

Sensorless Wavefront Correction Algorithms for Free-Space Optical Communications

by

Ondrej Čierny

B.S., Slovak University of Technology (2015)

S.M., Julius Maximilians University of Würzburg (2017)

S.M., Luleå University of Technology (2017)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2023

© 2023 Ondrej Čierny. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

- Authored by: Ondrej Čierny
Department of Aeronautics and Astronautics
May 23, 2023
- Certified by: Kerri L. Cahoy
Associate Professor and Bisplinghoff Faculty Fellow
Thesis Committee Chair
- Certified by: Jeffrey H. Shapiro
Julius A. Stratton Professor of Electrical Engineering
Thesis Committee Member
- Certified by: Franco N. C. Wong
Senior Research Scientist, Research Laboratory of Electronics
Thesis Committee Member
- Certified by: Baris I. Erkmen
Chief Executive Officer, Hedron
Thesis Committee Member
- Accepted by: Jonathan P. How
R. C. Maclaurin Professor of Aeronautics and Astronautics
Graduate Program Committee Chair

This page intentionally left blank.

Sensorless Wavefront Correction Algorithms for Free-Space Optical Communications

by

Ondrej Čierny

Submitted to the Department of Aeronautics and Astronautics
on May 23, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Aeronautics and Astronautics

Abstract

Free-space optical communications (FSOC) technology facilitates high-throughput wireless links across large distances with low size, weight, and power (SWaP) terminals. However, it is difficult to design reliable, low-cost FSOC terminals for long-range links through the atmosphere. Even in clear conditions, the effects of air turbulence along such links usually necessitate active wavefront correction via adaptive optics (AO). Conventional AO algorithms rely on direct wavefront sensing, an approach that is high in cost and SWaP and usually degrades in strong atmospheric scintillation. Sensing methods that are more tolerant to scintillation have been developed, but they are often more challenging to implement and further increase cost and SWaP.

Sensorless wavefront correction algorithms, such as stochastic parallel gradient descent (SPGD), are preferable in terms of cost and SWaP, and have been used in FSOC terminals as methods to optimize the received signal strength indicator (RSSI). A key challenge with such algorithms, however, is that their convergence rate degrades as more atmospheric modes are optimized. This can lead to an inadequate correction rate due to limited bandwidth of the AO element and cause link interruptions. To maintain a sufficient link margin in such conditions, correction algorithms with better convergence properties are needed.

This thesis focuses on the development and testing of a new non-stochastic algorithm for multimodal wavefront correction and a more general analysis of the circumstances where sensorless algorithms attain adequate performance for FSOC, including in strong scintillation. An end-to-end simulation environment is built to compare SPGD with the developed non-stochastic algorithm over a range of atmospheric conditions and hardware configurations. We show that in identical conditions, the non-stochastic algorithm either improves the link margin by 2–3 dB or relaxes the AO element bandwidth requirement by a factor of 2–3 compared to SPGD. Finally, the simulation results are validated in the laboratory under simulated atmospheric turbulence and compiled into a useful design tool for predicting sensorless wavefront correction performance.

Thesis Supervisor: Kerri L. Cahoy

Title: Associate Professor and Bisplinghoff Faculty Fellow

This page intentionally left blank.

Acknowledgments

This work would have never been possible without the generous support and guidance from many people. I want to start by acknowledging my advisor, Prof. Kerri Cahoy, for her remarkable mentorship since taking me in as a visiting student almost six years ago. Without hesitation, Kerri always encouraged and supported my evolving research interests, for which I am deeply grateful. I would also like to thank my committee members, Prof. Jeffrey Shapiro, Dr. Franco Wong, and Dr. Baris Erkmén, as well as my thesis readers, Dr. Paul Serra and Dr. Stephen Palese, for providing insightful comments and sparking rich discussions, which helped improve the quality of this work.

A huge debt of gratitude is owed to multiple people on Project Taara at X, where this research was carried out. I thank Baris, who was with Taara previously, for fostering collaborations with academic labs such as Kerri's, which led to this research opportunity in the first place. I would also like to acknowledge Devin Brinkley, Jean-Laurent Plateau, Paul Csonka, and Paul Epp for supporting me on this journey and making sure I have the necessary resources and time to focus whenever needed. Finally, a special thanks to Andrei Kazmierski, Sanam Mozaffari, Klaus Ulander, Garth Brydon, and Chris Roelle, whose work on the hardware testbed and atmospheric simulator facilitated my research.

Before joining Taara, I was lucky to build invaluable inter-disciplinary engineering skills by working on several NASA-sponsored projects. I thank the NASA STMD for supporting me through the CLICK mission and the NASA ESI program for sponsorship of the MOSAIC project. My graduate school experience would never have been as exciting without the incredible journey of working alongside my friends Will Kammerer, Peter Grenfell, Paul Serra, Grant Gunnison, Faisal Fogle, Shreeram Kacker, and others on the CLICK and MOSAIC teams.

On the personal front, I am deeply grateful for my parents Edita and Jaroslav, who have always supported my ambitions and have been big role models for me. My interest in aerospace engineering can be traced back to breathtaking lectures in physics from my excellent college teacher, Dr. Pavol Valko, which ignited a passion for space exploration in me that persists to this day. I would also like to acknowledge my best college friend,

Richard, who motivated me to step outside my comfort zone and embark on studying and working abroad. Our adventures of living together in different parts of Europe are some of my fondest memories as a student. I would also like to thank all my other friends that have followed me on this journey, both from Europe and those I met in the US, for making this experience so much more enjoyable. A special thanks to Maya and Juanjo for the countless hours we spent on Zoom peer-pressuring each other to get our theses written.

Finally, I am beyond grateful for Alexa, who has been the most supportive and loving partner one could dream of during the past four years. Even though these years were some of the most stressful ones, she managed to help me stay sane and turn it into the most fulfilling and enjoyable period of my life, for which I owe her a great debt of gratitude.

Contents

1	Introduction	21
1.1	Motivation	21
1.2	Research Contributions	23
1.3	Thesis Organization	24
2	Background and Preliminaries	25
2.1	Complexity Drivers	25
2.1.1	Link Acquisition	25
2.1.2	Pointing and Tracking	28
2.1.3	Atmospheric Effects	30
2.2	Challenges with Turbulence Compensation	31
2.2.1	Conventional and Sensorless Architectures	33
3	Wavefront Correction via Optimization of Fiber RSSI	37
3.1	Stochastic Parallel Gradient Descent	38
3.1.1	Weighted Modal SPGD	40
3.2	Correction of Binary Modes via Time-Division	41
3.2.1	Walsh Modes	42
3.2.2	Estimation and Correction	42
3.3	RF-Inspired Zonal Algorithms	45
3.4	Discussion and Research Gaps	46

4	Generalized Non-Stochastic Modal Algorithm	49
4.1	Relationship for Arbitrary Bases	49
4.1.1	Individual Perturbations	50
4.1.2	Multiple Perturbations	52
4.1.3	Logarithmic Measurements	55
4.2	Coefficient Estimation	56
4.2.1	Step Dither	56
4.2.2	Sinusoidal Dither	57
4.2.3	Quadrature Dither	60
4.3	Multimodal Correction	61
4.3.1	Time-Division	61
4.3.2	Frequency-Division & Other Methods	64
4.3.3	Error Rejection	65
5	Simulation Studies	67
5.1	Simulation Environment	67
5.1.1	DM Model	68
5.1.2	Receiver Model	71
5.1.3	Atmospheric Data	74
5.2	Algorithm Configurations	76
5.2.1	Number of Modes	77
5.2.2	Dither Magnitude	77
5.2.3	Feedback Gain	78
5.2.4	Dither and Gain Weighting	78
5.2.5	System Dynamics	79
5.2.6	Joint Optimization	80
5.3	Issues with Non-Binary Modes	83
5.3.1	Basis Modifications	83
5.3.2	Local Maxima	86
5.4	Performance Evaluations	90

5.5	Discussion	93
5.6	Practical Case Studies	94
6	Experimental Validation	97
6.1	Laboratory Setup	97
6.2	Hardware Calibration	101
6.2.1	Turbulence Injection	101
6.2.2	Estimation Weights	103
6.3	Algorithm Validation	104
6.4	Discussion	109
7	Conclusion	111
7.1	Summary of Work and Contributions	111
7.2	Limitations and Future Work	113

This page intentionally left blank.

List of Acronyms

ADC	analog-to-digital converter
AM	amplitude modulation
AO	adaptive optics
APD	avalanche photodiode
API	application programming interface
BMC	Boston Micromachines Corporation
CBC	coherent beam combining
CD	code-division
CDMA	code-division multiple access
DAC	digital-to-analog converter
DM	deformable mirror
DR	dynamic range
DSP	digital signal processing
EDFA	erbium-doped fiber amplifier
EDRS	European Data Relay System
ESC	extremum seeking control
FD	frequency-division
FOR	field of regard
FOV	field of view
FSOC	free-space optical communications
FWHM	full width at half maximum
GEO	geostationary orbit
GPS	Global Positioning System

IF influence function
IMU inertial measurement unit
IQ in-phase & quadrature
LEO low-Earth orbit
log amp logarithmic amplifier
MEMS micro-electro-mechanical system
PD photodiode
PM polarization-maintaining
QAM quadrature amplitude modulation
RF radio frequency
RMS root mean square
RSSI received signal strength indicator
SHWFS Shack-Hartmann wavefront sensor
SLM spatial light modulator
SMF single-mode fiber
SNR signal-to-noise ratio
SPGD stochastic parallel gradient descent
SR Strehl ratio
SRI self-referencing interferometry
SWaP size, weight, and power
TD time-division
TDMA time-division multiple access
TIA transimpedance amplifier
WFS wavefront sensor

List of Symbols

Notation	Description	Page List
E	scalar electric field	31, 71
χ	field log-amplitude	31, 71, 73
Φ	field phase (phasefront)	31, 32, 40, 41, 42, 44, 71, 73
r_0	Fried parameter	32, 41, 47, 74, 75, 76, 80, 83, 85, 86, 87, 88, 90, 92, 93, 94, 95, 102, 104, 106, 107, 112, 116
τ_{atm}	atmospheric time constant	32, 38, 41, 44, 47, 56, 74, 76, 77, 78, 80, 81, 82, 90, 91, 92, 94, 95, 97, 104, 106, 109, 112, 113, 115, 116
σ_R^2	Rytov variance	32, 41, 47, 74, 75, 76, 83, 85, 87, 104, 106, 112, 115
D	receive aperture diameter	32, 41, 47, 71, 72, 74, 75, 76, 80, 83, 85, 86, 87, 88, 90, 92, 93, 94, 95, 102, 104, 106, 107, 112, 116
i	mode index	32, 38, 40, 41, 42, 44, 52, 79, 105
a_i	modal coefficient	32, 33, 40, 41, 42, 44, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60, 65, 77, 78, 79, 81, 86, 103
\mathcal{M}_i	2D spatial mode	32, 33, 40, 41, 50, 51, 52, 53, 56, 57
τ_m	measurement delay	38, 39, 41, 44, 56, 61, 76, 79, 81
u	control vector	38, 39, 65
ϕ	phase dither (scalar)	38, 39, 40, 42, 44, 45, 56, 57, 58, 59, 60, 61, 77, 78, 79, 80, 81, 103

Notation	Description	Page List
J	optimization metric	38, 39, 41
S	RSSI measurement	38, 41, 42, 44, 50, 52, 55, 56, 57, 58, 60, 61, 73, 77, 86, 103
ψ	gradient estimation noise	38, 39
μ	base SPGD gain	39, 41, 46, 78, 79, 81
γ	base weight	40, 41
N	total number of modes	41, 44, 45, 49, 53, 62, 65, 77, 80, 81, 83, 89, 113
k	algorithm iteration	41, 44, 66
φ	phasefront dither	41
\mathcal{W}_i	Walsh mode	42, 44, 50, 86
\mathcal{R}_n	Radermacher function	42
n	function order	42
ℓ	highest order (n_{\max})	42, 43, 44, 49
\mathbf{H}	Hadamard matrix	42, 43
b	peak RSSI (constant)	42, 44, 50, 52, 53, 55, 56, 57, 58, 59, 60
S_0	baseline RSSI	42, 44, 56, 57
S_1	dithered RSSI	42, 44, 56, 57
w_i	calibration weight (linear)	52, 53, 55, 56, 58, 59, 61, 104, 105
j	additional mode index	53, 60
c_i	calibration weight (log.)	55, 57, 58, 60, 61, 77, 103, 104, 105
$w_{1,i}$	1 st estimation weight (linear)	56, 59
$w_{2,i}$	2 nd estimation weight (linear)	56, 59
$c_{1,i}$	1 st estimation weight (log.)	57, 58, 60, 103
$c_{2,i}$	2 nd estimation weight (log.)	57, 58, 103
ω_i	dither frequency	57, 58, 59, 60, 61, 65
f_{est}	estimation bandwidth	61, 62, 66

Notation	Description	Page List
f_{dm}	DM bandwidth	61, 62, 64, 65, 70, 73, 74
τ_{dm}	DM actuator time constant	61, 69, 70, 74, 76, 77, 79, 80, 81, 82, 90, 91, 92, 94, 95, 97, 104, 106, 109, 112, 113, 115, 116
α	actuator settling point	61, 62, 64, 76, 79, 80, 81
T_{est}	estimation period	62
β	spectral efficiency	62, 64
f_s	sampling frequency	62, 64, 66, 81
g	feedback control gain	66, 78, 79, 80, 81
M	DM actuator across aperture	68, 69, 70, 80
η	fiber coupling efficiency	71, 72, 73, 80, 81, 83, 85, 88, 90
σ_n^{dB}	RMS RSSI noise	73, 78, 80, 82, 91, 92
f_{det}	RSSI detection bandwidth	73
λ	wavelength	76, 94, 95
γ_i	normalized atmo. weight	79, 81
e_ϕ	dither weighting exponent	79, 81
e_g	feedback gain weighting exp.	79, 81
e_μ	SPGD gain weighting exp.	79, 81
P	percentile	90, 94, 95

This page intentionally left blank.

List of Figures

Fig. 1-1	Conventional terminal architecture	22
Fig. 2-1	Beaconless acquisition	27
Fig. 2-2	Hybrid array-based receivers	29
Fig. 2-3	Dither-based alignment correction	30
Fig. 2-4	Impact of clear air turbulence	31
Fig. 2-5	AO with fiber-coupled receivers	33
Fig. 2-6	SHWFS versus SRI	35
Fig. 3-1	Walsh modes	43
Fig. 4-1	Fiber RSSI approximation error	51
Fig. 4-2	Approximation error under multiple perturbations	54
Fig. 4-3	Diagram comparing time-division architectures	63
Fig. 4-4	Diagram of modal frequency-division	65
Fig. 5-1	Subsystems of the simulation environment	68
Fig. 5-2	Example of fields from WaveTrain	75
Fig. 5-3	Example of joint parameter optimization	82
Fig. 5-4	Visual comparison of non-binary bases	84
Fig. 5-5	Quantitative comparison of non-binary bases	85
Fig. 5-6	Example of local maxima with non-binary mode adjustments	87
Fig. 5-7	Illustration of issues caused by local maxima	88
Fig. 5-8	Summary of correction performance across bandwidth ratios	91
Fig. 5-9	Summary of performance across all conditions	92

Fig. 5-10 Tail of the performance distribution	93
Fig. 6-1 Diagram of the experimental setup	98
Fig. 6-2 Annotated photo of the experimental setup	100
Fig. 6-3 Calibration of turbulence injection	102
Fig. 6-4 Calibration of estimation weights	105
Fig. 6-5 Comparison of experimental and simulated wavefront correction	107
Fig. 6-6 Summary of simulated versus measured performance	108

List of Tables

Tab. 2.1	Statistical properties of atmospheric turbulence	32
Tab. 5.1	Comparison of layers within the DM model	69
Tab. 5.2	Commonly used DM model parameters	69
Tab. 5.3	Example receiver noise calculation	74
Tab. 5.4	Summary of atmospheric dataset parameters	76
Tab. 5.5	Summary of algorithm configurations	81
Tab. 6.1	Summary of hardware components	98
Tab. 6.2	Comparison between simulated and verified scenarios	106

List of Algorithms

Alg. 1	Pseudocode for weighted modal SPGD	41
Alg. 2	Wang’s method for correction of Walsh modes	44

This page intentionally left blank.

Chapter 1

Introduction

1.1 Motivation

The interest in line of sight free-space optical communications (FSOC) technology continues to experience significant growth. Historically, FSOC links were built and deployed mainly as technological demonstrations. At present, FSOC technology is starting to be actively deployed for larger-scale commercial use. The early commercial adopters are spread across many sectors of the telecommunications industry: from satellite [1–3], middle-mile [4–7], and last-mile [7, 8] connectivity, all the way to data center interconnects [9, 10]. Although there are many benefits associated with FSOC (*e.g.*, directivity, bandwidth, jamming resistance, and less restrictive licensing), imperfect availability and higher costs hinder more ubiquitous deployment of this technology [11, 12].

One of the main challenges that complicates the design of low-cost and high-availability links is their susceptibility to alignment and atmospheric effects. Compensation of these effects usually necessitates active beam steering, and, for long-range atmospheric links, active wavefront correction via adaptive optics (AO) [13, 14]. Figure 1-1 depicts the high-level components that may be present in such long-range FSOC receivers. Since these compensation technologies typically rely on different optoelectronic sensors, there is a need for multiple conjugate optical sensing planes (besides the modem itself). Light is routed to these planes either from an independent channel (beacon laser), or, less commonly, by splitting the primary communications channel (beaconless systems) via bulk optics. This

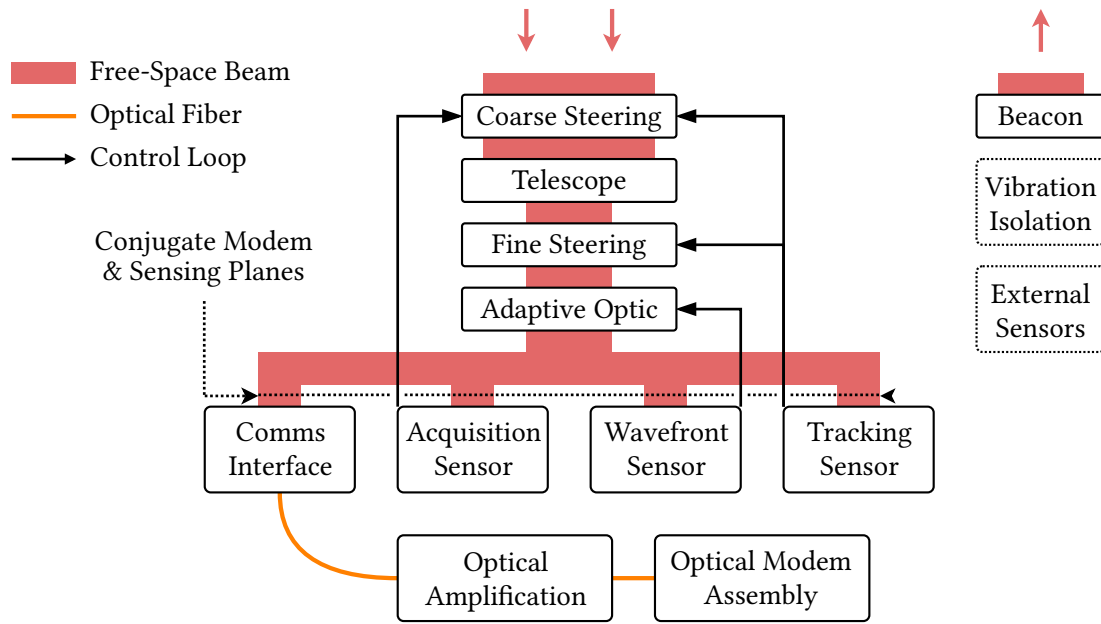


Figure 1-1: High-level depiction of conventional terminal architecture, assuming a longer-range atmospheric link. The need for dedicated optoelectronic sensors for different control loops results in the fragmentation of the device into multiple conjugate optical planes, which increases the hardware complexity and cost. Some blocks may be combined or unnecessary depending on the exact scenario and system requirements.

complicates not only the opto-mechanics but also the manufacturing, as these planes have to be co-aligned and the residual errors calibrated out. The extra complexity associated with the sensors (especially for wavefront sensing), the beacon channel, the optomechanics, and the manufacturing and alignment procedures can significantly contribute to the overall system cost.

To reduce the cost and complexity of these auxiliary systems, researchers have been exploring both beaconless and sensorless approaches. The basic idea behind a sensorless design is to eliminate the conjugate sensing plane and perform the sensor's function either: (a) via a hybrid sensor-receiver, or (b) entirely via the receiver. A simple example of (a) is a multi-element photodiode that is simultaneously used for alignment sensing and communications. For (b), perturbation-based algorithms for real-time optimization of the received signal strength indicator (RSSI) are a good example of utilizing the modem directly without custom sensing hardware. However, such algorithms are less effective for atmospheric links that depend on wavefront sensing for active AO correction. The core

issue with sensorless AO algorithms, such as stochastic parallel gradient descent (SPGD), is that their convergence rate degrades as more atmospheric modes are optimized (*i.e.*, due to more optimization degrees of freedom). This can lead to an inadequate correction rate due to limited bandwidth of the AO element and cause link interruptions, which has hindered the transition from dedicated sensing solutions to sensorless wavefront correction [15–18]. To maintain a sufficient link margin in such conditions, correction algorithms with better convergence properties are needed. While some non-stochastic algorithms with more promising performance have been proposed, they have only been demonstrated in a few special cases, and a generalized approach is missing.

1.2 Research Contributions

The aim of this thesis is to explore new approaches for wavefront correction without dedicated or hybrid wavefront sensors. More concretely, we investigate a generalized, non-stochastic algorithm based on optimization of RSSI at the modem. We also focus on improving the understanding of the circumstances under which these algorithms attain adequate performance, particularly in the context of atmospheric FSOC links. To this end, this work addresses three research contributions:

1. Development of a generalized, non-stochastic, sensorless wavefront correction algorithm. Inspired by state-of-the-art research, we devise a gradient-free method that can be calibrated to correct arbitrary modes using arbitrary AO hardware.
2. Verification of performance improvement over stochastic, gradient-based methods in terms of FSOC link margin and the AO bandwidth requirement.
3. Investigation of the algorithm’s performance bounds. By compiling performance results across different conditions, we improve the understanding of the circumstances under which sensorless correction is feasible and under which direct wavefront sensing is still preferable.

1.3 Thesis Organization

This thesis is organized as follows: Chapter 2 gives more background on the key factors that usually lead to complex sensor-driven solutions in FSOC terminals. We show why turbulence compensation, in particular, can be challenging to implement, and describe the drawbacks of conventional methods. The chapter also mentions relevant research in sensorless alternatives where applicable.

Chapter 3 summarizes state-of-the-art sensorless wavefront correction algorithms relevant to low-cost FSOC terminals. We cover three specific approaches from different research areas and discuss their pros and cons. The chapter concludes with a description of the research gaps that this work attempts to address.

The key contributions from this thesis begin in Chapter 4, which describes the development of a novel non-stochastic algorithm inspired by the state-of-the-art approaches discussed in Chapter 3. In the first half of the chapter, simple estimation equations are derived, which enables generalization of the algorithm to arbitrary adaptive optic elements and basis functions. In the second half of the chapter, we summarize different mechanisms through which these equations can be utilized for multi-modal wavefront correction.

Chapter 5 focuses on analyzing the non-stochastic algorithm in an end-to-end simulation environment. The simulations are used to better understand its behavior in different configurations and to rapidly evaluate its performance. The algorithm is then compared to the state-of-the-art stochastic algorithm over a range of atmospheric conditions and hardware parameters. We compile the results into a useful tool for performance prediction and show an improvement in performance over the state-of-the-art.¹

In Chapter 6, we describe an approach for hardware validation of the simulation results in the laboratory. The design of an experimental testbed is discussed, along with the results from its calibration. The chapter concludes with a comprehensive summary of the results from the performance validation.

Finally, Chapter 7 summarizes the work, the contributions, and outlines a few possible areas for future research.

¹The key findings from Chapters 4 and 5 are also reported in [19].

Chapter 2

Background and Preliminaries

2.1 Complexity Drivers

The conventional multi-sensor architectures stem from the need to meet a few distinct requirements related to bringing up and maintaining the FSOC link. In this section, we categorize the performance requirements between link acquisition, pointing & tracking, and compensation of atmospheric effects. Additionally, we highlight some alternative sensorless approaches or other proposed methods to reduce the associated systems' complexity and cost.

2.1.1 Link Acquisition

The first step in closing an FSOC link between two terminals is to establish sufficient alignment for power to be acquired in the modem's optical receivers. Since FSOC utilizes particularly narrow beams (commonly on the order of μrad), in practice, the initial alignment error is many times larger than the beam's angular extent. The signal acquisition is further complicated by the fact that the field of view (FOV) of the conventionally used single-mode fiber (SMF) coupled optical receivers is also very narrow (also commonly on the order of μrad , considering high-gain receivers with few μm cores). While the initial alignment uncertainty can be reduced at the expense of adding finer position and orientation sensors (*e.g.*, IMU, GPS, star tracker), acquisition usually demands even higher

accuracy, and other solutions need to be implemented.

The conventional approach is twofold: (1) use a dedicated alignment sensor (*e.g.*, a focal plane array) with large enough FOV to cover the alignment uncertainty, and (2) increase the angular extent covered by the beam (*e.g.*, using a beacon laser with a larger divergence angle). After the signal is acquired at the detector's plane, the feedback is used to correct the alignment error and enable signal acquisition at the SMF-coupled receiver with a narrow FOV. Link acquisition can be particularly challenging for terminals on moving platforms, such as aircraft or spacecraft, where the relative orientation changes at a high rate, and the acquisition window can be relatively short. In such scenarios, both high-rate and wide field of regard (FOR) sensing and beam steering solutions are needed.

Beaconless acquisition

While sensorless acquisition (*i.e.*, using only a narrow-FOV optical modem) has yet to be demonstrated in the literature, researchers have been exploring methods for beaconless acquisition to simplify the hardware design. Eliminating the beacon laser, however, comes with two challenges: (1) some fraction of the primary channel's power needs to be routed to the acquisition sensor, reducing the link margin, and (2) the typically smaller divergence angle of the primary beam complicates the alignment process. While there is no way around the loss of some signal with a dedicated sensor, two solutions for utilizing the smaller divergence beam have been discussed in the literature: beam scanning and beamwidth control. A high-level illustration of these methods is shown in Fig. 2-1.

Beam scanning, or search, utilizes steering actuators to probe the uncertainty cone, usually with a spiral or raster search pattern. The downside is that it can be more time-consuming than beacon-assisted acquisition. This approach has been successfully demonstrated in practice, for instance, by the European Data Relay System (EDRS) satellites, which provide low-Earth orbit (LEO) to geostationary orbit (GEO) relay services [20–24]. On the other hand, beamwidth control uses mechanisms that enable some adjustability of focus along the optical path. As illustrated in Figure 2-1, during acquisition, defocus is injected to expand the beam and, ideally, cover the entire uncertainty cone. The defocus is removed for communications, and the beam is as close to diffraction-limited as possible.

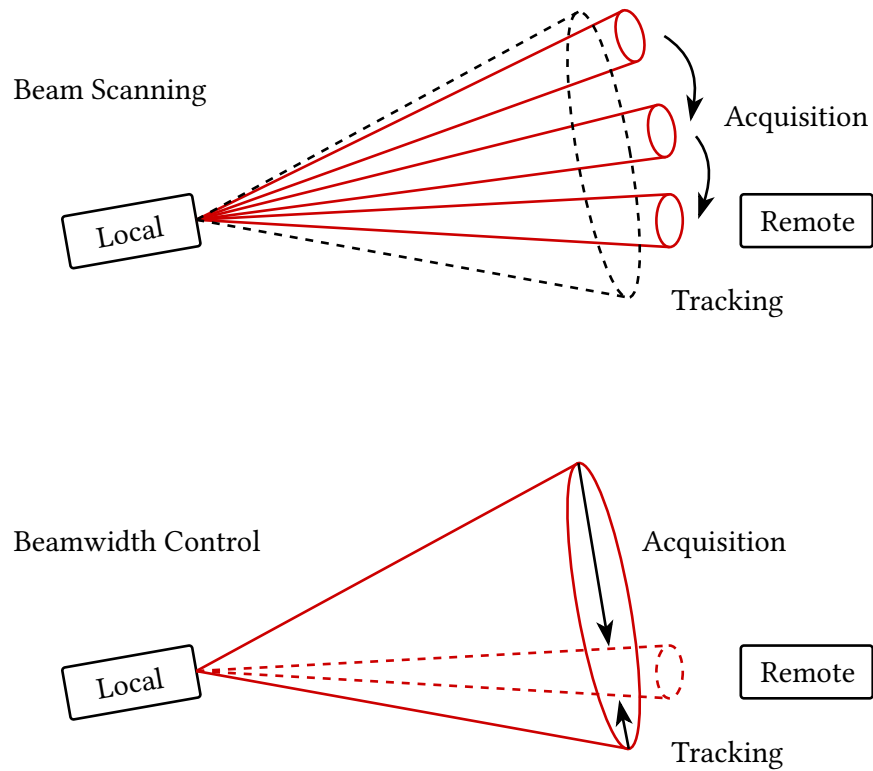


Figure 2-1: Depiction of beaconless probing of the uncertainty cone between a local and a remote terminal. The top graphic illustrates beam scanning, while the bottom case depicts how beamwidth control can enable both wide-beam acquisition and narrow-beam communications.

Various implementations of this concept were proposed, for example, using motorized lenses, fibers, or wedge pairs [25–29].

2.1.2 Pointing and Tracking

After acquisition, the terminals must continuously maintain alignment within a fraction of the beamwidth to minimize fluctuations in the received power. Any significant dynamic errors (*e.g.*, due to platform motion or vibration) need to be actively compensated, which is commonly accomplished via a closed-loop spatial tracking system. Since this loop’s effectiveness depends on the accuracy of beam steering and alignment sensing, the terminals require both precision sensors and fine-pointing mechanisms.

Unlike link acquisition, which necessitates larger FOR sensing but relatively low bandwidth, tracking systems utilize smaller FOR at a higher bandwidth. Depending on the spectrum of disturbances the device is subject to, multiple (staged) disturbance rejection systems may be necessary in parallel for different frequency bands (these might also include, for example, passive or active vibration isolation systems). Consequently, it is common to have a dedicated conjugate plane with a faster, narrower-FOV alignment sensor used exclusively for tracking.

Hybrid Array-Based Receivers

Due to the lower FOR requirement, it is possible to eliminate the conjugate tracking plane by using a hybrid sensor-receiver that combines the functions of fine alignment sensing and communications. These devices typically rely on some array-like arrangement of active surfaces (or waveguides), such as pairs or quadrants, to enable one- or two-axis position sensing via differentials in the amount of captured light. With fiberless coherent receivers, this approach has been demonstrated using balanced photodetectors mounted perpendicular to each other along separate interferometer arms [30], or using a quadrant detector on a single arm [31]. With incoherent (*i.e.*, based on direct detection) fiberless receivers based on avalanche photodiodes (APDs), simultaneous gigabit-rate communications and spatial tracking was demonstrated with hybrid multi-segment APD




Figure 2-2: Simplified visualizations of some hybrid array-based receivers: (a) coherent receiver with two arms and perpendicular balanced detectors (RX: received light, LO: local oscillator), (b) coherent receiver with one arm and a single quadrant detector, (c) hybrid APD array for incoherent detection, (d) tracking with a fiber bundle.

arrays [32, 33]. Finally, with fiber-coupled receivers, several investigations explored the use of fiber bundles to enable inference of position from differences in the fiber-coupled intensities [34–36].

Dither-Based Alignment Correction

Another sensorless approach explicitly developed for detectors coupled via SMF is based on direct optimization of the detected optical intensity (commonly referred to as modem RSSI in the context of FSOC). Since single-mode waveguides are particularly sensitive to alignment, hill-climbing optimization algorithms can effectively find the maximum coupling efficiency and, by extension, the optimal alignment. Extremum seeking control (ESC) is the most commonly used technique to implement such alignment correction feedback loop in practice. In ESC, a small sinusoidal dither is injected to perturb the system at some frequency, and the output at this frequency is demodulated to infer the appropriate direction of correction.

Since tracking is a two-dimensional problem, both azimuth and elevation have to be dithered simultaneously at different frequencies or using two 90° -offset sinusoids, resulting in a circular dither pattern. The latter case, often referred to as fiber nutation or nutation tracking in literature, uses an IQ demodulator to isolate the azimuth and elevation controls. Nutation tracking has been demonstrated with actuated fibers [37–39],

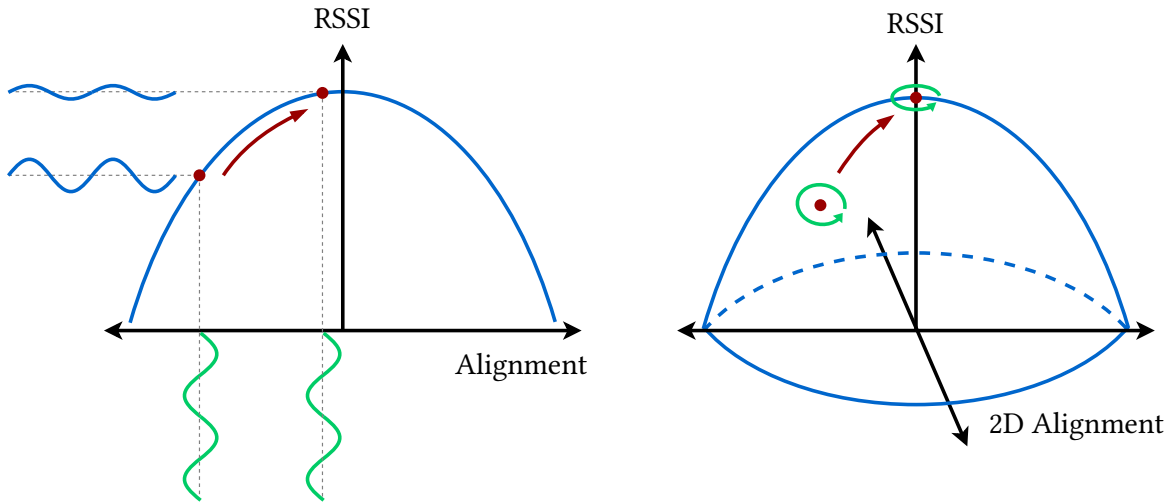


Figure 2-3: A high-level illustration of alignment correction via ESC in one dimension (left) vs. two dimensions, which is based on two IQ sinusoidal dithers forming a circular dither pattern (nutation tracking, right).

electro-optic modulators [29, 40], and actuated mirrors [41–43].

2.1.3 Atmospheric Effects

For FSOC links that involve atmospheric propagation, the beam may suffer additional distortion on top of tilt from platform motion and diffraction (spreading) from free-space propagation. Interaction with atmospheric particles through scattering and absorption leads to an irreversible transmission loss, which varies with distance, wavelength, and other environmental factors. In addition, random fluctuations of the local atmospheric refractive index (*i.e.*, clear air turbulence) lead to aberrations of the beam’s phase and amplitude. Depending on the exact link geometry, aperture size, and environmental factors, this may severely impact the efficiency of coupling into the modem at the receiver [13, 44].

Mitigations for absorption and scattering are usually passive and involve choosing an appropriate wavelength and utilizing approaches such as site diversity. In contrast, the higher-order wavefront aberrations from turbulence can be actively mitigated via wavefront correction at the receiver [13, 14]. As shown in Fig. 1-1, AO-based phase conjugation is most commonly employed to minimize the phasefront variation and maximize coupling into the modem. However, given the difficulties associated with the measurement and




Figure 2-4: Illustration of wavefront perturbations resulting from clear air turbulence [45]. A plane wavefront experiences spatiotemporal refraction and diffraction while passing through turbulent atmospheric eddies mixed by winds. These effects occur due to random variations in the local refractive index, which is a function of the local air temperature.

compensation of these variations, wavefront correction is typically the most complex (and costly) of the compensation systems used in FSOC terminals.

2.2 Challenges with Turbulence Compensation

From the perspective of an FSOC receiver, the effect of atmospheric turbulence can be characterized as a spatiotemporal variation in the amplitude and phase of the received wavefront. Mathematically, this is often described as a perturbation of the incident electric field using phasor notation [46]:

$$E(x, y, t) = E_0(x, y) \exp [\chi(x, y, t) + i\Phi(x, y, t)] , \quad (2.1)$$

where E is the scalar component of the electric field, x and y are the spatial coordinates, t is the time coordinate, E_0 is the (unperturbed) incident field, χ is the log-amplitude perturbation, and Φ is the phase perturbation. A few statistical properties are commonly used to quantify the average spatiotemporal variation of the received wavefront. The

Table 2.1: Key statistical properties that describe atmospheric wavefront perturbations.

Statistical property	Definition and remarks
Spatial coherence (Fried parameter, r_0)	Following notation in Andrews 2005 [46], $r_0 = 2.1\rho_0$, where ρ_0 is the spatial coherence radius, defined by the e^{-1} point of the normalized (second-order) complex spatial correlation function (ensemble-averaged over a long time).
Temporal coherence (Time constant, τ_{atm})	Similar to the definition of spatial coherence, we define τ_{atm} by the e^{-1} point of the complex temporal correlation function (in this case averaged spatially, not temporally).
Scintillation regime (Rytov variance, σ_R^2)	Following notation in Andrews 2005 [46], $\sigma_R^2 = 1.23C_n^2 k^{7/6} L^{11/6}$, where C_n^2 is the refractive index structure constant, k is the wavenumber, and L is the link distance. The value $\sigma_R^2 \approx 1$ is considered a threshold between weak and strong fluctuations.

Fried parameter, r_0 , is used to describe the average spatial coherence length of the fields. Similarly, the temporal coherence can be characterized using an atmospheric time constant τ_{atm} . Finally, the variance of the field amplitudes, also referred to as scintillation, is typically quantified using the Rytov variance, σ_R^2 . A summary of these three parameters and their definitions is given in Table 2.1. If the FSOC receiver has an optical aperture of diameter D , an important parameter in terms of wavefront correction is the ratio D/r_0 . The larger D/r_0 is, the higher the spatial frequencies one must compensate to reject the atmospheric perturbations.

While wavefront correction may generally comprise both amplitude and phase compensation (*i.e.*, full-wave correction), practical applications often only implement phase conjugation techniques [14, 47]. To this end, AO elements, such as deformable mirrors (DMs), are commonly utilized to facilitate coupling of the received wavefront into a SMF. The phase compensation can either be *zonal*, when the stroke of each DM actuator is optimized individually, or *modal*, when the correction is applied through some orthogonal basis functions and the modal coefficients are optimized instead. Assuming the phasefront perturbation can be decomposed as

$$\Phi(x, y, t) = \sum_{i=0}^{\infty} a_i(t) \mathcal{M}_i(x, y), \quad (2.2)$$




Figure 2-5: Comparison between a fiber-coupled FSOC receiver with AO based on dedicated wavefront sensing (left) and direct RSSI optimization at the modem (right).

where \mathcal{M}_i are modes from an orthonormal basis (with increasing spatial frequencies) and a_i the associated modal coefficients, the goal of modal compensation is to optimize a_i actuated by the DM such that the reflected phasefront is as flat as possible.

2.2.1 Conventional and Sensorless Architectures

To determine the corrections, a wavefront sensor (WFS) is typically placed behind the AO to sample the residual phasefront variation and provide feedback to the DM in a closed loop. This conventional approach is depicted on the left side in Fig. 2-5. The addition of a WFS complicates the terminal design for several reasons. Due to the difficulty of directly measuring phase at optical frequencies, a WFS is a non-trivial device from both a hardware and a software perspective, which has an impact on its cost. Splitting the wavefront between the conjugate planes of the modem and the WFS also penalizes the link budget (in beaconless systems), and necessitates additional optomechanical components as well as a more complex alignment procedure.

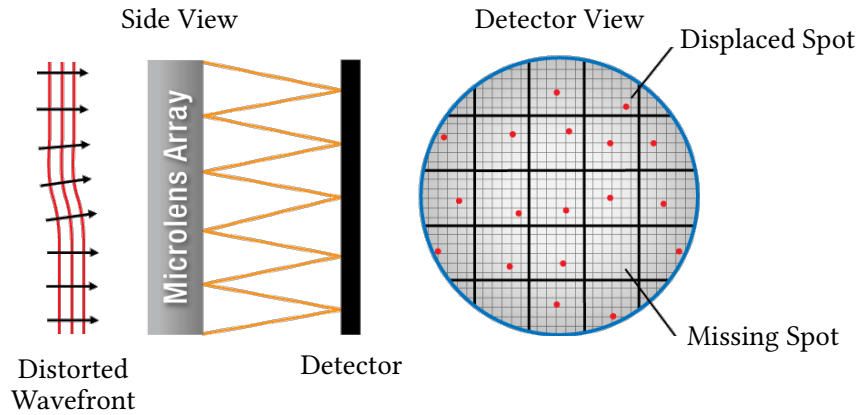
In addition, implementing a reliable WFS under strong atmospheric scintillation poses further technical challenges. For example, the most commonly used Shack-Hartmann wavefront sensor (SHWFS) suffers from degraded wavefront reconstruction in strong scintillation due to speckle noise and branch cuts (discontinuities) in the phase function [48, 49]. On the other hand, alternative sensing methods that are more resilient to scintillation, such as self-referencing interferometry (SRI), are often more difficult to implement and

further increase the system cost and size, weight, and power (SWaP) [50, 51]. A high-level illustration of these two sensing methods is shown in Fig. 2-6.

A lower cost and SWaP alternative to the conventional approach based on a dedicated WFS are iterative dither-based optimization algorithms. As with dither-based alignment correction (*e.g.*, fiber nutation), one can devise a multidimensional optimization algorithm for wavefront correction without having to depend on direct wavefront sensing. In literature, such approaches are sometimes referred to as multidither AO algorithms, wavefront-sensorless AO, or simply sensorless AO. Given that these methods can be fine-tuned to optimize different metrics, they have been explored not only in the context of FSOC, but also for high-contrast imaging and high-power delivery. A variety of these algorithms have been proposed and demonstrated in the literature, both in zonal [53–59] and modal [15, 57, 59–63] configurations.

With FSOC receivers coupled via a SMF, the target optimization metric can conveniently be the modem RSSI. This architecture is visualized on the right-hand side in Fig. 2-5. Since coupling to single-mode waveguides is particularly sensitive to phasefront perturbations [44], these algorithms can effectively find the peak waveguide coupling efficiency and, by extension, the optimal wavefront correction. However, a common issue with such approaches is that the number of optimization steps needed grows with the number of atmospheric modes corrected. This puts pressure on the system to run dithers and corrections as quickly as possible to keep up with dynamic wavefront changes, and the actuator bandwidth will typically bottleneck the realizable refresh rate in situations where multiple modes require compensation [15, 16]. We cover a subset of these algorithms that are relevant to low-cost FSOC receivers in more depth in Chapter 3.

Shack-Hartmann Wavefront Sensor (SHWFS)



Self-Referencing Interferometer (SRI): Top View




Figure 2-6: The working principle of a Shack-Hartmann wavefront sensor (SHWFS) versus a WFS based on self-referencing interferometry (SRI) [51, 52].

This page intentionally left blank.

Chapter 3

Wavefront Correction via Optimization of Fiber RSSI

Given several decades of research and development in fiber-optic communications, it is common for FSOC systems to utilize fiber-coupled photonic components. Apart from the main benefit of the general availability of high-performance, low-cost transceivers and amplifiers, they can also facilitate more efficient communication schemes, such as those based on optically pre-amplified receivers and coherent detection. In terms of system complexity, it is therefore desirable to use the fiber RSSI measurements for wavefront correction algorithms without relying on additional sensing hardware.

A further advantage of this approach stems from the fact that optical receivers utilize very sensitive detectors, which are commonly optimized for operation at multi-GHz bandwidth to support high-rate communications. Since atmospheric turbulence is a much slower process (τ_{atm}^{-1} is typically between tens to hundreds of Hz [14]), one can sample the RSSI at low bandwidth and achieve a very high signal-to-noise ratio (SNR) when the link is operational. This is in contrast to a dedicated WFS, which has access to only a fraction of the received power and further distributes it between many pixels. In this chapter, we give an overview of a few state-of-the-art sensorless wavefront correction algorithms proposed specifically for SMF-coupled FSOC receivers.

3.1 Stochastic Parallel Gradient Descent

One of the most extensively researched algorithms which has been demonstrated with SMF receivers is SPGD, an approach initially proposed for turbulence compensation by Vorontsov [55, 64]. To estimate a wavefront correction, small random step dithers are first simultaneously injected into all the available control channels. The change in RSSI in response to this parallel step perturbation is then measured. It should be noted that the perturbation and measurement must occur while the atmosphere is approximately correlated to obtain useful information, *i.e.*, $\tau_m \ll \tau_{\text{atm}}$, where τ_m is the measurement settling time after dither injection. Since the smallest practical τ_m value is constrained by the actuator bandwidth, which is usually on the order of tens of kHz, it is also assumed here that the RSSI noise is negligible due to the high SNR that is typically attainable at such low measurement bandwidths (in contrast to the GHz-rate communications).

The measurement is normally done after both a positive and a negative (flipped) dither pattern is applied. This approach enables the approximation of the gradients irrespective of the number of control channels involved. In its most basic form, dithers with random signs of equal probability are injected to every channel u_i (i can represent mode or actuator index) in parallel:

$$\phi_i = \pm|\phi|, \quad (3.1)$$

where ϕ is a scalar phase dither magnitude (in radians). Two optimization metrics $J(u_i)$, which in this case are RSSI measurements $S(t)$, are then compared:

$$J^+ = J(u_i + \phi_i), \quad (3.2)$$

$$J^- = J(u_i - \phi_i), \quad (3.3)$$

$$\delta J = J^+ - J^-. \quad (3.4)$$

It can be shown [64] that the expected values $\langle \delta J \phi_i \rangle$ are proportional to the true channel gradients:

$$\frac{1}{\phi^2} \langle \delta J \phi_i \rangle = \frac{\partial J}{\partial u_i} + \psi_i, \quad (3.5)$$

where ψ represents a zero-mean stochastic estimation noise. Since the products $\delta J \phi_i$ approximate the individual gradients, small corrections Δu_i can be applied to the control vector to climb towards the maximum RSSI:

$$\Delta u_i = \mu \delta J \phi_i, \quad (3.6)$$

where μ is a dimensionless feedback gain (sometimes referred to as *learning rate*), which can trade convergence rate for stability and is usually determined empirically.

Depending on the optimization metric, it may be helpful to normalize δJ . For instance, with J being the fiber RSSI, the following normalization can be used to obtain a dimensionless metric irrespective of the average received power:

$$\delta J = \frac{J^+ - J^-}{J^+ + J^-}. \quad (3.7)$$

SPGD has been evaluated with SMF receivers both in zonal [57, 65] and modal [15, 65] configurations. The only difference in implementation is that the zonal variant's u_i are the individual actuator strokes, whereas with the modal variant, u_i are coefficients of some basis functions. A fit between the basis functions and the actuator influence functions is used to transform the coefficients into DM voltages. While the zonal variant can directly optimize each segment, as the number of actuators increases, so does the estimation uncertainty and the number of corrections needed for the system to converge. This starts being an issue when the convergence time, which is a function of both τ_m and the number of actuators corrected, becomes comparable to the coherence time of the atmospheric turbulence.

The benefit of the modal variant is that one can more optimally balance the number of correction channels (*i.e.*, by reducing or expanding the subset of modes) and the convergence rate to maximize performance in specific atmospheric conditions. Indeed, research has shown that even with recent multi-kHz bandwidth DMs, modal SPGD outperforms the zonal variant in typical dynamic turbulence, assuming average wind speed [65]. Nonetheless, the conditions where modal SPGD provides adequate performance are limited due to the need to constrain the number of corrected modes. To extend these conditions,

researchers have focused on selecting only the dominant modes for correction, as well as optimizing the correction basis itself. For instance, utilizing a basis derived specifically to eliminate small correlations observed between a few Zernike polynomial functions, a performance improvement of about 5% has been shown [16, 66, 67].

3.1.1 Weighted Modal SPGD

Another approach that has been investigated to optimize the performance of modal SPGD for atmospheric compensation is the use of weighted (non-binary) dithers ϕ_i . This enables emphasizing or de-emphasizing individual modes in the parallel gradient estimation and correction. Assuming again that the phasefront Φ can be decomposed as

$$\Phi(x, y, t) = \sum_{i=0}^{\infty} a_i(t) \mathcal{M}_i(x, y), \quad (3.8)$$

Vorontsov proved that the modal SPGD convergence rate is theoretically optimal if and only if the dithers ϕ_i and the modal coefficients a_i are statistically dependent functions [64].

While this requirement is impossible to fully meet in practice (a_i are unknown random variables), it is possible to approximate the ensemble variances $\langle a_i^2 \rangle$ from models of the turbulence spatial power spectrum, assuming some expected (average) outdoor conditions. The predicted $\langle a_i^2 \rangle$ can then be used as weighting factors for ϕ_i , and de-emphasize correction of modes with higher spatial frequencies, which are expected to have a smaller contribution to the wavefront error. More specifically, recent works [16, 18] explore weighting the dithers as

$$\phi_i = \pm \gamma \sqrt{\langle a_i^2 \rangle}, \quad (3.9)$$

where γ is a dimensionless constant that scales the coefficient dithers according to the predicted modal root mean square (RMS) values. In this case, the respective $\langle a_i^2 \rangle$ are determined theoretically from a Kolmogorov power spectrum model. A summary of the algorithm based on such weighting is given in Alg. 1.

Dynamic simulations of this approach in the literature indicate an improvement in the average Strehl ratio (SR), however, the performance gain heavily depends on temporal

bandwidth constraints. In a hypothetical scenario where $N = 105$ modes are corrected, and the SPGD correction rate is about one tenth of τ_{atm} , an improvement in the average correction of up to 3 dB has been shown [18]. However, it is unclear how this result scales with N or how it changes with practical limitations such as receiver noise, which will constrain the smallest detectable modal perturbation. A broader analysis of the algorithm's performance across different atmospheric conditions (*i.e.*, τ_{atm} , D/r_0 , and σ_R^2) is also missing.

for $k \leftarrow 0$ to ∞ do	▷ Iteration k
$\varphi \leftarrow \gamma \sum_{i=0}^N \pm \sqrt{\langle a_i^2 \rangle} \mathcal{M}_i$	▷ Construct a parallel weighted dither
$\Phi \leftarrow \Phi + \varphi$	▷ Positive phasefront dither injection
$J^+ \leftarrow S(2k\tau_m)$	▷ RSSI after measurement settling
$\Phi \leftarrow \Phi - 2\varphi$	▷ Negative phasefront dither injection
$J^- \leftarrow S([2k + 1]\tau_m)$	▷ RSSI after measurement settling
$\delta J \leftarrow (J^+ - J^-)/(J^+ + J^-)$	▷ Normalized gradient calculation
$\Phi \leftarrow \Phi + \varphi (\mu\delta J + 1)$	▷ Phasefront correction (and dither removal)
end	

Algorithm 1: Pseudocode for weighted modal SPGD.

3.2 Correction of Binary Modes via Time-Division

Another modal algorithm developed specifically for segmented (piston-only) DMs with detectors coupled via SMF by Wang [61] utilizes serial correction of binary modes through time-division (TD). Unlike SPGD, Wang's method is non-stochastic, non-parallel, and gradient-free. In contrast to gradient-based methods, which do not rely on knowledge of the function being optimized, the algorithm exploits a deterministic relationship between coefficients of the binary Walsh modes and RSSI measurements. More concretely, changes in RSSI resulting from coefficient dithers are used to implement a Walsh mode estimator, which enables correction of the modes via feedback control.

3.2.1 Walsh Modes

The two-dimensional Walsh modes $\mathcal{W}_i(x, y)$ are outer products of one-dimensional Walsh functions, which themselves are products of Rademacher functions $\mathcal{R}_n(x)$. These are binary equivalents of pure spatial frequencies on a unit interval discretized in powers of two:

$$\mathcal{R}_n(x) = \text{sgn}[\sin(2^n \pi x)], \quad n \in \mathbb{N}, \quad x \in [0, 1], \quad \mathcal{R}_n(x) \in \{-1, 1\}, \quad (3.10)$$

where n is the function order, and x the normalized spatial coordinate. Performing all the possible permutations of products between \mathcal{R}_n up to a certain order $n_{\max} = \ell$ results in the one-dimensional set of Walsh functions, which make a complete orthonormal basis on a unit interval discretized to 2^ℓ pixels. The binary values of these functions can also be obtained from rows of a 2^ℓ -order Hadamard matrix \mathbf{H}_{2^ℓ} . Just as the one-dimensional functions, the two-dimensional modes allow decomposing phasefronts $\Phi(x, y)$ up to a spatial resolution of $2^\ell \times 2^\ell$ on a unit square. An $\ell = 3$ subset of the two-dimensional modes is pictured in Fig. 3-1 following the *sequency* ordering (when modes are sorted by increasing spatial frequency). In the limit of $\ell \rightarrow \infty$, an arbitrary phasefront may be decomposed:

$$\Phi(x, y, t) = \sum_{i=0}^{\infty} a_i(t) \mathcal{W}_i(x, y). \quad (3.11)$$

3.2.2 Estimation and Correction

To estimate the residual modal errors, a relationship between a_i and the RSSI measurement S is used. If the residuals are small, Wang has shown [61] that S varies with any individual Walsh coefficient a_i approximately as

$$S \approx b_i \cos^2 a_i, \quad (3.12)$$

where b_i is a function of all the coefficients but a_i . If we step a_i by a known dither magnitude ϕ , the pre- and post-dither measurements S_0 and S_1 can be used to solve for the unknowns




Figure 3-1: Illustration of how two-dimensional Walsh modes are constructed from outer products of one-dimensional Rademacher functions. The modes are displayed in *sequency* ordering (increasing spatial frequency) and up to $\ell = 3$. The white and black spaces correspond to values of +1 and -1 respectively. The two-dimensional binary pattern resembles a sequency-ordered H_{64} Hadamard (Walsh) matrix.

in the equation, which forms the basis of the estimation algorithm:

$$S_0 = b_i \cos^2 a_i, \quad (3.13)$$

$$S_1 = b_i \cos^2(a_i - \phi), \quad (3.14)$$

$$a_i = \arctan \left(\frac{\sqrt{S_0/S_1} - \cos \phi}{\sin \phi} \right). \quad (3.15)$$

A few requirements need to be met to implement this method successfully. Both the injection of ϕ and the correction of a_i depend on the ability to actuate the Walsh modes, which necessitates at least $2^\ell \times 2^\ell = 2^{2\ell}$ addressable piston segments in a square grid on the DM. As with SPGD, τ_{atm} must be longer than the measurement settling time τ_m , and the measurement must have a sufficient SNR. Finally, because the approximation error in Eq. 3.12 increases with large residuals, the dominant a_i should be corrected via feedback control, or else all the estimates may incur additional error.

In practice, if the first two requirements are met, the correction may be performed once an estimate is obtained, and involves removing ϕ and a_i from Φ at the same time. The process is then repeated for other \mathcal{W}_i up to some $i_{\text{max}} = N$, and started over, as summarized in Alg. 2. Due to this modal TD, it is important to choose N appropriately to balance temporal and spatial correction since N affects both the per-mode refresh rate and the residual error from uncorrected modes (*i.e.*, for $i > N$).

<pre> for $k \leftarrow 0$ to ∞ do for $i \leftarrow 0$ to N do $S_0 \leftarrow S(2k\tau_m)$ $\Phi \leftarrow \Phi + \phi \mathcal{W}_i$ $S_1 \leftarrow S([2k + 1]\tau_m)$ $a_i \leftarrow \arctan \left[(\sqrt{S_0/S_1} - \cos \phi) / \sin \phi \right]$ $\Phi \leftarrow \Phi - [a_i + \phi] \mathcal{W}_i$ end end end </pre>	<pre> ▷ Iteration k ▷ Mode i ▷ Baseline RSSI after settling ▷ Step Walsh dither injection ▷ Dithered RSSI after settling ▷ Modal coefficient estimation ▷ Correction and dither removal </pre>
--	--

Algorithm 2: Pseudocode for correction of Walsh modes via TD.

If the algorithm is actively compensating turbulence (*i.e.*, in tracking mode), the third requirement is implicitly satisfied. However, larger residuals may be present when the algorithm is first switched on or after a signal fade (*i.e.*, in AO acquisition mode), leading to estimation bias driven by the inaccuracy of Eq. 3.12. Nonetheless, simulations with $N = 255$ and $\phi = \pi/4$ in frozen atmospheric turbulence have shown that about 70% of the total compensation is usually obtained after the first iteration of modal corrections [61]. After two full iterations, 99% of the correction is achieved, demonstrating that Eq. 3.12 can also be used effectively in this scenario. However, a more comprehensive analysis of how performance is affected when N is practically constrained by the actuator bandwidth or when the dithered measurements are corrupted by noise is lacking. In addition, the approach has not been tested in time-varying turbulence or across different D/r_0 and σ_R^2 .

3.3 RF-Inspired Zonal Algorithms

Coherent beam combining (CBC) is another problem that has prompted a significant amount of research in sensorless phase optimization. A common CBC approach is to recombine multiple amplified SMFs-coupled laser sources, which necessitates phase control along each fiber to compensate for phase drift stemming from the amplifier and path-length variations [68]. Unlike atmospheric effects, which have some degree of spatial phase correlation, phase variations in CBC are typically uncorrelated, necessitating zonal algorithms. However, like with FSOC receivers, a single fiber-coupled detector is desired to simplify the system. Apart from SPGD, researchers analyzed more scalable non-stochastic algorithms, given the growing desire to combine more and more beams (commonly hundreds [68], even thousands proposed [69]).

One broader approach demonstrated in several variations involves dithering each phase tuner at a specific frequency, akin to a radio frequency (RF) carrier. It can be shown that the tuner phase error is proportional to the coherently demodulated power at the given frequency [70], which resembles a basic amplitude modulation (AM) receiver and enables feedback control. However, due to potentially hundreds of channels sharing the same detection chain, multiplexing must be utilized to demodulate all the error signals.

This led to demonstrations inspired by various RF multiple-access methods, such as time-division [71, 72], frequency-division [70, 73, 74], and code-division [75–77]. Apart from AM, the more spectrally efficient quadrature amplitude modulation (QAM) has also been proposed to double the number of tuners per slot, leading to better utilization of the available tuning bandwidth.

3.4 Discussion and Research Gaps

Reviewing the literature, we notice two high-level groups of approaches that are based on RSSI optimization: (1) generalized stochastic algorithms (SPGD variants), and (2) specialized non-stochastic algorithms (multiplexing of binary modes or individual actuators).

Although SPGD is more easily applied to different application cases (*e.g.*, it generalizes to arbitrary AO hardware and basis functions), its convergence rate degrades as more variables are optimized, and it often requires hundreds of iterations to converge in simulated frozen turbulence [57, 58, 65]. Since the dithers are generated randomly, their effect on RSSI is not deterministic, and the usefulness of each iteration varies. In the worst case, a phasefront perturbation may not affect RSSI altogether, and no useful information is measured. The number of iterations also depends on the gradient descent rate (determined by the gain μ), which must be optimized for both speed and long-term stability.

Non-stochastic algorithms, conversely, are based on the injection of deterministic perturbations. Since each optimized variable is dithered independently (through multiplexing), the RSSI measurements can be used to estimate optimal corrections without having to climb toward the maximum (*i.e.*, they enable direct gradient-free optimization). Although multiplexing divides the available bandwidth between all optimized variables, research indicates that the convergence rate may scale more favorably with the number of variables compared to stochastic methods [61, 68]. However, non-stochastic algorithms have only been demonstrated in a few special use cases. A generalized approach that applies to arbitrary AO elements, basis functions, and multiplexing methods is missing.

In general, the literature lacks an in-depth evaluation of different sensorless approaches in higher-fidelity simulations or experiments. Most research works analyze performance

in frozen atmospheric turbulence (without a temporal aspect) and consider only weak or no scintillation. Other important practical considerations, such as the actuator bandwidth or the measurement noise, are also often neglected. Consequently, it is not very clear what performance one can expect with different algorithms as the environmental and hardware parameters change.

This work addresses these gaps by directly comparing stochastic and non-stochastic algorithms over a broad range of atmospheric parameters and hardware configurations. Ultimately, our goal is to improve the understanding of the circumstances in which sensorless approaches attain adequate performance (*e.g.*, across properties such as D/r_0 , τ_{atm} , and σ_R^2), and in which dedicated wavefront sensing is still preferable. In Chapter 4, we address the lack of a generalized non-stochastic algorithm, while in Chapters 5 and 6 we focus on comparing the performance of the stochastic and the non-stochastic approaches across different conditions.

This page intentionally left blank.

Chapter 4

Generalized Non-Stochastic Modal Algorithm

Motivated by the convergence properties demonstrated by Wang's non-stochastic algorithm for binary modes, this chapter summarizes the derivation of a new, generalized, non-stochastic modal algorithm with two goals in mind. The first goal is to not restrict the AO system design to specific hardware and a specific basis (*i.e.*, only $2^\ell \times 2^\ell$ segmented DMs correcting Walsh modes). In particular, given the dominance of tilt in turbulence-induced aberrations, the Walsh basis is likely not the ideal choice for turbulence compensation ($N = \ell$ modes are needed to represent a discretized tilt fully). In addition, the basis is incompatible with continuous membrane DMs, which are very common in practical AO experiments. Finally, the second goal is to obtain a generalized form that can be used not only with step dither (as in SPGD and Wang's approach) but also with sinusoidal dither (as in the RF-inspired algorithms).

4.1 Relationship for Arbitrary Bases

As discussed in Sec. 3.2.2, the core element of direct modal estimation is a deterministic relationship between the modal coefficients and the RSSI measurements. With the binary

phase modes, the unique cyclic property of

$$a_i \mathcal{W}_i(x, y) = [a_i + k\pi] \mathcal{W}_i(x, y), \quad k \in \mathbb{Z}, \quad (4.1)$$

leads to a simple periodic relationship (Eq. 3.12) without local maxima. Unfortunately, with arbitrary non-binary modes, such simple relationships typically cannot be obtained. One can, however, use first-order relationships under the assumptions of small residual wavefront errors (*i.e.*, in tracking mode). A common choice in first-order aberration analysis is the exponential form of Maréchal’s approximation of the SR (also referred to as Mahajan’s approximation, or extended Maréchal’s approximation):

$$\text{SR} \approx e^{-\Phi_{\text{rms}}^2}, \quad (4.2)$$

where the SR is typically defined as the normalized on-axis intensity (*i.e.*, a ratio of the on-axis point spread function peak in an aberrated and an aberration-free wavefront), and Φ_{rms} is the RMS wavefront error in radians. Apart from being a good first-order estimate of imaging quality, this relationship has also been shown useful in the approximation of SMF coupling efficiency for FSO receivers in atmospheric turbulence [78].

4.1.1 Individual Perturbations

Considering a flat wavefront perturbed by a single mode $a_i \mathcal{M}_i$ and assuming \mathcal{M}_i has a unit variance, it follows that $\Phi_{\text{rms}} = a_i$. To approximate the RSSI measurement S we can therefore use a modification of Eq. 4.2:

$$S \approx b e^{-a_i^2}, \quad (4.3)$$

where b is a constant that defines the peak RSSI. Fig. 4-1 examines the mean approximation error of this equation across the first 50 modes from both the Walsh (Fig. 4-1a) and the normalized Zernike bases (Fig. 4-1b, both of which have modes with unit variances by definition). We observe that for $|a_i| \leq 0.5$ rad, which can be considered a tracking-like range (with less than 1 dB of loss), the approximation errors are typically below 2%.

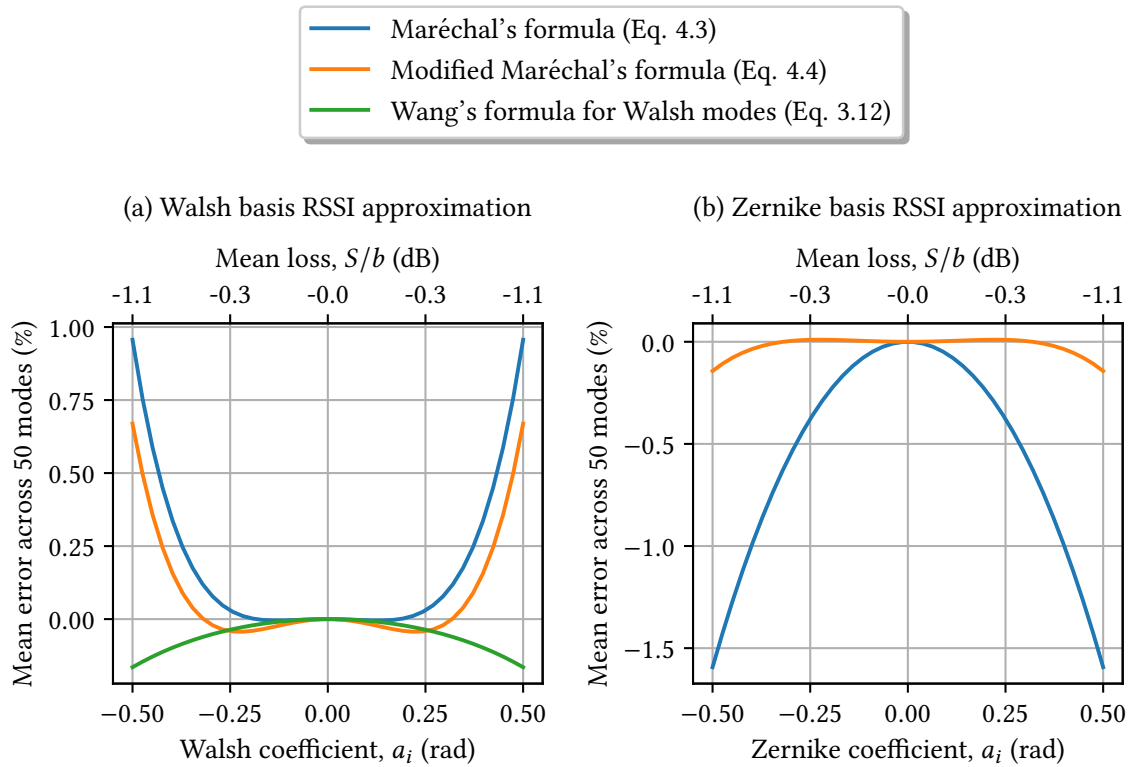


Figure 4-1: Numerical evaluation of how different formulas approximate the relationship between RSSI measurements and modal coefficients. The true RSSI values are first calculated using an overlap integral between $a_i \mathcal{M}_i$ and the fiber mode. The mean RSSI approximation errors are then evaluated across the first 50 Walsh modes (left) and the first 50 Zernike modes (right). The percentage error with respect to the true RSSI is depicted, averaged over all 50 modes. The top axis also shows the mean loss resulting from the respective perturbations (to emphasize a tracking-like range).

While it is convenient to have a single relationship for any mode \mathcal{M}_i , there are a few benefits in per-mode calibration. It has been shown that even with modes normalized to unit variance, the approximation accuracy of Eq. 4.2 does vary with their spatial structure, with higher spatial frequency modes generally having better results [79, 80]. In addition, in a practical AO system, the fit between the actuator influence functions (IFs) and the corrected modes should be taken into account. Depending on the shape of the IF and the correction basis used, different modes may be reproduced with different errors. To help mitigate these effects, the relationship may be augmented with a simple weighing constant. We refer to this as the modified Maréchal's formula:

$$S \approx be^{-w_i a_i^2}, \quad (4.4)$$

where the weight w_i is found experimentally through a fit on RSSI measurements. A simple way to perform this calibration is to sweep the desired modes $a_i \mathcal{M}_i$ with the specific DM on top of a flat reference wavefront. As seen in Fig. 4-1b, in particular for the non-binary basis, the calibration can lead to an improvement in the approximation by several percent, with the mean error typically being smaller than 0.2%.

4.1.2 Multiple Perturbations

While Eqs. 3.12 and 4.4 give optimal results when only one mode is present, the wavefront is generally perturbed by a random superposition of multiple modes during dynamic turbulence compensation [16, 81]. This is the case especially during AO acquisition, but also during tracking, partly due to the limitations of closed-loop rejection. In such a scenario, Φ_{rms}^2 may be expressed as a sum of the modal variances:

$$\Phi_{\text{rms}}^2 = \sum_i a_i^2, \quad (4.5)$$

and the approximation becomes

$$\begin{aligned} S &\approx b e^{-\sum_i a_i^2} \\ &\approx b \prod_i e^{-a_i^2}. \end{aligned} \quad (4.6)$$

Finally, considering per-mode calibration, we get

$$S \approx b \prod_i e^{-w_i a_i^2}. \quad (4.7)$$

If a single mode \mathcal{M}_i is adjusted while the coefficients $a_{j \neq i}$ remain unchanged, we can treat their contribution as a constant:

$$\begin{aligned} S &\approx b \left(\prod_{j \neq i} e^{-w_j a_j^2} \right) e^{-w_i a_i^2} \\ &\approx b_i e^{-w_i a_i^2}, \end{aligned} \quad (4.8)$$

where b_i is a function of all the coefficients but a_i , exactly as in Eq. 3.12.

To quantify the validity of Eq. 4.8, we numerically evaluate its error for a few values of a_i as a function of Φ_{rms} . The aberrated wavefronts are first generated from uniform residuals $a_{j \neq i}$ with random signs as:

$$\Phi_i(x, y) = \frac{\Phi_{\text{rms}}}{\sqrt{N}} \sum_{j \neq i}^N \pm \mathcal{M}_j(x, y), \quad (4.9)$$

where, as in the previous analysis, we use the first $N = 50$ modes from both the Zernike and the Walsh basis. The true S is then calculated for the given a_i based on $\Phi(x, y) = \Phi_i(x, y) + a_i \mathcal{M}_i(x, y)$, and compared with the approximation from Eq. 4.8. We repeat this process for each mode and evaluate the mean error across N . As seen in Fig. 4-2, under larger Φ_{rms} , the error can increase by several percent. The relationship will thus work best in tracking mode, when Φ_{rms} is actively minimized. On the other hand, during the AO acquisition phase, the errors may necessitate more than one estimation and correction cycle per mode to reach full compensation (as also observed with Wang's formula [61]).

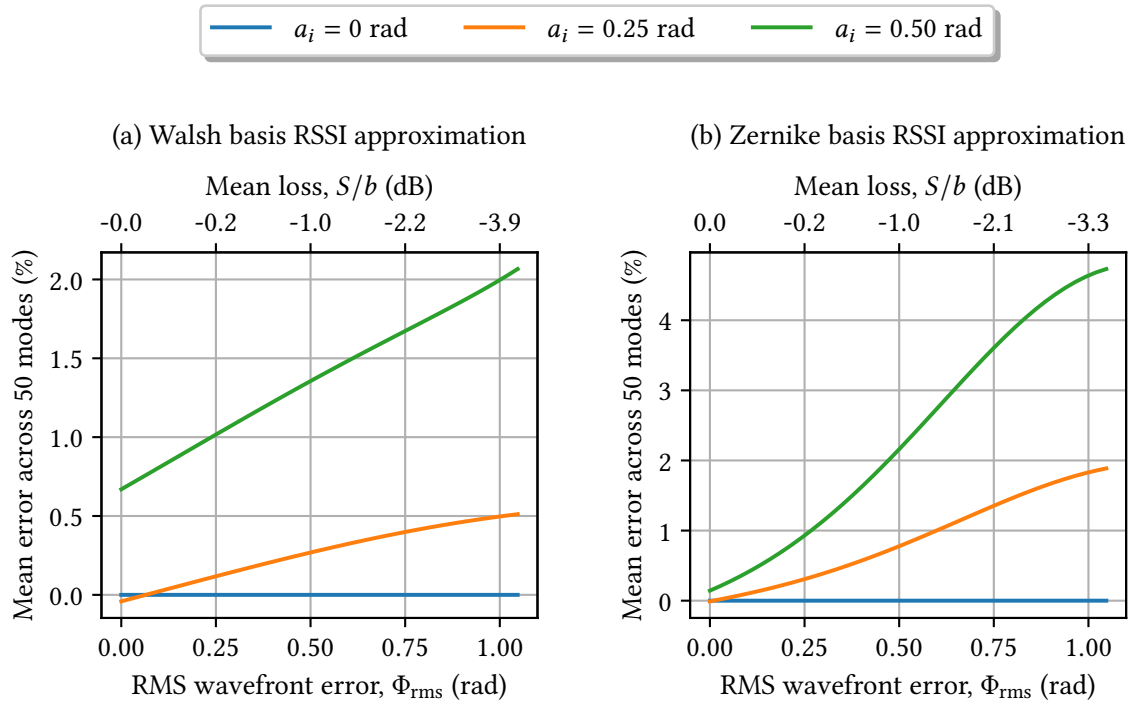


Figure 4-2: The error of Eq. 4.8 in increasing residual wavefront error, which is simulated via a randomly generated multi-mode perturbation such that $\Phi_{\text{rms}}^2 = \sum_{j \neq i} a_j^2$. The mean approximation error across the first 50 modes is shown for a few values of a_i , which helps visualize the possible variations in estimation performance between the AO acquisition and tracking mode. This initial estimation error typically contributes to the need to perform more than one iteration to reach full modal correction (during acquisition).

4.1.3 Logarithmic Measurements

The exponential relationship becomes particularly convenient when RSSI measurements in decibels are available, for instance, as a byproduct of having a logarithmic amplifier in the detection electronics, or through signal processing. Considering a relative power measurement

$$S^{\text{dB}} = 10 \log_{10} S, \quad (4.10)$$

we can switch the base in Eq. 4.8:

$$\begin{aligned} S &\approx b_i e^{-w_i a_i^2} \\ &\approx b_i 10^{-w_i a_i^2 / \ln 10}, \end{aligned} \quad (4.11)$$

and plug into Eq. 4.10 to obtain

$$\begin{aligned} S^{\text{dB}} &\approx 10 \log_{10} (b_i 10^{-w_i a_i^2 / \ln 10}) \\ &\approx b_i^{\text{dB}} - \frac{10}{\ln 10} w_i a_i^2, \end{aligned} \quad (4.12)$$

where b_i^{dB} represents the peak RSSI in units of dBm or similar. To further simplify the equation, we can also define a new weight

$$c_i = \frac{10 w_i}{\ln 10}, \quad (4.13)$$

to finally obtain

$$S^{\text{dB}} \approx b_i^{\text{dB}} - c_i a_i^2. \quad (4.14)$$

In contrast to Eqs. 4.8 and 3.12, the logarithmic RSSI measurements give a computationally simpler relationship without loss of approximation accuracy and enable calibration of the weights w_i through a quadratic fit.

4.2 Coefficient Estimation

To enable correction of the unknown coefficients a_i , we first need to estimate them from measurements of RSSI. In this section, we use the obtained relationships (Eqs. 4.8 and 4.14) to solve for a_i based on changes in $S(t)$ driven by modal dither of some magnitude ϕ_i . Note that this analysis still considers RSSI to be noise-free for the sake of simplicity (we focus on noise sensitivity in Chapter 5 as part of higher-fidelity dynamic simulations).

4.2.1 Step Dither

Like in SPGD and Wang's algorithm, dither injection may be performed as a step change of an unknown mode \mathcal{M}_i . If we continue with the assumption from Chapter 3 that the wavefront remains correlated while the RSSI measurement settles (*i.e.*, $\tau_m \ll \tau_{\text{atm}}$, where τ_{atm}^{-1} is commonly tens to hundreds of Hz, while τ_m^{-1} can be tens to hundreds of kHz), we can express the ratio of the pre- and post-dither RSSI, S_0 and S_1 , via Eq. 4.8 as follows:

$$\frac{S_0}{S_1} = \frac{b_i e^{-w_i a_i^2}}{b_i e^{-w_i (a_i - \phi_i)^2}} \quad (4.15)$$

$$= e^{-w_i [a_i^2 - (a_i - \phi_i)^2]}, \quad (4.16)$$

where ϕ_i is the magnitude of the injected dither on top of the unknown coefficient a_i in radians. Taking the natural logarithm of each side, we get

$$\begin{aligned} \ln \frac{S_0}{S_1} &= -w_i [a_i^2 - (a_i - \phi_i)^2] \\ &= w_i \phi_i^2 - 2w_i a_i \phi_i, \end{aligned} \quad (4.17)$$

which allows us to solve for the unknown coefficient

$$\begin{aligned} a_i &= \frac{w_i \phi_i^2 - \ln(S_0/S_1)}{2w_i \phi_i} \\ &= w_{1,i} [w_{2,i} - \ln(S_0/S_1)], \end{aligned} \quad (4.18)$$

where we define new estimation weights $w_{1,i} = (2w_i \phi_i)^{-1}$ and $w_{2,i} = (w_i \phi_i^2)$.

If measurements in dB are available, the intensities are (based on Eq. 4.14)

$$S_0^{\text{dB}} = b_i^{\text{dB}} - c_i a_i^2, \quad (4.19)$$

$$S_1^{\text{dB}} = b_i^{\text{dB}} - c_i (a_i - \phi_i)^2. \quad (4.20)$$

We can take the difference to get an estimation equation of a similar form:

$$\begin{aligned} \Delta S^{\text{dB}} &= S_0^{\text{dB}} - S_1^{\text{dB}} \\ &= c_i \phi_i^2 - 2c_i a_i \phi_i, \end{aligned} \quad (4.21)$$

and finally solve for the unknown

$$\begin{aligned} a_i &= \frac{c_i \phi_i^2 - \Delta S^{\text{dB}}}{2c_i \phi_i} \\ &= c_{1,i} [c_{2,i} - \Delta S^{\text{dB}}], \end{aligned} \quad (4.22)$$

where the analogous log-estimation weights are $c_{1,i} = (2c_i \phi_i)^{-1}$ and $c_{2,i} = (c_i \phi_i^2)$.

4.2.2 Sinusoidal Dither

Similar to nutation tracking and CBC algorithms based on zonal sinusoidal dither, we seek to derive an estimator based on modal sinusoidal dither to enable the RF-inspired correction methods. In this case, dither injection is not in the form of a step change but rather a continuous sinusoidal perturbation of the given mode \mathcal{M}_i with amplitude ϕ_i at some frequency $\omega_i \gg \tau_{\text{atm}}^{-1}$.

With logarithmic measurements of RSSI, solving for a_i based on the detected AM at ω_i turns out to be straightforward. Extending Eq. 4.14 with the sinusoidal coefficient dither,

we get

$$\begin{aligned}
S^{\text{dB}}(t) &= b_i^{\text{dB}} - c_i (a_i + \phi_i \sin \omega_i t)^2 \\
&= b_i^{\text{dB}} - c_i (a_i^2 + 2a_i \phi_i \sin \omega_i t + \phi_i^2 \sin^2 \omega_i t) \\
&= b_i^{\text{dB}} - c_i \left(a_i^2 + 2a_i \phi_i \sin \omega_i t + \frac{1}{2} \phi_i^2 - \frac{1}{2} \phi_i^2 \cos 2\omega_i t \right). \tag{4.23}
\end{aligned}$$

Isolating the relevant frequency component of RSSI, $S_{\omega_i}^{\text{dB}}$, we can directly solve for a_i :

$$S_{\omega_i}^{\text{dB}} = -2c_i a_i \phi_i, \tag{4.24}$$

$$\begin{aligned}
a_i &= -\frac{S_{\omega_i}^{\text{dB}}}{2c_i \phi_i} \\
&= -c_{1,i} S_{\omega_i}^{\text{dB}}. \tag{4.25}
\end{aligned}$$

In practice, $S_{\omega_i}^{\text{dB}}$ is obtained via coherent demodulation of RSSI at ω_i (e.g., by digital down-conversion and filtering), as in a standard AM receiver. In addition, due to the quadratic response, a harmonic will be present at $2\omega_i$ with an amplitude given by

$$\begin{aligned}
S_{2\omega_i}^{\text{dB}} &= \frac{1}{2} c_i \phi_i^2 \\
&= \frac{1}{2} c_{2,i}. \tag{4.26}
\end{aligned}$$

On the other hand, with linear measurements, Eq. 4.8 by itself does not give a simple closed-form solution for a_i . However, a good approximation can be obtained from its Taylor expansion:

$$\begin{aligned}
S &\approx b_i e^{-w_i a_i^2} \\
&\approx b_i \left[1 - w_i a_i^2 + \mathcal{O}(a_i^4) \right] \\
&\approx b_i \left(1 - w_i a_i^2 \right). \tag{4.27}
\end{aligned}$$

In contrast to the standard Maréchal's approximation (Eq. 4.2), the form of Eq. 4.27 is sometimes referred to as the *simplified* Maréchal's approximation [80, 82], and assuming

no per-mode calibration (*i.e.*, $w_i = 1$), this form is equivalent for both Maréchal's and Wang's formulas. Adding in the sinusoidal coefficient dither, we get

$$\begin{aligned}
S(t) &= b_i \left[1 - w_i (a_i + \phi_i \sin \omega_i t)^2 \right] \\
&= b_i \left[1 - w_i (a_i^2 + 2a_i \phi_i \sin \omega_i t + \phi_i^2 \sin^2 \omega_i t) \right] \\
&= b_i \left[1 - w_i \left(a_i^2 + 2a_i \phi_i \sin \omega_i t + \frac{1}{2} \phi_i^2 - \frac{1}{2} \phi_i^2 \cos 2\omega_i t \right) \right]. \tag{4.28}
\end{aligned}$$

Isolating the relevant frequency component of RSSI, S_{ω_i} , we can express a_i as follows:

$$S_{\omega_i} = -2b_i w_i a_i \phi_i, \tag{4.29}$$

$$\begin{aligned}
a_i &= -\frac{S_{\omega_i}}{2b_i w_i \phi_i} \\
&= -\frac{w_{1,i} S_{\omega_i}}{b_i}. \tag{4.30}
\end{aligned}$$

While this result is similar to its logarithmic variant seen in Eq. 4.25, it has some drawbacks. Most importantly, since a_i is a function of b_i , another measurement is needed to estimate b_i . One option is to measure it via the harmonic at $2\omega_i$, which goes as

$$\begin{aligned}
S_{2\omega_i} &= \frac{1}{2} b_i w_i \phi_i^2 \\
&= \frac{1}{2} b_i w_{2,i}. \tag{4.31}
\end{aligned}$$

Another option is to utilize a low-frequency measurement at some ω_0 , assuming

$$\tau_{\text{atm}}^{-1} \ll \omega_0 \ll \omega_i, \tag{4.32}$$

where any variation in a_i due to turbulence or modulation is negligible. Collecting such terms from Eq. 4.28 and assuming $|a_i| \ll 1$, we get

$$\begin{aligned}
S_{\omega_0} &= b_i \left(1 - w_i a_i^2 - \frac{1}{2} w_i \phi_i^2 \right) \\
&\approx b_i \left(1 - \frac{1}{2} w_{2,i} \right), \tag{4.33}
\end{aligned}$$

which also enables estimation of b_i .

However, even if b_i is known, the estimation suffers from additional approximation error due to the simplification used in Eq. 4.27. Although this is negligible while tracking (*i.e.*, when $|a_i| \ll 1$), first-order analyses show it can increase the estimation error by several percent during initial AO acquisition. To avoid these complexities, it may be preferable to utilize logarithmic measurements and the simple linear coefficient estimation via Eq. 4.25.

4.2.3 Quadrature Dither

The sinusoidal dither can be extended to include a quadrature component, allowing estimation of two modes per frequency with an estimator akin to a QAM receiver. For instance, expanding Eq. 4.14 as a function of two modal coefficients a_i and a_j , we get

$$S^{\text{dB}} \approx b_{ij}^{\text{dB}} - c_i a_i^2 - c_j a_j^2. \quad (4.34)$$

Injecting two orthogonal sinusoidal dithers at a common frequency ω_{ij} with amplitudes ϕ_i and ϕ_j , the fiber RSSI will go as

$$\begin{aligned} S^{\text{dB}}(t) &= b_{ij}^{\text{dB}} - c_i (a_i + \phi_i \sin \omega_{ij} t)^2 - c_j (a_j + \phi_j \cos \omega_{ij} t)^2 \\ &= b_{ij}^{\text{dB}} - c_i \left(a_i^2 + 2a_i \phi_i \sin \omega_{ij} t + \frac{1}{2} \phi_i^2 - \frac{1}{2} \phi_i^2 \cos 2\omega_{ij} t \right) \\ &\quad - c_j \left(a_j^2 + 2a_j \phi_j \cos \omega_{ij} t + \frac{1}{2} \phi_j^2 + \frac{1}{2} \phi_j^2 \cos 2\omega_{ij} t \right), \end{aligned} \quad (4.35)$$

which permits estimation of both coefficients from the IQ components at ω_{ij} :

$$a_i = -c_{1,i} S_{I,\omega_{ij}}^{\text{dB}}, \quad (4.36)$$

$$a_j = -c_{1,j} S_{Q,\omega_{ij}}^{\text{dB}}. \quad (4.37)$$

In this case, the amplitude of the second harmonic is

$$S_{2\omega_i}^{\text{dB}} = \frac{1}{2} (c_i \phi_i^2 - c_j \phi_j^2), \quad (4.38)$$

where it is evident that if the ratio of amplitudes is set appropriately, the harmonic will vanish. Specifically, if we maintain

$$\frac{\phi_i}{\phi_j} = \sqrt{\frac{c_j}{c_i}} = \sqrt{\frac{w_j}{w_i}}, \quad (4.39)$$

then $S_{2\omega_i}^{\text{dB}} \rightarrow 0$, which may unlock more interference-free spectrum for multiplexing methods such as frequency-division.

4.3 Multimodal Correction

For optimal atmospheric turbulence compensation, a sufficient number of modes must be corrected at an adequate rate to minimize the spatiotemporal wavefront errors. Although the spatial bandwidth of commercially available DMs supports the correction of hundreds to thousands of modes [83], the number of modal coefficients that can be estimated via dither at a given rate f_{est} is practically constrained by the DM temporal bandwidth f_{dm} , which is usually on the order of a few to tens of kHz. To ensure that an optimal number of modes is estimated based on the spatiotemporal characteristics of the turbulence, a few different architectures may be utilized.

4.3.1 Time-Division

A simple architecture for simultaneous estimation and correction of multiple modes is TD, where each mode is addressed within a short time bin. This approach can be utilized both with step dither (as suggested by Wang) and sinusoidal or quadrature dither (similar to TDMA in RF communications).

As shown on top of Fig. 4-3, with step dither, each bin is split into a dither and a correction phase. Assuming a first-order actuator response with a time constant $\tau_{\text{dm}} = (2\pi f_{\text{dm}})^{-1}$ and a negligible RSSI latency, the measurement can be triggered in

$$\tau_m = \alpha \tau_{\text{dm}}, \quad (4.40)$$

where α defines how many DM time constants are spent waiting for a sufficiently settled actuator response. Considering periodic estimation and correction of N modes, the estimation period of any given mode is

$$T_{\text{est}} = 2N\tau_m \quad (4.41)$$

$$= 2N\alpha\tau_{\text{dm}}, \quad (4.42)$$

or, equivalently,

$$f_{\text{est}} = \frac{f_{\text{dm}}}{4\pi\alpha N} \quad (4.43)$$

$$= \frac{f_{\text{dm}}}{\beta N}, \quad (4.44)$$

where f_{dm}/β represents the effective estimation bandwidth shared among the N modes.

On the other hand, with sinusoidal dither, the perturbation is active throughout the whole time bin. As seen at the bottom of Fig. 4-3, once an estimate of the coefficient is obtained from the demodulator, a correction is applied while the mode being dithered is switched. The demodulation time T_{bin} is primarily driven by the dither frequency ω and filtering overhead. Since ω is inherently constrained by f_{dm} , we can express the attainable estimation update rate again as

$$f_{\text{est}} = \frac{f_{\text{dm}}}{\beta N}, \quad (4.45)$$

where β is the ratio between f_{dm} and the effective demodulation bandwidth, similar to Eq. 4.43. Finally, if quadrature dither is implemented, each bin can be shared by two modes, which doubles the estimation rate to

$$f_{\text{est}}^{\text{IQ}} = \frac{2f_{\text{dm}}}{\beta N}. \quad (4.46)$$

In practice, implementation complexity is an important factor in choosing between the dither methods and optimizing the efficiency-related constant β . The step dither approach is more straightforward from an implementation perspective since the injection and measurement can occur at a lower sampling rate $f_s = \tau_m^{-1} < f_{\text{dm}}$. However, low-

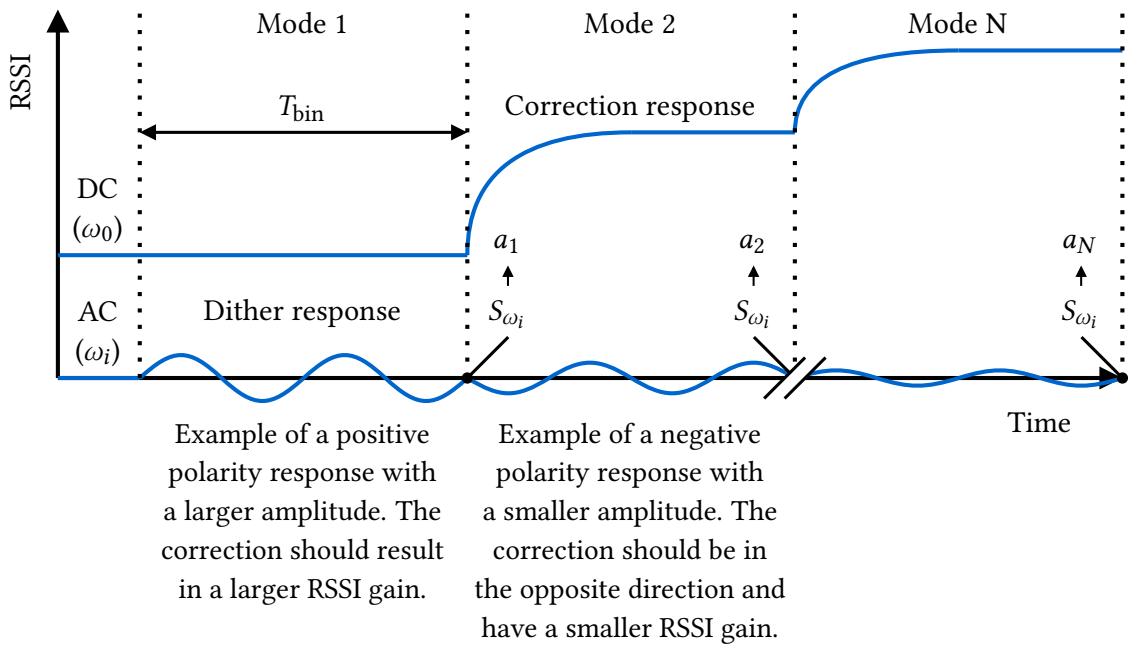
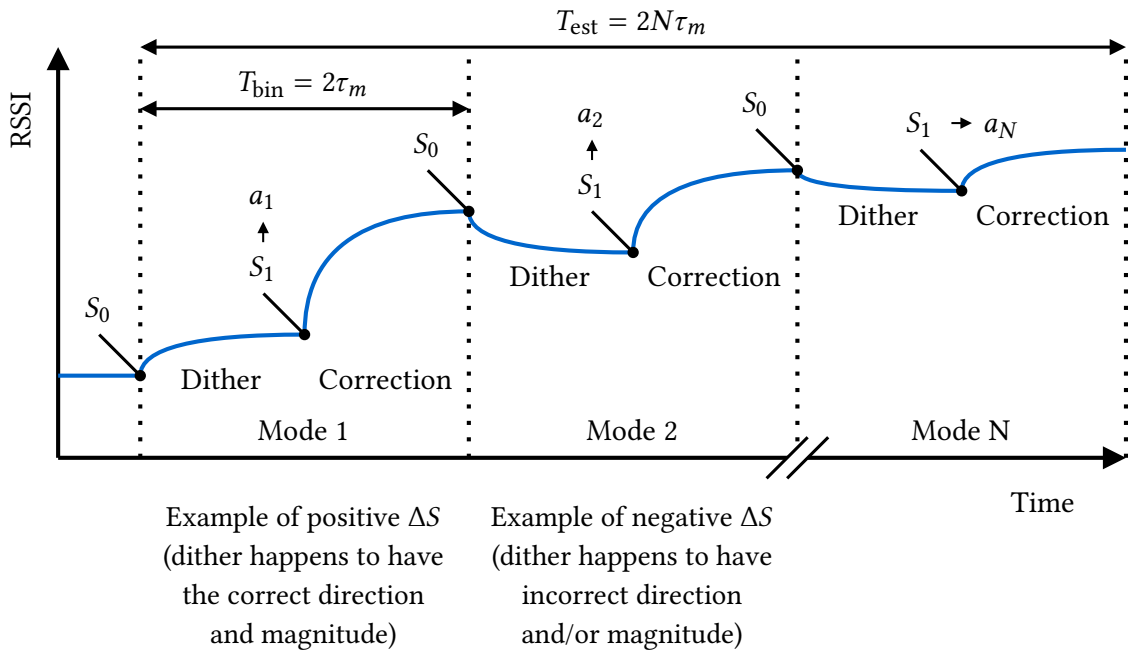


Figure 4-3: A high-level diagram of the estimation and correction flow with TD with step dither (top) and sinusoidal dither (bottom). An illustration of what the RSSI may look like during a short cycle of dithers and corrections is shown. Note that the graphs are dimensionless and highly simplified for the sake of simple illustration.

frequency measurement noise and the $4\pi\alpha$ factor in Eq. 4.43 may reduce the performance. While α can be reduced via loop shaping methods (e.g., by adding a pre-emphasis impulse to the DM control inputs), this is also at the expense of a faster, more complex actuator drive.

In contrast, estimation via sinusoidal dither requires both higher-rate dither injection and RSSI measurement (i.e., f_s may be larger than f_{dm} due to the oversampling associated with sinusoidal signals). In addition, the smaller β is, the more distorted the dither signal gets by the DM frequency response, which necessitates more stringent calibration. This may be especially challenging with quadrature dither, where full amplitude and phase recovery are needed. However, given the higher spectral efficiency from Eq. 4.46, a balance between desired performance and implementation complexity should be considered.

4.3.2 Frequency-Division & Other Methods

Another RF-inspired architecture that can be utilized for mode multiplexing is frequency-division (FD). In this case, each mode is assigned a bin in the frequency domain, and the DM is continuously perturbed by a sum of multiple sinusoidal dithers. Since each mode has its demodulator, the coefficients can be estimated and corrected simultaneously. However, as visualized in Fig. 4-4, the effective demodulation bandwidth is constrained by both f_{dm} and the number of frequency bins used. Hence, Eq. 4.45 is still applicable in terms of estimation rate. In this case, β is equal to the ratio of f_{dm} to the total available modulation bandwidth, which is driven by the maximum frequency ω_N and the bin filtering overhead (this is explored in more depth in Chapter 5). Finally, as in Eq. 4.46, the estimation rate can be doubled with quadrature dither on each frequency.

While the implementation of FD is more challenging than that of TD due to the complexities associated with multi-frequency dither injection and parallel demodulation, it does offer some advantages. Since each injector-demodulator block is an independent unit, the dither magnitude and estimation bandwidth allocation can be optimized in real-time for each mode based on the atmospheric conditions without the need to adjust a switching sequence. Additionally, in contrast to sinusoidal TD, which necessitates a separate demodulation and correction loop running at different rates, these two can be




Figure 4-4: A high-level diagram of multimodal estimation via FD. The dotted curve at baseband illustrates the spectrum of the atmospheric turbulence, which is decomposed into N modes dithered at different frequencies ω_i , as illustrated by the vertical arrows. The maximum dither frequency ω_N is constrained by the DM temporal bandwidth f_{dm} . Note that the graph is dimensionless and highly simplified for visual clarity.

unified in FD given that the coefficients are estimated continuously.

Finally, apart from TD and FD, a few other methods may be considered to divide the DM bandwidth among multiple modes. Code-division (CD) with a common sinusoidal dither (akin to CDMA) is a technique that has already been demonstrated with zonal algorithms for CBC (Sec. 3.3) and applies to modal multiplexing in the same fashion. For instance, one could use a set of the one-dimensional Walsh functions (Sec. 3.2.1) in the time domain to code different modes on a single dither frequency. In this case, the temporal switching would not be used to adjust modes directly but rather to reproduce the codewords, which also spreads the spectrum for better utilization of f_{dm} . In addition, hybrid architectures that combine FD with spectrum spreading via TD or CD may also be considered, however, these are outside the scope of this thesis.

4.3.3 Error Rejection

Regardless of the estimation method used, a feedback control mechanism is required to reject the modes from the wavefront as they are estimated. Here, we consider a basic parallel integration scheme, where the correction $u_i(t)$ for each mode is driven by the integral of the respective estimated coefficients $a_i(t)$. The standard form of such an integral

controller gives

$$u_i(t) = -\frac{g}{T_i} \int_0^t a_i(t) dt, \quad (4.47)$$

where g is a dimensionless controller gain and T_i is the integrator time constant. A discretized variant of the equation may be written as

$$u_i(k) = -g \frac{T_s}{T_i} \sum_{l=0}^k a_i(k-l) \quad (4.48)$$

$$= u_i(k-1) - g \frac{T_s}{T_i} a_i(k), \quad (4.49)$$

where k is the iteration or sample number and $T_s = f_s^{-1}$ is the control loop sampling time. In our case, since we are constrained by f_{est} , we assume

$$u_i(k) = -g \frac{f_{\text{est}}}{f_s} \sum_{l=0}^k a_i(k-l) \quad (4.50)$$

$$= u_i(k-1) - g \frac{f_{\text{est}}}{f_s} a_i(k). \quad (4.51)$$

Finally, if f_s and f_{est} are equivalent (e.g., in TD, where the sampling and correction is synchronous), the equation simplifies to

$$u_i(k) = -g \sum_{l=0}^k a_i(k-l) \quad (4.52)$$

$$= u_i(k-1) - ga_i(k), \quad (4.53)$$

in which case g can also be thought of as a damping factor for the instantaneous coefficient corrections. In an ideal noiseless and latency-free system, a unity gain may be assumed (e.g., as in Wang's algorithm, described in Alg. 2). However, in practice, the hardware imperfections must be considered while optimizing g to ensure a stable control loop response.

Chapter 5

Simulation Studies

To enable quick assessment of how various configurations of the algorithm perform in different atmospheric conditions, we developed an end-to-end software simulator that models a simplified FSOC receiver with AO for turbulence compensation. This chapter covers the simulation environment, the learnings from initial algorithm testing, and finally, a comprehensive algorithm performance analysis across a broader set of environmental and hardware parameters.

5.1 Simulation Environment

The simulation environment is built upon simplified hardware models interacting with a beam corrupted by simulated atmospheric turbulence. Distinctions between the core parts of the environment are visualized in Fig. 5-1. The high-level building blocks include a DM model, a receiver model, an aberrated wavefront generator, and a digital signal processing (DSP) block with the algorithm logic described in the previous chapter. Apart from the aberrated wavefront generator, which is based on a third-party tool [84] and further described in Sec. 5.1.3, all models are written from the ground up using the Python programming language with the help of open-source simulation packages such as NumPy [85], Numba [86], and SciPy [87]. This section gives more detail on each of the core simulation building blocks.

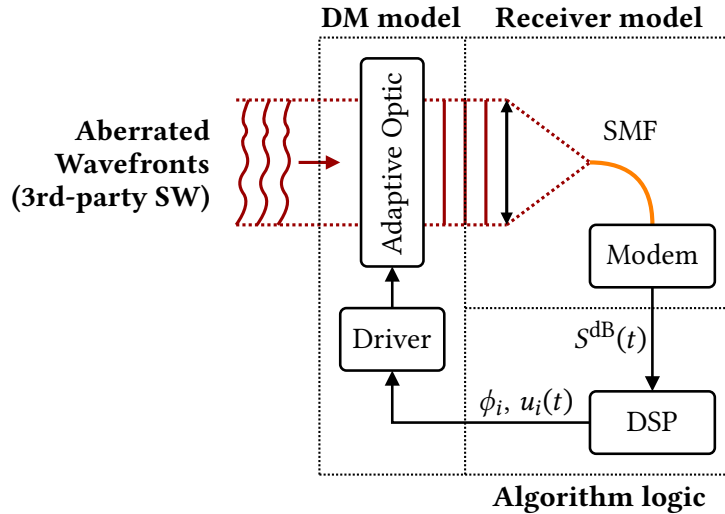


Figure 5-1: The primary building blocks of the algorithm simulation environment.

5.1.1 DM Model

To enable testing of the algorithm with both binary and continuous modes, we implement a configurable DM model that can act either as a segmented or as a continuous membrane DM with $M \times M$ actuators. The model is divided into three layers: a driver, an actuator, and a device layer. A high-level description of the layers, including the effects they model, is given in Table 5.1, and some typical values of the model parameters (including rationale) are summarized in Table 5.2.

The purpose of the driver layer is to simulate the interface between the DSP logic and the analog signals controlling the actuators. The primary effects considered are limited digital-to-analog converter (DAC) resolution and slight output variations between the individual DACs. The former is addressed by quantizing the actuator stroke to a configurable number of bits, and the latter by adding a random full-scale and bias error to each DAC output, given as a percentage of the nominal range. However, for the sake of model simplicity, a few aspects of the driver are not considered in the simulator. We assume the response is linear and significantly faster than the actuator itself, hence the DAC dynamics are not modeled. We also assume the RSSI measurements are the main contributor of noise in the algorithm (we further discuss this in Sec. 5.1.2) and therefore do not model DAC noise (flicker). A higher-fidelity treatment of noise from different sources

Table 5.1: A high-level description of the three DM model layers, with a summary of effects that were modeled versus excluded for the sake of simplicity.

Layer	Description	Modeled effects	Not modeled
Driver	Conversion between digital and analog domains	<ul style="list-style-type: none"> • DSP register sizes • DAC quantization • DAC static variation 	<ul style="list-style-type: none"> • Nonlinearity • DAC dynamics • DAC noise
Actuator	Dynamic response of DM actuators	<ul style="list-style-type: none"> • 1st-order response • Response variation 	<ul style="list-style-type: none"> • Higher-order dynamics
Device	Map between actuator strokes and phase function	<ul style="list-style-type: none"> • Piston segments • Gaussian phase influence functions 	<ul style="list-style-type: none"> • Fill factor (dead zone) • Edge effects

Table 5.2: Commonly used parameters in the different layers of the DM model.

Parameter	Typ. value	Rationale
Number of actuators across aperture, M	8 or 16 (segmented)	Facilitate mapping to 2^l sized Walsh modes (M has to be a power of 2)
	12 (membrane)	Simulate the lab device (BMC Multi-3.5-DM [83])
Number of DAC bits	14	Simulate the lab driver (BMC X-Driver [83])
DAC variation	2%	Assume full-scale and bias errors with an open-loop calibration within a few %
τ_{dm} baseline	35 μs (swept)	Baseline derived from rise time in [83], but swept through a wide range
τ_{dm} variation	2%	Assume fabrication tolerances of a few %
Influence function (IF) full width at half maximum (FWHM)	1.3 actuators	Derived from measurements in [88]

is considered for future work, which is discussed in Chapter 7.

The dynamic response of the actuators themselves is handled in the actuator layer. Given that the actuators are most commonly mechanical devices, it is assumed that their frequency response is dominant and is the bottleneck in the attainable dither and correction rate. Each actuator is modeled by a first-order transfer function with a time constant $\tau_{\text{dm}} = (2\pi f_{\text{dm}})^{-1}$, where f_{dm} is the corner (-3 dB) frequency. To include an option for simulating imperfect fabrication tolerances, a random variation of τ_{dm} among the actuators can be provided in the percentage of the nominal value.

Finally, the purpose of the device layer is to spatially transform the actuator strokes into a higher-resolution phase function that interacts with the received wavefront. With the segmented DM variant, the strokes are mapped to $M \times M$ square piston segments. In contrast, to simulate a membrane DM, each actuator's influence on the wavefront is modeled by a Gaussian influence function (IF), where the full width at half maximum (FWHM) defines the inter-actuator coupling.

Both device variants assume a square actuator arrangement for a few reasons. First, the availability of laboratory equipment for hardware validation is limited to a square 12×12 continuous-membrane DM. Hence the model is primarily used to simulate this device. Second, the square aperture facilitates mapping to the Walsh basis for algorithm testing with binary modes. And third, since the aberrated wavefronts are provided on a square grid, and all the computations are performed via square matrices, it is a more computationally and data-efficient approach. Other simplifying assumptions include a perfect fill factor (*i.e.*, no dead zone between segments is modeled), an equal spatial response for the device's edge actuators, and no modeling of diffraction at the segment edges or quilting on the membrane surface.

5.1.2 Receiver Model

The receiver model simulates a fiber optic modem connected via a conventional SMF coupling interface, which is geometrically optimized to maximize the coupling efficiency η between the fiber mode and the DM-corrected wavefront. Using the conventional Gaussian approximation of the fundamental fiber mode, we can express its spatial envelope as

$$E_{\text{SMF}}(x, y) = \sqrt{\frac{2}{\pi\sigma_0^2}} e^{-(x^2+y^2)\sigma_0^{-2}}, \quad (5.1)$$

where σ_0 is the fiber mode field radius. Note that here E_{SMF} is a normalized version of the field such that the total power is a unity for the sake of easier analysis, *i.e.*,

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |E_{\text{SMF}}(x, y)|^2 dx dy = 1. \quad (5.2)$$

To simplify the computation of η , we propagate this field from the fiber core through the coupling optic, which (under ideal optic assumptions) results in an expansion to some $\sigma_1 > \sigma_0$ at the coupling interface:

$$E_{\text{if}}(x, y) = \sqrt{\frac{2}{\pi\sigma_1^2}} e^{-(x^2+y^2)\sigma_1^{-2}}. \quad (5.3)$$

Let $E_{\text{rx}}(x, y)$ be the received DM-corrected wavefront, which is also assumed to be normalized such that its total power is unity. Assuming a square DM aperture (due to reasons given in Sec. 5.1.1), we can express it as

$$E_{\text{rx}}(x, y) = \frac{1}{D} \text{rect}\left(\frac{x}{D}, \frac{y}{D}\right) e^{\chi(x,y)+i\Phi(x,y)}, \quad (5.4)$$

where D is the aperture diameter, χ is the normalized log-amplitude, and Φ is the phasefront variation. If the received field is ideally relayed to the coupling optic, we can calculate η between the two fields via a coherent overlap integral as

$$\eta = \left| \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E_{\text{if}}^*(x, y) E_{\text{rx}}(x, y) dx dy \right|^2. \quad (5.5)$$

To geometrically optimize the interface, we first assume an ideal aberration-free beam is received. Considering $\chi \rightarrow 0$, and $\Phi \rightarrow 0$, we get a flat square wavefront:

$$E_{\text{flat}}(x, y) = \frac{1}{D} \text{rect}\left(\frac{x}{D}, \frac{y}{D}\right). \quad (5.6)$$

The calculation of η then simplifies to

$$\begin{aligned} \eta_{\text{flat}} &= \left| \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E_{\text{if}}^*(x, y) E_{\text{flat}}(x, y) dx dy \right|^2 \\ &= \left| \sqrt{\frac{2}{\pi}} \frac{1}{\sigma_1 D} \int_{-\frac{D}{2}}^{\frac{D}{2}} \int_{-\frac{D}{2}}^{\frac{D}{2}} e^{-(x^2+y^2)\sigma_1^{-2}} dx dy \right|^2 \\ &= 2\pi \left(\frac{\sigma_1}{D}\right)^2 \text{erf}^4\left(\frac{D}{2\sigma_1}\right). \end{aligned} \quad (5.7)$$

This equation has a global maximum at a specific D/σ_1 value determined numerically as

$$\arg \max_{D/\sigma_1} \eta_{\text{flat}} \approx 1.98, \quad (5.8)$$

and a corresponding maximum efficiency

$$\eta_{\text{max}} \approx 0.792, \quad (5.9)$$

$$\eta_{\text{max}}^{\text{dB}} \approx -1.01 \text{ dB}. \quad (5.10)$$

Given this result, we assume the coupling optic is purposefully selected such that $\sigma_1 \approx D/1.98 \approx 0.505 D$, which is equivalent to clipping of the propagated fiber mode's amplitude roughly at the e^{-1} point, considering the DM aperture as the constraint. Consequently, Eq. 5.5 can be expanded into the final form used in the receiver model:

$$\begin{aligned} \eta &\approx \left| \sqrt{\frac{2}{\pi}} \frac{1.98}{D^2} \int_{-\frac{D}{2}}^{\frac{D}{2}} \int_{-\frac{D}{2}}^{\frac{D}{2}} e^{-1.98^2(x^2+y^2)/D^2} e^{\chi(x,y)+i\Phi(x,y)} dx dy \right|^2 \\ &\approx \frac{2.5}{D^4} \left| \int_{-\frac{D}{2}}^{\frac{D}{2}} \int_{-\frac{D}{2}}^{\frac{D}{2}} e^{\chi(x,y)+i\Phi(x,y)-3.92(x^2+y^2)/D^2} dx dy \right|^2. \end{aligned} \quad (5.11)$$

In practical atmospheric conditions, both χ , Φ , and the total power in the received wavefront S_{rx} are functions of time. The instantaneous RSSI, S , at the fiber receiver is then

$$S = \eta S_{\text{rx}}. \quad (5.12)$$

Conventional optical receivers are based on a photodiode (PD) that converts S into a photocurrent, which is detected through a high-speed transimpedance amplifier (TIA). For the purposes of modem RSSI monitoring, it is also common to mirror the photocurrent into a logarithmic amplifier (log amp). Given that logarithmic measurements facilitate the algorithm implementation, we also follow this approach. Since the log amp circuit, including the associated analog-to-digital converter (ADC), typically run at a significantly higher bandwidth than f_{dm} , we do not consider their dynamics in this model. For simplicity, we assume the digital measurements are predominantly corrupted by additive white Gaussian noise stemming from the amplifier, and quantized to a certain number of ADC bits. Since the algorithm logic is invariant to average power, we consider the obtained measurements to be dimensionless and model them as

$$S^{\text{dB}} = 10 \log_{10}(\eta S_{\text{rx}}) + n^{\text{dB}}, \quad (5.13)$$

$$n^{\text{dB}} \sim \mathcal{N}(0, \sigma_n^{\text{dB}}), \quad (5.14)$$

where σ_n^{dB} is the RMS noise after the amplification stage expressed in dB.

To get a sense of what this variable may be in practice, we approximate it for a few values of RSSI based on specifications of an off-the-shelf InGaAs PD and log amp (part numbers from Chapter 6 are used). The results are shown in Table 5.3 for two possible settings of the log amp's dynamic range (DR), and conservatively assume the log amp detection bandwidth is configured to $10 \times f_{\text{dm}}$, which according to Table 5.2 gives $f_{\text{det}} = 10 (2\pi 35 \mu\text{s})^{-1} \approx 45 \text{ kHz}$. We can see that even with an RSSI that is significantly below the sensitivity of common optical transceivers (typically -30 to -40 dBm), the expected noise is very low. This primarily stems from the fact that the bandwidths involved in wavefront control (kHz range) are much smaller than typical communication bandwidths (GHz range).

Table 5.3: Expected amplification noise with an off-the-shelf InGaAs PD and log amp.

RSSI	Photocurrent	Dark current	RMS noise ($10 \times f_{\text{dm}}$)	
			DR=100 dB	DR=25 dB
-50 dBm	10 nA	<50 pA	0.003 dB	0.0008 dB
-60 dBm	1 nA		0.02 dB	0.005 dB
-70 dBm	100 pA		0.21 dB	0.05 dB

5.1.3 Atmospheric Data

As previously shown in Fig. 5-1, the primary dynamic input into the model is a set of aberrated wavefronts which simulate a received beam distorted by atmospheric turbulence. We rely on a validated 3rd-party tool from MZA called WaveTrain [84] to generate this data. The tool is configured to propagate a beam at 1.55 μm through horizontal multi-layer turbulence with a Kolmogorov spectrum and a uniform wind model. The turbulence layers are configured to obtain a desired D/r_0 and Rytov number σ_R^2 at the receiver aperture, while the wind speed is adjusted to control the temporal coherence time τ_{atm} .

To assess wavefront correction across a broader set of atmospheric conditions, we export datasets with various parameter combinations, which are summarized in Table 5.4. In terms of spatial aberration, D/r_0 values between 2 and 5 are considered under approximate weak, medium, and strong scintillation regimes. To simulate different temporal regimes, we fix τ_{atm} and instead adjust τ_{dm} in the simulator to effectively control the actuator-to-atmospheric bandwidth ratio $\tau_{\text{atm}}/\tau_{\text{dm}}$. A dataset with the evolution of the complex electric field at the receiver is exported for each parameter combination. To verify the parameters, an offline post-processing script is also used to determine D/r_0 and τ_{atm} from the e^{-1} points in the respective (ensemble-averaged) spatial and temporal autocorrelation functions. An example of the generated field is shown in Fig. 5-2.




Figure 5-2: Examples of field amplitude and phase in two randomly selected instants exported from WaveTrain. The plots on the left depict a weaker turbulence scenario in terms of D/r_0 and σ_R^2 , while the plots on the right show an example of stronger turbulence. Apart from larger intensity scintillation, the higher $\sigma_R^2 \approx 3$ on the right also leads to more pronounced branch cuts in the phase function, which can complicate wavefront sensing.

Table 5.4: A summary of parameters that characterize the atmospheric datasets exported from WaveTrain.

Wavelength, λ	1.55 μm		
Aperture vs. Fried parameter, D/r_0	2, 3, 4, 5		
Coherence time, τ_{atm}	20 ms		
Approximate Rytov number, σ_R^2	0.75	1.5	3
Irradiance scintillation regime	weak	medium	strong
Temporal & spatial sampling	5 kHz, 128×128 pixels		
Dataset length	1 sec (5k samples)		

5.2 Algorithm Configurations

With the simulation environment prepared for wavefront correction experiments, we focus on selecting and configuring the algorithms to be tested. To compare the non-stochastic algorithm with the state-of-the-art and assess the dither and correction mechanisms outlined in Chapter 4, we test three unique approaches. The first variant combines step dither with TD, which is most easily implemented. At a high level, this approach is similar to Wang’s algorithm but is generalized here to both segmented and membrane DMs. The second variant uses quadrature dither with FD, which is the spectrally-efficient RF-inspired method. Finally, as the third option, we test the optimized modal SPGD variant described in Sec. 3.1.1, which we consider a state-of-the-art reference.

While there is some subtlety in the configuration of each of the three algorithms, at a high level, a few similar parameters drive their behavior. With each variant, one must select the number of corrected modes, the phase dither magnitude, and the feedback gain. For algorithms based on parallel correction, such as FD and SPGD, additional weighting may be considered to optimize the dither magnitudes and gains for each mode. Lastly, with both step and sinusoidal dither, the implementation necessitates some configuration related to the system dynamics. For TD and SPGD, the constant $\alpha = \tau_m/\tau_{\text{dm}}$ (defined in Sec. 4.3.1) needs to be chosen to ensure stable measurements with sufficient settling time. With FD, on the other hand, the maximum dither frequency ω_N has to be chosen to ensure an adequate response and help define the distribution of the frequency bins. This

section briefly describes the effect of each parameter and concludes with a summary of a joint optimization process developed to facilitate the configuration of each algorithm for different conditions.

5.2.1 Number of Modes

Considering the limited bandwidth of DM actuators, the number of corrected modes N is one of the most important parameters that helps balance temporal and spatial wavefront errors. The core issue stems from the fact that with each additional dithered mode, the effective correction rate of all the modes is decreased. To achieve an improvement in the wavefront error, the modes must be corrected within some fraction of the atmospheric coherence time τ_{atm} . Given that the convergence time of each algorithm is constrained by the DM time constant τ_{dm} , it follows, therefore, that N should be some function of the temporal bandwidth ratio $\tau_{\text{atm}}/\tau_{\text{dm}}$ in order to achieve optimal spatiotemporal correction.

5.2.2 Dither Magnitude

Configuration of the dither magnitudes ϕ_i is also key because they not only drive the estimation error, but also inherently contribute to the wavefront error. Considering, for instance, non-stochastic coefficient estimation via TD, we can expand Eq. 4.22 from Chapter 4 with a noisy RSSI measurement to see that

$$\begin{aligned} a_i &\propto \frac{S^{\text{dB}} + n^{\text{dB}}}{2c_i\phi_i} \\ &\propto \frac{S^{\text{dB}}}{2c_i\phi_i} + \frac{n^{\text{dB}}}{2c_i\phi_i}, \end{aligned} \quad (5.15)$$

where the second term is a noise-driven a_i estimation error (in units of radians). It is evident that if ϕ_i is below some detection threshold, the measurements will have insufficient contrast to achieve stable correction, which will lead to an increase in residual wavefront error. The same principle applies to SPGD, where the gradients are determined from a change in RSSI as well. Optimization of ϕ_i is therefore important to minimize the sum of the residual correction error and the error injected due to the dithers themselves. At a

minimum, ϕ_i should be a function of the expected receiver noise magnitude σ_n^{dB} (estimated in Table 5.3) to ensure that the estimation error is below a certain acceptable threshold.

5.2.3 Feedback Gain

The integral feedback gain g (μ for SPGD) is the last knob in the correction chain, and it helps balance the convergence rate and the stability of the closed-loop response. Even if, hypothetically, the obtained estimates were noise-free, hardware effects (*e.g.*, controller latency) may lead to closed-loop stability issues if the gain is too large. In addition, the gain may also be used as a way to reduce the effect of estimation noise (*i.e.*, as a damping factor) without increasing dither magnitude, albeit at the expense of reduced correction bandwidth. This may be useful when τ_{atm} is particularly large and the system is not bandwidth-constrained, or if the estimation noise is too high despite large-magnitude dithers. The latter may occur during deep signal fades or link interruptions, in which case the gain should tend to zero since no useful signal is measured. However, in nominal tracking conditions, a fixed gain is commonly employed.

5.2.4 Dither and Gain Weighting

Given that the spatial power spectrum of turbulence is non-uniform, there is a benefit in weighting the individual dithers and corrections according to the expected ensemble variances $\langle a_i^2 \rangle$ of the respective modes. This principle has already been used in research on SPGD optimization, which is summarized in Sec. 3.1.1. This approach is particularly useful for the non-stochastic algorithm with FD, where dither injection and coefficient correction also occur in parallel. Through weighting, the combined dither magnitude can be more optimally distributed and de-emphasized for modes that are expected to contribute less to the wavefront error, which helps reduce the overall sum of residual and self-injected errors. Likewise, slightly higher feedback gains may help with convergence on the lower-noise (as a result of dither weighting), larger-variance $\langle a_i^2 \rangle$, while smaller gains can be used for the higher-noise, lower-variance modes. To facilitate the testing of different weighting

distributions, we first define a normalized atmospheric weight as

$$\gamma_i = \sqrt{\frac{\langle a_i^2 \rangle}{\langle a_0^2 \rangle}}, \quad (5.16)$$

where the i -th weight is proportional to i -th's mode expected RMS value normalized by the dominant mode. The $\langle a_i^2 \rangle$ values can be obtained either through mathematical models of the turbulence power spectrum or calculated offline by decomposing the aberrated wavefronts into the basis of interest, which is the approach we use here. Once γ_i are determined, we can achieve different weighting distributions for dithers or gains by exponentiating the weights as follows

$$\phi_i = \phi \gamma_i^{e_\phi}, \quad (5.17)$$

$$g_i = g \gamma_i^{e_g}, \quad (5.18)$$

$$\mu_i = \mu \gamma_i^{e_\mu}, \quad (5.19)$$

where ϕ , g , and μ are the baseline dither magnitudes and feedback gains respectively, and $e_{\phi,g,\mu}$ are configurable exponents that allow shifting between more weighted ($e \rightarrow 1$) and more uniform ($e \rightarrow 0$) distributions.

5.2.5 System Dynamics

Apart from parameters inherent to the algorithms, some hardware-specific configuration related to the system dynamics must also be considered. In our model, the frequency response is dominated by the DM actuators, which are modeled with some variation across their time constants τ_{dm} (summarized in Table 5.2), and small errors in their full range and bias. With TD and SPGD, which rely on RSSI measurements after discrete step dithers, an adequate measurement delay τ_m is needed to ensure the modes are sufficiently realized, and the RSSI is stable. We use the parameter $\alpha = \tau_m / \tau_{\text{dm}}$ to define this delay in units of the baseline time constant. Like the feedback gain, this parameter trades stability for bandwidth and should therefore be configured to minimize the total wavefront error resulting from these effects. An analogous parameter with FD, which is based on sinusoidal

dither, is the maximum dither frequency ω_N . On one hand, since the dither amplitude rolls off with frequency, ω_N is fundamentally limited by the range of the DM driver. On the other hand, the variance between each actuator's amplitude and phase response also increases with frequency, which will affect the stability of mode actuation, and further constrain ω_N .

5.2.6 Joint Optimization

Noticing that every parameter can affect the correction performance in some way and that adjusting one parameter may influence the optimal setting of another, we attempt to find a few optimal sets of parameters for different conditions through an offline gradient-based optimization procedure. Each optimization step involves running a full algorithm simulation through one of the 1-second-long atmospheric snapshots (summarized in Table 5.4) with a given configuration. After the run, we evaluate the median of the fiber coupling efficiency $\eta(t)$, which is used as a performance optimization metric. Starting with an initial guess of the parameter values, the parameters are randomly perturbed, and their gradients with respect to the median of $\eta(t)$ are estimated. A correction is then applied to the configuration in the direction of the estimated gradients, with the purpose of maximizing $\eta(t)$. This process continues iteratively until we observe parameter convergence and is then restarted ten extra times to get a sense of its uncertainty. The scale of the perturbations and corrections is determined empirically such that stable convergence occurs within a few thousand iterations.

An example of this joint optimization is visualized in Fig. 5-3 for the TD algorithm. In this particular scenario, we utilize the $M = 8$ (*i.e.*, 8×8 actuators) segmented DM model (with the default model parameters according to Table 5.2) for correction of dominant Walsh modes in $D/r_0 \approx 3$ conditions. The optimization process is started with an initial rough guess of $N = 10$, $\phi = \pi/10$, $g = 0.75$, $\alpha = 1.75$, and is repeated for three different receiver noise assumptions with an increasing value of σ_n^{dB} . We can see that with $\tau_{\text{atm}}/\tau_{\text{dm}} \approx 120$, the value of N converges to 11 in all cases. As is expected, the optimal value of ϕ increases with higher noise, starting at about 0.12π in a nominal noise case (chosen conservatively from Table 5.3), to about 0.16π in an extreme case. In addition, the results confirm that

Table 5.5: A summary of the optimized algorithm parameters that are used in further performance evaluations of the three algorithm variants.

Var.	SPGD	TD	FD	Notes
N	$\lceil 0.77 (\tau_{\text{atm}}/\tau_{\text{dm}})^{0.55} \rceil$		$\lceil 0.58 (\tau_{\text{atm}}/\tau_{\text{dm}})^{0.45} \rceil$	• Empirical fits on batch results
ϕ	$\pi/30$	$\pi/8$	$\pi/30$	• Assuming nominal noise case • Larger in TD (non-parallel alg.)
μ/g	$\mu = 1 \text{ dB}^{-1}$		$g = 0.95$	• SPGD: $\Delta a_i = \mu_i \phi_i \delta J$ (δJ is in dB) • Apart from f_s , controller is the same in TD and FD (Sec. 4.3.3)
e_ϕ	$e_\phi = 0.5$	$e_\phi = 0$	$e_\phi = 0.5$	• SPGD: $\Delta a_i \propto \mu_i \phi_i \propto \gamma_i^{e_\mu + e_\phi}$ ($= \gamma_i^{0.75}$)
$e_{\mu,g}$	$e_\mu = 0.25$	$e_g = 0$	$e_g = 0.4$	• Uniform in TD (non-parallel alg.)
α/ω_N	$\alpha = \tau_m/\tau_{\text{dm}} = 2.5$		$\omega_N \tau_{\text{dm}} = 2$	• FD: $N/2$ uniformly distributed ω_{ij} frequencies with IQ dithers

a smaller g may be beneficial in very noisy environments by reducing the correction uncertainty and avoiding excessively large dithers. The optimal value of α converges to about 2.5, corresponding to roughly 90% settled actuators (considering 1st-order dynamics).

After running this procedure for the other algorithm variants, we compile a baseline set of optimized parameters for each algorithm. In addition, we also run special batches of simulations that sweep N in varying $\tau_{\text{atm}}/\tau_{\text{dm}}$, since the optimal value of N is expected to heavily depend on the bandwidth ratio. These results are evaluated offline to find the N that produces the highest median $\eta(t)$, which is then used to empirically fit the relationship between the two variables. A summary with the results of this fit, as well as the other optimized parameters, is shown in Table 5.5. Apart from N , we do not further adapt the parameters in different simulated conditions. While some of them may benefit from adaptation in specific corner cases (e.g., ϕ and g during a deep fade or link interruption), we are mostly interested in performance during nominal tracking conditions.

Parameter optimization with TD on a segmented DM (Walsh modes)
 Scenario conditions: $D/r_0 \approx 3$, $\sigma_R^2 \approx 0.75$, $\tau_{\text{atm}}/\tau_{\text{dm}} \approx 120$

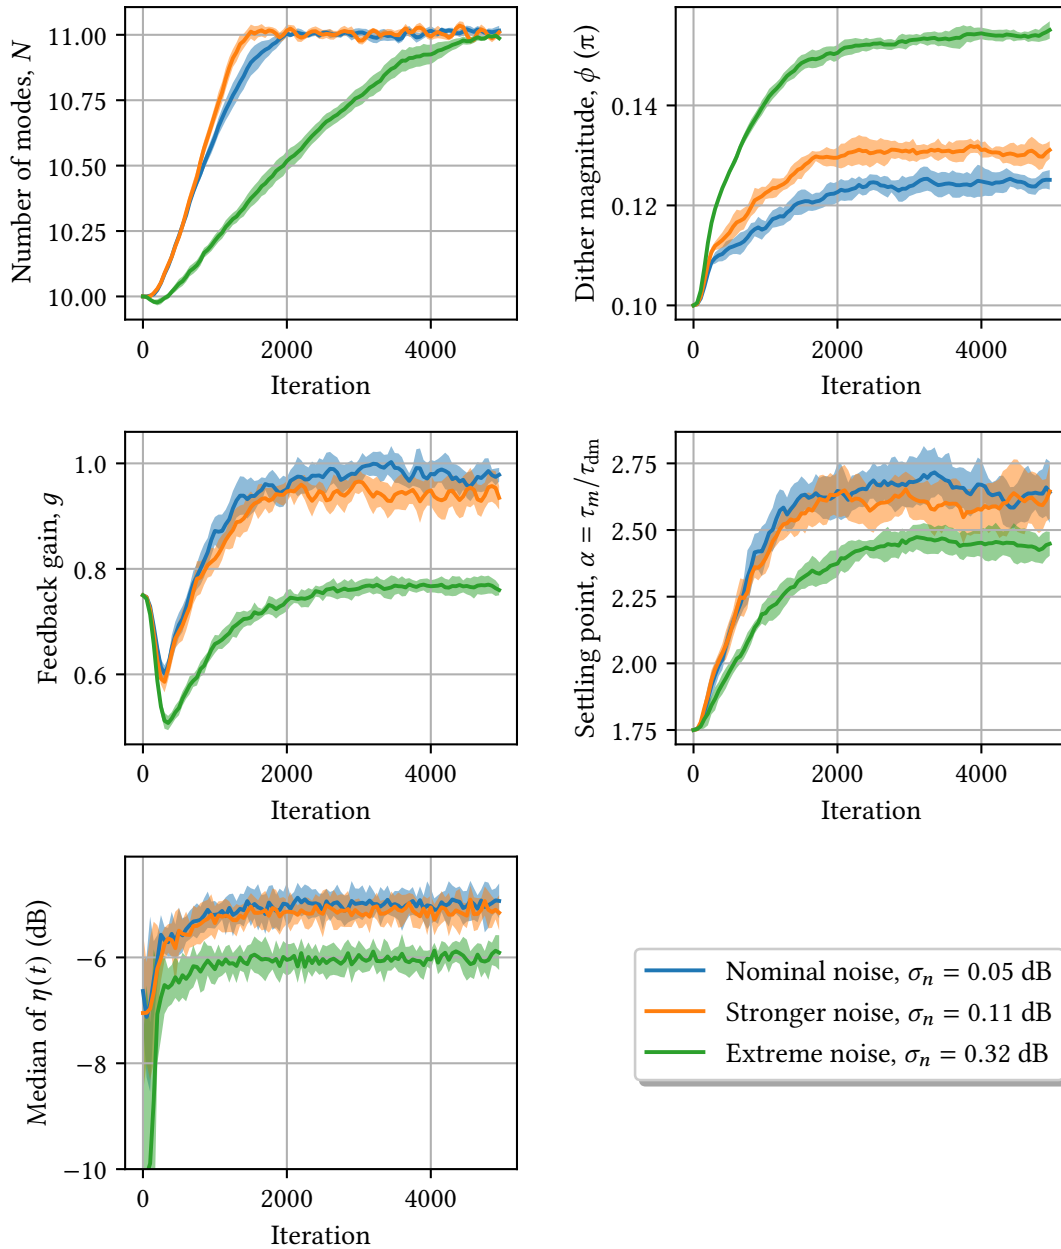


Figure 5-3: An example of results obtained through a joint parameter optimization with the TD algorithm. The evolution of the parameters (top and center) and the performance metric (bottom left) is shown. In this $\tau_{\text{atm}}/\tau_{\text{dm}} \approx 120$ scenario, a segmented DM is configured to correct dominant modes from the Walsh basis for three different cases of the RMS noise σ_n^{dB} . The shaded regions represent the standard deviation between ten repetitions of the optimization process, and the solid lines represent the mean values.

5.3 Issues with Non-Binary Modes

After compiling a reasonable set of parameters, we focus on testing the correction of non-binary (continuous) modes using the non-stochastic algorithm, which is enabled by the new approach developed in Chapter 4. This section describes some learnings related to (a) optimization of the basis itself, and (b) local maxima that may occur during correction of an arbitrary non-binary basis.

5.3.1 Basis Modifications

Given the square aperture of the DM, we assess the use of two correction bases that are orthonormal in a square geometry. First, we test the commonly used two-dimensional Legendre modes (Fig. 5-4a), which are constructed from outer products of Legendre polynomials and are orthonormal on a unit square by definition. Second, we also test a variant of the Zernike polynomials (Fig. 5-4b) that is orthonormalized to a discrete square grid using a matrix-based method [89, 90].

However, as seen in Figs. 5-4a and 5-4b, the RMS of a majority of these modes is dominated by their edge amplitude, in contrast to the uniform binary modes. While this is not an issue from an algorithmic standpoint, it can lead to a large peak-to-peak actuation requirement near the DM edges, which is an ineffective use of its range given the limited overlap with the fiber mode in this region. In an attempt to better weight the modes according to the waveguide intensity distribution, we perform Gaussian apodization of the modes followed by the same matrix orthonormalization procedure. This results in a new set of modes we refer to as apodized Legendre (Fig. 5-4c) and apodized Zernike (Fig. 5-4d) modes, which resemble higher-order waveguide propagation modes.

To quantitatively compare these bases, we focus on two metrics: (1) the gain $\Delta\eta$ after a limited set of modes is corrected (*i.e.*, the aberration fit efficiency), and (2) the maximum actuator stroke needed to perform the full correction. To decouple these metrics from the effects of dynamics, we run the algorithm on 500 frozen atmospheric aberrations picked from the $D/r_0 \approx 3$, $\sigma_R^2 \approx 0.75$ dataset used in the previous section. In each of these simulations, we (a) start with a flat DM state, (b) run two correction cycles with $N = 11$

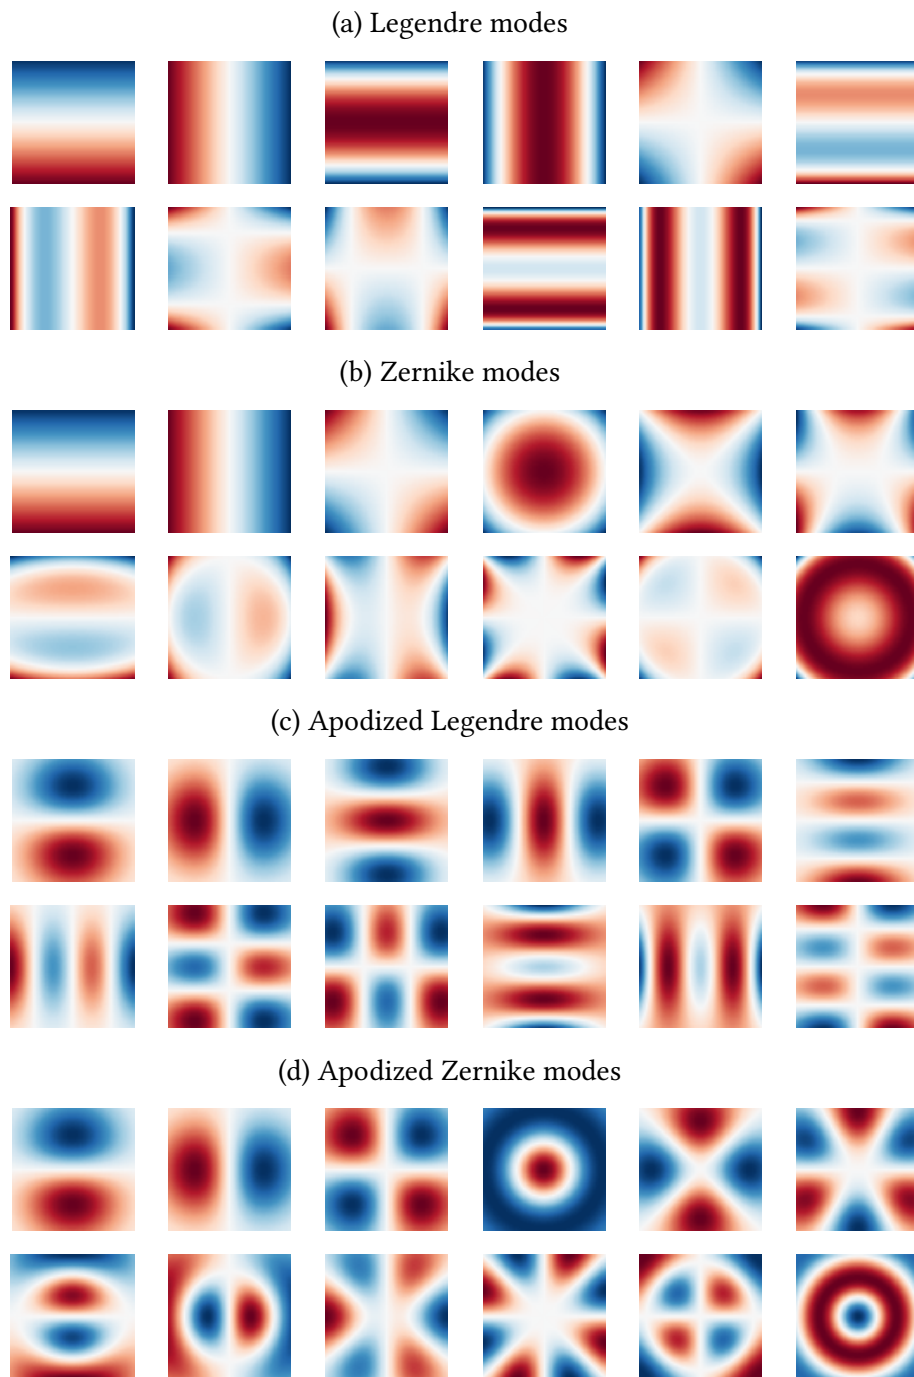


Figure 5-4: A visual comparison of the first 12 modes from the explored non-binary bases. The modes are ordered by increasing spatial frequency. We can see that the Gaussian apodization helps distribute the peaks and valleys more uniformly within each mode, which helps relax the DM actuation range requirement, as demonstrated in Fig. 5-5.

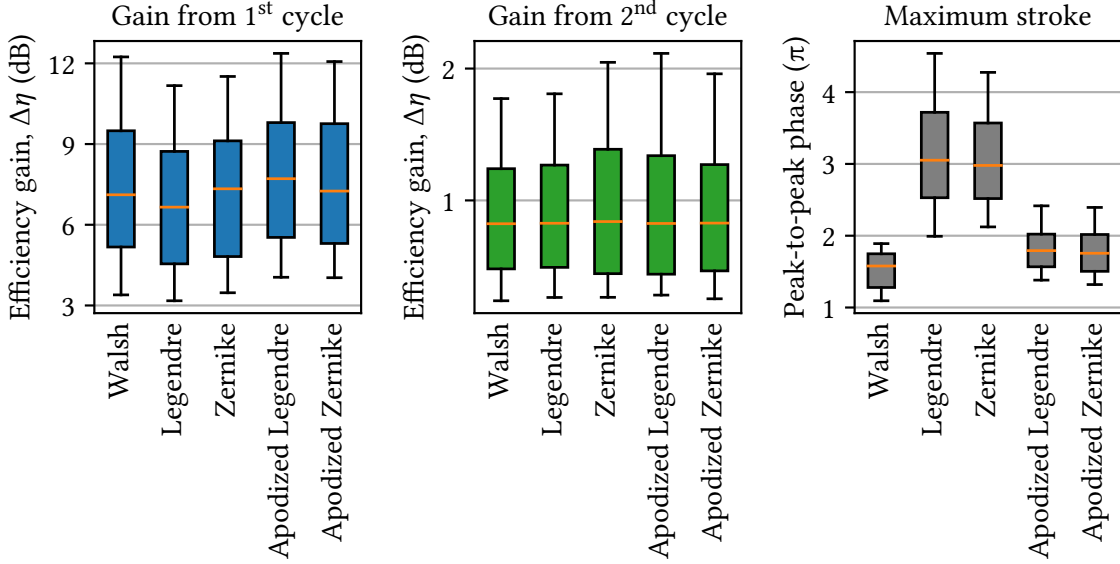


Figure 5-5: Distributions of $\Delta\eta$ and maximum phase stroke when different bases are used to correct 500 frozen atmospheric aberrations from the $D/r_0 = 3$, $\sigma_R^2 \approx 0.75$ dataset.

dominant modes from the given basis, and (c) evaluate $\Delta\eta$ and the peak-to-peak actuator phase after the correction. We collect these metrics for both the 8×8 segmented DM with Walsh modes and the 12×12 membrane DM calibrated with the four continuous bases pictured in Fig. 5-4.

The aggregated simulation results can be seen in Fig. 5-5, showing the distributions of $\Delta\eta$ after the first and second correction cycles together with the maximum actuator stroke at the end of the second correction cycle. In terms of $\Delta\eta$, we observe only subtle differences, with the apodized modes providing marginally higher correction of about 0.5 dB on average. Unlike membrane DMs, segmented DMs can preserve their phase response through 2π wraps of individual actuators, so no more than 2π stroke is needed to represent an arbitrary combination of Walsh modes. On the other hand, using the apodized modes helps reduce the needed stroke by almost a factor of two, which confirms they significantly relax the requirement on the DM actuation range. Given that the most favorable metrics are obtained with the apodized Legendre basis, we decide to use it as the default for correction on the membrane DM.

5.3.2 Local Maxima


Binary phase functions, such as the Walsh modes, have the unique property of

$$\mathcal{W}_i a_i = \mathcal{W}_i (a_i + k\pi) , \quad k \in \mathbb{Z} . \quad (5.20)$$


This $a_i + k\pi$ phase equivalence leads to the periodic $S(a_i)$ relationship seen in Eq. 3.12, which is favorable because it has no local maxima by definition. Given that non-binary modes generally do not have such property, their $S(a_i)$ functions are not periodic and produce similar shapes only around the central lobe. As seen in an example in Fig 5-6a, the functions typically decay exponentially, or have a sinc-squared-like response with small local maxima (side lobes). If a_i are sufficiently small (*e.g.*, in tracking mode), this is not a concern since the algorithms are calibrated for (and meant to operate near) $a_i \approx 0$.

However, this assumption may not be valid during acquisition or after a RSSI fade or link interruption. Fig. 5-6b illustrates this on a specific frozen aberration from the dataset with $D/r_0 \approx 3$ in strong scintillation (assuming no correction). Apart from minor bias errors, we observe that the response to Walsh corrections still qualitatively follows the cosine-squared-like shape. On the other hand, with Zernike modes, some deviations from the original curves and more pronounced local maxima are observed. Evidently, precautions should be taken with modal corrections if membrane DMs are used. While simulations on frozen atmospheric aberrations do not show any indication of performance degradation due to this effect (Fig. 5-5, in contrast, demonstrates performance improvement), the issue does manifest itself in some time-varying simulations. Unlike the former, where the DM is initialized with zero-valued coefficients, the latter can incur larger residual errors due to outdated DM states (*e.g.*, from low correction rate, or after an RSSI fade), which increases the likelihood of cases akin to Fig. 5-6b or worse.

This issue is further explored in a batch of dynamic simulations, where the correction of Walsh modes is directly compared against apodized Legendre modes. For simplicity, we use the same atmospheric and algorithm parameters as in Sec. 5.2.6, but test in both weak and strong scintillation, and perform 100 repetitions of the simulation to understand the uncertainty in the correction trajectory. The results of this batch are visualized in



(a) Response to coefficient adjustments on a corrected wavefront



(b) Response to coefficient adjustments on a highly aberrated wavefront with $D/r_0 \approx 3$

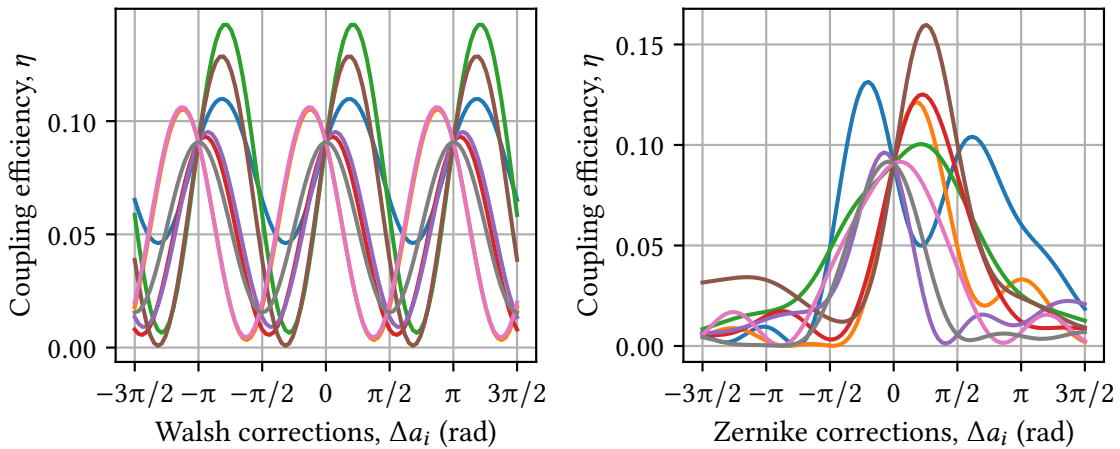


Figure 5-6: Change in coupling efficiency in response to Walsh mode adjustments (left) and Zernike mode adjustments (right). Each color corresponds to an individual perturbation of a specific mode, with the first eight low-order modes shown. The top plots depict a scenario with no residual wavefront error, while the bottom plots show a case where the beam is distorted by a stronger $D/r_0 = 3$, $\sigma_R^2 \approx 3$ atmospheric aberration.

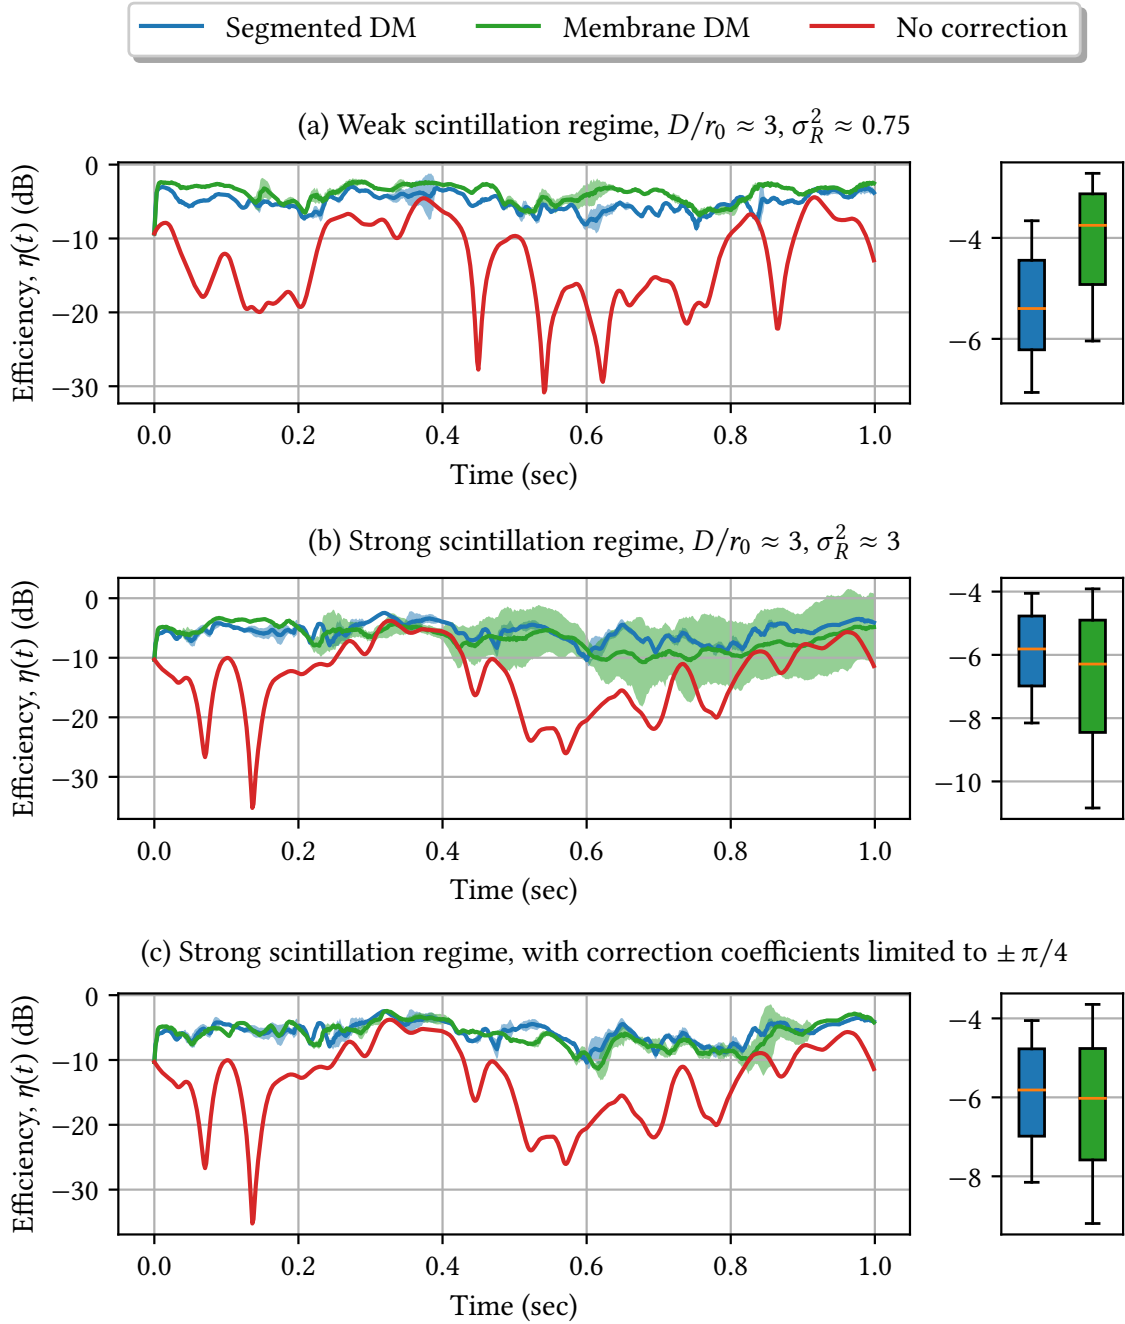


Figure 5-7: An illustration of the issues caused by local maxima during correction of $D/r_0 = 3$ aberrations: (a) in weak scintillation, (b) in strong scintillation, and (c) in strong scintillation, but with the cumulative modal corrections limited to $\pm \pi/4$ for the membrane DM case. The shaded regions correspond to 2σ uncertainty over 100 simulation repetitions. The box plots on the right visualize the aggregated $\eta(t)$ across all repetitions.

Fig. 5-7. We continue to see no sign of issues in weak scintillation, with the $N = 11$ apodized Legendre modes outperforming the Walsh modes with a margin of about 1.5 dB. The correction trajectories are mostly deterministic, with only small regions of uncertainty. However, we observe a large qualitative difference in the results when we run the algorithm in strong scintillation (Fig. 5-7b). At around 0.4 seconds, a permanent algorithm divergence occurs on the membrane DM, indicating that one or multiple Legendre modes may be stuck in a local maximum. Interestingly, even weak noise can significantly influence the correction trajectory, as suggested by the wide uncertainty region of more than 10 dB. On multiple occasions, the diverged state of the membrane DM becomes worse than having no phase correction at all, as illustrated by the red curve crossings. Overall, this leads to significantly degraded performance compared to the segmented DM with Walsh modes, which does not show this behavior, demonstrating the advantage of the binary modes' immunity to local maxima.

In an attempt to mitigate this issue on membrane DMs, we enforce an artificial limit on the individual modal correction coefficients actuated by the DM. While this may hinder correction performance, it reduces the maximum possible instantaneous error, which reduces the likelihood of coefficient divergence. Specifically in this example, we test a limit of $\pm \pi/4$ for every actuated Legendre mode. This is chosen as sufficient based on the predicted spectrum of aberrations in all the considered datasets. The result with this mitigation in place can be seen in Fig. 5-7c, showing better correction behavior with only brief phases of smaller uncertainty but without long-term divergence. While the performance is more comparable to that of the segmented DM, its uncertainty is still more pronounced, which indicates that more sophisticated precautions against local maxima may be needed to achieve a more reliable, practical implementation.

5.4 Performance Evaluations

To assess the performance of each algorithm in a variety of conditions, we run multiple batches of simulations across all the exported atmospheric datasets summarized in Table 5.4. In terms of hardware configurations, we evaluate both the segmented DM with Walsh modes and the membrane DM with the apodized Legendre modes. Since the actuator bandwidth is commonly the primary bottleneck when correction of multiple modes is needed, we also focus on evaluating the performance as a function of $\tau_{\text{atm}}/\tau_{\text{dm}}$ to determine the most bandwidth-efficient variant. The higher the ratio, the more modes can be corrected at a sufficient rate, which is a metric we automatically configure according to Table 5.5. In addition, we maintain the nominal noise assumption established in Sec. 5.2.6, which is conservatively taken to be the worst case during nominal tracking conditions.

We continue to record the coupling efficiency $\eta(t)$ to quantify the wavefront correction performance. However, to isolate the efficiency of the algorithm itself, we additionally normalize $\eta(t)$ by the ideal phase-conjugate coupling efficiency. Effectively, this removes the contribution of intensity aberrations and gives the residual phase loss due to uncorrected spatiotemporal phase errors. Apart from the median phase loss over time, we also calculate the 10th percentile P_{10} to quantify the extent of the distribution's tail. The log-percentile P_r^{dB} can be a useful metric in link budgets, assuming some fraction r (e.g. 10%) of packets can be recovered via error control mechanisms. Finally, we repeat all the simulations 20 times to obtain more statistically significant results.

Fig. 5-8 depicts the aggregated values of the median of phase loss across $\tau_{\text{atm}}/\tau_{\text{dm}}$ and different scintillation regimes in $D/r_0 = 3$. For better visualization of the results from all the remaining atmospheric datasets, we find the threshold $\tau_{\text{atm}}/\tau_{\text{dm}}$ that gives a reasonable performance across all tested D/r_0 values. In this analysis, we determine the threshold for a median phase loss of -3 dB, which results in the set of graphs seen in Fig. 5-9. Lastly, an example side-by-side comparison of the P_{10}^{dB} metric is shown in Fig. 5-10 for the two DM models.



Median performance in $D/r_0 = 3$ and different scintillation regimes

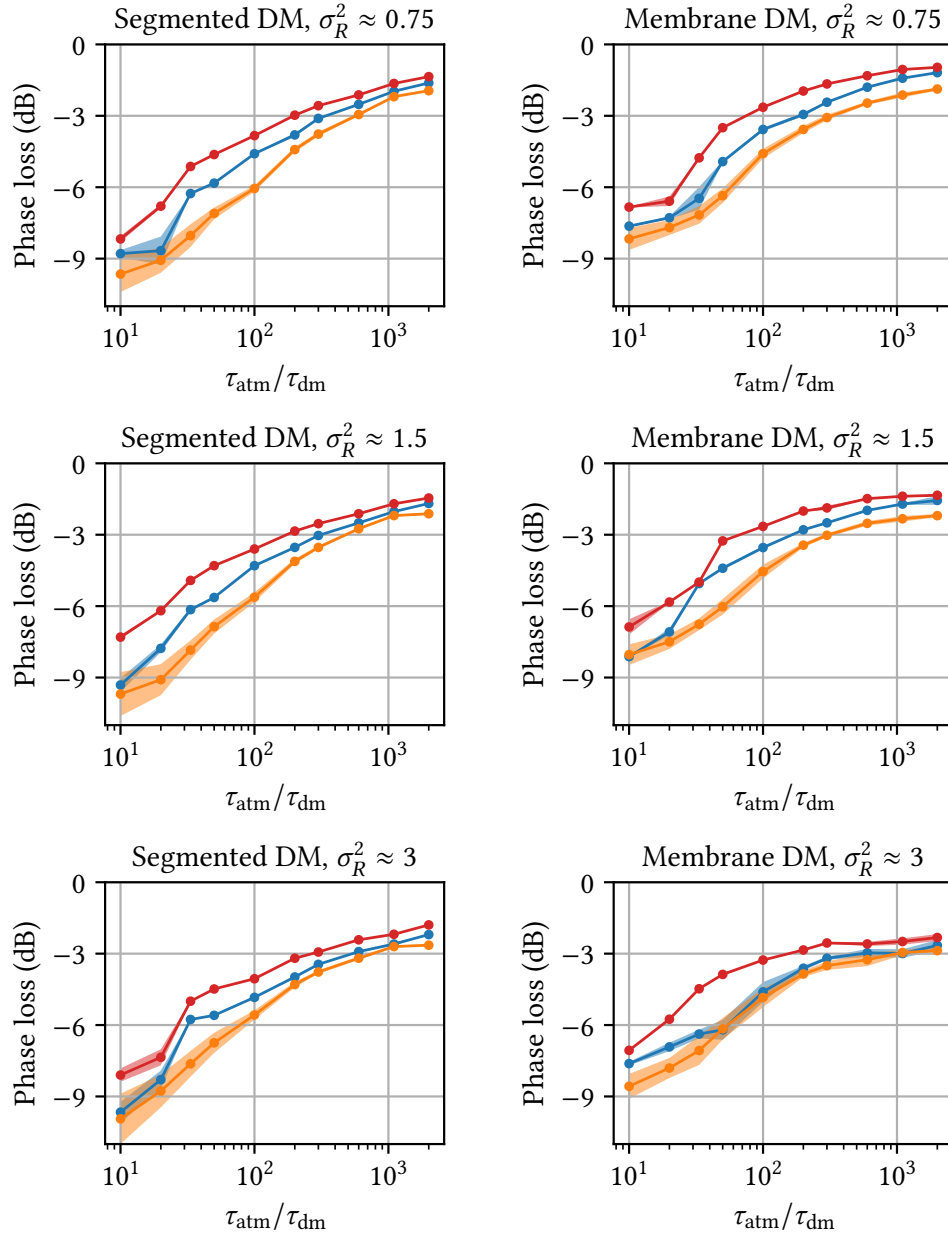


Figure 5-8: A summary of performance results for each algorithm across different scintillation regimes and bandwidth ratios. The algorithms are configured according to Table 5.5, and the nominal noise case of $\sigma_n^{\text{dB}} = 0.05$ dB is assumed. The median of the phase loss is evaluated as a function of $\tau_{\text{atm}}/\tau_{\text{dm}}$, illustrating how higher bandwidth ratios enable dithering and correction of more modes, which leads to better performance. The solid lines and shaded regions correspond to the mean $\pm 1\sigma$ obtained from 20 simulation repetitions.



Algorithm $\tau_{\text{atm}}/\tau_{\text{dm}}$ requirement for -3 dB median efficiency

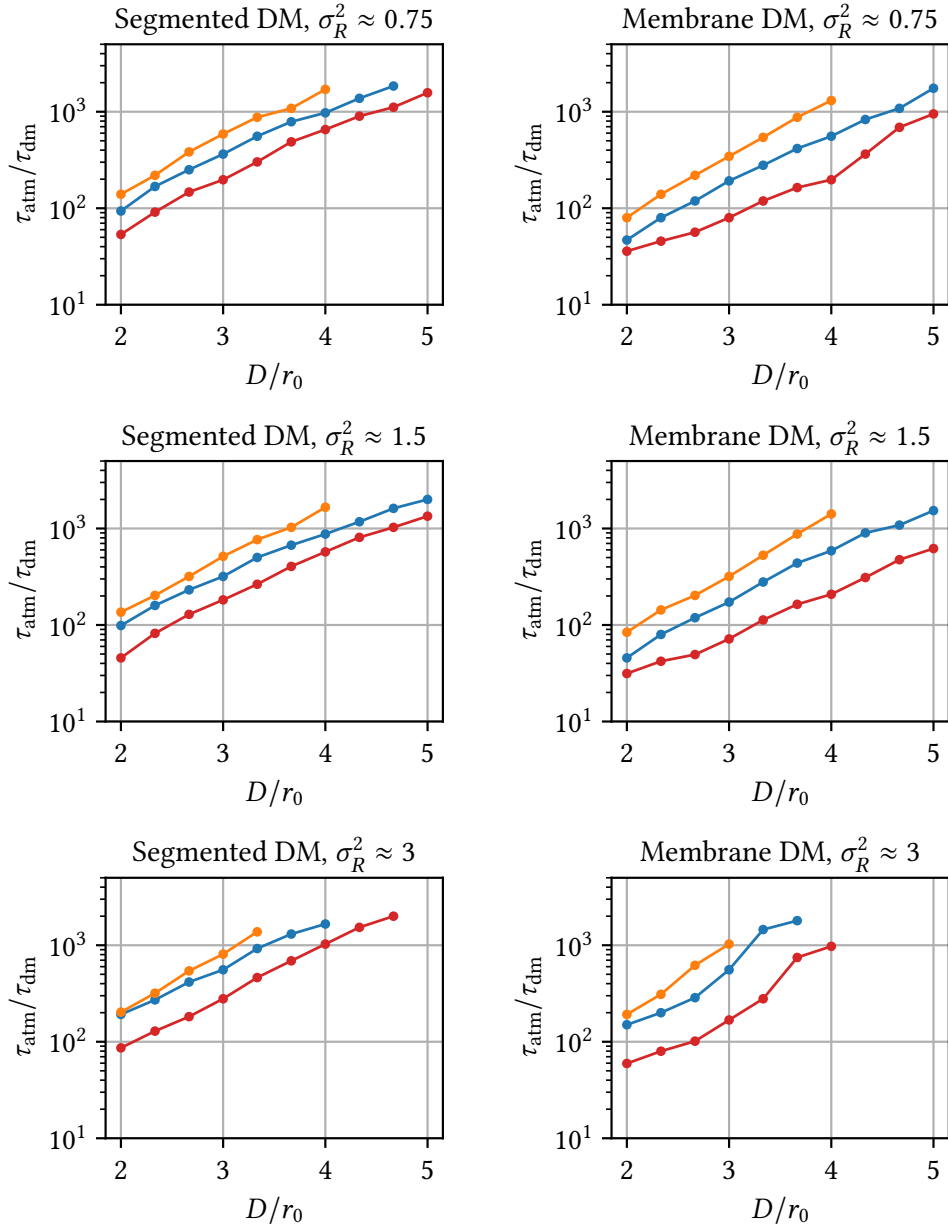


Figure 5-9: A summary of the bandwidth ratios required to achieve -3 dB median phase loss with different algorithms. The thresholds are evaluated from batches of simulations across $\tau_{\text{atm}}/\tau_{\text{dm}}$ in a variety of conditions and graphed against D/r_0 . The algorithms are configured according to Table 5.5, and the nominal noise case of $\sigma_n^{\text{dB}} = 0.05$ dB is assumed.

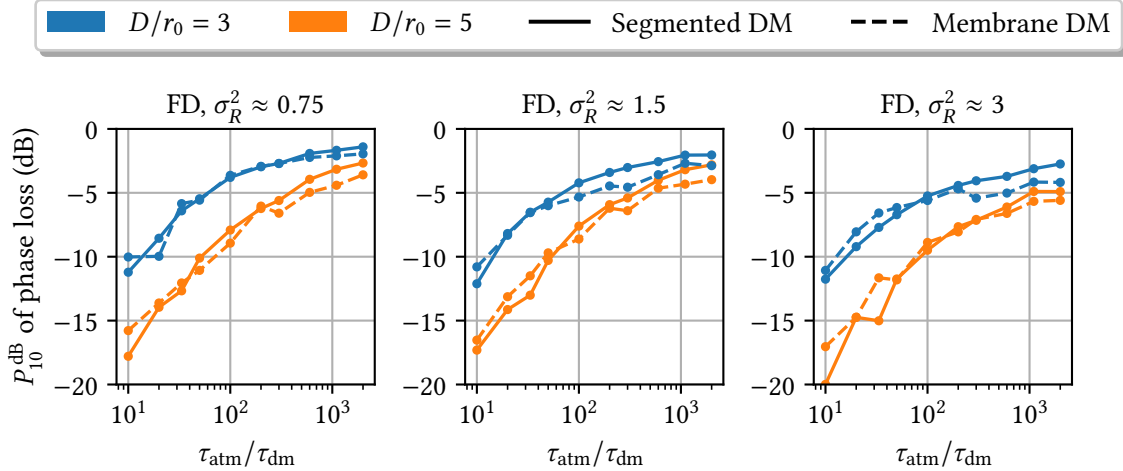


Figure 5-10: Estimated 10th percentile of phase loss with the FD algorithm in different scintillation regimes. Even in weak scintillation, the binary modes achieve a slightly narrower distribution of loss, which can be advantageous in terms of link margin.

5.5 Discussion

Given the trends in Figs. 5-8 and 5-9, we see that, in general, the non-stochastic algorithms outperform SPGD at equivalent bandwidth ratios, demonstrating better convergence properties in comparison to the stochastic correction. The most significant quantitative differences (up to 2–3 dB) can be seen at ratios ranging from a few tens to a few hundreds, which is a common region of operation with commercially available DMs. In addition, taking into account the assumptions used in our analysis, we observe that the QAM version of the FD algorithm consistently outperforms TD, which indicates a higher spectral efficiency was achieved in its configuration.

As previously noted in Fig. 5-7, another trend we see in simulations with weaker scintillation is that the continuous modes provide about 1 dB better median phase loss than the segmented DM with Walsh modes. However, as seen in the bottom graphs in Figs. 5-8 and 5-9, this gain is no longer present in the strong scintillation regime. This suggests either insufficiently mitigated issues with local maxima or possibly degraded fitting due to the presence of more pronounced branch cuts in the phase function. We can see that for D/r_0 larger than about 3.5, the membrane DM, as modeled here, is no longer able to meet

the desired performance metric with a reasonable bandwidth ratio (assuming τ_{atm}^{-1} on the order of hundreds of Hz, $\tau_{\text{atm}}/\tau_{\text{dm}} > 1000$ suggests actuators with >100 kHz bandwidth, which are not commercially available).

This issue is even more evident when looking at the P_{10}^{dB} metric. Fig. 5-10 shows the evolution of this metric specifically for the best-performing FD algorithm. Even in weak scintillation, the metric falls behind by about 0.5–1 dB with membrane DMs, implying that efficiency distributions generally have longer tails. As in Section 5.3, we attribute this effect to poor local maxima mitigation. This suggests that better precautions, in addition to coefficient limits, may be needed for applications where little link margin exists and where binary modes cannot be utilized.

5.6 Practical Case Studies

While the compiled graphs (Figs. 5-8 and 5-9) are valuable for predicting performance across generalized atmospheric parameters, it can be helpful to know what results can be expected in a specific (practical) FSOC link. This section analyzes two interesting practical case studies to demonstrate how the aggregated results can be utilized in such a way.

Example 1: Point-to-Point Link

The first case we analyze is a 5 km horizontal point-to-point link through the atmosphere, which is a common FSOC deployment scenario for middle-mile connectivity [5]. Assuming a transmitter with $\lambda = 1.55 \mu\text{m}$, a receiver with $D = 75 \text{ mm}$, and a moderately-strong turbulence with a uniform refractive index structure constant $C_n^2 = 5 \times 10^{-14} \text{ m}^{-2/3}$, we find that such link will experience strong scintillation and a $D/r_0 \approx 3$ (following the approach outlined in [46]). In terms of AO hardware at the receiver, we assume the use of a commercial-off-the-shelf segmented DM from Boston Micromachines Corporation (BMC) with $\tau_{\text{dm}} = 20 \mu\text{s}$ [83].

Given that this is an FSOC link, we are more interested in evaluating the log-percentile P_{10}^{dB} of phase loss because it can be useful in link budget calculations (in this case assuming that up to 10% of packets can be recovered via error control mechanisms). Considering

these parameters, we can use Fig. 5-10 to approximate the performance of the FD algorithm (from the solid blue curve on the right-hand side plot, which assumes strong scintillation). We find that this configuration can achieve a -3 dB performance in terms of the P_{10}^{dB} metric up to a $\tau_{\text{atm}} \approx 20$ ms, and a -6 dB performance up to $\tau_{\text{atm}} \approx 2$ ms.

Since τ_{atm} in a point-to-point link is primarily a function of the wind speed, it can vary substantially with the weather but is commonly on the order of a few to tens of milliseconds [14]. Thus, given that the -3 and -6 dB thresholds fall approximately in this range, we can conjecture that this algorithm should achieve acceptable performance in this application case if the link budget allocation for AO is at least on the order of 3–6 dB.

Example 2: LEO Downlink

Another interesting application case is that of a low-Earth orbit (LEO) downlink. In this example, we assume a transmitter at an altitude of 400 km with $\lambda = 1.55$ μm and a larger ground receiver with $D = 40$ cm. There are two key factors that make a difference in this scenario. First, the atmospheric properties significantly change with the elevation angle as the satellite passes over the ground station. The strongest turbulence occurs at the lowest elevation, when the beam travels through the longest swath of near-ground atmospheric eddies, leading to a very low r_0 (commonly less than 5 cm). In contrast, the biggest r_0 occurs at the highest elevation, when the link is the most vertical (up to $r_0 \approx 20$ cm, assuming a nominal C_n^2 model [91]). Second, because the slew rate of LEO satellites is very rapid (about 1 deg/s), τ_{atm} ends up mainly being a function of the satellite motion instead of wind. In the worst case, it may be on the order of 1 millisecond [91], which puts a lot more pressure on a sensorless AO algorithm.

Assuming the link can close at an elevation angle as low as 30 degrees, we find that the link will experience weak scintillation and a $D/r_0 \approx 3$. Given these factors, we assume a membrane DM may be a good choice of hardware given its slightly improved performance in weak scintillation (as discussed in Sec. 5.4). Considering this configuration, we can use Fig. 5-10 again to predict the performance of the FD algorithm (from the dashed blue curve on the left-hand side plot, which assumes weak scintillation). We find that with $\tau_{\text{dm}} = 20$ μs and $\tau_{\text{atm}} = 1$ ms, we obtain approximately a P_{10}^{dB} metric of -6 dB. While this

is a fairly large phase loss, it could be acceptable if the link has a sufficient margin. For reference, the simulation of the SPGD algorithm leads to a performance of about -10 dB in this scenario, which indicates that non-stochastic algorithms could enable a broader set of practical application cases.

Chapter 6

Experimental Validation

To help validate the results from the simulation studies, we build a testbed for hardware-in-the-loop experiments under laboratory-simulated atmospheric turbulence. The main objective of these experiments is to confirm the algorithm’s performance across a representative sample of the atmospheric and hardware parameters considered in Chapter 5. Since the performance is sensitive to both spatial and temporal turbulence characteristics, the experiments necessitate both proper spatial reproduction of the aberrated wavefronts (Sec. 5.1.3), and control of the bandwidth ratio $\tau_{\text{atm}}/\tau_{\text{dm}}$. This chapter gives an overview of the testbed design, its calibration to meet these requirements, and finally, validation of the algorithm’s wavefront correction performance.

6.1 Laboratory Setup

The experimental setup is split into two major subsystems: an atmospheric turbulence simulator and a SMF-coupled receiver with a DM for wavefront correction. The high-level design of the testbed is depicted in Fig. 6-1, with a summary of its primary components given in Table 6.1. In addition, an annotated photo of the setup is shown in Fig. 6-2.

The atmospheric turbulence simulator is based on a collimated 1.55 μm source relayed onto two reflective SLMs. To reproduce arbitrary aberrations, the first SLM adjusts the wavefront amplitude, and the second adjusts its phase. We utilize two Thorlabs liquid crystal SLMs with 1920×1200 pixels, which are optimized for 8-bit, 2π phase control




Figure 6-1: A high-level diagram of the experimental setup. Starting on the bottom right, a fiber-amplified 1.55 μm source is spatially modulated in amplitude and phase by two SLMs to simulate wavefronts corrupted by atmospheric turbulence. The left-hand side shows the DM, which is used for wavefront correction, and the fiber optic receiver. An auxiliary WFS is placed right before the SMF coupling interface to facilitate alignment and debugging.

Table 6.1: A summary of the core hardware components used in the setup.

Purpose	Component / Part number
Turbulence simulation	2× Thorlabs EXULUS-HD4 SLM
Wavefront correction	BMC Multi-3.5-DM with X-driver
Diagnostics and alignment	Phasics SID4-SWIR WFS
Optics, opto-mechanics, fibers	Common Thorlabs components
Fiber power measurement	Thorlabs FGA01FC InGaAs PD
Amplification electronics	Analog Devices ADL5304 log amp
Digital conversion	Analog Devices AD7980 16-bit ADC

at $\lambda=1.55 \mu\text{m}$. For amplitude modulation, the liquid crystal birefringence is exploited to interfere the two reflected polarization components (the theory behind this approach is documented in [92, 93]). To achieve this effect, we use a polarization-maintaining (PM) fiber to launch the beam with an approximate 45° polarization offset relative to the working direction of the first SLM, and a linear polarizer to calibrate the interference between the reflected components. The polarization is then rotated using a $\lambda/2$ wave plate into the working direction of the second SLM, where the beam is relayed for phase modulation.

The aberrated beam is then relayed onto the correction element, which is a 12×12 actuator, continuous-membrane MEMS DM from Boston Micromachines Corporation (BMC). A 2:1 relay is used here to approximately match the SLM and DM aperture footprints. The actuators have a maximum mechanical stroke of $3.5 \mu\text{m}$, which, after reflection, corresponds to an optical phase shift of about 9π at $\lambda=1.55 \mu\text{m}$. The Boston Micromachines Corporation (BMC) X-driver is used to control the strokes with a 14-bit resolution. After correction, the beam is finally relayed into the SMF coupling lens, which was chosen to optimize the coupling efficiency according to the process described in Sec. 5.1.2. The fiber receiver consists of a fast InGaAs PD from Thorlabs attached to a custom electronic board with a log amp and an ADC.

Given the number of beamsplitters and optical surfaces between the source and the receiver SMF, a significant fraction (around -40 dB) of the launched optical power is lost, even without amplitude or phase modulation. To obtain an RSSI comparable to the sensitivity of common fiber optic receivers, we use an erbium-doped fiber amplifier (EDFA) at the source to mitigate the large system loss. Finally, to configure and control all the devices in the testbed, one central Linux-based desktop computer is set up with all the necessary hardware and software interfaces.




Figure 6-2: An annotated photo of the setup in the laboratory.

6.2 Hardware Calibration

Apart from optimizing static alignment between the components, we also perform device-level calibration of the SLMs and the DM to prepare the setup for wavefront correction experiments. The purpose of the SLM calibration is to ensure that the atmospheric wavefronts are reproduced as closely to the simulated conditions as possible. On the other hand, the DM is used to calibrate the algorithm for modal wavefront correction with the apodized Legendre basis, which is used in simulations with the DM model in Chapter 5.

6.2.1 Turbulence Injection

Since the SLMs are particularly sensitive to polarization, we use a polarimeter for coarse calibration and linear polarizers after each device to fine-tune the dynamic range (DR) of both amplitude and phase modulation. The amplitude SLM is calibrated in its fully *off* state (*i.e.*, maximum interference) by rotating the polarizer until the minimum power in the fiber is measured. Following this procedure, we obtain a DR of about 50 dB for intensity modulation, which is sufficient to realize the fields with strong scintillation considered in the simulation datasets. For phase modulation, a similar approach is used to align the polarization with the working direction of the second SLM. The full 2π DR of phase modulation is then verified by applying a test gradient and confirming that the change in fiber-coupled power is as expected.

After optimization of the DR, the datasets with the aberrated electric fields (the generation of which is described in Sec. 5.1.3) are split into time series of phase and intensity at each SLM pixel, and converted to 8-bit depth according to the measured SLM response. This converted data is then streamed to each SLM via an HDMI interface at a frame rate of 20 Hz. Since the original fields are sampled at 5 kHz, the 1-second-long atmospheric snapshots take 250 seconds to execute on the testbed, which is later considered in the determination of the temporal bandwidth ratio. To verify that the wavefronts are reproduced correctly, we record the RSSI while the SLMs are active and compare it with the expected trend from the software model.

This comparison is summarized in Fig. 6-3 for two datasets with different scintillation

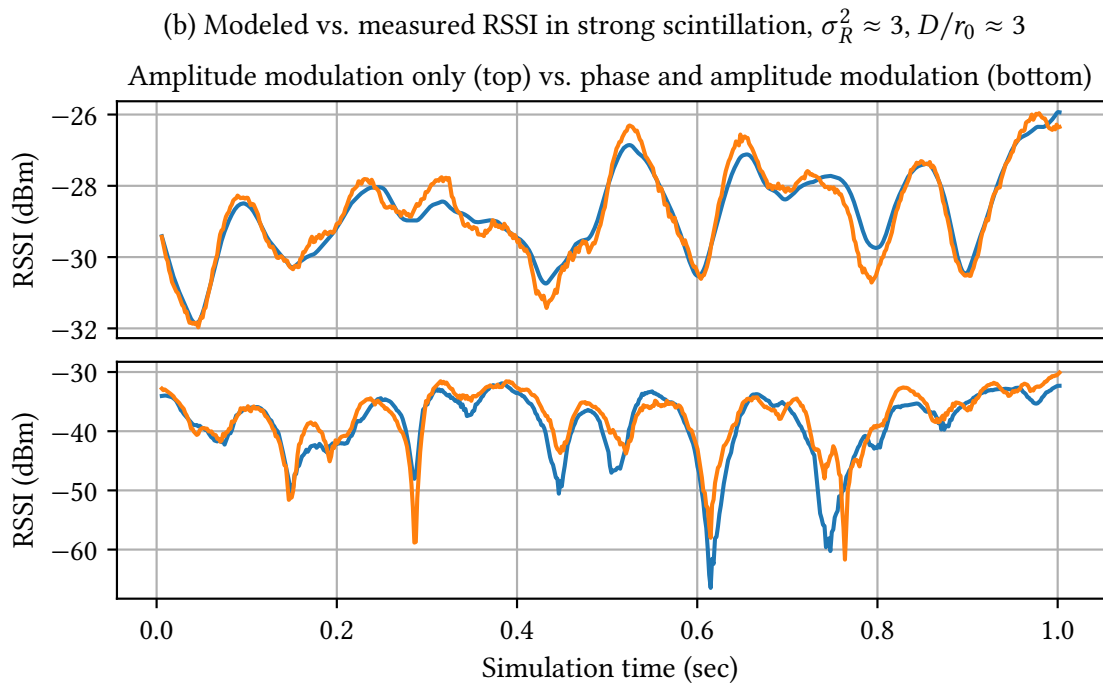
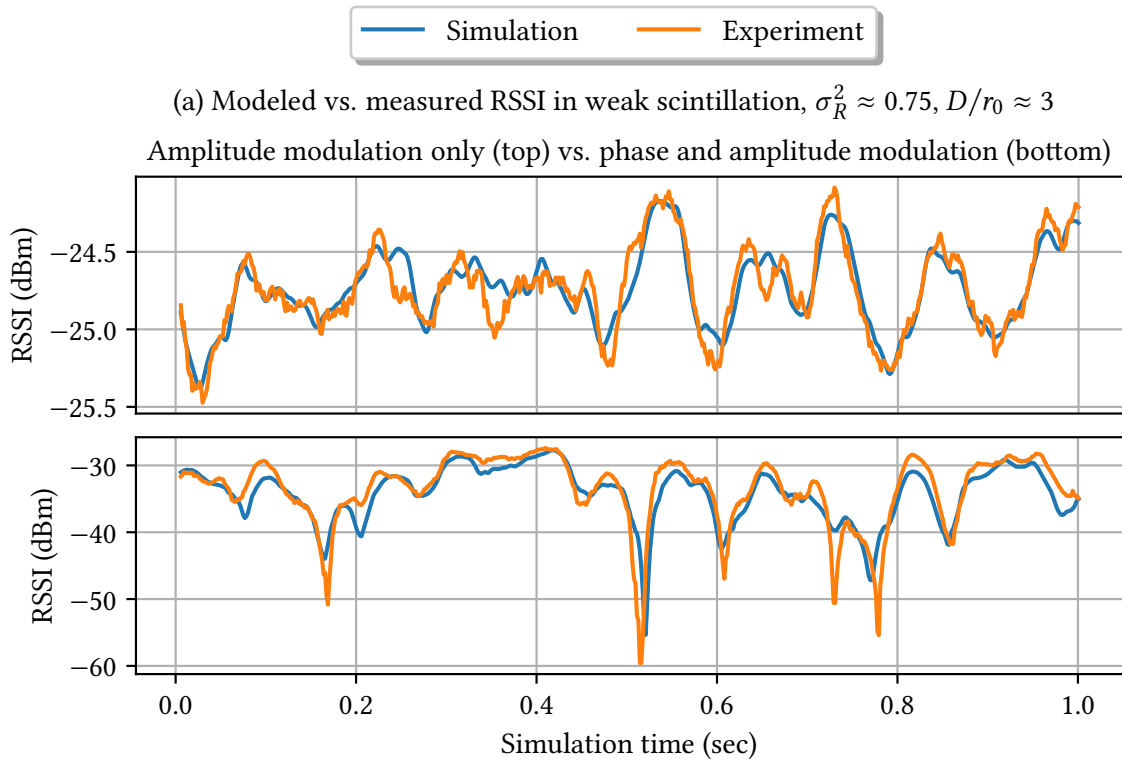


Figure 6-3: Comparison between the modeled and the measured impact of turbulence on RSSI. The recorded power versus time is shown for $D/r_0 \approx 3$ in weak scintillation (top figures) and strong scintillation (bottom figures). Both amplitude-only and joint amplitude-phase modulation (*i.e.*, both SLMs are engaged) are compared.

regimes. Given that the modeled S^{dB} is dimensionless, the simulation results are offset by a constant to match the average measured RSSI, which allows better visualization of the data correlation. While some small local discrepancies between the predicted and the measured RSSI can be seen, we generally observe good agreement in the trend and the DR of the reproduced disturbances. The deviations are primarily attributed to residual alignment errors, low-frequency environmental effects in the laboratory, and open-loop control of the SLMs without hardware synchronization (which is partially mitigated by the slower software frame rate). In particular, imperfect alignment may lead to both active aperture offsets and diffraction of the wavefront (due to some residual propagation occurring between the pupil relays), which is expected to lead to small errors in the realization of the aberrated wavefront.

6.2.2 Estimation Weights

As discussed in Chapter 4, the algorithm calibration reduces to the experimental measurement of the weights c_i for each mode that is to be corrected. To run the algorithm, these weights are augmented with the known dither magnitudes ϕ_i to obtain the necessary log-estimation weights $c_{1,i}$ and $c_{2,i}$ defined in Sec 4.2. A straightforward way to find c_i in the laboratory¹ is to disengage the SLMs and individually adjust the modal coefficients a_i with the DM while recording S^{dB} . The weights are then readily obtained via a quadratic fit on these measurements (following Eq. 4.14). The benefit of this approach is that it not only reduces the modal approximation errors (as predicted in Fig. 4-1), but also accounts for some setup-specific effects that are challenging to model, such as alignment errors between the DM and the SMF coupling lens, or spatial phase errors resulting from imperfect mapping of the phase function to the actuator strokes.

Given that the main purpose of the experimental validation is to confirm the performance results obtained in Chapter 5, the algorithm is configured exactly according to the simulations, and the apodized Legendre basis is thus used for all algorithm tests. To actuate the modes, we use an application programming interface (API) provided by BMC, which

¹This calibration procedure could also be performed outside of laboratory conditions but would necessitate a flat reference wavefront.

transforms the desired phase function to actuator voltages according to the open-loop characterization of the device. A subset of the measurements which are used to determine c_i (and its linear variant, w_i) is summarized in Fig. 6-4. While most of the obtained w_i are close to unity (as predicted by the standard Maréchal's approximation, given in Eq. 4.2), we observe a few outliers, some with values as low as 0.84. However, in all cases, the weighted quadratic fit produces very good agreement with the measurements, confirming the utility of Eq. 4.14 and the simulation results in Fig. 4-1.

6.3 Algorithm Validation

The primary objective of the hardware-in-the-loop algorithm testing is to confirm the performance statistics obtained in Chapter 5 across a representative sample of the considered atmospheric parameters and system configurations. Given the available laboratory equipment, we primarily focus on the validation of results obtained with the membrane DM model. In addition, since the FD algorithm necessitates custom DM drive electronics and real-time DSP, we limit our testing to the TD and SPGD algorithms (based on step dither), which are more easily prototyped.

Considering that the final results in Sec. 5.4 are compiled from hundreds of simulations across a large spectrum of D/r_0 , σ_R^2 , and $\tau_{\text{atm}}/\tau_{\text{dm}}$ values, we decided to prioritize the permutations for hardware testing such that the total number of experiments is manageable. We choose the datasets with $D/r_0 \approx 3$ in weak and strong scintillation regimes, which are considered medium-difficulty scenarios to exercise the algorithms based on the simulation results. In terms of temporal regimes, we replicate a few different $\tau_{\text{atm}}/\tau_{\text{dm}}$ on the testbed by injecting artificial delays between the DM updates. A complete summary of the differences between the configurations used in the simulation studies (Chapter 5) and in the experimental studies (Chapter 6) is given in Table 6.2.

The selected permutations are executed on the testbed with the algorithms configured exactly according to the respective simulation cases. This involves the number of corrected modes, the dither magnitudes, and the feedback gains, all of which are jointly optimized in the prior simulation studies. To assess the wavefront correction performance, we calculate

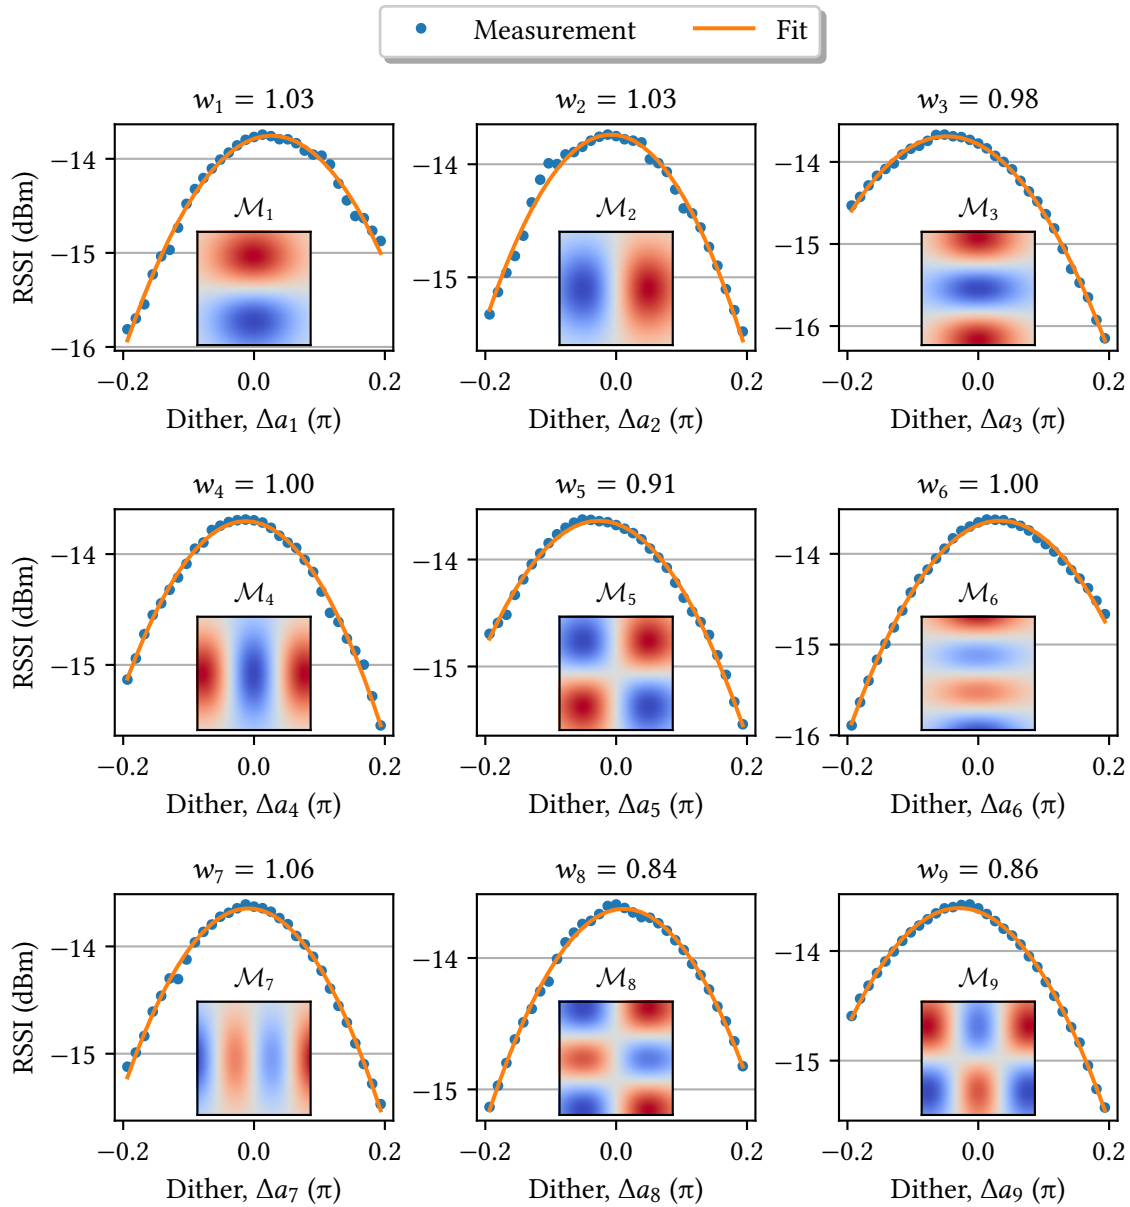


Figure 6-4: A summary of the measurements collected to help determine the algorithm calibration weights. Each mode is individually adjusted via the DM API while measurements of RSSI are collected. A quadratic fit is then performed to obtain the respective log-weights c_i . The plots depict the fits obtained up to $i = 9$, with the respective modes in the insets. For reference, the linear weights w_i are calculated via Eq. 4.13 and shown above the plots.

Table 6.2: A summary of system configurations and atmospheric parameters used in the simulations studies, including a subset that was selected for experimental validation.

Category	Simulation studies	Experimental studies
DM	12 × 12 membrane 16 × 16 piston	12 × 12 membrane —
Algorithm	Time-division (TD) Frequency-division (FD) SPGD (reference)	Time-division (TD) — SPGD (reference)
σ_R^2	0.75 (weak) 1.5 (medium) 3 (strong)	0.75 (weak) — 3 (strong)
D/r_0	2–5	3
$\tau_{\text{atm}}/\tau_{\text{dm}}$	10–2000	40, 85, 250

the phase loss over time by normalizing the recorded RSSI with the ideal phase-conjugate RSSI, which is recorded as part of the SLM tests outlined in Sec. 6.2.1 (*i.e.*, amplitude-only SLM modulation is considered equivalent to ideal phase correction). An example of the measured phase loss while the algorithms are correcting turbulence on the testbed is shown in Fig. 6-5 alongside the predicted performance from the software simulation. As in the analysis done in Chapter 5, we compile both the median and the 10th percentile of the phase loss over the duration of the turbulence snapshot for all the experiments. A comprehensive summary of these statistics is visualized in Fig. 6-6, where we overlay all the experimental results on top of the trends predicted in Chapter 5.

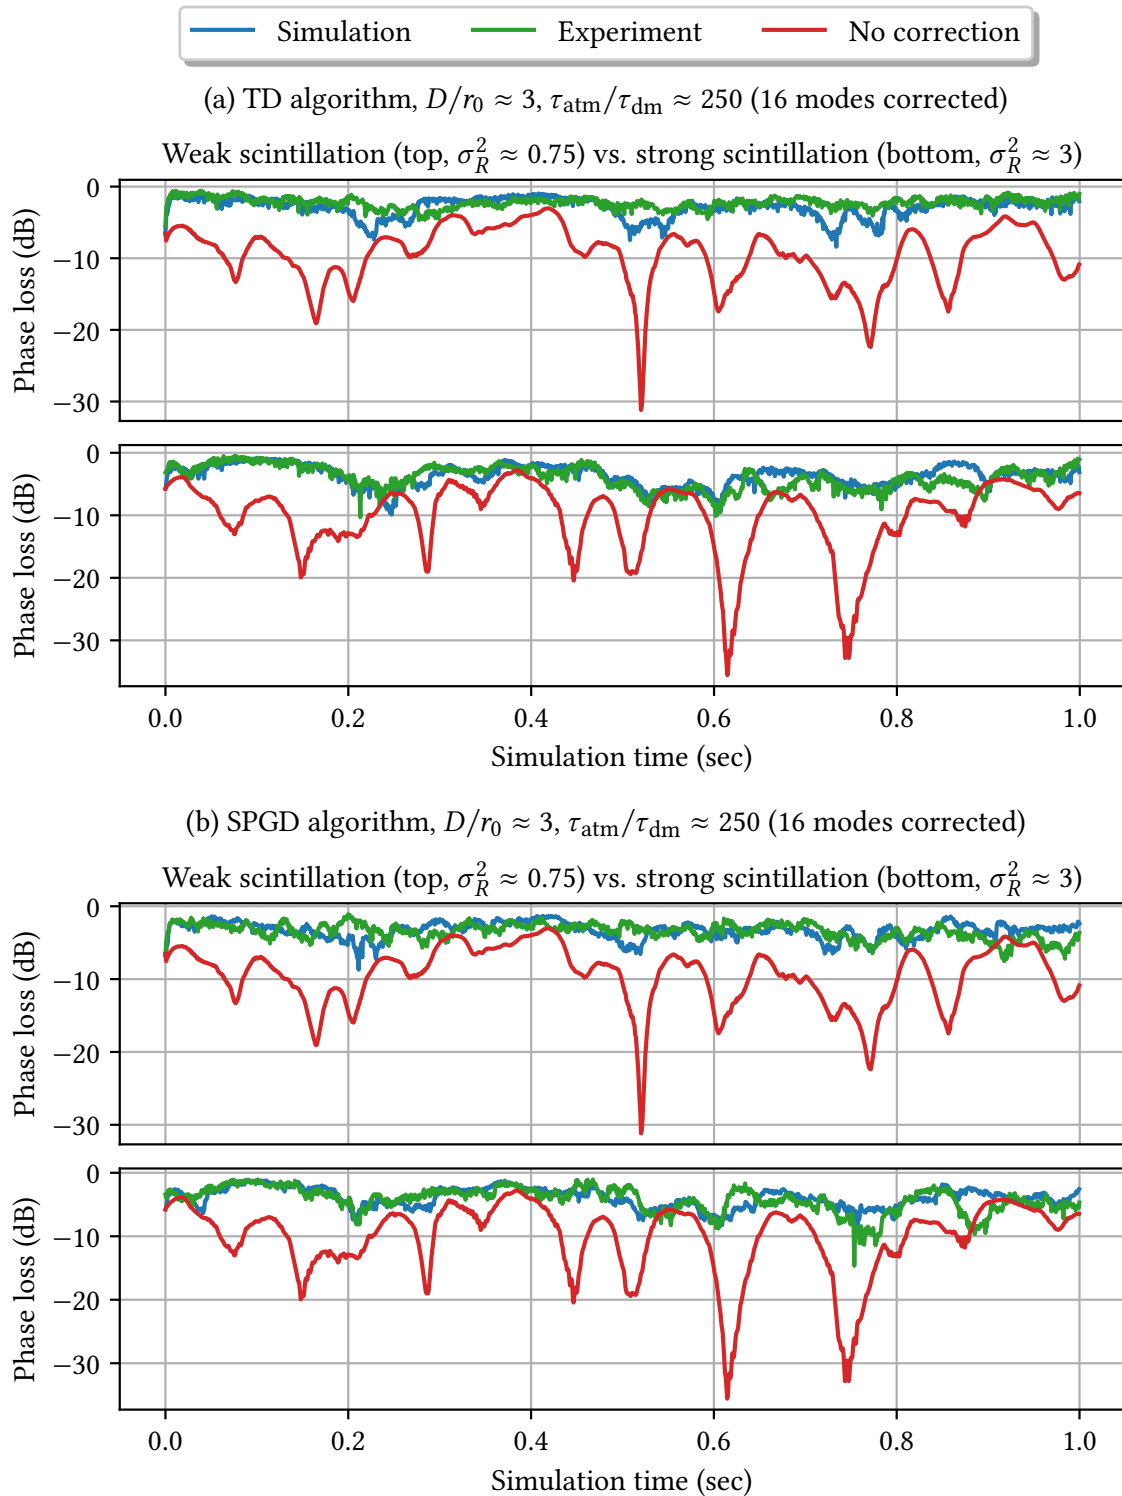

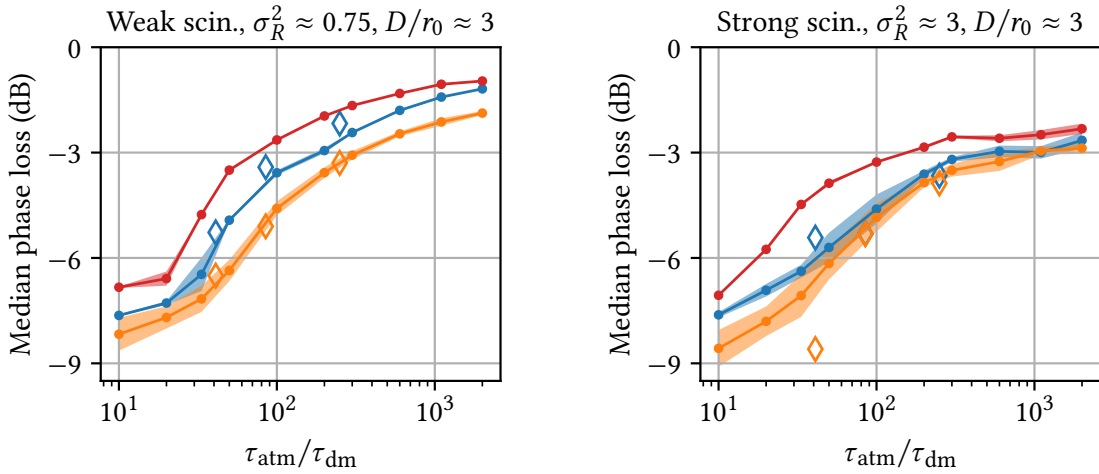


Figure 6-5: Comparison of the simulated and the measured trend in correction performance for the TD algorithm (top plots) and the SPGD algorithm (bottom plots). In each case, a $D/r_0 \approx 3$ scenario with both weak and strong scintillation is shown. For reference, we also depict the phase loss assuming no correction, which is adapted from data in Fig. 6-3.



(a) Median of phase loss vs. bandwidth ratio for different algorithms



(b) 10th percentile of phase loss vs. bandwidth ratio for different algorithms

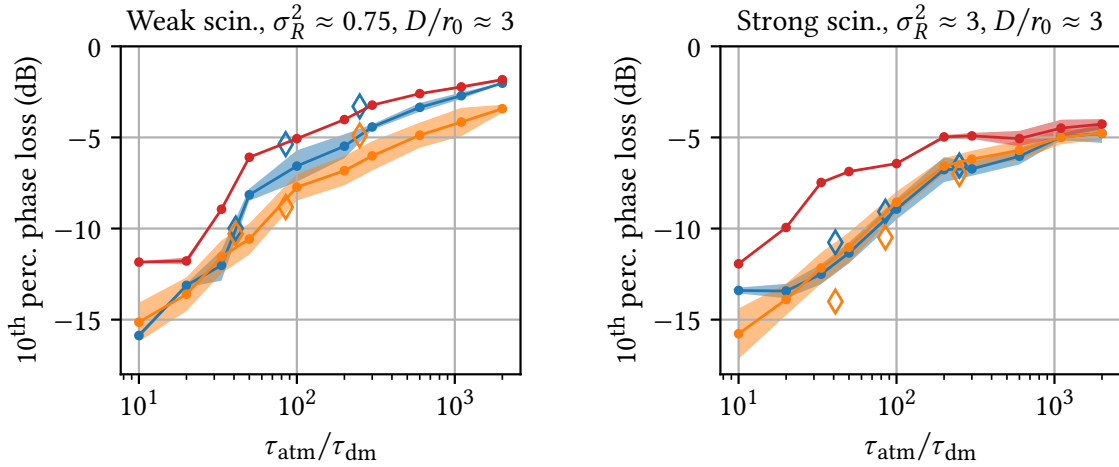


Figure 6-6: A statistical overview of the measured performance (diamonds) overlaid on top of the results from Chapter 5 (solid lines and shaded regions correspond to the average $\pm 1\sigma$ simulation results). In terms of the TD algorithm, we observe slightly better experimental results in weak scintillation, and good agreement with simulations in strong scintillation. The SPGD hardware results generally match the simulation results, apart from a single outlier in strong scintillation, where the algorithm could not achieve stable convergence at the lowest tested bandwidth ratio.

6.4 Discussion

Based on the data shown in Figs. 6-5 and 6-6, we generally observe good agreement between the simulated and the measured correction performance with both algorithms. More specifically, the measured SPGD performance, which is tested as a state-of-the-art reference, falls within the $\pm 2\sigma$ simulation range in all but one case, which indicates that the simulation models have likely not omitted any significant effects. The singular SPGD outlier is only observed in strong scintillation at the lowest temporal bandwidth ratio of $\tau_{\text{atm}}/\tau_{\text{dm}} = 40$, where stable convergence of the algorithm was not successfully achieved, and several spikes in phase loss occurred. However, given that this is not a very likely operation regime due to the large phase loss, it was not investigated further.

In terms of the TD algorithm, we observe slightly better hardware performance in weak scintillation (by roughly 0.5 dB) and excellent agreement with simulations in strong scintillation. In addition, looking at the weak scintillation example in Fig. 6-5a, we notice more robust tracking performance on hardware, which is in contrast with the occasional performance dips seen in the simulation results. This leads to an improvement of the 10th percentile metric by approximately 1–2 dB, as seen on the left of Fig. 6-6b. However, since this improvement is not observed in strong scintillation, it is inconclusive whether it is inherent to the algorithm itself or is an artifact of imperfect reproduction of turbulence (1–2 dB local errors can be observed in tests of turbulence injection in Fig. 6-3 as well).

In addition, with both algorithms, we notice a few instances where the phase loss dips below what is expected without correction (*i.e.*, running the algorithm produces a worse result than not running it). This is visible in two instants for the TD algorithm in strong scintillation (around $t = 0.25$ s and $t = 0.9$ s on the second plot in Fig. 6-5a) and two instants with the SPGD algorithm (around $t = 0.9$ s in Fig. 6-5b). While this may seem unexpected, it is likely a combination of at least two effects that have also been observed in simulations. On one hand, these instances occur only when the turbulence happens to be momentarily very weak (*i.e.*, during the peaks of the uncorrected phase loss), in which case dither injection may make a corrected wavefront noticeably worse than an uncorrected wavefront. On the other hand, as discussed in Sec. 5.3.2, correction of non-binary modes

may still lead to brief algorithm divergence due to local RSSI maxima, even with the mitigation we tested in simulation (the same mitigation is utilized in experiments).

Ultimately, minor deviations from the simulation results are expected for several reasons. On one hand, several hardware effects, such as SLM/DM fill-factors, nonlinearities, and surface edge effects, were not modeled for the sake of implementation simplicity. On the other hand, static hardware errors, such as component misalignment, imperfect pupil relays, and aperture mismatch, can also contribute to deviations from the modeled spatial phase structures. Finally, the lack of a more precise hardware synchronization between the SLMs, and between the DM updates (including due to the use of a non-real-time operating system) may also affect the evolution of the executed modal corrections.

Chapter 7

Conclusion

7.1 Summary of Work and Contributions

This work addressed the development and testing of a generalized, modal, non-stochastic, sensorless wavefront correction algorithm, as well as a broader analysis of the circumstances under which sensorless algorithms attain sufficient performance for FSOC.

While generalized algorithms that are agnostic to the choice of hardware or modal correction basis exist, they are normally based on stochastic estimation. This hinders their convergence rate, which in turn may limit the practical circumstances under which they are useful. Research indicates that non-stochastic algorithms can achieve more promising convergence rates, however, they have only been demonstrated in a few special configurations, and a generalized modal algorithm is missing. In addition, it is generally difficult to predict how different algorithms perform as the atmospheric and hardware parameters change due to a lack of higher-fidelity simulations and experiments.

We addressed the lack of a generalized modal algorithm in Chapter 4, which is the first contribution from this work. Inspired by a modal approach based on time-division (TD) with binary step dithers, and coherent beam combining (CBC) methods utilizing frequency-division (FD) with zonal sinusoidal dithers, we developed a new algorithm that can be calibrated for correction of arbitrary modes through TD or FD. We also presented a simple approach to perform this calibration in practice. The calibration allows the use of linear estimation equations to run the algorithm, which are less complex in terms of

digital signal processing (DSP).

In Chapters 5 and 6, we addressed the second and the third thesis contributions, which focus on evaluating performance across different conditions and proving that the non-stochastic approach achieves better efficiency. We developed an end-to-end simulation environment to help understand the performance of different algorithms under varying atmospheric conditions and hardware configurations. To make sure each test case is configured optimally, we developed tools for offline optimization of the algorithm parameters, and solved issues related to running the algorithm on membrane deformable mirrors (DMs). We also designed, built, and calibrated a testbed to experimentally reproduce a subset of the simulations.

After compiling a large number of simulation results, we demonstrated that the non-stochastic algorithm consistently outperforms the stochastic one in terms of convergence rate. This was evident in situations where the AO actuator bandwidth was the performance bottleneck, in which case the non-stochastic algorithm achieved on average 2–3 dB better link margin with the same hardware configuration. In addition, we also analyzed the bandwidth needed to meet a specific performance target and found that using the non-stochastic algorithm can relax the AO actuator bandwidth requirement by a factor of 2–3 in comparison to the stochastic approach. We were able to successfully replicate these results on the hardware testbed as well.

Most importantly, the compiled results helped us understand the performance of sensorless wavefront correction algorithms under a wider range of circumstances than previously considered in the literature. Using the aggregated results across different key hardware and atmospheric turbulence parameters (such as D/r_0 , $\tau_{\text{atm}}/\tau_{\text{dm}}$, and σ_R^2), we can more easily predict the necessary link margin in a variety of practical application cases. We demonstrated this in two example case studies of interest in Sec. 5.6. Ultimately, the compiled performance graphs turned out to be a valuable tool for predicting the conditions in which sensorless approaches attain adequate performance in the context of FSOC, and in which it is still preferable to use conventional approaches based on dedicated wavefront sensing.

7.2 Limitations and Future Work

There are several interesting directions in which this work could be further expanded. While our analysis through simulations and experiments led to promising results, our approach had a few limitations, and some of our observations deserve more in-depth analysis and higher-fidelity modeling. In addition, there are multiple ways in which the presented stochastic and non-stochastic methods could be further optimized, both for performance and for robustness in terms of practical use. This section concludes this thesis with a summary of these ideas.

Adaptive Algorithm Configuration

Our analysis assumed that some key properties are known and that they do not change over time. This includes, for example, the atmospheric time constant τ_{atm} and the statistics of the RSSI noise, which we used for offline optimization of the algorithm configuration. However, since these properties may change over time in practice, it would be helpful to estimate them *online*. This would allow for adaptive configuration of the algorithm as the underlying conditions change, which could lead to an improvement in performance.

One specific example of how this could be helpful is related to the parameter N , which controls how many modes are corrected by the algorithm. As extensively discussed in this work, if the actuator bandwidth is the main performance bottleneck, N must be chosen carefully to minimize the spatiotemporal wavefront error. If τ_{atm} was estimated online, it could be used to re-configure N in real time. This could be done, for instance, via some function $N \propto f(\tau_{\text{atm}}/\tau_{\text{dm}})$, which is the same approach we followed in our offline algorithm configuration process.

Local Maxima Analysis

One area that would undoubtedly benefit further investigation is that of local RSSI maxima, which complicate the optimization of non-binary modes. While we partially mitigated the algorithm divergences observed with membrane DMs due to this effect, we used a fairly naive approach that constrained the corrections to decrease the likelihood of long-term

divergence. A more in-depth analysis of why and when these divergences occur, as well as how to most effectively prevent them, would be helpful to increase the robustness of modal correction with membrane DMs. More specifically, a literature review on non-convex and global optimization techniques would help assess more advanced approaches for local maxima mitigation. Evaluating the algorithm behavior with and without such mitigations in place (both on hardware and in simulations) would be beneficial to improve the understanding of this problem. Finally, it would be valuable to experimentally test the optimization of binary modes on segmented DMs, given that this configuration has been shown to avoid local RSSI maxima altogether.

Algorithm Improvements

A few modifications could be explored to further optimize both the stochastic and the non-stochastic algorithms. For example, our analysis assumed a fixed gain (learning rate) for the stochastic algorithm, which is commonly the case in literature on SPGD. However, recent advances in machine learning have produced many new promising gradient-based optimization algorithms with adaptive learning rates. It would be interesting to test if such methods improve stochastic wavefront correction performance, which is a more dynamic (evolving) optimization problem in contrast to the offline optimization problems common in machine learning.

With the non-stochastic algorithm, we implemented simple integral feedback controllers to reject the measured modal errors. It would be useful to investigate how much can the error rejection improve by using more modern feedback control methods. Specifically, it might be interesting to evaluate techniques from adaptive and optimal control theory, as well as more advanced predictive filters. Another modification that is easily applicable to the frequency-division (FD) algorithm would be to adapt not only the dither magnitude but also the bandwidth for each mode. This would allow, for example, estimating and correcting tilt at a different rate than higher-order aberrations, which are expected to occur faster according to turbulence models.

Finally, both the stochastic and the non-stochastic algorithms could benefit from loop-shaping filters, which would allow faster dither injection by pre-emphasizing the dither

signal amplitudes. While this would increase the implementation complexity and put more pressure on the DACs and actuators, it would also reduce the required $\tau_{\text{atm}}/\tau_{\text{dm}}$ needed to achieve a given performance goal.

Additional Simulations and Higher-Fidelity Modeling

Additional interesting insights might be obtained by running more simulations and by improving the simulator's fidelity. For instance, the small discrepancies observed between experimental and simulated results could be better analyzed if the software models were improved. This would help understand whether the discrepancies stem from hardware imperfections that were not modeled or from other effects, such as uncertainty due to local maxima. A higher-fidelity simulator could also provide insights on how other DM imperfections, such as quilting, the restricted motion of the edge actuators, or atypical influence functions, affect the algorithm's performance. This may help inform the design of future AO hardware so that it is more optimized for sensorless wavefront correction.

Running more simulations across different values of σ_R^2 would also help understand why and how much does phase correction (and reconstruction) degrade in stronger scintillation. While it is evident from our simulations that stronger scintillation leads to more pronounced local maxima issues with non-binary modes, we also observe some increase in phase loss with binary modes (even though this metric is normalized according to the best possible phase correction). A more in-depth analysis of this issue may help increase the fundamental understanding of how branch cuts in the phase function (due to the emergence of branch points, or wavefront nulls) affect the behavior of sensorless wavefront correction.

Higher-Fidelity Hardware Validation

Several improvements could be made to the hardware testbed and the testing approach to improve the validation of the simulation results. Improving the SLM alignment and implementing hardware synchronization between the two SLMs would help reproduce the simulated fields with a smaller spatiotemporal error, which might mitigate some discrepancies we observed in the results. In addition, the testbed fidelity could be improved by implementing real-time digital signal processing (DSP) to run the algorithm and by

not simulating different τ_{dm} to control the $\tau_{\text{atm}}/\tau_{\text{dm}}$ ratio. Ideally, one would adjust τ_{atm} , and the algorithm would be configured for the actual hardware τ_{dm} . This may, however, require more capable SLMs or DMs with faster refresh rates for turbulence injection. It would also be helpful to repeat the hardware experiments across more D/r_0 values, more average RSSI levels, and perform a larger set of repetitions to obtain a more statistically significant validation. Ultimately, running outdoor measurements in actual atmospheric turbulence would be the best way to verify the algorithms experimentally.

Bibliography

- [1] Nils Pachler et al. “An Updated Comparison of Four Low Earth Orbit Satellite Constellation Systems to Provide Global Broadband”. In: *2021 IEEE International Conference on Communications Workshops, ICC Workshops 2021 - Proceedings*. IEEE, June 2021. ISBN: 9781728194417. DOI: [10.1109/ICCWORKSHOPS50388.2021.9473799](https://doi.org/10.1109/ICCWORKSHOPS50388.2021.9473799).
- [2] Carlos Carrizo et al. “Optical inter-satellite link terminals for next generation satellite constellations”. In: *Free-Space Laser Communications XXXII*. Vol. 11272. 2. SPIE, Mar. 2020, p. 2. ISBN: 9781510633070. DOI: [10.1117/12.2545629](https://doi.org/10.1117/12.2545629).
- [3] Doug Mohny. *Space lasers, smart antennas are key enablers for satellite broadband*. Mar. 2021. URL: <https://datacenterfrontier.com/space-lasers-smart-antennas-are-key-enablers-for-satellite-broadband/> (visited on 07/07/2021).
- [4] Mahesh Krishnaswamy. *Bringing light-speed internet to Sub-Saharan Africa*. Nov. 2020. URL: <https://blog.x.company/bringing-light-speed-internet-to-sub-saharan-africa-4e022e1154ca> (visited on 05/15/2021).
- [5] Baris Erkmen. *Beaming broadband across the Congo River*. Sept. 2021. URL: <https://x.company/blog/posts/taara-beaming-broadband-across-congo/> (visited on 12/09/2022).
- [6] Bruce Moision et al. “Demonstration of free-space optical communication for long-range data links between balloons on Project Loon”. In: *Free-Space Laser Communication and Atmospheric Propagation XXIX*. Vol. 10096. SPIE, Feb. 2017, pp. 259–272. ISBN: 9781510606333. DOI: [10.1117/12.2253099](https://doi.org/10.1117/12.2253099).
- [7] Catherine Shu. *Singapore-based Transcelestial uses lasers to build affordable internet networks*. Feb. 2023. URL: <https://techcrunch.com/2023/02/16/transcelestial/> (visited on 02/21/2023).
- [8] Rohit Jha. *Transcelestial raises Series A Funding to tackle Global Last Mile Connectivity Challenges*. July 2020. URL: <https://www.transcelestial.com/blog/transcelestial-raises-series-a-funding-to-tackle-global-last-mile-connectivity-challenges> (visited on 12/09/2022).
- [9] Shaojuan Zhang, Eduward Tangdionga, and Nicola Calabretta. “A Novel Optical Wireless Data Center Network Architecture Based on Passive Diffractive Optics and Fast Tunable Transmitters”. In: *OSA Advanced Photonics Congress*. Optica Publishing Group, July 2020. DOI: [10.1364/NETWORKS.2020.NEM2B.3](https://doi.org/10.1364/NETWORKS.2020.NEM2B.3).

- [10] Abdelbaset S. Hamza et al. “OWCell: Optical wireless cellular data center network architecture”. In: *IEEE International Conference on Communications*. IEEE, July 2017. ISBN: 9781467389990. DOI: [10.1109/ICC.2017.7996501](https://doi.org/10.1109/ICC.2017.7996501).
- [11] Magdalena Garlinska et al. “From Mirrors to Free-Space Optical Communication—Historical Aspects in Data Transmission”. In: *Future Internet* 12.11 (Oct. 2020), p. 179. ISSN: 1999-5903. DOI: [10.3390/FI12110179](https://doi.org/10.3390/FI12110179).
- [12] Mohamed Alzenad et al. “FSO-Based Vertical Backhaul/Fronthaul Framework for 5G+ Wireless Networks”. In: *IEEE Communications Magazine* 56.1 (Jan. 2018), pp. 218–224. ISSN: 15581896. DOI: [10.1109/MCOM.2017.1600735](https://doi.org/10.1109/MCOM.2017.1600735). arXiv: [1607.01472](https://arxiv.org/abs/1607.01472).
- [13] John E. Kaufmann. “Performance limits of high-rate space-to-ground optical communications through the turbulent atmospheric channel”. In: *Free-Space Laser Communication Technologies VII*. Vol. 2381. 20. SPIE, Apr. 1995, pp. 171–182. DOI: [10.1117/12.207402](https://doi.org/10.1117/12.207402).
- [14] Hamid Hemmati. *Near-earth laser communications*. 2nd ed. CRC Press, 2020. ISBN: 9781420015447. DOI: [10.1201/9780429186721](https://doi.org/10.1201/9780429186721).
- [15] Jason Stewart et al. “Comparing adaptive optics approaches for NASA LCRD Ground Station 2”. In: *Free-Space Laser Communication and Atmospheric Propagation XXV*. Vol. 8610. SPIE, Mar. 2013, pp. 126–136. ISBN: 9780819493798. DOI: [10.1117/12.2010126](https://doi.org/10.1117/12.2010126).
- [16] Max Segel et al. “Optimal modal compensation in gradient-based wavefront sensorless adaptive optics”. In: *Laser Communication and Propagation through the Atmosphere and Oceans VIII*. Vol. 11133. 6. SPIE, Sept. 2019, pp. 185–193. ISBN: 9781510629592. DOI: [10.1117/12.2531121](https://doi.org/10.1117/12.2531121).
- [17] Italo Toselli and Szymon Gladysz. “Influence of bandwidth error on the performance of adaptive optics systems for uncooperative beacons”. In: *Applied Optics* 60.22 (Aug. 2021), F118–F125. ISSN: 2155-3165. DOI: [10.1364/AO.427678](https://doi.org/10.1364/AO.427678).
- [18] Max Segel and Szymon Gladysz. “Optimal, blind-search modal wavefront correction in atmospheric turbulence. Part I: simulations”. In: *Optics Express* 29.2 (Jan. 2021), pp. 805–820. ISSN: 1094-4087. DOI: [10.1364/OE.408682](https://doi.org/10.1364/OE.408682).
- [19] Ondrej Čierny, Bruce Moision, and Kerri Cahoy. “Evaluating the performance of a sensorless wavefront correction algorithm for turbulent horizontal point-to-point links”. In: *Free-Space Laser Communications XXXV*. Vol. 12413. 15. SPIE, Mar. 2023, pp. 296–310. DOI: [10.1117/12.2651076](https://doi.org/10.1117/12.2651076).
- [20] Harald Hauschildt et al. “European data relay system goes global”. In: *2017 IEEE International Conference on Space Optical Systems and Applications, ICSOS 2017*. Institute of Electrical and Electronics Engineers Inc., May 2018, pp. 15–18. ISBN: 9781509065110. DOI: [10.1109/ICSOS.2017.8357204](https://doi.org/10.1109/ICSOS.2017.8357204).
- [21] Uwe Sterr, Mark Gregory, and Frank Heine. “Beaconless acquisition for ISL and SGL, summary of 3 years operation in space and on ground”. In: *2011 International Conference on Space Optical Systems and Applications, ICSOS’11*. 2011, pp. 38–43. ISBN: 9781424496853. DOI: [10.1109/ICSOS.2011.5783704](https://doi.org/10.1109/ICSOS.2011.5783704).

- [22] Daniel Dallmann et al. “GEO-LEO beaconless spatial acquisition reality in space”. In: *2015 IEEE International Conference on Space Optical Systems and Applications, ICSOS 2015*. IEEE, Mar. 2016. ISBN: 9781509002818. DOI: [10.1109/ICSOS.2015.7425068](https://doi.org/10.1109/ICSOS.2015.7425068).
- [23] Frank F. Heine et al. “In orbit performance of Tesat LCTs”. In: *Free-Space Laser Communications XXXI*. Vol. 10910. SPIE, Mar. 2019, p. 29. DOI: [10.1117/12.2510721](https://doi.org/10.1117/12.2510721).
- [24] Frank F. Heine et al. “Status of Tesat lasercomms activities”. In: *Free-Space Laser Communications XXXIII*. Vol. 11678. 5. SPIE, Mar. 2021, p. 8. ISBN: 9781510641914. DOI: [10.1117/12.2577609](https://doi.org/10.1117/12.2577609).
- [25] M. Guelman et al. “Acquisition and pointing control for inter-satellite laser communications”. In: *IEEE Transactions on Aerospace and Electronic Systems* 40.4 (Oct. 2004), pp. 1239–1248. ISSN: 00189251. DOI: [10.1109/TAES.2004.1386877](https://doi.org/10.1109/TAES.2004.1386877).
- [26] Peter LoPresti et al. “Adaptive divergence and power for improving connectivity in free-space optical mobile networks”. In: *Applied Optics* 45.25 (Sept. 2006), pp. 6591–6597. ISSN: 15394522. DOI: [10.1364/AO.45.006591](https://doi.org/10.1364/AO.45.006591).
- [27] K. H. Heng et al. “Adaptive beam divergence for inter-UAV free space optical communications”. In: *2008 IEEE Photonics Global at Singapore, IPGC 2008*. 2008. ISBN: 9781424429059. DOI: [10.1109/IPGC.2008.4781473](https://doi.org/10.1109/IPGC.2008.4781473).
- [28] Keith M. Hinrichs et al. “Continuous beam divergence control via wedge-pair for laser communication applications”. In: *Free-Space Laser Communication and Atmospheric Propagation XXVII*. Vol. 9354. SPIE, Mar. 2015, p. 93540L. ISBN: 9781628414448. DOI: [10.1117/12.2084557](https://doi.org/10.1117/12.2084557).
- [29] Jeffrey M. Roth et al. “Variable, two-color acquisition beam for free-space laser communication terminals”. In: *Free-Space Laser Communications XXXI*. Vol. 10910. SPIE, Mar. 2019, p. 25. ISBN: 9781510624627. DOI: [10.1117/12.2513964](https://doi.org/10.1117/12.2513964).
- [30] Dirk Giggenbach, Anton Schex, and Bernhard Wandernoth. “Prototype of a coherent tracking and detection receiver with wideband vibration compensation for free-space laser communications”. In: *Free-Space Laser Communication Technologies VIII*. Ed. by G. Stephen Mecherle. Vol. 2699. 22. SPIE, Apr. 1996, p. 186. DOI: [10.1117/12.238411](https://doi.org/10.1117/12.238411).
- [31] Eric A. Swanson and Vincent W.S. Chan. “Heterodyne spatial tracking system for optical space communication”. In: *IEEE Transactions on Communications* 34.2 (1986), pp. 118–126. ISSN: 00906778. DOI: [10.1109/TCOM.1986.1096501](https://doi.org/10.1109/TCOM.1986.1096501).
- [32] Mike S. Ferraro et al. “Position sensing and high bandwidth data communication using impact ionization engineered APD arrays”. In: *IEEE Photonics Technology Letters* 31.1 (Jan. 2019), pp. 58–61. ISSN: 10411135. DOI: [10.1109/LPT.2018.2882886](https://doi.org/10.1109/LPT.2018.2882886).
- [33] Taylor A. Page et al. “A system overview of a small form factor free space optical communication prototype”. In: *Free-Space Laser Communications XXXII*. Ed. by Hamid Hemmati and Don M. Boroson. Vol. 11272. SPIE, Mar. 2020, p. 20. ISBN: 9781510633070. DOI: [10.1117/12.2551310](https://doi.org/10.1117/12.2551310).
- [34] Shannon N. Suddath. “Fiber-coupled wide field of view optical receiver for high speed space communication”. S.M. thesis. The University of Tulsa, 2017. URL: <https://www.proquest.com/docview/1894626538>.

- [35] Peter Barcik et al. “Concept of a fiber-based laser beam tracking system for a free space optical link”. In: *2018 11th International Symposium on Communication Systems, Networks and Digital Signal Processing, CSNDSP 2018*. IEEE, Sept. 2018. ISBN: 9781538613351. DOI: [10.1109/CSNDSP.2018.8471850](https://doi.org/10.1109/CSNDSP.2018.8471850).
- [36] Vuong V. Mai and Hoon Kim. “Wide field-of-view transceiver design for bidirectional free-space optical communication systems”. In: *OECC/PSC 2019 - 24th OptoElectronics and Communications Conference/International Conference Photonics in Switching and Computing 2019*. IEEE, July 2019. ISBN: 9784885523212. DOI: [10.23919/PS.2019.8818046](https://doi.org/10.23919/PS.2019.8818046).
- [37] Eric A. Swanson and Roy S. Bondurant. *A space-based optical communication system utilizing fiber optics*. Tech. rep. MIT Lincoln Laboratory, Nov. 1989. URL: <https://apps.dtic.mil/sti/citations/ADA217333>.
- [38] Ulrich A. Johann, Klaus Pribil, and Heinz Sontag. “Novel optical-fiber-based conical scan tracking device”. In: *Optical Space Communication II*. Ed. by Juergen Franz. Vol. 1522. SPIE, May 1991, p. 243. DOI: [10.1117/12.46079](https://doi.org/10.1117/12.46079).
- [39] Xueqiang Zhao et al. “Experimental verification of coherent tracking system based on fiber nutation”. In: *Optics Express* 27.17 (Aug. 2019), p. 23996. ISSN: 1094-4087. DOI: [10.1364/oe.27.023996](https://doi.org/10.1364/oe.27.023996).
- [40] Todd E. Knibbe, Eric A. Swanson, and James K. Roberge. “Spatial tracking using an electro-optic nutator and a single-mode optical fiber”. In: *Free-Space Laser Communication Technologies IV*. Ed. by David L. Begley and Bernard D. Seery. Vol. 1635. SPIE, June 1992, p. 309. DOI: [10.1117/12.59284](https://doi.org/10.1117/12.59284).
- [41] Robert J. Murphy et al. “A conical scan free space optical tracking system for fading channels”. In: *Free-Space Laser Communications IX*. Ed. by Arun K. Majumdar and Christopher C. Davis. Vol. 7464. 21. SPIE, Aug. 2009, 74640P. DOI: [10.1117/12.826299](https://doi.org/10.1117/12.826299).
- [42] Chenzhe Lao et al. “Large field of view beaconless laser nutation tracking sensor based on a micro-electro-mechanical system mirror”. In: *Applied Optics* 59.22 (Aug. 2020), p. 6534. ISSN: 1559-128X. DOI: [10.1364/ao.396490](https://doi.org/10.1364/ao.396490).
- [43] Ondrej Čierny, Paul C. Serra, and Kerri L. Cahoy. “Beaconless pointing and tracking for bidirectional optical links using MEMS mirror nutation”. In: *Free-Space Laser Communications XXXIII*. Vol. 11678. 5. SPIE, Mar. 2021, p. 14. ISBN: 9781510641914. DOI: [10.1117/12.2582615](https://doi.org/10.1117/12.2582615).
- [44] Cyril Ruilier. “A study of degraded light coupling into single-mode fibers”. In: *Astronomical Interferometry*. Vol. 3350. 24. International Society for Optics and Photonics, July 1998, pp. 319–329. DOI: [10.1117/12.317094](https://doi.org/10.1117/12.317094).
- [45] David F. Buscher. *Practical Optical Interferometry: Imaging at Visible and Infrared Wavelengths*. Cambridge University Press, Jan. 2015, pp. 1–267. ISBN: 9781107615281. DOI: [10.1017/CBO9781107323933](https://doi.org/10.1017/CBO9781107323933).
- [46] Larry C. Andrews and Ronald L. Phillips. *Laser Beam Propagation through Random Media*. 2nd ed. SPIE Press, Sept. 2005. ISBN: 9781510643703. DOI: [10.1117/3.626196](https://doi.org/10.1117/3.626196).

- [47] Jeffrey D. Barchers and David L. Fried. “Optimal control of laser beams for propagation through a turbulent medium”. In: *Journal of the Optical Society of America A* 19.9 (Sept. 2002), pp. 1779–1793. ISSN: 1520-8532. DOI: [10.1364/JOSAA.19.001779](https://doi.org/10.1364/JOSAA.19.001779).
- [48] Jeffrey D. Barchers, David L. Fried, and Donald J. Link. “Evaluation of the performance of Hartmann sensors in strong scintillation”. In: *Applied Optics* 41.6 (Feb. 2002), pp. 1012–1021. ISSN: 2155-3165. DOI: [10.1364/AO.41.001012](https://doi.org/10.1364/AO.41.001012).
- [49] Joshua H. Follansbee et al. “Measuring phase errors in the presence of scintillation”. In: *Optics Express* 28.25 (Dec. 2020), pp. 37721–37733. ISSN: 1094-4087. DOI: [10.1364/OE.408825](https://doi.org/10.1364/OE.408825). arXiv: [2012.12695](https://arxiv.org/abs/2012.12695).
- [50] James Notaras, Carl Paterson James Notaras, and Carl Paterson. “Direct wavefront sensing for atmospheric adaptive optics in strong scintillation”. In: *Advanced Wavefront Control: Methods, Devices, and Applications III*. Vol. 5894. SPIE, Aug. 2005, pp. 229–237. DOI: [10.1117/12.617077](https://doi.org/10.1117/12.617077).
- [51] Michael J. Steinbock. “Implementation of Branch-Point-Tolerant Wavefront Reconstructor for Strong Turbulence Compensation”. S.M. thesis. Air Force Institute of Technology, 2012. URL: <https://apps.dtic.mil/sti/citations/ADA562776>.
- [52] Thorlabs Inc. *Shack-Hartmann Wavefront Sensors*. URL: https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=5287 (visited on 12/08/2022).
- [53] T. R. O’Meara. “The multidither principle in adaptive optics”. In: *Journal of the Optical Society of America* 67.3 (Mar. 1977), p. 306. ISSN: 0030-3941. DOI: [10.1364/josa.67.000306](https://doi.org/10.1364/josa.67.000306).
- [54] Dimitrios N. Loizos, Paul P. Sotiriadis, and Gert Cauwenberghs. “A robust continuous-time multi-dithering technique for laser communications using adaptive optics”. In: *IEEE International Symposium on Circuits and Systems*. IEEE, 2006, pp. 3626–3629. ISBN: 0780393902. DOI: [10.1109/ISCAS.2006.1693412](https://doi.org/10.1109/ISCAS.2006.1693412).
- [55] M. A. Vorontsov, G. W. Carhart, and J. C. Ricklin. “Adaptive phase-distortion correction based on parallel gradient-descent optimization”. In: *Optics Letters* 22.12 (June 1997), p. 907. ISSN: 0146-9592. DOI: [10.1364/ol.22.000907](https://doi.org/10.1364/ol.22.000907).
- [56] R. Mukai, K. Wilson, and V. Vilnrotter. “Application of genetic and gradient descent algorithms to wavefront compensation for the deep-space optical communications receiver”. In: *The Interplanetary Network Progress Report*. Vol. 42-161. 2006. URL: https://ipnpr.jpl.nasa.gov/progress_report/42-161/161U.pdf.
- [57] Thomas Weyrauch et al. “Fiber coupling with adaptive optics for free-space optical communication”. In: *Free-Space Laser Communication and Laser Imaging*. Vol. 4489. 22. SPIE, Jan. 2002, pp. 177–184. DOI: [10.1117/12.453227](https://doi.org/10.1117/12.453227).
- [58] Piotr Piatrou and Michael Roggemann. “Beaconless stochastic parallel gradient descent laser beam control: Numerical experiments”. In: *Applied Optics* 46.27 (Sept. 2007), pp. 6831–6842. ISSN: 15394522. DOI: [10.1364/AO.46.006831](https://doi.org/10.1364/AO.46.006831).
- [59] Luca Rinaldi et al. “Sensorless adaptive optics for optical communications”. In: *Free-Space Laser Communications XXXIII*. Vol. 11678. 5. SPIE, Mar. 2021, p. 26. ISBN: 9781510641914. DOI: [10.1117/12.2582942](https://doi.org/10.1117/12.2582942).

- [60] Martin Booth. “Wave front sensor-less adaptive optics: a model-based approach using sphere packings”. In: *Optics Express* 14.4 (Feb. 2006), p. 1339. ISSN: 1094-4087. DOI: [10.1364/oe.14.001339](https://doi.org/10.1364/oe.14.001339).
- [61] Feiling Wang. “Control of deformable mirror with light-intensity measurements through single-mode fiber”. In: *Applied Optics* 49.31 (Nov. 2010), G60. ISSN: 0003-6935. DOI: [10.1364/ao.49.000g60](https://doi.org/10.1364/ao.49.000g60).
- [62] Ping Yang et al. “A novel wavefront sensing technique for high speed atmospheric measurement based on digital micromirror device”. In: *Optics in Atmospheric Propagation and Adaptive Systems XIV*. Vol. 8178. SPIE, Oct. 2011, 81780N. ISBN: 9780819488053. DOI: [10.1117/12.897566](https://doi.org/10.1117/12.897566).
- [63] Feiling Wang. “High-contrast imaging via modal convergence of deformable mirror”. In: *Astrophysical Journal* 751.2 (June 2012), p. 83. ISSN: 15384357. DOI: [10.1088/0004-637X/751/2/83](https://doi.org/10.1088/0004-637X/751/2/83).
- [64] M. A. Vorontsov and V. P. Sivokon. “Stochastic parallel-gradient-descent technique for high-resolution wave-front phase-distortion correction”. In: *Journal of the Optical Society of America A* 15.10 (Oct. 1998), p. 2745. ISSN: 1084-7529. DOI: [10.1364/josaa.15.002745](https://doi.org/10.1364/josaa.15.002745).
- [65] Esdras Anzuola et al. “Performance of wavefront-sensorless adaptive optics using modal and zonal correction”. In: *Optics in Atmospheric Propagation and Adaptive Systems XIX*. Vol. 10002. SPIE, Oct. 2016, pp. 133–140. ISBN: 9781510604087. DOI: [10.1117/12.2241067](https://doi.org/10.1117/12.2241067).
- [66] Guang-Ming Dai. “Wavefront simulation for atmospheric turbulence”. In: *Image Reconstruction and Restoration*. Vol. 2302. SPIE, Sept. 1994, pp. 62–72. DOI: [10.1117/12.188066](https://doi.org/10.1117/12.188066).
- [67] J. Y. Wang and J. K. Markey. “Modal compensation of atmospheric turbulence phase distortion”. In: *J. Opt. Soc. Am.* 68.1 (Jan. 1978), pp. 78–87. ISSN: 0030-3941. DOI: [10.1364/JOSA.68.000078](https://doi.org/10.1364/JOSA.68.000078).
- [68] Charles Lailabai Linslal et al. “Challenges in coherent beam combining of high power fiber amplifiers: a review”. In: *ISSS Journal of Micro and Smart Systems* 11.1 (Mar. 2022), pp. 277–293. ISSN: 2509-7997. DOI: [10.1007/S41683-022-00099-4](https://doi.org/10.1007/S41683-022-00099-4).
- [69] Gregory W. Allan. “Phasing of Ground-based Optical Arrays for Space Applications”. PhD thesis. Massachusetts Institute of Technology, 2022. URL: <https://dspace.mit.edu/handle/1721.1/145073>.
- [70] Thomas Shay and Vincent Benham. “First experimental demonstration of phase locking of optical fiber arrays by RF phase modulation”. In: *Free-Space Laser Communications IV*. Vol. 5550. SPIE, Oct. 2004, pp. 313–319. DOI: [10.1117/12.597455](https://doi.org/10.1117/12.597455).
- [71] Yanxing Ma et al. “Active phase locking of fiber amplifiers using sine-cosine single-frequency dithering technique”. In: *Applied Optics* 50.19 (July 2011), pp. 3330–3336. ISSN: 2155-3165. DOI: [10.1364/AO.50.003330](https://doi.org/10.1364/AO.50.003330).

- [72] Xiaolin Wang et al. “Coherent beam combination of 1.08 kW fiber amplifier array using single frequency dithering technique”. In: *Optics Letters* 36.6 (Mar. 2011), pp. 951–953. ISSN: 1539-4794. DOI: [10.1364/OL.36.000951](https://doi.org/10.1364/OL.36.000951).
- [73] Thomas M. Shay et al. “Self-synchronous and self-referenced coherent beam combination for large optical arrays”. In: *IEEE Journal on Selected Topics in Quantum Electronics* 13.3 (May 2007), pp. 480–486. ISSN: 1077260X. DOI: [10.1109/JSTQE.2007.897173](https://doi.org/10.1109/JSTQE.2007.897173).
- [74] T. M. Shay et al. “High-power phase locking of a fiber amplifier array”. In: *Fiber Lasers VI: Technology, Systems, and Applications*. Vol. 7195. SPIE, Feb. 2009, pp. 365–372. DOI: [10.1117/12.809416](https://doi.org/10.1117/12.809416).
- [75] L. Lombard et al. “Orthogonal coding methods for increasing the number of multiplexed channels in coherent beam combining”. In: *Applied Optics* 53.8 (Mar. 2014), pp. 1493–1502. ISSN: 2155-3165. DOI: [10.1364/AO.53.001493](https://doi.org/10.1364/AO.53.001493).
- [76] Min Jiang et al. “Coherent beam combining of fiber lasers using a CDMA-based single-frequency dithering technique”. In: *Applied Optics* 56.15 (May 2017), pp. 4255–4260. ISSN: 2155-3165. DOI: [10.1364/AO.56.004255](https://doi.org/10.1364/AO.56.004255).
- [77] Paul G. Sibley et al. “Crosstalk reduction for multi-channel optical phase metrology”. In: *Optics Express* 28.7 (Mar. 2020), pp. 10400–10424. ISSN: 1094-4087. DOI: [10.1364/OE.388381](https://doi.org/10.1364/OE.388381).
- [78] Franz Fidler and Oswald Wallner. “Application of Single-Mode Fiber-Coupled Receivers in Optical Satellite to High-Altitude Platform Communications”. In: *EURASIP Journal on Wireless Communications and Networking* 2008.1 (Feb. 2008), pp. 1–7. ISSN: 1687-1499. DOI: [10.1155/2008/864031](https://doi.org/10.1155/2008/864031).
- [79] Virendra N. Mahajan. “Strehl ratio for primary aberrations in terms of their aberration variance”. In: *Journal of the Optical Society of America* 73.6 (June 1983), pp. 860–861. ISSN: 00303941. DOI: [10.1364/JOSA.73.000860](https://doi.org/10.1364/JOSA.73.000860).
- [80] T. Sean Ross. “Limitations and applicability of the Maréchal approximation”. In: *Applied Optics* 48.10 (Apr. 2009), pp. 1812–1818. ISSN: 2155-3165. DOI: [10.1364/AO.48.001812](https://doi.org/10.1364/AO.48.001812).
- [81] G. Schitter et al. “Model-based aberration correction in a closed-loop wavefront-sensor-less adaptive optics system”. In: *Optics Express* 18.23 (Nov. 2010), pp. 24070–24084. ISSN: 1094-4087. DOI: [10.1364/OE.18.024070](https://doi.org/10.1364/OE.18.024070).
- [82] Vladimir Sacek. *Strehl ratio*. July 2006. URL: <https://www.telescope-optics.net/Strehl.htm> (visited on 03/05/2023).
- [83] Boston Micromachines Corporation. *Standard Deformable Mirrors*. URL: <https://bostonmicromachines.com/products/deformable-mirrors/standard-deformable-mirrors/> (visited on 01/31/2023).
- [84] MZA. *Publications by MZA*. URL: <https://www.mza.com/technology.php?page=technologyPubs> (visited on 03/23/2023).
- [85] Charles R. Harris et al. “Array programming with NumPy”. In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. ISSN: 1476-4687. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2). arXiv: [2006.10256](https://arxiv.org/abs/2006.10256).

- [86] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. “Numba: A LLVM-based Python JIT Compiler”. In: *Proceedings of LLVM-HPC 2015: 2nd Workshop on the LLVM Compiler Infrastructure in HPC*. Vol. 2015-Janua. Association for Computing Machinery, Nov. 2015. ISBN: 9781450340052. DOI: [10.1145/2833157.2833162](https://doi.org/10.1145/2833157.2833162).
- [87] Pauli Virtanen et al. “SciPy 1.0: fundamental algorithms for scientific computing in Python”. In: *Nature Methods* 17.3 (Feb. 2020), pp. 261–272. ISSN: 1548-7105. DOI: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2). arXiv: [1907.10121](https://arxiv.org/abs/1907.10121).
- [88] Katie Morzinski et al. “Performance of MEMS-based visible-light adaptive optics at Lick Observatory: closed- and open-loop control”. In: *SPIE Astronomical Telescopes + Instrumentation, Adaptive Optics Systems II*. Vol. 7736. 27. SPIE, July 2010, pp. 639–654. ISBN: 9780819482266. DOI: [10.1117/12.857444](https://doi.org/10.1117/12.857444).
- [89] Guang-ming Dai and Virendra N. Mahajan. “Nonrecursive determination of orthonormal polynomials with matrix formulation”. In: *Optics Letters* 32.1 (Jan. 2007), pp. 74–76. ISSN: 1539-4794. DOI: [10.1364/OL.32.000074](https://doi.org/10.1364/OL.32.000074).
- [90] Jingfei Ye et al. “Comparative assessment of orthogonal polynomials for wavefront reconstruction over the square aperture”. In: *JOSA A* 31.10 (Oct. 2014), pp. 2304–2311. ISSN: 1520-8532. DOI: [10.1364/JOSAA.31.002304](https://doi.org/10.1364/JOSAA.31.002304).
- [91] Carlos E. Carrizo. “Advanced adaptive compensation system for free-space optical communications”. PhD thesis. Universitat Politècnica de Catalunya, Sept. 2019. URL: <http://hdl.handle.net/2117/168867>.
- [92] Long Zhu and Jian Wang. “Arbitrary manipulation of spatial amplitude and phase using phase-only spatial light modulators”. In: *Nature Scientific Reports* 4.1 (Dec. 2014), pp. 1–7. ISSN: 2045-2322. DOI: [10.1038/srep07441](https://doi.org/10.1038/srep07441).
- [93] Paul Balondrade and Sébastien M. Popoff. *Calibration of Linearly Aligned Nematic Liquid Crystal based Spatial Light Modulators*. Tech. rep. Sept. 2018. DOI: [10.5281/ZENODO.1423222](https://doi.org/10.5281/ZENODO.1423222).