D, is pixels. The diameter of the pupil is constant. Constant noise and transmission signals chosen for SNR = 1 at D = 511 pixels. Pupil diameter is a constant 32 m.

$$\sigma_e^2 = \alpha\sigma_n^2 \tag{4.23}$$

where σ_e^2 is the MS error, σ_n^2 is the MS noise and α is the amplification of the MS noise. Figure 4.19 shows how much the MS noise is amplified when extrapolated by the diagonal-octagonal method followed by full reconstruction by the Hampton-Agathoklis method. This shows that when the MS of noise is maintained at a constant for different

resolutions, the MS error increases nearly linearly with respect to the logarithm of the number of sensors. The quadratic fit of data points in Figure 4.19 is (4.24).

$$\alpha(N) = 0.0261 \ln^2 N + 0.2405 \ln N + 0.6833 \tag{4.24}$$

where $\ln N$ is the natural logarithm of the number of sensors on the square, N. The RMS error is the square root of the MS error, so the amplification of a noise signal into RMS error would increase logarithmically with the number of sensors.

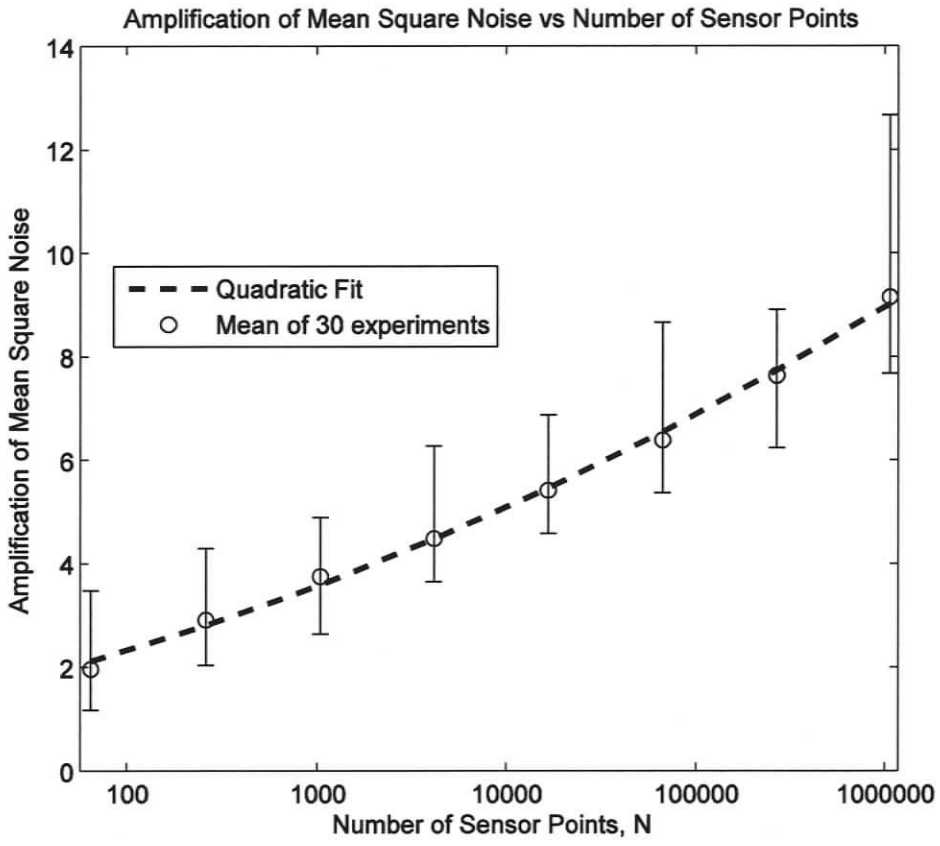

Figure 4.19: Amplification of ms noise vs. the number of sensor points. Error bars indicate maximum and minimum amplification in 30 noisy data sets at each resolution.

4.6 Conclusions

Techniques for extrapolating gradient measurements obtained over a circular aperture to a square aperture have been presented. The gradient data on a fictitious square aperture is

used to evaluate the performance of the proposed extrapolation processes with respect to computation time and reconstruction error inside the circular aperture.

Results indicate that the diagonally flat extrapolation method was virtually indistinguishable from the unrealistic ideal case of having a square aperture when applied to the Hampton-Agathoklis method. The extrapolation using the diagonally flat techniques does not cause errors when there is no noise, as it is evident from the lack of a knee point in the residual error plots. Further, the computation cost of this extrapolation technique is $O(N^{1/2})$ and is expected to be insignificant compared to the $O(N)$ reconstruction method.

The curved mirror extrapolation techniques are shown to be better than even to the ideal square aperture case when the gradient SNR is very low. This is due to the noise filtering effect of interpolating data in the interior of the pupil to obtain the extrapolated data. This approach comes with a cost that is $O(N)$ flops.

The octagonal mirror reflection reduces the error of the diagonal extrapolation approach with no additional computation cost. This combination of methods performs best for high gradient SNR. The octagonal mirror reflection combined with the curved mirror extrapolation leads to a reduction in computation cost by at least a factor of 5 and results in reduced error for high SNR gradient data and amplified error in the low SNR gradient data.

It was found that when the gradient SNR is constant as resolution increases, the amount of reconstruction error decreases. If only the noise is kept constant instead, the noise amplification into rms error increases logarithmically with the number of sensors.

The examples indicate that Hampton-Agathoklis method (Hampton, Agathoklis and Bradley 2008) can be used for the reconstruction of gradient data from circular apertures with results as good as if data from a square aperture were available.

Chapter 5

Applications of Numerical Integration to Adaptive Optics

5.1 Abstract

It was shown in Chapter 3 and in (Hampton, Agathoklis and Bradley 2008) that the method to estimate the wavelet decomposition of the integral of partial derivatives provides lossless reconstructions from gradient data measured as numerical partial differences in the absence of noise. It was also shown that the method could attenuate the input noise and modelling errors. Chapter 4 and (Hampton, Agathoklis and Bradley 2009) showed methods of expanding the gradient data set from a circle to a square. In this chapter, the application of these algorithms (collectively referred to as wavelet integration) to an AO system model is described.

This chapter presents the application of wavelet integration to an AO simulator, with the objective of determining the conditions where the wavelet integration method is preferable to standard pseudo-inverse techniques. It is determined that the new control scheme produces more steady-state error than techniques in the literature. However, the amount of computational savings available gives the new approach an advantage when following a dynamic input. It is also shown that the decomposition form leads to a computationally efficient method for controlling multiple deformable mirrors (DMs) in an AO configuration dubbed a Woofer-Tweeter system.

5.2 Introduction to adaptive optics system components

The purpose of AO is to correct the aberrations on star light that are caused by the turbulent atmosphere of the Earth. A seminal work on AO, describing systems and atmospheric turbulence has been written by (F. Roddier 1999). In an AO system, the correction is typically performed by a combination of at least one deformable mirror (DM), wave front sensor (WFS) and a control computer arranged in the order shown in Figure 5.1.

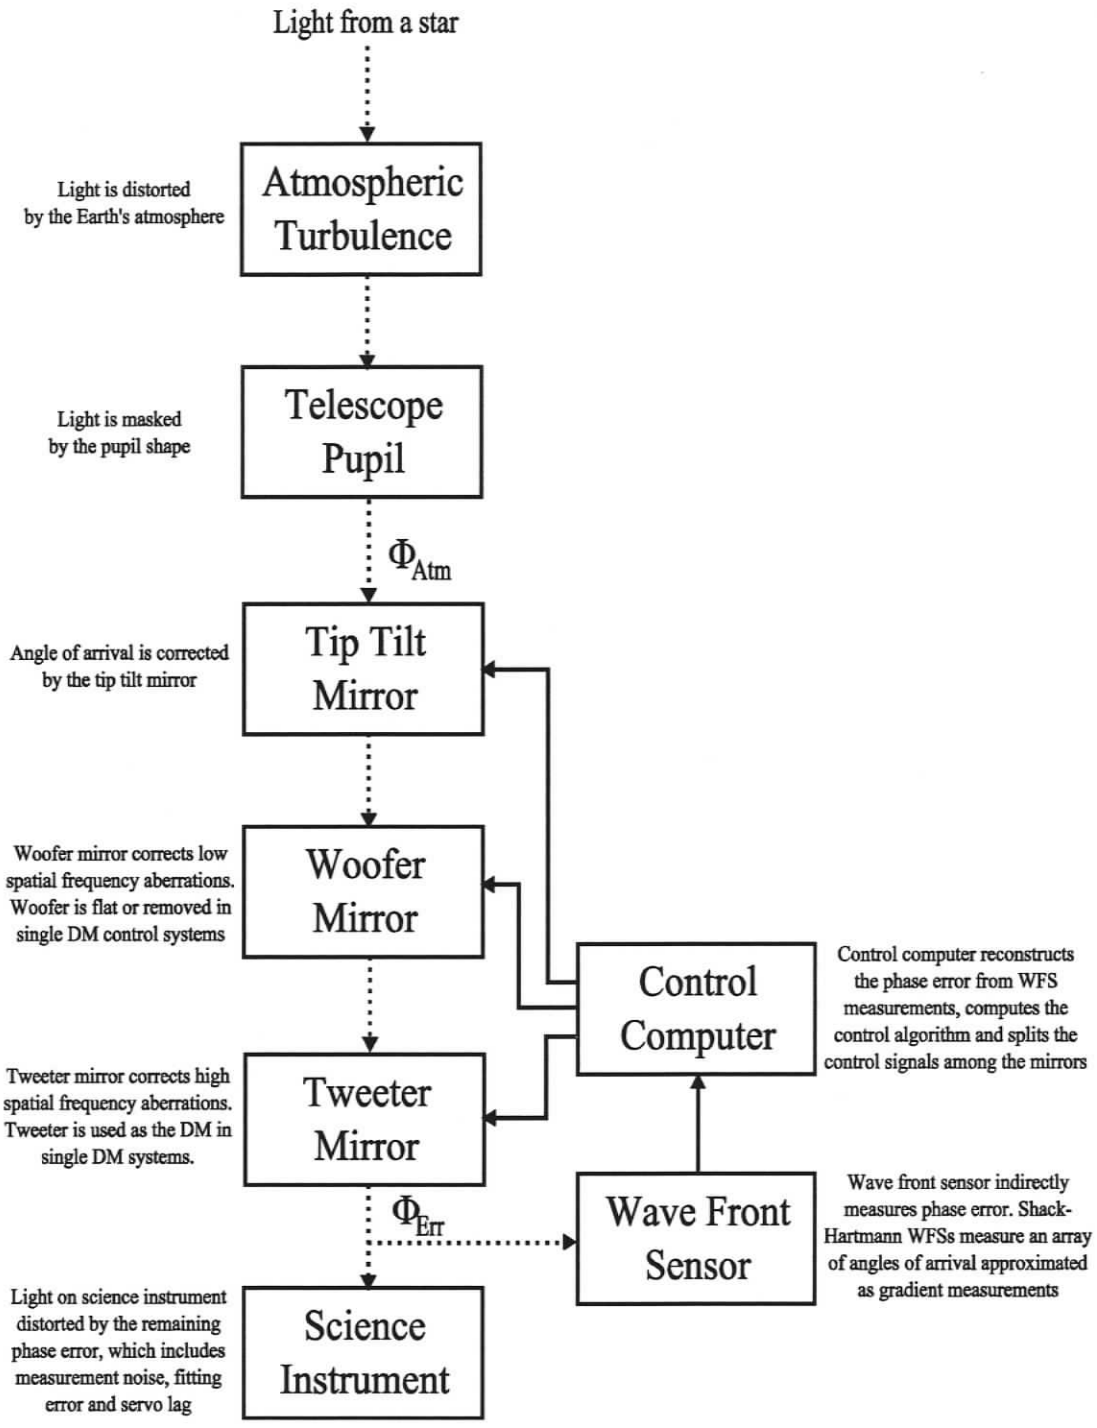

Figure 5.1: Block diagram of an adaptive optics system, including optional woofer mirror. The solid lines are computer signals and the dotted lines are beams of light. Variable Φ is wave front phase and subscript 'Atm' refers to the atmospheric turbulence and 'Err' is the residual error.

5.2.1 Deformable mirror

The DM is composed of a grid of actuators under a bendable membrane. A voltage applied to the actuator pushes the actuator normal to the plane of the actuator grid. This force bends the reflective surface. The goal of AO systems is to bend DMs in a manner that will correct for the phase distortions in the wave front.

Multiple DMs may be used in an AO system and one such configuration is named a Woofer-Tweeter system. This system uses two mirrors acting as one, with both high-stroke and high actuator count, to provide correction of the distortions caused by the Earth's atmospheric turbulence. This system is discussed in further detail in Section 5.4.

5.2.2 Shack-Hartmann wave front sensor

In this work, a simulated Shack-Hartmann WFS (SH-WFS) (Widiker, Harris and Duncan 2006) is used to measure the gradient of a distorted wave front. The data measured by a SH-WFS is approximated as discretely sampled gradient measurements. The measurements can be approximated by the Fried model (Fried 1977) discussed in Section 2.2.1. However, this chapter will utilize a SH-WFS model based on Fourier optics described in Section 5.6.1.3.

These measurements are masked with a circular pupil. Extrapolation methods such as those in Chapter 4 are necessary to fill the missing data prior to the application of wavelet integration.

5.2.3 System dynamics

An overview of an AO system is shown as Figure 5.1. This section introduces the simple dynamic model used for DMs and SH-WFSs.

An important concept in modelling discretely sampled systems is the model of a sample delay. In the control systems in this chapter, the delay is applied in the temporal domain. A single delay in the continuous temporal domain is $\exp(-j\omega\Delta t)$, where Δt is the sample time and ω is the frequency in radians per second. A single delay in the discrete domain is equivalently modelled as z^{-1}.

A method to measure the dynamic response of AO systems is given in Appendix A, which was published as (Hampton et al. 2008). In Appendix A, it was measured that the response of the DM and the SH-WFS can each be modelled as a single delay. This model is supported by (Poyneer, Macintosh and Véran 2007). In this chapter, all simulated DMs will use this model.

This system model was employed in (Hampton et al. 2004) and (Hampton 2005) in order to increase the control system bandwidth. However, this chapter will only consider the classic AO control system defined in Section 5.3.2.

5.3 Single deformable mirror systems

This section describes in detail the classic approach to controlling an AO system as well as a new approach that incorporates wavelet integration. A classical AO system uses only one DM.

5.3.1 Interaction matrix

The interaction matrix, D, is a direct measurement of how a DM and SH-WFS interact. There are approximately two measurements from the SH-WFS for every DM actuator because the SH-WFS measures the gradient of the wave front in two directions. This leads to a sparse rectangular matrix shown in Figure 5.2. The pseudo-inverse (Watkins

2002) of D is called the reconstruction matrix, M, or the reconstructor. It is not sparse and converts error, measured by the SH-WFS, into actuator error signals used by the control system to remove wave front error.

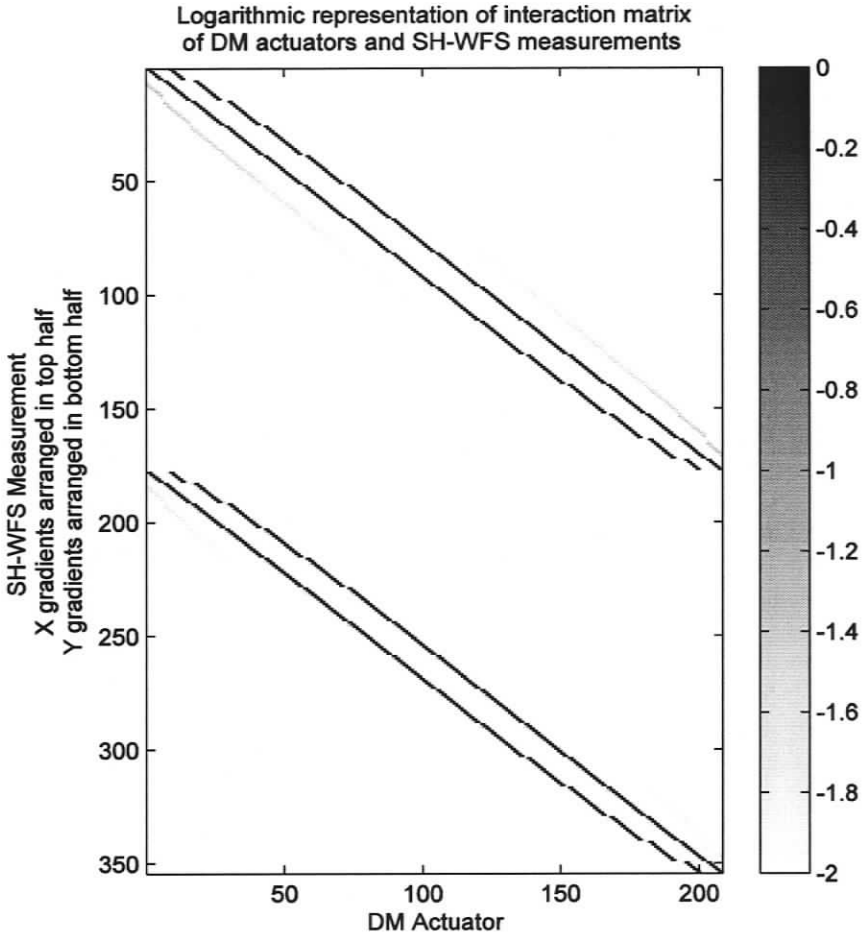

Figure 5.2: Interaction matrix between 16 x 16 DM and 15 x 15 SH-WFS

5.3.2 Classical adaptive optics control system

In many on-sky AO systems, the control system consists primarily of a single integrator with a gain that has been tuned to an optimal value. It was reported in (Poyneer, Macintosh and Véran 2007) that the general form of this control system is given by

$$C(z) = \frac{g}{1 - cz^{-1}} \tag{5.1}$$

The value of c is 1 for an integrator and less than 1 for a "leaky" integrator. Invisible modes are shapes that the sensor cannot detect. Using a value of less than one can be useful to prevent invisible modes from accumulating (Poyneer, Macintosh and Véran 2007). In the following work, c will be equal to 1. It was shown in (Hampton 2005) that when c is 1, the system is stable for $0 < g < 1$ and it is critically damped when $g = 0.25$. It takes several samples for such configurations to settle on a steady state change, which will be taken advantage of in later sections.

The state space control system used for the interaction matrix based system is derived from (5.1) with $c = 1$. This is the form of the controller regardless of the number of delays in the system.

$$\vec{a}[k] = \vec{a}[k-1] + gM_{a,d}\vec{d}[k] \tag{5.2}$$

where $M_{a,d}$ is the pseudo-inverse of the interaction between actuator commands, DM and the SH-WFS of Figure 5.2. Vector $\vec{d}$ is the noise-contaminated measurement of the phase slopes as defined later in (5.23) when the simulation apparatus is defined.

5.3.3 Closed loop system with two delays

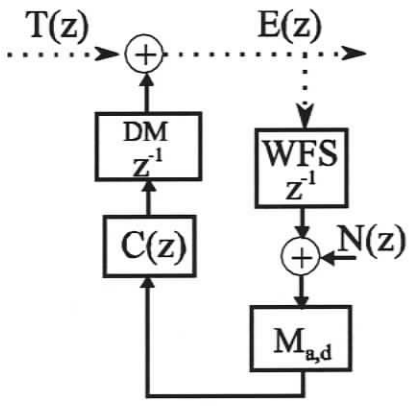

Figure 5.3: Two delay single mirror AO system

Assuming a 1-D model of the control system, the system may be represented in the z-domain. The turbulence, Φ_{Atm}, is represented here by $T(z)$. The residual phase, Φ_{Err}, is presented by $E(z)$. The noise signal, $N(z)$, is additive to the WFS measurement. The SH-WFS and the DM are each modelled as single delays and are further explained in Section 5.6.1. The error for a given turbulence can be determined with (5.3).

$$\begin{aligned}\frac{E(z)}{T(z)} &= \frac{1}{1+C(z)z^{-2}} \\ &= \frac{\left(1-cz^{-1}\right)}{1-cz^{-1}+gz^{-2}}\end{aligned} \tag{5.3}$$

where $E(z)/T(z)$ is also known as the rejection ratio.

$$\begin{aligned}\frac{E(z)}{N(z)} &= -\frac{C(z)z^{-2}}{1+C(z)z^{-2}} \\ &= -\frac{gz^{-2}}{1-cz^{-1}+gz^{-2}}\end{aligned} \tag{5.4}$$

where $E(z)/N(z)$ is the SH-WFS noise filter. The residual error becomes

$$E(z) = \frac{\left(1-cz^{-1}\right)}{1-cz^{-1}+gz^{-2}}T(z) - \frac{gz^{-2}}{1-cz^{-1}+gz^{-2}}N(z) \tag{5.5}$$

5.3.4 Closed loop system with one delay

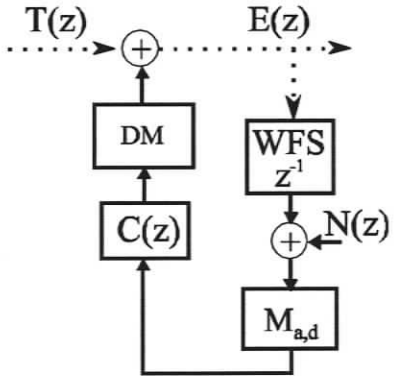

Figure 5.4: One delay single mirror AO system.

The DM model includes the response of the control computer. When the time to process the control algorithm and the time to reach steady state on the DM are insignificant compared to the time between samples of the SH-WFS, the system delay reduces by up to one sample. If the system delay is one, the closed loop equations become

$$\begin{aligned}\frac{E(z)}{T(z)} &= \frac{1}{1+C(z)z^{-1}} \\ &= \frac{\left(1-cz^{-1}\right)}{1-(c-g)z^{-1}}\end{aligned} \tag{5.6}$$

$$\begin{aligned}\frac{E(z)}{N(z)} &= -\frac{C(z)z^{-1}}{1+C(z)z^{-1}} \\ &= -\frac{gz^{-1}}{1-(c-g)z^{-1}}\end{aligned} \tag{5.7}$$

$$E(z) = \frac{\left(1-cz^{-1}\right)}{1-(c-g)z^{-1}}T(z) - \frac{gz^{-1}}{1-(c-g)z^{-1}}N(z) \tag{5.8}$$

The limit of gain when the delay is two samples is $g=1$ and it is $g=2$ when there is only one delay. Therefore, it is reasonable to double the gain when the delay is reduced to one sample since the same percentage of the maximum gain is being used. Further,

when $g = 0.25$ and $c = 1$ in (5.5), the two closed loop poles are located at $z = 0.5$. To place the single closed loop pole of (5.8) at $z = 0.5$ when $c = 1$, g must equal 0.5. So, when a two delay system is controlled, the gain will be 0.25 and it will be 0.5 when a one delay system is controlled for a fair comparison of performance.

5.3.5 Wavelet integration control system

The wavelet integration method has been shown to be useful as a mathematical tool when the source of the gradient data is the Fried or Hudgin modelled partial difference operators. Here it is shown as a reconstruction method for AO using a simulated SH-WFS.

A standard interaction matrix for AO (i.e., interaction between DM actuators and SH-WFS) is rectangular, as illustrated in Figure 5.2. It is important to note that an interaction matrix between actuators and phase, $D_{\varphi,a}$, is square with a strong main diagonal due to the one-to-one correspondence and alignment of the phase and actuators, as shown in Figure 5.5.

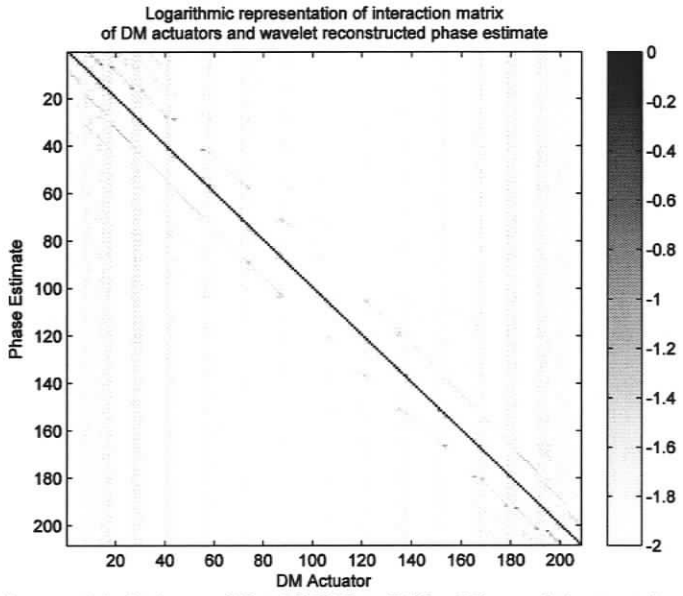

Figure 5.5: Interaction matrix between 16 x 16 DM and 16 x 16 wavelet reconstructed phase estimate

One method for converting the phase error into actuator error is to measure the interaction matrix between phase and actuators and then use the inverse of this matrix in (5.9), in a similar fashion to (5.2). An example interaction matrix between actuators is shown in Figure 5.5. This interaction includes the SH-WFS and wavelet integration.

$$\vec{a}[k] = \vec{a}[k-1] + gM_{a,\varphi}\tilde{\varphi}[k] \tag{5.9}$$

where $\vec{a}[k]$ is the vector of actuator commands at time sample k. The matrix $M_{a,\varphi}$ is a matrix that converts phase domain to actuator domain. Matrix $M_{a,\varphi}$ is the pseudo-inverse of the matrix shown in Figure 5.5. The variable g is the gain and $\tilde{\varphi}[k]$ is the estimate of the phase screen obtained via O(N) wavelet reconstruction.

Unfortunately, this approach removes many of the advantages of having an O(N) phase estimation method if this is followed by an O(N^2) vector-matrix multiply and an O(N^3) method to determine the inverse.

The computationally simplest approach is to use the phase error directly as actuator error and ignore details such as mechanical coupling. This approach requires that the phase estimate be scaled by the inverse of the peak phase influence of an actuator for a command of 1 unit. By this approach, the state equation can be given as (5.10).

$$\vec{a}[k] = \vec{a}[k-1] + \frac{g}{\max(D_{\varphi,a})}\tilde{\varphi}_{err}[k] \tag{5.10}$$

The measured estimate of the interaction between the actuator commands and the phase influence is a square matrix that has a width equal to the number of actuators that influence the area within the telescope pupil. In the 32×32 MEMS DM case, 812 actuators influence the pupil area. Therefore $D_{\varphi,a}$ is 812×812.

Another approach that does not require additional on-line computations is to divide the gain by the diagonal of $D_{\varphi,a}$, rather than the maximum value of $D_{\varphi,a}$. Then multiply this vector, element by element, with the phase error estimate. This allows fluctuations in manufacturing and influence due to proximity to the pupil edge to be taken into account. The approach of (5.11) is what will be used for the 32×32 MEMS DM model. The layout of this control system is given as Figure 5.6.

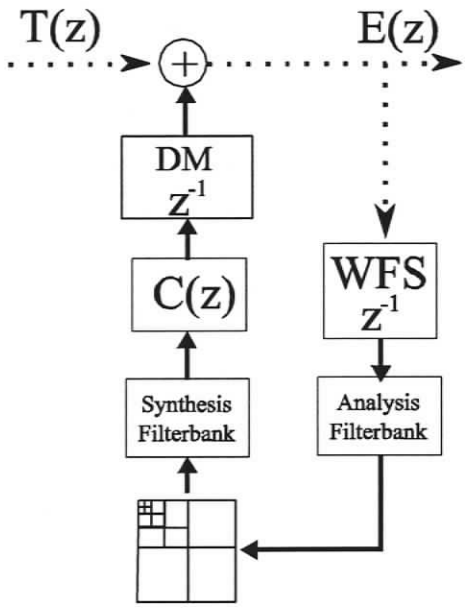

Figure 5.6: Single DM control system layout when using wavelet reconstruction.

$$\vec{a}[k]=\vec{a}[k-1]+\left(\frac{g}{diag\left(D_{\varphi,a}\right)}\right)\tilde{\varphi}_{err}[k] \tag{5.11}$$

The performance of wavelet based phase estimation will be tested directly against an inverse interaction matrix as a base line.

5.4 Woofer-Tweeter configuration

Extremely large telescopes (ELTs) require high actuator stroke and a large number of actuators, which is a contradictory situation for state-of-the-art DM technology. One

initiative to design and build a very large optical telescope is the Thirty Meter Telescope (TMT), named for the diameter of its primary mirror. This project is a collaborative effort between Canada and the United States and has merged several ELT design groups.

An important characteristic in an AO system is the relationship between the DM's stroke and the telescope's primary mirror diameter; the DM stroke requirements increase by the 5/6 power of the primary mirror diameter (F. Roddier 1999). Furthermore, the required number of DM actuators increases by the square of the primary's diameter. Therefore, an increase in primary mirror diameter from 8 m (e.g., Gemini North) to 30 m (e.g., TMT) increases the stroke requirement of each DM actuator by a factor of 3. The diameter increase also dictates that the number of actuators must increase by a factor of 14 to maintain equivalent correction.

The TMT design requires the NFIRAOS system to compensate for up to 11 microns of stroke, after tip-tilt correction (Dekany, et al. 2004), and require 5 thousand actuators (Ellerbroek, Boyer, et al. 2006) with plans of upgrading to 15,000. This stroke and actuator count both reach beyond the capabilities of current high actuator MEMS DM technology which are composed of 4096 actuators with 4 microns of stroke (Cornelissen, Bierden and Bifano 2006).

It is a fortunate property of atmospheric turbulence that the low spatial frequency aberrations are also the ones that statistically have the largest magnitude (Noll 1976). Therefore, pairing a high stroke-low spatial frequency DM to correct for the high magnitude aberrations with a low stroke-high spatial frequency DM to correct the higher frequency aberrations has the potential to meet the DM requirements of TMT and other ELTs.

This alternative to a single mirror is to use a high stroke, low density mirror (Woofer) in cooperation with a low-stroke, high-density mirror (Tweeter). The goal is to make this system of mirrors respond as though it was a single high-stroke, high-density and highly responsive mirror.

The potential drawbacks of the system are from two of the system characteristics: (i) there is only one SH-WFS to monitor the sum of corrective effects of these DMs, and (ii) the tweeter is capable of reproducing most, if not all, woofer shapes. Therefore, it was simple to envision a situation where the tweeter wastes its stroke by opposing a woofer mode and the SH-WFS would be blind to this effect as the net aberration is zero.

A method based on ideal Zernike polynomials (Noll 1976) has been presented in the literature (Hu et al. 2006) but ignores the dynamics of the mirrors and has not been shown on an experimental system. An earlier work that showed that two deformable mirrors could operate in closed loop has been presented (Barchers 2002) but the method requires a sensor after each mirror, which reduces the number of photons that reach the imaging camera. Systems for telescopes must operate with only one sensor placed after the mirrors to maximize photon flux on the science camera.

The report to TMT on the proof of concept of the W-T system is included in Appendix B; it was published as (Hampton et al. 2006). This solution separates the signals for the mirrors via the Singular Value Decomposition (SVD) to ensure that the mirrors correct an orthogonal set of modes. A method of separating these signals with the SVD is also published as (Conan et al. 2007). It was shown in (Keskin et al. 2008) that the control signals of a closed loop W-T system might be used to estimate the point-spread function (PSF) for use in the post-processing of astronomical images.

For control approaches that use interaction matrices, the alignment between mirrors is not critical. A particular alignment of these two mirrors is given as Figure 5.7, which is beneficial to wavelet integration.

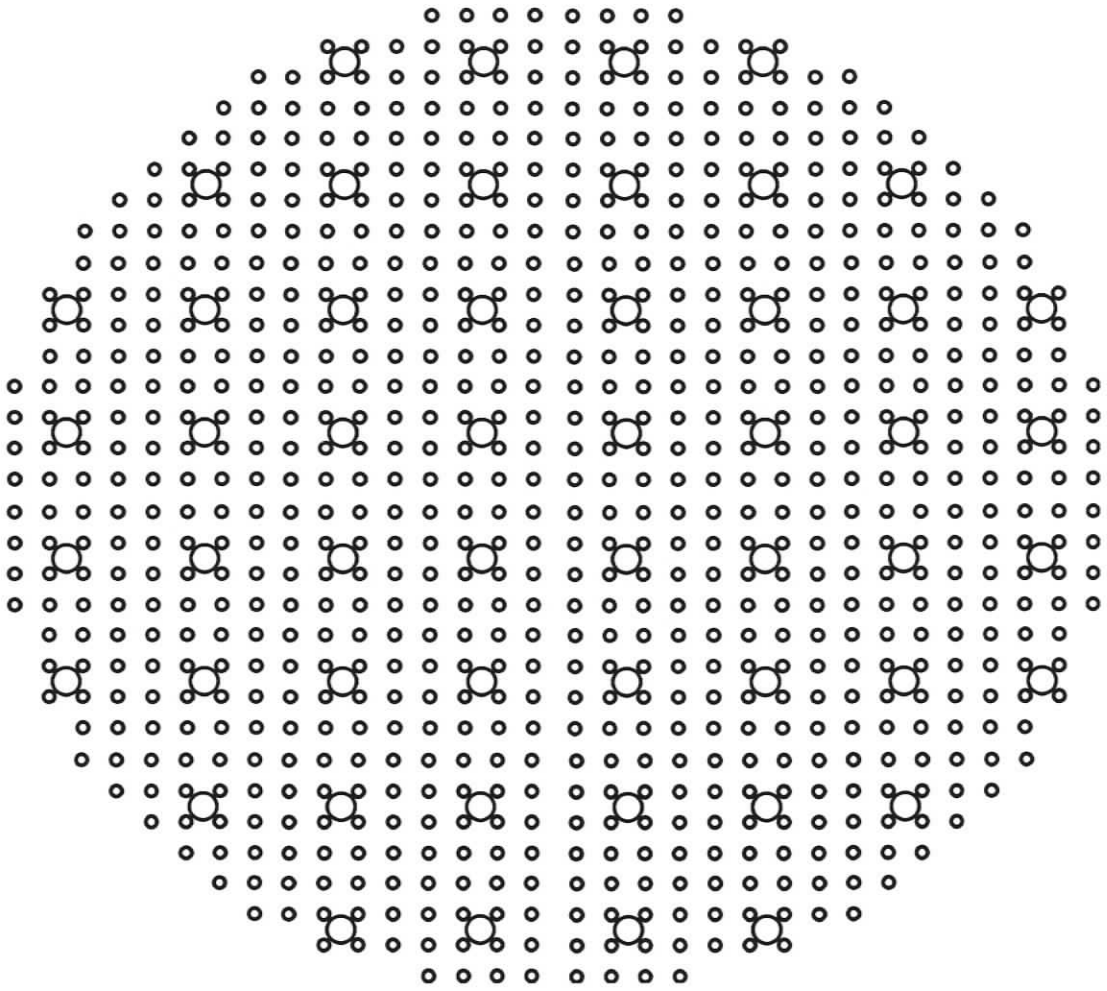

Figure 5.7: Alignment of woofer actuators (large circles) to tweeter actuators (small circles).

5.4.1 Signal separation with SVD

In order to control the W-T system, the error signal must be separated into tweeter error and woofer error in some manner. The method in Appendix B relies on matrix-vector multiplies and the Singular Value Decomposition (SVD). It is common to use the SVD to develop the pseudo-inverse of interaction matrices. If $D_{WFS,WR}$ and $D_{WFS,TR}$ are the interaction matrices of the woofer and the tweeter respectively with the WFS, then the pseudo-inverses are $M_{WR,WFS} = D^{\dagger}_{WFS,WR}$ and $M_{TR,WFS} = D^{\dagger}_{WFS,TR}$. The interaction matrix between the woofer and the tweeter can be determined from these matrices as

$$\begin{aligned} D_{TR,WR} &= M_{TR,WFS} D_{WFS,WR} \\ &= U_{TR} \Sigma_{TR,WR} V^{T}_{WR} \end{aligned} \tag{5.12}$$

This matrix projects woofer actuators onto the tweeter. Performing the SVD on this matrix creates a set of orthogonal modes in each of the actuator spaces to create identical modes in WFS space. U_{TR} is a matrix of orthogonal tweeter modes, V_{WR} is a matrix of woofer modes and matrix $\Sigma_{TR,WR}$ has non-zero values on its diagonal that represent the visibility of each of the mode pairs. These modes can then be used to project tweeter modes onto the woofer and vice versa as matrix operations.

The size of the largest matrix created in this way is $N_{wm} \times N_{ta}$ where N_{wm} is the number of woofer modes and N_{ta} is the number of tweeter actuators. This computational cost was not considered high in comparison to the $O(N^2)$ pseudo-inverse at the time of writing Appendix B but is now at least an order of magnitude slower than the wavelet integration technique.

5.4.2 Signal separation with wavelets

In this section, a computationally efficient W-T approach is developed that takes advantage of the structure of the wavelet decomposition associated with the proposed integration method. The performance is compared with the SVD-based method. For these comparisons, an 8×8 woofer and 32×32 tweeter are used. There are 52 woofer actuators and 812 tweeter actuators that significantly affect the area within the pupil. The alignment of these actuators is shown by Figure 5.7.

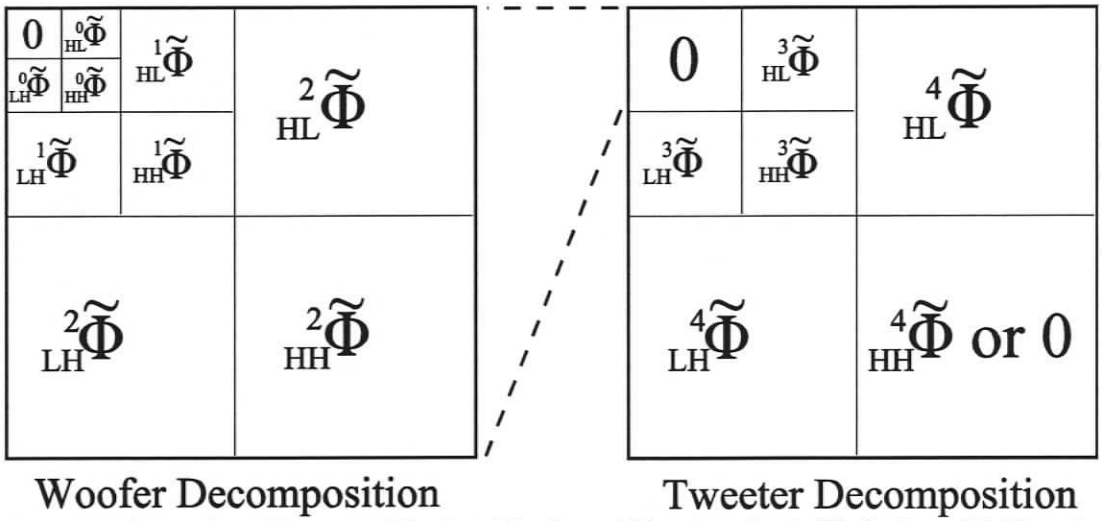

Figure 5.8: Separation of decomposition into Woofer and Tweeter subsets. The bottom right Tweeter quadrant is optionally 0 to suppress waffle modes.

This computationally efficient approach has a specific requirement that each woofer actuator must be centered in an independent square of $2^k \times 2^k$ tweeter actuators where k is integer and greater than 1. This alignment takes advantage of the multi-grid nature of the proposed integration process.

If the woofer is 8×8 and aligned as suggested above, an 8×8 phase estimate is available as an intermediate step in solving the 32×32 phase estimate for the tweeter. It is then a simple matter to divert the 8×8 phase screen to the woofer. For computational

simplicity, the region of the decomposition associated with the woofer is simply set to 0, as shown in Figure 5.8, for further reconstruction to the tweeter resolution.

5.5 Woofer-Tweeter control system

The control system used for the W-T has the same transfer function as the single DM case discussed in Section 5.3.5. The tweeter error signal is in the phase domain and the control system is left to converge on the correct signal in actuator domain. The woofer error in the phase domain is converted to woofer actuator domain with a relatively small reconstruction matrix. This signal is fed to the woofer controller in a manner no different from the single DM case. This process is shown in the diagram of Figure 5.9. It will be compared to the SVD based method based on Figure 5 in Appendix B with $g_L = 1$.

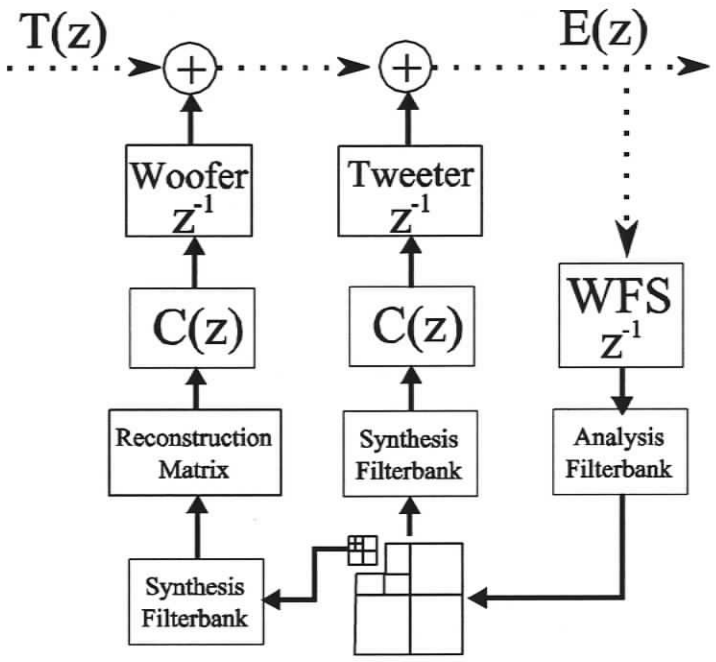

Figure 5.9: Layout of W-T control system using the Haar decomposition form for signal separation

It is assumed for this presentation that the woofer mirror is capable of responding at the same rate as the tweeter. This removes the considerations of filtering the woofer and

tweeter signals in the temporal domain that are discussed in Appendix B. The separation of the signals in the spatial domain was discussed in Section 5.4.2 and is represented in Figure 5.9. Further, the woofer handles the tilt correction as well to simplify the presentation since there are no stroke limitations in a simulation. The control system using the SVD of (5.12) is shown as Figure 5.10.

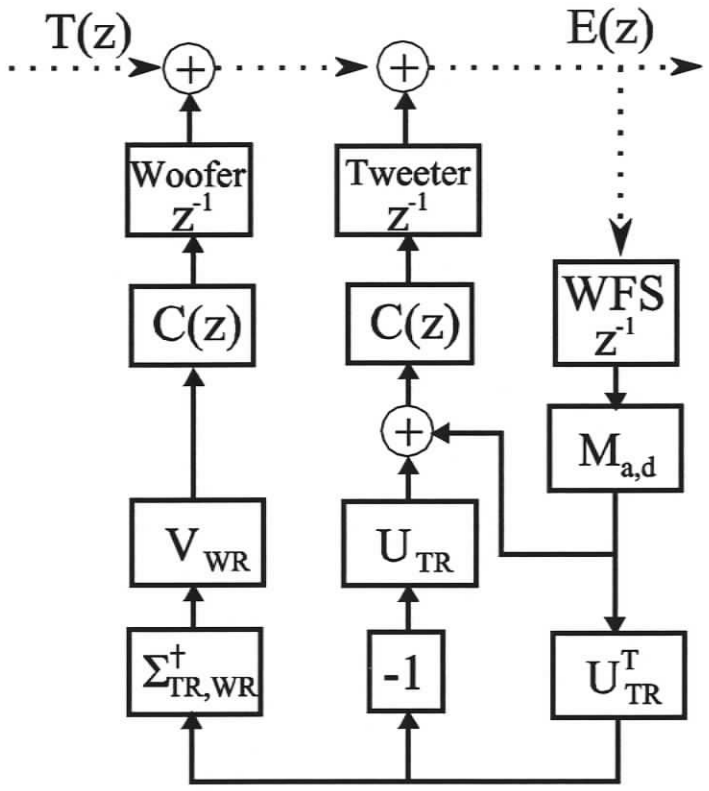

Figure 5.10: Simplified layout of W-T control system using the SVD for signal separation. U_{TR}^{T} is a projection to the W-T modes that both mirrors can form. U_{TR} is a projection back to Tweeter space to remove these modes from the Tweeter's error signal. $V_{WR}\Sigma^{\dagger}_{TR,WR}$ projects these modes into Woofer space.

5.5.1 Computational complexity

It was shown that the woofer signal is organized into the upper 8×8 region of the intermediate wavelet decomposition. The separation of the woofer signal from the tweeter signal is then a computationally free aspect of the multi-grid nature of the wavelet

reconstruction method due to this convenient organization of signals. Further, it is obvious that the woofer synthesis and reconstruction could occur in parallel with the tweeter synthesis since the region of the decomposition used for the woofer is replaced by zeros for the tweeter.

Since the woofer has much fewer actuators than the tweeter, it is not a large computational burden to use a measured woofer interaction matrix between its actuators and this lower resolution phase estimate. For the 52-actuator woofer, this results in a matrix of size 52×52 and would cost 5408 flops to compute. The SVD-based method uses up to three 52×812 sized matrices for a total cost of 253,344 flops. The computational savings are obvious. Further, if the number of tweeter actuators is N, then the SVD-based method is greater than $156N$ whereas the 52×52 reconstruction matrix is a flat computational fee of 5408 flops independent of N.

The computational cost of the algorithms is shown in Figure 5.11. The $O(N^2)$ pseudo-inverse reconstructor matrix is used for the SVD-based method since this was one of its assumptions. This shows that the new wavelet-based woofer tweeter control system can operate many times faster than the matrix-based control algorithm in Appendix B. When comparing the control algorithms independently of the reconstruction methods, splitting the signal in the wavelet decomposition space is nearly 40 times faster than splitting the signals in actuator space using matrix multiplies. The lower limit of the computation cost of this newer approach is approximately 14 flops per tweeter actuator, including a full reconstruction.

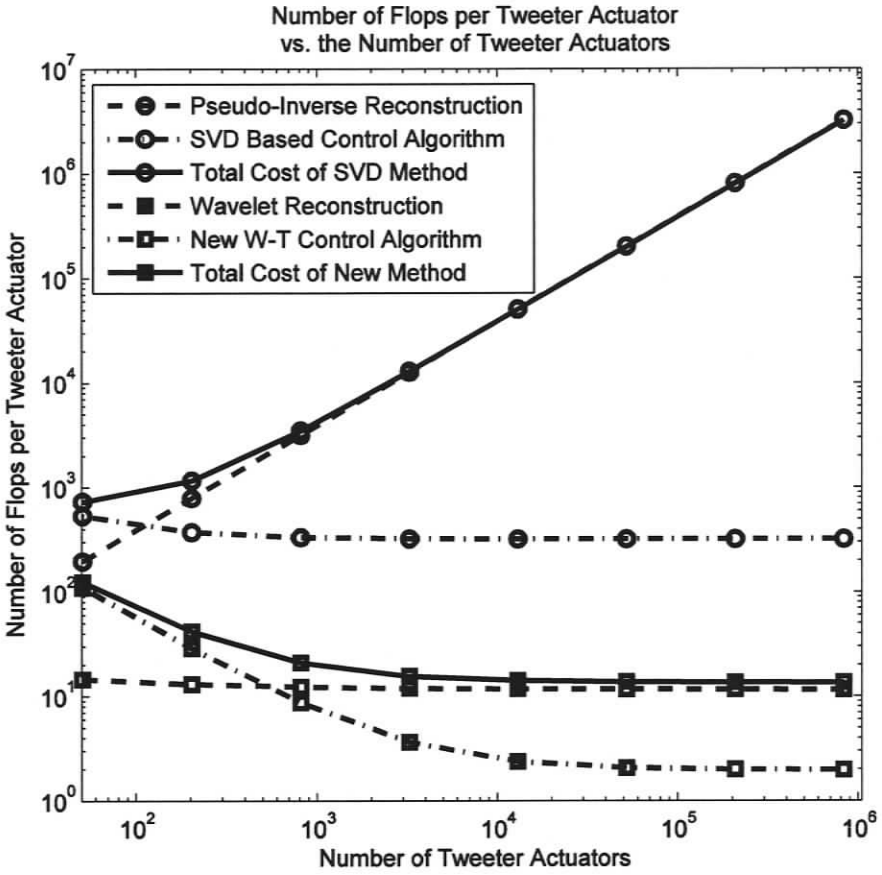

Figure 5.11: Computational cost per tweeter actuator.

5.5.2 Theoretical stroke requirements

The stroke and actuator requirements for a telescope can be derived from the works of (Noll 1976) and later (F. Roddier 1999). The mean square phase of the atmosphere is given in (5.13). The residual mean square phase after perfect tilt correction is given by (5.14).

$$\sigma_1^2 = 1.03\left(D/r_0\right)^{5/3}\ rad^2 \tag{5.13}$$

$$\sigma_3^2 = 0.134\left(D/r_0\right)^{5/3}\ rad^2 \tag{5.14}$$

It was presented by (D. Gavel 2008) that the DM fitting error for continuous membrane mirrors can be given as

$$\sigma_{DM}^2 = 0.221\left(D/r_0\sqrt{N_{Sq}}\right)^{5/3} rad^2 \tag{5.15}$$

where N_{Sq} is the number of actuators of the square DM. Alternatively, $\sqrt{N_{Sq}}$ is the DM diameter measured in number of actuators. The form of (5.15) is also shown in (F. Roddier 1999) with a statement that the coefficient can be in the range of 0.2 to 0.4 depending on the characteristics of the mirror. The rms stroke required for each device can then be determined as being at most the residual error of the previous mirror. So to correct ±3 standard deviations of atmospheric phase, the stroke of the woofer is approximately $6\sigma_3$ and the tweeter stroke is $6\sigma_{woofer}$. For an 8×8 woofer, the mean square error after woofer correction is ideally $\sigma_{woofer}^2 = 0.0069\left(D/r_0\right)^{5/3} rad^2$.

5.6 Simulations

Simulations of deformable mirrors (DMs) and a Shack-Hartmann Wave Front Sensor (SH-WFS) are used in the following experiments to determine closed loop stability and accuracy. The simulation tools were developed by Dr. Rodolphe Conan based on the presentation in (Nakajima 1988) for the SH-WFS and uses Gaussian influence functions for the DM in a manner similar to (Vogel and Yang 2006).

5.6.1 Simulation Apparatus

This simulation is conducted in MatLab and is composed of three distinct models: (i) the turbulent phase screen, (ii) the DM, and (iii) the SH-WFS.

5.6.1.1 Phase screen model

The phase screen is generated as in Chapters 3 and 4. Again, r_0 is 0.25 m, the outer scale, L_0, is 22 m and the wave length is 500 nm. The wind speed is 25 m/s when applicable.

The telescope dimensions are chosen to be smaller than in other chapters due to the computational complexity of simulating a 30 m telescope with 128×128 actuators with 64 pixels per actuator. The highest resolution mirror tested is 32×32. If there is to be one actuator per r_0, then the diameter of the simulated telescope should be 8 meters (i.e., 0.25 m per actuator multiplied by 32 actuators). In this chapter, the diameter of the telescope is taken to be 8 m to be comparable to papers written on the 8 m class telescopes, such as (Rigaut, Ellerbroek and Northcott 1997) for the Gemini telescopes.

In Fried alignment, the 32×32 MEMS DM corresponds to a 31 x 31 lenslet SH-WFS. There are 8×8 phase pixels per lenslet, which leads to the requirement of a 248×248 pixel phase screen. The generated phase screens are denoted as Φ_{Atm}.

5.6.1.2 Deformable mirror model

The DM is modelled as a Gaussian influence function where the mechanical coupling between the actuators is defined as a percentage influence of one actuator at a displacement of one actuator pitch from the center of the given actuator. A MEMS DM (Cornelissen, Bierden and Bifano 2006) was defined for the simulation to have 32×32 actuators with 8 phase pixel pitch, and mechanical coupling with its neighbour actuators of 13%. This mirror is the tweeter and is used for single mirror tests. The simulated woofer DM is defined as an 8×8 mirror with 32 phase pixel pitch and mechanical coupling of 24%. The model is linear and superposition is considered valid. This is contrary to the measurements in (Wallace et al. 2006) for MEMS DMs, but provides an

adequate simulation for closed loop systems. The combined influence functions provide the corrective phase, Φ_{Cor}. This variable may be the combination of several mirrors in more complicated systems, such as the W-T.

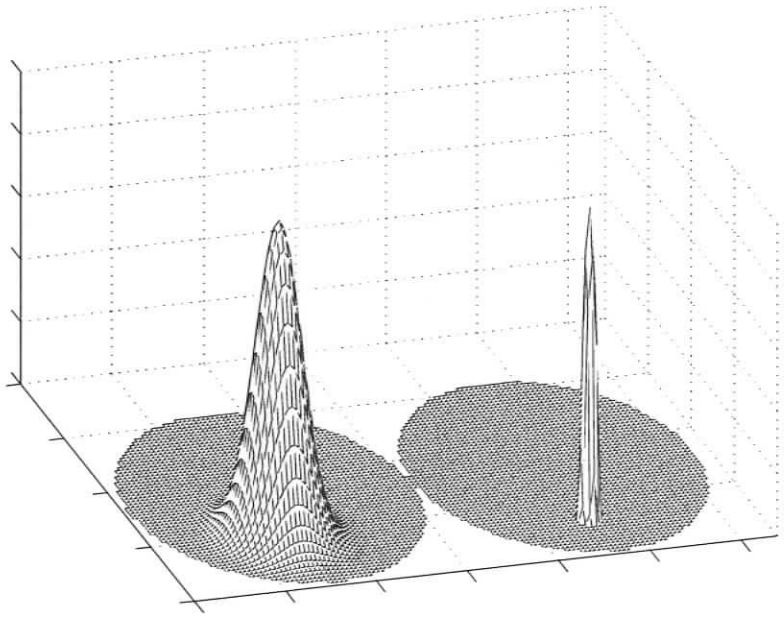

Figure 5.12: Actuation of a single actuator of the Woofer (left) and Tweeter (right). Diagram shows the difference in influence width of each mirror. Diagram is not to scale.

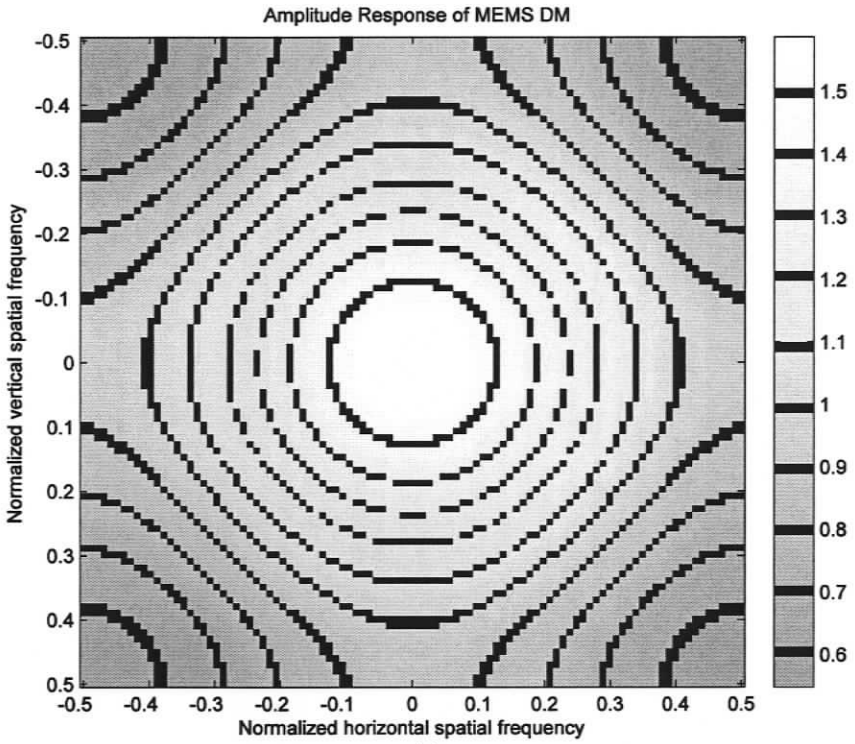

(a)

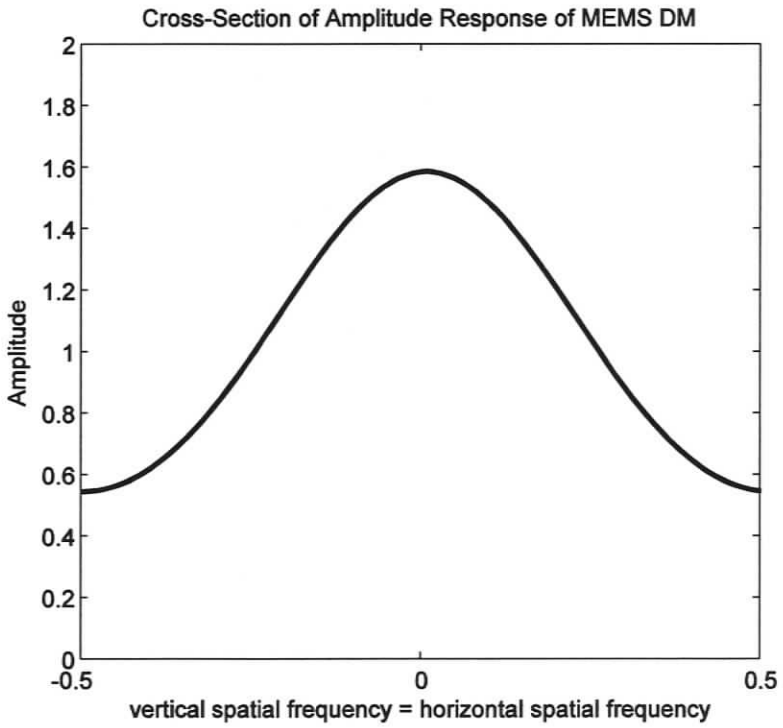

(b)

Figure 5.13: Approximate amplitude response of the MEMS DM model. (a) is a 2-D representation and (b) is the diagonal cross-section. This shows that low frequency actuator modes are amplified and high frequency modes are naturally attenuated.

The frequency response of the MEMS DM can be approximated as a 2-D filter composed of separable horizontal and vertical filters. The single dimensional filter is approximated as (5.16) in the z-domain and the magnitude of this with respect to spatial frequency, v, is (5.17).

$$F_{MEMS}(z) = 0.13z^{1} + 1 + 0.13z^{-1} \tag{5.16}$$

$$\left|F_{MEMS}(v)\right| = 1 + 0.26\cos 2\pi v \tag{5.17}$$

When a horizontal vector and a vertical vector that each contain (5.17) are convolved, the result is that of Figure 5.13. This shows that when actuator commands are sent to the DM, the low spatial frequencies are amplified and the high spatial frequencies are attenuated. Reconstruction matrices include an inverse of this spatial filtration property and cancel this potentially beneficial aspect.

5.6.1.3 Wave front sensor model

The SH-WFS model represents a 31×31 array of lenslets that focus the incoming phase screen onto a CCD camera. The approach of the SH-WFS model is to take the 248×248 pixel phase screen and separate it into 31×31 blocks of 8×8 pixels. Then the Fourier transform is computed on each block to model the phenomena of Fourier optics associated with each lenslet. This creates a focused spot location within the block. Noise can be added to this to simulate the electron noise of the camera CCD. A spot position is calculated using a center of gravity algorithm. The displacement of the spot from the center of the block becomes the measurement of the average angle of incidence. These slope measurements are taken as the gradient field for reconstruction algorithms. The equations that represent this description are as follows.

$$\Phi_{Err} = \Phi_{Atm} + \Phi_{Cor} \tag{5.18}$$

where Φ_{Cor} is the corrective phase provided by the DMs and includes the times 2 multiplier associated with a mirror reflection, Φ_{Atm} is the phase distortion caused by the Earth's atmosphere and Φ_{Err} is the residual error after the light of the guide star reflects off all the corrective elements.

$$W = P_{tel} \exp\left(j\Phi_{Err}\right) \tag{5.19}$$

where P_{tel} is the 8 m diameter telescope pupil and W is the resulting wave. The wave, W, is divided into 31×31 blocks and each has a 2-D index of p,q. There are 749 out of 961 blocks that are within telescope pupil and only these are computed.

$$\begin{bmatrix} d_{x,p,q} \\ d_{y,p,q} \end{bmatrix} = centroid\left(\left|F\left\{W_{p,q}\right\}\right|^2 + \eta_{p,q}\right) \tag{5.20}$$

where $\vec{d}$ is the vector of measured of the slopes. $\left|F\{\ \}\right|^2$ is the squared magnitude of the Fourier transform. Variable $\eta_{p,q}$ represents the electron noise in the SH-WFS and are each an 8×8 block of zero mean Gaussian white noise with standard deviation of σ_e. As shown in (F. Roddier 1999), the centroids are calculated as in (5.21) and (5.22).

$$d_x = \frac{\sum_{i,j} x_{i,j} I_{i,j}}{\sum_{i,j} I_{i,j}} \tag{5.21}$$

$$d_y = \frac{\sum_{i,j} y_{i,j} I_{i,j}}{\sum_{i,j} I_{i,j}} \tag{5.22}$$

where x and y is the horizontal or vertical displacement from the center of the lenslet for pixel (i,j). $I_{i,j}$ is the intensity of pixel (i,j). This measurement of the slopes, $\vec{d}$, is considered proportional to the gradients as shown in (5.23).

$$\begin{bmatrix} \dot{\varphi}_{x,p,q} \\ \dot{\varphi}_{y,p,q} \end{bmatrix} \approx \alpha \begin{bmatrix} d_{x,p,q} \\ d_{y,p,q} \end{bmatrix} \tag{5.23}$$

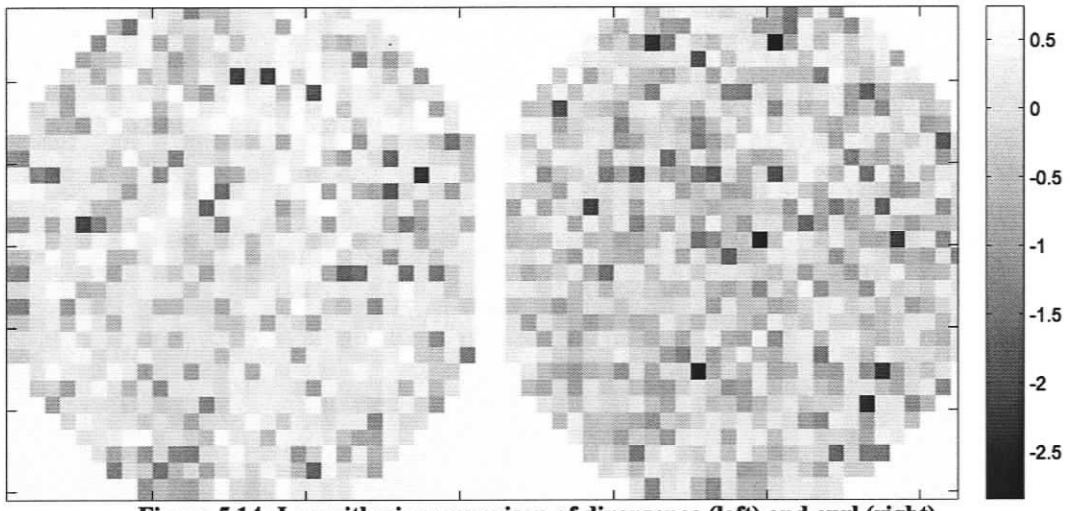

Figure 5.14: Logarithmic comparison of divergence (left) and curl (right). The rms of curl is 42.5% of the divergence.

This simulated sensor measures the gradient at discrete points in the Fried geometry. However, there is no overlap of phase among the lenslets as was modelled by Fried because each block is independent. The result is similar to a down-sampled Fried model due to this lack of overlap. This causes the measured gradient field to not have a curl of zero for all paths since there are effectively missing gradient measurements between the lenslets. This is shown in Figure 5.14 and it was calculated from this data that the rms of the curl is 42.5% of the rms of the divergence. This indicates significant error compared to a true gradient field with a curl of zero. Therefore, it is important to show the effect of such modelling errors. This will be studied in 5.6.2 and 5.6.3.

5.6.2 Single DM simulation

5.6.2.1Rejection ratio

The first performance metric studied is the power spectrum of the ratio of the residual error to the input phase strength, also known as the rejection ratio given in (5.3). A zero mean Gaussian white noise random signal was generated for 1050 time samples for each of 20 Zernike modes (Noll 1976). These were used to generate 1050 random phase screens that contain only the 20 lowest order polynomial shapes, except piston.

Let Θ_{in} be the white-noise input signal and Θ_{out} be the corresponding output signal. Let $D_{\varphi,z}$ be an interaction matrix between the first through fifth order Zernike modes (Noll, 1976) for a total of 20 modes.

$$\varphi_{atm}[l] = D_{\varphi,z}\theta_{in}[l] \tag{5.24}$$

The closed looped control systems of (5.11) and (5.2) each operated on this data to generate 1050 residual phase screens.

$$\theta_{out}[l] = D^{\dagger}_{\varphi,z}\varphi_{err}[l] \tag{5.25}$$

The Fourier transform is computed for each signal. The last 1024 samples of each of these data sets were used in order to ignore the initial settling time of the system.

$$\Theta_{in}[\beta,\psi] = F\{\theta_{in}[\beta,l]\} \tag{5.26}$$

$$\Theta_{out}[\beta,\psi] = F\{\theta_{out}[\beta,l]\} \tag{5.27}$$

where β represents the particular Zernike mode, with piston as 0. In order to avoid divide by zero or very small numbers, the signals in (5.26) and (5.27) are each binned such that each point is the average power of the small number of surrounding points. The subscript

of 'x' in (5.28) is a placeholder for 'in' or 'out'. The power is also averaged among the 20 Zernike modes.

$$P_x[\psi] = \frac{1}{(20)(11)} \sum_{\beta=1}^{20} \sum_{\tau=-5}^{5} \left|\Theta_x[\beta, \psi + \tau]\right|^2 \tag{5.28}$$

The plot is the ratio of the output power divided by the input power, measured in decibels. These measurements are compared to the power of (5.3) as the ideal model.

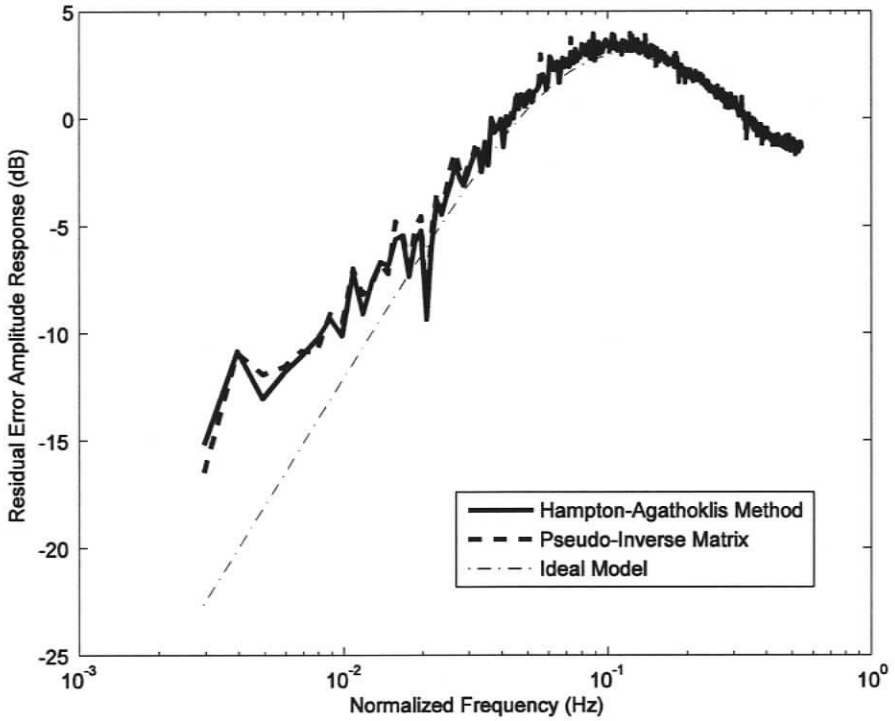

Figure 5.15: Average rejection ratios for first 20 Zernike modes when $g = 0.25$. Plot shows the Hampton-Agathoklis integration method is similar to the pseudo-inverse approach for these modes.

The critically damped case of Figure 5.15 and the underdamped case where $g = 0.6$ of Figure 5.16 is studied and it is determined that there is little difference between the reject ratios of using the O(N^2) reconstruction matrix and of using the O(N) wavelet-based phase estimate. Further, there was no appreciable difference between a full phase reconstruction or of one where the waffle modes were suppressed, so it was not displayed. This was expected since the first 20 Zernike modes do not contain waffle at

32×32 resolution. Divergence of the measurements from the ideal model of (5.3) is caused by at least two sources: (i) the averaging of the spikes of the 20 rejection ratio plots increases the overall result, and (ii) the amount of data collected which results in poor estimates of lowest frequency results.

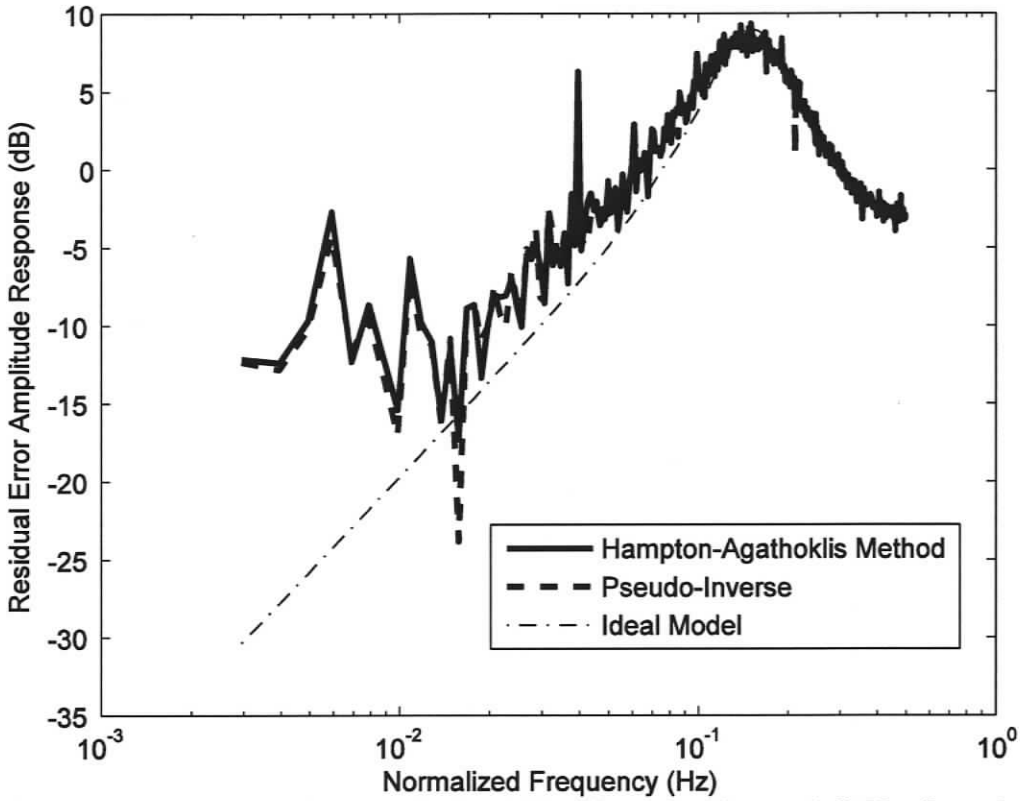

Figure 5.16: Average rejection ratios for first 20 Zernike modes when $g = 0.60$. Plot shows the Hampton-Agathoklis integration method is similar to the pseudo-inverse approach for these modes.

5.6.3 Woofer-Tweeter simulations

The W-T experiments will demonstrate the comparison between the wavelet reconstruction technique and the SVD-based technique to obtain W-T control. The statistical parameters used to generate the phase screen used for these experiments is defined in Section 5.6.1.1. The alignment between the woofer and the tweeter is that of Figure 5.7. The control system for the wavelet reconstruction method is shown in Figure 5.9. The value of the gain, g, is chosen to the critically damped case of 0.25. The control

system for the SVD-based method splits the signals between the mirrors as in Figure 5 in Appendix B. However, the control algorithm has been modified to match C(z). The sample rate of the control system is simulated to be 1kHz for the purposes of determining the motion of the phase screens, though the actual processing rate is approximately 2 Hz.

5.6.3.1 Correction of dynamic phase

The residual error associated with a non-zero wind velocity will be studied. In this experiment, the wind speed is chosen to be 25 m/s as in (Vogel and Yang 2006).

The SVD-based W-T controller is operated with a delay of 2 samples. The performance of this baseline method is compared to the wavelet reconstruction-based W-T controller for delays of one sample and two samples. This is to show (i) the benefit of reducing the delay by using a fast reconstruction method, and (ii) the comparison of quality when the delays are unchanged. In order for the closed loop poles to be placed at the same point in the z-domain (i.e., $z = 0.5$), the gain for one delay controller in (5.8) is 0.5 and the gain for two delay controller in (5.5) is 0.25.

When there is no noise for the SVD or wavelet reconstruction W-T methods, the means are 0.3802 and 0.3367 respectively for two delays, as seen in Table 1. However, one of these delays is caused by the delay of the computations. If the computations are considered instantaneous, the system delay reduces to one. When the system can respond this fast there is an increase in Strehl ratio to a mean of 0.4311. This is a 13.4% improvement over the two delay SVD approach and 28% increase over the two delay wavelet reconstruction. Even though there is a loss when the delays of the two approaches are considered equal, performance can be greatly increased when one can take advantage of the speed of the algorithm.

The Strehl ratio is the ratio of peak light intensity of a corrected point source image to the peak of the image of an ideal point source with no turbulence. The Strehl ratio was recorded for a wavelength of 500 nm.

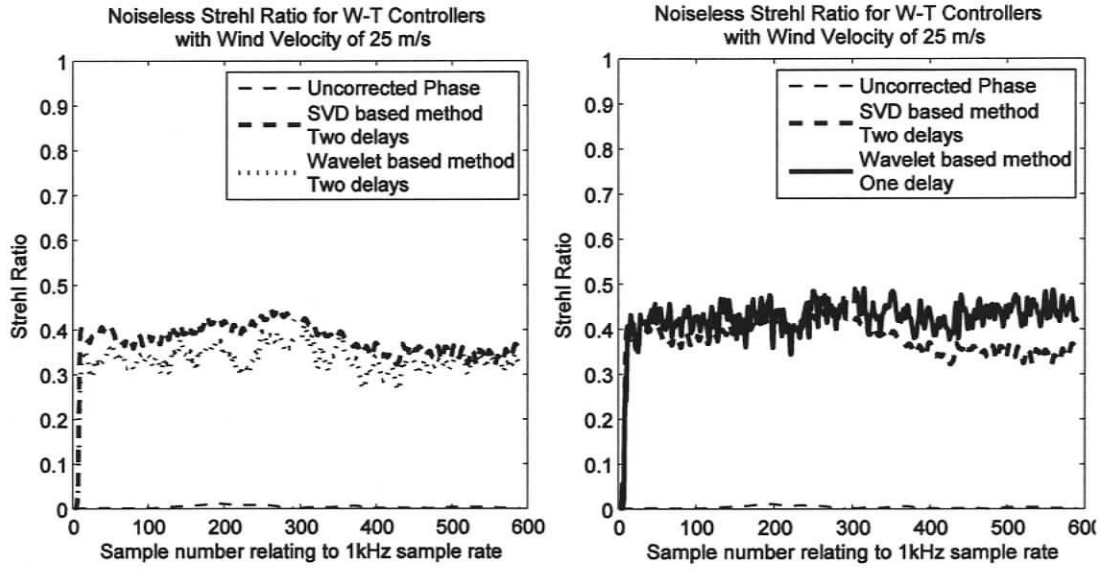

Figure 5.17: Noiseless Strehl ratio when turbulence velocity is 25 m/s. Sample rate is 1 kHz. The one delay case on the right shows the performance upper limit when computation and transmission delays are zero with closed loop pole z = 0.5.

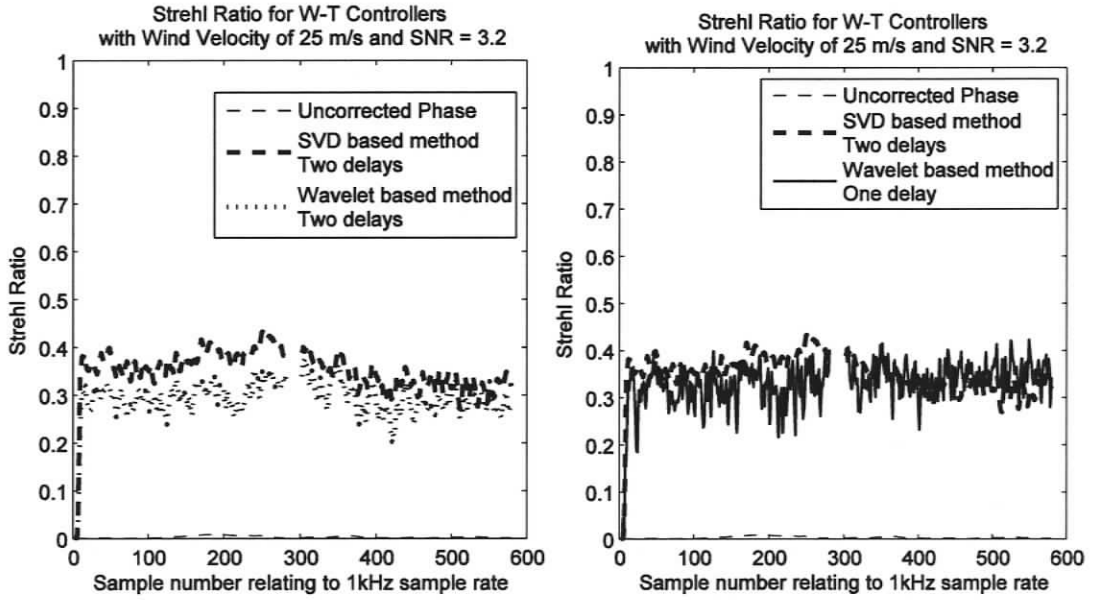

Figure 5.18: Strehl ratio when turbulence velocity is 25 m/s and SNR = 3.2. The right hand side shows that the faster system's performance is not an advantage at this noise level.

Table 1: Measured statistics of Strehl ratio for each tested W-T method. Wind Velocity is 25 m/s. Uncorrected mean = 0.003.

SNR	Control Method	Mean	Standard Deviation
∞	SVD Two Delays	0.3802	0.0297
	Wavelet Two Delays	0.3367	0.0257
	Wavelet One Delays	0.4311	0.0263
3.2	SVD Two Delays	0.3549	0.0336
	Wavelet Two Delays	0.2984	0.0280
	Wavelet One Delays	0.3357	0.0375

When an electron readout noise of $\sigma_e = 10$ per pixel per sample is applied to the SH-WFS, the performance of the two delay SVD-based W-T method is nearly the same as the one delay wavelet reconstruction-based W-T method. According to (Rigaut, Ellerbroek and Northcott 1997), the number of photons from a magnitude 0 source in the R band is 8.2×10^{11} photons per second. There are 749 lenslets within the circular pupil and 1000 samples per second. This becomes 1.1×10^{6} photons per lenslet per sample. The chosen photon signal was 256 photons per lenslet per sample. This means that the star can be dimmer by a factor of ~4300. A change of five magnitudes is a factor of 100 for astronomy, so each magnitude is a change of 2.512. Therefore, the results of Figure 5.18 for SNR = 3.2 would be similar to a magnitude 9 star with a SH-WFS read-out noise of $\sigma_e = 10$ per pixels per sample.

The following two experiments determine the sources of error caused by the static fitting error and of the noise.

5.6.3.2 Correction of static phase

This test of the system is to determine its performance when correcting a static phase disturbance. The W-T AO system is permitted to settle for 20 samples before sampling the stroke and error. The phase screen is shifted by 0.5 m 29 times to obtain 30 simulated measurements. This test is conducted free of noise to isolate the fitting error of the W-T system.

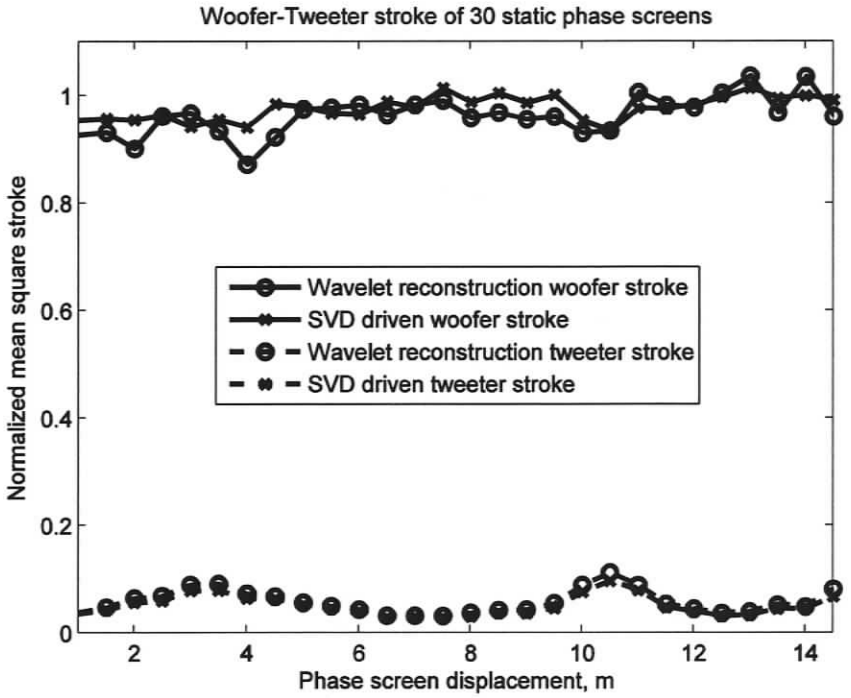

Figure 5.19: W-T stroke of 30 static phase screens

The normalized mean squared stroke was plotted as Figure 5.19. It was determined from this data that the mean square woofer stroke is 0.3% less for the wavelet reconstruction method in comparison to that of the SVD-based method. This reduction has the negative effect of increasing the mean square tweeter stroke by 3%. It was expected that the tweeter would require only 22.1% of the woofer stroke as determined by (5.15). It was determined via the histogram of Figure 5.20 that the core 99.4% of the

tweeter required 34% of the stroke of the core 99.4% of the woofer stroke. Since the SVD-based method is such that 100% of woofer mode stroke is corrected by the woofer (i.e., tweeter stroke is minimized), measuring an increase of the required tweeter stroke on the order of 50% suggests that using the residual error of the woofer as the stroke requirement of the tweeter may be optimistic in determining the tweeter stroke.

The normalized mean squared error is plotted as Figure 5.21. This shows that the error for both methods follows the same trend of residual error for identical phase screens. The normalized mean square residual error after W-T correction was expected to be .00603 from (5.15) divided by (5.13) with $N_{sq} = 1024$. This is close to the dash-dot line in Figure 5.21 measured at 0.00635.

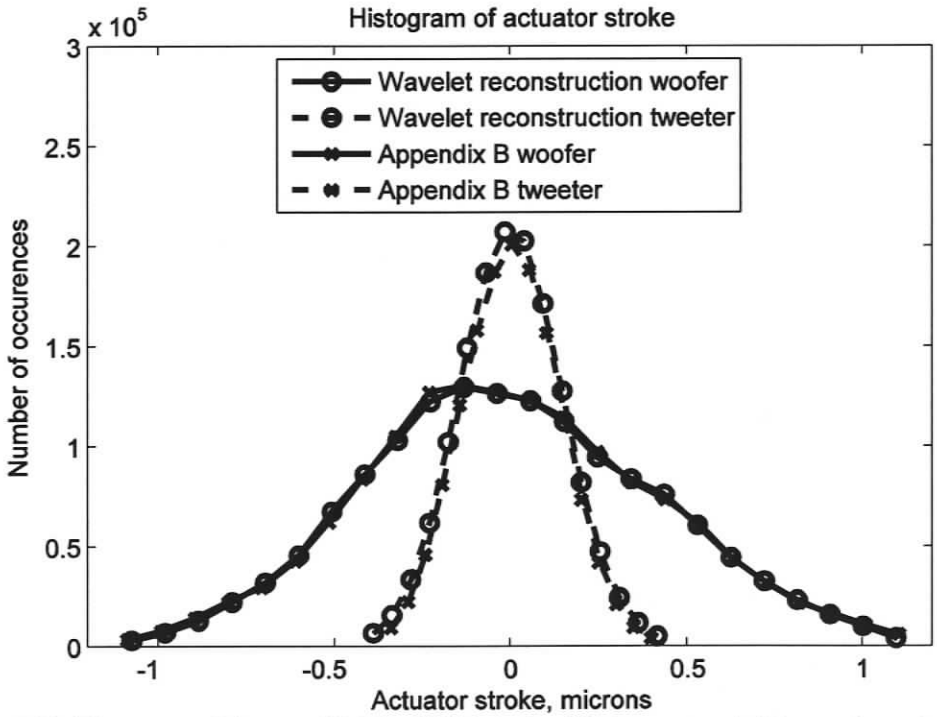

Figure 5.20: Histogram of the core 99.4% of W-T stroke. Tilt is uncorrected by woofer or tweeter. Plot shows that the two methods have nearly overlapping stroke requirements.

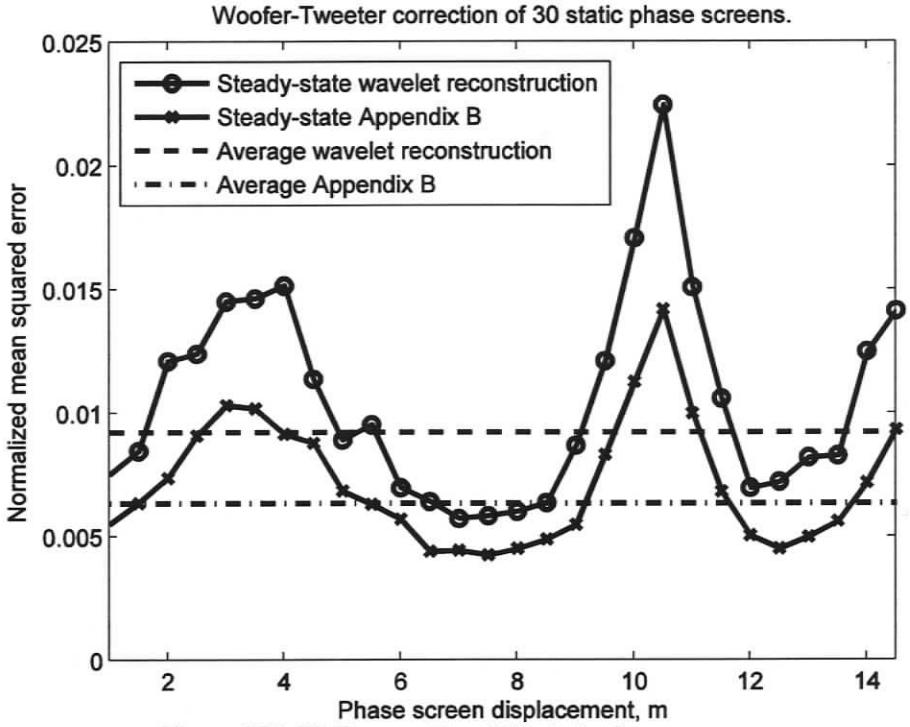

Figure 5.21: W-T correction of 30 static phase screens

It is calculated from Figure 5.21 that the wavelet reconstruction method results in an average of 45% higher error over that of using the SVD. However, in terms of percentage rms error, the SVD approach is 92.25% correct and the wavelet method is 90.5% correct.

In Figure 5.22, there are two pseudo-inverse methods used. The "SH-WFS Pseudo-Inverse Matrix" is based on WFS slopes as input. The "Phase Pseudo-Inverse Matrix" uses the phase estimate of the wavelet reconstruction method as the input. "Phase Gain Vector" uses the reciprocal of the main diagonal of the phase interaction matrix as shown in (5.11). This was tested to determine if there was any advantage to spending the computations to use a phase-based reconstruction matrix.

It was determined that there is no advantage for the static case when using this pseudo-inverse since the Strehl performance overlaps the computationally simple

approach. It is determined that the classic pseudo-inverse approach had a mean 21% increase in Strehl ratio over the wavelet approach.

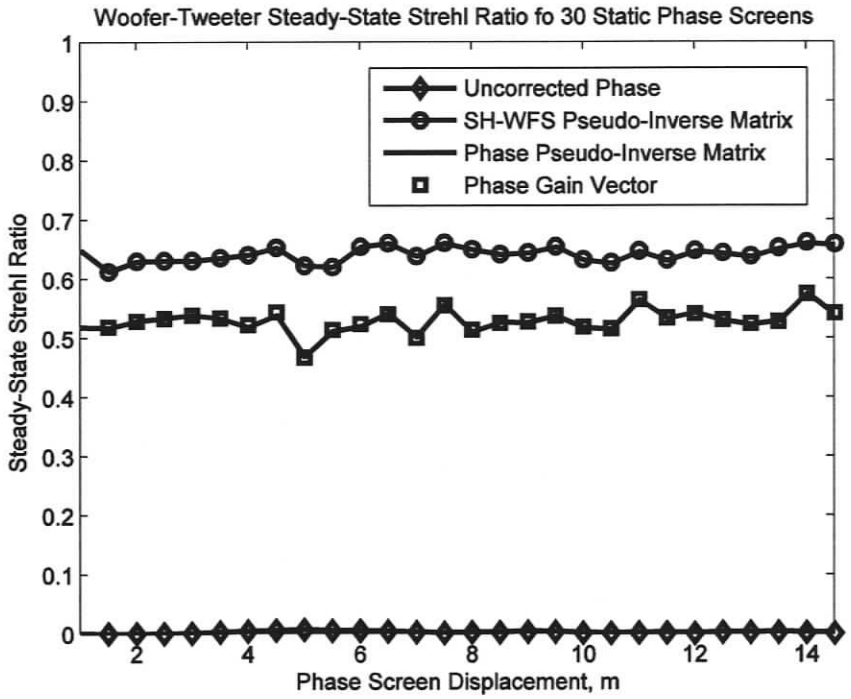

Figure 5.22: Strehl ratio steady-state correction of W-T systems for 30 static phase screens

In summary, the wavelet reconstruction method can be orders of magnitude faster than that of an SVD-based controller and results in similar stroke requirements. However, there is a trade-off seen in the increase in residual error and loss of Strehl ratio. It was presented in Section 5.6.3.1 that this problem can be overcome by using the speed of the wavelet method.

5.6.3.3 Correction of noise

To determine the effect of noise in the phase correction, this simulation sets the phase to zero and the Gaussian white-noise standard deviation of the electron noise is set to σ_e per pixel per sample. Similarly to (Rigaut, Ellerbroek and Northcott 1997), the SNR can be calculated as in (5.29).

$$SNR = \frac{N_{pe}}{\sqrt{N_{pe} + N_{pix}\sigma_e^2 + N_b}} \tag{5.29}$$

where N_{pe} is the number of photons per lenslet per sample and is also a noise term in the denominator due to the Poisson distribution nature of photons. N_{pix} is the number of pixels under a lenslet and is equal to 64. N_b is associated with the background noise and is zero for these experiments. The N_{pe} term in the denominator is insignificant compared to $N_{pix}\sigma_e^2$ and so the photon signal is considered a constant of $N_{pe} = 256$ rather than Poisson noise. These considerations reduce the SNR to be approximated by (5.30).

$$SNR \approx \frac{N_{pe}}{8\sigma_e} = \frac{32}{\sigma_e} \tag{5.30}$$

The phase screens used for this simulation are composed only of zeros (i.e., flat wave fronts) to isolate the effect of noise. This simulation used 600 samples where the closed loop system operates on electron noise. It was determined that there is no significant change in noise amplification when the waffle modes are suppressed. This is in contrast to previous chapters where it was shown that the error due to noise is reduced when the waffle modes are removed. The source of this difference is the noise in this experiment is on CCD pixels prior to calculating the centroid. The application of noise in previous chapters is equivalent to post-centroid noise.

A soft threshold of σ_e is applied to the SH-WFS image (i.e., the image is subtracted by σ_e and all negative values are set to 0) prior to centroiding. The error caused by the noise is normalized by the rms amplitude of the phase screen for Figure 5.23. The measured noise contaminated intensities are soft thresholded and then centroided as in (5.21) and (5.22).

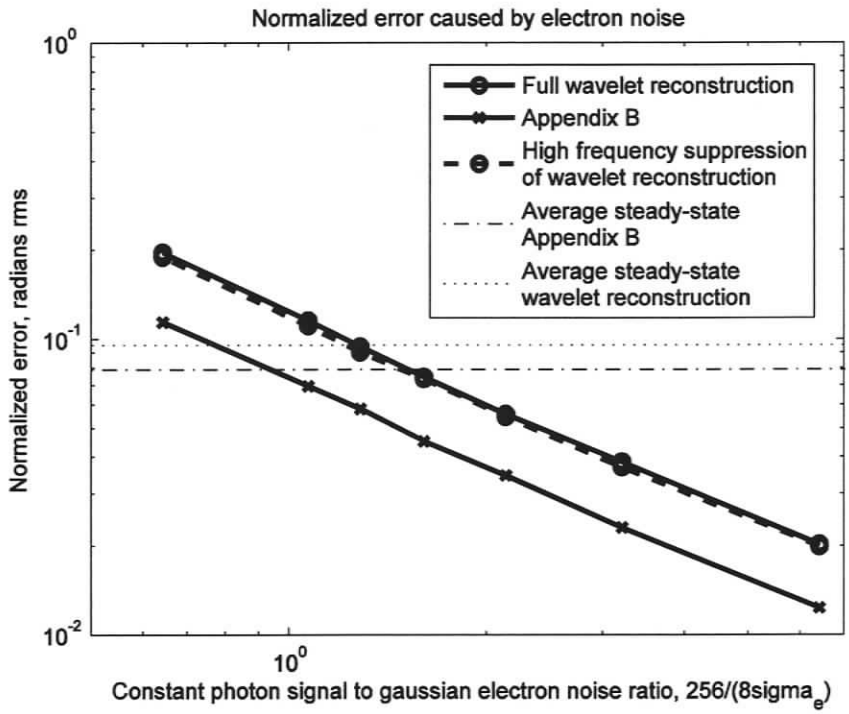

Figure 5.23: Normalized rms error caused by CCD readout noise

The plot of Figure 5.23 shows that the wavelet reconstruction method reacts to electron read-out noise approximately 70% more for all points on the curve than the SVD-based methods discussed in Appendix B. For convenience, the horizontal lines from Figure 5.21 have been superimposed in this plot to compare error due to noise and fitting errors. These lines have been converted from mean square error to rms error.

5.7 Conclusions

It was shown that the SH-WFS model used is not an ideal gradient sensor. The expectation of a measurement with a curl of zero is not necessarily applicable to SH-WFSs. This provided a realistic sensor model to test the robustness of the wavelet reconstruction approach. It has been shown in all experiments that this method is stable in closed loop despite the large modelling errors of the sensor.

The average rejection ratios for the first 20 lowest order Zernike modes (except piston) were nearly overlapping for an SVD-based pseudo-inverse and the wavelet reconstruction method. The W-T methods had a higher steady-state error and a higher amplification of sensor noise when splitting the signals in the wavelet decomposition space compared to the SVD-based signal separation method given in Appendix B. This baseline method fully measures the interactions between the mirrors and the sensor and this method has no dependence on the curl being zero.

When comparing the computation costs of the methods, the SVD-based method (including a pseudo-inverse reconstruction matrix) requires approximately 160 times the number of flops necessary for the wavelet-based approach for a 32×32 tweeter and 8×8 woofer. When comparing only the control algorithms (since there are much faster algorithms than using large reconstructor matrices), splitting the W-T signals with matrices requires 40 times the number of flops as the method that splits the signals in Haar wavelet decomposition space. The required stroke for both approaches was determined to be nearly equal as shown in Figure 5.20.

The wavelet reconstruction approach provides opportunities to reduce the system delay, which can improve the quality of correction for moving atmospheric turbulence. This improvement is despite the higher error found when converging on static phase screens and a higher amplification of SH-WFS noise. When the control system lag dominates the error (such as wind velocity = 25 m/s), faster algorithms that obtain an approximate answer can provide better correction than those that take longer to calculate the least-squares answer. This benefit is cancelled if the SNR is reduced to 3.2, which is the SNR from a magnitude 9 star in the R band sampled at 1kHz by 749 lenslets on an 8 meter class telescope with an electron read noise of 10 electrons rms per pixel per sample. Reducing the magnitude (i.e., brighter guide star), sample rate, number of lenslets or electron read noise all contribute to higher SNR.

Chapter 6

Conclusions

6.1 Summary

This dissertation has shown that it is possible to compute the integration of 2-D partial differences in as few as 10.17 flops per pixel of the result. This can be reduced to 5.33 operations per pixel when the high frequency component is ignored. The method is to find an estimation of the Haar wavelet transform of the integrated data as a preliminary step to obtaining the integrated data. This method is defined for data sets that are of size $2^M \times 2^M$. In the case where the data set smaller than $2^M \times 2^M$ or if the data set is convex, methods have been provided to extrapolate the missing data.

This method is well suited for finding the numerical integral of partial differences, as demonstrated by the photograph examples. Its robustness was tested in the case of AO simulations to determine whether it is useful for integrating discretely sampled gradient measurements as well.

6.2 Conclusions

The Haar wavelet filterbank can be factored in a manner that allows computed partial differences to be used as input rather than a data set of intensities. The method is lossless in the absence of noise or sensor modelling errors. The method is robust as shown by the reconstruction of photographs in conditions of extreme noise. A 14 dB improvement of the reconstruction to error ratio compared to the signal to noise ratio was achieved for a 4-megapixel image.

The partial difference input may be replaced by discretely sampled gradient measurements to obtain an approximate result. It has been demonstrated that this numerical integration method converges in closed loop systems that use gradient sensors, further demonstrating the robustness of the algorithm.

The integration has been obtained without the need to define the Dirichlet boundary condition (i.e., knowing the edge values of the integration) (Tumblin, Agrawal and Raskar 2005) or the Neumann boundary condition (i.e., forcing the edge gradients to be zero in the direction normal to the boundary) (Fattal, Lischinski and Werman 2002). Only the gradient measurements are necessary to obtain the result of integration. Only the constant offset and the highest frequency waffle mode are lost in the integration due to invisibility of these modes to a Fried aligned partial difference operation. Hudgin aligned detectors may detect the waffle mode, but the constant term is still lost.

The performance of the extrapolation reconstruction process is evaluated with respect to computation time and reconstruction error inside the circular aperture. Results indicate that the diagonally flat extrapolation method was virtually indistinguishable from the unrealistic ideal case of having a square aperture when applied to the Hampton-

Agathoklis integration method. The extrapolation using the diagonally flat techniques does not cause errors when there is no noise, evidenced by the lack of a knee point in the residual error plots. Further, the computation cost of this extrapolation technique is $O(N^{1/2})$ and is expected to be insignificant cost compared to the $O(N)$ reconstruction method. It is not a requirement of this method that the boundary be circular. A convex boundary is sufficient for the equations provided.

The curved mirror extrapolation techniques are shown to be better than even the ideal square aperture case when the gradient SNR is very low. This is due to the noise filtering effect of interpolating data in the interior of the pupil to obtain the extrapolated data. This approach comes with a cost that is $O(N)$ flops.

The octagonal mirror reflection improves the diagonal extrapolation approach with no additional computation cost. This is the recommended combination of methods for high SNR. When applied to the circular extrapolation method, the computation cost is decreased by at least a factor of 5 for these examples.

The examples indicate that the Hampton-Agathoklis integration method can be used for reconstruction of gradient data from circular apertures with results as good as if a square aperture were available.

A simulated SH-WFS, either by the methods of (Vogel and Yang 2006) or by modelling the Fourier optics, produce gradient measurements that have significant curl terms. This causes reconstruction errors in this numerical integration process since the method expects the curl to be zero. If it is desirable to improve the reconstruction error of this algorithm, the estimation found by integration may be used as the input to iterative Poisson solvers similarly to the MGCG algorithm of (Vogel and Yang 2006).

The quality of the integration can be less important than the speed of the integration when the input changes relatively rapidly, as can be the case in AO or when the computer hardware is not powerful for cases like a gradient camera. It is shown that if the speed of the algorithm can significantly reduce the delay of the control system, improved results can be obtained when following a dynamic input.

6.3 Future Work

6.3.1 Error reduction

It was shown by (Vogel and Yang 2006) that using a small number of iterations of the Poisson solver at each resolution of the MGCG method leads to improved performance. It is expected that using the estimation of the integral found by the methods of this dissertation will also be improved by using this widely accepted method to converge closer to the optimal result. The computational cost can be kept low by applying this algorithm to each resolution of the synthesized result of the wavelet integration process.

6.3.2 Adaptive optics systems

6.3.2.1 Prediction

It was shown that the system delay in AO can dominate the residual error when the turbulence is moving quickly. A common area of research in video image processing is to develop compression algorithms that allow for predicting the motion of blocks within the image. It has been shown by (Park and Kim 2000) and (Li 2003) that this may be accomplished in a wavelet decomposition space directly despite the shift variant nature of wavelets. This could lead to predictive AO control systems using a wavelet space.

6.3.2.2 Tomography

One research area for advanced AO systems is to use tomography (Tokovinin and Viard 2001) to reconstruct the volume of turbulence above the telescope from the measurements of several SH-WFSs. An intermediate step in this process is to integrate the gradient measurements of the SH-WFSs into phase (D. Gavel 2004). An area of research would be to determine if this integration method is adequate for tomography.

6.3.2.3 Control algorithm on wave front sensors

Since this test is simulation-based, one area of future work is to implement this on real AO hardware. In order to reduce data transmission from the SH-WFS, AO systems may perform centroiding on board the camera. The amount of computations necessary to centroid even 2×2 pixels is at least 10 flops per lenslet, including thresholding. When a lenslet has 4×4 pixels, there can be as many as 78 flops per lenslet.

The entire wavelet based W-T method can be as few as 14 flops per active tweeter actuator. There is a strong possibility that if a centroiding algorithm can be implemented on a camera, then the entire control system can also be implemented in the camera. This would entirely remove the necessity of an additional control computer. The implication of this is that the only data transmission would be directly to the mirrors and this would reduce transmission time to 33% of systems that require a computer.

6.3.2.4 Compression

Even when a control computer is considered necessary, reduction in transmitted bits is possible. The Haar wavelet decomposition is a method used for lossless data compression (Vetterli and Kovacevic 1995). The data may be compressed further if small values are approximated as zeros. Performing the gradient analysis component to obtain the Haar

decomposition on board the camera would reduce the transmitted data from the WFS to the computer by at least 50% before compression. Then, synthesis of the decomposition is performed in the computer for use in the control algorithm.

6.3.2.5 Woofer-Tweeter

It is known that the wavelet reconstruction-based W-T method does not split the woofer and tweeter signals in an orthogonal manner. There may be further performance gains by separating the signals at the 16×16 resolution or higher, rather than 8×8. Methods would also be required for using woofer mirrors that do not directly fit one of the resolutions of the wavelet synthesis process.

6.3.3 Increase of dimensions

The method of factoring the Haar wavelet filterbanks shown in Chapter 3 is applicable to any higher dimension of filterbanks. The difficulty is that the number of filters that must be determined increases rapidly with the number of dimensions. Not counting redundant filters to improve performance, the number of filters necessary is $\left(2^d - d\right)\left(2^d - 1\right)$ due to the branching structure, where d is the number of dimensions. The 1-D numerical integration is solved with 1 partial difference-based filter. For 2-D, there are a minimum of 6 filters required (though 8 were used to provide redundancy). For 3-D there are at least 35 filters and for 4-D, 180, with many more if all possible redundancies are included. The algebra can be worked out for any dimension, but a more useful product would be an algorithm that can determine these filters automatically.

Bibliography

Andersen, D. R., M. Fischer, R. Conan, M. Fletcher, and J.-P. Véran. "VOLT: the Victoria Open Loop Testbed." *Proceedings of SPIE 7015.* 2008. 70150H.

Barchers, J. D. "Closed-loop stable control of two deformable mirrors for compensation of amplitude and phase fluctuations." *Journal of the Optical Society of America* 19, no. 5 (2002): 926-945.

Conan, R., C. Bradley, P. Hampton, O. Keskin, A. Hilton, and C. Blain. "Distributed modal command for a two-deformable-mirror adaptive optics system." *Applied Optics* 46, no. 20 (2007): 4329-4340.

Cornelissen, S. A., P. A. Bierden, and T. G. Bifano. "Development of a 4096 element deformable mirror for high contrast astronomical imaging." *Proceedings of SPIE 6306.* 2006. 630606.

Dekany, R., et al. "Adaptive optics requirements for TMT." Edited by D. B. Calia, B. L. Ellerbroek and R. Ragazzoni. *Proceedings of SPIE 5490.* 2004. 879-890.

Donoho, D. L. "De-noising by soft thresholding." *IEEE Transactions on Information Theory* 41, no. 5 (May 1995): 613–627.

Dowla, F. U., J. M. Brase, S. S. Olivier, and C. A. Thompson. "Fast Fourier and wavelet transforms for wavefront reconstruction in adaptive optics." *Proceedings of SPIE 4124.* 2000. 118-127.

Ellerbroek, B. L. "Efficient computation of minimum-variance wave front reconstructors with sparse matrix techniques." *Journal of the Optical Society of America A* 19, no. 9 (Sept. 2002): 1803-1816.

Ellerbroek, B. L., et al. "A conceptual design for the Thirty Meter Telescope adaptive optics systems." *Proceedings of SPIE 6272.* 2006. 62720D.

Fattal, R., D. Lischinski, and M. Werman. "Gradient domain hih dynamic range compression." *Association for Computin Machinery*, 2002: 249-256.

Fernández, E. J., I. Iglesias, and P. Artal. "Closed-loop adaptive optics in the human eye." *Optics Express* 26, no. 10 (2001): 746-748.

Frankot, R. T., and R. Chellappa. "A method for enforcing integrability in shape from shading." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10 (1988): 439-451.

Freischlad, K., and C. Koliopoulos. "Modal estimation of a wave-front from difference measurements using the discrete Fourier transform." *Journal of the Optical Society of America. A* 3 (Nov. 1986): 1852-1861.

Fried, D. L. "Least-square fitting a wave-front distortion estimate to an array of phase-difference measurements." *Journal of the Optical Sociey of America* 67 (Mar. 1977): 370-375.

Gavel, D. "Performance of Adaptive Optics Systems." *Center for Adaptive Optics Summer School.* UC Santa Cruz, 2008.

Gavel, D. T. "Suppressing anomalous localized waffle behavior in least-squares wavefront reconstructors." Edited by P. L. Wizinowich and D. Bonaccini. *Proceedings of SPIE 4839.* 2003. 972-980.

Gavel, D. "Tomography for multiconjugate adaptive optics systems using laser guide stars." *Proceedings of SPIE 5490.* 2004. 1356-1373.

Gilles, L. "Order-N sparse minimum-variance open-loop reconstructor for extreme adaptive optics." *Optics Letters* 28, no. 20 (Oct. 2003): 1927-1929.

Gilles, L., B. Ellerbroek, and C. Vogel. "Multigrid preconditioned conjugate-gradient method for large-scale wave-front reconstruction." *Journal of the Optical Society of America A* 19 (Sept. 2002): 1817-1822.

Hampton, P. J. *Design and Implementation of Type 1 and Type 2 control systems with residual feedback for modal adaptive optics.* Victoria, BC: University of Victoria, 2005.

Hampton, P. J., C. Bradley, A. Hilton, and P. Agathoklis. "Multiple Mirror Control System." Edited by D. B. Calia, B. L. Ellerbroek and R. Ragazzoni. *Proceedings of SPIE 5490.* 2004. 1374-1383.

Hampton, P. J., P. Agathoklis, and C. Bradley. "A new wave-front reconstruction method for adaptive optics systems using wavelets." *IEEE Journal of Selected Topics in Signal Processing* 2 (Oct. 2008): 781-792.

Hampton, P. J., P. Agathoklis, and C. Bradley. "Wave Front Reconstruction over a Circular Aperture using Gradient data Extrapolated via the Mirror Equations." *Applied Optics* 48, (2009): 4018-4030

Hampton, P. J., R. Conan, C. Bradley, and P. Agathoklis. "Control of a woofer tweeter system of deformable mirrors." Edited by H. Lewis and A. Bridger. *Proceedings of SPIE 6274.* 2006. 62741Z.

Hampton, P. J., R. Conan, O. Keskin, P. Agathoklis, and C. Bradley. "Self characterization of linear and non-linear adaptive optics systems." *Applied Optics* 47 (2008): 126-134.

Hartmann, J. "An Improvement of the Foucalt Knife-Edge Test in the Investigation of Telescope Objectives." *Translated from Sitzungsberichte der K. Preussichen Akademie der Wissenshaften Session.* 1907.

Herrmann, J. "Least-squares wavefront errors of minimum norm." *Journal of the Optical Society of America* 70 (1980): 28-35.

Hockney, R. W. "A fast direct solution of Poisson's equation using Fourier analysis." *Journal of the Association for Computing Machinery* 12, no. 1 (1965): 95-113.

Hu, S., B. Xu, X. Zhang, J. Hou, J. Wu, and W. Jiang. "Double-deformable mirror adaptive optics system for phase compensation." *Applied Optics* 45, no. 12 (2006): 2638-2642.

Hudgin, R. "Wave-front reconstruction for compensated imaging." *Journal of the Optical Society of America* 67 (Mar. 1977): 375-378.

Keskin, O., R. Conan, P. Hampton, and C. Bradley. "Derivation and experimental evaluation of a point spread function reconstruction from a dual deformable mirror adaptive optics system." *Optical Engineering* 47 (2008): 046601:1-10.

Kolaczyk, E. D. "Methods for analyzing certain signals and images in astronomy using Haar wavelets." *Conference Record of the 31st Asilomar Conference on Signals, Systems & Computers.* 1997. 80–84.

Kovesi, P. "Shapelets correlated with surface normals produce surfaces." *Proceedings of the 10th IEEE International Conference on Computer Vision.* 2005. #1550-5499/05.

Li, X. "New Results of Phase Shifting in the Wavelet Space." *IEEE Signal Processing Letters* 10, no. 7 (2003): 193-195.

Martin, F., and R. Conan. "Optical parameters relevant for high angular resolution at Paranal from GSM instrument and surface layer contribution." *Astrophysics Supplemental Series* 144 (2000): 39-44.

Miller, L. R. "Magnetic Gradient Sensor." *IEEE Transactions on Geoscience Electronics*, 1965: 10-13.

Nakajima, T. "Signal-to-noise ratio of the bispectral analysis of speckle interferometry." *Journal of the Optical Society of America A* 5 (1988): 1477-.

Noll, R. J. "Zernike polynomials and atmospheric turbulence." *Journal of the Optical Society of America* 66, no. 3 (Mar. 1976): 207-211.

Park, H.-W., and H.-S. Kim. "Motion Estimation Using Low-Band-Shift Method." *IEEE Transactions on Image Processing* 9, no. 4 (2000): 577-587.

Peréz, P., M. Gangnet, and A. Blake. "Poisson Image Editing." *Association for Computing Machinery*, 2003: 313-318.

Poyneer, L. A., and J.-P. Véran. "Optimal modal Fourier transform wave-front control." *Journal of the Optical Society of America A* 22 (2005): 1515-1526.

Poyneer, L. A., B. A. Macintosh, and J.-P. Véran. "Fourier-transform wavefront control with adaptive prediction of the atmosphere." *Journal of the Optical Society of America A* 24, no. 9 (Sept. 2007): 2645-2660.

Poyneer, L. A., D. T. Gavel, and J. M. Brase. "Fast wave front reconstruction in large adaptive optics systems with the Fourier transform." *Journal of the Optical Society of America A* 19 (Oct 2002): 2100-2111.

Ragazzoni, R. "Pupil plane wavefront sensing with oscillating prism." *Journal of Modern Optics* 43 (1996): 289-293.

Rigaut, F., B. L. Ellerbroek, and M. J. Northcott. "Comparison of curvature-based and Shack-Hartmann based adaptive optics for the Gemini telescope." *Applied Optics* 36, no. 13 (1997): 2856-2868.

Roddier, C., and F. Roddier. "Wave-front reconstruction from defocused images and the testing of ground-based optical telescopes." *Journal of the Optical Society of America A* 10 (Nov. 1993): 2277-2287.

Roddier, F. *Imaging through the atmosphere.* Cambridge, UK: Cambridge University Press, 1999.

Talmi, A., and E. N. Ribak. "Wavefront reconstruction from its gradients." *Journal of the Optical Society of America A* 23, no. 2 (Feb. 2006): 288-297.

Tanaka, N. "A simple but efficient preconditioning for conjugate gradient Poisson solver using Haar wavelet." *International Journal of Wavelets, Multiresolution and Information Processing* 4, no. 2 (2006): 273–284.

Tippur, H. V., S. Krishnaswamy, and A. J. Rosakis. "A coherent gradient sensor for crack tip deformation measurements: analysis and experimental results." *International Journal of Fracture* 48 (1991): 193-204.

Tokovinin, A., and E. Viard. "Limiting precision of tomographic phase reconstruction." *Journal of the Optical Society of America A* 18 (2001): 873-882.

Tumblin, J., A. Agrawal, and R. Raskar. "Why I want a gradient camera." *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2005).* 2005. 103-110.

Tyson, R. K. "Adaptive optics and ground-to-space laser communications." *Applied Optics* 35, no. 19 (1996): 3640-3646.

van Oudheusden, B. W. "The behaviour of a thermal-gradient sensor in laminar and turbulent shear flow." *Journal of Physics E: Scientific Instruments* 22 (1989): 490-498.

In *Wavelets and Subband Coding*, by M. Vetterli and J. Kovacevic, Ch. 6. Upper Saddle River, NJ: Prentice-Hall, 1995.

Vogel, C. R., and Q. Yang. "Multigrid algorithm for least-squares wave front reconstruction." *Applied Optics* 45, no. 4 (Feb. 2006): 705-715.

Wallace, B. P., P. J. Hampton, C. H. Bradley, and R. Conan. "Evaluation of a MEMS deformable mirror for an adaptive optics testbench." *Optics Express* 14, no. 14 (2006): 10132-10138.

Watkins, D. S. *Fundamentals of Matrix Computations.* New York, USA: John Wiley & Sons, Inc., 2002.

Widiker, J. J., S. R. Harris, and B. D. Duncan. "High speed Shack-Hartmann wave front sensor design with commercial off-the-shelf optics." *Applied Optics* 45, no. 2 (Jan. 2006): 383-395.

Zou, W., X. Qi, and S. A. Burns. "Wavefront-aberration sorting and correction for a dual-deformable-mirror adaptive-optics system." *Optics Letters* 33, no. 22 (2008): 2602-2604.

Appendix A: Self Characterization of Linear and Non-Linear AO Systems

The material in the following manuscript was originally published as

Self characterization of linear and non-linear adaptive optics systems
Hampton, P. J., R. Conan, O. Keskin, P. Agathoklis, and C. Bradley.
Applied Optics 47 (2008): 126-134.

Self Characterization of Linear and Non-Linear Adaptive Optics Systems

Peter J. Hampton[1], Rodolphe Conan[2], Onur Keskin[2], Colin Bradley[2], Pan Agathoklis[1]

University of Victoria, Adaptive Optics Laboratory, Engineering Lab Wing B 133,

PO Box 3055 STN CSC, Victoria, BC, Canada, V8W 3P6

[1] Department of Electrical and Computer Engineering

[2] Department of Mechanical Engineering

phampton@engr.uvic.ca

This paper presents methods used to determine the linear or non-linear static response and the linear dynamic response of an Adaptive Optics system. This AO system consists of a non-linear MEMS Deformable Mirror, a linear Tip-Tilt Mirror, a control computer and a Shack Hartmann Wave Front Sensor. The system is modeled using a single-input-single-output structure to determine the one dimensional transfer function of the dynamic response of the chain of system hardware. An AO system has been shown to be able to characterize its own response without additional instrumentation. Experimentally determined models are given for a TTM and a DM.

OCIS Codes: 010.1080, 230.4320, 010.7350, 230.3990

1. Introduction

A deformable mirror (DM) is situated in an adaptive optics (AO) system between the observed object (such as a star) and the science camera, as shown in Fig. 1. The wave front, aberrated by the Earth's atmosphere φ_{Atm}, reflects off the DM and becomes a residual wave front error, φ_{Err}. It is the goal of the control system to minimize the variance of φ_{Err}.

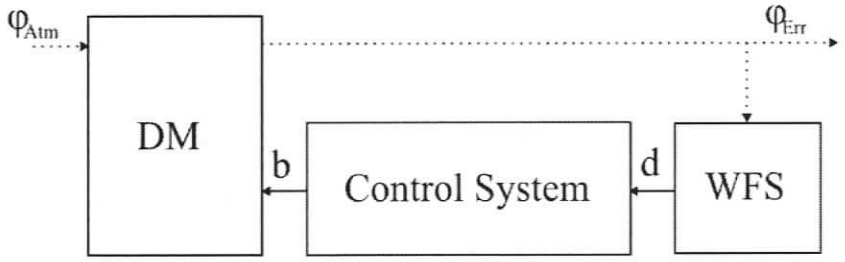

Figure 1: Layout of a simple AO system

The signal vector *b* is the output of the computer used to drive the DM. The signal vector *d* is the measurement of spot displacements from the Shack Hartmann Wave Front Sensor (WFS). The block between *d* and *b* is called the 'Control System' rather than the 'Computer' since there are WFS image processing algorithms conducted within the computer that convert the image to a measurement vector.

Design of sophisticated control algorithms requires an accurate model of the controlled system. Models of the system could be constructed by combining individually modeled responses of the DM and WFS as in [1, 2]. It has also been shown in [3] that the tip-tilt mirror (TTM) would have a different model than the DM. Instead of modeling the mirror and sensor separately, this paper describes an experimental method for determining a single dynamic system model of the WFS, computation delays and the DM. This process is repeated replacing the DM with the TTM. Furthermore, this method does not require additional measurement equipment beyond the functional AO system.

In order to simplify the following derivations, it is assumed that the mid-point of DM actuator stroke (i.e. piston) is at 0.

2. Linear Model of the Static Interaction

The first step in modeling the external AO system, i.e. exclusive of the Control System, is to determine the static model of the components between the actuator voltage vector, *b*,

and the steady-state WFS sensor measurement, d, of Fig. 2. Vector b is an $n\times 1$ vector where n is the number of actuators of the DM. Vector d is a $w\times 1$ vector where w is the number of WFS measurements; this is also twice the number of WFS lenslets [4]. At this stage, the DM is modeled as two operators: (i) matrix N, which converts b to a continuous mirror surface shape as in Eq. 1, and (ii) addition, which models the reflection of the incoming light wave front off the mirror surface as in Eq. 2. For simplicity, the WFS will be modeled as a single operator, P, which converts the residual error of the wave-front into the WFS measurements. This is added to the sensor noise, q_{wfs}, and the bias of the WFS measurement, d_0. The result is the final measurement, d, which can be expressed as a function of b and time, t, as in Eq. 3.

$$\varphi_{Cor}(b) = Nb \tag{1}$$

$$\varphi_{Err}(b,t) = \varphi_{Atm}(t) + \varphi_{Cor}(b) \tag{2}$$

$$d(b,t) = d_0 + P\varphi_{Err}(b,t) + q_{wfs}(t) = d_0 + PNb + P\varphi_{Atm}(t) + q_{wfs}(t) \tag{3}$$

The combination of PN acts as a $w\times n$ matrix, even though the intermediate space between P and N is a continuous surface. It is necessary to measure PN in order to understand how changes to the actuator vector, b, alter the WFS measurement, d. The ideal measurement would be

$$\frac{\Delta d}{\Delta b_i} = PNI_i \tag{4}$$

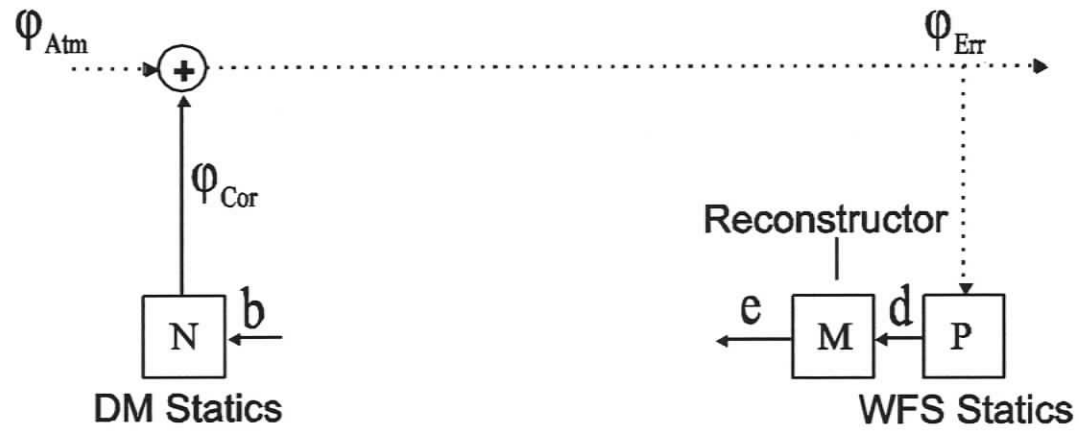

Figure 2: Block model of the static DM, WFS and Reconstructor.

In Eq. 4, Δb_i is the change in displacement of the i^{th} DM actuator; Δd is the change in the WFS sensor measurement caused by the actuator displacement; and I_i is the i^{th} column of an $n \times n$ identity matrix. As shown in Eq. 3, there are disturbances that will distort the ideal measurement defined in Eq. 4. Namely, the sensor noise, q_{wfs}, and the influence of the atmospheric turbulence, $P\varphi_{Atm}$. To minimize these detrimental issues, this measurement is made with a 639 nm laser diode at the input of the AO system. This laser diode is operated at a current that is insufficient to cause lasing. This creates non-coherent red light and speckles are avoided. This light is passed through a 5 μm pin hole, collimated and then passed through an 18 mm diameter pupil. This beam is resized to 2.7 mm at the DM and 1.7 mm at the WFS [6]. This reduces the intervening turbulence to only the volume of air inside the AO system. Further, the effect of q_{wfs} can be reduced by using a large signal-to-noise ratio (SNR). In this system, the SNR can be increased by increasing the power of the laser, increasing the integration time of the WFS camera or using large changes in the actuator displacement, Δb_i.

One method for measuring PN is to find a linear approximation from two different voltage measurements. This is a similar process as finding the slope of a line by

determining two data points. For convenience, the two levels, $\pm h$, are chosen to eliminate actuator quantization error. The i^{th} actuator is pushed to a value of h while all other actuators are set at 0, which can be represented by $h\mathrm{I}_i$. The WFS measurement is recorded after a delay of 100 ms to ensure that the steady-state position is reached. This measurement is then repeated when the actuator is pulled to $-h$.

$$\begin{aligned}\frac{\Delta\hat{d}}{\Delta b_i} &= \frac{d\left(hI_i, t_{2i}\right) - d\left(-hI_i, t_{2i+1}\right)}{2h} \\ &= PNI_i + \frac{P\phi_{Atm}(t_{2i}) - P\phi_{Atm}(t_{2i+1}) + q_{wfs}(t_{2i}) - q_{wfs}(t_{2i+1})}{2h} \\ &\approx PNI_i\end{aligned} \tag{5}$$

where d is shown as a function of b and t that was defined in Eq. 3 and h is the quantized magnitude of the actuator. The approximation is due to the controlled environment where $\varphi_{Atm}(t)$ is negligible and a high SNR is used to make $q_{wfs}(t)$ negligible.

$$\hat{D} = \begin{bmatrix} \frac{\Delta\hat{d}}{\Delta b_1} & \frac{\Delta\hat{d}}{\Delta b_2} & \cdots & \frac{\Delta\hat{d}}{\Delta b_{n-1}} & \frac{\Delta\hat{d}}{\Delta b_n} \end{bmatrix} \tag{6}$$

The matrix in Eq. 6 is referred to as the "Zonal Interaction Matrix" and is a w x n matrix that is an estimate of PN. This could also be found using a Hadamard-type technique [5] which moves many actuators at once to increase SNR. The structure of $\hat{D}$ as a model of the conversion from actuator space to sensor space leads to an interesting structure of its Singular Value Decomposition (SVD) [7]. The SVD decomposes a single matrix into three matrices; (i) U, (ii) Σ and (iii) V. Applying the SVD to $\hat{D}$ provides

$$U\Sigma V^T = \hat{D} \tag{7}$$

where the U and V matrices are orthonormal as shown by the properties of their column vectors in Eq. 8 and 9. The Σ matrix is a sparse rectangular matrix as shown by the properties of its elements in Eq. 10.

$$u_i^T u_j = \begin{cases} 1 & for \quad i = j \\ 0 & for \quad i \neq j \end{cases} \tag{8}$$

$$v_i^T v_j = \begin{cases} 1 & for \quad i = j \\ 0 & for \quad i \neq j \end{cases} \tag{9}$$

$$\sigma_{i,j} = \begin{cases} non-negative & for \quad i = j \\ 0 & for \quad i \neq j \end{cases} \tag{10}$$

where u_i and u_j are the i^{th} and j^{th} column vectors of the w x w U matrix. Similarly, v_i and v_j are the i^{th} and j^{th} column vectors of the n x n V matrix. The element on the i^{th} row and j^{th} column of the w x n Σ matrix is $\sigma_{i,j}$. The definition in Eq. 10 shows that the Σ matrix is sparse since it only has the potential to have non-zero values on its main diagonal. Since w is greater than n in this application, there are at most n non-zero elements of Σ. These values are arranged in descending order as is the definition of the SVD.

The properties shown in Eq. 8 to 10 allow for $\hat{D}$ to be represented as a sum of matrices that are each determined from the vector pairs as shown in Eq. 11.

$$\hat{D} = \sum_{i=1}^{n} u_i \sigma_{i,j} v_i^T \tag{11}$$

This representation shows that if a single column vector of V is multiplied by $\hat{D}$ then the result is the respective column of U that is scaled by the respective singular value. If the change in the actuator voltage vector, b, is equal to the i^{th} column of V then

$$\sigma_{i,j} u_i = \hat{D} v_i \approx PNv_i = PN(\Delta b) \tag{12}$$

Since v_i is in actuator space and u_i is in the WFS sensor space, one interpretation is that the application of a step change of v_i to b will cause a steady-state change of d equal to $\sigma_{i,i}u_i$. Since u_i and v_i are unit vectors, this means that $\sigma_{i,i}$ represents the magnitude of the change in the sensor space for a given unit magnitude change in actuator space. Very small singular values relative to the first singular value indicate that the respective shape applied to the actuators is difficult to detect. Retaining the small singular values is a problem as can be seen by the definition of the pseudo-inverse of $\hat{D}$.

$$\sum_{i=1}^{n-m} \frac{v_i u_i^T}{\sigma_{i,i}} = V\Sigma^{+}U^T = \hat{D}^{\dagger} \tag{13}$$

where m is the number of modes dropped due to the inability of the WFS to detect them. As can be seen from Eq. 13, the singular value is now in the denominator. The very small singular values are discarded because they would dominate the pseudo-inversion. These correspond to actuator modes that are difficult to detect, so the signal of the respective sensor modes would be dominated by sensor noise.

If the measurement process of $\hat{D}$ is such that the disturbances in Eq. 5 can be considered insignificant and the steady-state model, PN, of the interaction between the actuator voltage, b, and the sensor measurement, d, is linear then the terms of Eq. 12 can be considered equal. If this is the case, then it is known that multiplying the sensor signal by the pseudo-inverse provides a positive unit gain feedback path for all the retained modes of $\hat{D}^{\dagger}$, as shown in Eq. 14.

$$\hat{D}^{+}\hat{D}v_i = \begin{cases} v_i & for \quad 1 \le i \le n-m \\ 0 & elsewhere \end{cases} \tag{14}$$

However, in control system development it can be more desirable to have the system defined to incorporate a negative feedback path, rather than positive. For this reason, the Reconstructor Matrix, M, will be defined to be the negative of $\hat{D}^{\dagger}$.

$$M = -\hat{D}^{\dagger} \tag{15}$$

$$e = Md \tag{16}$$

where the Reconstructor Matrix, M, is the conversion from the measured WFS error, d, into an actuator error signal, e. Then, the model of how the change of actuator voltage, Δb, changes the actuator error measurement, Δe, is shown as Eq. 17 and is based on the definitions of Eq. 12 to 16.

$$\begin{aligned} M_b PN(\Delta b) &= \Delta e \\ M_b PN v_{b,i} &= \begin{cases} -v_{b,i} & for \quad 1 \le i \le n-m \\ 0 & elsewhere \end{cases} \end{aligned} \tag{17}$$

where the subscript b indicates that the respective matrix and vectors are defined with respect to the actuator voltage vector, b. This naming convention will be important in the next section since it was found that the MEMS DM is not linear, so its final Reconstructor will not be with respect to b. Section 3 presents the method used to ensure linearity for the MEMS DM employed in this experiment. The definition in Eq. 17 is sufficient for this TTM since it has a linear displacement response to the input voltage.

3. Non-linear Voltage Response

Linear control theory is developed assuming the closed loop gain of the system is not dependant on the state of the system components, i.e. the displacement of an actuator. When a device is non-linear, the gain of that device is dependant on its state. Subsequently, the closed loop gain of an AO system would change depending on the displacement of non-linear DM actuators.

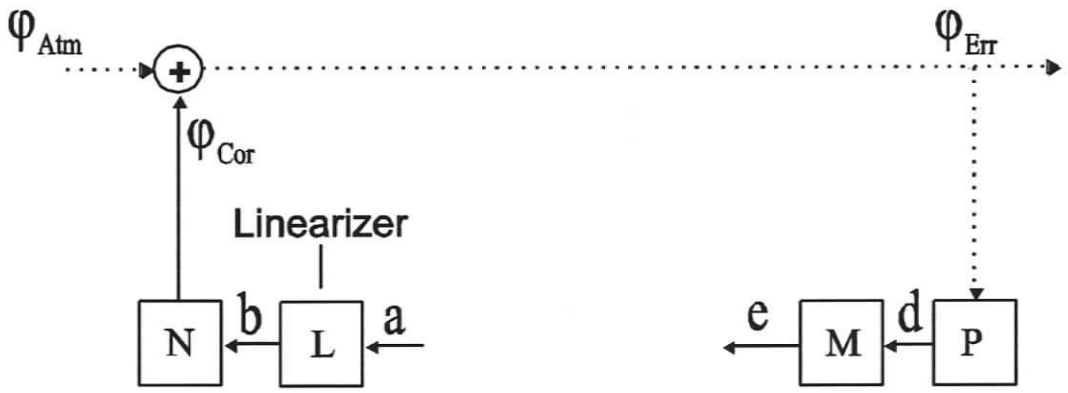

Figure 3: Inclusion of linearization block, L

This section describes a process for developing the L block in Fig. 3 to convert the non-linear DM response, N, into a linear response of NL. This process uses the Reconstructor matrix, M_b, derived in the previous section using the SVD of the Interaction matrix. It has been shown [8] that the response of a MEMS DM is quadratic with respect to voltage by measuring the peak displacement with an interferometer. The displacement of another MEMS DM was measured with the WFS. The actuators were individually pushed to 15 different positions and the respective actuator error, e_i, was measured using an initial Reconstructor matrix, M_b. Furthermore, at high actuator voltage, the displacement response diverges from the single quadratic fit of [8]. A similar result was found for this system using interferometer measurements [6]. This paper does not dispute the claim that MEMS DMs have a quadratic displacement characteristic. Rather, the divergence may be unique to this power supply. The result is plotted in Fig. 4 (left). Two quadratic polynomials were then fit to these 8 data points to determine the coefficients in a least squares sense.

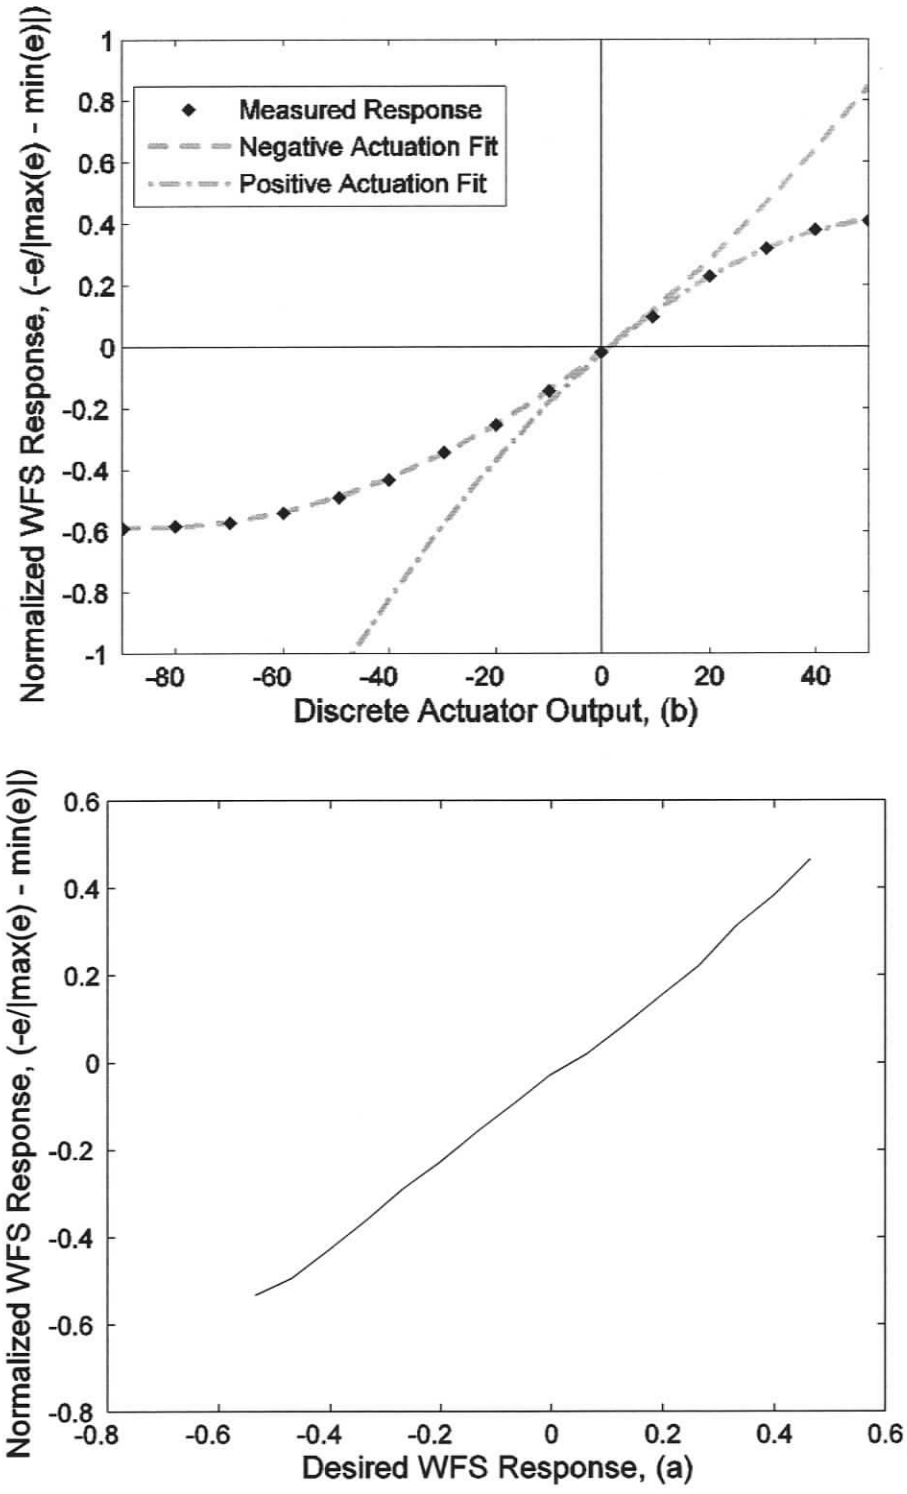

Figure 4: (Top) Non-linear response of the MEMS DM, (Bottom) Linear response when applying Eq. 21 to output signal

$$\begin{aligned} e_i(b_i,t) &= I_i^T M_b d(b_i,t) \\ e_i(b_i,t) &= I_i^T M_b \left(PNI_i b_i + P\varphi_{Atm}(t) + q_{wfs}(t) + d_0\right) \\ e_i(b_i) &\approx I_i^T M_b \left(PNI_i b_i + d_0\right) \end{aligned} \tag{18}$$

where d and e_i are shown as functions of b_i and t. The vector d_0 is the WFS measurement when the actuators are at the chosen 0 position. It does not need to be directly measured since the constant offset is ignored in Eq. 20. The reason for the approximation in Eq. 18 is the same as of Eq. 5. The values of $e_i(b_i)$ are somewhat arbitrary because the Reconstructor, M_b, was made using a non-linear MEMS DM. This measurement is normalized to make the measurement range equal to 1 so that the result is not dependant on the scaling of M_b. One quadratic function is fit to the positive actuation of b_i and another is fit to the negative actuation of b_i as shown in Eq. 19.

$$a_i = \begin{Bmatrix} c_{2,i}b_i^2 + c_{1,i}b_i + c_{0,i} \\ c_{5,i}b_i^2 + c_{4,i}b_i + c_{3,i} \end{Bmatrix} \begin{matrix} for \\ \end{matrix} \begin{matrix} b_i > 0 \\ b_i \le 0 \end{matrix} = -\frac{e_i(b_i)}{\left|\max(e_i) - \min(e_i)\right|} \tag{19}$$

The measurements are also shifted vertically due to d_0. Therefore, the measurements of $c_{0,i}$ and $c_{3,i}$ are related to d_0 and are forced to 0 to remove this sensor bias and it follows that when $a_i = 0$, b_i will also equal 0. The negative value of $e(b_i)$ is used so that MPNL will be negative as is MPN. Normalization by the measured range of e_i (denominator of the right hand side of Eq. 19) is used so that a_i can be later converted to a measure of the actuator displacement (i.e. micrometers) if necessary.

The output to the MEMS DM is then found by Eq. 21, which are the solutions for b_i in Eq. 20 with the constraint that b_i is equal to 0 when a_i equals zero.

$$0 = \begin{Bmatrix} c_{2,i}b_i^2 + c_{1,i}b_i - a_i \\ c_{5,i}b_i^2 + c_{4,i}b_i - a_i \end{Bmatrix} \begin{matrix} for \\ \end{matrix} \begin{matrix} a_i > 0 \\ a_i \le 0 \end{matrix} \tag{20}$$

$$b_i = \left\{ \begin{array}{ll} \dfrac{-c_{1,i} + sign(c_{1,i})\sqrt{c_{1,i}^2 + 4c_{2,i}a_i}}{2c_{2,i}} & for \quad a_i > 0 \\ \dfrac{-c_{4,i} + sign(c_{4,i})\sqrt{c_{4,i}^2 + 4c_{5,i}a_i}}{2c_{5,i}} & \quad a_i \le 0 \end{array} \right\} = L(a_i) \tag{21}$$

where a_i is the desired displacement of the actuator, b_i is the respective output signal.

Once this linearization is complete, M and V must be re-evaluated using the linear actuator response by the process explained in Section 2, but with respect to a_i instead of b_i. Once M_a and V_a have been determined, it can be shown that

$$\begin{aligned} M_a PN(\Delta b) &= \Delta e \\ M_a PN(L(\Delta a)) &= \Delta e \\ M_a PN(L(v_{a,i})) &= \Delta e = \left\{ \begin{array}{ll} -v_{a,i} & for \quad 1 \le i \le n-m \\ 0 & \quad elsewhere \end{array} \right. \end{aligned} \tag{22}$$

Eq. 22 shows that the system is configured to have a linear unit gain negative feedback path for the first n-m modes. This path connects the output, a, of the controller to its input, e.

Now that M_a is evaluated, the relationship between a and e can be measured for the full range of a. The measured actuator error, $-e$, is plotted against the linearized actuation, a, for a single actuator in Fig. 4 (right). This shows that using the non-linear response of Eq. 21 to cancel the non-linearity in Eq. 19 provides a nearly linear response.

4. Channel Identification

The introduction of the L block, as discussed in the previous section, had the effect that an input to the DM actuators corresponding to an actuator eigenvector, $v_{a,i}$, will lead to an output from the Reconstructor matrix block, M_a, which at steady state will be $-v_{a,i}$ or null. Consider an input signal to the actuator represented by

$$a[k] = v_{a,1}\theta_{in}[k] \tag{23}$$

where $\theta_{in}[k]$ is a scalar value at time sample k. The corresponding scalar output signal from the reconstructor matrix is

$$\theta_{out}[k] = -v_{a,1}^{T} e[k] = -v_{a,1}^{T} M_a d[k] \tag{24}$$

where the dynamics between signals $a[k]$ and $d[k]$ are unknown. Using $\theta_{in}[k]$ as a single dimensional input and using $\theta_{out}[k]$ as a single dimensional output has configured the AO system into a single-input-single-output (SISO) structure. This SISO behavior is studied by moving the entire mirror as a single mode, rather than as single actuators. Any non-linearity in the cross coupling of actuators may cause excitement of other modes. This effect is ignored in this process by using Eq. 24. Then the dynamics of the system can be represented using a linear system as shown in Fig. 5, where $\Theta_{in}(z)$ and $\Theta_{out}(z)$ are the z-transforms [10] of the sequences $\theta_{in}[k]$ and $\theta_{out}[k]$ respectively. Thus

$$\begin{aligned} \frac{\Theta_{out}(z)}{\Theta_{in}(z)} &= G_{Loop}(z) \\ &= \frac{\beta_w z^{-w} + \beta_{w+1} z^{-w-1} + \cdots + \beta_{x-1} z^{-x+1} + \beta_x z^{-x}}{1 - \alpha_1 z^{-1} - \alpha_2 z^{-2} - \cdots - \alpha_{y-1} z^{-y+1} - \alpha_y z^{-y}} \end{aligned} \tag{25}$$

The linear system assumption will be validated and the values of α_i, β_i, x, y and w will be determined using measurements of $\theta_{out}[k]$ given $\theta_{in}[k]$.

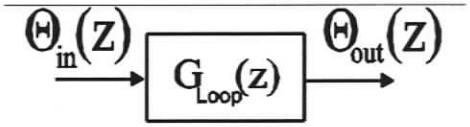

Figure 5: Concept for identifying an AO system's dynamics

Consider the difference equation corresponding to Eq. 25.

$$\theta_{out}[k] = \sum_{i=1}^{y} \alpha_i \theta_{out}[k-i] + \sum_{i=w}^{x} \beta_i \theta_{in}[k-i] \tag{26}$$

Where $\theta_{in}[k]$ must be a zero mean white noise input signal. Then, Eq. 26 is multiplied by $\theta_{in}[k - j]$ and the expected value operator, E{ } [9], is applied to this product. The result is the notation of Eq. 27 to Eq. 29. Variables i, j and k are all integer indices.

$$p(j) = \sum_{i=1}^{y} \alpha_i p(j-i) + \sum_{i=w}^{x} \beta_i c(j-i) \tag{27}$$

where

$$p(j) = E\{\theta_{out}[k]\theta_{in}[k-j]\} \tag{28}$$

$$c(j) = E\{\theta_{in}[k]\theta_{in}[k-j]\} = E\{\theta_{in}[k-j]\theta_{in}[k]\} = c(-j) \tag{29}$$

Eq. 27 can be written in a square matrix form for $w \le j \le x + y$ as

$$\begin{bmatrix} \alpha_1 & \cdots & \alpha_y & \beta_w & \cdots & \beta_x \end{bmatrix}^T = \begin{bmatrix} p(w-1) & \cdots & p(w-y) & c(0) & \cdots & c(w-x) \\ \vdots & & \vdots & \vdots & & \vdots \\ p(x+y-1) & \cdots & p(x) & c(x+y-w) & \cdots & c(y) \end{bmatrix}^{-1} \begin{bmatrix} p(w) \\ \vdots \\ p(x+y) \end{bmatrix} \tag{30}$$

The above equation can be used to obtain values for α_i, β_i for a given selection of the w, x, and y parameters. There can be many valid choices for these system parameters. However, the following conditions have to be satisfied;

(i) Causality: For an input signal with enough excitation, the correlation between the input and output signals should be zero until the input has traveled through the system. It follows that

$$w > 0 \tag{31}$$

$$x \ge w \tag{32}$$

$$y \ge 0 \tag{33}$$

$$p(j) = 0 \quad for \quad j < w \tag{34}$$

(ii) $G_{Loop}(z)$ shown in Eq 25 must be stable

(iii) The steady-state value of θ_{out} due to a unit step

$$\theta_{in}[k] = \begin{cases} 1 & for \quad k \geq 0 \\ 0 & elsewhere \end{cases} \tag{35}$$

should be 1 to conform with the definition of M. Using the z-transform, this condition implies

$$\sum_{i=w}^{x} \beta_i = 1 - \sum_{i=1}^{y} \alpha_i \tag{36}$$

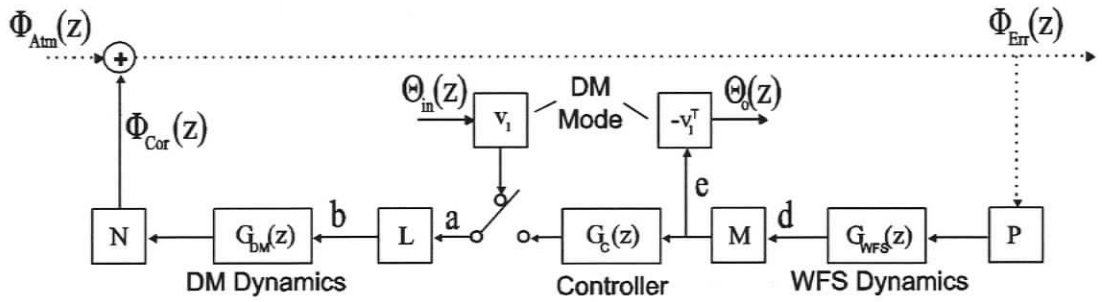

Figure 6: AO system as a SISO system wrt signals $\Theta_{in}(z)$ and $\Theta_{out}(z)$

5. Experimental Results

Based on the conditions of Eq. 30 listed at the end of Section 4, the following approach will be used to obtain values for the system parameters [w, x, y, α_i, β_i];

(i) using a zero-mean white noise signal as input $\theta_{in}[k]$, measure the output, $\theta_{out}[k]$,

(ii) compute $p(j)$ using

$$p(j) = \begin{cases} \dfrac{1}{K-j} \displaystyle\sum_{k=0}^{K-j-1} \theta_{out}[k+j]\theta_{in}[k] & for \quad j \geq 0 \\ 0 & elsewhere \end{cases} \tag{37}$$

(iii) compute $c(j)$ using

$$c(j) = c(-j) = \frac{1}{K-|j|} \sum_{k=0}^{K-|j|-1} \theta_{in}[k+|j|]\theta_{in}[k] \tag{38}$$

(iv) choose the minimum value for w that satisfies $p(w\text{-}1) \approx 0$ while $|p(w)| >> 0$ since the physical meaning of w is the delay of the response.

(v) find a solution for Eq. 30 with minimal $x + y$ so that Eq. 35 is approximately satisfied and

(vi) measure the output and validate the model when using different signals for $\theta_{in}[k]$.

The above approach was used to obtain the system parameters for the TTM and DM separately. Results for the DM were also found for two different sampling rates. The results are discussed in the following sub-sections.

5.1 Tip-tilt mirror parameter identification

A pseudo-random signal with zero mean and variance of 0.04 was generated and used as $\theta_{in}[k]$. This relates to a standard deviation of 0.2 V rms. The length of this signal was 2000 sample points. The sample rate of the system was 66 Hz. The output signal $\theta_{out}[k]$ was measured and the following values were computed for $p(j)$ and $c(j)$:

$$[p(0) \quad p(1) \quad p(2) \quad p(3)] = [8.8\text{e-}5 \quad 0.0212 \quad 0.0180 \quad -2.9\text{e-}4] \tag{39}$$

$$[c(0) \quad c(1)] = [0.0401 \quad -5.6\text{e-}4] \tag{40}$$

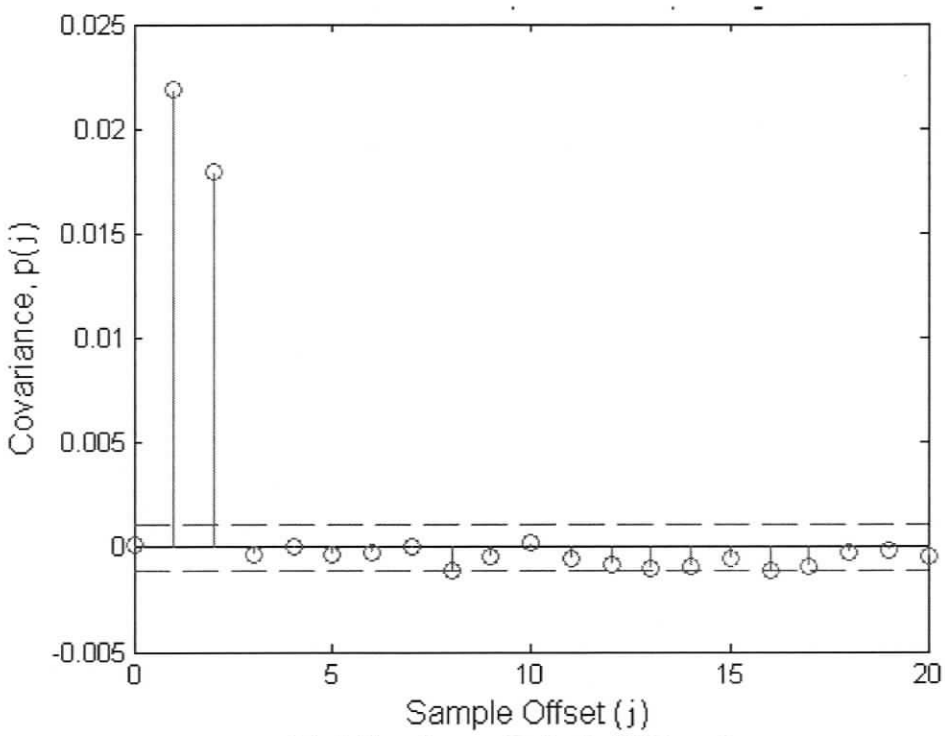

Fig 7: Covariance, $p(j)$, for the TTM x-axis

Since $p(0) \approx 0$ and $p(1) = 0.0212$ is one of the two biggest measured values of $p(j)$, $w = 1$ is a good choice. The measurements of $p(3)$ and higher are negligible, as shown in Fig. 7. Due to the use of white noise where $c(j)$ is ideally 0 for $j \neq 0$, the values of α_i are ideally independent of all the values of β_i when they are solved by Eq. 30. Then, the values of the p vector used to solve for α are $p(x+1)$ through $p(x+y)$. Therefore, if $p(x+1)$ through $p(x+y)$ are insignificant then y should be restricted to 0 since the solution of α_i would otherwise be based entirely on measurement noise. This, as well as condition (i) of Eq. 30, place limits on x and y of $w = 1 \leq x \leq 2$ as well as $y > 0$ only when $x = 1$ in this case. To determine the model parameters α_i, β_i with $w = 1$, the following possible combinations of x and y were considered;

(a) $x = 1$, $y = 0$, $\alpha_i = 0$ for $i > 0$, $\beta_i \neq 0$ for $i = 1$

(b) $x = 1, y = 1, \alpha_i \neq 0$ for $i = 1, \beta_i \neq 0$ for $i = 1$

(c) $x = 1, y = 2, \alpha_i \neq 0$ for $i = 1$ or $2, \beta_i \neq 0$ for $i = 1$

(d) $x = 2, y = 0, \alpha_i = 0$ for $i > 0, \beta_i \neq 0$ for $i = 1$ or

The specific values of these solutions are provided in Table 1. The solutions for (a), (b) and (c) contradict the steady state condition of Eq. 30. The only case that does not contradict the conditions of Eq. 30 is case (d). This leads to the following model

$$G_{Loop}(z) = 0.5359z^{-1} + 0.4559z^{-2} \tag{41}$$

Table 1: Experimentally determined models for TTM with 66 Hz sample rate

Axis	x	y	α_1	α_2	β_1	β_2	Steady State
TTM:X	1	0	0	0	0.5287	0	0.5287
	1	1	0.863	0	0.5287	0	3.8591
	1	2	0.863	-0.7464	0.5287	0	0.5985
	2	0	0	0	0.5351	0.4563	0.9914
TTM:Y	1	1	0.9329	0	0.5474	0	8.1580
	2	0	0	0	0.5483	0.5107	1.0590

* Any coefficient not listed is equal to 0

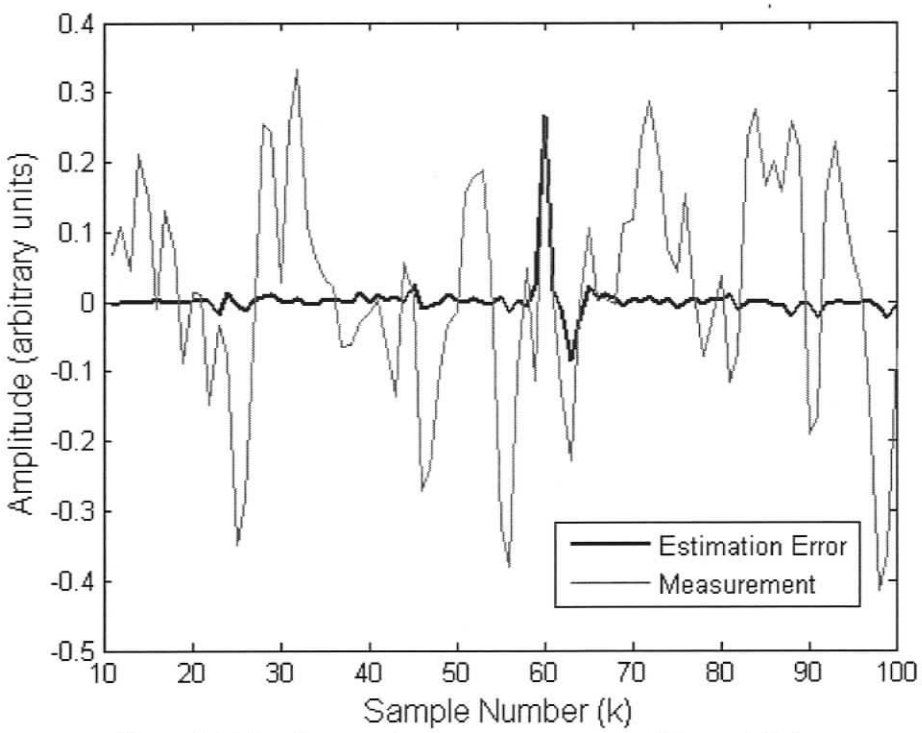

Figure 8: Estimation error between measurement and the modeled response

To validate the model, another pseudo-random input signal was generated and used as the input signal. Then the measured output was compared with the output obtained from the model. Fig. 8 shows an example of the measured output and the error between this signal and the output of the model of $G_{Loop}(z)$. This shows that the model of Eq. 41 nearly matches the actual system dynamics.

An anomaly in the comparison is that occasionally the model was not able to accurately model the actual measurement, as shown from sample number 60 to 65 of Fig. 8. There are two ways that the measurement will change in an AO system; (i) the shape or angle of the mirrors change, and (ii) the shape of the incoming wave front changes. There are two ways to make the angle of the mirrors change; (i) change the voltage delivered to the actuators or (ii) an external vibration. The anomaly appears as though it

is a damped oscillation that only lasts 75 ms, so this is an indication that the cause was a vibration that was sensed by the WFS.

5.2 MEMS deformable mirror parameter identification

The DM was driven as in Eq. 23 where $\theta_{in}[k]$ was a white noise signal of zero mean and 0.25 variance. This variance was chosen to give a high signal-to-noise ratio (SNR) compared to actuator quantization without being large enough to reach the physical limits of the DM stroke. The length of the sequence was 2000 sample points. The sample rate for this experiment was 261 Hz. The values of p(j) were calculated and plotted in Fig. 9 and shown in Eq. 42.

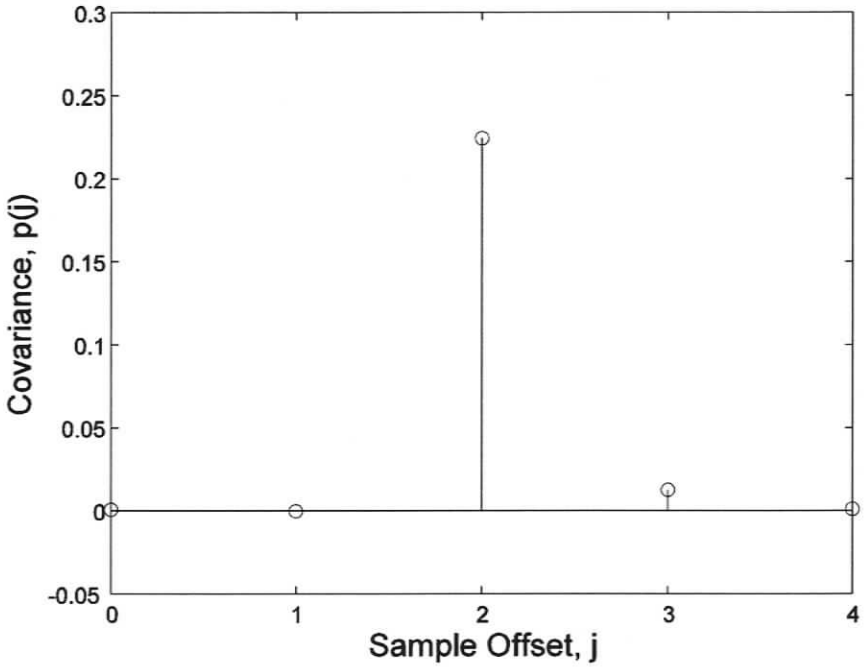

Figure 9: Covariance vs. sample offset for a MEMS DM at 261 Hz sample rate

$$\begin{aligned}&[p(0) \quad p(1) \quad p(2) \quad p(3) \quad p(4)]\\&=[0.0007 \quad -0.0002 \quad 0.2243 \quad 0.0125 \quad 0.0011]\end{aligned} \tag{42}$$

$$[c(0) \quad c(1)]=[0.25 \quad -0.0018] \tag{43}$$

As can be seen from Fig. 9 and Eq. 42, there are only 2 offsets, j, that produce any significant correlation between the input and output signals. The majority of the response of the WFS measurement occurs on the second sample after the DM signal is generated within the computer. Therefore, w is chosen to be 2. Assuming $p(3)$ is significant, x and y are limited to $w = 2 \leq x \leq 3$ and $y > 0$ only when x is 2. This leads to the following cases to be considered;

(a) $x = 2, y = 0, \alpha_i = 0$ for $i > 0, \beta_i \neq 0$ for $i = 2$

(b) $x = 2, y = 1, \alpha_i \neq 0$ for $i = 1, \beta_i \neq 0$ for $i = 2$

(c) $x = 2, y = 2, \alpha_i \neq 0$ for $i = 1$ or $2, \beta_i \neq 0$ for $i = 2$

(d) $x = 3, y = 0, \alpha_i = 0$ for $i > 0, \beta_i \neq 0$ for $i = 2$ or 3

The results of these calculations can be found in Table 2.

Table 2: Experimentally determined models for MEMS DM

Sample Rate	w	x	y	α_1	α_2	β_1	β_2	β_3	Steady State
66 Hz	1	1	0	0	0	0.942	0	0	0.9420
		1	1	0.0169	0	0.942	0	0	0.9582
		2	0	0	0	0.942	0.0159	0	0.9579
261 Hz	2	2	0	0	0	0	0.8972	0	0.8972
		2	1	0.0629	0	0	0.8972	0	0.9574
		2	2	0.0629	0.0014	0	0.8972	0	0.9589
		3	0	0	0	0	0.8972	0.0565	0.9537

* Any coefficient not listed is equal to 0

Unlike the TTM dynamics, all of the possibilities approximately satisfy the steady state condition of Eq. 30. The models of cases (b), (c) and (d) essentially provide the same response. These models only differ on scales that are less than or equal to the 8-bit DM quantization steps. The model of case (a) is the same as the first term of the other cases. Therefore, it needs to be determined if using additional terms are necessary. This is dependant on the desired accuracy of the application and the coarseness of the

quantization involved. Case (a) can provide a lower order model at the cost of approximately 5% error compared to using any of the other models. It has been found through closed loop control experiments that

$$G_{Loop}(z) = z^{-2} \tag{44}$$

is a sufficient model of this MEMS DM AO test bench system [11]. This is a scaled version of case (a) in order to maintain the constraint that the steady state response is equal to 1. The models for cases (b) and (d) would also be good choices. Case (c) does not provide a significant increase in accuracy so the added complexity is not beneficial.

Non-linear cross coupling causes a change of magnitude of the given mode and an excitation of additional modes. The effect is present in this mirror but it is assumed to be a small effect for the chosen mode. This is because the steady state values shown in Table 2 are approximately 95% of the theoretically expected value of 1. This leaves only 5% error caused by the combined effects of actuator quantization and this non-linear cross coupling.

This process was repeated for 66 Hz. The results are shown in Table 2. As can be seen from Table 2, w has been reduced to 1. This was caused by the specific architecture of the lower sample rate implementation. The 66 Hz implementation operates the WFS and Controller in series. This implies that a signal travels through the entire loop in one sample, hence there is only one delay. The 261 Hz implementation operates these two processes in parallel. Since there are two processes in parallel where the output of the WFS is the input of the controller and vice versa, two samples are required for a signal to travel through the entire system. The result is a measurement of two delays.

With $G_{Loop}(z)$ determined, the transfer function of the closed loop system can be found for any given controller, $G_c(z)$. Using the system layout of Fig. 6 (with the switch moved to the output of $G_c(z)$) and the definition of MPNL, the transfer function of the controlled *n-m* modes of the AO system is defined as

$$\frac{\Phi_{Err}(z)}{\Phi_{Atm}(z)} = \frac{1}{1+G_{Loop}(z)G_C(z)} \tag{45}$$

where $G_c(z)$ is the transfer function of the developed controller, which is outside the scope of this paper. However, the inclusion of the $G_c(z)$ block in Fig. 6 is so that the computer still performs all the necessary computations for the controller. This way, the processing time of the computer becomes part of the $G_{Loop}(z)$ model, even though the result of $G_c(z)$ is not used. This is optional, but may be important if the final controller performs a significant amount of computations.

6. Conclusions

An AO system can be inherently utilized as a reliable measurement tool for determining the response of its discrete components. The experimental approach described here provides accurate AO system models, in the discrete domain, without additional measurement equipment. Two general form equations are derived that allow for the dynamics of different devices and for fitting various models to the data. Adequate models would have a steady-state response of nearly 1 to a unit step input and a stable transfer function. The better model is a compromise between accuracy and complexity. An acceptable model can be normalized by its steady state value so that the final model has a steady state of 1, as described in Sections 2 and 3. The responses of the particular AO components tested are fast enough, compared to the sample rate, that low order models of

the response are sufficient. However, this may change for increased sample rates or for AO components that have different dynamic responses.

7. References

1. D.P. Looze, "Realization of Systems with CCD-based measurements" Automatica **41**, 2005-2009 (2005).
2. A. Wirth, J. Navetta, D. Looze, S. Hippler, A. Glindemann, D. Hamilton, "Real-time modal control implementation for adaptive optics", Applied Optics **37**, 4586-4596 (1998).
3. M.A. van Dam, D. Le Mignant, B.A. Macintosh, "Performance of the Keck Observatory adaptive-optics system", Applied Optics **43**, pp. 5458-5467 (2004)
4. J.J. Widiker, S.R. Harris, B.D. Duncan, "High-speed Shack-Hartmann wavefront sensor design with commercial off-the-shelf optics", Applied Optics **45**, pp. 383-395 (2006)
5. M. Kasper, E. Fedrigo, D. P. Looze, H. Bonnet, L. Ivanescu, and S. Oberti, "Fast calibration of high-order adaptive optics systems," J. Opt. Soc. Am. A **21**, 1004-1008 (2004)
6. R. Conan, C. Bradley, P. Hampton, O. Keskin, A. Hilton, and C. Blain, "Distributed modal command for a two-deformable-mirror adaptive optics system," Appl. Opt. **46**, 4329-4340 (2007)
7. D.S. Watkins, *Fundamentals of Matrix Computations*, John Wiley & Sons, Inc., New York, USA, 2002
8. B.P. Wallace, P.J. Hampton, C.H. Bradley, R. Conan, "Evaluation of a MEMS deformable mirror for an adaptive optics test bench" Optics Express **14** 10132-10138 (2006)
9. S. Haykin, *Adaptive Filter Theory*, Fourth Edition, Prentice-Hall, New Jersey, USA, 2002
10. K. Ogata, *Discrete-Time Control Systems, Second Edition*, Prentice-Hall, Inc., New Jersey, USA, 1995
11. P.J. Hampton, R. Conan, C. Bradley, P. Agathoklis "Control of a woofer-tweeter system of deformable mirrors", in *Proceedings of Advanced Software and Control for Astronomy*, Hilton Lewis, Alan Bridger, ed. (SPIE 6274, 2006), 1Z

Appendix B: Control of a Woofer Tweeter System of Deformable Mirrors

The material in the following manuscript was originally published as

Control of a woofer tweeter system of deformable mirrors
Hampton, P. J., R. Conan, C. Bradley, and P. Agathoklis
Edited by H. Lewis and A. Bridger.
Proceedings of SPIE 6274, 2006, 62741Z

Control of a Woofer Tweeter System of Deformable Mirrors

Peter J. Hampton*[ac], Rodolphe Conan[bc], Colin Bradley[bc], Pan Agathoklis[ac]

[a] Department of Electrical and Computer Engineering

[b] Department of Mechanical Engineering

[c] University of Victoria, British Columbia, Canada

ABSTRACT

This paper describes the control of two deformable mirrors (DM) and a tip tilt mirror for adaptive optics. The purpose of this experimental adaptive optics system at the University of Victoria is to prove the Woofer Tweeter concept for use in instruments for the Thirty Meter Telescope (TMT) which is currently under development. The first deformable mirror is a large stroke DM (Woofer) capable of lower frequency correction in both the temporal and spatial domains. The other DM (Tweeter) is capable of the high temporal and spatial frequency corrections of the turbulence. The response speed of the Woofer is incorporated into the Tweeter controller in order to allow for appropriate offloading from the Tweeter to the Woofer. In order to determine which Tweeter shapes must compensate for the slower Woofer and which are not coupled to the Woofer, the cross correlation of the devices is determined. The method of converting the wave front sensor (WFS) measurements to control signal error is given. The transfer functions of the controller are provided, along with rejection ratio plots, bandwidths and amplitude response to system noise.

Key words: Adaptive Optics, Control, Digital, Discrete, Woofer, Tweeter, TMT

INTRODUCTION

A simple adaptive optics system for astronomy uses a single wave front sensor (WFS) and a single deformable mirror (DM) to correct for the distortions imposed on light by the atmosphere and the static aberrations of the telescope optics[1]. In the next generation telescopes, both the actuator density and maximum actuator stroke requirements increase significantly due to the enormity of these very large telescopes. Current technology is cost prohibitive to design a single mirror that satisfies both of these requirements. Fortunately, the large stroke required is for the compensation of low spatial frequency distortion because the majority of the turbulence power is for the low spatial frequency distortions[2]. This allows the system to be designed with two DMs; (i) a high stroke, low actuator density DM named the Woofer and (ii) a low stroke, high actuator density DM named the Tweeter. The Woofer for this test bench was produced by the Laboratory of Astrophysics of the Observatory of Grenoble (LAOG) and has 52 actuators on an 8x8 grid. The Tweeter was produced by Boston Micromachines Corporation and has 140 actuators on a 12x12 grid.

The Adaptive Optics Laboratory at the University of Victoria has recently produced a test bench for this Woofer-Tweeter system. This project is part of the development of the Thirty Meter Telescope (TMT) that will be built in the next decade. It was feared that a potential problem for this approach is that the Woofer and Tweeter could form equal but opposite shapes. These distortions on the light beam would cancel with each other before reaching the WFS and be both invisible and a waste of the available stroke of the mirrors. This misbehavior has been avoided in the control scheme

presented later in this paper by determining the correlated shapes and an offloading scheme from the Tweeter to Woofer has been developed.

Initial simulated and experimental results have shown that the developed controller can appropriately split the correction between the mirrors and acts similarly to the single DM case[3]. The Woofer corrects for the low-spatial-low-temporal frequency disturbances and the Tweeter corrects for the remaining disturbance. It has been assumed that the Woofer can respond slower than the tweeter. The Woofer's impulse response is modeled as an exponential decay given by $e^{kT_s/\tau}$. A one dimensional representation of the controller approach is shown in Figure 1. The Woofer slowly approaches the steady state of the input signal. During this time, the Tweeter compensates for the residual error. The combined response of the two DMs is then equal to how the single DM case would respond to the input. A simple layout of the single DM system components is shown in Figure 4. This paper focuses primarily on discrete control in the Z-domain[4].

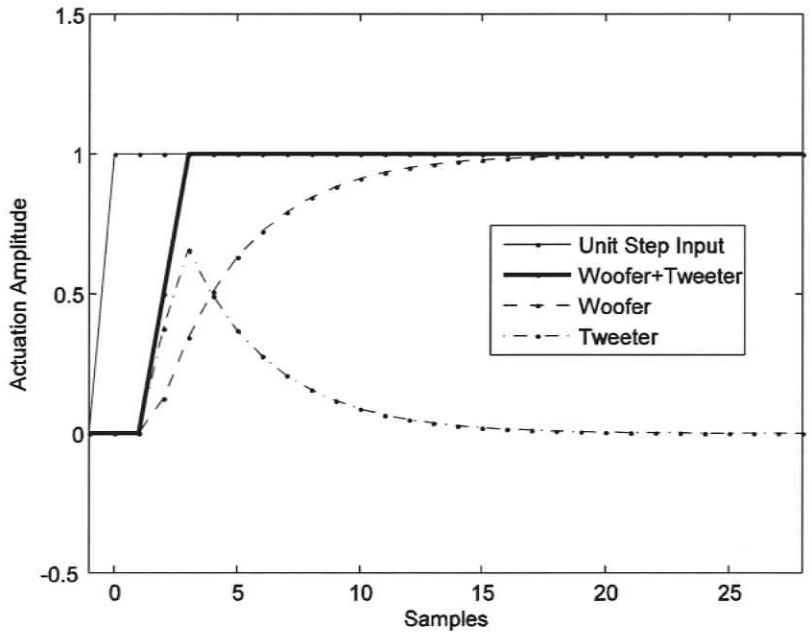

Fig. 1. One Dimensional Off-Loading Example

WOOFER-TWEETER MODE SELECTION

A straight forward approach to controlling this Woofer-Tweeter system would be to simply do a global interaction matrix and control the entire system as a single device as in Eq. 2.1. Although conceptually simple and shown to work in simulations, this method proved difficult, if not impossible to use in practice. It is believed that a strong reason for this is that the ratio of actuator motion to control signal for the Woofer is on the order of 1000 times that of the Tweeter. So any measurement noise on the Woofer portion of the interaction matrix would be 1000 times larger than the noise on the Tweeter portion. This could make the Woofer noise significant compared with the Tweeter portion of the global interaction matrix.

$$D_{Global} = \begin{bmatrix} D_{tt} & D_{wr} & D_{tr} \end{bmatrix} \tag{2.1}$$

Where D will represent the interaction matrix and the subscripts tt, wr, tr refer to the Tip Tilt, Woofer and Tweeter mirrors respectively. D^{+} represents the pseudo inverse of the interaction matrix (I.e. the reconstructor matrix).

An alternative approach was developed that treats the mirrors as separate devices. The applicable shapes from one mirror can be offloaded to another in order to reduce the stroke requirements of smaller devices. In order to offload shapes from one mirror to another, the shapes that are correlated between the mirrors must be determined. The first step is to find the separate interaction matrix for each individual mirror (D_{tt}, D_{wr} and D_{tr}). Zonal interaction matrices were chosen for this experiment.

2.1 Actuator Linearization

Before shapes can accurately be transferred between mirrors, the mirrors need to be linear, otherwise simple addition and subtraction does not apply. Fortunately the response of the Woofer actuators was found to already be linear in the necessary range of motion. However, the MEMS Tweeter has a quadratic response. This was measured by moving the given actuator to 8 different positions and measuring the respective actuator error using the initial Tweeter reconstructor matrix (D_{tr}^{+}). Then a quadratic function is fit to these 8 data points to determine the coefficients (a, b and c) in a least squares sense using a Vandermonde matrix.

$$f(x) = ax^2 + bx + c \tag{2.2}$$

The values of $f(x)$ are somewhat arbitrary because the reconstructor used was made using a non-linear Tweeter. They are also relative to the position that happens to cause 0 error. Once the coefficients are found, c is set to 0 so that when $f(x) = 0$, x will also equal 0. The output to the Tweeter is then found by Eq. 2.3. The sign of the square route has the same sign as b. Once this linearization is complete, D_{tr} must be re-evaluated using the linear actuator response.

$$x = \frac{-b \pm \sqrt{b^2 + 4ay}}{2a} \tag{2.3}$$

Where y is the desired actuation of the tweeter actuator and x is the output to the tweeter.

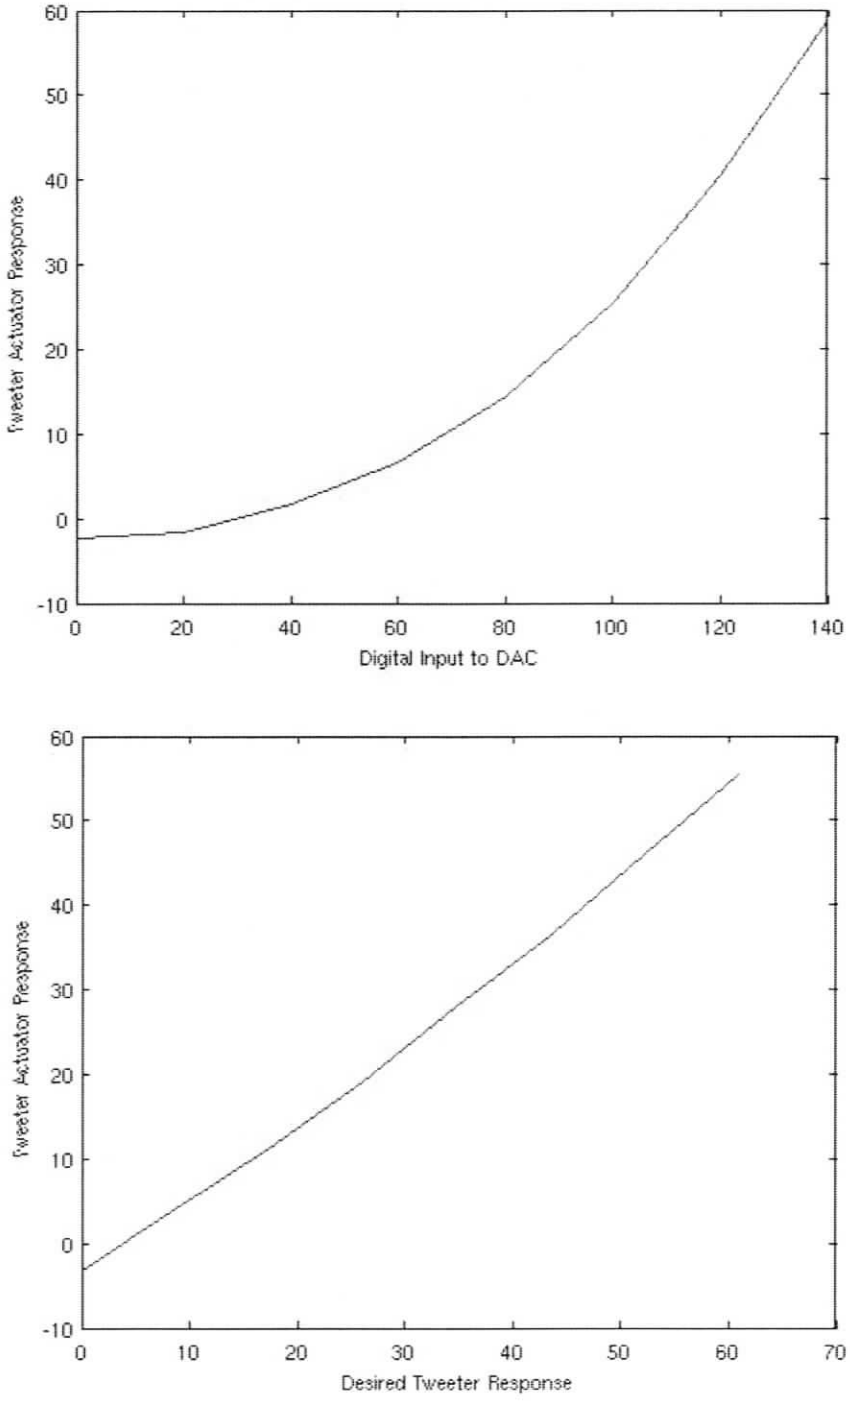

Figure 2: Quadratic actuator response of the Tweeter (top) and Linear response when applying Eq. 2.3 (bottom)

2.2 Tip Tilt Correlation

It is an eventual goal of this project to offload the average tilt to a tilt mirror or a tilt mount that can tilt a DM. This would be offloading at a very slow rate and would probably not occur each sample. In the mean time, there is full offloading to the tip tilt mirror and the Woofer and Tweeter are not to perform any tilt correction. This is accomplished by first determining what shapes are correlated between each DM and the Tip Tilt mirror by taking advantage of Singular Value Decomposition (SVD)[5]. Since D_{tt} is two vectors, there are only 2 non-zero singular values and the first 2 column vectors of $U_{<tr,tt>}$ will each be a set of Tweeter actuator commands that is the best fit to one of the two tilts. The same is repeated for the Woofer.

$$U_{<tr,tt>}W_{<tr,tt>}V^T_{<tr,tt>} = D^+_{tr}D_{tt} \tag{2.4}$$

$$B_{<tr,tt>} = [\vec{u}_{<tr,tt>,1} \quad \vec{u}_{<tr,tt>,2}] \tag{2.5}$$

$$U_{<wr,tt>}W_{<wr,tt>}V^T_{<wr,tt>} = D^+_{wr}D_{tt} \tag{2.6}$$

$$B_{<wr,tt>} = [\vec{u}_{<wr,tt>,1} \quad \vec{u}_{<wr,tt>,2}] \tag{2.7}$$

where $<,>$ denotes correlation of two mirrors, B is a subset of orthogonal mirror shapes and $\vec{u}_k$ is the k^{th} column vector of U

2.3 Woofer-Tweeter Correlation

The correlation between the Woofer and the Tweeter is conducted in a similar way. The difference is that the correlated shapes with the Tip Tilt mirror must be suppressed. Again, this can be achieved by taking advantage of the SVD.

$$U_{<tr,wr>}W_{<tr,wr>}V^T_{<tr,wr>} = \left(I - B_{<tr,tt>}B^T_{<tr,tt>}\right)D^+_{tr}D_{wr}\left(I - B_{<wr,tt>}B^T_{<wr,tt>}\right) \tag{2.8}$$

where $B_{<tr,wr>} = [\vec{u}_{<tr,wr>,1} \quad \vec{u}_{<tr,wr>,2} \quad \cdots \quad \vec{u}_{<tr,wr>,N}]$

$B_{<wr,tr>} = [\vec{v}_{<tr,wr>,1} \quad \vec{v}_{<tr,wr>,2} \quad \cdots \quad \vec{v}_{<tr,wr>,N}]$

$\left(I - B_{<wr,tt>} B^T_{<wr,tt>}\right) B_{<wr,tt>} = null$

$B^T_{<tr,tt>} \left(I - B_{<tr,tt>} B^T_{<tr,tt>}\right) = null$

N is the number of chosen correlated Woofer Tweeter modes

The matrix $V_{<tr,wr>}$ is a matrix of Woofer shapes that match the Tweeter shapes in matrix $U_{<tr,wr>}$. The match is made in WFS space and so is only dependant on the calibration data and does not require additional modeling. They are ordered left to right from highest correlation to least. All singular values beyond the N number cut off should be forced to 0. The shapes that are correlated with the Tip Tilt mirror result in a null matrix, so these are assigned the least amount of correlation with the other mirror and these shapes are automatically moved to the end of the set. The result is shown in the following figure. The first 12 shapes that have high correlation between the Woofer and the Tweeter are shown in order from left to right and top to bottom. This figure was generated by applying the modes to a virtual Woofer mirror in MatLab and generating a plot of the phase. The modes are very similar to the Zernike set of modes[2] starting with focus, then astigmatism followed by coma and trefoil and finally some 4th order shapes. The modes are off center and this may be due to a shifted alignment between the Woofer and the Tweeter.

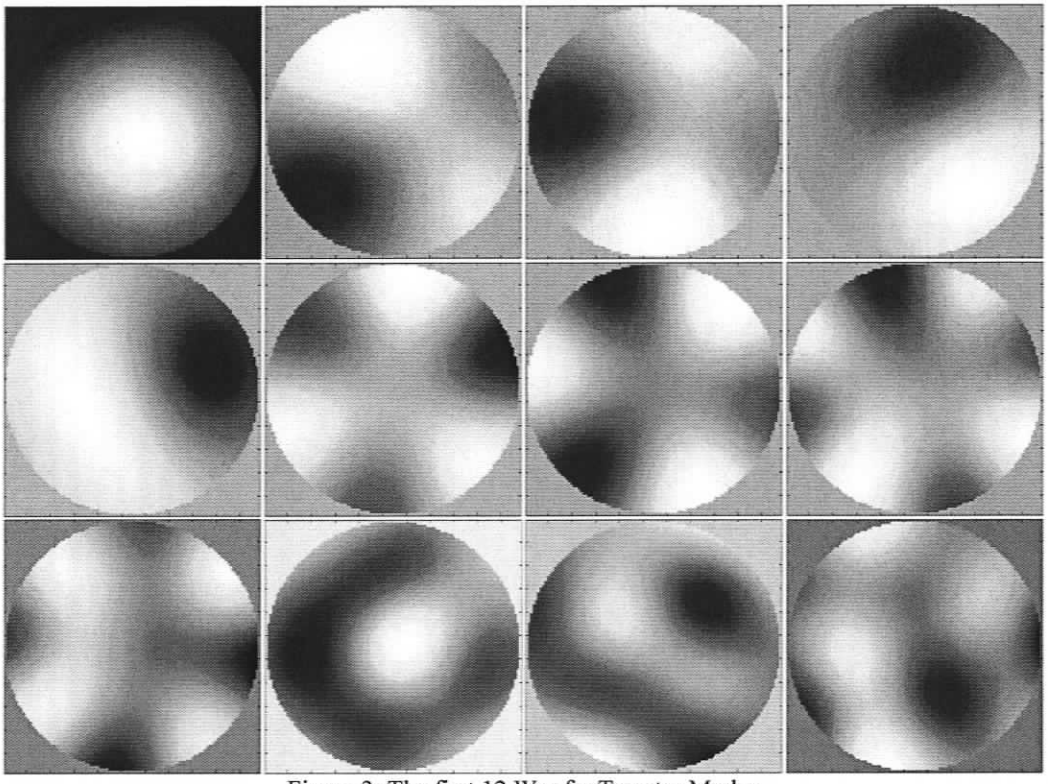

Figure 3: The first 12 Woofer Tweeter Modes

2.4 Reconstructors

The reconstruction of WFS measurements to mirror actuator commands will occur in stages. This allows the WFS error measurement to propagate to the mirrors without the ratio of real actuator motion to control signal being a factor as in the Global Interaction case. The first stage is to convert the measured WFS error to Tweeter actuator error with D_{tr}^{+}. Then Tweeter actuator error is converted to Tip Tilt error using Eq. 2.10 and Woofer error using Eq. 2.11. The controllers that handle these signals are derived in section 3.

$$E_{tr} = D_{tr}^{+} E_{WFS} \tag{2.9}$$

$$E_{tt} = V_{<tr,tt>} W_{<tr,tt>}^{+} U_{<tr,tt>}^{T} E_{tr} \tag{2.10}$$

$$E_{wr} = V_{<tr,wr>} W_{<tr,wr>}^{+} U_{<tr,wr>}^{T} E_{tr} \tag{2.11}$$

CONTROL SYSTEM

3.1 Model of an adaptive optics system

The model of an Adaptive Optics system has been determined by several independent researchers in either the continuous or discrete domains. The common continuous model is one where the WFS is modeled as a zero order hold divided by the sample time. The Laplace transform of a zero order hold is given in Eq. 3.1 and that of the WFS in Eq. 3.2. A zero order hold is also necessary to model the conversion of the sampled data into a pulse with the sample time as its duration. The updating of the DM is modeled as a delay, which can be a fraction of the sample time (Eq. 3.3).

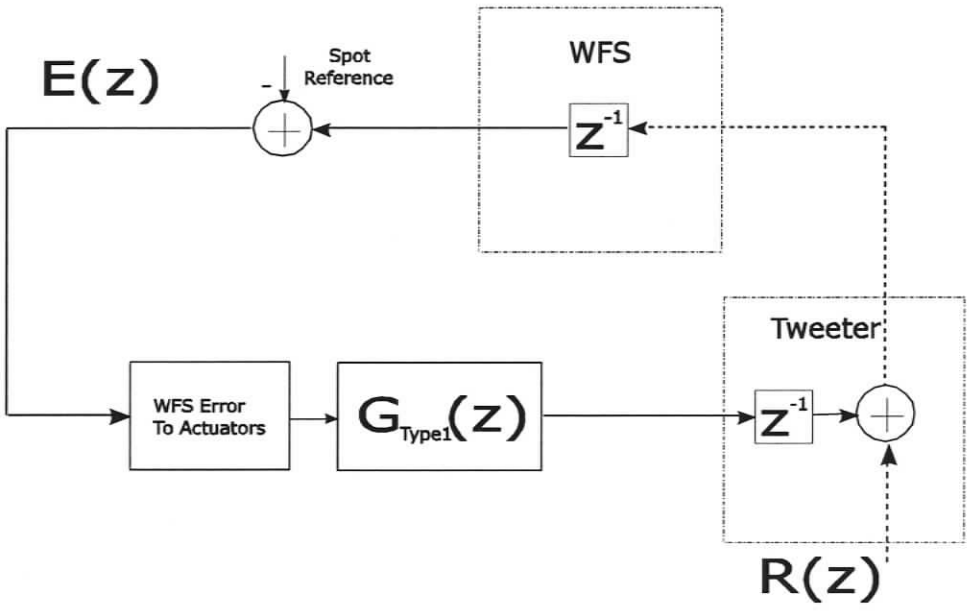

Figure 4: Discrete model of a single mirror AO system

$$G_{ZOH}(s) = \frac{1 - e^{sT_s}}{s} \tag{3.1}$$

$$G_{WFS}(s) = \frac{1 - e^{sT_s}}{sT_s} \tag{3.2}$$

$$G_{DM}(s) = e^{s\varphi} \tag{3.3}$$

Where T_s is the sample time and φ is the delay

The system that is controlled (i.e. the plant) is given as $G_P(s)$ in Eq. 3.5 for the continuous domain and as $G_P(z)$ in Eq. 3.8 for the discrete time model.

$$G_P(s) = G_{DM}(s)G_{WFS}(s)G_{ZOH}(s) \tag{3.4}$$

$$G_P(s) = e^{s\varphi}\left(\frac{1 - e^{sT_s}}{sT_s}\right)\left(\frac{1 - e^{sT_s}}{s}\right) \tag{3.5}$$

The experimental setup built at the University of Victoria is representing this model under the assumption that the DM update delay is considered to be one sample period. When φ is an integer multiple of the sample time, the transformation from the continuous domain to the discrete domain is simplified. The discrete model for the system tested at the University of Victoria is simply two sample delays, Eq. 3.8.

$$G_P(z) = Z\left\{e^{sT_s}\left(\frac{1 - e^{sT_s}}{sT_s}\right)\left(\frac{1 - e^{sT_s}}{s}\right)\right\} \tag{3.6}$$

$$G_P(z) = z^{-1}\left(1 - z^{-1}\right)^2 \frac{T_s z^{-1}}{T_s\left(1 - z^{-1}\right)^2} \tag{3.7}$$

$$G_P(z) = z^{-2} = G_{DM}(z)G_{WFS}(z) \tag{3.8}$$

where

$$G_{DM}(z) = z^{-1} \tag{3.9}$$

$$G_{WFS}(z) = z^{-1} \tag{3.10}$$

3.2 Single deformable mirror control

The simplest approach to controlling such a system is to use a single integrator and adjust the gain in order to make the system stable. This approach will be referred to as the Classic Adaptive Optics (CAO) controller.

$$G_{CAO}(z) = \frac{g}{1 - z^{-1}} \tag{3.11}$$

where g is a variable gain.

The closed loop transfer function of such a system is

$$T_{CAO}(z) = \frac{G_{CAO}(z)G_P(z)}{1 + G_{CAO}(z)G_P(z)} \tag{3.12}$$

$$T_{CAO}(z) = \frac{gz^{-2}}{1 - z^{-1} + gz^{-2}} \tag{3.13}$$

where $T_{CAO}(z)$ is stable for $0 \leq g < 1$

However, the control approach used here is to increase the bandwidth by increasing the complexity of the controller by 1 zero and 2 poles. The open loop transfer function of this controller is

$$G_{Type1}(z) = \frac{g\left(0.5 + 0.5z^{-1}\right)}{1 - 0.5z^{-2} - 0.5z^{-3}} \tag{3.14}$$

The closed loop transfer function of the system with this Type 1 controller is

$$T_{Type1}(z) = \frac{g\left(0.5z^{-2} + 0.5z^{-3}\right)}{1 + (g-1)\left(0.5z^{-2} + 0.5z^{-3}\right)} \tag{3.15}$$

Where $T_{Type1}(z)$ is stable for $0 \leq g < \sqrt{5}$

The closed-loop system for $g = 1$ using $G_{Type1}(z)$ has all poles at $z = 0$ and the system responds as the averaging filter of Eq. 3.16. The step response will settle in two

samples after the initial delay. It has been found using simulated and experimental results the CAO controller with g = 0.4 gives a similar performance as the Type 1 controller with $g = 1$.

$$T_{Type1}(z)\Big|_{g=1} = 0.5z^{-2} + 0.5z^{-3} \tag{3.16}$$

3.3 Slow Woofer, Fast Tweeter

It is assumed that the Tweeter responds the same as in the single DM case. One possibility is that the Woofer mirror will have a slower response time than of the Tweeter. If this effect is not considered for the controller design, the controller would over drive the Woofer and unexpected oscillations may occur.

$$G_{Tweeter}(z) = G_{DM}(z) = z^{-1} \tag{3.17}$$

$$G_{Woofer}(z) = \frac{\left(1 - e^{-T_s/\tau}\right)z^{-1}}{1 - e^{-T_s/\tau}z^{-1}} \tag{3.18}$$

Where τ is the time constant of the exponential decay

To avoid such a deterioration of performance, a first order response model is assumed for the Woofer. A filter is then developed for the Tweeter that will allow the Woofer-Tweeter combination to have the same response as a single fast DM. The subscript *HP* denotes that it is a High Pass filter.

$$G_{DM}(z) = G_{HP}(z)G_{Tweeter}(z) + G_{Woofer}(z) \tag{3.19}$$

$$G_{HP}(z) = \frac{G_{DM}(z) - G_{Woofer}(z)}{G_{Tweeter}(z)} \tag{3.20}$$

$$G_{HP}(z) = \frac{e^{-T_s/\tau}\left(1 - z^{-1}\right)}{1 - e^{-T_s/\tau}z^{-1}} \tag{3.21}$$

This high pass filter is applied to all mirror shapes that both the Woofer and Tweeter can produce. It is not applied to Tweeter shapes that the Woofer can not produce. This allows the Tweeter to temporarily correct for a new disturbance until the Woofer has enough time to produce the full correction of the given shape. This high pass filter has a zero at z = 1, which will cancel with the integrator in the controller. If this model is a good approximation of the real system, it is expected that the mirrors will seamlessly transfer low temporal frequency disturbance from the Tweeter to the Woofer.

3.4 Rejection Ratio

The rejection ratio is the amplitude response of the following error. It is often used in Adaptive Optics to assess the bandwidth of the system. This can be determined using

$$\text{Rejection Ratio} = \left|\frac{E(z)}{R(z)}\right| = \left|1 - T(z)\right| \tag{3.22}$$

Where E(z) is the measured error, R(z) is the optical disturbance input and T(z) is the closed loop transfer function.

This is usually plotted with a logarithmic scale for the frequency and decibel scale for the amplitude. The frequency of the first point of the plot that crosses the 0dB line (starting at 0 Hz and scanning towards the sample rate) is called the 0dB bandwidth.

3.5 Tweeter Controller

Even though the Woofer is assumed to respond slower than the Tweeter, the Woofer will be controlled by the same controller as in the single DM case. This is expected to work if the Tweeter corrects the error that remains because of the slow response of the Woofer.

For the shapes that overlap between the Woofer and Tweeter mirrors, the Tweeter will be controlled by the controller, $G_{TL}(z)$, given in Eq. 3.25. The integrator of $G_{Type1}(z)$ is cancelled by the zero at z = 1. The subscript *TL* denotes the Tweeter Low spatial frequency controller.

$$G_{TL}(z) = G_{Type1}(z)G_{HP}(z) \tag{3.23}$$

$$G_{TL}(z) = \frac{g\left(0.5+0.5z^{-1}\right)}{1-0.5z^{-2}-0.5z^{-3}} \frac{e^{-T_s/\tau}\left(1-z^{-1}\right)}{1-e^{-T_s/\tau}z^{-1}} \tag{3.24}$$

$$G_{TL}(z) = \frac{ge^{-T_s/\tau}\left(0.5+0.5z^{-1}\right)}{1+\left(1-e^{-T_s/\tau}\right)z^{-1}+\left(0.5-e^{-T_s/\tau}\right)z^{-2}-0.5e^{-T_s/\tau}z^{-3}} \tag{3.25}$$

The high spatial frequency shapes cannot be controlled by the Woofer. Therefore, for those frequencies, all the corrections must be provided by the Tweeter as though it was the only mirror in the system. So, the controller is the same as in the single mirror case.

$$G_{TH}(z) = G_{Type1}(z) \tag{3.26}$$

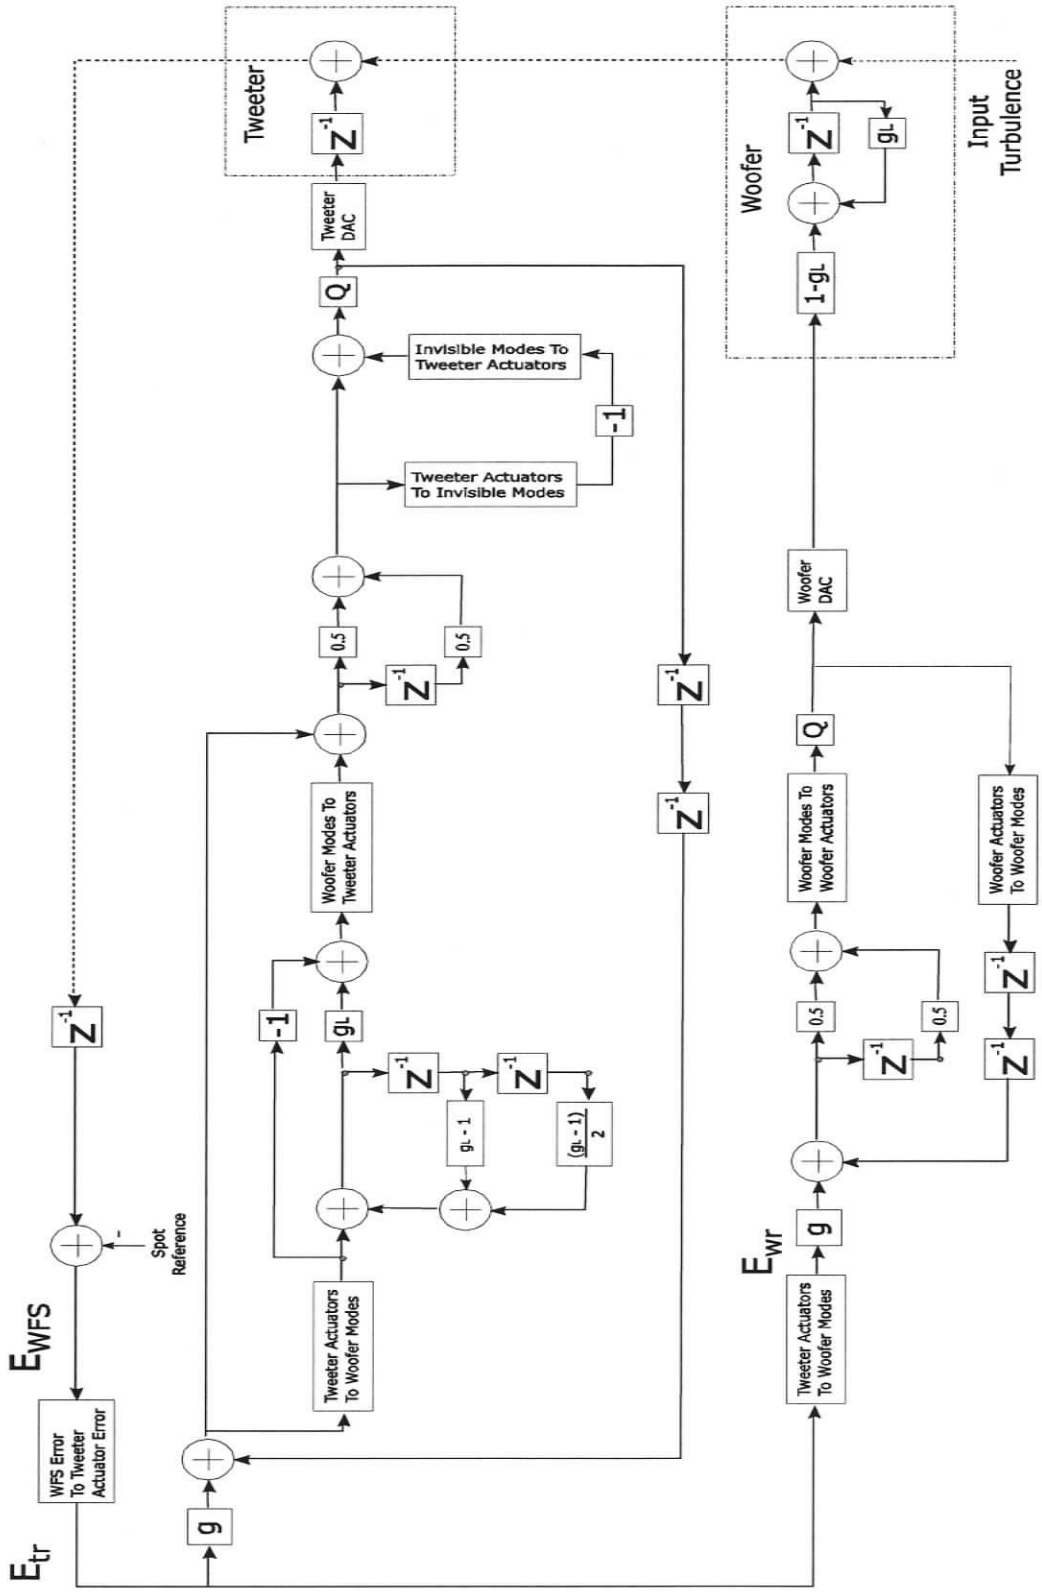

Figure 5: Woofer tweeter control system. Tip Tilt control omitted. $g_L = e^{-T_s/\tau}$

3.6 Mathematical complexity

In a fully implemented system, the largest cost for computation time in the controller will be spent on the Tweeter. It is expected that the Tweeter would have at least 1000 actuators and the Woofer would have less than 100. The controller provided in Figure 5 is O(n).

$$flops = n(11+4q+2m)+m(10+4p) \tag{3.27}$$

Where n is the number of Tweeter actuators, m is the number of Woofer modes, p is the number of Woofer actuators, q is the number of Tweeter invisible modes.

This formulation of the number of flops is based on Figure 5 but also includes the conversion of Tweeter actuator error to Tip Tilt actuator error. It does not include the linearization of the Tweeter actuators nor the quantization of any signal. It also does not include the computations to convert the WFS measurement to Tweeter actuator data because it will be done via a computationally optimal algorithm separate from the controller. In the case that the Woofer remains the same, the Tweeter is replaced with a 32x32 mirror, the Tweeter invisible modes remains at 12 (including tilts) and the Woofer modes are increased from 20 to 50, then it would be expected that this controller would use 173,716 flops. In comparison, the direct vector matrix multiplication to get the Tweeter actuator error from the WFS measurement would cost $4n^2$ flops, or over four million flops. So this controller is comparatively light on computations at less than 5% of an O(n^2) operation.

EXPERIMENTAL RESULTS

4.1 White noise input

The disturbance used in the initial experiments is a known but random set of actuator commands for all mirrors. This ensures that the disturbance is correctable since it is on the same scale as the mirrors. The data collected is plotted as rejection ratio plots. White noise input is desirable for constructing rejection ratio plots because it ensures that there is significant power at all frequencies of the spectrum. The expected rejection ratio when the system is operating well is given as Eq. 4.1. The result shown in Figure 6 shows that when the Woofer is as fast as the Tweeter, the measured rejection ratio matches the single fast DM model. The 0dB bandwidth of this system is approximately 6.5% of the sample rate for the given gain of 1. The error signal is E(z) and the input signal is R(z).

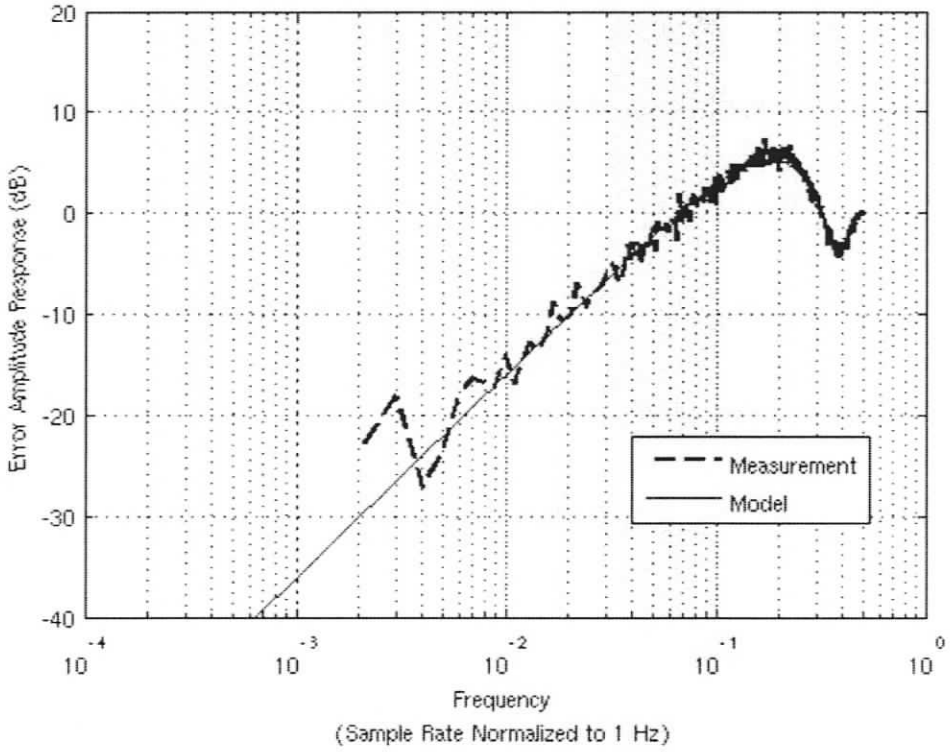

Figure 6: Rejection ratio of a Woofer mode for $\tau = 0$

$$\left|\frac{E(z)}{R(z)}\right| = \left|\frac{1}{1+\left(G_{TL}(z)G_{Tweeter}(z)+G_{Type1}(z)G_{Woofer}(z)\right)G_{WFS}(z)}\right| = \left|1-0.5z^{-2}-0.5z^{-3}\right| \quad (4.1)$$

The Woofer was then slowed by filtering the control signals through a first order system. The exponential decay constant, τ, was increased to $20T_s$. This means that the Woofer would settle to within 5% of the steady state output signal in 60 samples. In the mean time, the Tweeter assumes the responsibility for correcting the high frequency component of the Woofer shapes. The result is that the rejection ratio plots for all measured speeds of the Woofer remained consistent with the fast single mirror model. An example of this is given as Figure 7 (left). This means that the combination of the Tweeter controller and Woofer controller are providing the correct signals to their respective mirrors to achieve the desired correction.

The case when the slow Woofer operates without the aid of the Tweeter was also considered. It is expected that if the Tweeter does not correct for the high frequency disturbance that the rejection ratio would be given by Eq. 4.2. In this case, $\tau = 20T_s$. This model is fully supported by the measurements taken and displayed in Figure 7 (right). It is shown to illustrate the improvement provided by the Tweeter.

$$\left|\frac{E(z)}{R(z)}\right| = \left|\frac{1}{1+G_{Type1}(z)G_{Woofer}(z)G_{WFS}(z)}\right| = \left|\frac{\left(1-0.5z^{-2}-0.5z^{-3}\right)\left(1-e^{-T_s/\tau}z^{-1}\right)}{1-e^{-T_s/\tau}z^{-1}-0.5e^{-T_s/\tau}\left(z^{-2}-z^{-4}\right)}\right| \quad (4.2)$$

This controller is set up so that the closed loop response is the same as an averaging filter. This is shown in Figure 8 with an amplitude response plot of the Woofer output signal ,C(z), compared with the input signal, R(z). The measured data follows the model of an averaging filter very closely.

$$\left|\frac{C(z)}{R(z)}\right| = \left|0.5z^{-2} + 0.5z^{-3}\right| = \left|0.5 + 0.5z^{-1}\right| \tag{4.3}$$

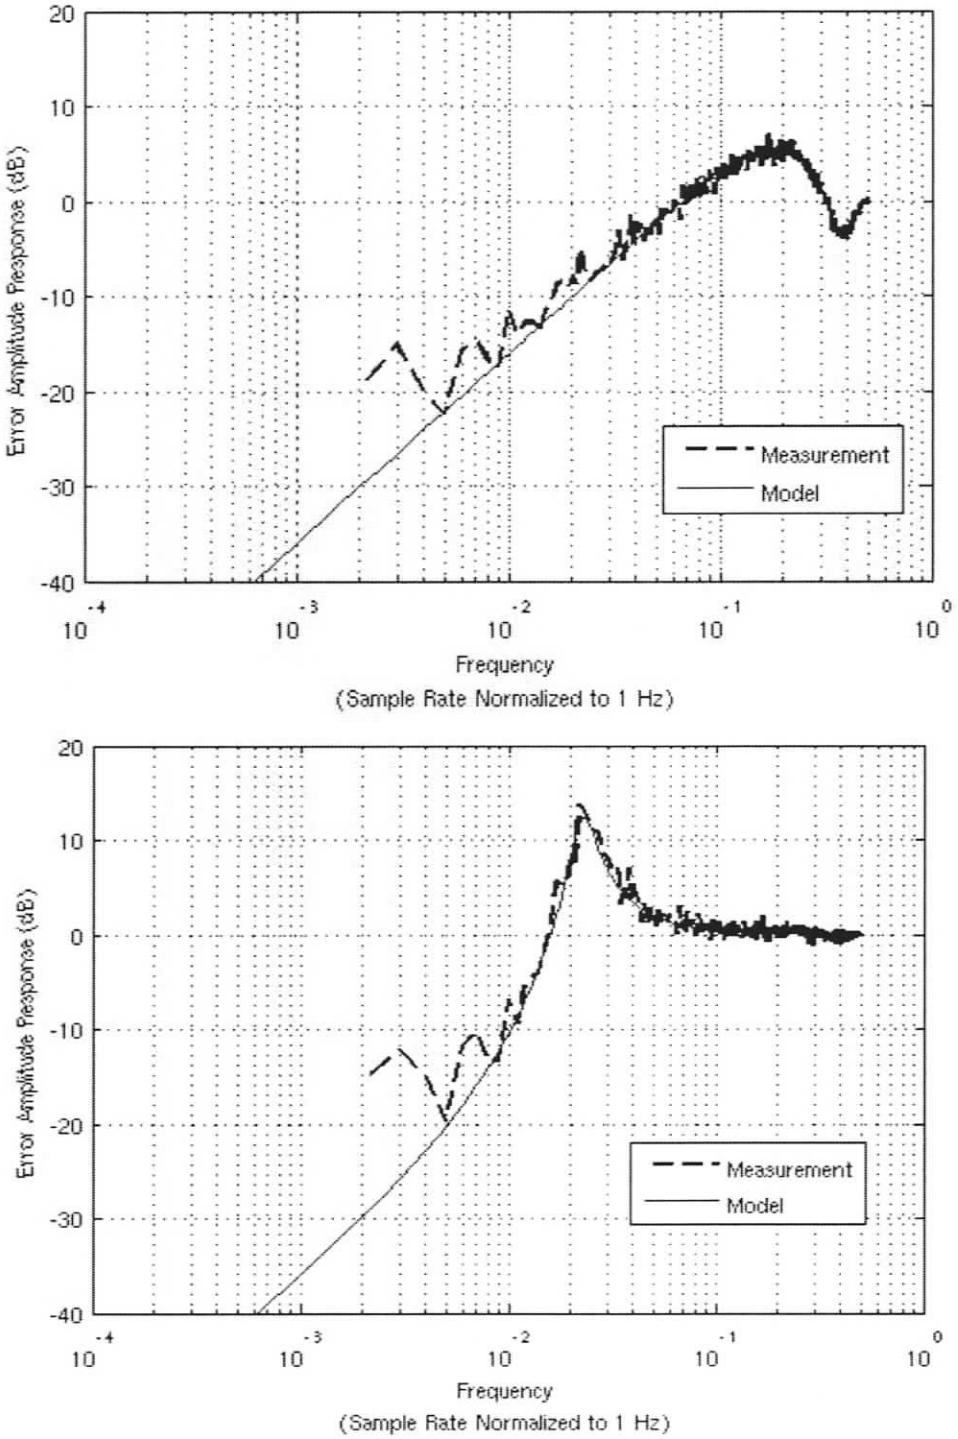

Figure 7: Rejection ratio of a Woofer mode with (top) and without (bottom) Tweeter assistance for $\tau = 20T_s$

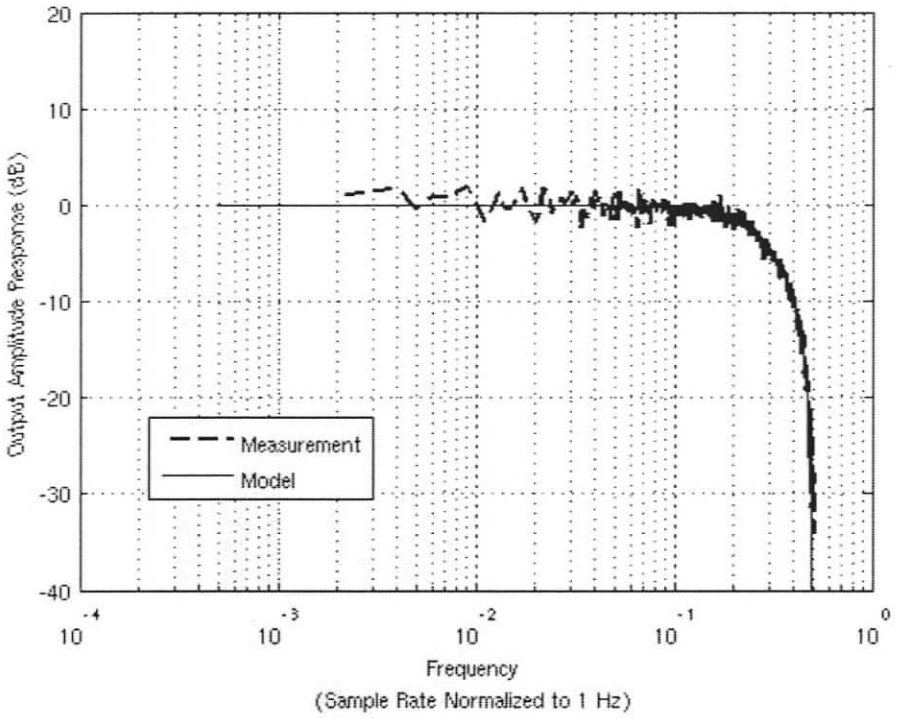

Figure 8: Example output amplitude response

4.2 Simulated turbulence input

Another issue that needed to be studied is the amount of stroke that is offloaded to the Woofer from the Tweeter in this scheme. A simulated phase screen for a turbulence of $D/r_o = 6.25$ was generated and the simulated Tip Tilt mirror and Tweeter corrected it in a simulated AO system. This was to determine appropriate actuator signals that would correspond to such turbulence. These actuator commands were then scaled to generate a significant disturbance and applied to the real Tweeter without the controller's knowledge. Effectively, this generates turbulence that has the correct distribution of spatial and temporal frequency power. It was found that the maximum standard deviation of the Tweeter stroke was 4.859 discrete steps when the maximum standard deviation of the input turbulence was 8.971 discrete steps. The minimum standard deviation of

Tweeter actuation was 1.527 discrete steps when the turbulence minimum standard deviation of the turbulence was equivalent to 4.594 discrete steps. So the fast Woofer took almost half of the work, which is similar to the results of the simulation when the Woofer assumed the low order polynomial shapes that were similar to Figure 3. It is important to note that there was no tilt in this input, so a significant amount of offloading to the tilt mirror is not taken into account. Also, since the two mirrors are close to the same dimensions, the correlation between the mirrors is poor after the first 9 or 10 modes. It is expected that when the Tweeter has significantly more actuators than the Woofer that the matching of shapes will be improved and more modes can be offloaded to the Woofer.

CONCLUSIONS

The Adaptive Optics test bench at the University of Victoria is used to design and test an Adaptive Optics control system consisting of a slow Woofer mirror and a fast Tweeter mirror. The correlated shapes are determined using WFS measurements from the calibration step. The resultant shapes are similar to low order Zernike modes. Linearization of the Tweeter actuators requires additional computation power but allows linear mathematics to apply for offloading to the Woofer. It also maintains the desired gain value for any level of actuation. Using an appropriate high pass filter in the Tweeter controller provides the necessary offloading of low frequency correction from the Tweeter to the Woofer. With this control scheme, the combination of a slow Woofer with a fast Tweeter has the same dynamic response as a single fast mirror. This allows for the large amplitude, low spatial frequency disturbances to be corrected by the Woofer without degrading the performance of the system. The measured Woofer Tweeter

rejection ratios matched the rejection ratio expected from a fast single mirror system. The bandwidth of this system is approximately 6.5% of the sample rate regardless of Woofer response speed. The computation cost of the proposed controller implementation increases linearly with the number of Tweeter actuators.

FUTURE WORK

Now that the basic components of the Woofer-Tweeter system are operational and behaving according to the expected models, the controller can be expanded to meet the needs of the TMT project. The mirrors should be upgraded to have actuator density on the same order as what would be used in a TMT instrument so that the desired amount of offloading can be verified. It is planned that a 1k actuator mirror will replace the 140 actuator Tweeter and a 9x9 Woofer to replace the 8x8 LAOG Woofer.

The initial TMT plans for the wave front reconstruction is a tomographic reconstruction followed by a filter given in Eq. 6.1 then given to a controller that is probably of the form given in Eq. 3.11. Rather than filter the error signal, then pass the result to the controller, the control approach provided in this paper performs both the filtering and the control at once using a specific case of $\alpha_0 = \alpha_1 = 0.5$ and $\beta_1 = \beta_2 = 0$. In other words, the closed loop response of the AO system will have the form given in Eq. 6.2, rather than using $f(z)$ as a sub component of the controller as shown in Eq. 6.3. This will accomplish the desired filtering of the WFS data prior to the mirror without excessive sacrifices to bandwidth. Note that the values of α and β are expected to be different depending if Eq. 6.2 or Eq. 6.3 is used. The control scheme for the general case of Eq. 6.1 has been developed but has yet to be experimentally tested. It is expected that

allowing β to be non-zero would increase the computational complexity by $4n+8m$, where n is Tweeter actuators and m is Woofer modes.

$$f(z)=\frac{\alpha_0+\alpha_1 z^{-1}}{1+\beta_1 z^{-1}+\beta_2 z^{-2}} \tag{6.1}$$

$$T_{current\ plan}(z)=f(z)G_{DM}(z)G_{WFS}(z) \tag{6.2}$$

$$T_{initial\ TMT\ plan}(z)=\frac{f(z)G_{CAO}(z)G_{DM}(z)G_{WFS}(z)}{1+f(z)G_{CAO}(z)G_{DM}(z)G_{WFS}(z)} \tag{6.3}$$

It is understood that the plan for TMT is to combine several NGS wave front error measurements to determine the necessary tilt and focus correction. The controller will be expanded to allow for multiple sources of wave front data for these 3 low order modes.

REFERENCES

[1] R. K. Tyson, *Introduction to Adaptive Optics*, SPIE Press, Bellingham, Washington, 2000

[2] R.J. Noll, Zernike polynomials and atmospheric turbulence, *J. Optical Society of America* (1976) 207-211.

[3] P.J. Hampton, *Design and Implementation of Discrete Type 1 and Type 2 Controllers with Residual Feedback for Modal Adaptive Optics,* University of Victoria, Victoria, BC, 2005.

[4] K. Ogata, *Discrete-Time Control Systems, Second Edition*, Prentice-Hall, Inc., New Jersey, USA, 1995

[5] D.S. Watkins, *Fundamentals of Matrix Computations*, John Wiley & Sons, Inc., New York, USA, 2002

Errata

The formula for the computational complexity of Eq. 3.27 is not a direct count of the flops associated with Fig. 5. The equation may reflect the computation cost of a more efficient implementation that was more difficult to show on a diagram. The computation cost directly from Fig. 5 is

$$flops=n(7+4q+6m)+m(11+4p)$$