

**The Effect of Differential Color Refraction on Astrometric
Observations of Solar System Bodies and Earth Satellites from
Ground-Based Optical Telescopes**

by

Roman O. Geykhman

B.S.E, B.A., University of Pennsylvania (2007)
S.M., Massachusetts Institute of Technology (2011)

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

September 2019

© Massachusetts Institute of Technology 2019. All rights reserved.

Author
Department of Aeronautics and Astronautics
August 1, 2019

Certified by.....
Kerri Cahoy
Associate Professor of Aeronautics and Astronautics
Thesis Supervisor

Certified by.....
Richard Binzel
Professor of Planetary Sciences

Certified by.....
Paul Chodas
Director, Center for Near Earth Object Studies, NASA Jet Propulsion Laboratory

Certified by.....
Joseph Scott Stuart
Technical Staff, MIT Lincoln Laboratory

Accepted by
Sertac Karaman
Associate Professor of Aeronautics and Astronautics
Chairman, Graduate Program Committee

The Effect of Differential Color Refraction on Astrometric Observations of Solar System Bodies and Earth Satellites from Ground-Based Optical Telescopes

by

Roman O. Geykhman

Submitted to the Department of Aeronautics and Astronautics
on August 1, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Aeronautics and Astronautics

Abstract

Earth's atmosphere is optically dispersive and subjects astrometric observations from ground-based optical telescopes to systematic bias from differential color refraction (DCR). This bias is evident in Minor Planet Center observations of asteroids with known spectral types and in observations of GPS and GLONASS satellites. DCR bias is on the order of 0.1 arcsec, and until recently, fixed-pattern star catalog errors exceeded this level. With the release of the Gaia DR2 star catalog in April of 2018, catalog error is no longer dominant and the systematic error floor in ground-based astrometry is defined by DCR.

Unaccounted-for DCR bias in observations can introduce a small but statistically significant bias into the estimate of Keplerian mean motion of inner Solar System asteroids, reduce the probability of successfully observing a stellar occultation by a Kuiper Belt Object, and in rare pathological cases can mean the difference between predicting an impact or a miss by a hazardous asteroid. DCR in observations of geostationary satellites can introduce a large bias into the estimate of solar radiation pressure area-to-mass ratio in a single-night orbit fit and tens of meters of error into an orbit prediction derived from several nights of observation.

Measurements of the 2017 near-Earth flyby of the asteroid 3122 Florence from MIT and MIT Lincoln Laboratory facilities in Westford, MA and Socorro, NM suggest that narrow passbands are insufficient to mitigate DCR, and measurements of a sample of geostationary satellites' spectra at the US Naval Observatory Flagstaff Station show that the DCR bias of active satellites can vary by up to 0.1 arcsec over half an hour. While the DCR of fiducial stars is predictable from catalog data, satellites' DCR must be measured directly. To that end, a slitless spectrograph was deployed at the Firepond Optical Facility in Westford, MA and observed GPS and GLONASS satellites over seven nights. Using that data, I demonstrate DCR compensation yielding a 60% reduction in bias and 30% reduction in noise in astrometric residuals relative to color-agnostic processing when all atmosphere-induced effects (stellar DCR, target DCR, and parallactic refraction) are accounted for.

Disclaimer: This material is based upon work supported by the National Aeronautics and Space Administration (NASA) and the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA) under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of NASA, ODNI, IARPA or the U.S. Air Force.

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

Thesis Supervisor: Kerri Cahoy

Title: Associate Professor of Aeronautics and Astronautics

Acknowledgments

This document includes content reproduced from the following sources, used with permission

- R. Geykhman and K. Cahoy. Spectroscopic measurements of asteroids allow mitigation of differential color refraction effects on ground-based astrometry and orbit prediction accuracy. In *Proc.SPIE* 10702, Ground-based and Airborne Instrumentation for Astronomy VII, 107024U (6 July 2018); <https://doi.org/10.1117/12.2312230>
- R. Geykhman and K. Cahoy. The effect of differential color refraction on short-arc estimation of the area-to-mass ratio of geostationary objects using ground-based telescopes. In *Astrody-namics 2018. Proceedings of the AAS/AIAA Astrodynamics Specialist Conference held August 19–23, 2018, Snowbird, Utah, U.S.A.*, pp 1975–1994. Univelt, 2019. ISBN 978-0-87703-657-9
- R. Geykhman and K. Cahoy. The Effect of Differential Color Refraction on Asteroid Orbital Parameter Estimates. *Submitted for publication*, 2019

Although these results have only been published in the past year, it would be correct to say I have been working on the research described in this document in one form or another since October 2008. Back then, as a kid barely out of school, I volunteered for the engineering team tasked with reactivating the Firepond Optical Facility, and all of this flowed from there. There was no straight path to here, it took a while for all of the pieces to come together, and I have many people to thank for their guidance, assistance, and support during the intervening years.

The longest list of names is that of the people whose effort was necessary to collect the many hundreds of gigabytes of raw data used for this thesis. This data was collected over many years and many nights at many sites. For the data from the MIT Lincoln Laboratory Experimental Test Site in Socorro, New Mexico, I would like to thank Greg Spitz and Matthew Blythe. For the data from the MIT George R. Wallace Astrophysical Observatory in Westford, Massachusetts, I am grateful to Tim Brothers and Mike Person. For the data from the US Naval Observatory Flagstaff Station, I would like to thank Dave Monet. For the data from the MIT Lincoln Laboratory Firepond Optical Facility in Westford, Massachusetts, the list of people is the longest and presented in no particular order: Armand O'Donnell, Pete Trujillo, Wayne Robertson, Julie Johnson, Alex Szabo, Joseph Swick, Rodney Case, Dan Bonhomme, Mark Derome, Jon Byford, Jim Traffie, Bruce Whittier, Steve Cummiskey, Nik Waggener, Rebecca Busacker, Dave Radue, Jan Kansky, Phil Maynard, Carlos Gonzalez Huertas, Paul Pepin, Dan Birchall, Meg Noah, Gene Rork, Jack Fleischman, Steve Hunt, and Larry Swezey.

I would also like to thank my doctoral committee and Lincoln Scholars advising team for their input over the past few years. To Kerri Cahoy, Paul Chodas, Rick Binzel, Scott Stuart, Arthur Lue, Deborah Woods, and Zach Folcik: Thank you all.

For the financial support of this thesis, I would like to thank the Lincoln Scholars Program and specifically Jim Ward, Tim Hall, Sumanth Kaushik, and Larry Candell for their vote of confidence in approving my dive back into school. I would also like to extend thanks to the David Karp Memorial Fund at Lincoln Laboratory for the support of my studies and to Jon Ashcom for helping to secure the funding for the construction and operation of the spectrograph at Firepond.

In closing, I would like to thank my parents Oskar and Eugenia Geykhman and my grandparents Viktor Serman and Liliya Kamchatova for raising me right, teaching me to work, and sending me to school, and my wife Rena Zheng for everything else that came after.

Contents

Front Matter	1
Abstract	3
Acknowledgements	5
Contents	7
List of Figures	15
List of Tables	23
List of Abbreviations	25
Notation	27
1 Introduction	29
1.1 Phenomenology	29
1.2 Evidence of DCR Bias and Its Significance	31
1.2.1 Astrometric Observations of Calibration Satellites in Earth Orbit	31
1.2.2 Astrometric Observations of Asteroids with Known Spectral Types	35
1.3 Motivation	37
1.3.1 Improvements in Star Catalog Accuracy	37
1.3.2 Asteroid Search	38
1.3.2.1 Hazardous Asteroid Impact Monitoring	38
1.3.2.2 Yarkovsky Effect Characterization	42
1.3.3 Kuiper Belt Object Occultations	43
1.3.4 Near-Earth SSA	44
1.3.4.1 Satellite Debris Monitoring in Earth Orbit	44
1.3.4.2 Resolved Imaging of Geostationary Satellites	45
1.4 Gap Analysis and Thesis Contributions	46
1.5 Thesis Outline	47

2 Literature Review	49
2.1 Astrometry Through Time	49
2.1.1 Eyeball Era	49
2.1.2 Photographic Plate Era	50
2.1.3 Electronic Era	50
2.1.4 Sources of Measurement Error in CCD Astrometry	52
2.2 Star Catalog Accuracy	53
2.2.1 Photographic Plate Catalogs	54
2.2.2 Space-Based Catalogs	54
2.2.3 Ground-Based CCD Catalogs	55
2.2.4 Observation Refinement With A-Posteriori Catalog Corrections	55
2.3 Spectroscopic Measurements of Solar System Bodies	56
2.3.1 Past and Current Surveys of Asteroids	56
2.3.2 Kuiper Belt Objects	59
2.4 Spectroscopic Measurements of Satellites	59
2.4.1 Satellite Materials	59
2.4.2 In-Situ Measurements	60
2.5 Atmospheric Refraction	61
2.5.1 Models of Refraction and Their Applicability	61
2.5.2 Differential Color Refraction (DCR)	61
2.6 Ephemeris Propagation	62
2.6.1 Heliocentric Force Model	62
2.6.2 Geocentric Force Model	63
3 Methods	65
3.1 Overall Theme	65
3.2 Data Sources	66
3.2.1 Reduced Data and Ephemerides	66
3.2.1.1 International Astronomical Union Minor Planet Center	66
3.2.1.2 Jet Propulsion Laboratory	66
3.2.1.3 International GNSS Service	66
3.2.2 On-Sky Measurements	67
3.2.2.1 MIT George R. Wallace, Jr. Astrophysical Observatory	67
3.2.2.2 United States Naval Observatory Flagstaff Station	68
3.2.2.3 MIT Lincoln Laboratory Experimental Test System	68

3.2.2.4 MIT Lincoln Laboratory Firepond Optical Facility	69
3.3 The OrbitTool Software Framework	70
3.3.1 Design Philosophy	71
3.3.2 Coordinate System Definitions	71
3.3.3 Propagation of Ephemerides	74
3.3.4 Optical Observation Model	76
3.3.4.1 Aberration	76
3.3.4.2 Refraction and Parallax	78
3.3.5 Orbit Fitting	79
3.3.6 Validation	82
3.4 Astrometric Algorithms	87
3.4.1 Image Segmentation	87
3.4.1.1 Pixel Categorization	87
3.4.1.2 Pixel Aggregation	89
3.4.2 Point Spread Function Fitting	90
3.4.3 Star Matching	92
4 Predictions, Analysis, and Preliminary Measurements	95
4.1 DCR Bias of Stars	95
4.1.1 Blackbody Approximation Method	96
4.1.1.1 Simplified Approach	96
4.1.1.2 Choosing The Blackbody Temperature	97
4.1.2 The Effect of Interstellar Reddening	102
4.1.3 Multi-Color Direct Polynomial Approximation	107
4.1.3.1 UCAC Color Indices	107
4.1.3.2 Gaia DR2 Three Color and Parallax	110
4.1.3.3 DCR of the Average Star and the Merits of Per-Star DCR Compensation	110
4.2 DCR Bias of SSA Targets	113
4.2.1 Asteroids	113
4.2.2 Kuiper Belt Objects	116
4.2.3 Earth Satellites	116
4.3 Extended Analysis of Minor Planet Center Observations of Asteroids With Known Spectral Types	117
4.4 Re-Reductions of LINEAR Imagery	126

4.5	Preliminary Measurements on 3122 Florence	128
4.5.1	Scenario	128
4.5.2	Timing Validation	129
4.5.3	Ephemeris Refinement	130
4.5.4	Results	131
4.6	Spectroscopic Measurement Requirements	135
4.6.1	Resolution Requirement for Satellites, Asteroids, and KBOs	135
4.6.2	Resolution Requirement for Asteroids with Unmodeled Absorption Bands	136
4.6.3	Accuracy Requirement in Terms of Time-Aperture Product	139
4.7	Chapter Summary	145
5	Simulations of Orbit Solutions from DCR-Biased Observations	147
5.1	Inner Solar System Asteroids	147
5.1.1	Simulation Scenario	147
5.1.2	Mean Motion Bias and Recovery Residual for Single-Apparition Fits	149
5.1.3	Masked Signature of DCR	150
5.1.4	Exploration of the Orbital Element Space	158
5.2	Spurious Detection of Yarkovsky Acceleration	159
5.2.1	Sensitivity of Yarkovsky Term to Observation Bias	159
5.2.2	Physically-Plausible Yarkovsky Accelerations	161
5.2.3	Existence of Orbits Susceptible to Spurious Yarkovsky Detections	163
5.2.4	Post-Fit Residuals of Orbit Fits with (Spurious) Yarkovsky Acceleration	166
5.2.5	Likelihood of a Scenario Susceptible to Spurious Yarkovsky Acceleration	168
5.3	Impact Monitoring of Potentially Hazardous Asteroids	169
5.3.1	False Confidence in Optics-Only Predictions of Near Misses: An Example	169
5.3.1.1	Bias in the B-Plane	170
5.3.1.2	Estimated Impact Probability vs. Observation Elevation Limit	172
5.3.1.3	Impact Probabilities with Albedo-Only DCR Estimation	173
5.3.2	Impact Point Bias of Known Impactors	174
5.3.2.1	Evolution of the Impact Point Estimate of the Apophis-Like Example After 2029	174
5.3.2.2	Short-Timescale Evolution of Impact Point for Newly Discov- ered Impactors	178
5.4	Prediction of Stellar Occultations by Kuiper Belt Objects	180
5.4.1	Scenario	181

5.4.2	Simulation Results	182
5.4.3	Conclusions	188
5.5	DCR Bias in Geostationary Orbits	188
5.5.1	Scenario	189
5.5.2	Cross-Track Prediction Accuracy	189
5.5.3	Area-to-Mass Ratio Bias	191
5.5.4	Sub-20 Meter Conjunction Predictions	196
5.5.5	Empirical Validation of Predictability of GEO Appulses	196
5.5.5.1	Ballistic Target Selection	197
5.5.5.2	One-Night AMR Estimation	197
5.5.5.3	Cross-Track Predictability	198
5.6	Chapter Summary	199
6	DCR Compensation of Observations of GPS and GLONASS Satellites	
	with Concurrent Spectroscopy at the Firepond Optical Facility	201
6.1	Slitless Spectroscopy	202
6.1.1	Spectrograph Design Considerations	202
6.1.1.1	Mechanisms to Divide Light	202
6.1.1.2	Grating Spectrographs	203
6.1.2	Transmission Grating Geometry	205
6.1.3	Grating Efficiency Calculation	206
6.2	Slitless Spectroscopy of Geostationary Satellites at the US Naval Observatory	
	Flagstaff Station	208
6.2.1	Slitless Spectroscopy Processing	209
6.2.2	Results from NOFS Measurements	211
6.2.3	Conclusions	213
6.3	Multi-band Photometry of GPS and GLONASS Satellites	213
6.3.1	Scenario	213
6.3.2	Results from Multiband Photometry Experiment	214
6.4	Design of the Firepond Slitless Spectrograph	216
6.5	Performance Predictions	218
6.5.1	SNR and Suppression of Contamination by Higher Spectral Orders	218
6.5.2	SNR of $m = 0$ PSF	219
6.5.3	Clutter	221

6.6	Assembly and Laboratory Testing	223
6.6.1	Test Overview	223
6.6.2	Spot Projector	225
6.6.2.1	Breadboard Free-Space OAP Spot Projector	225
6.6.2.2	Commercial Mirror Lens Spot Projector	227
6.6.2.3	All-Reflecting Amateur Telescope Spot Projector	228
6.6.3	Mechanical Stability Validation	228
6.6.4	Absolute Photometric Calibration	229
6.6.4.1	Measurement Through Achromat Lens	230
6.6.4.2	Measurement of Spot Projector Transmittance and Pinhole Spectrum	232
6.6.5	Relative Photometric Calibrations	233
6.6.5.1	Blocking Filter Transmittance Validation	234
6.6.5.2	Grating Efficiency Validation	234
6.7	Installation, Checkouts, and On-Sky Calibration	238
6.7.1	Kinematic Stability Relative to 48" Telescope	240
6.7.1.1	The Necessity of Mount Bias Models	240
6.7.1.2	Empirical Bias Model for the Spectrograph	241
6.7.2	Wavelength Scale	242
6.7.3	Seeing and Achieved Resolution	246
6.7.4	Grating Tilt Calibration	247
6.7.5	Photometry Validation	248
6.7.5.1	Filter Transmittance Check	248
6.7.5.2	Total Spectral Transmittance Validation	249
6.7.5.3	Results of On-Sky Photometric Validation	250
6.7.6	Refraction Model Validation	252
6.7.6.1	Horizon Reference	253
6.7.6.2	Results of Spectral Refraction Validation	254
6.7.7	Discussion of On-Sky Calibration	255
6.8	Satellite Measurements at the Firepond Optical Facility	259
6.8.1	Planning Observations	259
6.8.2	Summary of Data Collections	260
6.8.2.1	Brightness of Targets	260
6.8.2.2	Sky Brightness	262
6.8.2.3	SNR Requirement vs. Data Availability	264

6.8.3 Spectroscopy Results	266
6.8.3.1 DCR Compensation with Spectroscopy	266
6.8.3.2 Analysis of DCR Prediction Error Statistics	268
6.8.3.3 Difference in Atmosphere vs Difference in Targets	270
6.8.3.4 Blue Glint	273
6.8.4 Experiment Summary	275
6.9 Chapter Summary	276
7 Conclusion	279
7.1 Summary of Original Contributions	279
7.1.1 Dominance of DCR Bias in Ground-Based Optical Astrometry	279
7.1.2 Consequences of Unaccounted-for DCR Bias in Orbit Solutions	280
7.1.3 Compensation of DCR Bias with Concurrent Spectroscopy	281
7.2 Recommendations	282
7.2.1 Astrometric Data Collection	282
7.2.2 Astrometric Data Reduction	282
7.2.3 Orbit Fitting	283
7.3 Future Work	284
7.3.1 Further Spectroscopic Measurements	284
7.3.1.1 Asteroids and KBOs	284
7.3.1.2 Earth Satellites	285
7.3.2 Handling Existing Data	285
7.3.2.1 Reprocessing Raw Imagery of Solar System Bodies	285
7.3.2.2 Color-Aware Farnocchia-Style Corrections	286
7.3.2.3 Maximum-Likelihood Orbit Determination	286
References	289

List of Figures

1-1	Schematic of DCR Bias	30
1-2	Accuracy Assessment of GEODSS Observations in the Late 1990s	32
1-3	Firepond Recommissioning Team, June 12, 2009	33
1-4	Residuals of Observations of GPS and GLONASS Satellites Reduced with URAT1	34
1-5	Residuals of Archived Observations of S-type and C-type Numbered Asteroids Resolved into Station-Local Horizontal/Vertical Coordinates	36
1-6	Orbit Geometry for Two-Dimensional Analysis of Significance of DCR on Asteroid Impact Predictions	39
1-7	Impact Point Bias in Two-Dimensional Analysis of $q = 0.9, Q = 1.0$ Orbit	41
1-8	Schematic of Yarkovsky Effect	42
1-9	Schematic of Stellar Occultation	43
1-10	Appulses of Guidestars by a Geostationary RSO	45
2-1	Bus-DeMeo Asteroid Taxonomy	58
2-2	Asteroid Reflectance Spectra for Selected Taxonomic Classes	58
2-3	Quaoar Reflectance Spectrum	59
2-4	Spectral Reflectances of Common Spacecraft Materials	60
3-1	MIT George R. Wallace Astrophysical Observatory, Westford, MA	67
3-2	1.3 Meter Telescope at USNO, Flagstaff, AZ	68
3-3	MIT Lincoln Laboratory Experimental Test Site, Socorro, NM	68
3-4	MIT Lincoln Laboratory Firepond Optical Facility, Westford, MA	69
3-5	Schematic of Local Right/Down/Out and East/North/Out Coordinates	73
3-6	Relativistic Aberration of Light in Optics Observation Model	77
3-7	Refraction and Parallax Refraction Geometry	78

3-8	Heliocentric Force Model Validation	83
3-9	Geocentric Force Model Validation	85
3-10	High-Pass Kernel	88
3-11	Schematic of Calculation of Charge Density and Pixel Signal for Streaking PSF	91
4-1	Empirical Observation of DCR of Stars at Firepond	97
4-2	Predicted DCR Error from Blackbody Spectrum Approximations at Westford, MA and Socorro, NM	98
4-3	Color-Color Diagram of Pickles Atlas Stars	100
4-4	Residual Error of DCR for Pickles Stars Approximated with DCR-Optimal Blackbody Temperature	100
4-5	DCR Error of Blackbody Spectrum Approximations Derived from Multi-color Photometry	101
4-6	Biases from Neglecting Interstellar Extinction	103
4-7	URAT1 Brightness and Color Statistics	106
4-8	DCR of Reddened Stellar Spectra vs. DCR of Blackbody Spectra Inferred from $B - V$ with Bare Silicon Detector at Sea Level	106
4-9	DCR Prediction Error vs. Elevation for Polynomial Function of UCAC5 Photometry	109
4-10	DCR Prediction Error vs. Elevation for Polynomial Function of Gaia DR2 Photometry and Stellar Parallax	111
4-11	DCR Bias for the Average Catalog Star Relative to Solar-Analog	112
4-12	DCR Bias for Asteroids vs. Elevation Angle	113
4-13	Sensitivity of Difference in S-type and C-type Asteroid DCR Bias to Meteorological Conditions and Detector Type	114
4-14	DCR Accuracy with Albedo-Based Classification	115
4-15	DCR Bias for KBOs vs. Elevation Angle.	116
4-16	DCR Bias Predictions for Satellites at Sea Level	117
4-17	Residuals of Bare CCD Farnocchia-Corrected Observations of Asteroids with Radar-Constrained Orbits and Known Spectral Types	119
4-18	Residuals of Narrow Passband and Broad Passband Farnocchia-Corrected Observations of Asteroids with Well-Constrained Optics-Only Orbits and Known Spectral Types	122

4-19	Residuals of Observations of Optics-Only Objects with Unspecified Pass- band Broken Down by Sites	123
4-20	Residuals of Observations of Optics-Only Objects with Johnson V Pass- band Broken Down by Sites	124
4-21	Residuals of Observations of Optics-Only Objects with Johnson R Pass- band Broken Down by Sites	125
4-22	Mean $\pm 3\sigma$ Horizontal and Vertical Residuals of S- and C-type Asteroids in Subset of LINEAR Data Reprocessed with UCAC5	126
4-23	Observing Geometry of 3122 Florence Flyby in 2017	128
4-24	Timing System Checkout for MIT LL Sensors and the MIT Wallace Observatory . .	130
4-25	Residuals of 3122 Florence Against JPL Ephemeris Dated Oct 21, 2017	131
4-26	Residuals of 3122 Florence Against Ephemeris Fit to Recent Data Only	132
4-27	Predictions and Observations of DCR for 3122 Florence with Narrow Passbands . . .	133
4-28	Reflectance Spectra of 3122 Florence and DCR in Sloan Passbands	134
4-29	Vertical Residuals of Firepond Bare Silicon Observations of 3122 Florence and Predicted DCR Bias	134
4-30	Spectroscopic Resolution Requirement Simulations for Satellites	135
4-31	Spectroscopic Resolution Requirement for Compensation of DCR Bias in Ground-Based Observations of Solar System Bodies with Bare-Silicon Detectors . . .	137
4-32	Resolution Required for Asteroid Spectra with Absorption Bands	138
4-33	Accuracy Requirement for Spectroscopic Compensation of DCR Bias in Ground-Based Observations with a Bare-Silicon Detector	144
5-1	Simulated Observing Network	148
5-2	Distribution Statistics of Simulated Post-Fit Residuals of Orbit Class- Representative Asteroids for Single-Apparition Fit	152
5-3	Distribution Statistics from a Single Station of Simulated Post-Fit Residuals	153
5-4	Single Apparition Fit Mean Motion Bias and Postfit Residuals for Asteroid Orbits with DCR Bias	154
5-5	Trade-off in Accuracy and Precision of Mean Motion Estimate in a Single- Apparition Fit	155
5-6	Trade-off in Accuracy of Mean Motion Estimate by Discarding Elevation Data	156

5-7	DCR-Induced Bias in In-Track Component of Recovery/Precovery Residual from a Single-Apparition Optics-Only Fit vs. SNR of Bias	157
5-8	Orbits with Statistically Significant Mean Motion Bias vs. Inclination, Semi-Major Axis, and Observation Lower Elevation Limit	158
5-9	Orbits with Statistically Significant Mean Motion Bias vs. Semi-Major Axis, Offset Between Ascending Node and Opposition, Season of Opposition, and Observation Lower Elevation Limit	159
5-10	Spurious Yarkovsky Signal in Two-Apparition Optics-Only Fit	160
5-11	Yarkovsky Effect Magnitudes of Asteroids in the JPL Database vs. Data Arc	161
5-12	Maximum Plausible Yarkovsky Acceleration vs. Brightness and Physical Size	163
5-13	Evolution of Best-Fit Yarkovsky Accelerations and Their 1σ Uncertainties with Increased Lower Elevation Limit for the Example Cases in Table 5.3	166
5-14	Mean and Standard Deviation of Mean Post-Fit Residual of Three-Apparition Optics-Only Asteroids in Table 5.3 vs. Lower Elevation Limit on Input Observations and Inclusion or Exclusion of Yarkovsky Acceleration, in Inertial and Station-Local Coordinates	167
5-15	Distribution in Time and Elevation of Observations Selected for Maximal Spurious Yarkovsky Detection Relative to Distribution of All Possible Observations	168
5-16	Monte Carlo Simulations of Optics-Only Predictions of 2036 B-Plane of Apophis-Like Impactor Derived from 13 Year Data Arc Prior to 2029	171
5-17	Observation Volume and Impact Estimate Probability vs. Elevation Limit for Hypothetical Apophis-Like Impactor	172
5-18	Asteroid Impact Point Uncertainty Calculation	176
5-19	Simulated Drift of Apparent Impact Point with True Impact of Apophis-Like Asteroid at Mumbai, India	177
5-20	Evolution of Impact Bias and Uncertainty with Time for 30 Day Running Window of Impact Point Estimates for a Hypothetical Newly-Discovered Asteroid	179
5-21	Evolution of Impact Bias and Uncertainty with Time for Florence-Like Impactor	180
5-22	Time-Evolution of Bias of Cross-Track Prediction and Formal its 3σ Bounds for Zero-Inclination KBO	183
5-23	Time-Evolution of Bias of Cross-Track Prediction and Formal its 3σ Bounds for 20° Inclined KBOs	184

5-24	Probability of Shadow Path Interception by $\pm 3\sigma$ Wide Line of Telescopes vs. KBO Orbital Parameters and Observation Lower Elevation Limit	186
5-25	Observation Volume and Occultation Measurement Success Probability vs. Elevation Limit for Selected $i = 30^\circ$ KBO Orbits	187
5-26	Time-Evolution of $\pm 3\sigma$ Along-Track and Cross-Track Bounds of GEO Or- bit Predictions from Four-Hour Arc vs. Observing Duty Factor and Orbit Geometry	190
5-27	GEO Next-Night Cross-Track Prediction Error for 2 Minutes Per Hour Observation Duty Cycle Over Four-Hour Arc	192
5-28	GEO Next-Night Cross-Track Prediction Error for 32 Minutes Per Hour Observation Duty Cycle Over Four-Hour Arc	193
5-29	Bias and $\pm 3\sigma$ Formal Uncertainty in AMR With 8 Hz, 250 mas RMS Astrometry . .	194
5-30	Geometric Intuition for Bias in Geostationary Satellite Solar Radiation Pressure Coefficient Estimate	195
5-31	Ballistic Targets Time and Cross-Track Errors for One-Night Orbit Fit	198
6-1	Dispersive Optical Elements	202
6-2	Slit Spectrograph	204
6-3	Slitless Spectrograph	205
6-4	Transmission Grating Geometry	206
6-5	Schematic of NOFS Spectrograph Data Product	209
6-6	Grating Efficiency for NOFS Slitless Spectrograph	210
6-7	Reconstruction of NOFS Spectra from Basis Functions for Two Satellites	211
6-8	NOFS Spectroscopy of Selected Geosynchronous Satellites	212
6-9	Mean $\pm 3\sigma$ Astrometric Residuals of Two GLONASS Satellites vs. Eleva- tion Angle	214
6-10	Multiband Photometry and DCR Estimation of Two GLONASS Satellites	215
6-11	Equipment Available at the Firepond Facility	216
6-12	SNR Along Spectrum for 12.5 mag 80% Solar Panel Target	218
6-13	Comparison of SNR Along Spectrum for 12.5 mag 20% and 80% Solar Panel Target	220
6-14	Simulated Effect of Star Clutter on Slitless Spectroscopy of a Satellite	221

6-15	Fraction of Spectrum Contaminated by Star Clutter	222
6-16	Bench Test Setup Overview	224
6-17	Photographs of First Attempted Spot Projector	225
6-18	Free-Space OAP Spot Projector Imaged Through 16-inch Telescope	226
6-19	Samyang 500 mm Mirror Lens Spot Projector	227
6-20	Chromatic Aberration in Samyang Mirror Lens	227
6-21	Spot Projector Made with Orion Telescope	228
6-22	Repeatability of $m = 0$ Centroid Location Under Cycling of Filter Wheel	229
6-23	Radiometric Calibration Source Set Up at Firepond	230
6-24	Setup for Measurement of Spot Projector Spectrum	231
6-25	Aperture Mask for Measuring Calibration Lamp	231
6-26	Setup for Piecewise Measurement of Spot Projector Spectrum	233
6-27	Blocking Filter Bench Calibration	234
6-28	Ratio of Responses of 606 and 906 Grating vs. Prediction	235
6-29	Optical Path Difference Calculation for Titled Grating	236
6-30	Influence of Grating Tilt on Ratio of Grating Efficiencies	237
6-31	Estimation of Grating Tilt Using Response Ratio	237
6-32	Installation of the Spectrograph at Firepond, August 16, 2018	238
6-33	Post-Fit Pointing Residuals of Mount Bias Models Computed for the 48" Telescope and 16" Spectrograph Telescope from Data Collected on the Night of Nov 7-8, 2018	241
6-34	Time-History of Difference Between 16" Telescope Pointing Measured from Bare-CCD Astrometry Frames and Pointing Computed from Mount Bias Model . . .	243
6-35	MODTRAN Atmospheric Spectral Transmittances and Sample Stellar Spec- tra for a Warm Wet Night and a Cold Dry Night	244
6-36	Stability of Wavelength Scale Over Operational Nights with Star Measurements . . .	245
6-37	Achieved Optical Quality and Spectral Resolution	246
6-38	On-Sky Measurements of Grating Tilt	247
6-39	On-Sky Measurements of Blocking Filter Transmittance	248
6-40	Absolute Photometry Mismatch	250
6-41	Spectral Type Mismatch	250
6-42	One vs. Two AR Coatings	251

6-43	Grating Images of HD198001 Observed at 21° Elevation During Checkouts	252
6-44	Grating Efficiency of Spectral Orders Suitable for Horizon Reference	253
6-45	Magnitude of Observed DCR Error Compared to Absolute DCR	256
6-46	Example of DCR Verification Data	257
6-47	Measured Deviation from Refraction Model vs. Elevation	258
6-48	Satellite Illumination Conditions	259
6-49	Elevation vs. Phase Angle of Satellites Observed with Spectrograph	260
6-50	Measured Brightness of Satellites Observed During Spectroscopy Measurements . . .	261
6-51	Measured Sky Brightness at Firepond During Spectroscopy Measurements	262
6-52	SNR vs. Predicted DCR Estimation Accuracy for Spectroscopy Measure- ments at Firepond for 40 nm Spectral Resolution	264
6-53	Data Volume vs. Desired DCR Estimation Accuracy	265
6-54	Predicted vs. Observed Vertical Astrometric Residual of Satellites with Measured Spectra	267
6-55	Distributions of Residuals Under Different Processing	268
6-56	Observed DCR vs. Predicted Separated by Elevation Range	271
6-57	Observed DCR vs. Elevation	272
6-58	Distributions of Observations Corrected with Spectroscopic Measurements and with Solar Panel/Kapton Mixture Model	272
6-59	Measured Brightness of SCC 32660 Exiting Eclipse	273
6-60	Reconstructed Spectra of SCC 32660 Exiting Eclipse	273
6-61	Estimate of Atmospheric Transmission into Partial Eclipse Region	275
7-1	Error Distributions in Orbit Determination Routines	286

List of Tables

3.1	Approximate Lines of Code in OrbitTool by Functional Area	71
3.2	Geocentric Force Model Accuracy	86
3.3	Geocentric Empirical Force Magnitudes	87
4.1	Farnocchia-Correctable CCD Observations of Asteroids with Known Spectral Types and Well-Constrained Orbits in the MPC Database as of December 2018	118
4.2	Top Five Sites Observing Long-Arc Optics-Only S- and C-Type Objects with Selected Passbands	118
4.3	Calculated Resolution Requirements to Correct Observations of Asteroids to 25 mas	139
5.1	Asteroid Family Parameters Used in Simulation	149
5.2	Simulated Orbits with Large DCR-Induced Mean Motion Biases	150
5.3	Simulated Orbits with DCR-Induced Spurious Yarkovsky Detections	165
5.4	Simulated Hazardous Asteroid Orbital Parameters	172
5.5	2036 B-Plane Miss Distance Estimate vs. Fidelity of DCR Correction and Observation Quality	174
5.6	Simulated Mumbai Impactor Orbital Parameters	174
5.7	Best-Fit Multi-Night Area-to-Mass Ratio	197
5.8	One-Night AMR for Ballistic Targets	198
6.1	Geostationary Satellites Observed by NOFS Spectrograph	209
6.2	Summary of Spectrograph Operational Nights	240
6.3	Distribution Statistics of Residuals Under Different Processing	269

List of Abbreviations

AMR	Area-to-mass ratio
BCRF	Barycentric Celestial Reference Frame
CCD	Charge-coupled device
COTS	Commercial off-the-shelf
DCR	Differential color refraction
Dec	Declination
EOPs	Earth orientation parameters
ETS	MIT LL Experimental Test Site, Socorro, NM
FWHM	Full-width half-max
GCRF	Geocentric Celestial Reference Frame
GEO	Geosynchronous/geostationary orbit
GEODSS	Ground Based Electro-Optical Deep Space Surveillance
GLONASS	Russian GLObal NAvigational Satellite System
GNSS	Global navigational satellite sytem, used generically
GPS	United States Global Positioning System
GSL	GNU Scientific Library
HAMR	High area-to-mass ratio
ICRF	International Celestial Reference Frame
IERS	International Earth Rotation and Reference Systems Service
IGS	International GNSS Service
ILRS	International Laser Ranging Service
ITRF	International Terrestrial Reference Frame
KBO	Kuiper Belt Object

LINEAR	Lincoln Near Earth Asteroid Research
MEO	Medium Earth orbit, i.e. GPS/GLONASS altitude
MIT LL	MIT Lincoln Laboratory
MJD	Modified Julian Date
MPC	International Astronomical Union Minor Planet Center
NOFS	United States Naval Observatory Flagstaff Station
OAP	Off-axis paraboloid
OTA	Optical tube assembly
PHA	Potentially hazardous asteroid
PSF	Point spread function
PSM	Parameter sensitivity matrix
QE	Quantum efficiency
RA	Right ascension
RAAN	Right ascension of the ascending node
RMS	Root mean-square
RSO	Resident space object
SCC	Space Command Catalog
SLR	Satellite laser ranging
SMASS	Small Main-Belt Asteroid Spectroscopic Survey
SNR	Signal-to-noise ratio
SRP	Solar radiation pressure
SSA	Space situational awareness
SST	Space Surveillance Telescope
STM	State transition matrix
TAI	International Atomic Time
TCS	Telescope Control System
TDB	Barycentric Dynamical Time
UCAC	USNO CCD Astrograph Catalog
URAT	USNO Robotic Astrometric Telescope Catalog
USNO	United States Naval Observatory
UTC	Coordinated Universal Time

Notation

q	Perihelion
Q	Aphelion
a	Semi-major axis
e	Eccentricity
i	Inclination
Ω	RAAN
ν	True anomaly
ω	Argument of perihelion
m_0	Mean anomaly at epoch
$n, \Delta n$	Keplerian mean motion, bias in mean motion
A_2	Normalized Yarkovsky acceleration
R_E	Earth radius
ω_E	Earth angular velocity
π	Stellar parallax, as well as 3.1415...
\mathbf{A}	Least-squares design matrix
$\mathbf{W}^{\frac{1}{2}}$	Least-squares weighting matrix
Σ	Formal covariance matrix
\mathbf{R}	Cartesian rotation matrix
$\mathbf{P}(t)$	Non-Cartesian parameter sensitivity matrix
$\Phi(t)$	Cartesian state transition matrix
\mathbf{x}	Cartesian state
\mathbf{y}, θ	Observation vector, usually angles
$z, \Delta z, \delta z, y$	Angle from or along local vertical or vertical focal plane offset
Φ, ϕ, ψ	Detectable photoelectron flux, flux density, or signal
λ	Wavelength
σ	Square root of the second moment of a Gaussian distribution

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

Earth’s atmosphere is an optically dispersive medium. This subjects observations of Solar System bodies and Earth satellites made with ground-based telescopes to systematic bias from differential color refraction (DCR). DCR bias is on the order of 100 mas and affects both the targets and the fiducial stars used to form the astrometric solution in telescope images. The majority of data for space situational awareness (SSA) of satellites and orbital debris in geosynchronous (GEO) orbit comes from observations made with optical telescopes, as do nearly all observations of asteroids and all observations of Kuiper Belt Objects (KBOs).

Until recently, the systematic bias from DCR had been overshadowed by the fixed-pattern error in the star catalogs used to form astrometric solutions. Over the past several years, the results of the European Space Agency’s Gaia satellite mission have resulted in a dramatic improvement in the accuracy of star catalogs. In January 2017, the Gaia DR1-registered UCAC5 star catalog was released with a 25 mas RMS all-sky accuracy and in April 2018, the Gaia DR2 catalog was released with sub-milliarcsecond accuracy. This reduction in catalog error has exposed underlying DCR bias as a dominant contributor to the error floor in ground-based optical astrometry.

1.1 Phenomenology

DCR is shown schematically in Figure 1-1. Short wavelengths refract more than long wavelengths and appear from the ground to be incident from a higher elevation. Target signals composed mostly of short wavelengths appear higher up than target signals composed mostly of longer wavelengths. The bias affects both fiducial stars and astrometry targets. The effect on the astrometric solution for a target is inverted for the DCR of fiducial stars: targets registered against bluer stars appears

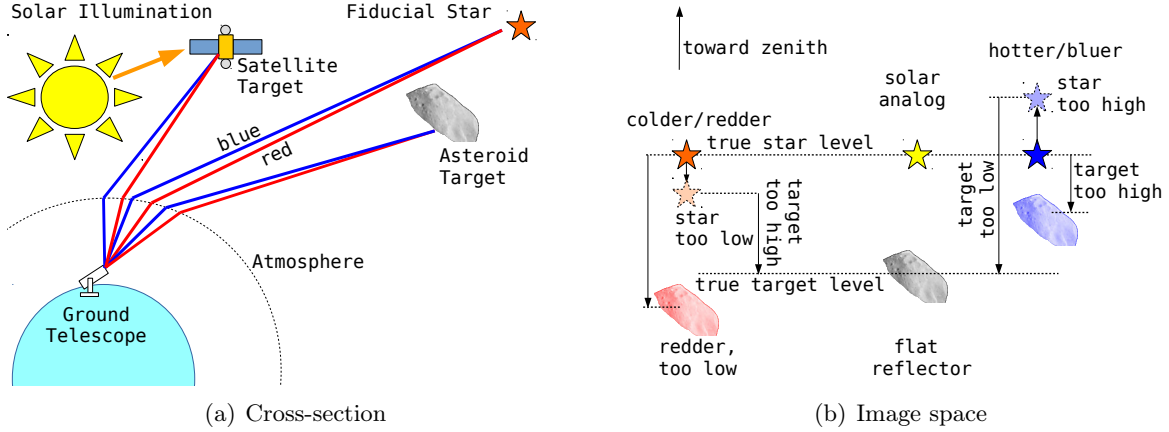


Figure 1-1: Schematic of DCR Bias

to originate from a lower elevation because the fiducial sources appear to be at a higher elevation and targets registered against redder stars (most stars are redder than Sun-like) appear to originate from a higher elevation because the stars defining the plate coordinate system appear systematically lower.

Proceeding mostly from Filippenko [1], the computation of the DCR bias for a target is as follows. At a zenith distance z , let the target detectable photoelectron spectral flux density be $\phi(\lambda, z)$ and let the total refraction as a function of wavelength at that zenith distance be $y = \Delta z(\lambda, z)$. The symbol y is the vertical angular coordinate on the focal plane corresponding to wavelength λ . DCR transforms the spectral flux density into an angular flux density on the telescope's focal plane. This density is given by

$$\psi(y) = \phi(\Delta z^{-1}(y, z), z) \left(\frac{dy}{d\lambda} \right)^{-1} \quad (1.1)$$

The Jacobian term $\left(\frac{dy}{d\lambda} \right)^{-1}$ approaches infinity in the red portion of the silicon band where the refraction versus wavelength curve is relatively flat, and for numerical reasons it is preferable to compute the the vertical centroid δz_0 of the angular distribution in wavelength coordinates via:

$$\delta z_0 = \frac{\int \psi(y) y dy}{\int \psi(y) dy} = \frac{\int \phi(\lambda, z) \Delta z(\lambda, z) d\lambda}{\int \phi(\lambda, z) d\lambda} \quad (1.2)$$

This is the DCR bias with respect to some zero point of $\Delta z(\cdot)$ at zenith distance z . For convenience, that zero point is taken to be the vertical centroid of a solar-analog star.

The relevant quantity is the spectral flux of detectable photoelectrons. For a solar-illuminated target, that spectrum is the product of the solar illumination spectrum with the target spectral

reflectance, atmospheric spectral losses and overall instrument spectral quantum efficiency. For a relative astrometric measurement against a fiducial star, the atmosphere and detector are common between the spectra of the target and star. That common path renders the DCR bias only weakly sensitive to atmospheric conditions (though it is somewhat sensitive to meteorological visibility). The greatest sensitivity is to the difference between the exoatmospheric spectrum of fiducial stars and the exoatmospheric spectrum of the target. DCR bias is only weakly sensitive to the absolute amount of refraction at a given zenith distance and to the method used to compute refraction. Both a simple $R \propto \sec z$ approximation of refraction and the more involved Auer-Standish [2] algorithm yield nearly indistinguishable values for DCR bias. What is important for accurate modeling of the DCR of one source with respect to another is not the calculation of the total bending but the fidelity of the model for the wavelength-dependence of the refractive index of air and the target's detected photoelectron spectrum.

1.2 Evidence of DCR Bias and Its Significance

Though cognizance of DCR bias is not new, evidence that it is a significant contributor to the error floor in achievable astrometric accuracy had been limited until the start of the 2010s. Chapter 2 will go into the details of the error sources contributing to ground-based astrometry, focusing mostly on the astronomical community. This section will relate some history from the near-Earth SSA side of the story to explain why.

1.2.1 Astrometric Observations of Calibration Satellites in Earth Orbit

In the 1990s, MIT Lincoln Laboratory (MIT LL) had completed for the United States Air Force the design for the upgrade of the Ground Based Electro-Optical Deep Space Surveillance (GEODSS) telescopes from tube cameras to large format charge-coupled devices (CCDs). These devices had 1960×2560 2.2 arcsec square pixels. A primary purpose of the upgrade was to improve detection sensitivity over the tube cameras. The deployment of the large-format CCDs enabled the use of in-frame astrometry for SSA data reductions in place of mechanical measurements of telescope pointing. The software change from relying on mechanical angle measurements to in-frame astrometric solutions occurred in the late 1990s and improved the absolute accuracy of the GEODSS systems from a performance of 20–40 arcsec RMS down to about 4–5 arcsec RMS. A further upgrade in 2003 pushed that down to 2 arcsec RMS (Figure 1-2).

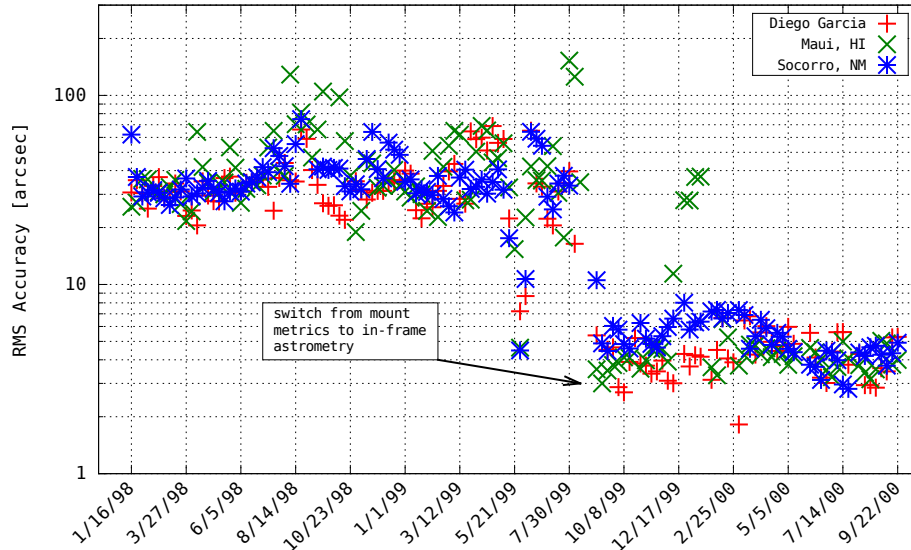


Figure 1-2: Accuracy Assessment of GEODSS Observations in the Late 1990s. Data from [3]. The US Air Force GEODSS system began operation in the 1970s with vidicon detectors capable of only single point source measurements and relied on electromechanical measurements of physical gimbal angles to generate observations. CCD detectors enabled in-frame astrometric reductions on satellite targets moving rapidly against the star background, improving accuracy to the 5 arcsecond level [3]. Further refinements to the software lowered errors down to 2 arcsec RMS in the early 2000s [4].

In 2009, several related efforts were underway at MIT Lincoln Laboratory. The MIT LL Experimental Test Site, Socorro, NM (ETS) in Socorro, New Mexico had been operating the Lincoln Near Earth Asteroid Research (LINEAR) systems [5] for over a decade and was preparing for the delivery of the Space Surveillance Telescope (SST) [6]. That effort involved adapting and extending the astrometric software that had been developed for the LINEAR program, the GEODSS astrometric processing upgrade [3, 4], and the Morón Optical System [7] for the SST.

At the same time, the MIT Lincoln Laboratory was reactivating the Firepond Optical Facility in Westford, Massachusetts to serve as an additional testbed for ground-based SSA techniques [8]. I was on the engineering team tasked with the recommissioning (Figure 1-3) and part of my job was to port that ETS astrometric reduction software for use at Firepond. At that time, the state-of-the-art for LINEAR and for MIT LL ground-based SSA was to use the USNO-B star catalog [9]. The USNO-B catalog had a systematic error floor on the order of a few tenths of an arcsecond, but in the context of the GEODSS upgrade which saw accuracy go from many tens to single-digit arcseconds, this was an improvement over past performance for ground-based SSA.

The Firepond recommissioning achieved initial operating capability in early July 2009. Throughout July and August 2009, several nights were spent validating the performance of the newly-




Figure 1-3: Firepond Recommissioning Team, June 12, 2009. Image Credit: Steve Hunt, MIT LL. Counter-clockwise from right: Jan Kansky, Lawrence Swezey, Roman Geykhman, Wayne Robertson.

reactivated system. As part of the checkout procedures, several full nights were spent observing GPS satellites for the purpose of a metric accuracy assessment. This was meant to verify the proper operation of the entire system: adequate optical quality, focus, and collimation; mechanical stability during satellite tracking; image timing system integrity; and correct astrometric processing software configuration.

The data collections consisted of several thousand distinct observations of GPS satellites observed with 120 msec exposures. The assessed accuracy, computed as the running mean astrometric residual averaged over five-second periods, was between 0.1 and 0.2 arcsec, while the seeing-induced short-timescale astrometric noise was on the order of 0.3 arcsec. This was consistent with the claimed accuracy of the USNO-B catalog, and there did not appear to be any preferential directionality to the error statistics. In some parts of the sky, the error was aligned with the local horizontal, in other parts it was aligned with the local vertical, or with the RA/Dec axes.

The Firepond system continued to operate as a developmental testbed, and by 2012 the camera control and timing capture software had reached its intended configuration. Checkout observations of GPS and GLONASS satellites continued to be made as part of normal operations in order to validate system performance.

In 2013, the UCAC3 star catalog [10] became available. Metric accuracy assessments computed against observations reduced with this new catalog exposed underlying systematic error sources that had been masked by the fixed-pattern astrometric biases present in the USNO-B catalog.

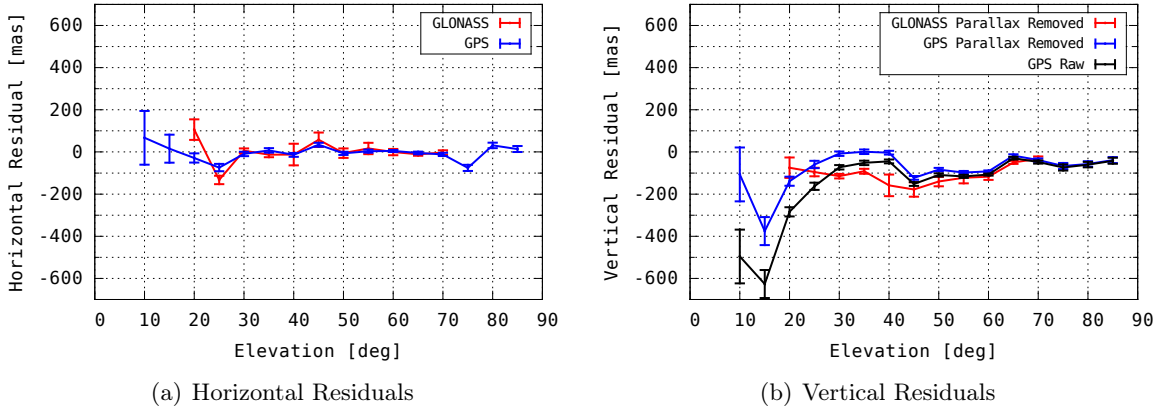


Figure 1-4: Residuals of Observations of GPS and GLONASS Satellites Reduced with URAT1. Mean and $\pm 3\sigma$ of mean residuals vs. elevation, $N=130,000$ observations of GPS satellites and 35,000 observations of GLONASS satellites. Star DCR correction not applied. GPS satellites' vertical residuals are plotted with parallactic refraction (Section 3.3.4.2) both removed and left in to demonstrate the relative magnitude of parallax and DCR at GPS distances.

The most apparent revelation from the new analysis was that if astrometric residuals statistics were tabulated in station-aligned horizontal/vertical coordinates rather than sky-aligned RA/Dec coordinates, then the nonzero component of the mean astrometric error of all observed calibration satellites was confined entirely to the local vertical coordinate, while the mean horizontal error was near zero, suggesting strongly that all of the systematic errors in SSA astrometry were coming from atmosphere-induced biases.

By the end of 2016, approximately 165,000 distinct astrometric observations of GLONASS and GPS satellites had been made as part of routine system checkout observations. At the time, the 2015-released URAT1 star catalog [11] was the most accurate available, and it was possible to analyze the error statistics from the metric accuracy assessments of data reduced with these newer catalogs in bulk. Figure 1-4 shows the mean astrometric residuals for all observed satellites as a function of elevation angle. On average, the mean horizontal residual is near zero in all elevation bins. The outlier peak at 45 degrees is consistent with a catalog bias of no more than the 30 mas claim of the URAT1 catalog. While other outliers in the horizontal coordinate do exist at high elevation that cannot be attributed to small-number statistics, the presence of consistent near zero-mean errors for much of the span of elevation indicates good overall astrometric data quality.

Confidence in the astrometric reductions allows the claim to be made that the behavior of the vertical residual in the same region is strongly indicative of unambiguous DCR bias. The variability in the GPS satellites' vertical residual curve with elevation is indicative of stabilized spacecraft

with Sun-tracking solar panels whose apparent aspect to the observing site varies with elevation angle. GPS is in a repeating twice-a-day orbit and one would expect fairly strong coupling between phase and elevation angles from a single ground site. The separation in traces between the vertical residuals of the GPS and GLONASS satellites is indicative of spacecraft in different orbits presenting different aspect angles and possibly of the difference in surface materials of two different spacecraft constellations.

1.2.2 Astrometric Observations of Asteroids with Known Spectral Types

The analysis in Section 1.2.1 presents strong evidence that atmosphere-induced bias is dominant in ground-based optical observations of calibration targets. The persistence of the bias after removal of parallactic refraction suggests strongly that DCR is the cause of the remaining detectable astrometric error. To strengthen that claim, it is necessary to compare the astrometric residuals of objects with a known color difference. The Minor Planet Center (MPC) archive of astrometric observations of asteroids offers can support this type of analysis. Of the several hundred thousand distinct asteroids in the MPC database, several thousand are spectroscopically-characterized and fall into either the C-type of flat gray reflectors or the S-type of slightly red objects [12].

Figure 1-5 shows the horizontal and vertical residuals for the two asteroid classes from all contributing stations and separately from the original LINEAR program. The first thing to notice is that there is, indeed, a separation between the redder S-type asteroids and the C-type asteroids in the vertical coordinate that does not appear to exist to the same level of statistical significance in the horizontal coordinate. However, the common-mode structure in both the vertical and horizontal residuals makes it difficult to attribute “all” of the systematic error to atmospheric biases.

A large fraction of the Minor Planet Center archive consists of observations made by the LINEAR program. Considering only observations from that sensor, one sees that there is similar structure in both horizontal and vertical residuals. This is where the star catalog biases and DCR are difficult to dissociate from other effects. The bump in the vertical residual at an elevation of about 40 degrees corresponds to a declination of about -20 degrees. The LINEAR system came online in the mid 1990s and used the USNO-A astrometric catalog available at that time. This catalog is made from scans of photographic plates from surveys conducted in two halves, one for the Northern Hemisphere collected at Palomar, and one for the Southern Hemisphere collected in Chile [13]. The seam between the two halves occurred around that declination band, and the error in joining the two datasets propagated into the LINEAR data and dominated the statistics of the systematic error




Figure 1-5: Residuals of Archived Observations of S-type and C-type Numbered Asteroids Resolved into Station-Local Horizontal/Vertical Coordinates. Mean $\pm 3\sigma$ of the mean residual versus elevation with respect to JPL best estimate ephemerides retrieved Nov 2017 [15]. Observations retrieved from IAU Minor Planet Center [16] Jan 2017. Aggregate data from all stations spans calendar year 2000 to 2016. The LINEAR system observations span the calendar years 2000 – 2012 when the original LINEAR program was shut down and replaced by the Space Surveillance Telescope. The total number of LINEAR observations consists of approximately 50,000 observations of C-type asteroids and 120,000 observations of S-type asteroids.

for the entire Minor Planet Center archive.

With the publication of the Farnocchia catalog corrections [14] to the USNO catalog, it became possible to remove this error from the LINEAR data and from other early observations. When catalog error is removed, the structure in the horizontal coordinate of the residuals disappears. Also, spurious separation in the mean horizontal residuals for the two classes of asteroids disappears, and the common mode error in the vertical coordinate is greatly attenuated, while the separation in the vertical residuals remains. That remaining separation of traces is the key to the argument that DCR is significant because demonstrates that when all else is held equal: same sensor, same average atmosphere, and same elevation angle, but the color of the target varies, a measurable astrometric bias occurs in the vertical coordinate. The analysis of asteroid residuals is revisited again in Section 4.3.

1.3 Motivation

Astrometry has historically been plagued by low data availability, which in turn has limited the motivation and capacity to improve reduction techniques and investigate the limits of observation accuracy in SSA. The availability of inexpensive image detector technology, in combination with improvements in star catalog accuracy, has removed this limitation, leaving fertile ground for investigation. Quoting (with emphasis added) from the National Academy’s report *Continuing Kepler’s Quest* [17], published in 2012,

The committee believes that the primary limitation in the current system for objects not experiencing significant drag is not the accuracy of the algorithms, but rather the *quantity and the quality of the sensor tracking data*. The key system limitations are current sensor coverage, *understanding of the quality of the observations*, and the challenge of fusing disparate data from different systems and phenomenology. *Understanding the quality or statistics of the observations is necessary for obtaining a realistic covariance, which is needed for computing an accurate probability of collision*. For near-Earth orbiting satellites another limitation is understanding and modeling of the atmosphere.

1.3.1 Improvements in Star Catalog Accuracy

The largest motivating factor for improvements in measurement techniques and data processing for ground-based SSA is the fact that star catalog accuracy has improved by several orders of magnitude over the past several decades. Where once systematic catalog errors were measured in fractions of an arcsecond, they are now measured in fractions of a milliarcsecond. The removal of this once-dominant error source in the astrometric measurement exposes other underlying physical effects.

DCR bias coexists with other sources of systematic error in CCD astrometry. These error sources are generally zero-mean and are a product of mechanical perturbations to the observing instrument that spoil the optical quality in a long-exposure image relative to an ideal in the absence of disturbances. Recent work with high frame-rate detectors and image stacking techniques [18, 19] has made it possible to push past these effects and realize the full benefit of modern star catalogs. DCR bias is a real physical effect rather than a deviation from a theoretical ideal in the way mechanical perturbations are, and this thesis shows that DCR bias is the remaining dominant error source once catalog bias and instrument non-idealities are surmounted.

1.3.2 Asteroid Search

Nearly all Solar System bodies are observed exclusively with optical telescopes and the majority of those telescopes are ground-based. Accurate observations of asteroids are necessary for both planetary defense purposes and in order to enable accurate characterization of non-hazardous asteroids for scientific purposes. Because the majority of these sensors operate with bare silicon detectors to optimize detection probability for small or distant objects, DCR is the dominant long-timescale error source in these observations after star catalog error is mitigated.

Accurate predictions of Earth impacts and Earth impact points are crucial for long-term planning. Low-accuracy and low-volume data can make accurate forecasting difficult. The evolution of the trajectory predictions for 99942 Apophis and the case of 2014 AA provides an illustrative example. After the discovery of 99942 Apophis with optics-only observations in 2004, an optics-only orbit estimate yielded a potential impact in 2036, but a wide range of potential impact points wrapped around almost half the Earth [20]. Alternatively, 2014 AA was discovered less than a day before impact and the predicted uncertainty in the impact points was also large, though the actual impact point was right in the center of the predicted distribution [21]. While the case of 2014 AA and the early estimates of the Apophis impact may be pinned on shortness of the arc (relative to the prediction time span), a driving source of statistical uncertainty in both of these cases is optics error model. This thesis argues that DCR bias is a major contributor to the optics error budget and thus a driving source of error in long-term asteroid impact monitoring accuracy.

1.3.2.1 Hazardous Asteroid Impact Monitoring

Because almost all observations of asteroids are made with ground-based optics, systematic errors in those observations and unmodeled effects in the ground optics observation model will dominate the error in orbit predictions. Because of the low data volumes involved and because of the near-lack of independent measurements by ground-based radars, these can be “silent” errors that are not detectable in the statistics of the post-fit residuals of the observations used to construct the orbit estimate. When an asteroid does occasionally pass into radar range, it is possible to be surprised by a prediction from an optics-only orbit estimate that differs significantly from a direct range measurement. For example, during the 2004 pass of the asteroid 3908 Nyx, the target was off-nominal by over 100 km due to a combination of unmodeled astrometric errors and unmodeled non-gravitational accelerations from the Yarkovsky effect [22].

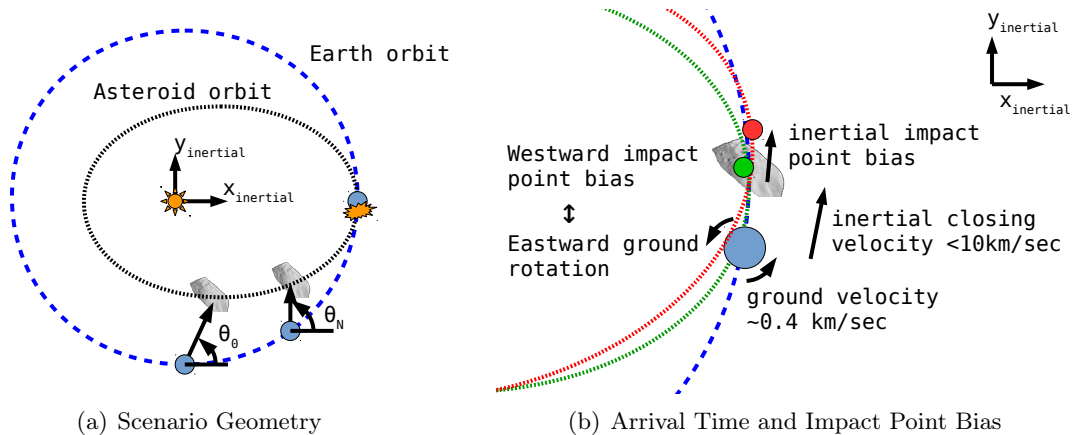


Figure 1-6: Orbit Geometry for Two-Dimensional Analysis of Significance of DCR on Asteroid Impact Predictions. An asteroid with aphelion at exactly 1 AU is chosen to minimize the effect of the curvature of the Earth’s orbit, thereby making arrival time difference primarily an East-West bias in the impact point.

Due to other systematic error sources in ground-based optical astrometry, it is not appropriate to ascribe all mismatches between prediction and observation to DCR bias in the input observations. It is, however, illustrative to analyze the degree to which DCR bias can contribute to asteroid orbit prediction errors, and whether any such contribution can be significant.

Experience with both near-Earth SSA and asteroid search quickly imparts the lesson that the orbital element with the highest uncertainty and the greatest sensitivity to observation error is Keplerian mean motion. In the case of an asteroid that is known to impact the Earth, a bias in mean motion, will, at the very least, bias the arrival time, and allow the Earth to rotate under the nominal impact point, thereby introducing an East/West bias in the predicted impact point.

To gauge the significance of this bias, consider the scenario depicted in Figure 1-6, where two-dimensional, two-body physics applies and a single station on the Earth makes a single observation of the asteroid per night. The asteroid’s orbit forces an observation from the night side of the Earth to be made at a low elevation angle, imparting the maximum DCR bias to the observation. The geometry forces the observations to always be made either near dusk or near dawn for long stretches of time, lowering the chance for observation errors to cancel.

With purely two-body physics in two dimensions, the orbital state of the asteroid is a four-dimensional quantity characterized by four invariant Keplerian elements:

$$\mathbf{k} = [a, e, \omega, M_0] \tag{1.3}$$

which are related to the Cartesian state by a function of time:

$$\mathbf{x}(t) = [x(t), y(t), \dot{x}(t), \dot{y}(t)] = f(t, \mathbf{k}) \quad (1.4)$$

The single inertial angle observation on each night is a function of the Cartesian state,

$$\theta_t = \tan^{-1}(\Delta y_{\text{inertial}}/\Delta x_{\text{inertial}}) \quad (1.5)$$

The observation Jacobians, given by:

$$\Delta\theta_t = \frac{\partial\theta}{\partial\mathbf{x}} \frac{\partial\mathbf{x}}{\partial\mathbf{k}}(t)\Delta\mathbf{k} \quad (1.6)$$

may then be used to compute iterative corrections to the state estimate. Beginning with the sensitivity of observations to perturbations in the Keplerian state,

$$\begin{bmatrix} \vdots \\ \Delta\theta_t \\ \vdots \end{bmatrix} = \begin{bmatrix} \vdots \\ \frac{\partial\theta}{\partial\mathbf{x}} \frac{\partial\mathbf{x}}{\partial\mathbf{k}}(t) \\ \vdots \end{bmatrix} \Delta\mathbf{k} \quad (1.7)$$

The sensitivity of the Keplerian state to perturbations in the observations away from their error-free, noise-free values is:

$$\Delta\mathbf{k} = \begin{bmatrix} \vdots \\ \frac{\partial\theta}{\partial\mathbf{x}} \frac{\partial\mathbf{x}}{\partial\mathbf{k}}(t) \\ \vdots \end{bmatrix}^+ \begin{bmatrix} \vdots \\ \Delta\theta_t \\ \vdots \end{bmatrix} \quad (1.8)$$

where $[\]^+$ refers to the left pseudoinverse. If $\Delta\theta = [\Delta\theta_0, \dots, \Delta\theta_N]^T$ is the *uncorrected bias* in the observations, then to first order, the Cartesian error in the orbit estimate propagated out to the nominal impact time is given by:

$$\Delta\mathbf{x}(t_{\text{impact}}) = f(t_{\text{impact}}, \mathbf{k}_0 + \Delta\mathbf{k}) \quad (1.9)$$

Referring to Figure 1-6(b), asteroids with semi-major axes less than 1 AU and aphelion Q at exactly 1 AU dwell in front of the Earth during their aphelion. If the dominant bias is in some combination of mean motion, mean anomaly at epoch, or in the argument of perigee, this has the effect of moving the dwell location further in front of the Earth or closer to the Earth as it catches

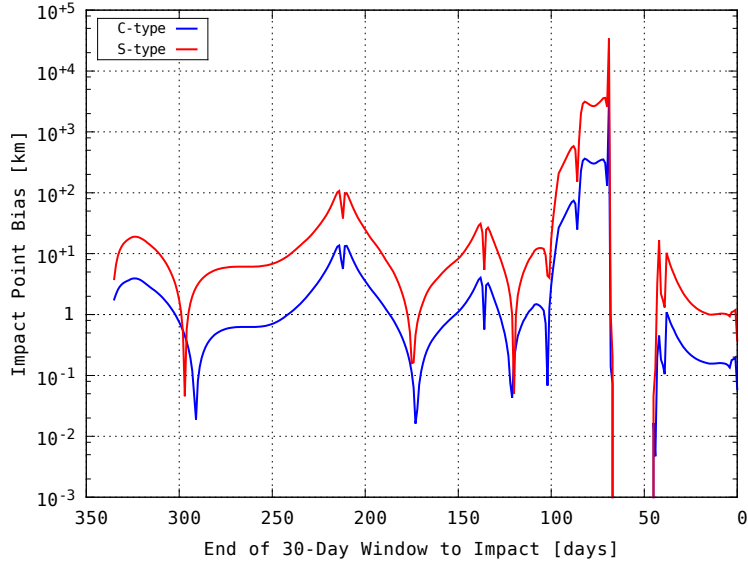


Figure 1-7: Impact Point Bias in Two-Dimensional Analysis of $q = 0.9, Q = 1.0$ Orbit. The gap between 50 and 70 days to go corresponds to the asteroid passing too close to the Sun for ground-based observations to be made from the night side.

up to the asteroid. Neglecting the details of gravity capture, the Cartesian state error of the asteroid translates into an impact time bias via:

$$\Delta t_{\text{impact}} \approx \Delta y / \dot{y} \quad (1.10)$$

which in turn becomes an East/West bias via scaling by the Earth's equatorial surface velocity:

$$\Delta s \approx R_E \omega_E \Delta y / \dot{y} \quad (1.11)$$

Taking this all together, an (admittedly contrived) scenario can be found where DCR bias in observations imparts a significant bias to the predicted impact point. Consider the case where the perihelion $q = 0.9$ AU and an impact point is computed after thirty nights of observations. The bias in the impact point as a function of when that thirty-day window ends is shown in Figure 1-7. The bias is smaller for a relatively flat-reflecting C-type asteroid than for a red S-type asteroid. There are times where the impact point bias computed from a 30-day window of observations can exceed several hundred kilometers. Though this scenario was deliberately constructed to maximize the effect of DCR, this preliminary analysis shows the existence of pathological cases where DCR bias in observations can impart significant bias to asteroid orbit estimates, and motivates a further analysis of the implications of DCR for asteroid orbits. This analysis is developed in Chapter 5.

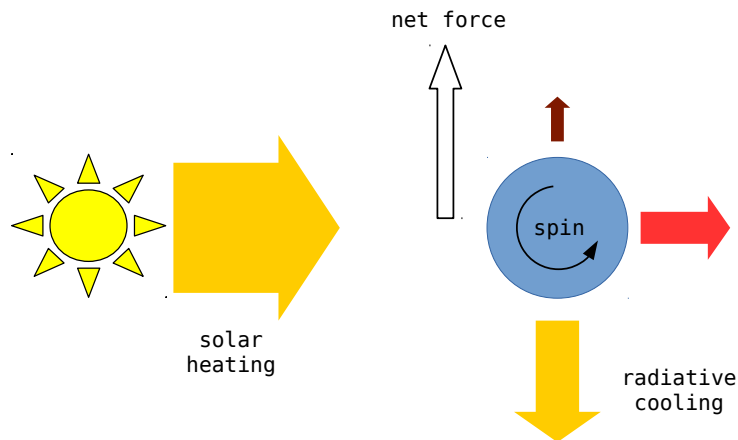


Figure 1-8: Schematic of Yarkovsky Effect. The orientation of the object’s spin axis relative to its orbital velocity vector determines the degree to which Yarkovsky acceleration perturbs the object’s orbit and at what time scale. A spin axis directly perpendicular to the orbital plane aligns the Yarkovsky acceleration with the in-track component and maximizes the measured along-track deviation of the object from a trajectory modeled with no non-gravitational perturbations.

1.3.2.2 Yarkovsky Effect Characterization

The Yarkovsky effect is a non-gravitational acceleration caused by non-isotropic thermal re-emission by the surface of a spinning asteroid heated by sunlight (Figure 1-8). The Yarkovsky effect contributes an acceleration on the order of 10 to $100 \times 10^{-15} \text{AU/day}^2$ to asteroids where it has been measured (though on some low-density objects it is much higher) [23], meaning that for an asteroid with an unmodeled Yarkovsky acceleration, Cartesian position error between ground truth and the orbit estimate accumulates at a rate of between 100 and 1000 meters per year.

Astrometric errors can mask this physical effect if the non-gravitational accelerations result in observable deviations from nominal angular position that are on the same order of magnitude as the astrometric error. In the case of zero-mean astrometric noise, it is possible for observation error to attenuate with the incorporation of more data and still measure the desired effect. DCR, however, is not zero-mean and has long correlation times, meaning its effect on orbit solutions does not attenuate as quickly as purely random noise. At typical observation distances on the order of 1 AU, systematic biases on the order of 50 mas correspond to about 36km of apparent cross-range displacement. This is about an order of magnitude higher than the Yarkovsky-induced displacement from an orbit solution free of non-gravitational perturbations.

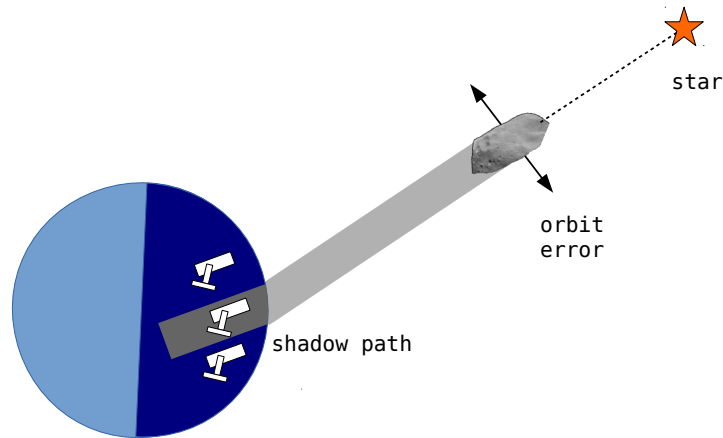


Figure 1-9: Schematic of Stellar Occultation. The primary observable is the duration of the occultation. If measured from multiple stations or against multiple stars, the physical size of the object may be measured. The ability to make this measurement is predicated on the ability to predict the ground-track of the shadow in advance. The accuracy of the prediction of the ground-track of the shadow is identical with the accuracy of the orbit prediction of the occulting object.

1.3.3 Kuiper Belt Object Occultations

The Kuiper Belt consists of objects beyond the orbit of Neptune. The first object discovered in this orbital regime was Pluto in 1930. With the advent of CCD detectors in the 1990s, additional KBOs were discovered beginning in the early 1990s [24], with several thousand more discovered over the past two decades [15]. Given their location at the outer edges of the Solar System, KBOs are studied primarily with stellar occultations [25]. The measurement of stellar occultations is shown schematically in Figure 1-9.

The ability to make occultation measurements is predicated on the ability to make accurate predictions of the orbits of the objects to be studied. Because nearly all KBOs are relatively recent discoveries, their orbital solutions are based on arcs that are short relative to their orbital periods. This implies that the accuracy in the estimate of the orientation of the orbital plane of these objects is relatively coarse and a systematic bias in the observations used to estimate the orbit will tilt estimate of the orientation of their orbital plane. At a typical observation range of 40 AU, a 0.1 arcsec observation bias aligned perpendicular to the orientation of the orbital plane of a KBO can introduce approximately 3000 km of bias, which will directly translate into a comparable bias in the prediction of the ground-track of the occultation, potentially leading to a missed measurement.

1.3.4 Near-Earth SSA

1.3.4.1 Satellite Debris Monitoring in Earth Orbit

Earth orbit is an increasingly crowded environment and accurate tracking of resident space objects (RSOs) is critical to maintaining safe separation between active satellites and debris fragments [26]. While the operator of an active satellite with a working transponder can determine its location and command its maneuvers with high accuracy, uncontrolled debris must be tracked non-cooperatively. In the GEO orbital regime in particular, nearly all of this tracking will occur with ground-based telescopes, meaning the systematic error floor on the the orbit solution for these potential hazards to active satellites is determined by the systematic error floor on the input observations.

A recent and public example illustrating the situation occurred in the summer of 2017 when ExoAnalytic Solutions detected a breakup event affecting SCC 27820 [27]. Multiple fragments separated from the satellite, each potentially a hazard to other RSOs in GEO. Most of these objects were small, meaning they cannot be tracked with ground-based radars, and logically, at least some of these would be high area-to-mass ratio (HAMR) objects.

Faint HAMR objects are readily detectable in orbits in and near the geosynchronous belt [28]. A major source of uncertainty in the orbital predictions for this class of objects is the amount of non-gravitational perturbation to their trajectory from direct solar radiation pressure (SRP) [28, 29]. Most active satellites' AMRs are significantly under $0.1 \text{ m}^2/\text{kg}$ and the effect of SRP is small, taking many hours of observation to measure (see Section 5.5). While for objects with extremely high area-to-mass ratio (AMR), such as pieces of Kapton ($\text{AMR} \gg 1 \text{ m}^2/\text{kg}$), the effect of radiation pressure is readily noticeable above the noise in astrometric observations over a relatively short time, for objects with an intermediate AMR¹, the magnitude of the observable effect of AMR is on par with the magnitude of DCR bias.

Many HAMR objects have complex shapes and attitude dynamics, suggesting strongly that their individual spectral reflectances could vary depending on the material composition and orientation they present to an observing telescope. Potentially, these observations can be subject to large and time-varying amounts of DCR bias. Thus, for this class of objects, DCR bias defines the error floor in orbit prediction accuracy.

¹An example of an object with an intermediate AMR is the MSG-2 cooler cover (SCC 29106), which has an AMR of approximately $0.2 \text{ m}^2/\text{kg}$ [29].

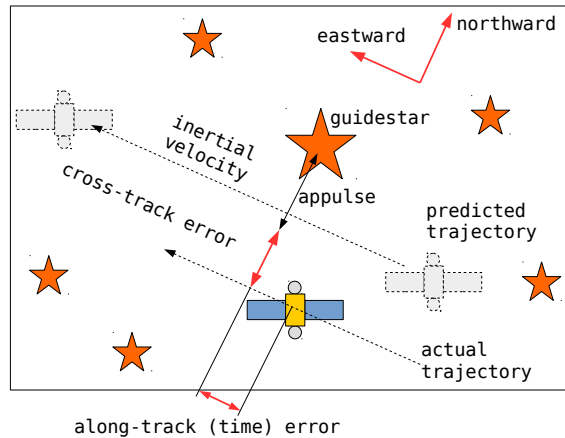


Figure 1-10: Appulses of Guidestars by a Geostationary RSO. Direct imaging of geostationary RSOs requires predicting appulses of sufficiently-bright natural guidestars in advance. While timing accuracy is less important, the ability to make predictions accurate in the cross-track coordinate determines the time horizon over which imaging operations can be planned in advance.

1.3.4.2 Resolved Imaging of Geostationary Satellites

Another motivating factor for high-accuracy astrometry, and one that is the focus of the near-Earth SSA portion of this thesis, is non-cooperative resolved imaging of geostationary satellites. Recent interest in direct-imaging of geostationary satellites and debris using natural guide stars [30] has created the necessity to accurately predict orbits of geostationary satellites.

Direct imaging using adaptive optics or interferometric techniques is restricted to occur within the isoplanatic angle of the reference guidestar (Figure 1-10). Isoplanatic angles are measured in single-digit arcseconds [31], which imposes the cross-track prediction requirement. Geostationary objects move at 15 arcsec/sec with respect to the star background and the desired timing accuracy of the appulse events imposes the in-track prediction requirement.

The arc lengths of the observing periods and data volumes required to achieve the systematic accuracy and uncertainty levels necessary to meet these requirements necessitate dedicated ground-based tracking. The lower limit on arc length is driven by the need to constrain the formal error bounds on the prediction. The upper limit is imposed by station-keeping and momentum dump maneuver schedules of the observed satellites, which can invalidate predictions derived from pre-maneuver observations. The susceptibility of ground-based data to DCR bias, as well as the volume and kind of data required to overcome it, motivate the near-Earth SSA portion of this thesis.

1.4 Gap Analysis and Thesis Contributions

Analysis of the implications of DCR bias for astrometric accuracy is not new, nor is spectroscopic characterization of Solar System bodies and Earth satellites. But what is new at the time of this writing is the reduction of sources of systematic bias in ground-based optical astrometry that, until recently, had been dominant drivers of observation and orbit prediction error and had masked the signature of DCR. In light of these recent developments, the gap in the literature is:

- DCR bias has not yet been identified as a dominant long-timescale systematic error source in astrometric observations made from ground-based telescopes operating bare-silicon detectors.
- There has been no comprehensive analysis of the sensitivity of orbit solution data products to DCR bias in the input observations and the significance of that sensitivity to the applicable mission area.
- No known effort has been made to measure SSA targets' spectra for the purpose of compensating for DCR bias in ground-based observations made with bare-silicon sensors.

Given this gap, this thesis makes the following original contributions:

1. DCR bias is definitively shown to be the dominant long-timescale systematic error in astrometric observations reduced with modern star catalogs. The analysis includes Solar System bodies as well as Earth satellites and shows that observing with narrow passbands or at low zenith distances are not always suitable or sufficient mitigation strategies.
2. The consequences of leaving DCR bias uncorrected in orbit solutions are examined for Earth satellites, inner Solar System asteroids, and Kuiper Belt Objects. It is shown that there exist certain use-cases where DCR bias in observations can significantly affect the resulting orbit estimate in all of these orbital regimes.
3. A DCR mitigation technique based on concurrent spectroscopic measurements is developed. The efficacy of the technique is demonstrated on measurements of GPS and GLONASS satellites at the MIT Lincoln Laboratory Firepond Optical Facility in Westford, MA. In data collected over seven nights, a reduction in the magnitude of residual DCR bias by 60% with respect to color-agnostic processing is shown.

1.5 Thesis Outline

This thesis is organized as follows. Chapter 2 is a literature review, summarizing the evolution of astrometric accuracy over time, reviewing the history of spectroscopic characterization of SSA targets, and modeling and measurement of astronomical refraction. Also included is necessary context for the achievable accuracy one can expect in an orbit solution given the uncertainty present in the forces acting on both geocentric and heliocentric objects.

Chapter 3 describes the technical approach of this thesis. For this thesis, I wrote a fully-featured orbit determination and astrometric data reduction software package, and Chapter 3 describes the design philosophy, assumptions, and algorithms in that code, and the methods used to validate it against existing baselines.

Chapter 4 is a deep dive into DCR bias, with first-principles analysis backed by measurements of the asteroid 3122 Florence. The predictability of DCR bias of fiducial stars from cataloged quantities is examined in depth, as is the expected DCR bias of typical targets under a variety of atmospheric conditions. The efficacy of standard mitigation techniques such as restricting observations to high elevations and observing through narrow passbands is examined, and shown to be insufficient. Given this insufficiency, the requirements for DCR compensation with real-time spectroscopic measurements are derived.

Chapter 5 builds on Chapter 4 to examine the consequences of uncompensated DCR bias on orbit solutions derived from biased data. For inner Solar System objects, the bias in observations projects mostly into a biased estimate of Keplerian mean motion. Although small in absolute magnitude, for certain orbits the bias in mean motion from a short-arc estimate can be many σ 's outside of its formal error bounds and can look like a statistically significant Yarkovsky acceleration in optics-only fits spanning several apparitions. Though this contribution to the error in an orbit solution for an asteroid is never the difference between recovery and a lost asteroid, for certain pathological cases of near-Earth asteroids it can make the difference between predicting a miss and predicting an impact. For Kuiper belt objects, it can be the difference between successfully predicting a stellar occultation and entirely missing the occultation path. Chapter 5 also examines the effect of DCR bias on predicting orbits of geostationary satellites with the specific motivation of predicting appulses of natural guidestars. In this orbital regime, the effect of DCR projects into a biased estimate of SRP AMR, which can both skew appulse predictions and bias the AMR estimate of HAMR objects.

Chapter 6 describes the main contribution of this thesis: a demonstration of DCR bias compensation using concurrent spectroscopic observations. This chapter begins with motivating measurements made on geostationary satellites with the slitless spectrograph at United States Naval Observatory Flagstaff Station (NOFS) over part of the silicon passband and a proof-of-concept multiband photometry measurement over the full silicon passband on GLONASS satellites at the Firepond Optical Facility. Building on these experiments, a slitless spectrograph instrument was designed and deployed at the Firepond Facility to make measurements of GPS and GLONASS satellites over the full silicon passband. The design, testing, and operation of this instrument is described, and the chapter concludes with results demonstrating the efficacy of concurrent spectroscopic compensation for DCR bias.

Chapter 7 summarizes the thesis, makes recommendations for the operation of and upgrade paths for SSA telescopes, and presents several potential avenues for future research.

Chapter 2

Literature Review

2.1 Astrometry Through Time

One of the main gaps in the literature occurs at the intersection of DCR bias with the context of overall achievable accuracy in ground-based astrometry. Part of the reason that this gap exists is that until recently, other sources of systematic error defined the achievable accuracy of observations made from ground-based optical telescopes. Some of that error was driven by the limitations of the then-available observing instruments: a telescope meant to be used by an astronomer bundled up for a cold winter night and looking through the eyepiece is going to have design compromises that favor human-machine interaction over mechanical precision and measurement accuracy. This section examines the evolution of achievable astrometric accuracy from the time of eyeball observations to the current era of highly-sensitive electronic detectors.

2.1.1 Eyeball Era

From the ancient Greek and Mayan astronomers, to William and Caroline Herschel in the 18th century, measurements of the positions of stars relied on visually aligning the observing instrument to the target and recording a mechanical measurement of its configuration. This process was limited to the sensitivity of the human eye and the mechanical repeatability and accuracy of the measuring device over variation in temperature and orientation with respect to the mechanical load imposed by gravity. While the mechanical repeatability and accuracy of well-built instruments can be quite good, as Figure 1-2 indicates, it is often measured in tens of arcseconds.

2.1.2 Photographic Plate Era

The introduction of photography to astronomy in the 19th century enabled an improvement in the process of making astrometric measurements. Although an astronomer still had to dress warmly on a cold night and manually adjust the photographic plate to keep it aligned to a target star while the telescope mechanism experienced mechanical perturbations [32], the resulting long exposure enabled measurement of targets too faint to detect with the human eye.

The ability to record an image eliminated the need to accurately read telescope setting circles and account for misalignment and deflection of the telescope mechanism. In a photograph, the astrometric measurement can be based on relative measurements between stars captured within a single photographic plate. The accuracy of the measurement now depended on the accuracy in the measurement of the centers of images of stars in the developed image rather than the accuracy in the measurement of the pointing of a large telescope. High-precision mechanisms were manufactured to enable the accurate measurement of coordinates of targets detected on the plate. The operation of these measuring engines entailed placing the plate on a moving counterweighted table beneath a microscope and measuring the mechanical position of the plate underneath the microscope when the image of a star was centered in the objective lens [33]. Extensive calibration procedures were required to null out systematic mechanical biases from the measuring apparatus. With careful measurement, error levels on the order of several tens of milliarcseconds could be achieved on plates from telescopes with large focal lengths [34].

The start of the use of photographic plates coincided with the end of the era of refracting telescopes in professional astronomy in the early 1900s. Refracting telescopes can be subject to chromatic aberration, and for this reason many photographic plates were manufactured with photochemicals sensitive to only a narrow passband of light. This fixed the image quality problems that arose from the use of refractive optics, and it also mitigated DCR at the expense of sensitivity.

2.1.3 Electronic Era

The development of television cameras in the middle of the 20th century provided a way of removing the observer from the telescope, thereby enabling the faster and more automated operation required for near-Earth SSA. These detectors consisted of a photocathode scanned by an electron beam in a vacuum tube. The deposition of photoelectrons onto the detector surface modulated the scanning beam, thereby forming an electronic image. Though suitable for remote operation and electronic

processing, the tube detectors by design ran hot and were subject to high levels of dark current [35].

As such, tube cameras were suitable for detecting individual bright targets but were not sensitive enough to regularly detect enough stars within a single frame to form astrometric solutions. The beginning of the space age in the late 1950s saw the use of wide angle telescopes with photographic film for satellite tracking. The Baker-Nunn cameras developed by the Smithsonian Astrophysical Organization were used for SSA from low-Earth orbit up to GEO until the 1980s [36]. The fields of view of the film cameras were wide enough to enable (manual) registration of the images against bright stars. The GEODSS system fielded by the United States Air Force in the 1980s as a replacement for the Baker-Nunn cameras used tube cameras to detect large targets in geostationary orbit, but had to rely on mechanical measurements of telescope pointing to form the observation.

The development of CCDs during the latter half of the 20th century overcame the deficiency in sensitivity of tube cameras and resulted in sensors with low thermal noise and high quantum efficiency (QE) over a wide passband. Silicon, the most inexpensive and widespread detector technology, is sensitive to wavelengths from around 300 nm to 1100 nm, with peak QE exceeding 90% in designs for scientific applications. This is in contrast to single-digit percent peak QE of photographic film. Although the transition away from film and toward electronic detectors was well underway in the 1980s, the early model CCDs were small relative to the photographic plates they replaced. When GEODSS was upgraded to use CCD cameras in the mid 1990s, the MIT LL CCID-16 design used in the upgrade delivered a two-degree field of view with what was then one of the largest monolithic integrated circuits manufactured in the world. As a result, wide angle telescopes designed during the photographic era, such as the Samuel Oschin Telescope¹ at the Palomar Observatory, kept using photographic plates into the early 2000s.

It was not until the late 2000s when the size of imaging surface available with photographic plates came to be matched by monolithic CCD imagers. However, CCDs trade area against readout speed. In applications that image static or slow-moving targets, larger fields of view are attainable at the expense of readout rate. In near-Earth SSA, targets move at several arcseconds per second or faster, necessitating faster readouts of detectors at the expense of contiguous area.

¹This telescope, using photographic plates, was used to discover several KBOs during the late 1990s and early 2000s [37].

2.1.4 Sources of Measurement Error in CCD Astrometry

DCR bias coexists with other error sources in CCD astrometry and useful orbits can be computed from observations in spite of observational errors. Another way of stating this is that the cumulative effect of all error sources attenuate to zero with more data and more arc length. It is important to place DCR into the context of these other error sources in terms of both amplitude and timescale of attenuation. The error sources in CCD astrometry can be thought of in terms of short, medium, and long timescales. On the shortest timescale are atmospheric turbulence, pixel phasing and point spread function (PSF) streaking effects, on the medium timescale are phase angle-dependent illumination effects, and on the long timescale are “fixed pattern” effects from exposure timing bias, constant star catalog bias and DCR.

The random error from atmospheric seeing can be on the order of several hundred milliarcseconds on each exposure, but this error is uncorrelated across (suitably long) exposures. The per-exposure magnitude of this error is inversely proportional to the square root of the exposure time [38] and because it is uncorrelated across multiple exposures, can be thought of as averaging to zero quickly. A second short timescale error source is the combined effect of how the target and fiducial star PSF morphology interacts with the CCD pixel grid. The severity of this error source depends on the characteristics of the particular imaging chip and how wide the PSF is relative to the pixel grid. The data collection for the UCAC series of star catalogs (with a 905 mas wide pixels and seeing-limited PSFs a few pixels wide) measured this effect to be on the order of 12 mas worst-case [39]. In terms of CCD astrometry of moving targets, the timescale over which this error source averages to zero depends on the speed of the target across the detector. If the observing telescope perfectly points at the same star field while the target moves through the field, then the phasing of the target PSF across the detector pixels may be thought of as being drawn from a quickly varying random distribution while the phasing of the fiducial stars’ PSFs across the detector comes from a single draw. For each individual star, the pixel phasing error is constant, but it is averaged over many tens or hundreds of stars, meaning the dominant error is in the localization of the target’s PSF. The target PSF center error averages to zero roughly as fast as the target moves across the pixel array. For asteroids observed near opposition with typical telescopes, this timescale is on the order of hours per pixel, and for geosynchronous satellites it is on the order of many pixels per second.

Carpino’s 2003 analysis places the amplitude of the combined effect of short-timescale errors at roughly 250 mas for exposures suitable for single-frame detection of typical asteroids [40]. Experience

with short (< 10 sec) exposures on small (< 2 m) telescopes implies that 250 mas RMS is a good conservative estimate of this noise term for asteroid observations. It is conservative in the sense while many active surveys under-perform this threshold, many over-perform it as well [14, 41], and it is possible to build and deploy instruments that meet or exceed this accuracy specification with a combination of thermal/optical design, detector choice, and site location.

A medium-timescale systematic error is the difference between observed photocenter and target center of mass that comes from observing an object at nonzero phase angle or an object with a non-uniform mass distribution or surface reflectivity. The upper bound on this error is roughly the object radius divided by the target range. For a 500-meter diameter object observed at 0.01 AU, this quantity is no more than about 35 mas, and is near zero at opposition. How quickly (or if at all) this error “averages to zero” depends entirely on the observing geometry. While this error is not insignificant, by restricting the analysis of asteroids to observing distances of greater than 0.01 AU, it can be neglected with respect to the much larger systematic contribution from DCR. The error can be larger than 35 mas for Earth satellites (see Section 6.8.2.1), but in the mean-value sense averages to zero if satellites are observed over a sufficiently wide diversity of geometries, as shown in the near-zero mean horizontal residual in Figure 1-4.

Lastly there is clock bias. This error source is not considered in this analysis. The choice to neglect it is based on the fact that clock bias is an engineering problem that can be and has been satisfactorily resolved in many sensors. Though it can be a major systematic error source in existing archived data, the flavor of this analysis is to consider the importance of physics-based, rather than engineering-based, systematic observation errors. In the same vein, there is no analysis of any potential charge transfer effects that introduce brightness-dependent and image chip location-dependent astrometric biases [10], as these effects may be attenuated by post-processing calibrations or reduced in newer imaging chips.

2.2 Star Catalog Accuracy

In order to be useful for astrometric reduction, star catalogs must provide a sufficient density of stars so that a target is always in the vicinity of enough stars to form an astrometric solution. The brightness and angular velocity of the target impose constraints on this density. The stars’ brightness cannot differ from that of the target by more than the instrument’s dynamic range in order to maintain comparable localization accuracy for both.

Faint targets are observed with large-aperture telescopes with long exposures. Thus the fiducial catalog must contain a high density of (necessarily) faint stars in order to be useful. Fast-moving satellites, on the other hand, must either be observed with short exposures so that they do not streak very far, or else they must be rate-tracked by the observing telescope. Both scenarios reduce the detection probability of fiducial stars and necessitate the use of brighter stars as fiducials, thereby forcing a larger field of view. Given the constraints on CCD size, larger fields of view generally imply smaller apertures and/or faster f-numbers. Smaller apertures reduce target signal, while faster f-numbers raise background levels, and both reduce SNR on targets and stars. Thus, while the fact of a moving target pushes the desired field of view up, the maximum field of view is limited by detection sensitivity concerns for observing faint satellites.

The density of stars on the sky and telescope and detector sizes being what they are, a good rule of thumb is that star catalogs useful for SSA of GEO need to be complete to about 16th or 17th magnitude in order to consistently generate multi-star astrometric solutions for targets in that orbital regime.

2.2.1 Photographic Plate Catalogs

A major deep-sky astrometric catalog is derived from the Palomar All-Sky Survey, conducted in the 1950s. The Schmidt plates from that survey were digitized in the early 1990s as the USNO-A and USNO-B series of star catalogs, yielding a catalog complete to $V=20$ and accurate to 0.2 arcsec [9]. This catalog was used in asteroid search and near-Earth SSA throughout the 1990s and early 2000s.

2.2.2 Space-Based Catalogs

The first all-sky space-based astrometric catalog, from the European Space Agency's Hipparcos mission, became available in the late 1990s. While accurate to single-digit milliarcseconds, limitations of the detector technology onboard the Hipparcos satellite meant the the resulting catalog was only complete to $V=12$, or a density of only tens of stars per square degree or less [42]. This was a suitable density to register large fields of view, but not suitable for the fields of view and exposures in sensors used for near-Earth SSA or asteroid tracking.

Throughout the 2010s, the European Space Agency's Gaia mission has been operating as a follow-on to the Hipparcos satellite. This mission promises to yield stellar parallaxes as well as proper motions and aims for sub-milliarcsecond accuracy to 19th magnitude. The final data release with full astrometric solutions is expected in the 2020s. When this catalog is released, the density

of high-precision fiducial stars will approach that of the USNO-B catalog. The first data release containing positions, but not proper motions or parallaxes for faint stars, appeared in late 2016 [43].

The Gaia mission’s second data release (DR2) in April of 2018 [44] provides a more complete astrometric solution, including preliminary estimates of proper motions and parallaxes for stars beyond 16th magnitude. Gaia DR2 is a marked improvement in accuracy and its complete (if preliminary) astrometric solution makes it directly usable for SSA.

2.2.3 Ground-Based CCD Catalogs

Densities of hundreds of stars per square degree are required for practical astrometry with telescope fields of view on the order of half a degree or less, and through much of the early 2000s, the Schmidt plate scans were the only practical option. To fill the gap between the accuracy of space-based catalogs and the star density of ground-based catalogs, a series of ground-based catalogs tied to the space-based optical reference frame were compiled in the 2000s. The catalogs were registered against the celestial coordinate frame of the Hipparcos mission on the brighter end of the scale, while ground-based measurements were used to fill in the astrometric parameters of the fainter stars detectable only with the ground-based observing program.

Ground-based observations were conducted by the US Naval Observatory with CCD astrographs in the early 2000s and yielded the USNO CCD Astrograph Catalog (UCAC) series of star catalogs, which became available in 2009 with a completeness to 16th magnitude [45]. Continued re-reductions of this dataset in combination with other ground-based surveys was completed in 2013 with the UCAC4 catalog with a claimed accuracy of under 40 mas [41]. A second data collection campaign in the early 2010s with improved instrumentation yielded the URAT1 catalog, comprising only the Northern Hemisphere sky, with more accurate proper motions and slightly fainter cut-off [11]. Re-reduction of the UCAC dataset using both Hipparcos and Gaia DR1 data yielded the UCAC5 catalog in 2017, containing the same coverage to $V=16$ or so, but astrometric accuracy improved to the sub-30 mas level [46]. It is with the publication of these catalogs in the mid 2010s that the long-timescale systematic error floor in ground-based astrometry stopped being dominated by catalog bias and dropped down to where it was dominated by DCR.

2.2.4 Observation Refinement With A-Posteriori Catalog Corrections

Related to the question of star catalogs for reduction of newly-collected raw imagery or re-reduction of archived raw imagery with improved star catalogs is the question of what to do with older

observations for which raw imagery was not retained. In the mid 2010s, the available ground-based catalogs attained a sufficient improvement in accuracy to enable after-the-fact correction of these archival observations if the catalog used to reduce them was known.

In 2015, Farnocchia published [14] a set of mean corrections to positions and proper motions for star catalogs used to reduce the majority of MPC archived observations. In the absence of knowledge of which particular stars were used to reduce each individual observation, the corrections had to assume that observations were formed from suitable nearby stars, and as such what was published was a set of mean corrections for 0.8 square degree tiles across the sky.

While an improvement in post-fit residuals was realized using these approximate area corrections, as seen in Farnocchia’s paper and in the difference between Figures 1-5(b) and 1-5(c), the process does have limitations. Individual observations can only be corrected using the correction for the tile they belong in. For sensors with larger fields of view than covered by the individual tiles with astrometric solutions derived from those entire fields of view, the correction would be mismatched. Furthermore, as information is generally not retained about which stars were used to form individual MPC observations, and the Farnocchia corrections are color-agnostic by virtue of being defined as fixed RA/Dec space corrections without regard to sensor QE or target elevation, stellar DCR cannot be compensated for using these mean corrections. In Section 6.8.3.2, the necessity of accounting for stellar DCR will be quantified.

2.3 Spectroscopic Measurements of Solar System Bodies

2.3.1 Past and Current Surveys of Asteroids

The reflectance spectra of asteroids are well-characterized. Surveys with multiband photometry conducted by Chapman and Gaffey [47] and Tholen [48] in the 1970s and 1980s ranged from 300 nm to the upper limit of photocathode detectors of the era. Because there is relatively little energy detectable on the ground in the blue, more recent studies have focused on wavelengths above 450 nm and the near infrared. Measurements in near infrared wavelengths capture characteristic absorption features, enabling discrimination of different surface chemistries, but do not provide the information necessary for compensation of astrometric observations made with bare silicon detectors whose detection sensitivity cuts out redward of about 1.1 micron.

The Small Main-Belt Asteroid Spectroscopic Survey (SMASS) [49, 50] provides archived reflectance spectra on many asteroids. The Bus-DeMeo taxonomy derived from this data set (Figure

2-1) is used to classify asteroids by their spectral type. While there is a wide diversity of features in these spectra at near-infrared and mid-infrared wavelengths, what is of interest for analysis of DCR in observations made with silicon detectors is the behavior of the spectra shortward of 1.1 microns.

Of the cataloged asteroid classes populating the main belt, the C-type and S-type asteroids stand out as the two extremes of spectral slope in the silicon passband. Average reflectance spectra of S-type and C-type asteroids are shown in Figure 2-2. C-type asteroids have nearly flat spectral response in the silicon band and S-type asteroids are the most abundant type with the greatest spectral slope in the silicon band.

Because bare silicon has sensitivity below the 450 nm cutoff of SMASS and the atmosphere has appreciable throughput below that wavelength at high altitudes and low zenith distances, it is necessary to extrapolate the published spectra. Examination of the spectral reflectances for the C-type and S-type asteroids from the early multiband photometric surveys indicates that linear extrapolation down from 450 nm to 300 nm is valid for the purpose of predictions, with the caveat that individual asteroids' true reflectance may vary from that of the extrapolated mean class reflectance.


More recent work by Thomas [51], Sanchez [52], and others has focused on understanding the behavior of absorption bands in the reflectance spectra of asteroids as a function of observing phase angle and asteroid surface temperature. While the surface temperature of a main-belt asteroid would be expected to remain nearly constant, near-Earth asteroids' eccentric orbits would subject them to repeated heating/cooling cycles and the depth and shape of the absorption band would vary depending on where the asteroid is in its orbit and not just on the observing geometry. Work on understanding the phase dependence of asteroid spectral reflectances over the silicon passband is limited by the small number of spectroscopic observations on any one asteroid under a wide-enough variety of phase angles and surface temperatures and the significance of this effect on DCR bias cannot be rigorously characterized with presently available data.

It is important to note that of the hundreds of thousands of numbered asteroids in the Solar System, only a few thousand have been spectroscopically classified, whereas over 100,000 have had their albedo measured using multiband photometry extending out to infrared wavelengths [53]. From albedo, it is possible to coarsely infer taxonomic class of an asteroid without a spectroscopic measurement. Objects with low albedo (≈ 0.06) are likely to belong to the C-complex while objects with albedo around 0.2 are likely to belong to the S-complex [54]. There is, however, ambiguity in this coarse classification scheme if the performance metric is accuracy in estimation of DCR bias in observations made with bare-silicon detectors. This is explored in Section 4.2.1.

S-Complex



C-Complex



X-Complex



End Members




Figure 2-1: Bus-DeMeo Asteroid Taxonomy. SMASS and the Bus-DeMeo taxonomy is based on features in objects' reflectance spectra between $0.45 \mu\text{m}$ and $2.45 \mu\text{m}$. The reflectance spectra of many distinct taxonomic classes are similar over the silicon passband. The C-type and S-type spectra represent the most numerous examples of two extremes over the silicon passband: spectrally-flat and red, respectively. Figure credit: DeMeo, Binzel, Slivan, and Bus [50], used with permission.




Figure 2-2: Asteroid Reflectance Spectra for Selected Taxonomic Classes. Reflectance spectra linearly extrapolated down to 300 nm from SMASS [50].

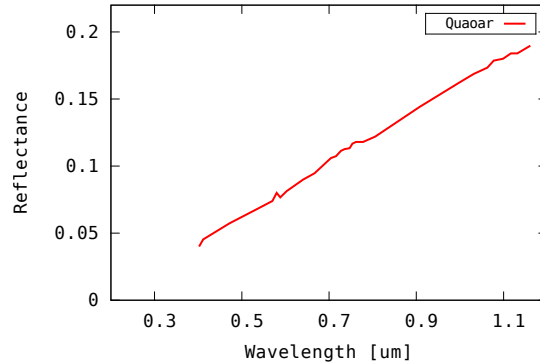


Figure 2-3: Quaoar Reflectance Spectrum. Representative of the extreme red KBO reflectance spectra. Source: Jewitt and Luu [57]. The spectrally-flat Barucci Group I objects (not shown) are the other extreme used for analysis in this thesis.

2.3.2 Kuiper Belt Objects

Many Kuiper belt objects' reflectance spectra are also well-characterized. Unlike most asteroids, their reflectance spectra in the silicon band are largely free of absorption features and are near linear ramps [55, 56]. Luu and Jewitt's measurement of the reflectance spectrum of Quaoar [57] (Barucci Group IV) will serve as a worst-case example of the DCR bias of red KBOs (Figure 2-3). The neutral color Barucci Group I objects [58] have relatively flat reflectance. As will be discussed in Section 4.2.2, the amount of DCR bias one would see in astrometric observations of more color-neutral objects depends on the conventions used in the astrometric processing.

2.4 Spectroscopic Measurements of Satellites

Satellites' spectra have not been as well-studied as those of Solar System bodies. Being composed of man-made materials, there is not the same motivation for scientific discovery. Rather, all spectroscopic measurement campaigns of Earth satellites have been targeted at distinguishing individual RSOs from each other and in identifying debris fragments. Similar to the approach taken by the asteroid community, many of these measurements extend into the near IR and do not always cover the full silicon band.

2.4.1 Satellite Materials

Laboratory measurement of common spacecraft materials provide the majority of the available spectral data over the entire silicon band. The measurements of common materials published Luu

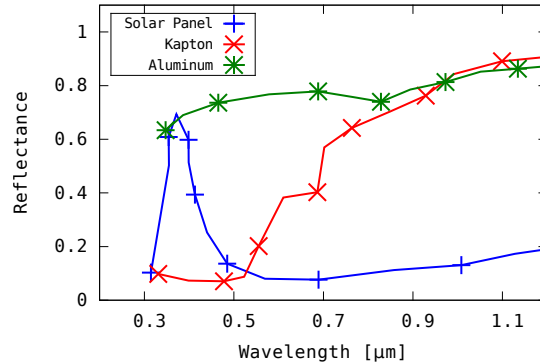


Figure 2-4: Spectral Reflectances of Common Spacecraft Materials. Sources: [59] and [60]. The variation of reflectances among different material samples makes predicting the spectrum of any individual satellite challenging.

et al. [59] and French et al. [60] (see Figure 2-4) provides the basis for the analysis of DCR in observations of satellites in this thesis.

Caution must be exercised in using laboratory measurements of man-made materials. Flat surfaces and thin films exhibit highly variable reflectances depending on phase angle [61], and on the thickness of the laminations in a particular batch of material [60]. In particular, solar panels' reflectance is highly specular and the angular extent of the specular glint off of the solar panel material measured by Bowers appears to be only a few degrees wide.

2.4.2 In-Situ Measurements

In the early 2000s, Jorgensen Abercromby obtained on-orbit measurements of reflectance spectra of payloads and rocket bodies [62]. Comparisons of predictions derived from laboratory measurements and spectra of on-orbit objects do not always agree [63]. Measurements of reflectance spectra of returned flown samples show evidence of weathering that modifies the spectral reflectance over time.

Recent work has focused more on rapid characterization of the composition of newly-launched RSOs and newly-discovered debris for the purpose of identification of object type. Measurements are mostly in the near infrared [64]. Vananti has made on-orbit measurements over the Silicon band for select objects, but again only for characterization purposes and in a piecemeal fashion [65]. Additionally, the NOFS operated a low-resolution slitless spectrograph, again for identification purposes rather than astrometry [66], through 2018. Representative data from NOFS is presented in Chapter 6.

2.5 Atmospheric Refraction

2.5.1 Models of Refraction and Their Applicability

The standard method of computing refraction is via a ray-trace from the ground station out into space. Typically, a spherically symmetric Earth and atmosphere are assumed. This enables use of the invariant refraction relation:

$$r_0 n_0 \sin z_0 = r_1 n_1 \sin z_1 \quad (2.1)$$

where r is the distance to the geocenter, z is the local angle with respect to zenith, and n is the refractive index at each point along the ray-trace. The ray trace for the total bending is computed efficiently via the Auer-Standish algorithm [2]. Estimates of refractive index as a function of local meteorological conditions (temperature, pressure, relative humidity) are well-defined handbook quantities, as are methods to estimate these parameters in the upper atmosphere given local meteorological conditions at the ground site [67].

The ease of calculation of refraction in a well-defined way masks natural variability at low elevation angles. It is difficult to measure this quantity precisely as all measurements must be made of a well-defined target against a well-defined angle reference. One such measurement is the timing of sunrise and sunset. Schaefer and Liller [68] conducted a global measurement campaign in the late 1980s to determine how well refraction models work near the horizon. The resulting estimate is that systematic variability at the horizon is significant and the variability at low elevation is roughly proportional to the amount of total bending. Differential astrometry greatly attenuates the effect of much of that variability, but caution should be exercised in assuming that predictions made using models developed for low zenith distances will hold closer to the horizon.

While some of the simulations and analysis in this thesis includes elevation limits as low as 5° , most of the analysis, and all of the measurements, stays above 20° elevation.

2.5.2 Differential Color Refraction (DCR)

The consideration of DCR in astronomical observing is not new. The implications of DCR were studied in the early 1980s. Filippenko [1] calculated the wavelength-dependent signal loss in a slit spectrograph oriented toward the celestial pole (and not toward zenith) in an equatorial telescope. During the planning of the Voyager flyby of Neptune, DCR bias compensation was necessary in the determination of the orbit of Triton [69, 70]. Stone [71] studied the effect of DCR bias on the

relative astrometry of stars using narrow passbands and predicted a bias of no more than 50 mas in the angle between solar-analogs and O and M type stars. Stone predicted a reduction of that bias by approximately 10 to 20 mas due to interstellar reddening (which to first order makes stellar spectra look more similar to each other). This meant that DCR was of little practical importance in astrometry given the several hundreds of milliarcseconds worth of errors in the fiducial star catalogs available in the 1980s. Stone’s 30 – 40 mas error bound applies to observations made with narrow passbands of targets that have approximate blackbody spectra. It does not hold for broad-band silicon detectors and Sun-illuminated Solar System bodies reduced against distant non-solar-analog stars.

Given that the study of DCR is not new, many astronomical observing programs attempt to mitigate its effect. Common practices include observing through narrow bands and at low zenith distances [72, 73]. Larger telescopes with slower optics can be designed to directly attenuate DCR by introducing dispersion compensating prism pairs into the optical path [74]. All of these mitigation strategies trade sensitivity, search rate, cost, or observing opportunities for accuracy, and are not uniformly instituted across all asteroid observing programs. A majority of observations in the Minor Planet Center archive come from telescopes operating unfiltered silicon detectors, including data from LINEAR, SST, and the Catalina Sky Survey [5, 75, 76], that took no measures to compensate for DCR bias, and a similar situation exists in SSA of the geostationary belt.

2.6 Ephemeris Propagation

The accuracy of orbit predictions is a function of both the accuracy of the observations and the fidelity of the force model used to propagate both over the data fit span and past the last observation. As with potential improvements to account for DCR in astrometric observations also subject to other error sources, the key question to ask is whether an improvement in compensating for DCR bias will be masked by other systematic errors in the orbit prediction process.

2.6.1 Heliocentric Force Model

The heliocentric force model can be accurate to sub-kilometer levels for many decades of integration for “well-behaved” targets in orbits of interest. The force model defined by the DE431 Ephemeris claims sub kilometer-level accuracy on well-constrained targets over integration spans measured in years [77]. This is an assertion to be made with caveats in that this accuracy can only be assured for

objects with a large baseline of high-quality observations. However, as discussed in Section 1.3.2.2, it is not unreasonable to expect orbits on individual asteroids without radar observations to be off by several kilometers at any point in time.

The two fundamental drivers of this error are not only biased observations or small-number statistics, but also uncertainties in non-gravitational forces acting on individual asteroids. For comets, and for asteroids in eccentric orbits with any volatiles on their surface, the dominant non-gravitational perturbation is out-gassing. For asteroids devoid of any volatiles, the Yarkovsky effect is the dominant non-gravitational perturbation and can result in kilometers per year worth of drift with respect to the gravity-only force model [22, 23], which is nearly an order of magnitude lower than the apparent Cartesian shift from unaccounted-for DCR at typical observation distances, meaning DCR compensation would not wash out in the presence of force model uncertainty.

2.6.2 Geocentric Force Model

The force model for geocentric objects also has an effective error floor stemming from two sources. The first and seemingly dominant error is the uncertainty in non-gravitational perturbations from drag in low orbits (not considered in this thesis) and radiation pressure in medium Earth orbit (MEO) and higher. Artificial satellites have an order of magnitude lower density than natural Solar System bodies and radiation pressure dominates the non-gravitational forces acting on them.

The second source of error is uncertainty in the geopotential. While the static portion of the gravity field is known to order and degree in excess of 2000 from decades of surface and satellite measurements [78], corresponding to accelerations well below of 10^{-20} m/sec², the accuracy the time-varying component of the gravity field is an area of active research [79], and modeling radiation pressure from non-isotropic thermal emission of stabilized and active satellites is difficult [80].

The LAGEOS and Etalon satellites are passive spherical satellites with known surface properties, masses, and dimensions, placed in orbits where atmospheric drag is vanishingly small. Laser ranging observations of these satellites provide a well-defined performance metric to evaluate the quality of the geocentric force model [81]. Centimeter-accurate orbits for these objects can only be synthesized if empirical accelerations are included in the model to absorb mismodeling of the gravity field and Earth albedo pressure [82]. The magnitude of the empirical acceleration on the LAGEOS satellites is on the order of 10^{-9} m/s² [83], setting the reference for what the expected error floor should be in a geocentric force model. Further elaboration on this latter point and how it extends to geostationary orbit will be deferred to Section 3.3.6.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Methods

3.1 Overall Theme

At the most basic level, the technical content of this thesis consists of four parts:

1. Making predictions about the effects of DCR on astrometric observations;
2. Analyzing existing astrometric observations and making targeted astrometric measurements of specific targets in order to validate the predictions made in step 1;
3. Simulating scenarios in asteroid search and near-Earth SSA to probe the consequences of unaccounted-for DCR;
4. Demonstrating the compensation of DCR bias in astrometry with concurrent spectroscopic measurements of calibration targets with well-defined ephemerides.

The intent of this work is to ensure that the predictions and simulations are grounded in empirical measurements. The mechanics of this entails writing software to generate large volumes of simulated data and to ingest large volumes of real data. It is therefore necessary to ensure that the underlying algorithms and coordinate transformations in that software agree with standard baselines adopted by organizations such as JPL, IGS, USNO, and IAU.

An example of this appeared in Chapter 1 in the analysis of MPC observations against JPL ephemerides and of Firepond observations of GPS and GLONASS satellites against IGS ephemerides. Validation was evident in the fact that horizontal residuals were zero to within the URAT1 star catalog accuracy, implying that the nonzero signal in the vertical residual was meaningful. The remainder of this chapter will describe data used and the software written for this task.

3.2 Data Sources

This section describes the real data used for this thesis. The data consists of both archived reduced observations of Earth satellites and Solar System bodies, and the third-party reference ephemerides for those objects (used for the analysis portion of the thesis) and raw image and signature data. Several dedicated measurement campaigns to collect observations on Earth satellites, and a single near-Earth asteroid flyby, were executed in support of this thesis.

3.2.1 Reduced Data and Ephemerides

3.2.1.1 International Astronomical Union Minor Planet Center

The MPC hosts a repository of all published optical observations of Solar System bodies. The convention for those observations is ICRF-aligned coordinates in a zero-velocity frame. That is, the effects of both diurnal and annual aberration of light are removed prior to submission to the MPC.

While there are millions of distinct observations in the database from the CCD era, there are at most a few thousand observations on any one asteroid, and many numbered asteroids have fewer than one thousand observations. Further, the observations are stored in a truncated-precision format. Right Ascension precision is limited to 15 mas and declination precision is limited to 10 mas, but many of the observations in the database are specified to a precision of 150 mas in RA and 100 mas in Dec. This truncation error is larger than the expected effect of DCR on any individual observation, thus the analysis of observations of real asteroids must be confined to examining population statistics over large numbers of asteroids observed from many observing stations.

3.2.1.2 Jet Propulsion Laboratory

Ephemerides of real asteroids used in this thesis come from the Jet Propulsion Laboratory Solar System Dynamics Group’s Horizons system [84]. In early 2017, the binary asteroid ephemerides from the Horizons system came as snapshots of modified divided difference polynomials from the 13th order Adams integrator used for propagating Solar System ephemerides. The algorithm for interpolating these snapshots was incorporated from the SPICE toolkit source code [85, 86].

3.2.1.3 International GNSS Service

The International GNSS Service (IGS) maintains a global network of multifrequency receivers to track broadcast signals from GPS and GLONASS satellites [81]. The distribution of the network

is much denser than that of the two constellations' control stations. This enables the IGS to form ephemerides for GPS and GLONASS satellites and ionospheric corrections that, while being delayed with respect to the real-time broadcast ephemeris, are significantly more accurate.

Finalized ephemerides of the centers of mass of active GPS and GLONASS satellites are accurate to several centimeters and are published by the IGS several days after the fact. The finalized orbits and finalized Earth orientation parameters (EOPs) give the nominal inertial angles from a ground station to the center of mass of the satellite to an accuracy of about one milliarcsecond.

The ephemerides are distributed as three-dimensional Cartesian positions in the Earth-fixed ITRF frame at fifteen-minute intervals. The recommended approach [87] to interpolating this data into a six-dimensional state vector is to first convert the Earth-fixed ephemeris into inertial coordinates and to fit a trigonometric series to a window of datapoints near the desired interpolation time. The fundamental frequency of the trigonometric series is the (roughly) half-day orbital period of GPS and GLONASS satellites. A third-order fit to the seven nearest datapoints was used.

3.2.2 On-Sky Measurements

3.2.2.1 MIT George R. Wallace, Jr. Astrophysical Observatory



Figure 3-1: MIT George R. Wallace Astrophysical Observatory, Westford, MA. Altitude near sea level. Image credit: Google Earth

The Wallace Astrophysical Observatory in Westford, MA (Figure 3-1), is MIT's local facility for instruction in observational astronomy and astronomical research. The facility has several telescopes suitable for observing slowly-moving astronomical targets. The August/September 2017 near-Earth flyby of the asteroid 3122 Florence was observed from this site using narrow passband Sloan filters frequently used in observations of asteroids.



Figure 3-2: 1.3 Meter Telescope at USNO, Flagstaff, AZ. Image credit: USNO

3.2.2.2 United States Naval Observatory Flagstaff Station

The 1.3 m telescope at the NOFS in Arizona (Figure 3-2) is equipped with a slitless spectrograph for making measurements to support satellite material identification of RSOs along the geostationary belt [66]. Signature data from this telescope is used in this thesis to survey the variation in DCR bias over the geostationary belt.

3.2.2.3 MIT Lincoln Laboratory Experimental Test System




Figure 3-3: MIT Lincoln Laboratory Experimental Test Site, Socorro, NM. Altitude 1.6 km, MPC Code 704. Image credit: MIT LL

MIT Lincoln Laboratory's ETS in Socorro, New Mexico (Figure 3-3), has supported the Laboratory's optical SSA mission since the 1970s. The ETS was the site of the LINEAR program from 1996 to 2016. Imagery collected using the original GEODSS Test System between 2000 and 2012 was archived and a subset of this imagery containing asteroids with known spectral types was

retrieved for re-analysis for this thesis. The site remains active and one of its telescopes was tasked with observations of geostationary satellites and the asteroid 3122 Florence for this thesis.

3.2.2.4 MIT Lincoln Laboratory Firepond Optical Facility



(a) Image Credit: Wikimedia

(b) Image Credit: MIT LL

Figure 3-4: MIT Lincoln Laboratory Firepond Optical Facility, Westford, MA. Altitude near sea level.

The majority of raw imagery used to develop the techniques in this thesis comes from the Firepond Optical Facility in Westford, MA (Figure 3-4). The main attraction at the facility is a 48” f/3 telescope on a high-speed and high-accuracy tracking mount with space for numerous auxiliary instruments. The mount is capable of tracking both astronomical targets and Earth satellites.

The 48” telescope has had a storied history. In 1962, when it was located in Lexington, MA, it was used in the first-ever Lunar laser ranging experiment [88]. In the late 1970s, it was relocated to Westford and used to develop satellite tracking techniques [89], and throughout the 1980s it was used in the development of coherent laser radar technologies for ballistic missile defense [90] and in support of Lincoln Laboratory’s efforts in the development of adaptive optics [31, 91].

The facility was mothballed in the mid 1990s following the end of the Cold War and was re-activated in 2009 as testbed to support Lincoln Laboratory’s SSA research. Since 2009, the 48” telescope has had an Andor 888 electron-multiplying CCD camera. Timing pulses from the camera tag each individual exposure to a GPS-referenced system clock, eliminating the major source of systematic error commonly found in astronomical images¹. This setup is suitable for reliably generating multi-star astrometric solutions on satellites above about 5000 km altitude at up to 8 Hz.

¹Dave Tholen tells me that some of the data collected from the top of Mauna Kea had a frame timing uncertainty of several seconds.

3.3 The OrbitTool Software Framework

Success in astrodynamics and SSA is 5% concept and 95% rigorous accounting of coordinate transformations, International Earth Rotation and Reference Systems Service (IERS) conventions, and standard algorithms. For the purpose of analyzing observations, astrometric reductions, correcting DCR bias, and injecting synthetic bias to simulate the effects on orbit solutions for asteroids and Earth satellites, I have written a (mostly) from-scratch precision orbit determination, astrometric image processing, and coordinate transformation toolbox called OrbitTool. The primary motivation for writing a purpose-built software package was to enable rapid analysis of DCR bias in the coordinate systems appropriate for that analysis, and to do so in a transparent manner with as few pieces of black-box software packages as possible in the critical path.

Other alternatives to writing a new codebase were considered and rejected for the following reasons:

- GTDS: Written in FORTRAN, hard-coded to assume diagonal covariances on optical observation types, no good way to inject/remove bias aligned with local vertical from optical observations [92].
- GMAT: Support for radio-frequency observations only, no angle observations, propagator implementation favors heliocentric over geocentric environments, and had other missing features as of 2017.
- STK/ODTK: Black-box commercial package, cumbersome to make it do what it is not intended to already do.
- OrbFit: Written in FORTRAN, limited documentation, assumed diagonal covariances on optical observation types, limited input/output precision, number of input observations per run limited to 10,000, does not support geocentric orbits.

By the time of completion of this thesis, OrbitTool consists of about 160,000 lines of C++ code, of which 85,000 are the off-the-shelf SOFA [93] and NOVAS [94] packages for implementing the terrestrial-to-celestial coordinate transformations of the IAU Conventions and approximately 75,000 lines specifically written for this thesis. The breakdown of those 75,000 lines is given in Table 3.1. Note that of that code, only approximately 5,000 lines pertain to atmospheric refraction, while the rest of it is for coordinate transformations, book-keeping, and input data handling.

Table 3.1: Approximate Lines of Code in OrbitTool by Functional Area

Function	Lines
Image processing	17,000
Observation and geometry handling	14,000
Force models	8,000
Simulation tools	7,000
Orbit fitting	6,000
Data ingestion	6,000
Refraction	5,000
Visualization and analysis	4,000
Glue code	8,000

This section will describe some of the features of the OrbitTool code as they pertain to the study of DCR in this thesis.

3.3.1 Design Philosophy

The design philosophy for OrbitTool was centered around its envisioned usage in support of the research objectives of this thesis. The result was a variant of the Unix Philosophy: the creation of a large number of small command-line programs that individually accomplish one well-defined task, and utilize a common set of validated library routines for coordinate transformations, orbit propagation and fitting, atmospheric refraction, and data handling. The software also had to be runnable on a wide variety of computers, some of which by necessity run older operating systems, and thus needed to have a minimum number of external dependencies.

The choice of a compiled programming language like C++ over a scripting language like Python or MATLAB was dictated as much by performance considerations as it was by the desire to enforce the design philosophy expressed in the previous paragraph. Scripting environments and interpreted languages without strong datatype enforcement allow a programmer to get lazy, especially if working alone on a product with a user base of one. A compiled language helps to enforce the objective of making small standalone programs whose invocation in a repeated fashion on different datasets generates the desired research results.

3.3.2 Coordinate System Definitions

All orbit propagation and observation is done in Cartesian coordinates aligned with The International Celestial Reference Frame (ICRF). Barycentric coordinates and TDB time are used for asteroids and geocentric coordinates and TAI time are used for Earth satellites. Simple Galilean transformations transform Cartesian states between the two systems and the NOVAS function

`tt2tdb()` is used to transform between geocentric and barycentric times in compliance with Section 10.1 of the IERS 2010 conventions [79].

For the purposes of generating observations and computing residuals for angle observations, the transformations between inertial GCRS coordinates and the Earth-fixed (ITRF) station coordinates are implemented as follows:

Station survey coordinates are defined in geodetic coordinates assuming the WGS-84 ellipsoid. These geodetic coordinates are used to compute both the geocentric Cartesian vector of the station and to define a local North-East-Down (NED) coordinate system. The Cartesian coordinates are computed using a mix of the NOVAS `terra()` function and the SOFA `gd2gc()` function with the WGS-84 ellipsoid. The two agree to centimeters. For a site located at a geodetic longitude, latitude (θ, ϕ) , the transformation between the ITRF and the local NED frame is given by:

$$\mathbf{R}_{\text{ned2itrs}} = \begin{bmatrix} -\sin \phi \cos \theta & -\sin \theta & -\cos \phi \cos \theta \\ -\sin \phi \sin \theta & \cos \theta & -\cos \phi \sin \theta \\ \cos \phi & 0 & -\sin \phi \end{bmatrix} \quad (3.1)$$

The transformation between the ITRF and GCRF systems is defined as:

$$\mathbf{R}_{\text{itrf2gcrf}} = \mathbf{PNSW} \quad (3.2)$$

Where \mathbf{P} is the precession matrix and \mathbf{N} is the nutation matrix, which are defined according to the IAU 2000/2006 precession-nutation theory, \mathbf{S} is the rotation about the positive Z direction of the Greenwich apparent hour angle, as corrected by the UT1-UTC measurement, and \mathbf{W} is the pole-wander correction. The NOVAS `precession()`, `nutatation()`, `spin()` and `wobble()` functions are used to populate these matrices.

For the purposes of this work, station displacement due to solid Earth tides, which is on the order of a 60 cm worst-case, is not considered, and the UT1-UTC correction and pole wander variables are taken as-is from the `finals2000A.all` file as distributed by the United States Naval Observatory (USNO). Precision beyond that is not necessary for accuracy to the 10-milliarcsecond level for angle measurements of targets beyond the GPS and GLONASS orbital altitude.

The transformation between relative Cartesian states, image offsets aligned to GCRS RA and Dec axes, and image offsets aligned with local horizontal and vertical directions at each station is defined as follows, in reference to Figure 3-5. Given a relative Cartesian vector $[X, Y, Z]^T$, expressed

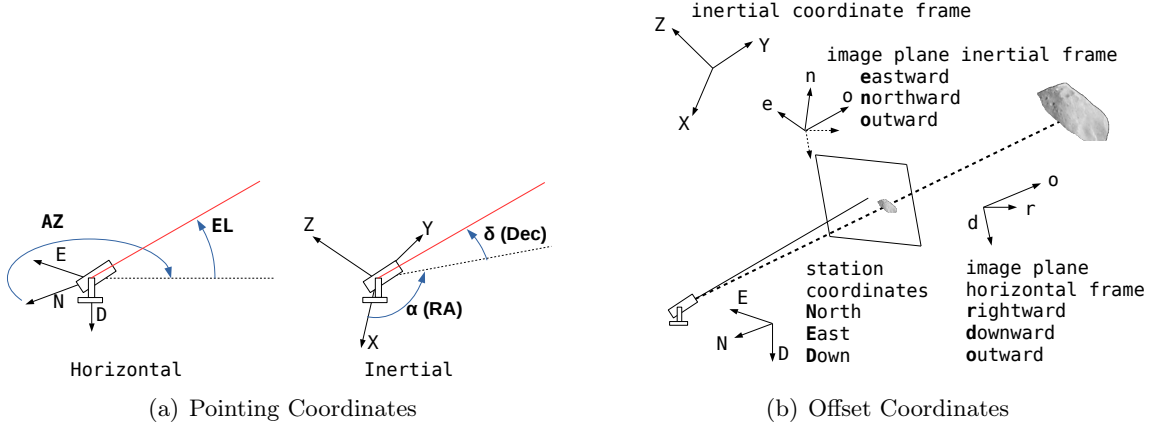


Figure 3-5: Schematic of Local Right/Down/Out and East/North/Out Coordinates. (a) Definition of azimuth/elevation and right ascension/declination for the same observation. (b) Definition of local ENO and RDO coordinates for residuals and offsets from a nominal astrometric observation.

in inertial GCRF coordinates, right ascension (RA) α and declination (Dec) δ are defined as:

$$\alpha = \text{atan2}(Y, X) \quad (3.3)$$

$$\delta = \text{atan2}\left(Z, \sqrt{X^2 + Y^2}\right) \quad (3.4)$$

Given the same vector transformed into station-local North-East-Down coordinates:

$$[N, E, D]^T = \mathbf{R}_{\text{ned2itrs}}^T \mathbf{R}_{\text{itrf2gcrf}}^T [X, Y, Z]^T \quad (3.5)$$

azimuth and elevation are defined as:

$$\text{AZ} = \text{atan2}(E, N) \quad (3.6)$$

$$\text{EL} = \text{atan2}\left(-D, \sqrt{N^2 + E^2}\right) \quad (3.7)$$

where for simplicity, the negative of elevation is used internally in the code so that the same code that transforms between RA and Dec and Cartesian inertial coordinates can transform to azimuth and negative elevation in station-local coordinates.

The image space offset coordinate system is defined as being tangent to the unit sphere in the direction of the observation with axes aligned to either the direction of the GCRS Eastward and Northward for inertial offsets or local downward and local rightward (Westward) for DCR analysis purposes. The terms “horizontal” and “vertical” are used throughout this document to refer to

components in the RDO frame.

$$\mathbf{R}_{\text{rdo2ned}} = \begin{bmatrix} -\sin AZ & \sin EL \cos AZ & \cos EL \cos AZ \\ \cos AZ & \sin EL \sin AZ & \cos EL \sin AZ \\ 0 & \cos EL & -\sin EL \end{bmatrix} \quad (3.8)$$

$$\mathbf{R}_{\text{eno2gcrf}} = \begin{bmatrix} -\sin \alpha & -\sin \delta \cos \alpha & \cos \delta \cos \alpha \\ \cos \alpha & -\sin \delta \sin \alpha & \cos \delta \sin \alpha \\ 0 & \cos \delta & \sin \delta \end{bmatrix} \quad (3.9)$$

Star catalog corrections and most error statistics are defined in the inertial ENO frame. Atmospheric bias physics lives in the downward component of the RDO frame. To study the effects of DCR in simulation, synthetic *upward/downward* biases $\Delta\theta$ can be injected into observation unit vectors as follows:

$$\mathbf{x}_{\text{gcrs,baised}} = \mathbf{x}_{\text{gcrs}} + \mathbf{R}_{\text{ned2gcrf}}\mathbf{R}_{\text{rdo2ned}} [0, \Delta\theta, 0]^T \quad (3.10)$$

3.3.3 Propagation of Ephemerides

Simulated ephemerides are generated by numerical propagation. All integration of ephemerides is done in Cartesian coordinates with a force model appropriate to the orbit.

The geocentric force model contains point mass Earth, Moon, Sun and Planets, as well as up to 100×100 harmonics of the geopotential (in practice 8×8 is used to model the geosynchronous regime), a time-varying solid and ocean tide potential in line with the IERS 2010 conventions [79], and a general a relativity correction. Empirical accelerations defined in the spacecraft-local in-track, cross-track, and radial directions and either constant in magnitude or periodic at once-per-orbit [82] are also implemented in the geocentric force model to match the conventions of the International GNSS Service (IGS) and International Laser Ranging Service (ILRS).

The heliocentric model contains point-mass major bodies, the ‘‘N16’’ minor planet perturbers, and the parametrized post-Newtonian correction for the Sun (including the minor planet perturbers in the potential) described by Moyer [95]. A Yarkovsky acceleration of the form:

$$\ddot{\mathbf{r}} = A_2 \left(\frac{r_o}{r}\right)^2 \mathbf{a}_t \quad (3.11)$$

is optionally included in the force model in order to gauge the effect of DCR bias in astrometric observations on the estimation of non-gravitational acceleration.

For numerical propagation, I use the multi-step Adams integrator implementation in the GNU Scientific Library [96]. Unfortunately the GSL as implemented does not provide a mechanism for storing dense output of an integrated quantity, and it is necessary to bypass the GSL’s API to directly access the intermediate states for later recall in the orbit fitting loop.

The GSL’s implementation of the Adams-Moulton PECE algorithm is based on the algorithm as described by Brown, Byrne, and Hindmarsh [97, 98] and stores each component of the state as a variable-coefficient polynomial array:

$$z_n = \left[y_n, h_n \dot{y}_n, \dots, h_n^q y_n^{(q)} / q! \right] \quad (3.12)$$

at the end of each propagation step h_n (Equation 2.1 of [98]). Dense output of the variable y for a time in the interval $(-h_n, 0]$ requires rescaling the above vector for the desired timestep. OrbitTool stores snapshots of these z arrays for all integrated quantities and recalls them for interpolation within the integrated span to compute model measurements, measurement partials, and residuals.

This integrator is hard-coded to a maximum of 8th order, and in practice does not exceed 5th or 6th order for most cases. The choice of polynomial basis functions rather than divided differences limits the order of the integrator, but for the purposes of this thesis, it is sufficient.

The integrated quantities are the six elements of the Cartesian state vector, additional invariant elements characterizing the non-gravitational perturbations or empirical force magnitudes as necessary, the columns of the state transition matrix (STM), and the columns of the parameter sensitivity matrix (PSM).

The STM $\Phi(t)$ is the partial of the state at time t with respect to the initial seed state of the propagator. Its evolution is defined by [80]:

$$\dot{\Phi} = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} \\ & \frac{\partial \mathbf{a}}{\partial \mathbf{x}} \end{bmatrix} \Phi \quad (3.13)$$

where $\frac{\partial \mathbf{a}}{\partial \mathbf{x}}$ are the instantaneous partials of the force model with respect to the Cartesian state. Similarly, if there are n_p invariant non-Cartesian parameters in the fit, such as SRP AMR, the sensitivity of the propagated state to these parameters may be appended as columns to the STM giving the combined STM/PSM matrix $[\Phi(t), \mathbf{P}(t)]$. The evolution of this combined STM/PSM is

a modified version of Equation 3.13,

$$\frac{d}{dt} [\Phi, \mathbf{P}] = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} \\ & \frac{\partial \mathbf{a}}{\partial \mathbf{x}} \end{bmatrix} [\Phi, \mathbf{P}] + \begin{bmatrix} \mathbf{0}_{3 \times 6} & \mathbf{I}_{3 \times n_p} \\ \mathbf{0}_{3 \times 6} & \frac{\partial \mathbf{a}}{\partial \mathbf{p}} \end{bmatrix} \quad (3.14)$$

where $\frac{\partial \mathbf{a}}{\partial \mathbf{p}}$ are the instantaneous partials of the force model with respect to the invariant parameters.

In the GSL numerical integrator framework, the state vector and the columns of the STM and the PSM are concatenated into a single large vector and treated the same way as any other quantity to be integrated would be.

3.3.4 Optical Observation Model

If dealing purely with simulated observations, it is not necessary to exactly match the computation of finer-grained effects like light travel time, relativistic aberration of light, and gravitational deflection of light. However, this thesis work is intended to be rigorously grounded against empirical measurements and existing state-of-the-art astrodynamics software packages. As such, it is necessary to exactly model the behavior of all standard data reduction steps as implemented by actual observers.

3.3.4.1 Aberration

In reference to Figure 3-6, all targets are subject to light travel time correction. A further correction is necessary to compute the apparent direction of incident light at a moving observer \mathbf{p}_0 to account for relativistic aberration of light [67] via:

$$\mathbf{p}_1 = \frac{\beta^{-1} \mathbf{p}_0 + \left(1 + \frac{(\mathbf{V} \cdot \mathbf{p}_0)/c}{1 + \beta^{-1}}\right) \mathbf{V}/c}{1 + (\mathbf{V} \cdot \mathbf{p}_0)/c} \quad (3.15)$$

where the observer velocity \mathbf{V} includes both the velocity of the geocenter in barycentric coordinates and the Earth surface velocity at the station, or only the Earth surface velocity for observation and orbit determination of objects in geocentric orbit.

When performing data reduction and measuring relative angles between fiducial stars and the target, the transformation of stars into focal plane coordinates must include the full aberration correction to go from the barycentric coordinates in which stars are defined to station coordinates on a rotating Earth.

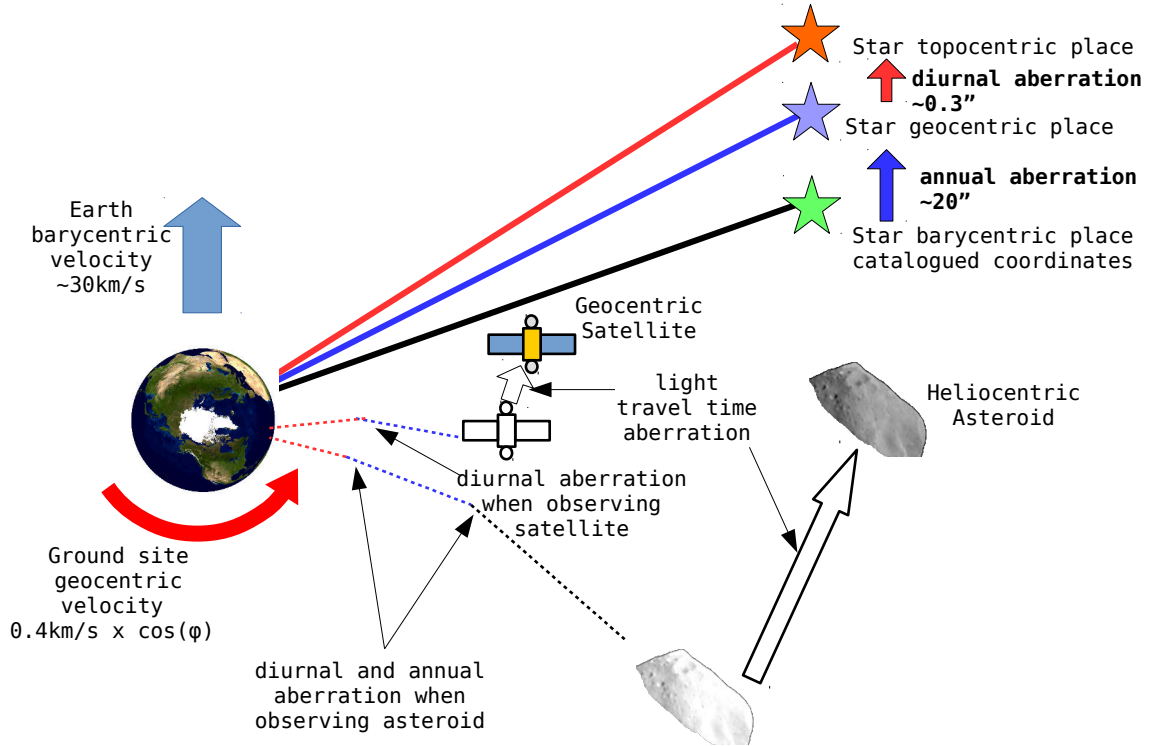


Figure 3-6: Relativistic Aberration of Light in Optics Observation Model

The convention of how to store reduced observations varies by context. When observing satellites whose orbital motion is defined in GCRF and not BCRF, the annual aberration must be left in. The diurnal aberration may be left in, necessitating its inclusion in the observation model, or it may be undone, necessitating no aberration correction in the observation model when executing an orbit fit. For asteroids whose orbital motion is defined in the BCRF, both the annual and diurnal aberration may be left in, necessitating their inclusion in the observation model, or they may both be left out, necessitating only light travel time correction when doing a fit.

Different pieces of software choose different conventions and it is necessary to be mindful of them when working with real data. The Minor Planet Center's optical observations have both diurnal and annual aberration removed from them, meaning it is not necessary to apply them again when computing observation residuals. Some MIT LL software removes diurnal (but not annual) aberration to generate observations of satellites, and some leaves it in.

The terminology is a little loose, but *geometric* or *astrometric* coordinates exclude any aberration correction, but might include light travel time correction for a moving target, *geocentric* observations include annual aberration, and *topocentric* observations additionally include diurnal aberration. Astrometric reductions performed with OrbitTool output *topocentric* observations.




Figure 3-7: Refraction and Parallax Geometry. A ground site is located at an altitude h_o above a spherical Earth. The geometric zenith distance to the target is z_g and the geometric slant range is r_g . True path of light from target to ground site is shown schematically in red and the apparent zenith distance of the target z_a is the zenith distance of a fictitious star directly behind the target. The parallax is a *downward* bias for near-Earth targets. The total bending angle Δz is defined as the distance between the (unmeasurable) observed zenith distance at the ground site z_o and z_a the apparent zenith distance of a target at infinity. The exit zenith angle at the conventionally-established boundary of the atmosphere z_v is computed via the refraction invariant equation $(R_E + h_o)n_o \sin z_o = (R_E + h_s) \sin z_v$.

3.3.4.2 Refraction and Parallax

A digression is called for to lament on the overloaded term “parallax.” “Stellar parallax” or just “parallax” to astronomers means the apparent shift over the year in the direction of incidence of light from a star outside the Solar System, and is typically measured in milliarcseconds. “Parallactic refraction” [99], or “refractive parallax,” or just “parallax” refers to the offset between the geometric vector from observing station to the target and the apparent vector to a hypothetical distant star behind the object caused by the bending of the light from the object downward to observing station (Figure 3-7). For much of the discussion in this thesis, the term “parallax” will have the latter meaning when discussing observations of satellites and the former meaning when discussing stars.

In reference to Figure 3-7, assuming circular symmetry and by application of the refraction invariant, and similar triangles, it can be shown that for a ground station at height h_o observing a

target at a *geometric* zenith distance z_g and slant range r_g , the parallax p obeys the relation:

$$\sin p = \frac{R_E + h_o}{r_g} (n_o \sin z_o - \sin z_a) \quad (3.16)$$

The true observed zenith distance z_o cannot be measured directly and must be inferred from a measurement of the apparent RA/Dec of a star transformed into an apparent zenith distance using IERS conventions. Thus while z_a can be considered a measured quantity, z_o is inferred from the refraction model.

Let the total bending be Δz . Substituting $z_o = z_a - \Delta z$, and $z_a = z_g + p$ we get the following relationship, which is easier to handle in an observation model that begins with geometric angles and slant ranges:

$$\tan p = \frac{\frac{R_E + h_o}{r_g} (n_o \sin (z_g + \Delta z) - \sin z_g)}{1 - \frac{R_E + h_o}{r_g} (n_o \cos (z_g + \Delta z) - \cos z_g)} \quad (3.17)$$

The total bending Δz is rightly calculated at the actual observed zenith distance z_o , which requires knowledge of z_a . This requires an iterative computation to solve for the observed zenith distance given an apparent zenith distance. When computing the parallax bias in an observation model, it is sufficient to substitute z_g for z_a in this iteration.

For GPS satellites, the refractive parallax can be as much as several hundred milliarcseconds at low elevations, but for geostationary objects it rarely goes above 50 mas, and it is completely negligible for asteroid observations.

For high elevation angles, it is not necessary to use the Auer-Standish algorithm to compute the total bending Δz . For elevations above about 23 degrees, it is possible to compute the total bending by lumping the entire atmosphere into a single layer about 8 km (the scale height) high with constant refractive index. This approach yields a parallax calculation consistent with the full refraction integral to about 10 mas. The agreement degrades to 20 mas at an elevation of 20 degrees and 75 mas at 15 degrees.

3.3.5 Orbit Fitting

The orbit determination component of OrbitTool follows mostly the conventions of Montenbruck and Gill [80]. Orbit fitting is run iteratively using weighted least squares. There is no assumption that the covariance of each set of input observations is diagonal.

The least squares formulation begins with the partials of the observation vector \mathbf{y} with respect

to the initial state \mathbf{x}_0 . The observation partials are computed separately for each independent measurement y_i at times t_i as:

$$\frac{\partial y_i}{\partial \mathbf{x}_0} = \frac{\partial y_i}{\partial \mathbf{x}(t_i)} \frac{\partial \mathbf{x}(t_i)}{\partial \mathbf{x}_0} \quad (3.18)$$

where $\frac{\partial \mathbf{x}(t_i)}{\partial \mathbf{x}_0} = [\Phi(t_i), \mathbf{P}(t_i)]$ is the combined state transition and parameter sensitivity matrix.

In standard unweighted least squares, the partials of the observation vector \mathbf{y} with respect to the initial state \mathbf{x} and the observation residuals $\Delta \mathbf{y}$ are computed for each observation time and stacked into a single large matrix:

$$\Delta \mathbf{y} = \begin{bmatrix} \Delta y_1 \\ \vdots \\ \Delta y_N \end{bmatrix} = \underbrace{\begin{bmatrix} \frac{\partial y_1}{\partial \mathbf{x}(t_1)} \Phi(t_1) \\ \vdots \\ \frac{\partial y_N}{\partial \mathbf{x}(t_N)} \Phi(t_N) \end{bmatrix}}_{\mathbf{A}} \Delta \mathbf{x} \quad (3.19)$$

The differential correction to the state estimate $\Delta \mathbf{x}$ is computed by:

$$\Delta \mathbf{x} = \mathbf{A}^+ \Delta \mathbf{y} \quad (3.20)$$

where \mathbf{A}^+ is the left pseudoinverse of \mathbf{A} and $\Delta \mathbf{y}$ are the observation residuals with respect to the current state estimate. In OrbitTool, the result of Equation 3.20 is obtained by computing the scalar value decomposition of \mathbf{A} using the routines in the GNU Scientific Library (GSL).

Unweighted least squares is fine if all the observations are homogeneous and have the same covariance, but if not, then the fit needs to be weighted by the formal error in the observations to take full advantage of the information content in the data [100]. In weighted least squares, Equation 3.19 is premultiplied by a weighting matrix $\mathbf{W}^{\frac{1}{2}}$,

$$\mathbf{W}^{\frac{1}{2}} \Delta \mathbf{y} = \mathbf{W}^{\frac{1}{2}} \mathbf{A} \Delta \mathbf{x} \quad (3.21)$$

and the resulting step is computed as:

$$\Delta \mathbf{x} = \left[\mathbf{W}^{\frac{1}{2}} \mathbf{A} \right]^+ \mathbf{W}^{\frac{1}{2}} \Delta \mathbf{y} \quad (3.22)$$

I have intentionally written the weighting matrix as $\mathbf{W}^{\frac{1}{2}}$ to emphasize that it needs to be given as the symmetric square root inverse of the formal covariance matrix of the input observations $\Sigma_{\mathbf{yy}}$.

That is, if the eigendecomposition of Σ_{yy} is:

$$\Sigma_{yy} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T \quad (3.23)$$

then,

$$\mathbf{W}^{\frac{1}{2}} = \mathbf{V}\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{V}^T \quad (3.24)$$

where $\mathbf{\Lambda}^{-\frac{1}{2}}$ is a diagonal matrix containing the inverse square roots of the eigenvalues of Σ_{yy} . This is a well-defined quantity because covariance matrices of real random variables are by definition symmetric positive definite and have real and positive eigenvalues.

At the terminal iteration, Equation 3.22 defines the sensitivity of the state estimate to the input observations. Assuming only Gaussian noise in the observation and no unmodeled bias², the relationship between the formal covariance of the state Σ_{xx} and Σ_{yy} is defined by [101]:

$$\Sigma_{xx} = \left[\mathbf{A}^T (\mathbf{W}^{\frac{1}{2}})^T \mathbf{W}^{\frac{1}{2}} \mathbf{A} \right]^{-1} \mathbf{A}^T (\mathbf{W}^{\frac{1}{2}})^T \mathbf{W}^{\frac{1}{2}} \Sigma_{yy} (\mathbf{W}^{\frac{1}{2}})^T \mathbf{W}^{\frac{1}{2}} \mathbf{A} \left[\mathbf{A}^T (\mathbf{W}^{\frac{1}{2}})^T \mathbf{W}^{\frac{1}{2}} \mathbf{A} \right]^{-1} \quad (3.25)$$

If $(\mathbf{W}^{\frac{1}{2}})^T \mathbf{W}^{\frac{1}{2}} = \Sigma_{yy}^{-1}$ then the interior of Equation 3.25 cancels, and:

$$\Sigma_{xx} = \left[\mathbf{A}^T \Sigma_{yy}^{-1} \mathbf{A} \right]^{-1} \quad (3.26)$$

$$= \left[\mathbf{A}^T (\mathbf{W}^{\frac{1}{2}})^T \mathbf{W}^{\frac{1}{2}} \mathbf{A} \right]^{-1} \quad (3.27)$$

$$= \left[(\mathbf{W}^{\frac{1}{2}} \mathbf{A})^T (\mathbf{W}^{\frac{1}{2}} \mathbf{A}) \right]^{-1} \quad (3.28)$$

where in the last form it is evident that the state covariance is an inverse of an inner product with itself of a matrix that was already computed for Equation 3.21.

I have implemented observation partials for two-way radar range and Doppler measurements, three- and six-dimensional Cartesian states (to be able to fit orbits to reference trajectories specified as a series of three- or six-dimensional Cartesian vectors) and for angle measurements from optical sensors. The purpose of this thesis and this software was to study biases that are confined to the local vertical coordinate. As such, in the implementation of the measurement partials for optical sensors, I have included an option to rotate observation residuals from ICRS-aligned northward/eastward coordinates to station-aligned horizontal/vertical coordinates, and have also included an option to discard the vertical part of the measurement. This rotation happens in the same place where

²The Gaussian noise and zero-bias assumption is common in orbit determination [14, 40], and the point of this thesis is to examine what happens when this assumption is violated.

the weight matrix is applied in Equation 3.21. That is, the relevant block in the $\mathbf{W}^{\frac{1}{2}}$ matrix is premultiplied with a rotation from ICRS to station-aligned coordinates given by:

$$\mathbf{R}_{\text{eno2rdo}} = \mathbf{R}_{\text{rdo2ned}}^T \mathbf{R}_{\text{ned2itrs}}^T \mathbf{R}_{\text{itrf2gcrf}}^T \mathbf{R}_{\text{eno2gcrf}} \quad (3.29)$$

as defined in Section 3.3.2.

3.3.6 Validation

DCR bias for typical stars and SSA targets is on the order of 50 – 100 mas, and in order to identify evidence of this bias in real observations of real targets and to analyze the effect of this bias on realistic orbit solution scenarios, it is necessary to ensure that the coordinate transformations, observation model, numerical propagation, force model, and fitting routines are all accurate to a level commensurate with the size of this angular bias for both geocentric and heliocentric objects. The validation for heliocentric propagation and observation is made against the JPL Horizons system and the validation of the geocentric force model is made against IGS [81] and ILRS [102] precision ephemerides.

The process of evaluating the heliocentric force model consists of selecting real asteroid targets from the Horizons system, seeding the OrbitTool propagator with a reference state from one of the snapshots in the Type 1 SPICE kernel for that object, and comparing the output of the propagator against the reference SPICE kernel from Horizons. The evaluation is made against an object both with and without the non-gravitational accelerations in its JPL force model. Residuals for a five-year propagation beginning around epoch 2010.0 are shown in Figure 3-8.

Agreement is better than three meters with respect to the purely gravitational and relativistic forces in the Horizons system, and to better than six meters over five years when Yarkovsky acceleration is included. In the worst case of making an observation of an asteroid at the edge of the Earth’s gravitational sphere of influence, a Cartesian state error of 6 m at a range of 1.5×10^6 km is an angle error of about 1 mas. This is more than an order of magnitude smaller than the typical effect of DCR, meaning the OrbitTool propagator is sufficiently accurate to analyze the effect of DCR bias in Solar System bodies.

The process of evaluating the geocentric force model is more involved. The geocentric force model implemented in the OrbitTool largely follows the IERS 2010 Conventions [79] for the static and time-varying gravity field and relativistic correction, uses the FES2004 ocean tide model [104],

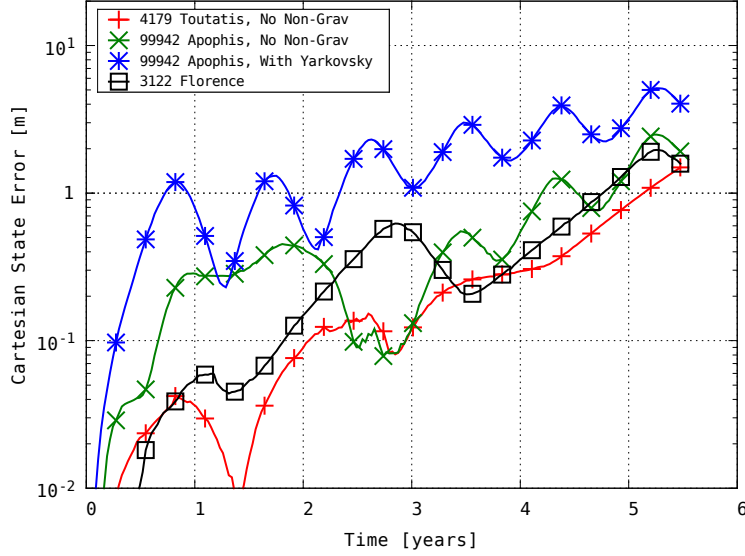


Figure 3-8: Heliocentric Force Model Validation. Binary ephemerides were requested from the JPL Horizons system. The ephemerides were used to seed the OrbitTool propagator and propagated for 2000 days. The low magnitude of the Cartesian difference between the OrbitTool heliocentric propagator and the Horizons ephemeris validates the implementation of the heliocentric force model and propagator inside OrbitTool. For asteroids with defined Yarkovsky acceleration, both the real asteroid and an equivalent ephemeris with the Yarkovsky term set to zero were compared. The agreement to better than 10 m over five years is sufficient to enable the analysis performed in this thesis. The source of the mismatch is likely a combination of the fact that Horizons uses a higher order integrator than what is available in the GNU Scientific Library and performs all calculations using long double precision floating point [103], rather than the double precision in OrbitTool.

and includes direct solar radiation pressure and Earthshine/albedo pressure as a single area-to-mass ratio parameter. The albedo pressure model uses a simple seasonal model for zonal surface emissivity and reflectivity described by Knocke [105]. All major planets’ gravitational contributions are also modeled.

ILRS and IGS reference ephemerides are claimed to be accurate to the centimeter level in posterior analysis using integration spans on the order of many days. It is critical to note that this accuracy is only achievable in posterior analysis as it relies on empirical accelerations to make up for force modeling errors [82]. For LEO orbits such as the TOPEX and JASON missions, the magnitudes of these empirical accelerations are on the order of 10^{-8} m/s² [106], attributed mainly to radiation pressure mismodeling, while for the LAGEOS satellites the empirical accelerations are on the order of 10^{-9} m/s², with approximately 10^{-11} m/s² sensitivity to the particular ocean tide model used in the integration [83].

This situation makes it difficult to replicate the sort of test made to validate the heliocentric


force model. In the heliocentric case, the form of the force model is simply gravitational attractions with post-Newtonian corrections. The distances to the gravitational perturbers are large enough to treat as point masses and the area-to-mass ratios are nearly always small enough to neglect non-gravitational forces. This makes for a model intended to be used for predictive purposes and it is a valid operation to seed the integrator with a reference state and compare it with the JPL reference.

In contrast, the geocentric force model, while defined to the 10^{-15} m/s² level, is expected to be accurate to no more than 10^{-9} m/s², with the empirical accelerations, which vary from day-to-day, making up the difference. In order to validate the geocentric force model in OrbitTool against the well-defined reference orbits that are grounded in satellite laser ranging (SLR) and GPS measurements, it is necessary to fit an orbit over a comparable arc using both the defined force model as well as empirical terms. This also implicitly validates the fitter framework for Cartesian state datatypes.


It is also instructive to execute the fit while excluding empirical accelerations to gauge the realistic prediction accuracy that can be expected for a geocentric precision orbit determination scenario. The sources of the empirical acceleration are both mismodeled radiation pressure effects (albedo, Earthshine, as well as non-isotropic emission/Yarkovsky effect, antenna thrust) and errors in the dynamic component of gravity. As such, one would expect that the magnitude of the error term would remain constant for at-the-satellite radiation pressure errors such as Yarkovsky effect and antenna thrust and diminish at roughly $1/R^2$ for Earth albedo mismodeling, $1/R^3$ for errors in the degree/order 2 gravity field, $1/R^4$ for errors in the degree/order 3, and so on.

The Cartesian residuals of this fit for four satellites are shown in Figure 3-9 and tabulated in Table 3.2. As a convenient performance metric for practical prediction purposes, I evaluate the worst-case Cartesian residual at the end of the seven-day fit span using the force model without empirical accelerations. The empirical forces for the same fit span are tabulated in Table 3.3 and agree in magnitude with typical values found in the literature.


The GPS and GLONASS objects used for the validation are both active payloads whose primary objective is to broadcast a roughly 500 Watt [107] signal straight down. The antenna thrust, like any radiation pressure, generates an acceleration of $\frac{P}{mc}$, where P is the radiated power, m is the mass of the satellite, and c is the speed of light. Satellites with masses on the order of 1000 kg will therefore experience a radial acceleration on the order of 10^{-9} m/s². This acceleration is on par with Earth albedo pressure and direct SRP, and is not necessarily reasonable to include as part of



(a) ILRS Geodesy Spheres



(b) GPS and GLONASS Satellites



(c) GPS and GLONASS Satellites Effect of Antenna Thrust

Figure 3-9: Geocentric Force Model Validation. Orbits fit to ILRS and IGS truth ephemerides for a seven-day span beginning 2016:353. All orbit fits include a 30×30 gravity model with solid body tides as per the 2010 IERS conventions [79], the FES2004 ocean tide model [104], direct solar radiation pressure and Earth albedo pressure [105]. For the ILRS geodesy spheres in (a), empirical accelerations include a constant in-track term and once-per-orbit periodic terms in-track and cross-track. For the GPS and GLONASS satellites in (b), there is also a constant radial acceleration. The inclusion of the constant radial empirical acceleration does not appear to make much of a difference. Panel (c) examines the effect of including the radial acceleration only on top of the physical gravity and radiation pressure model, resulting in only a marginally better fit.

Table 3.2: Geocentric Force Model Accuracy

Object	Semi-Major Axis [km]	3-Day Accuracy [m]
LAGEOS-1	12,200	11.7
Etalon-1	25,400	3.7
GLONASS 39155	25,500	5.5
GPS 28474	26,500	5.5

the non-empirical forces acting on the satellite. The marginal contribution of including an antenna thrust in the physical model is small compared to the performance improvement realized with the inclusion of empirical accelerations (Figure 3-9(c)). Including a radial term in the empirical forces and comparing its sign and magnitude (last column of Table 3.3) with this expected acceleration further validates the fitter framework in the face of widely-spaced magnitudes of the numerical values of Cartesian state elements and force model parameters. This validation is useful because while the GPS and GLONASS satellites used for the force model validation broadcast at only a few hundred watts, communications satellites in geostationary orbit broadcast at several kilowatts. At that point, it may make more sense to include a radial antenna thrust in the force model.

We can infer some trends from the best-fit empirical forces in Table 3.3 and estimate the worst-case magnitude of an unmodeled force for a geosynchronous satellite. An approximate doubling of the semi-major axis from the LAGEOS orbit to the Etalon / GLONASS / GPS altitude reduces the magnitude of the cross-track force by a factor of about 10. This is consistent with a $1/R^3$ reduction, implying an error in the implementation of the time-varying degree/order 2 geopotential coefficients in the OrbitTool force model. We would expect another reduction by a factor of 8 or so out at geostationary altitude. A greater reduction in the magnitude of the in-track acceleration might imply the same thing except it is not sensible to compare forces on active satellites with passive geodesy spheres with no 3-axis stabilization and no radiators. A Yarkovsky effect may be at play on the spinning spheres and non-isotropic thermal dumping is likely the case for the stabilized GPS/GLONASS satellites. Lastly, the radial force (not fit for in the case of the geodesy spheres) is consistent with several hundred watts of RF power from a satellite on the order of 1000 kg. Overall, it is probably safe to estimate that the dominant unmodeled forces at geostationary altitude will be on the order of 10^{-10} m/s² in-track and cross-track and no higher than 10^{-8} m/s² in the radial direction for a multi-kilowatt antenna.

Lastly, it is necessary to evaluate the orbit fitting framework for astrometric angle datatypes, which is the primary focus of this thesis. This is done with a two-stage process. First, the code for generating optical observations is validated against the Observer Table facility of the JPL Horizons

Table 3.3: Geocentric Empirical Force Magnitudes

Object	Semi-Major Axis [km]	In-Track [m/s ²]	Cross-Track [m/s ²]	Radial [m/s ²]
LAGEOS-1	12,200	6.5×10^{-9}	3.6×10^{-8}	
Etalon-1	25,400	4.9×10^{-11}	3.7×10^{-9}	
GLONASS 39155	25,500	2.7×10^{-10}	3.7×10^{-9}	1.8×10^{-9}
GPS 28474	26,500	5.3×10^{-10}	3.2×10^{-9}	1.4×10^{-9}

system. This validates the computation of angles in inertial space and the light travel time, diurnal and annual aberrations, and the transformation from inertial angles into station-aligned azimuth and elevation. Second, observations of synthetic geocentric and heliocentric objects are generated and the validation is a “Does in equal out?” test of recovering the original synthetic state from those synthetic observations, both noiseless and noised-up with zero-mean two-dimensional Gaussian errors with $\sigma = 0.2$ arcsec. The short answer is that everything agrees to the precision of the Horizons output.

3.4 Astrometric Algorithms

While a large portion of the original contributions of this thesis involves analysis of existing observations of SSA targets and simulation of hypothetical scenarios, the thesis also involves targeted data collections against asteroids and satellites. The OrbitTool software package incorporates image processing routines to extract astrometric observations from raw imagery collected during these measurement campaigns. This section describes the most salient features of the algorithms used for this purpose.

3.4.1 Image Segmentation

Image segmentation means separating the pixels in a telescope image into two categories: detections and empty sky. The algorithms described here are based on a mix of what I have seen described in open literature, what I have seen in open-source software like SExtractor [108], and what I have seen in closed-source software packages at MIT Lincoln Laboratory, and what I have found to work on ground-based SSA images from MIT LL sensors operating with short (0.1 – 1.0 sec) exposures.

3.4.1.1 Pixel Categorization

The image is first convolved with a high-pass filter to remove the effect of nonuniform illumination or background. The cameras used in this thesis are all thermoelectrically cooled, take short exposures, and have good QE uniformity, meaning that darks and flats aren’t strictly necessary, and a high-pass

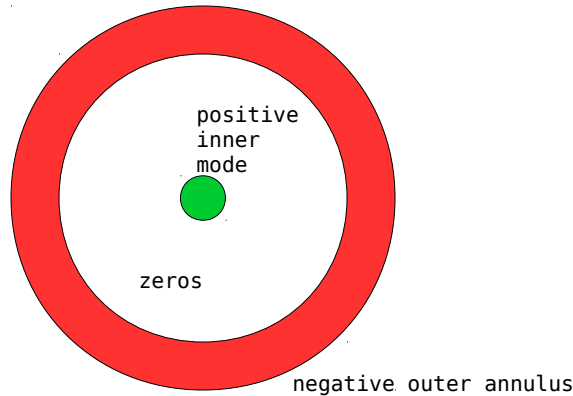


Figure 3-10: High-Pass Kernel. The negative outer annulus is meant to estimate the local background for the center pixel, and its sum is $-1/N$ for the N pixels that it covers. The spacing between the inner mode and the outer mode is meant to match the size of the typical PSF and the approximate spacing between stars so that the local background is meaningful for brighter stars and in relatively star-dense regions. Very star-dense regions that one might find in long exposure astronomical images are extremely uncommon in SSA images collected at exposures under 1 second with the 1-meter or smaller telescopes used to collect the data for this thesis.

is sufficient. The high-pass kernel is shown schematically in Figure 3-10. Its purpose is to estimate the local background level around each pixel and subtract that from the center.

For most data processing, the outer annulus runs from 6 to 12 pixels away from the center and radius of the central mode is two pixels. It is sometimes beneficial to increase the size of the central mode to match that of the nominal PSF for improved detection sensitivity on stars. This was only necessary in the processing of the data from MIT Wallace.

The background noise level is estimated from the background pixels and pixels 3σ above that level are classified as detections. This proceeds iteratively for four cycles. In the initial step, the background noise level is estimated by sorting the high-passed pixels and computing the standard deviation of all pixels between the 2.5th and 97.5th percentile. In subsequent steps, the background noise level is computed from the standard deviation only the 2.5th to 97.5th percentile of remaining background pixels. After the initial step, pixels are only flagged as belonging to detections if they are both 3σ above the background noise and belong to continuous groups of four or more eight-connected pixels³.

³Four-connected pixels share an edge, eight-connected pixels share an edge or a corner.

3.4.1.2 Pixel Aggregation

Both during the iterative flagging and at the end of the processing, the pixels flagged as detections are aggregated into contiguous eight-connected groups. During the iterative pixel flagging, this is used as for false alarm suppression. At the end of the processing, the groups' centroids and total enclosed signal are either used directly for astrometric reduction or passed down the processing chain for further morphological refinement.

The pixel aggregating algorithm described here probably goes back in some form to the first efforts at computer image processing. I learned it informally as an undergraduate and have seen equivalent implementations in every piece of astronomical image processing software where I could examine the source code.

The first step in the algorithm scans along each image row. At each transition from a pixel marked as background to a pixel marked as detection, a new data structure called a *run* is instantiated and is maintained until a transition from a detection pixel back to a background pixel. At a minimum, the run contains its left-most and right-most boundaries, but it is also useful to maintain other statistics such as total enclosed signal, and a (possibly empty) pointer to a parent run. The row in which the run exists combined with its left-most boundary column implicitly form a monotonically increasing unique identifier.

The second step involves stitching together runs in adjacent rows. This process involves incrementing along two rows at once and setting the parent pointer of each run to point to an adjacent run in the previous row if the two overlap horizontally or touch at a corner. During the stitching process, if the parent pointer of an adjacent run in a higher row is empty, then it is effectively its own parent and the parent pointer of the new run is set to it. If, however, the parent pointer of one or both rows is already set, then the parent links of one or both runs are traversed to the root to find the element with the lowest unique identifier. That element with the lowest unique identifier is assigned to be the parent of runs' parents.

The parent pointers are set only for elements with empty parent pointers to ensure that no runs get lost from the growing aggregate. Because unique identifiers are assigned monotonically when the horizontal runs are created, this is always a well-defined process, and for any shape of eight-connected aggregate, with any number of holes or branches, all members will be deterministically assigned the same unique identifier at the end of the stitching process.

This algorithm is efficient in the sense that any individual pixel is visited only once, and subse-

quent passes operate only on runs rather than on individual pixels. It is also efficient in the sense that quantities such as maximum left-right or vertical extent, total enclosed signal, and horizontal and vertical centroids may be computed incrementally as the runs are built left-to-right and as they are stitched together.

3.4.2 Point Spread Function Fitting

In all of the imagery collected on satellites for this thesis, the target moves against the star background at no more than about six pixels per exposure. This is comparable with seeing, meaning that star streaks are mostly circular rather than elongated in the images. Nevertheless for streak images, higher precision localization may sometimes be achieved by using fits of modeled PSFs rather than using image centroids [109], and the astrometric reductions used in this thesis use Gaussian streak PSFs.

The model of the Gaussian streak suggested by Koupryanov [109] is not strictly correct from a first-principles perspective. Whatever the nominal charge density PSF of a source with zero velocity relative to the focal plane is, then the charge density deposited by a linearly moving PSF is the integral over time of the charge density along the focal plane, and the pixel signal is the integral of that density over the pixel grid.

For a Gaussian PSF, the charge density on the focal plane for a source with strength ϕ_0 photo-electrons per second in the telescope aperture is given by:

$$I(x, y) = \phi_0 \frac{1}{2\pi\sigma^2} e^{-\frac{x^2}{2\sigma^2} - \frac{y^2}{2\sigma^2}} \quad (3.30)$$

where x and y are spatial coordinate offsets from source center axis in units of pixels and σ is the width of a symmetric point-spread function.

In reference to Figure 3-11, for a source moving in the x direction at a velocity v , the instantaneous amount of charge deposited in the focal plane is given by:

$$\Phi(x, y, t) = \frac{\phi_0}{2\pi\sigma^2} e^{-\frac{1}{2} \frac{(x-vt)^2}{\sigma^2} - \frac{1}{2} \frac{y^2}{\sigma^2}} \quad (3.31)$$

Let $\rho(x, y)$ be the charge density on the detector at the end of the exposure. This is the integral




Figure 3-11: Schematic of Calculation of Charge Density and Pixel Signal for Streaking PSF

over time of the instantaneous rate of charge deposition onto the focal plane, and is given by

$$\rho(x, y) = \int_{t=0}^T \Phi(x, y, t) dt = \frac{\phi_0}{2\pi\sigma^2} \left(e^{-\frac{1}{2}\frac{y^2}{\sigma^2}} \right) \int_{t=0}^T e^{-\frac{1}{2}\frac{(x-vt)^2}{\sigma^2}} dt \quad (3.32)$$

Making the substitution $\tau = (vt - x)/(\sigma\sqrt{2})$,

$$\rho(x, y) = \int_{t=0}^T \Phi(x, y, t) dt = \frac{\phi_0}{2\pi\sigma^2} \left(e^{-\frac{1}{2}\frac{y^2}{\sigma^2}} \right) \left(\frac{\sigma\sqrt{2}}{v} \right) \frac{\sqrt{\pi}}{2} \frac{2}{\sqrt{\pi}} \int_{\tau=-x/(\sigma\sqrt{2})}^{(vT-x)/(\sigma\sqrt{2})} e^{-\tau^2} d\tau \quad (3.33)$$

In this form, the integral becomes the error function, and using the anti-symmetry of the error function, the pixel charge density function becomes [110]

$$\rho(x, y) = \frac{\phi_0}{2\sigma v\sqrt{2\pi}} \left(e^{-\frac{1}{2}\frac{y^2}{\sigma^2}} \right) \left(\operatorname{erf} \left(\frac{x}{\sigma\sqrt{2}} \right) - \operatorname{erf} \left(\frac{x - vT}{\sigma\sqrt{2}} \right) \right) \quad (3.34)$$

It is convenient to recenter the coordinate system to the middle of the streak

$$\rho(x, y) = \frac{\phi_0}{2\sigma v\sqrt{2\pi}} \left(e^{-\frac{1}{2}\frac{y^2}{\sigma^2}} \right) \left(\operatorname{erf} \left(\frac{x + vT/2}{\sigma\sqrt{2}} \right) - \operatorname{erf} \left(\frac{x - vT/2}{\sigma\sqrt{2}} \right) \right) \quad (3.35)$$

Unless the streak direction is aligned with the pixel grid, Equation 3.35 has no closed form integral over arbitrarily oriented rectangles, and it is necessary to compute the integrated pixel signal from this charge density numerically. A 3×3 grid within each pixel provides an adequate solution given the optical quality of the data collected for this thesis.

To simplify processing, the same PSF is used to fit both streaking stars and rate-matched targets.

3.4.3 Star Matching

The star matching algorithm used in the code is based on the Valdes triangle matching algorithm [111], with modifications suggested by Heyl [112]. It is assumed that the nominal position of the telescope and its approximate plate scale is known. The parameters described in this section are tuned to work well with fields of view under half a degree in size.

All possible triangles of the 48 brightest stars and focal plane detections are formed and represented as the ratio of their two smaller sides to their largest side. Near-equilateral triangles and triangles whose largest side is larger than 500 arcsec or smaller than 10 arcsec are excluded. The remaining normalized triangles of bright stars are placed into a two-dimensional kd-Tree data structure.

Each of the remaining normalized triangles of bright focal plane detections are used to search the kd-Tree of normalized star triangles for candidate matches to a relative tolerance of 2×10^{-3} in the two-dimensional normalized triangle coordinates, with matches whose pixel scale differs by more than 10% rejected.

Following a simplified version of Heyl’s idea, each of the remaining star-to-detection triangle match candidates are used to form a rigid body transformation from focal plane pixel coordinates to sky-fixed northward/eastward coordinates in a plane tangent to the celestial sphere at the nominal center of the field of view. Given detection pixel coordinates r, c and star coordinates N, E , the form of the rigid body transformation is:

$$r = A_1 N + A_2 E + A_3 \tag{3.36}$$

$$c = -A_2 N + A_1 E + A_4 \tag{3.37}$$

A pair of triangles provides six independent measurements to solve for four unknowns in the least-squares sense. Heyl’s algorithm allows for shear, non-square pixels, and unknown orientation of focal plane raster scanning, but for this application where the parameters of all contributing sensors are known, a four-parameter rigid-body transformation is sufficient.

The vectors of the best-fit rigid body transformations for each candidate match are placed into a four-dimensional kd-Tree for the purpose of finding clusters of similar rigid-body transformations that correspond to the true match between focal plane and celestial coordinates. To regularize the distance metric between the pixel scale/rotation coordinates (A_1, A_2) and origin coordinates (A_3, A_4) , the scale/rotation coordinates are multiplied by $\frac{512 \text{ [pixels]}}{\sqrt{A_1^2 + A_2^2}}$ before placement in the kd-

Tree. This has the effect of placing two pixel grid points into the kd-Tree: one at the nominal origin and one at the right edge of a 1024×1024 focal plane. Thus, a distance metric that compares points to each other means the same thing.

Each point in the four-dimensional kd-Tree of point pairs generated by the rigid body transformation derived from candidate triangle matches is then assessed for popularity by querying the number of its neighbors within a 0.9 unit radius of its four-dimensional coordinate⁴. The rigid body transformation with the largest number of neighbors is taken as the center of the biggest cluster of rigid body transformations, corresponding to the true best-fit rigid body transformation between focal plane pixel coordinates and celestial coordinates. This is the improvement Heyl proposed to Valdes’s popularity array.

The constituent triangles of the rigid body transformations belonging to this largest cluster are then used to form star-to-detection matches. Any duplicates within this match list are removed and the resulting list is used to form a candidate polynomial plate model solution. The order of the polynomial is linear in each coordinate if the number of recovered matches is under ten and the order of the polynomial is cubic in each coordinate if the number of recovered matches is greater than ten. These candidate plate models are then iteratively used to find any remaining matches of stars and sources that were not bright enough to make the cut for the initial triangle match using the 48 brightest sources or belonged to triangles that were too large, too small, or too close to equilateral. This search is done using nearest-neighbor search to a tolerance of 2 arcseconds using the nominal plate model computed from the triangle match as a seed.

The iterative process of finding nearest neighbors and rejecting outliers is based on the assumed distribution of the post-fit plate residuals. It is assumed that the distribution of the magnitudes of post-fit residuals of the plate solution obeys a Chi distribution with two degrees of freedom. That is, the probability that the magnitude of a residual $\sqrt{\Delta x^2 + \Delta y^2}$ is less than r is given by:

$$F(r) = P\left(\sqrt{\Delta x^2 + \Delta y^2} \leq r\right) = 1 - \exp\left(-\frac{r^2}{2\sigma^2}\right) \quad (3.38)$$

which implies that the underlying Gaussian distribution’s σ may be estimated from the cumulative distribution:

$$\sigma^2 = -0.5 \frac{r^2}{\log(1 - F(r))} \quad (3.39)$$

⁴In a four-dimensional manifold defined by the pixel coordinates of two points, a four-dimensional search tolerance of 0.9 approximately means a threshold of a little under $0.9/\sqrt{2}$ pixels away from each point in the match.

The cumulative distribution $F(r)$ is approximated by the sorted list of post-fit residuals. That is, for a sorted list of post-fit residuals r_i of size N ,

$$F(r_i) \approx i/N \tag{3.40}$$

The value of σ from the Chi distribution is estimated from the interior of the sorted distribution. It is computed as the arithmetic average of all σ 's computed via Equation 3.39 between the 25th and 75th percentiles of the post-fit residuals. In order to flag outliers, this estimate of σ is compared with the sample σ derived from the post-fit residuals⁵. The comparison is made for all sorted residuals, in increasing order. The first residual that increases the sample σ beyond the Chi distribution σ derived from the interior of the distribution is flagged as an outlier, as are all larger residuals. This process is repeated until too many outliers are rejected, in which case the match is declared a failure, or until no more outliers are flagged. It is implicitly assumed that no more than 25% of the post-fit residuals belong to outliers in the beginning.

This iterative flagging heuristic is more rigorous than the sigma clipping criterion proposed by Valdes as it is rigorously grounded in the (mostly true) assumption that the inlier residuals are drawn from a two-dimensional Gaussian error distribution dominated by seeing noise while the outlier residuals are drawn from a different distribution.

⁵For a Chi distribution with two degrees of freedom, the standard deviation of the one-dimensional magnitudes of the residuals is related to the underlying σ by a factor of $\sqrt{2 - \frac{\pi}{2}}$.

Chapter 4

Predictions, Analysis, and Preliminary Measurements

4.1 DCR Bias of Stars

Not all fiducial stars are Sun-like and those that are Sun-like are sufficiently far away to be subject to non-negligible interstellar reddening. In the consideration of DCR compensation for solar-illuminated Earth satellites and Solar System bodies, it is necessary to analyze the degree to which stellar DCR contributes to the overall error in astrometric solutions derived from bare-silicon detectors. Stone [71] calculated a systematic bias on the order of several tens of milliarcseconds for early-1980s era tube detectors used with narrow passband filters and parametrized the bias by star spectral type. It is necessary to reproduce this calculation for modern CCD detectors with wider passbands and to parametrize the result by the color indices available in astrometric catalogs.

The approach that worked best was through analysis of synthetic photometry with a spot-check comparison against measurements where suitable. Using standard relations between visible band extinction $A(V)$ and per-wavelength extinction [113] and a representative collection of stellar spectra such as the Pickles Atlas [114], one may simultaneously vary spectral type and interstellar reddening for all possible stars, compute synthetic photometry in relevant passbands for any particular approach, and compare the DCR for the actual spectrum of the star with the DCR inferred from the synthetic photometry.

Prior to converging on this technique, several approaches were investigated during this dissertation research toward the objective of estimating stellar DCR from information available in all-sky

catalogs. These were:

1. Assume that stars are well-approximated by blackbody spectra and compute the DCR for an unreddened blackbody spectrum corresponding to each star’s color index;
2. Compute a best-fit blackbody temperature as a function of available color indices where the performance metric is DCR accuracy, rather than a physical match to effective temperature;
3. Computing DCR bias directly as a polynomial function of available photometric color indices (and stellar parallax) without a direct relationship to any stellar physics.

4.1.1 Blackbody Approximation Method

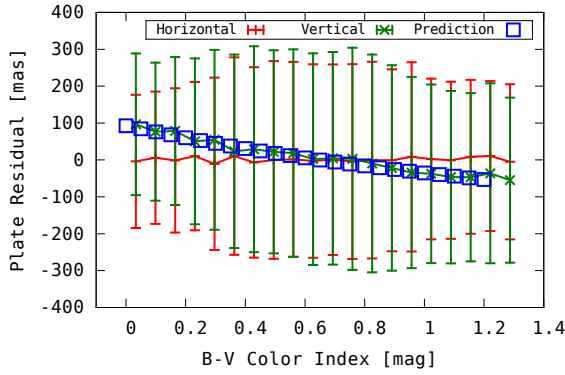
4.1.1.1 Simplified Approach

Stone’s analysis of stellar DCR parametrized stars by spectral type, and the most natural numerical quantification of spectral type is effective temperature. Under idealized conditions (and neglecting interstellar reddening), temperature may be inferred from the multicolor photometry observables compiled in all-sky star catalogs. A single two-color photometric measurement such as a $B - V$ or $B - R$ color index may then be used to estimate the blackbody temperature for each catalog star. A relation of the form $T \approx 10^{c_0 + c_1(B-V) + c_2(B-V)^2 + c_3(B-V)^3 \dots + g_1 \log_{10} g + h_1[\text{FeH}]}$ [115] can be used to map catalog color indices to blackbody temperatures and that spectrum may be used to compute DCR bias for any particular star. The metallicity [FeH] and log surface gravity $\log_{10} g$ are not available in star catalogs as they must be inferred from spectral measurements and stellar physics models, and nominal values for Main Sequence stars are used in this analysis¹.


This avenue of approach was attractive not only because of its relative simplicity but also because of its direct testability with available data. The UCAC4 catalog used for data reductions for data collected at the Firepond Facility includes Johnson B and R bands for nearly all of its stars. $B - R$ may be inferred from $B - V$ under the assumption that stars burn as blackbodies, thus making it possible to directly plot astrometric plate solution residual versus catalog color index of each star in aggregate.

The localization error for each individual star for a typical exposure used at Firepond (120–200 msec) has a $1-\sigma$ bound of about 0.3 arcsec, so it is only possible to compute this validation in the mean-value sense over many hundreds of individual astrometric solutions. When this aggregation

¹From the perspective of DCR accuracy, the temperature relation is weakly sensitive to deviations of these parameters from nominal values for a typical Main Sequence star.



(a) Color Slope Point Measurement



(b) Color Slope vs. Elevation


Figure 4-1: Empirical Observation of DCR of Stars at Firepond. (a) Observations collected at about 60 degrees elevation and overlaid with predictions. (b) Variability of color slope vs. elevation. This empirical observation seemingly validates the method of predicting DCR bias for stars, but it is not sufficient to indicate whether individual stars' DCR bias is skewed by the approximation of neglecting the effect of interstellar reddening.

is computed, it is possible to see good agreement between the DCR-versus-color slope predicted from the blackbody-only approximation to stellar spectra and the slope of the averaged aggregate empirical observations (Figure 4-1(a)). The agreement between predicted and observed slope over a wide range of elevations (Figure 4-1(b)) would seem to indicate that the theoretical framework is sound.


4.1.1.2 Choosing The Blackbody Temperature

The validity of the blackbody approximation to actual stellar spectra in terms of the performance metric of DCR prediction accuracy may be further evaluated by comparing the DCR computed from a blackbody spectrum approximation versus the DCR computed from physically-plausible stellar spectra with the same color indices over a variety of elevation angles and atmospheric conditions. Figure 4-2 makes this comparison for all 131 stellar spectra in the Pickles Atlas [114].


When the stars are near Sun-like, the blackbody approximation is valid to under 50 mas over a wide range of elevation angles, but at low elevations and for bluer stars, the validity of the approximation breaks down. The approximation also breaks down for red giants, whose spectra are sufficiently different from that of a blackbody to yield a high error when computing DCR. An additional complication comes from the fact that effective temperature must be inferred from observables such as $B - V$. Using a standard color-temperature relation and only $B - V$ yields a worse approximation for DCR, as seen in panels (c) and (d) of Figure 4-2.



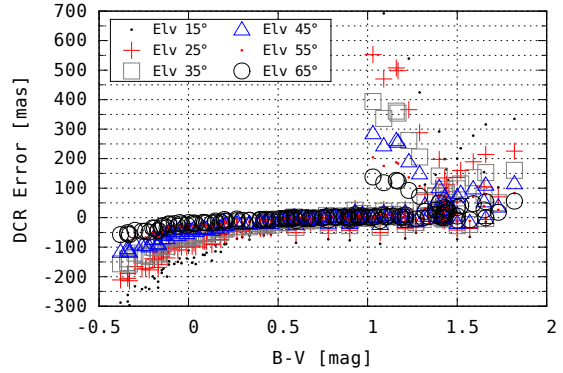
(a) BB Temp. = Canonical T_{eff} , Westford, MA



(b) BB Temp. = Canonical T_{eff} , Socorro, NM



(c) BB Temp. from $B - V$ Color, Westford, MA



(d) BB Temp. from $B - V$ Color, Socorro, NM

Figure 4-2: Predicted DCR Error from Blackbody Spectrum Approximations at Westford, MA and Socorro, NM. In panels (a) and (b), the spectra of the 131 stars in the Pickles Atlas [114] are used to compute baseline truth DCR. A blackbody approximation to the same spectra is computed using the canonical effective temperature for the spectrum listed in the Pickles Atlas. In panels (c) and (d), the blackbody temperature is computed from the canonical $B - V$ color index of the Pickles stars using a color-temperature relation [115] and nominal assumptions for a typical main sequence star ($g = 10^{4.3}$, $[\text{FeH}] = 0$). Using the approximation based on canonical temperature yields a better result in terms of DCR for the majority of stars than does basing the blackbody temperature on the $B - V$ color index, but the structure in panels (c) and (d) implies that a refined estimate of effective temperature may be computed to better match the DCR of the original stellar spectra.

The errors in panels (c) and (d) of Figure 4-2 are larger in magnitude than the errors in panels (a) and (b), but appear to have structure. This implies that it should be possible to compute corrections beyond the standard color-temperature relation that improve accuracy in computing stellar DCR. Intuition for the existence of this correction comes from two observations. The first observation is that the red giant branch, which has worse DCR prediction performance, may be identified in the color-color diagram of the Pickles Atlas (Figure 4-3) and excluded from an astrometric solution to avoid absorbing that error into the solution. The second observation is that for every Pickles star, observed at any elevation, there exists a blackbody spectrum that gives the exact DCR as the original spectrum, and furthermore, a blackbody spectrum with a temperature computed as the logarithmic average over elevation of those optimal blackbody temperatures gives a very good approximation to the DCR of that star observed over all elevations (Figure 4-4).

The straightforward approach to computing this optimal temperature is to use a polynomial function of observable color indices $B - V$ and $J - K$. Note that the optimal mean temperature depends on atmospheric conditions, and this approach is a departure from using only atmosphere- and detector-agnostic physics to predict stellar spectra and a step toward having to recompute atmosphere- and site-dependent quantities to correct for stellar DCR.

Panels (a) and (b) in Figure 4-5 show the resulting DCR error for all Pickles stars over elevation when a cubic polynomial is used to estimate the optimal temperature for each star. A cubic polynomial is insufficient to achieve correction to below 50 mas, but appears to capture elevation-dependence of DCR very well. The structure in the residual as a function of observable color implies that the approach is valid. Using the color-color diagram as a guide, we can exclude the red giant branch from the fit by restricting the fit to the stars that lie along the linear trunk of the color-color diagram (Panels (c) and (d) of Figure 4-5). This improves performance in the red, but not in the blue end of possible stars. Noting that most stars are redder than Sun-like, we can also exclude the very hot blue stars from consideration and compute a blackbody temperature that produces an approximation to DCR for most of the stars we would expect to see that is accurate to about 20 mas even down to 15 degrees elevation (Panels (e) and (f) of Figure 4-5).




Figure 4-3: Color-Color Diagram of Pickles Atlas Stars. The red giant branch is evident in this plot. DCR prediction performance for red giants based on blackbody approximations derived from multicolor photometry is worse. Using the color-color diagram, red giants can be identified and excluded from an astrometric solution to avoid the DCR prediction difficulties for that type of star.

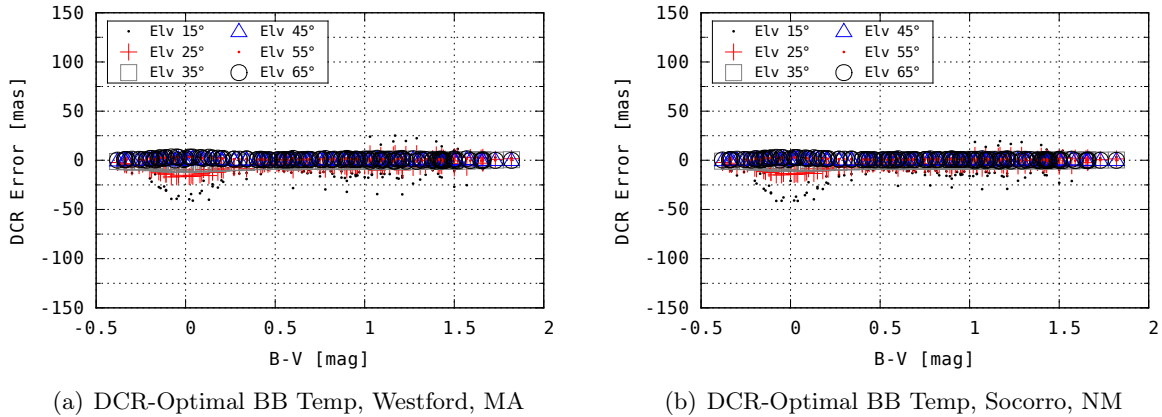


Figure 4-4: Residual Error of DCR for Pickles Stars Approximated with DCR-Optimal Blackbody Temperature. For each star at each elevation, there exists a blackbody spectrum that gives identical DCR. Logarithmically averaging this optimal temperature for each star over elevation yields a blackbody spectrum that approximates each star well over the full range of elevations. This mean blackbody temperature optimized for DCR predictions is not strictly a physical quantity in that it may bear little relation to the actual temperature of the star. This temperature is not directly accessible using the observables tabulated in star catalogs, but the utility of this non-physical quantity in generating a spectrum that gives a good DCR approximation over a large range of elevation angles led to the approach of empirically fitting DCR as a polynomial function of observable photometry without regard to grounding the DCR prediction in plausible stellar spectra (Section 4.1.3).

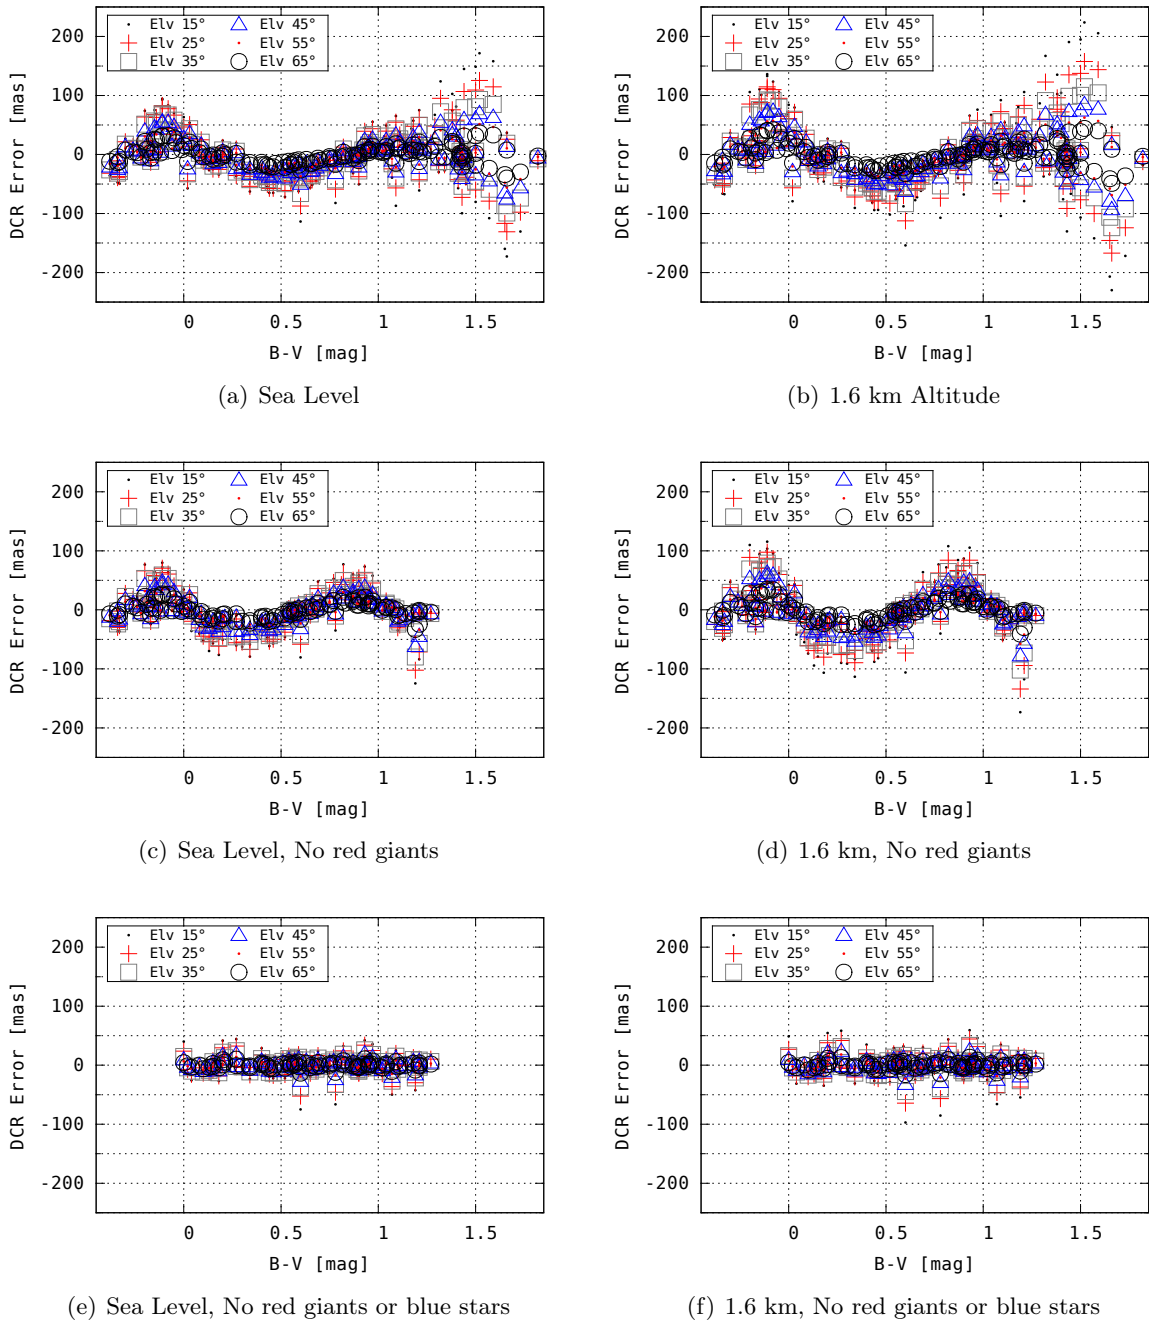


Figure 4-5: DCR Error of Blackbody Spectrum Approximations Derived from Multicolor Photometry. The DCR prediction performance of a blackbody approximation with the blackbody temperature computed as a constant order polynomial in observable color indices $B - V$ and $J - K$ is compared against DCR computed directly from Pickles Atlas spectra. Panels (a) and (b): the worst-case performance for all possible Pickles stars is as much as 100 mas. Panels (c) and (d): Excluding the red giant branch from consideration improves worst-case performance by about 30%. Panels (e) and (f): Excluding very blue stars which are not typically found in astrometric catalogs improves worst-case performance to at worst a few tens of milliarcseconds at low elevation angles.

4.1.2 The Effect of Interstellar Reddening


At this point, there appears to be a contradiction between the seeming validation of blackbody approximations derived solely from color temperature relations in Section 4.1.1.1 and Figure 4-1 and the necessary corrections to the color-temperature relation developed in Section 4.1.1.2. The resolution to this contradiction comes from considering the effect interstellar reddening.

This realization took several iterations of analysis to reach over the course of this thesis. As a first-cut to address the question of whether interstellar reddening contributes a significant error to the estimate of DCR bias, the DCR bias of an unreddened solar-analog spectrum having the same observable color indices as a reddened spectrum (i.e. lower blackbody temperature) was compared with the DCR bias of the reddened spectrum.


The extinction at any wavelength $A(\lambda)$ is characterized relative to the overall extinction in the V band $A(V)$, while the excess in color indices is defined for given values of the quantity $R_V = \frac{A(V)}{E(B-V)}$. In most directions, $R_V \approx 3.1$ [116] meaning that if the extinction in the V band is known, then the bias in the color index can be estimated. The estimation of absolute extinction in the visible band is also necessary to recover the absolute extinction curve, which is measured and tabulated as $A(\lambda)/A(V)$ in the literature. While there are approaches to estimating the absolute extinction from multicolor visible and infrared photometry [117], an approximation of about 1.8 magnitudes / kpc [118] is used in this analysis.

A typical Sun-like star has a V magnitude of about 4.7 at 10 pc. In the absence of extinction, a Sun-like star at the $V \approx 17$ magnitude cutoff of the URAT1 and UCAC5 catalogs is about 12 magnitudes fainter meaning it must lie at a distance of no more than $\sqrt{10^{(0.4 \times 12)}} \times 10 \text{ pc} \approx 2.5 \text{ kpc}$. At that distance it would have another 4.5 magnitudes of extinction, meaning its apparent magnitude would be too faint for the cutoff. From this reasoning, we can estimate the maximal distance of a Sun-like star as about 1 kpc, which would yield an apparent magnitude of about 14.7 with no extinction and 16.5 with extinction.


At an absolute extinction of 1.8 magnitudes, the expected the color excess would be about $1.8/R_v = 0.58$ to the red. In order to determine whether catalog $B - V$ is an adequate substitute for inferring absolute extinction for each star, it is necessary to compare the DCR bias for a Sun-like star ($B - V = 0.625$, 5700 K blackbody spectrum) with 1 kpc worth of extinction applied to the DCR bias of an unreddened spectrum for a $B - V = 1.225$ (4440 K blackbody) star. The results of this calculation are shown in Figure 4-6.




(a) Reddening vs. Wavelength and Star Distance



(b) Absolute DCR of 1 kpc Reddening and Equivalent $B - V$ Shift, Socorro, NM



(c) DCR Difference of Reddening and Equivalent $B - V$ Shift, Socorro, NM



(d) DCR Difference of Reddening and Equivalent $B - V$ Shift, Westford, MA

Figure 4-6: Biases from Neglecting Interstellar Extinction. Comparison of DCR bias of a Sun-like star ($B - V = 0.626$) subject to interstellar extinction at distances up to 1 kpc (“true DCR”) and a corresponding blackbody spectrum with $B - V = 0.626 + E(B - V) = 0.626 + A_V/3.1$ assuming 1.8 mag/kpc extinction (“approximate DCR”). (a) Extinction in assuming $R_V = 3.1$ at three distances. (b) Absolute DCR for Sun-like star subject to that extinction and for a blackbody with temperature computed using $B - V = 0.626 + E(B - V) = 0.626 + A_V/3.1$. Overall, the true DCR is more downward than the approximate DCR. This means that a DCR correction using the approximate method will result in a corrected focal plane measurement of the star’s location that is too low. Measurements of satellites and asteroids made against this star will thus appear to be too high and have a positive elevation bias (panels (c) and (d)). Calculations assume back-illuminated silicon detectors, and nominal meteorological conditions.

The results in Figure 4-6 indicate that in the worst-case, interstellar reddening can introduce a significant amount of bias. The degree to which the underlying physics of this bias causes the resulting DCR calculation to deviate from the results of the simplified approach was at first an open question given the agreement between the empirically-observed vs. predicted DCR-versus-color slope shown in Section 4.1.1.1. There is a tension between the fact that the simplified physics neglects the particulars of the interplay and ambiguity between reddening and spectral type, the fact that reddening makes stellar spectra more similar to each other, and the fact that most catalog stars are not at the worst-case extreme.

In the first cut of this work, the agreement between the observed and predicted stellar DCR vs. color slope was taken as evidence only of the fact that the underlying assumptions about instrument spectral quantum efficiency, atmospheric transmittance, and atmospheric refraction were valid. It was tacitly assumed that an average stellar DCR could be removed from archived MPC observations of asteroids with known spectral types. That is, the assumption was that the DCR vs. elevation curves for S-type and C-type asteroids presented in Section 1.2.2, which will be shown in Section 4.2.1 to not match DCR predictions that do not account for fiducial stars' DCR bias, could be made to match by an after-the-fact addition of a suitable stellar DCR correction. This turned out to be a blind alley, and the fundamental takeaway from that investigation is that while interstellar reddening does indeed tend to make stellar spectra look more similar to each other, to the point where, especially when observed with narrow passbands (as Stone found), the difference between stellar DCR for an O and M type star is under a few tens of milliarcseconds, this is not true when considering the difference between the DCR of a star and a solar-illuminated target not subject to interstellar reddening.

The more rigorous approach toward examining the efficacy of the blackbody approximation in the presence of interstellar reddening is to compare the DCR of a spectrum synthesized with a realistic stellar spectrum as its starting point and subjected to the effect of interstellar reddening.

Rather than making this computation for all 131 possible Pickles spectra at all physically-plausible distances within the Milky Way galaxy, it is sensible to restrict the analysis to stars that match the color and brightness distributions of real all-sky fiducial catalogs. Figure 4-7 shows the brightness and $B - V$ statistics of the URAT1 catalog. Though this is not an all-sky catalog, it is complete for the northern hemisphere, and being an entirely ground-based catalog, is representative of the statistics of fiducial stars detectable from a ground-based instrument. Its statistics indicate that we can capture all the necessary variation in the combined effect of spectral type and reddening

by considering for analysis synthetic spectra whose $B - V$ ranges between 0 and 2 and whose faint cutoff is about 18th magnitude. For the bright cutoff we choose 10th magnitude, to reflect the typical saturation point of the Firepond 48" telescope in its typical SSA operating mode. This corresponds to a minimum distance of a few tens of parsecs for some of the fainter star types in the Pickles catalog and a maximum distance of 4 kpc for the brighter stars.

Using these cutoffs, we can compute the synthetic spectra for all stars, their associated multicolor photometry, and their true DCR with respect to the unreddened solar spectrum, at any arbitrary zenith distance. This can then be used as a baseline against which the DCR computed from the blackbody approximation derived from $B - V$, or any other photometric observables, may be evaluated. This is what is shown in Figure 4-8, where it is evident that while the DCR-versus-color slope of the blackbody-derived DCR may indeed match the aggregate slope of the synthetic photometry-derived DCR baseline, there is in fact a large (> 100 mas) bias between the two methods' results. Further, this bias is largest in the region of $B - V$ where most of the stars in the fiducial catalogs reside.

This analysis reveals a key pitfall. On the one hand, interstellar reddening makes stars' spectra look more similar to each other to the point that in narrow passbands, the difference in DCR between extremes of the stellar spectra is small, as Stone found. On the other hand, when comparing stellar spectra against solar-illuminated spectra, the overall trend is the same but indexing the spectral types by a single number like $B - V$ color, or effective temperature, misses the common-mode offset introduced by the *lack* of reddening in the baseline unreddened solar spectrum against which the comparison in Figure 4-8 is made. It is the narrowness of the B and V passbands with respect to the full silicon passband that allows the ambiguity between reddening and spectral type to come through.

To belabor the point further, a blackbody spectrum is a fairly good approximation for the unreddened solar spectrum over the silicon passband, and most stars' $B - V$ is redder than the Sun. There are going to be relatively few blue stars encountered in the dataset used to generate Figure 4-1, as they will be outnumbered by redder stars. Thus, for the majority of stars, a $B - V$ to unreddened blackbody DCR calculation will simply slide down the curves in Figure 4-8, but without a true solar-analog to compare in the astrometry, it will not be evident that the DCR calculation does not slide far enough down the curve, and it will not be evident that the blue stars' spectra do not look bluer than Sun-like when evaluated in the context of DCR.




Figure 4-7: URAT1 Brightness and Color Statistics




Figure 4-8: DCR of Reddened Stellar Spectra vs. DCR of Blackbody Spectra Inferred from $B - V$ with Bare Silicon Detector at Sea Level. All DCR is computed with respect to an unreddened solar spectrum. Pickles Atlas spectra are subjected to interstellar reddening between 0.01 and 4.0 kpc. Synthetic photometry in Johnson B and V bands is computed from those reddened spectra and a color-temperature relation [115] is used to synthesize an (unreddened) blackbody spectrum. The DCR for the blackbody spectrum and all DCRs for true spectra are plotted assuming sea-level atmosphere and a standard silicon quantum efficiency. The blackbody approximation is best for near Sun-like stars at short distances, and this is indicated by the small number of points tangent to the blackbody baseline near $B - V \approx 0.6$. The narrowness of the B and V passbands with respect to the full silicon passband allows the ambiguity between reddening and spectral type to come through in the offset between the DCR computed from the blackbody approximation that ignores interstellar reddening and the DCR of the actual reddened spectra. The slope of the blackbody baseline and the slope of the locus of true DCR points vs. color index is approximately equivalent, resulting in the agreement between blackbody-derived prediction and measurement shown in Figure 4-1.

4.1.3 Multi-Color Direct Polynomial Approximation

Given that blackbody distributions can make for good approximations to stellar DCR, it is tempting to continue a search for a transformation from information tabulated in star catalogs to expected DCR bias that is grounded in heliophysics. However, the combined effects of variation in stellar types and variation in amount of interstellar extinction introduce an ambiguity into any attempt to compute this transformation. This ambiguity complicates physics-based approaches dependent on narrow passband photometry.

There exist standard techniques to approximate effective temperature as a function of color index, metallicity, and surface gravity, but absent spectra, surface gravity and metallicity cannot be deterministically inferred from the ensemble of color indices available in astrometric catalogs. These parameters must be determined by a direct measurement of spectral lines [119] or from absolute luminosity [120]. The Gaia spacecraft’s Blue Pass and Red Pass spectrometers provide the necessary spectrographic coverage to estimate the spectral type of fiducial stars, but calibrated spectra from these instruments is not incorporated into Gaia DR2 and will not be released until at least the third data release in 2021 [121]. Estimates of extinction and effective temperature that are included with Gaia DR2 are derived from parallax and two-color photometry rather than direct use of the spectra and users are specifically cautioned against relying on the accuracy of a stellar parameter estimate for any particular star [122].

Pending the release of spectra from the Gaia low resolution spectrometer instruments, and more refined estimates of stellar parameters from that data, DCR bias for stars must be inferred from multicolor photometry. Given that the transformation from cataloged quantities to stellar spectra to DCR bias is going to necessarily be missing some physics, the approach adopted in this thesis is to make a purely empirical mapping from the observables tabulated in astrometric catalogs to stellar DCR.

4.1.3.1 UCAC Color Indices

The following procedure, similar in concept to the process used to estimate stellar parameters in Gaia DR2, is used with UCAC5 to estimate the relation between catalog photometry and DCR in astrometric reduction. UCAC5 is chosen because of its completeness to around $V=17$, its accuracy by virtue of registration to Hipparcos and Gaia, and the large fraction of its stars that contain visible and near-IR photometry.

Synthetic photometry is generated for all stars in the Pickles Atlas. The synthetic photometry includes the effects of interstellar reddening assuming $R_V = 3.1$, a constant value of V-band extinction of 1.8 magnitudes/kpc, atmospheric attenuation computed via MODTRAN [123], two near-normal reflections from aluminum-coated mirrors, and silicon-band quantum efficiency typical of back-illuminated devices.

The exoatmospheric spectra of stars were generated at distances of 0.1 to 4.0 kpc. Synthetic photometry for Bessel [124] and Gaia passbands was computed to synthesize catalog photometry and the synthetic spectra were pruned with a distance and spectral type cutoff such that the Gaia G magnitude of the synthetic spectra stayed in the range (10, 18), corresponding to the useful dynamic range of a 1-meter-class telescope and such that $B - V$ remained in the range (0, 2), corresponding to the majority of stars in URAT1 and UCAC5.

Detectable photoelectron spectra for each Pickles star were computed from zenith down to an elevation of 5 degrees in 5 degree increments. Using the Auer-Standish wavelength-dependent refraction and atmosphere models [2], the DCR bias with respect to the solar spectrum is computed from these spectra at all heliocentric distances that meet the $B - V$ and G magnitude criteria. A third order polynomial function of combinations of color indices available in UCAC5 is fit with linear least squares to the computed DCR bias from synthetic photometry at each of the 5 degree elevation increments. The polynomial approximation to DCR in 2MASS J and K , Gaia G , and UCAC U_5 magnitudes is of the form:

$$\begin{aligned} \Delta z = & a_0 + a_1G + a_2(J - K) + a_3G^2 + a_4(J - K)^2 + a_5G(J - K) \\ & + a_6(G - J) + a_7(J - G)(J - K) + a_8U_5 + a_9(J - U_5) + a_{10}(J - U_5)^2 \\ & + a_{11}(J - U_5)^3 + a_{12}(J - U_5)^2(J - K) + a_{13}(J - U_5)(J - K)^2 + a_{14}(J - K)^3 \end{aligned} \quad (4.1)$$

The specific terms in this polynomial were chosen by starting with linear functions of available color indices and manually adding higher order terms by trial and error. The performance metric was the worst-case prediction error of DCR from the polynomial approximation with respect to DCR from synthetic photometry. The end result polynomial approximation is the cubic polynomial above, whose worst-case predictive performance is not distinguishable from a polynomial with more terms.

Rather than a generalized global approximation valid over the whole sky, the set of a_i coefficients is recomputed for each five-degree elevation increment, for each site and combination of optical

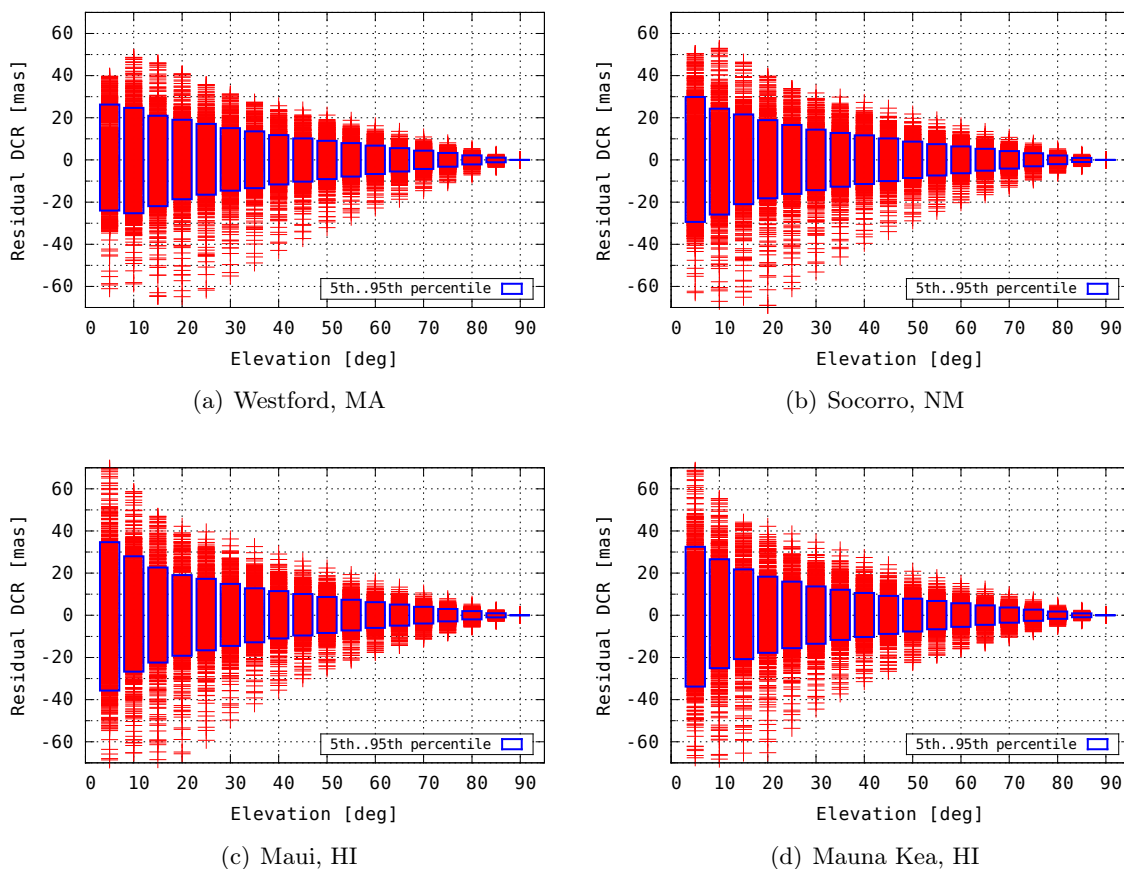


Figure 4-9: DCR Prediction Error vs. Elevation for Polynomial Function of UCAC5 Photometry. Synthetic photometry generated from Pickles Atlas stars is used to compute UCAC5 catalog photometry and DCR bias at the four sites. A cubic polynomial of computed synthetic photometry approximates the DCR bias at each elevation increment, with residual error in DCR shown in the figure. Empirically-fitted DCR approximations can reduce the stellar DCR contribution to astrometric solution error by a significant amount. The worst-case prediction error here is significantly below the natural variation in stellar DCR about its mean value (Figure 4-11).

passbands used at each site. When using this method to compensate for stellar DCR in astrometric reductions, linear interpolation of DCR values from the 5° increments to the observed elevation of each star is performed. The polynomial coefficients are fitted without weighting for the relative frequencies of different types of Pickles stars in space but are weighted in proportion to distance squared to capture the likelihood of encountering larger numbers of more distant stars subject to greater amounts of interstellar reddening. The residual errors in DCR from making this polynomial approximation with respect to the DCR computed from the synthetic photometry spectra are shown in Figure 4-9, and shows a worst-case prediction error of ± 20 mas above 15 degrees elevation for 90% of all simulated stellar spectra.

4.1.3.2 Gaia DR2 Three Color and Parallax

The same analysis is done for Gaia DR2 stars incorporating three-color Gaia photometry in the G , RP and BP bands and astrometric parallax estimate π into the empirical polynomial. The polynomial used was of the form:

$$\begin{aligned} \Delta z = & a_0 + a_1 G + a_2(BP) + a_3(RP) + a_4(BP - G)^2 + a_5 G(BP - G)^3 \\ & + a_6(RP - G)^2 + a_7(RP - G)^3 + a_8(RP - BP)^2 + a_9(RP - BP)^3 \\ & + a_{10} \frac{1}{\pi} + a_{11} \left(\frac{1}{\pi}\right)^2 + a_{12} \left(\frac{1}{\pi}\right)^3 \\ & + a_{13} \left(\frac{15}{\pi} - G\right) + a_{14} \left(\frac{15}{\pi} - G\right)^2 + a_{15} \left(\frac{15}{\pi} - G\right)^3 \end{aligned} \quad (4.2)$$

The results of this are shown in Figure 4-10 and are only marginally better than what is achievable with UCAC5 photometry. While direct measurement of parallax is useful, the reduction from four-color to three-color photometry reduces the benefit with respect to UCAC5 photometry. Subsequent Gaia releases incorporating higher fidelity calibrated spectra will aid in the per-star compensation of DCR bias.

4.1.3.3 DCR of the Average Star and the Merits of Per-Star DCR Compensation

One may also ask about the average DCR bias of all stars and its variation at a given elevation. This quantity is computed directly for every star in the UCAC5 catalog for each of the 5 degree elevation increments. The resulting mean and standard deviation for sites from sea level to 4 km altitude are shown in Figure 4-11. An exoatmospheric blackbody spectrum for 4400 K, corresponding to a $B - V = 1.4$, with no interstellar reddening, yields a very good approximation for the DCR bias for an average catalog star for all station altitudes. This approximation depends on the total instrument quantum efficiency.

The $1-\sigma$ standard deviation of DCR biases of cataloged stars is on the order of 150 mas at 45 degrees elevation for all sites in Figure 4-11. While the reconstruction of DCR from catalog photometry is imperfect, the worst-case prediction error of at most 20 mas for all elevation bands and for all sites is significantly below the $\sigma \approx 150$ mas variation in the DCR of all possible catalog stars. This demonstrates that there is a significant benefit to the quality of the astrometric solution in making per-star DCR bias corrections.

Much of the remaining analysis will discuss the DCR bias of the average star, but the $> \pm 150$ mas

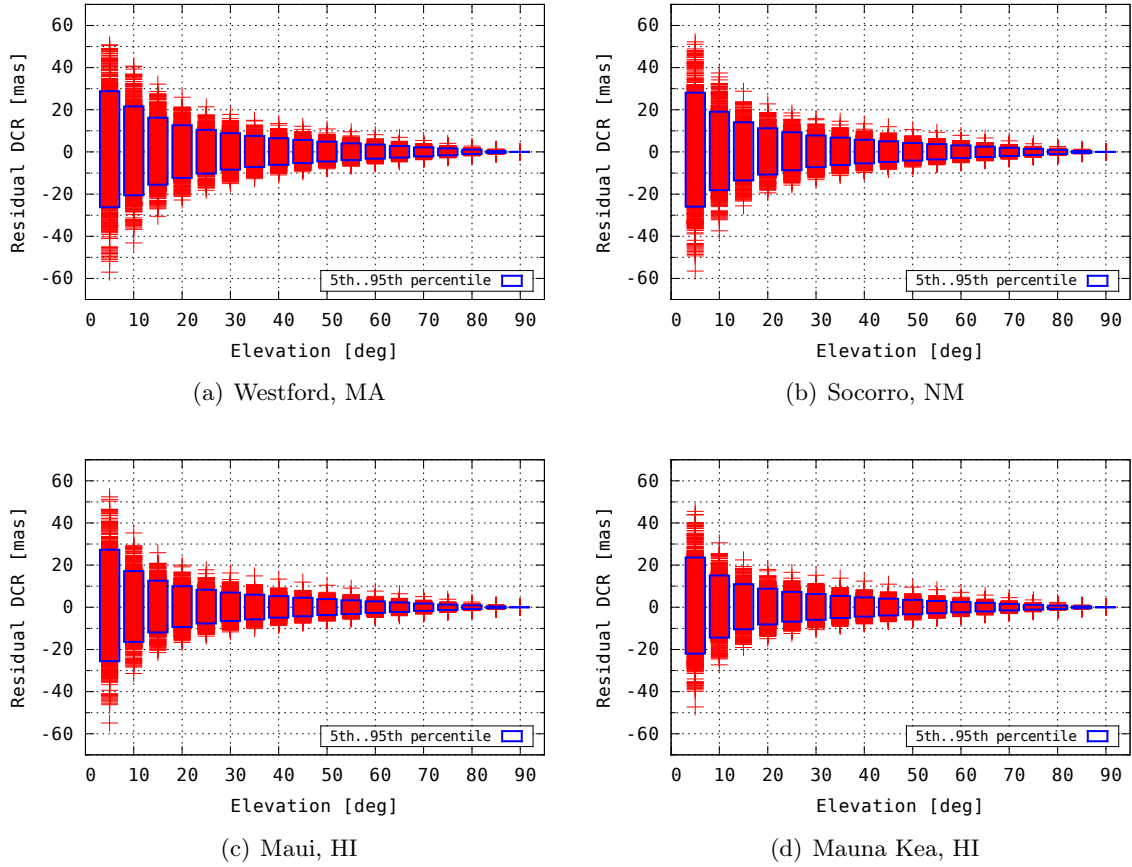


Figure 4-10: DCR Prediction Error vs. Elevation for Polynomial Function of Gaia DR2 Photometry and Stellar Parallax. The three-color photometry (G, BP, RP) and parallax measurements from Gaia DR2 provide a marginally better estimate of DCR bias than UCAC5 color photometry (Figure 4-9), but absent higher fidelity information on individual stars’ spectra, this information is insufficient to fully null out DCR bias.

variation about the average of the catalog implies that mean-value analysis may have little bearing on actual stellar DCR in any given field or the amount of stellar DCR bias imparted into any particular astrometric observation. The removal of the stellar contribution to DCR bias requires knowledge of the actual ensemble of stars used to perform the astrometric reduction in the first place. This limits the ability to make after-the-fact Farnocchia-style catalog-based corrections to already-reduced data in a color- and elevation-agnostic fashion. In order make after-the-fact catalog corrections rigorously, it would be necessary to know the spectral QE of the sensor and to estimate which stars were likely to be used in the reduction of each observation and to account for elevation at which it occurred.




Figure 4-11: DCR Bias for the Average Catalog Star Relative to Solar-Analog. Mean and $\pm 1\sigma$ bounds of DCR from the ensemble of all stars in the UCAC5 catalog. While the DCR bias of the average star in the UCAC5 catalog may be well-approximated by the DCR bias of an unreddened $B - V = 1.4$ blackbody spectrum, the $\pm 1\sigma$ variation about the mean is of a comparable magnitude, meaning that the worst-case error in per-star DCR corrections (Figure 4-9) would still yield an improvement in astrometric accuracy over observations that do not account for stellar DCR bias. The upward curl at low elevations at lower elevation sites and with lower meteorological visibility is a result of the difference between the spectra of the solar-analog baseline and the average catalog star diminishing when shorter wavelengths are attenuated for both spectra by the atmosphere. Thus, when blue light is cut out, both targets' spectra look more similar and the difference in their DCR diminishes. This effect is less pronounced at the 4 km altitude of Mauna Kea where the solar-analog baseline retains a large-enough portion of its short wavelengths to maintain the difference between the average (redder) star and the solar spectrum.




Figure 4-12: DCR Bias for Asteroids vs. Elevation Angle. Atmospheric conditions were simulated near sea level with a back-illuminated CCD, all-reflective optics with two aluminum-coated mirrors, and nominal meteorological conditions with 15 km visibility. Note that at low elevations, the same upward curl in the DCR bias that was seen in the DCR of stars is present in the DCR of S-type objects. At low elevations the loss of blue wavelengths to atmospheric attenuation makes the spectrum of a solar-analog star more similar to the spectrum of a Sun-illuminated S-type object. The effect is present, but less pronounced, for C-type objects.

4.2 DCR Bias of SSA Targets

4.2.1 Asteroids

Figure 4-12 shows the DCR bias computed for the nominal asteroid reflectance spectra of the SMASS C and S spectral types. These are the most numerous representatives of relatively spectrally-flat and red objects, respectively. Of the several hundred thousand numbered asteroids in the MPC database, only a few thousand are spectroscopically characterized. Given the lack of information, it would be reasonable to assume that any given asteroid would lie somewhere between these two extremes.

It is worth noting that the prediction for the DCR of S-type objects is negative (because the objects are red). However, in Section 1.2.2, the observed DCR bias for both S-type (and C-type) objects was positive. In light of the discussion of the DCR bias of stars, it is evident that not only is the separation in DCR between S-type and C-type objects evident in real MPC observations, but also the signature of stellar DCR. The average star is much redder than an S-type asteroid and the redness of the majority of fiducial stars used to reduce the MPC observations causes the observed DCR bias of both object types to be positive.

For the purpose of the analysis in this thesis, I consider the effect of a worst-case mismatch: that is, what happens if an asteroid that in reality is close to S-type is assumed to be C-type, or vice-versa. The relevant quantity for this analysis is the difference between the two traces in Figure

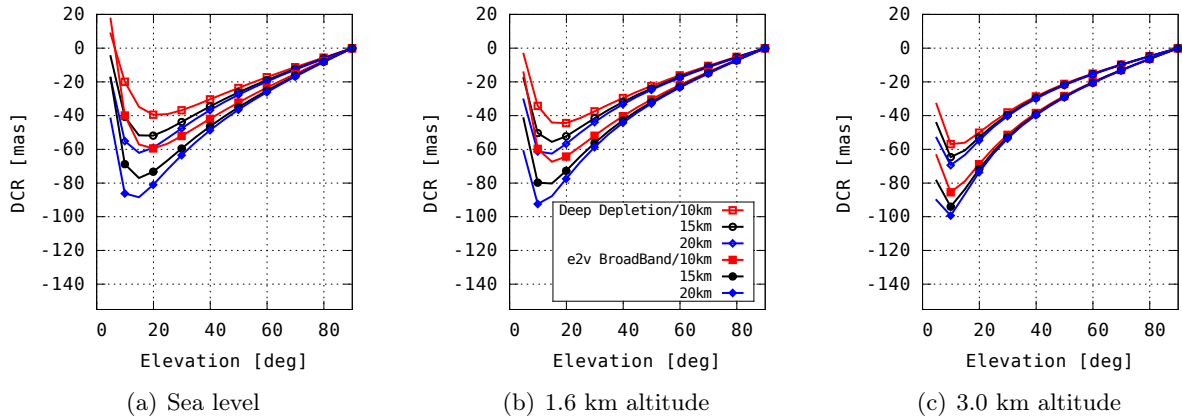


Figure 4-13: Sensitivity of Difference in S-type and C-type Asteroid DCR Bias to Meteorological Conditions and Detector Type. Deep depletion and ordinary (“broadband”) silicon detector quantum efficiencies taken from e2v Corporation datasheets [125]. The three altitudes represent the range of altitudes of stations contributing observations into the Minor Planet Center and the three meteorological visibility conditions (10 km, 15 km, 20 km) are representative of real-world extremes. There is little sensitivity to surface temperature, pressure, and humidity; traces examining variation in those parameters are not shown.

4-12. In Figure 4-13, this difference is shown as a function of station altitude, detector technology, and lower troposphere visibility.

While there is some sensitivity of this difference to detector technology and to altitude, for a fixed station operating with the same detector, there is relatively little sensitivity to temperature, pressure, or humidity (not shown in the plot). At low elevations, and especially for lower altitude sites, there is some sensitivity to lower troposphere visibility. This implies that while DCR compensation does not require knowledge of temperature pressure or humidity, it does require an estimate of meteorological visibility.

What is being shown here is the difference in DCR bias of an S-type asteroid with respect to a hypothetical solar-analog fiducial star and a C-type asteroid with respect to a hypothetical solar-analog star. But as seen in Section 1.2.2, and as will be seen again in Sections 4.3 and 4.4, the average star looks pretty close to an average S-type asteroid in terms of its DCR vs. elevation curve. This means that the DCR difference between an average star and a relatively flat reflector like a C-type asteroid, or for that matter, any object whose spectrum differs by a comparable amount from that of a typical star, is also sensitive to detector quantum efficiency, site altitude, and lower troposphere meteorological visibility at the time of observation.

In Section 2.3.1, it was stated that while few asteroids have been spectroscopically classified, over 100,000 have albedo measurements. Because albedo can be used to infer likely spectral type, it

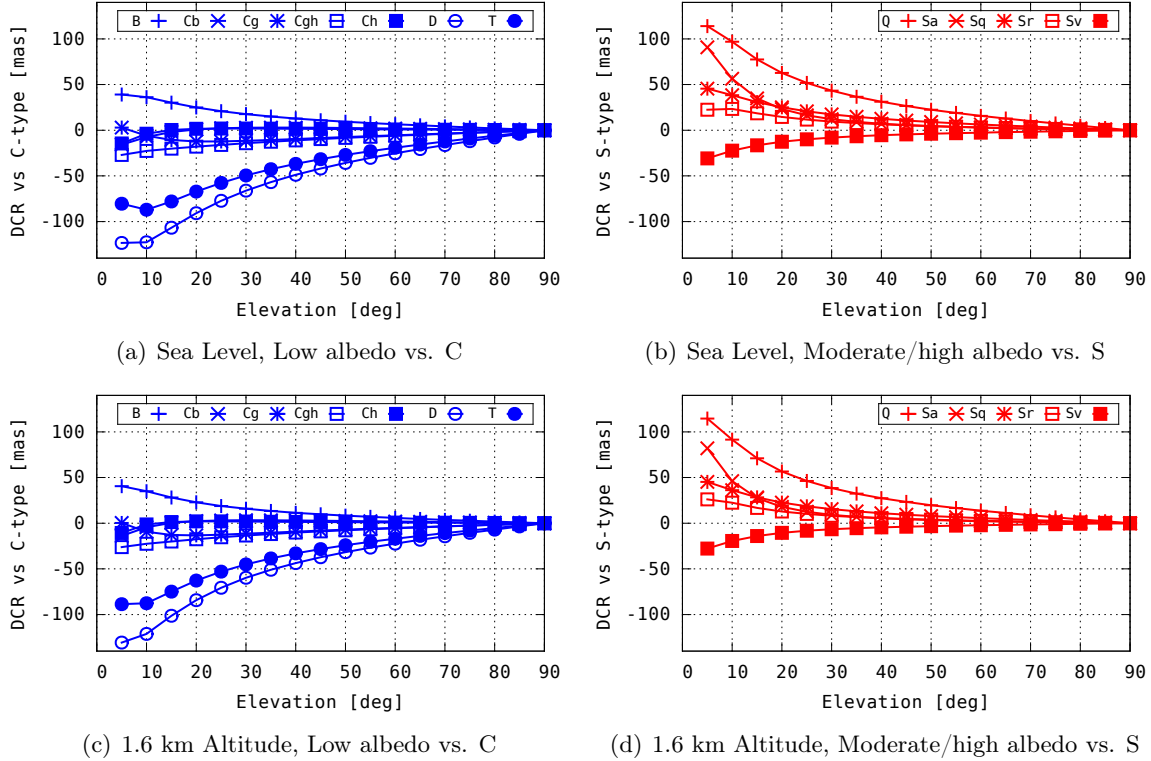


Figure 4-14: DCR Accuracy with Albedo-Based Classification. While for many classes within the C-complex and S-complex there is little difference in DCR among the members, there can be ambiguity in albedo with other taxonomic classes whose DCR is significantly different from the complex members’ DCR. Thus, while albedo can provide a reasonable estimate of DCR for many objects, in the worst case the DCR estimation error can be as large as the difference between a C-type and an S-type object.

is necessary to quantify the degree to which residual ambiguity exists in albedo-based classification in terms of DCR prediction accuracy. This comparison is made in Figure 4-14, comparing the DCR of low albedo B, Cb, Cg, Ch, Cgh, D, and T objects against the DCR of a C-type object and intermediate-to-high albedo Q, Sa, Sr, Sv, and Sq objects against S-type.

The error in DCR from assuming all low albedo objects are C-type and all intermediate-to-high albedo objects are S-type is often less than the error from assuming an S-type object is a C-type object or vice-versa, but not always. Among low albedo objects, D and T-type objects are not members of the C-complex and have larger spectral slope in the silicon passband than C-type. Among intermediate-to-high albedo objects, the difference between Q-type and S-type objects’ DCR is almost as large as the difference between S-type and C-type. Thus, while albedo information often provides an improvement in DCR correction over complete ignorance of target color, spectroscopic observations are necessary to guarantee accurate DCR estimation for inner Solar System objects.




Figure 4-15: DCR Bias for KBOs vs. Elevation Angle. Same assumptions as Figure 4-12.

4.2.2 Kuiper Belt Objects

Figure 4-15 shows the worst-case possible DCR bias for very red Barucci Group IV objects and spectrally-flat reflector Group I objects. It is necessary to recall that the DCR bias in reduced observations depends on the post-processing corrections made for fiducial stars. While the DCR bias of Group IV objects against solar-analog stellar spectra is significant, the DCR bias with respect to an “average” reddened and non-solar-analog star is nearly zero. Thus, if stellar DCR bias is corrected for on a per-star basis, the bias for Group IV objects would be large, but if it were neglected, the bias would on average be zero. Similarly, if observations of Group I objects were reduced with correction for stellar DCR bias, their DCR bias would be zero, but if those observations were reduced with stellar DCR neglected, their bias would be large and positive.

With increasing accuracy of star catalogs and the consequent ability to null out stellar DCR to high accuracy, the appropriate course of action in astrometric reduction is to correct for the effect of stellar DCR bias on a per-field basis. This means that the worst-case Quaoar bias and S-type bias would be expected in real data, and the simulations in this thesis are based on this assumption.

4.2.3 Earth Satellites

The DCR bias of a satellite target depends on the satellite’s surface materials and the particular cross-section the target presents to the observer at any point in time. Laboratory measurements of spectra of common satellite surface materials are useful in estimating the order of magnitude of satellite DCR bias and its sensitivity to atmospheric conditions and variability across satellites. DCR bias is largely insensitive to atmospheric temperature, pressure, and relative humidity but is somewhat sensitive to meteorological visibility and is very strongly sensitive to the materials of the




Figure 4-16: DCR Bias Predictions for Satellites at Sea Level

satellite being observed.

Figure 4-16 shows a calculation of DCR bias versus elevation angle of a solar panel/Kapton satellite observed from a ground site near sea level for two extremes of meteorological visibility, ratio of materials, and for small variations in material reflectance spectra. The spectrum is perturbed by varying the ratio of solar panel to Kapton and by moving the wavelength of the solar panels' blue glint left and right by 7.5 nm. At elevation angles typical of GEO SSA, the average DCR and the variation in DCR are on the order of 100 mas. The sensitivity to composition is greater when the satellite's spectrum skews blue.

4.3 Extended Analysis of Minor Planet Center Observations of Asteroids With Known Spectral Types

Here I expand on the analysis of asteroid observations from Section 1.2.2 to show that not only is DCR bias a significant contributor to the systematic error floor in ground-based optical astrometry of asteroids, but to also demonstrate that the standard proscription to observe with narrow passbands does not provide sufficient attenuation of DCR bias. As before, I examine astrometric residuals of C-type and S-type asteroids with well-constrained orbits derived from long observation arcs or multiple radar measurements and consider the residuals in station-local horizontal/vertical coordinates.

As of December 2018, there were 2117 asteroids with known spectral types in the JPL Horizons small body database, of which 913 are classified as belonging to the S complex by either Tholen or SMASS and 483 belonging to the C complex. Of these, 33 S-type objects and 10 C-type objects have at least two radar observations and a total observing arc covering at least two orbital periods.

Table 4.1: Farnocchia-Correctable CCD Observations of Asteroids with Known Spectral Types and Well-Constrained Orbits in the MPC Database as of December 2018

Passband	2 or more radar obs		Optics-only, 20 year arc	
	S-type	C-type	S-type	C-type
Unspecified	12774	3954	452312	261594
Clear	509	161	3578	1942
g'	28	12	2761	1515
r'	83	30	8353	4874
i'	184	51	8266	4233
V	14881	4033	369526	230605
R	16142	3195	184883	107995
I	26	21	1624	1041

Table 4.2: Top Five Sites Observing Long-Arc Optics-Only S- and C-Type Objects with Selected Passbands

Passband unspecified			V passband			R passband		
Site	No. S-type obs	No. C-type	Site	No. S-type	No. C-type	Site	No. S-type	No. C-type
704 (LINEAR)	336446	187690	703	151898	86524	D29	37834	23377
699	40607	24133	689	79208	68609	699	13798	8117
608	18500	10674	G96	49546	28334	Z22	12782	8130
644	13575	8484	E12	29239	17477	Y00	11764	7632
673	6681	22076	691	20876	12649	644	9391	5361

Replacing the requirement for two radar observations with a requirement for a total observation arc of at least twenty years yields 865 S-type objects and 483 C-type objects. The number of CCD observations in the Minor Planet Center archive of these objects as of December 2018 with available Farnocchia catalog corrections are shown in Table 4.1, broken down by number of observations made in passbands within the sensitive range of silicon detectors.

There are far fewer observations of radar-tracked objects than there are of objects with optics-only orbit estimates. As such, the bulk of the analysis is confined to the latter. I assume that where a magnitude is reported to the Minor Planet Center, the filter band associated with that magnitude is the filter through which the observation was made. Under this assumption, most of the observations of objects with optics-only orbits occurred with an unspecified passband or with Johnson V or R passbands. Table 4.2 further breaks down observations of those objects made with V, R, and unspecified passbands into the top five sites in number of archival observations of S and C-type objects. Although very few observations are reported to have occurred with a “Clear” passband, I assume that most of the observations with unreported passband occurred with a bare CCD. This is true for observations made with the LINEAR system (site 704).

Figure 4-17 shows astrometric residuals versus elevation of corrected MPC archival observations of radar-tracked objects observed with unspecified passbands and Figure 4-17 shows objects with optics-only orbits observed with V, R, and unspecified passband. Figures 4-19, 4-20, and 4-21 further break down the observations made in specific passbands by contributing sensor. Residuals




Figure 4-17: Residuals of Bare CCD Farnocchia-Corrected Observations of Asteroids with Radar-Constrained Orbits and Known Spectral Types. Horizontal and vertical residuals (with respect to JPL Horizons orbit) of MPC archival observations with available Farnocchia catalog corrections of asteroids categorized as belonging to the C or S spectral types and having well-constrained orbits defined as having two or more radar observations in concert with at least two revolutions’ worth of observation arc. The $1\text{-}\sigma$ error bars within each elevation band are computed as the RMS of per-observation formal uncertainty assuming that the per-observation formal uncertainty is a circularly-symmetric zero-mean Gaussian with $\sigma = 1$ arcsec. This is a conservative estimate for nearly all observations corrected for catalog bias. There is an insufficient amount of observations of radar-tracked asteroids for mean-value analysis of residuals to indicate the presence or absence of a statistically-significant DCR effect confined to only the local vertical coordinate.

are aggregated into five-degree elevation bins and a formal $\pm 1\sigma$ error bound for the mean residual in each bin is computed as the RMS of all the observations’ formal errors that fall into the bin. Following the convention of [126], the formal error on each observation is assumed to be 1 arcsec in both RA and Dec. This is a conservative assumption but the point of this exercise is to show that DCR bias is the dominant systematic error source and not to demonstrate the accuracy of DCR bias predictions against available data². What I aim to show here is that there is a nonzero mean error in the astrometry that is confined to the local vertical, but not local horizontal, coordinate and that this error is statistically significant relative to a conservative assumption on the formal uncertainty in the underlying observation. An expected separation of the mean vertical, but not horizontal, residuals of S and C-type asteroids, with redder S-type asteroids having a more negative elevation residual, serves as further evidence for the claim that DCR bias is the dominant systematic error source in corrected ground-based astrometry.

As expected, the low numbers of observations of objects with radar-constrained orbits (Figure 4-17) make it impossible to claim that DCR bias is a dominant error source in that dataset, but analysis of observations in all passbands of objects (with long-arc optics-only orbits) with large numbers of observations shows that systematic errors from DCR bias exceed systematic errors from sources

²Section 4.4 will address the evaluation of DCR models against real data.


other than DCR bias. When resolving observation residuals into horizontal and vertical coordinates, any systematic observation biases caused by atmospheric refraction will be mostly confined to the vertical coordinate. Systematic errors caused by uncorrected catalog biases, centroiding, streaking, and pixel phasing effects would not be confined to only the vertical coordinate over the full range of observation geometries and target locations on the focal plane. Thus, the mean value of the horizontal residual may be used as a proxy for the strength of systematics not pertaining to DCR.

The horizontal component of residuals of observations made in the Johnson V passband is indistinguishable from zero to better than a 3σ level while observations made in the Johnson R and unspecified (assumed bare CCD) passbands are indistinguishable from zero to better than a 1σ level. Furthermore, the horizontal residuals of S-type versus C-type asteroids are in all cases indistinguishable from each other to better than 1σ for objects with long-arc optics-only orbit estimates. A reasonable maximum estimate for the level of non-DCR long-term systematic error can be taken as approximately 25 mas from the peak horizontal residual of observations made with the Johnson V band in Figure 4-18(a), although most of this error appears to come from site 703 while other sensors observing with this passband appear to have horizontal systematics indistinguishable from zero (Figure 4-20).


In contrast, the vertical components of the residuals of the same observations show statistically significant amplitudes or statistically significant *separation* between observations of objects with different spectral types. Over a wide range of elevations, the mean residuals are nonzero to more than a 3σ level and the separation between residuals of different classes of objects are nonzero to more than 2σ over several adjacent elevation bins. This is a good time to recall that the σ 's here are computed with the assumption that the random error in each individual observation is drawn from a 1 arcsec RMS distribution, while in reality many Farnocchia-corrected observations are drawn from tighter distributions. In all passbands shown in Figure 4-18, there exist many continuous regions where the separation between the vertical residuals of the two classes of objects exceeds the 25 mas worst-case level of long time-scale systematics defined above. This is particularly true of observations made with the Johnson R passband, which covers the region where the spectral slope difference between S-type and C-type reflectance curves is large, and of observations made with unspecified (assumed bare CCD) passbands.

This is a somewhat incestuous comparison given that the orbit against which the residuals are computed is derived in large part from those very observations, and would thus be expected to be mostly zero-mean. However, the imposition of a twenty-year minimum data arc length imposes


a strong dynamic constraint on the orbit solution. In Section 5.2, I examine the distribution of post-fit residuals of simulated biased observations of multi-apparition asteroids and find that the vertically-aligned bias on the input observations is, in fact, reflected in the post-fit residuals of a fit to biased data, rendering valid the analysis presented in this section. Further, the fact that there is clear separation of vertical residuals with target spectral type but no such separation in horizontal residuals implies that the systematic effect of DCR is measurable in spite of the challenge of not having an independent orbit solution derived from bias-free observations, and the fact that this signature exists even in observations made with narrow passbands makes the point that narrow passbands alone do not provide sufficient suppression of DCR bias.



(a) Optics-only objects observed with Johnson V




(b) Optics-only objects observed with Johnson R




(c) Optics-only objects observed with unspecified passband, assumed bare CCD


Figure 4-18: Residuals of Narrow Passband and Broad Passband Farnocchia-Corrected Observations of Asteroids with Well-Constrained Optics-Only Orbits and Known Spectral Types. Horizontal and vertical residuals (with respect to JPL Horizons orbit) of MPC archival observations with available Farnocchia catalog corrections of asteroids categorized as belonging to the C or S spectral types and having well-constrained orbits defined at least twenty years' worth of observation arc. Error bars are $1\text{-}\sigma$ assuming 1 arcsec formal error of the underlying observations. Unlike radar objects, there is a sufficient number of observations of long-arc optics-only asteroids to state unambiguously that DCR bias is the dominant systematic error in asteroid observations. Comparison of observations of S-type and C-type asteroids' vertically-aligned residuals indicate the presence of a difference in the mean vertical residual between the average S-type and C-type asteroid that is nonzero and significant to at least 3σ over a sizable band of elevation angles while the difference in horizontal residuals is effectively zero.



(a) Optics-only objects observed by LINEAR with bare CCD




(b) Optics-only objects observed with unspecified passband by all sites other than LINEAR




(c) Optics-only objects observed with unspecified passband by site 699


Figure 4-19: Residuals of Observations of Optics-Only Objects with Unspecified Passband Broken Down by Sites. LINEAR observations with unspecified passband are known to have occurred with an unfiltered silicon CCD. Error bars are $1\text{-}\sigma$ assuming 1 arcsec formal error of the underlying observations. Corrected LINEAR observations in the aggregate are free from systematic error in the horizontal coordinate but contain statistically significant bias *and* separation in bias between S- and C-type objects in the vertical coordinate. Data from LINEAR dominates the corpus of observations of these objects, meaning the LINEAR residuals would be expected to have zero mean with respect to orbits computed from the same dataset. The separation of vertical residual traces corresponding to objects of different spectral classes implies the existence of statistically significant DCR bias in the observations. Furthermore, the presence of a separation of residuals between the two types in data from other sites to a greater level in the vertical than in the horizontal coordinate adds to the evidence that DCR bias is a significant contributor to the observation error of asteroids, even in the presence of other systematic errors evident in the nonzero mean horizontal residuals from site 699.



(a) Optics-only objects observed by site 703 in V passband

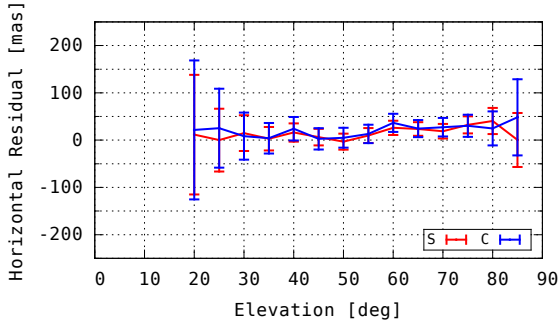


(b) Optics-only objects observed by site 689 in V passband





(c) Optics-only objects observed by site G96 in V passband



Figure 4-20: Residuals of Observations of Optics-Only Objects with Johnson V Passband Broken Down by Sites. Error bars are $1\text{-}\sigma$ assuming 1 arcsec formal error of the underlying observations. Little separation in the vertical is expected in the V passband and the purpose of this analysis is to gauge the level of site-specific idiosyncrasies and medium timescale systematic errors.



(a) Optics-only objects observed by site D29 in R passband



(b) Optics-only objects observed by site 699 in R passband



(c) Optics-only objects observed by site Z22 in R passband




Figure 4-21: Residuals of Observations of Optics-Only Objects with Johnson R Passband Broken Down by Sites. Error bars are $1\text{-}\sigma$ assuming 1 arcsec formal error of the underlying observations.




Figure 4-22: Mean $\pm 3\sigma$ Horizontal and Vertical Residuals of S- and C-type Asteroids in Subset of LINEAR Data Reprocessed with UCAC5. $N \approx 1400$ C-type and ≈ 3500 S-type objects. While the overall astrometric quality is high as inferred from the horizontal residuals, the vertical residuals do not match the prediction for S-type asteroids and the difference between the curves with and without stellar DCR correction in the data does not match the same difference in the prediction. This indicates that while DCR is unambiguously present, the spectral quantities required to predict it are not sufficiently well-known to enable after-the fact correction on a per-target basis. Further, the worse prediction accuracy for astrometry without per-star DCR compensation (that is, an assumption that a mean star looks like a 4400K blackbody) implies that color-agnostic after-the-fact corrections are going to perform measurably worse than color-aware astrometric solution because the DCR of the “average star” is not a physically meaningful quantity given the large variation in stellar DCR within the UCAC5 catalog, and in all catalogs with an equivalent ensemble of stars.

4.4 Re-Reductions of LINEAR Imagery

Having shown the presence of DCR bias in archival data, I proceed to examine the consistency between observed bias and predicted bias in high-quality asteroid astrometry. The purpose of this exercise is to gauge the degree to which astrometric observations of asteroids can be corrected after-the-fact using an assumed model of spectral reflectance and atmospheric spectral refraction. I show that it is difficult to achieve agreement between theoretical predictions of DCR bias and observed DCR bias in the absence of per-asteroid spectroscopic observations, even with spectrally-aware astrometric reductions.

A subset of raw image data from the LINEAR system was obtained from MIT LL archives [127]. A subset of several dozen nights between 2002 and 2012 were re-reduced with the UCAC5 star catalog. The reduction procedure was run both with the fiducial star’s DCR correction applied and with it neglected to highlight the effect of, and evaluate the accuracy of, the correction. About 3500 distinct observations of S-type asteroids and 1400 distinct observations of C-type asteroids are

present in the dataset. Both runs' residuals against Horizons reference ephemerides are plotted in Figure 4-22.

Overall, this reprocessing was of higher quality than the original astrometry. Individual formal errors on the astrometric re-reduction were on the order of 50 mas RMS with many tens or hundreds of fiducial stars. The reason this lower formal error was achievable in post-processing was because more accurate (but slower-executing) algorithms were used than on the real LINEAR pipeline and data was freely discarded if it did not meet goodness conditions on the asteroid SNR, the ensemble of available fiducial stars, and the star clutter environment around the asteroid PSF. Fiducial stars were selected to remain within the dynamic range of the detector and circular-symmetric Gaussian point-spread functions were fit to both fiducial stars and asteroids³. Independent plate solutions were computed for each of the readout channels of the CCID-16 chip, rather than the global model plate solution used in the original data reductions submitted to the MPC, once it was realized that there was a small systematic horizontal or vertical shift across channel boundaries between image regions sent through adjacent readout circuitry⁴. The lower formal error on each observation enables meaningful analysis with a reduced data volume over that present in the full MPC archive. A reduction from a 1000 mas RMS to 50 mas RMS is roughly equivalent to a 400-fold increase in data volume, bringing the thousands of reprocessed observations to par with the tens of thousands of observations analyzed in Section 4.3. This improved accuracy over the 1 arcsec RMS assumption on MPC data enables analysis with a much smaller number of distinct observations for each spectral type than is available in the MPC archive.

The quality of the astrometric solution is best evaluated in the local horizontal coordinate. The more numerous S-type asteroids' mean residuals remain within the 25 mas accuracy claim of UCAC5. For evaluating the quality of stellar DCR correction, the S-type asteroids provide the better dataset⁵. Figure 4-22(c) shows the S-type vertical residuals with and without stellar DCR correction applied and overlaid with the DCR bias prediction for a generic S-type asteroid and the DCR bias of the mean catalog star as discussed in Section 4.1.3.3. The empirical data does not match the prediction for S-type asteroids and the difference between the curves with and without

³Unlike the processing of near-Earth SSA data for this thesis that necessitates use of Gaussian streak PSFs, LINEAR objects' motion is negligible in a single exposure and the optical quality is uniformly good, enabling use of circular-symmetric Gaussian PSFs.

⁴The discontinuity at the boundaries was on the order of a quarter pixel. One possible source of this error is clock phasing errors in the readout circuit. Eric Pearce believes another non-negligible effect is charge drift in the frame transfer area of the chip caused by spurious static fields present during the readout process.

⁵This may just be luck-of-the-draw: perhaps the particular C-type asteroids in this dataset had lower-quality orbits.

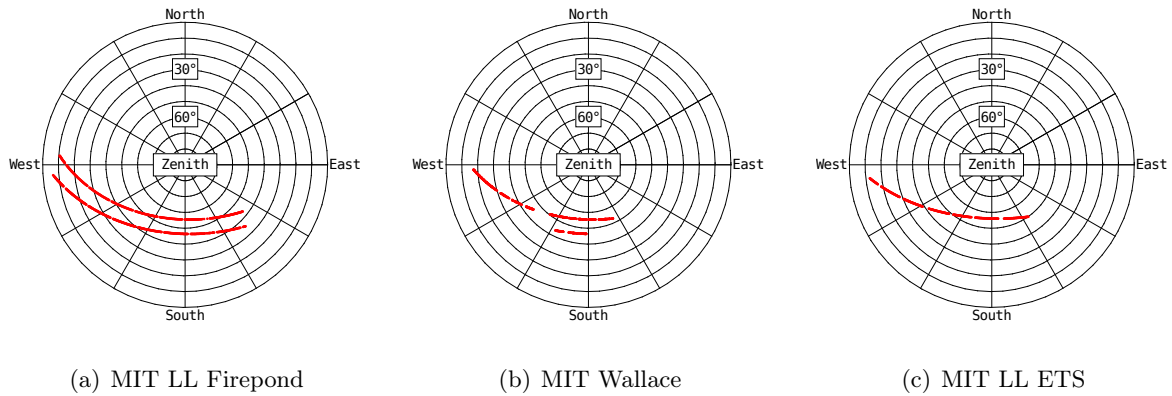


Figure 4-23: Observing Geometry of 3122 Florence Flyby in 2017

stellar DCR correction in the empirical data does not match the same difference in the prediction (predicated on the assumption that the DCR of the mean star is equivalent to that of a 4400 K blackbody spectrum). This indicates that while DCR is unambiguously present, the per-asteroid spectral quantities required to predict it are not sufficiently well-known (or the asteroid orbits for this small number of targets are not sufficiently constrained) to enable after-the fact correction of existing observations or to match up observed DCR with predictions. Aside from being an argument in favor of per-star DCR correction, this is an argument in favor of more spectroscopic measurements of asteroids concurrent with astrometric observation and in favor of a more concerted effort at spectral characterization of observing instruments and of real-time atmospheric conditions, as only mean meteorological visibility corrections were applied in this reprocessing.

4.5 Preliminary Measurements on 3122 Florence

4.5.1 Scenario

On the nights of August 31, 2017 and September 1, 2017, MIT LL sensors in Socorro, NM and Westford, MA, and the MIT Wallace Observatory in Westford, MA were tasked to observe the near-Earth flyby of the asteroid 3122 Florence. The MIT LL sensors observed with bare silicon detectors and the Wallace site observed through Sloan r' and i' passbands. The nearest approach distance was about 7 million kilometers and the object peaked at a brightness of about 8th magnitude. Target motion was about 0.5 arcsec / sec with respect to the star background and a wide diversity of elevation angles were covered by the pass (Figure 4-23).

4.5.2 Timing Validation

One of the obstacles to collecting data on asteroid passes is that object velocity against the star background can introduce an error source from instruments that do not adequately time-tag their exposures. Because most astronomical targets move slowly relative to the sky background, it is rare for instrumentation not specifically fielded for Earth-orbit SSA to time-tag exposures to the kinds of millisecond accuracies required to accurately constrain in-track motion for targets that move at many arcseconds per second.

Even though almost all near-Earth asteroid passes do not move at nearly the rate that satellites do, they nevertheless can move at significant fractions of an arcsecond per second. If that motion were to be aligned with the local vertical at any time, any unknown time-tagging biases in the exposures (even long exposures) would either mask or skew the apparent error in the vertical direction by the product of velocity and timing error.

As the point of this thesis is to measure systematic vertical errors on the order of several tens of milliarcseconds, it is necessary to ensure accurate timing calibration of all observing equipment not specifically designed for sub-millisecond accuracy. This requires observing GPS and GLONASS calibration satellites with the same instrument, with camera electronics operating in the same mode, and at the same exposure times, as used to observe the actual asteroid pass. This is a restrictive criterion, as many astronomical telescopes are not capable of smoothly following a GPS satellite while operating at the exposure necessary to capture an asteroid and many more do not have a field-of-view that is sufficient to detect streaking fiducial stars and the satellite or a streaking satellite and stationary fiducial stars at the same time. Fortunately, the Wallace telescopes had sufficient field of view and frame rate to observe these targets in sidereal stare.

In order to decouple any DCR bias of the calibration satellite, only horizontal motion and horizontal error must be used. The system timing bias may be estimated as:

$$\Delta T \approx \frac{\Delta AZ \times \cos EL}{\dot{A}Z \times \cos EL} \quad (4.3)$$

where $\Delta AZ \times \cos EL$ is the horizontal component of an astrometric residual and $\dot{A}Z \times \cos EL$ is the horizontal component of the *inertial* angular velocity vector in the focal plane. This is *not* a full transformation between inertial and rotating coordinate frames. Nothing is gained by mixing in the rotational velocity of the Earth. It is best to think of this process as drawing from a family of inertial coordinate systems, each aligned to the local station horizon at the time it is selected.

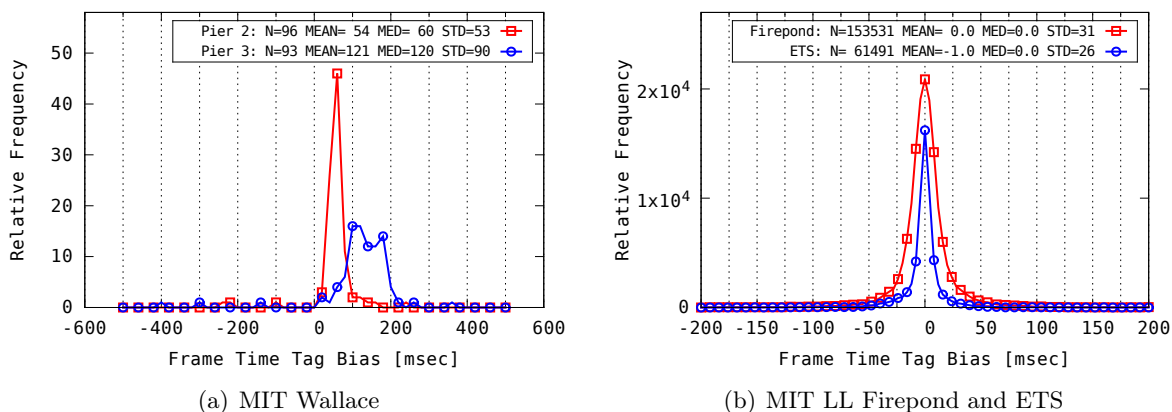


Figure 4-24: Timing System Checkout for MIT LL Sensors and the MIT Wallace Observatory. Time biases are referenced against observations of GPS and GLONASS satellites. The MIT LL sensors use sub-millisecond hardware timing capture, as demonstrated by the essentially zero-mean time bias in those sensors. The Wallace observatory uses software synchronization to an NTP server, and this process is both noisier and systematically biased. Time biases of 100 msec on 3122 Florence correspond to angle errors of about 50 mas at closest approach. Note also the small number of observations generated at the Wallace site. The mounts at the Wallace facility are not capable of tracking GPS satellites smoothly enough to generate good astrometry. Consequently, the mounts were repeatedly repositioned into sidereal track and the satellite was allowed to streak through the field of view. This is a very labor-intensive process, thus the low number of observations.

Obviously, better timing system performance data is collected when the inertial velocity of the target is most closely aligned with the horizontal. This alignment does not always exist and while many GPS/GLONASS satellites are visible at any time from any site, it is necessary to plan the calibration observations to occur at those times where the satellites are grazing the Eastern or Western horizons.

Figure 4-24 shows timing system biases computed on GPS and GLONASS satellite observations from the Wallace observatory and the two MIT LL facilities in preparation for the observing 3122 Florence. While the MIT LL sensors' time-tagging error is basically zero-mean, the Wallace sensors are off by many tens of milliseconds. This timing error was corrected prior to generating the astrometric observations from that sensor.

4.5.3 Ephemeris Refinement

Observation residuals for all stations' observations of 3122 Florence were first calculated against the JPL reference orbit from October 21, 2017 and found to deviate from the JPL orbit in the horizontal coordinate (Figure 4-25) by more than 0.2 arcsec. The same error was not present in calibration

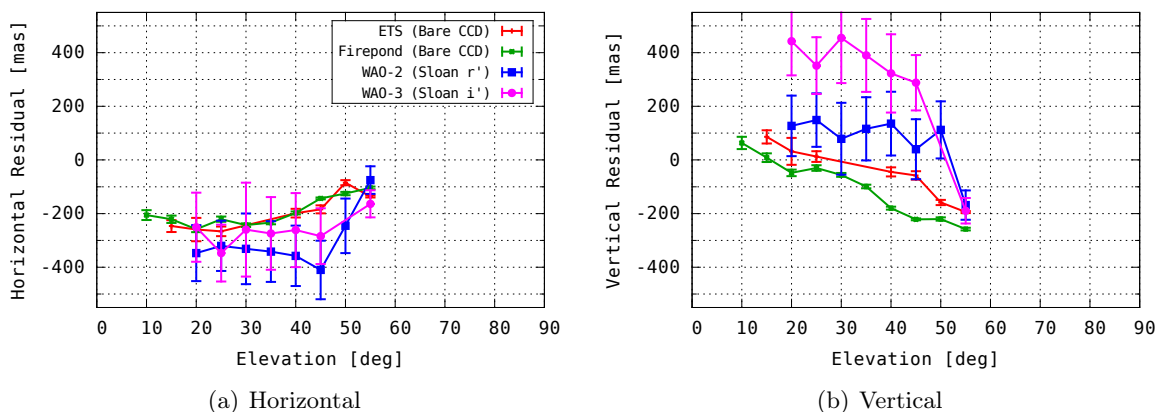


Figure 4-25: Residuals of 3122 Florence Against JPL Ephemeris Dated Oct 21, 2017. Mean residual and its uncertainty vs. elevation. The $\pm 3\sigma$ error bars from each sensor are computed using the per-image astrometric solution fit RMS residual as the nominal uncertainty of each observation, attenuated by the square root of the number of observations taken within each elevation bin. The error budget on the measurement is driven by defect of illumination, which for the range of this pass was predicted to be no more than 100 mas, in the worst-case at a high phase angle, and less than that for the near-zero phase angle during this flyby. Thus the high residual in the horizontal coordinate from all sensors is evidence that the target is off-nominal with respect to the JPL ephemeris. The amount of offset is consistent with the formal uncertainty bounds of the JPL solution, but renders meaningless any analysis that attempts to measure DCR on the order of 100 mas. That analysis requires an ephemeris refinement.

measurements of GPS and GLONASS satellites taken during the observing session and the bias was attributed to a real error in the JPL orbit estimate. The error in the horizontal bias indicated a cross-range shift of a few kilometers, which was consistent with JPL’s published uncertainty of the RAAN of the orbit. The OrbitTool was used to compute a refined trajectory using data published from other sensors before and shortly after the flyby and the resulting horizontal residuals are more in line with expectations (Figure 4-26).

4.5.4 Results

Overall, approximately 250,000 distinct observations were generated across the three sites, with the overwhelming majority coming from bare CCD observations at Firepond and ETS. This is several orders of magnitude more data than is available for many long-arc numbered asteroids, and because it comes from the same sensors and observes the same target with the same reflectance spectrum, the low formal error in the residual-versus-elevation curve allows more concrete conclusions to be drawn than can be drawn from reprocessing low-volume LINEAR search imagery of many distinct targets.

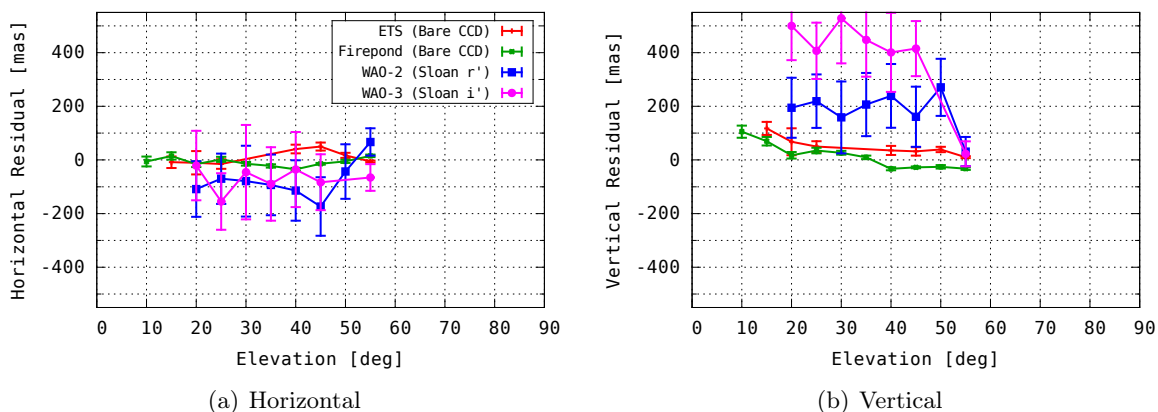


Figure 4-26: Residuals of 3122 Florence Against Ephemeris Fit to Recent Data Only. Fit spanned from April 2017 to October 2017. Observations included all ground-based data published to the Minor Planet Center and the single radar observation available at CNEOS at the time. Elevation angles were excluded from the fit to remove the effect of DCR bias. Farnocchia catalog corrections were applied to all ground-based data. A subset of the observations collected in this experiment were also used to generate the fit, but their number was kept low to avoid skewing the orbit in their favor over other sensors. The lower systematic error in the horizontal coordinate relative to the residuals against the JPL ephemeris corresponds to several kilometers worth of cross-range shift of the asteroid at its perihelion. This adjustment was consistent with the published uncertainty in the RAAN of the orbit at the time.

With system timing calibrated, the systematic error budget for the measurements is driven by the defect of illumination. The maximum size of this error source is driven by the size of the target and the distance of closest approach. Florence is about 5 km diameter at 7×10^6 km giving ≈ 0.1 arcsec angular extent. The vertical residuals from the Wallace site collected with narrow filters exceed the worst-case bound for defect of illumination and greatly exceed the predicted DCR bias when observing with narrow bands (Figure 4-27).

As Figure 4-28 shows, one possibility (though unlikely) is that this error is consistent with absorption bands near the cutoffs of the Sloan filters. No such band exists in the published SMASS spectra for 3122, and other spectral measurements by Thomas [51] are too noisy to say for certain. Further spectroscopic measurements of 3122 Florence are called for in order to understand what is happening here, but the key takeaway is that caution should be exercised in assuming that narrow bands can mitigate DCR bias for asteroids.

Analysis of the higher-volume data from the bare-silicon sensor at Firepond enables evaluation of the quality of the DCR predictions made using published spectra, MODTRAN atmospheric transmittance models, and the stellar DCR correction as described earlier in this document. Residuals with and without stellar DCR correction using both the single-color reddening-unaware technique





Figure 4-27: Predictions and Observations of DCR for 3122 Florence with Narrow Passbands. The Thomas 2014 measurement of the spectrum of Florence [51] does not cover the full passband and is extended as a generic Q-type reflectance spectrum for this analysis. Observed DCR greatly exceeds predicted DCR, serving as a cautionary example against relying on narrow passbands to mitigate DCR bias.


(for URAT1) and multicolor reddening-aware technique (UCAC5) are plotted in Figure 4-29 against the predictions of DCR for bare silicon made using generic class spectra from SMASS as well as the specific SMASS spectrum for 3122 Florence and the Thomas spectrum. The elevation residual from single-color reddening-unaware star DCR corrections tracks the particular spectrum for 3122 Florence best, but is systematically biased from it by about 100 mas. This bias could be attributed to either defect of illumination or a systematic error in the computation of DCR bias for that target, given that its spectrum in the blue was linearly extrapolated. Using the reddening-aware star DCR correction with the higher-quality UCAC5 catalog yields a larger negative residual that does not track any of the predicted DCR traces, implying Florence was redder than expected during the pass.

Taken together with the unexpectedly-high residuals from observations made with narrow passbands from the Wallace Site, the large residual with stellar DCR corrected UCAC5 residuals from bare CCD observations serves as evidence in support of the assertion that spectroscopic measurement of asteroids are necessary over the full silicon passband in concert with astrometric observations over the same passband.

Another lesson learned from this exercise is that asteroids make for poor absolute astrometric references in the evaluation of DCR models. However, the statement about the ineffectiveness of narrow passbands in mitigating DCR bias holds since the comparison of residuals from sensors operating with narrow passbands is made against comparable sets of residuals from Firepond and ETS as evidenced by equivalent mean horizontal error.




(a) Sloan Passbands and Refraction With Respect to 700nm at 30° Elevation




(b) Estimates of Florence Reflectance

Figure 4-28: Reflectance Spectra of 3122 Florence and DCR in Sloan Passbands. The elevation bias seen in the observations from the Wallace site would be consistent with absorption features near the edges of the Sloan bands. No such band exists in the published spectra and the Thomas spectrum in panel (b) is of insufficient quality to determine if such a feature exists.



(a) URAT1 Reductions



(b) UCAC5 Reductions

Figure 4-29: Vertical Residuals of Firepond Bare Silicon Observations of 3122 Florence and Predicted DCR Bias. Data from the flyby was reduced using both the URAT1 star catalog (with and without a blackbody-based single-color stellar DCR correction) and the UCAC5 catalog with and without a reddening-aware stellar DCR correction.

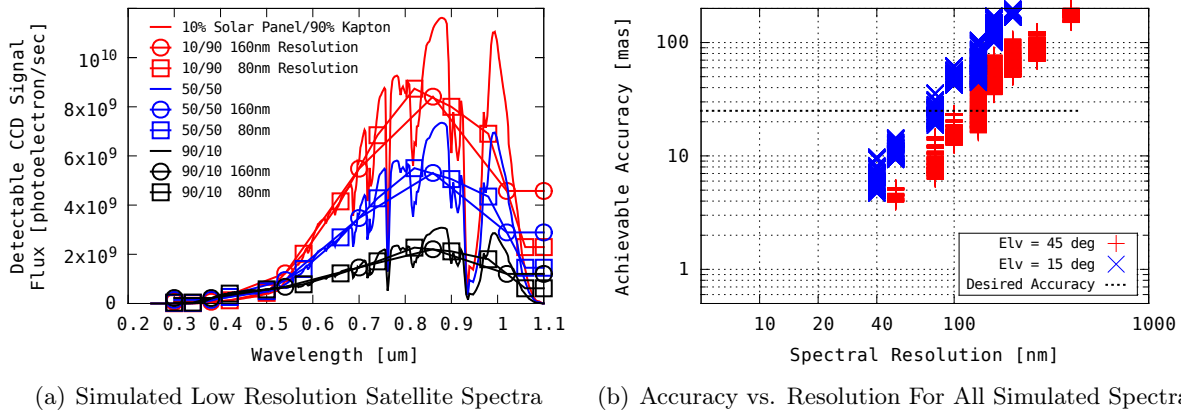


Figure 4-30: Spectroscopic Resolution Requirement Simulations for Satellites. Simulations for Firepond Facility in Westford, MA.

4.6 Spectroscopic Measurement Requirements

While the DCR of fiducial stars is predictable from information present in star catalogs, the mitigation of DCR bias on a per-target basis requires the measurement of each individual target’s spectrum. This section derives the spectral resolution and absolute accuracy requirements for that measurement that are necessary to perform DCR compensation to an accuracy commensurate with that of modern star catalogs. A target accuracy of 25 mas at an elevation of 15 degrees is chosen as the objective. This choice is driven by two considerations: the first is the 25 mas accuracy claim of the UCAC5 star catalog that was released when this thesis research began, and the second is the limitation on predicting stellar DCR at lower elevations discussed in Section 4.1.3.

4.6.1 Resolution Requirement for Satellites, Asteroids, and KBOs

The procedure to compute a resolution requirement for the spectral measurement is as follows. A high resolution baseline satellite reflectance spectrum is synthesized for the class of target being examined. For Solar System bodies, the baseline spectrum is either an S-type, C-type or Quaoar-type spectrum. For satellites, the target spectrum is modeled as a mixture of nominal material spectra, and the ratio of these materials is varied and the material spectra are perturbed by moving reflectance peaks in wavelength by several nanometers. This resulting baseline spectrum is attenuated by instrument quantum efficiency and atmospheric transmittance and then subsampled with a square boxcar sampling function to simulate the lower-resolution measurement. A DCR estimate is computed from that subsampled spectrum and compared with the DCR estimate of the original

baseline full-resolution spectrum. A spectral resolution is considered sufficient if the worst-case error in the DCR prediction from the low-resolution measurement with respect to the high-resolution baseline for all test spectra is under 25 mas at a zenith distance of 75 degrees.

For satellites the nominal spectrum used for this calculation is a linear mixture of solar panel and Kapton reflectances. The ratio of the mixture was varied by increments of 10% from 10/90 to 90/10. A further perturbation consisted of shifting the Kapton and solar panel reflectances left and right by up to 7.5 nm in wavelength. Laboratory measurements of large ensembles of both material types indicate that the onset of their reflectance peaks may vary by that amount depending on the particular thickness of the material in the case of Kapton film [60, 65], and the device chemistry and phase angle for solar panels [59]. Figure 4-30 shows an example of the full-resolution and subsampled spectra for a satellite as seen from Westford, MA and the resulting DCR estimation error from subsampling.

To achieve 25 mas accuracy at 15 degrees elevation in Westford, a resolution of 80 nm is sufficient for most cases, but 50 nm is required to cover all cases. This is between two and three times finer resolution than afforded by cycling through Sloan filters. The calculation was repeated for nominal S, C, and Quaoar-type reflectance spectra with no absorption bands or variations and for deep depletion silicon detectors. The results in Figure 4-31 indicate that 40 nm resolution is sufficient to achieve better than 25 mas correction accuracy from sites at all altitudes from sea level to the top of Mauna Kea.

4.6.2 Resolution Requirement for Asteroids with Unmodeled Absorption Bands

Taking as the reference point the experience with 3122 Florence, it is also necessary to consider the resolution requirement to compensate for DCR of asteroids observed with narrow passbands but with a possible unmodeled absorption band near the edges of the passband.

The resolution requirement for that measurement is derived by analyzing a hypothetical absorption feature in the continuum of the asteroid's reflectance spectrum. Starting with a nominal reflectance spectrum for an asteroid, an absorption feature is imposed onto that spectrum. As before, the true DCR bias of that perturbed spectrum serves as the baseline for the DCR estimated from a lower-resolution spectrum and a resolution is considered sufficient if the worst-case DCR bias estimation error is under 25 mas at a zenith distance of 75 degrees.

For unperturbed spectra, the mean reflectance spectra for SMASS C, S, and Q asteroid classes, linearly extrapolated to 300 nm served as the baseline. A single absorption feature was scanned over

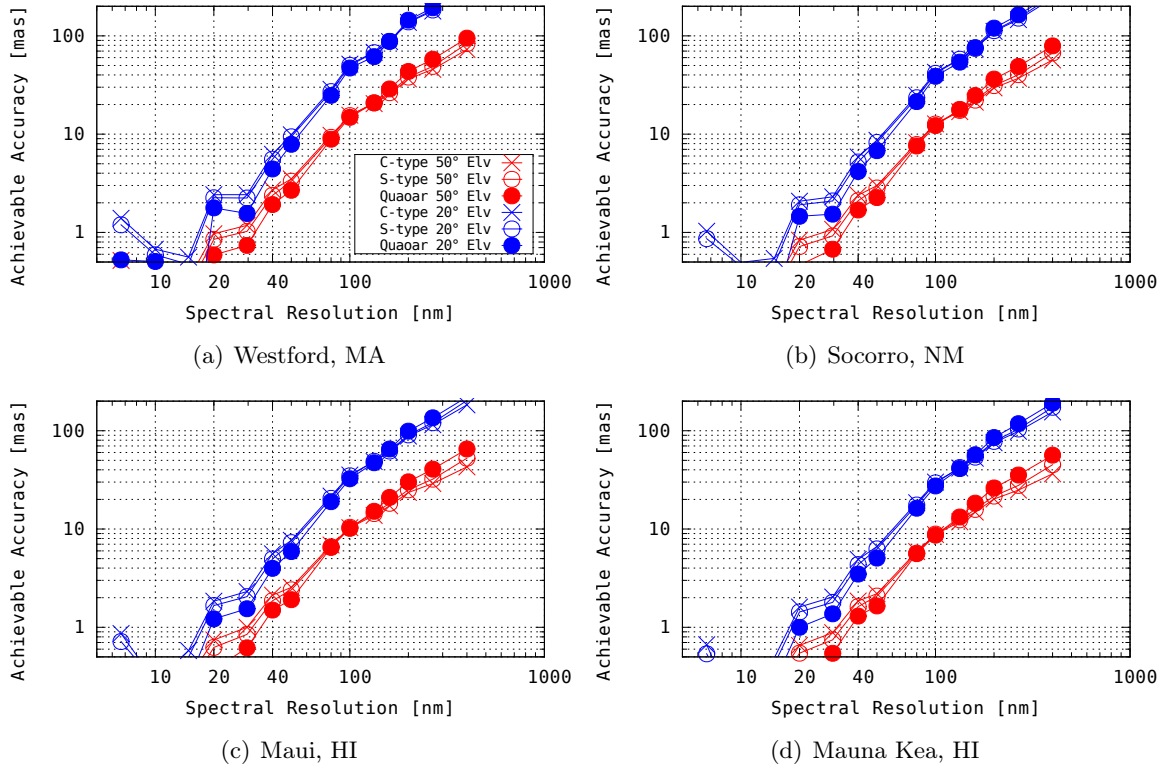


Figure 4-31: Spectroscopic Resolution Requirement for Compensation of DCR Bias in Ground-Based Observations of Solar System Bodies with Bare-Silicon Detectors. The resolution requirement is less stringent than that required for satellites because there is relatively little energy in the shorter wavelengths in spectra of natural objects than there can be in the spectra of satellites.

the entire sensitive band of silicon to construct the perturbed spectra. The boxcar averaging and subsampling operation is performed using the spectrum weighted by atmospheric transmittance and instrument overall quantum efficiency. The calculation is made for the bare CCD and for individual Sloan [128] bands for sea level altitude, 1.6 km altitude representing the LINEAR system's location in New Mexico and 3.0 km altitude representing Haleakala, Hawaii. An example intermediate calculation showing the sensitivity of DCR bias to the location of the absorption band is shown in Figure 4-32(a) with an exaggerated 55 nm wide 75% deep absorption band.

The absorption band used in this calculation was square with a depth of 16% and a width of 75 nm. The simulation showed little sensitivity to varying the width of the band while keeping the band area constant. An obvious question is whether these features actually occur in the reflectance spectra of real asteroids. In some C-type asteroids, small absorption bands in the continuum in the region between 700 and 900 nm are up to 50 nm wide and up to 5% deep [129]. Other common absorption features exist through much of the visible band redward of 500 nm [130], and for asteroids

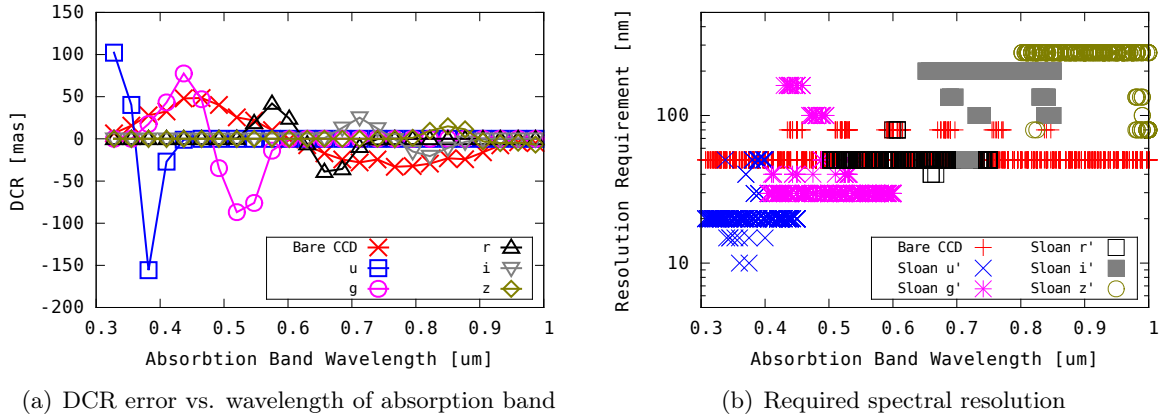


Figure 4-32: Resolution Required for Asteroid Spectra with Absorption Bands. (a) The effect of a 55 nm-wide, 75% absorption band on the DCR of an asteroid observed from the ground with respect to a reflectance spectrum without that band. (b) The spectral resolution required on the measurement of the target flux containing a 75 nm wide, 16% deep absorption band necessary to achieve a 25 mas accuracy on the DCR prediction as a function of the location of the absorption band and bands of smaller width but constant band area. The overall resolution requirement for any filter is the worst-case of an absorption feature within the filter passband excluding individual outlier points corresponding to sensitivity to the particular shape of the band. Calculations made for Westford, MA.

in eccentric orbits, temperature dependence and phase angle dependence may amplify them [51, 52]. While it is unlikely that individual absorption bands as strong as that used in this calculation exist, the experience with narrow-band observations of 3122 Florence indicates that caution is warranted in assuming narrow passbands can mitigate DCR bias and support the claim that measurements of the spectrum as detected on the ground are necessary at a resolution higher than that afforded by Sloan passbands.

Figure 4-32(b) shows an example of the exhaustive simulations of the resolution requirement for the Westford, MA. The figure plots for each filter the necessary resolution over which the averaging and subsampling of the ground spectrum is sufficiently fine to enable DCR estimation to within a 25 mas accuracy down to a zenith distance of 75 degrees as a function of the location of the single absorption band. The calculation varies the width and depth of the absorption band while keeping the area constant at $75 \text{ nm} \times 16\%$. Analogous simulations were generated for the 1.6 km altitude of Socorro, NM and the 3.0 km altitude of Haleakala, HI assuming both deep depletion and ordinary cooled silicon quantum efficiency.

Unsurprisingly, because the sensitivity of DCR estimation for a given pass-band from a coarsely sampled spectrum is greatest near the edges of the passband, we see an inverted U shape on the resolution requirement for each filter. Further, because the refraction slope $d\Delta z/d\lambda$ is greater in

the blue, the requirement is tighter for the u' and g' passbands than it is for the redder filters.

It is appropriate to quote the worst-case requirement for every filter over all three asteroid types observed at 15° elevation as the requirement for observing with that filter. The requirements are summarized in Table 4.3. As stated, there is little sensitivity of the resolution requirement to the shape of the absorption band so long as the depth-width product remains constant. An exception to that statement is the small number of simulation points corresponding to the u' and g' prime filter that lie below the overall resolution requirement shelf for those filters. I do not include those outlier points in the overall tabulation of requirements on the grounds that the square absorption band is only a calculation aid, not a physical phenomenon, and a different band shape would lead to a slightly finer or coarser requirement.

Overall, the simulations indicate that a 20 nm resolution is the bare minimum to observe in the u' band in order to achieve DCR compensation comparable to catalog accuracy of 25 mas down to a zenith distance of 75 degrees. Given the few outlier points mentioned in the previous paragraph, and the unexpectedly-large residuals of 3122 Florence observed with Sloan filters shown in Figure 4-27, it is appropriate to take the results of this simulation with a grain of salt and to append a safety factor. Thus a 10 nm resolution is recommended when observing with bluer filters, and a 25 – 40 nm resolution is recommended when observing with redder filters or bare CCDs.

Table 4.3: Calculated Resolution Requirements to Correct Observations of Asteroids to 25 mas

Site (Altitude)	Westford, MA (0.1 km)		Socorro, NM (1.6 km)		Hawaii (3.0 km)	
Passband	Regular Silicon	Deep Depletion	Regular	Deep	Regular	Deep
Bare CCD	50	50	50	50	50	80
u'	20	20	20	20	20	20
g'	30	30	30	30	40	40
r'	50	50	50	50	50	50
i'	133	200	20	200	200	200
z'	80	266	133	266	133	266

4.6.3 Accuracy Requirement in Terms of Time-Aperture Product

The definition of the spectroscopic accuracy requirement is predicated on the the desired accuracy. A spectroscopic measurement samples the target spectrum at over set of wavelength bins Δ_i centered on wavelengths λ_i . Transforming the continuous definition of DCR bias in Equation 1.2 into a discrete sum via the sampling approximation:

$$\int_{\lambda=\lambda_i-\Delta\lambda_i/2}^{\lambda_i+\Delta\lambda_i/2} \phi(\lambda, z) d\lambda \approx \phi(\lambda_i, z) \Delta\lambda_i \quad (4.4)$$

yields:

$$\delta z_0 \approx \frac{\sum_i \phi(\lambda_i, z) \Delta \lambda_i \Delta z(\lambda_i, z)}{\sum_i \phi(\lambda_i, z) \Delta \lambda_i} \quad (4.5)$$

The uncertainty in the predicted DCR bias may be computed directly from this equation under the assumption that the uncertainty of the flux sample measurement in each wavelength bin is dominated by Poisson noise. Assuming that the photoelectron counts in each measurement channel $\phi(\lambda_i, z) \Delta \lambda_i$ are large enough for Gaussian statistics to hold, the uncertainty in DCR bias is given by:

$$\sigma_z \approx \frac{\sqrt{\sum_i \text{var}(\phi(\lambda_i, z) \Delta \lambda_i) (\Delta z(\lambda_i, z))^2 - \delta z_0^2}}{\sum_i \phi(\lambda_i, z) \Delta \lambda_i} \quad (4.6)$$

It is worth noting that Equation 4.6 refers to an idealized spectroscopic measurement of the detectable photoelectron flux in the astrometric instrument at the time of observation. The flux in each wavelength bin is the product of solar illumination, target reflectance, atmospheric attenuation, and instrument quantum efficiency. If the spectroscopic measurement made to correct for DCR bias occurs concurrently with the astrometric observation and at the same ground site, then the uncertainties in target spectral reflectance and atmospheric transmittance are not relevant as independent quantities and only the uncertainty in the detectable flux matters. If, however, an attempt is made to correct for DCR bias using a prior measurement of target spectral reflectance combined with a model or measurement of atmospheric transmittance, then it is sensible to individually analyze the required accuracy of the spectral reflectance and atmospheric transmittance independently.

This analysis assumes that the spectroscopic measurement occurs simultaneously with the astrometric observation. This simplification boils everything down to the uncertainty of the detectable photoelectron flux on the ground. In a ground-based spectroscopic measurement, the noise in the measured signal in each wavelength bin is composed of the noise in the target signal (exoatmospheric signal attenuated by atmosphere) and the noise in a background term. For a slit spectrograph measurement, the background term is a function of the sky background spectrum and the angular extent of the slit. For slitless spectroscopy measurements the background term is the full sky background.

The exact analysis of the exposure time requirement for a given accuracy of either measurement type is dependent on many parameters of the measurement system including the seeing-limited PSF size, the shape of the sky background spectrum, and the particular design of the spectrograph. Let the telescope aperture area be denoted A , exposure time T , PSF solid angle P , sky background spectral flux density $\phi_b(\lambda)$ with overall brightness β visual magnitudes per square arcsecond, a

target with spectral flux density $\phi_t(\lambda)$ and overall brightness M visual magnitudes, and spectral resolution $\Delta\lambda$. The overall photometric SNR is given by:

$$\xi = \frac{T \cdot A \cdot 10^{-0.4M} \cdot \int_{\lambda} \phi_t(\lambda) d\lambda}{\sqrt{T \cdot A \cdot 10^{-0.4M} \cdot \int_{\lambda} \phi_t(\lambda) d\lambda + 10^{-0.4\beta} \cdot P \cdot T \cdot A \cdot \int_{\lambda} \phi_b(\lambda) d\lambda}} \quad (4.7)$$

By this convention, the flux densities ϕ_t and ψ_b integrate to the photoelectron count of 0th magnitude targets. In the high-SNR regime, the SNR is:

$$\xi = \sqrt{T \cdot A \cdot 10^{-0.4M} \int_{\lambda} \phi_t(\lambda) d\lambda} \quad (4.8)$$

while in the low-SNR regime, the rearranged expression (and recalling that both ϕ_t and ϕ_b integrate to the same quantity),

$$\xi = \frac{\sqrt{T \cdot A \cdot 10^{-0.4M} \int_{\lambda} \phi_t(\lambda) d\lambda}}{\sqrt{1 + P \cdot 10^{-0.4(\beta-M)}}} \quad (4.9)$$

highlights the interplay of the PSF solid angle and sky background brightness relative to target brightness.

For simplicity, this calculation assumes a slit spectrograph with uniform unity grating efficiency and a flat sky background spectrum. Let the slit be matched in width to the PSF and span H PSFs in height. The spectrographic measurement is made by subtracting an average sky background spectrum from the H regions surrounding the spectrum. The variance of the measured target spectrum photoelectron count in each wavelength bin is:

$$\begin{aligned} \text{var}(\phi(\lambda_i, z)\Delta\lambda_i) &= \text{var} \left(\int_{\lambda_i - \Delta\lambda/2}^{\lambda_i + \Delta\lambda/2} (T \cdot A \cdot 10^{-0.4M} \phi_t(\lambda) - T \cdot A \cdot 10^{-0.4\beta} \frac{1}{H} \sum_{i=1}^H P \phi_b(\lambda)) d\lambda \right) \\ &= (T \cdot A)^2 \cdot \left(10^{-0.8M} \int \phi_t(\lambda) + 10^{-0.8\beta} \frac{P^2}{H} \int \phi_b(\lambda) \right) \\ &= (T \cdot A \cdot 10^{-0.4M})^2 \cdot \left(\int \phi_t(\lambda) + 10^{-0.8(\beta-M)} \frac{P^2}{H} \int \phi_b(\lambda) \right) \end{aligned} \quad (4.10)$$

The full form of Equation 4.6 is now:

$$\sigma_z \approx \frac{\sqrt{\sum_i \text{var}(\phi(\lambda_i, z)\Delta\lambda_i)(\Delta z(\lambda_i, z)^2 - \delta z_0^2)}}{\sum_i \phi(\lambda_i, z)\Delta\lambda_i} \quad (4.11)$$

$$\approx \frac{1}{\sqrt{T \cdot A \cdot 10^{-0.4M}}} \frac{\sqrt{\sum_i \left(\phi_t(\lambda_i)\Delta\lambda_i + \frac{P^2}{H} 10^{-0.8(\beta-M)} \phi_b(\lambda_i)\Delta\lambda_i \right) (\Delta z(\lambda_i, z)^2 - \delta z_0^2)}}{\sum_i \phi_t(\lambda_i)\Delta\lambda_i}$$

Similar to the form of the SNR, in the high-SNR regime, the DCR accuracy is independent of sky background level, PSF size, and slit height, but in the low SNR regime these parameters are relevant. In the low-SNR regime, the DCR accuracy is nearly linear in PSF solid angle.


Given the many parameters for DCR accuracy, it is necessary to confine the analysis to a reference optical system. The reference telescope aperture is one square meter, the reference exposure time is one second, and the reference seeing disk is 2 arcseconds in diameter. The free parameter is the sky background expressed in visual magnitudes per square arcsecond, which varies between 17.0 VM/arcsec² at sea level and 18.5 VM/arcsec² at a high altitude dark-sky site⁶. The spectrograph slit is assumed to be matched in width to the target PSF and to span 100 PSFs in height. The relevant quantity is the required visual magnitude or SNR of the target required to achieve the desired accuracy at a specified spectral resolution. In the high-SNR regime, time-aperture product may be traded against target brightness to achieve the desired SNR. In the low-SNR regime, additional exposure time or aperture will be required to beat back the effects of sky background.

The result of this calculation for C, S, and Quaoar type spectra down to a target magnitude of $M = 18$ are shown in Figure 4-33 as a function of target magnitude and astrometric bare-CCD SNR for sites from sea level up to 4 km altitude for the reference telescope parameters at the 40 nm resolution required to achieve 25 mas astrometric accuracy. These calculations show that as one might expect, the problem of spectroscopic characterization is limited by exposure-aperture product at all elevation angles. While at the top of Mauna Kea, a 40 nm spectral resolution is sufficient to achieve milliarcsecond accuracy on a bright target at 50° elevation, a typical asteroid or KBO brightness of $M \approx 18$ would, under the dark sky conditions, necessitate an aperture-exposure of about 1600 m²-sec to achieve 25 mas accuracy to the 1 σ level. A correction to Gaia-level accuracy floor set by the 40 nm spectral resolution would require in excess of 40,000 m²-sec.


A similar calculation was performed for solar panel/Kapton mixture satellite spectra giving comparable results. Satellites in GEO are on the order of 12th magnitude in brightness while

⁶This assumption is justified with measurements at the Firepond Facility presented in Section 6.8.2.2.


geosynchronous debris objects may be as faint as 18th magnitude [131]. With a slit spectrograph, a 1-meter telescope would need less than a minute to measure the spectrum of an active satellite and approximately half an hour to measure the spectrum of a small debris fragment to correct its DCR to 25 mas. A 3-meter telescope would need approximately three minutes to make the same measurement. This may or may not be fast enough to observe a rapidly-tumbling piece of debris. The accuracy calculation for satellite measurements will be revisited in Chapter 6 when discussing the design and operation of the 16-inch spectrograph constructed to measure satellites' spectra for this thesis.




(a) Accuracy vs. SNR, Westford, MA




(b) Accuracy vs. Brightness, Westford, MA




(c) Accuracy vs. SNR, Socorro, NM




(d) Accuracy vs. Brightness, Socorro, NM




(e) Accuracy vs. SNR, Maui, HI



(f) Accuracy vs. Brightness, Maui, HI



(g) Accuracy vs. SNR, Mauna Kea, HI



(h) Accuracy vs. Brightness, Mauna Kea, HI

Figure 4-33: Accuracy Requirement for Spectroscopic Compensation of DCR Bias in Ground-Based Observations with a Bare-Silicon Detector. The requirement is formulated in terms of required SNR or apparent target V magnitude in a 1 square meter telescope at 1 second exposure with 2 arcsecond seeing at 40 nm resolution.

4.7 Chapter Summary

This chapter has addressed the first of the three original contributions of this thesis:

1. DCR bias is definitively shown to be the dominant long-timescale systematic error in astrometric observations reduced with modern star catalogs. The analysis includes Solar System bodies as well as Earth satellites and shows that narrow passbands are not always a sufficient mitigation strategy.

Evidence shown in Section 1.2 demonstrated that atmosphere-induced bias is the dominant error source in observations of GPS and GLONASS satellites, and that the effect of DCR bias is evident the difference between observations of neutral-colored C-type asteroids and red S-type asteroids. This chapter expanded on that analysis by demonstrating that the signature of DCR bias is evident in asteroid observations from all major contributors into the Minor Planet Center database, and that this signature remains even for sensors observing with narrow passband filters.

The limitations of the color-agnostic Farnocchia catalog corrections to observations reduced with older astrometric catalogs were demonstrated by the persistence of systematic target color-dependent errors confined to station-local vertical coordinates after application of catalog corrections. The positive common-mode error in residuals of both S-type and C-type asteroids, despite predictions that those errors should be negative, further demonstrated that the signature of *stellar*, as well as target, DCR bias is also a dominant driver of error in ground-based asteroid astrometry.

An attempt was made to overcome the limitations of after-the-fact catalog corrections of asteroid observations by reprocessing a subset of LINEAR raw images using the UCAC5 star catalog. Although the reprocessing achieved better astrometric precision than the original operational processing pipeline, and achieved better agreement with predictions than the MPC data set, the low number of observations on any particular asteroid made it difficult to draw definitive conclusions from this effort. A dedicated observing campaign during the 2017 flyby of the asteroid 3122 Florence with bare-CCD observations from New Mexico and Westford, supplemented with narrow-passband observing from the MIT Wallace Observatory generated over 250,000 distinct observations that provided a firm demonstration that narrow passbands alone do not always provide sufficient mitigation of DCR bias.

The limited efficacy of narrow passbands as a DCR mitigation strategy motivates the use of concurrent spectroscopy for per-target DCR compensation. Spectral resolution and time-aperture requirements were derived to enable spectroscopic compensation to a level of accuracy commensurate

with that of modern star catalogs. A 25 mas accuracy goal was chosen based on the worst-case predictability of stellar DCR based on cataloged quantities available in the UCAC5 and Gaia DR2 catalog. A 40 nm spectral resolution and a time-aperture product of 1 m²-sec is necessary to achieve this goal on a 6th magnitude target from near sea level or a 7th magnitude target observed from the top of Mauna Kea.

Chapter 5

Simulations of Orbit Solutions from DCR-Biased Observations

Having demonstrated that DCR bias is detectable in real observations of Solar System bodies and Earth satellites and is statistically significant, this chapter analyzes its effects on optics-only orbit solutions for those objects. The purpose of the analysis is to consider potential worst-case effects of uncompensated DCR bias on optics-only orbit solutions. The assumption is made that contributing sensors generate observations that are free from systematic errors that are endogenous to the sensor itself (such as a constant clock bias) and are subject only to short-timescale, zero-mean random error dominated by atmospheric seeing and a long-timescale, non-zero-mean DCR bias. Implicit in this assumption is that fixed-pattern catalog bias is negligible and that other error sources, such as pixel phasing, streaking, and jitter during the exposure are also zero-mean sources with short correlation timescales and that their effect on the orbit solution attenuates to zero quickly enough for them to be folded into the short-timescale random error component.

5.1 Inner Solar System Asteroids

5.1.1 Simulation Scenario

The simulated asteroid orbits enumerated in Table 5.1 were generated for the first part of this analysis. All orbits were generated such that the target was at opposition at the simulation epoch. The asteroids' inclinations with respect to the ecliptic were varied between zero and 15 degrees. For nonzero inclinations, trajectories were generated for all values of RAAN sampled at 30 degree

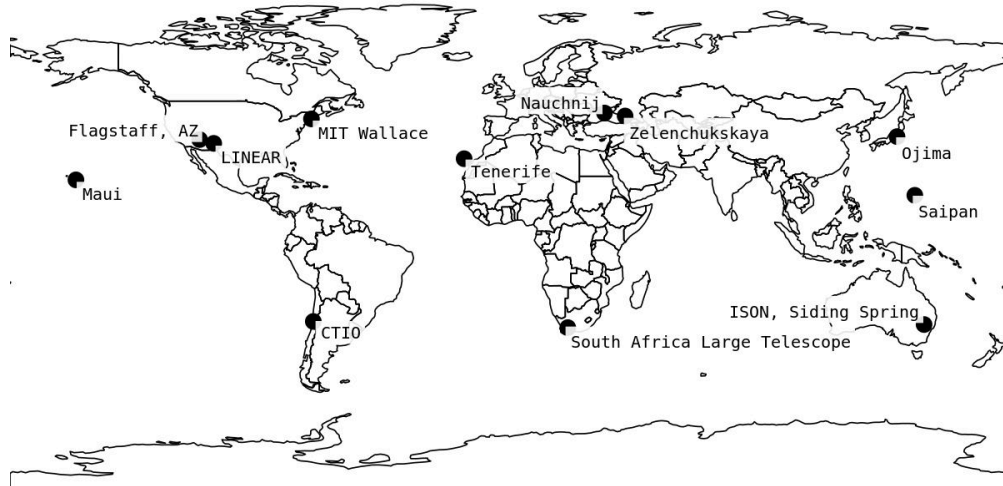


Figure 5-1: Simulated Observing Network

intervals. For eccentric orbits, the argument of perihelion ω was also varied. In Table 5.1, ω is specified in terms of the true anomaly at opposition ν . The purpose of this exercise was to sample the orbit space and identify worst-case effects of DCR bias on orbital elements. No effort was made to populate the simulation according to the actual population densities in these orbit classes or to capture all possible detection scenarios of near-Earth objects where detection near opposition may not be possible.

Observations of these orbits were simulated at one-hour intervals from the observing stations shown in Figure 5-1. The formal error on each asteroid observation was set to 0.25 arcsec in RA and Dec. This number was chosen as a conservative estimate folding in both short time-scale and medium timescale random error based on Carpino’s 2003 analysis and more recent work by Farnocchia showing sites that routinely outperform this level. DCR bias corresponding to the difference between an S-type and C-type asteroid was then injected into the observations along the local vertical at each observing site.

A Solar System force model incorporating all major gravitational perturbers and relativistic effects described by Moyer [95] is used for orbit propagation. Recall that the OrbitTool force model agrees with JPL Horizons to several meters over a decade (Figure 3-8). A Yarkovsky acceleration is optionally included in the force model in order to gauge the effect of DCR bias in astrometric observations on the estimation of non-gravitational acceleration (Section 5.2). No non-gravitational accelerations are included in the baseline truth orbits used in the simulation.

The observing arcs used in the simulation are single-apparition fits centered on the time of opposition and extend up to six months forward and backward in time. For some orbits, continuous

Table 5.1: Asteroid Family Parameters Used in Simulation

Class	Perihelion q range (AU)	Aphelion Q range (AU)	True Anomaly ν Range (deg)
Apollo	(0.5, 0.9)	(1.0, 1.3)	(150, 210)
Aten	period < 1 year	(1.1, 1.5)	(150, 210)
Amor	(1.0, 1.3)	(1.2, 2.4)	(0, 360)
Main Belt	(2.1, 3.3)	(2.1, 3.3)	(0, 360)

observations over the full arc are possible, while for others the actual data volume is reduced by geometric constraints. The constraints imposed on the observations are that data is only generated during astronomical twilight and that the target remain above the horizon. The cutoff for the minimum elevation angle of the target above the horizon is a parameter of the study. No modeling is made of reduced detection probability for weather outages. Modeling of reduced detection probability for small or distant targets is not performed for initial analysis of the effect of DCR on orbital parameters, but is performed for analysis of the DCR bias on optics-only measurement of non-gravitational acceleration.

5.1.2 Mean Motion Bias and Recovery Residual for Single-Apparition Fits

Orbits fit to the simulated biased observations were compared to the original simulated reference orbits in Keplerian coordinates. The single-apparition observing arcs centered on opposition in all of these scenarios are sufficiently long to constrain the orientation of the asteroids' orbits (inclination and RAAN) to near their true values even in the presence of observational bias. However, for all orbit classes sampled in this study, there exist instances where the DCR-induced bias in the estimate of Keplerian mean motion ($n = \sqrt{GM/a^3}$) differs from the mean motion of the simulated reference orbit by a value that is statistically significant with respect to the formal bounds on n attained by the orbit fit.

The mean motion of the fit to biased data can differ from the mean motion of the original orbit by as much as 10–15 σ with respect to its formal error (again, assuming underlying observation accuracy of 250 mas RMS). Figures 5-4(a) shows the mean motion bias of all simulated orbits broken down by orbit class and minimum allowable elevation limit of input observations. The deviation in the estimate of mean motion caused by DCR bias in the input observations is on the order of 0.5 arcsecond per year for Main Belt, Amor, and Apollo asteroids and a few arcseconds per year for Atens. Figure 5-4(b) shows this bias normalized to the formal uncertainty in mean motion given once-per-hour observations from the simulated network. It is important to note that the standard dictum to avoid observation at large zenith distances is not particularly helpful. While restricting

Table 5.2: Simulated Orbits with Large DCR-Induced Mean Motion Biases

Parameter	Main Belt	Aten	Amor	Apollo
Epoch (JD TDB)	2457012.50	2457103.5	2457012.5	2457012.5
a (AU)	2.9000000340533245e+00	9.0000000887734999e-01	1.1000000155675267e+00	1.0000000051667113e+00
e	1.3793102112057429e-01	2.2222221016656232e-01	4.5454530658963037e-02	4.999999224993313e-01
ω (deg)	2.6881621745621430e+02	1.8083676565900069e+02	2.0881621745621422e+02	2.6881621745621430e+02
i (deg)	1.4999999999999996e+01	4.999999999999947e+00	1.4999999999999996e+01	1.499999999999982e+01
Ω (deg)	2.1946739922783119e-16	1.8000000000000000e+02	5.999999999999993e+01	8.5486464551684185e-16
m_0 (deg)	1.799999999999997e+02	1.8000000000000000e+02	1.8000000000000011e+02	1.8000000000000000e+02
$\Delta n/\text{SNR } \text{el} > 20^\circ$ (mas/yr)	396 / 7.8	224 / 3.4	168 / 12.5	296 / 12.3
$\Delta n/\text{SNR } \text{el} > 30^\circ$ (mas/yr)	289 / 5.1	124 / 0.3	143 / 8.8	138 / 8.0
$\Delta n/\text{SNR } \text{el} > 40^\circ$ (mas/yr)	194 / 3.1	281 / 0.4	124 / 5.8	165 / 4.5

observations to high elevation does attenuate the effect of DCR bias on mean motion, there still exist cases for all orbit classes where an elevation cutoff of 40 degrees allows up to 5σ errors in mean motion.

5.1.3 Masked Signature of DCR

Given that DCR bias in observations can project into a large error in one orbital element, it is important to consider whether the bias will also manifest in the post-fit observation residuals in a detectable way. While in certain orbit geometries, DCR bias in observations can indeed bias the orbital solution away from its true value by a statistically significant amount, the signature of DCR bias is not evident in the post-fit residuals when the fit is performed using data from a single apparition only.

Table 5.2 lists the orbital parameters of four representative simulated asteroids with numerically large and statistically significant DCR-induced mean motion biases. Figure 5-2 shows the aggregate post-fit residuals for a single-apparition fit resolved into station-local horizontal and vertical coordinates and plotted versus observed elevation angle. Figure 5-3 shows representative residuals from a single station. The post-fit residuals in the horizontal coordinate are identically zero while the post-fit vertical residuals from nearly all stations are only a few milliarcseconds away from zero. Given the assumed formal error of the underlying observations, the DCR bias is not evident in post-fit residuals across a single apparition.

Although the post-fit residuals are zero-mean, the input bias may manifest as a larger than expected post-fit RMS. The contribution to the fit residuals of the DCR bias signal is under 50 mas in the worst case for nearly all fit simulations as shown in Figure 5-4(c). The resulting empirical fit residuals in the presence of noise would be the RSS of this quantity and the baseline per-observation accuracy (250 mas in this simulation) for about a combined 255 mas RMS error level in the presence

of systematic DCR bias and 250 mas RMS in the absence of systematic DCR bias. This is not a detectable signal given the heterogeneity of atmospheric conditions, sensor configurations, and noise levels encountered in real observations.

One common method to mitigate the effect of DCR bias on orbit solutions is to raise the elevation angle cutoff for accepting observations. This is, of course, impossible for ground-based observations of certain classes of asteroids. For orbits that are readily observable at low zenith distances, there is a direct trade between lowering data volumes to limit the effect of DCR bias and lower elevation cutoff to realize a \sqrt{N} improvement in the formal error bounds with more data. Figure 5-5 shows this trade for all simulated asteroids. The worst-case mean motion bias introduced by DCR bias drops nearly linearly with increasing elevation cutoff. However, the formal error increases as data volume decreases with higher elevation cutoffs. As Figures 5-5 and 5-6 show, in the case of large data volumes from the geographically-diverse set of ground stations used in this simulation, ignoring elevations and making the orbit fit to azimuth-only data yields marked improvement in the accuracy of the mean motion estimate for many asteroids in this simulation.

Although the bias in the orbital solution is hidden with respect to the original dataset from which it came, the fact that the bias accumulates in the mean motion coordinate implies that it will eventually build into a directly observable and statistically significant bias in Keplerian mean anomaly over time. In that vein, consider what happens on the next apparition of the asteroid (or alternatively on precovery observations from the prior apparition). The relevant quantity is the recovery or precovery residual in the along-track coordinate of the observation which for most orbit geometries coincides with the orbital velocity vector (where the mean motion bias accumulates). Panels (a) and (b) in Figure 5-7 show the in-track component of the DCR-induced prediction error on recovery/precovery for observations for each asteroid on its next apparition normalized to the formal error in that direction. The variation in the orbit geometries used in this study means that there is not a one-to-one mapping between normalized mean motion bias and normalized in-track residual on the next apparition, but there do exist cases where the recovery residual is statistically significant to a 5σ level at even conservative elevation cutoffs. The last four panels of Figure 5-7 show the numerical value of the recovery residual at the next opposition against its value normalized to its uncertainty with statistically significant examples evident even at high elevation cutoffs. None of these prediction errors make the difference between recovery and a lost asteroid, but they can be large enough in numerical value to be noticeable above the natural noise in the astrometric measurement, but only on subsequent apparitions.

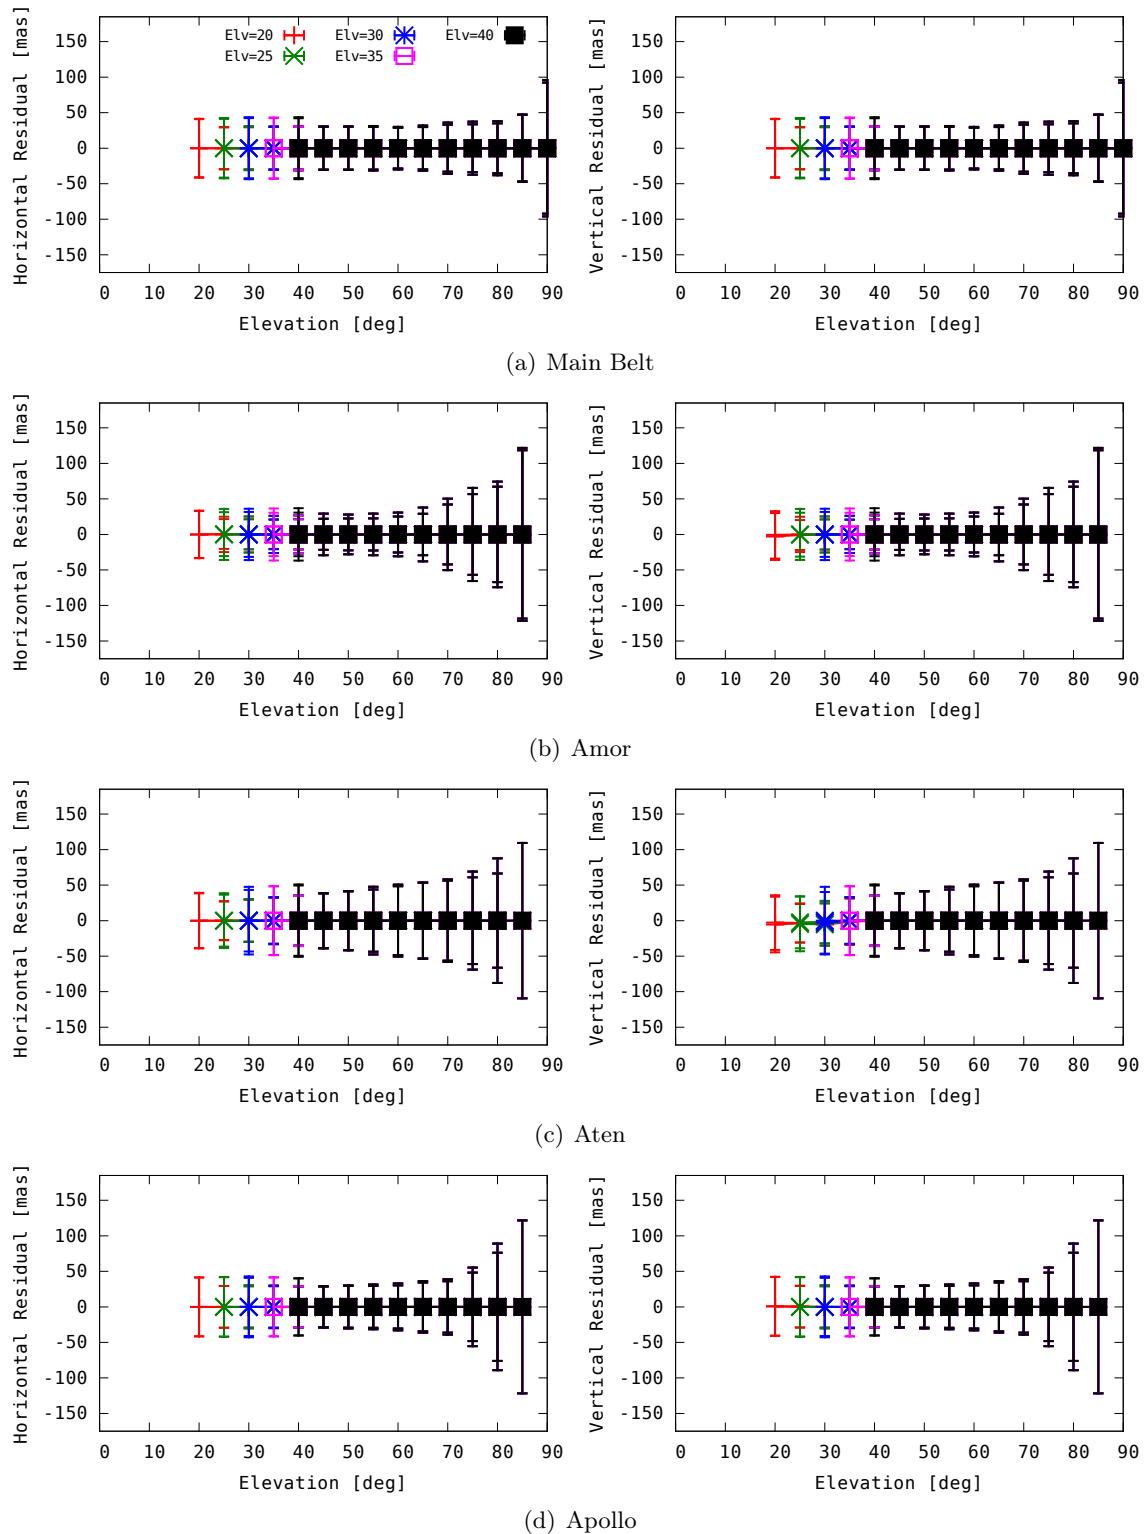


Figure 5-2: Distribution Statistics of Simulated Post-Fit Residuals of Orbit Class-Representative Asteroids for Single-Apparition Fit. Mean $\pm 3\sigma$ of distributions of aggregate horizontal and vertical single-apparition post-fit residuals of Main Belt, Amor, Aten, and Apollo asteroid examples. Post-fit residuals for fits with different allowable lower elevation limits are shown as separate traces. The distributions of post-fit residuals from all individual stations are all zero-mean, meaning the signature of DCR bias is not evident in the post-fit residuals across a single apparition.

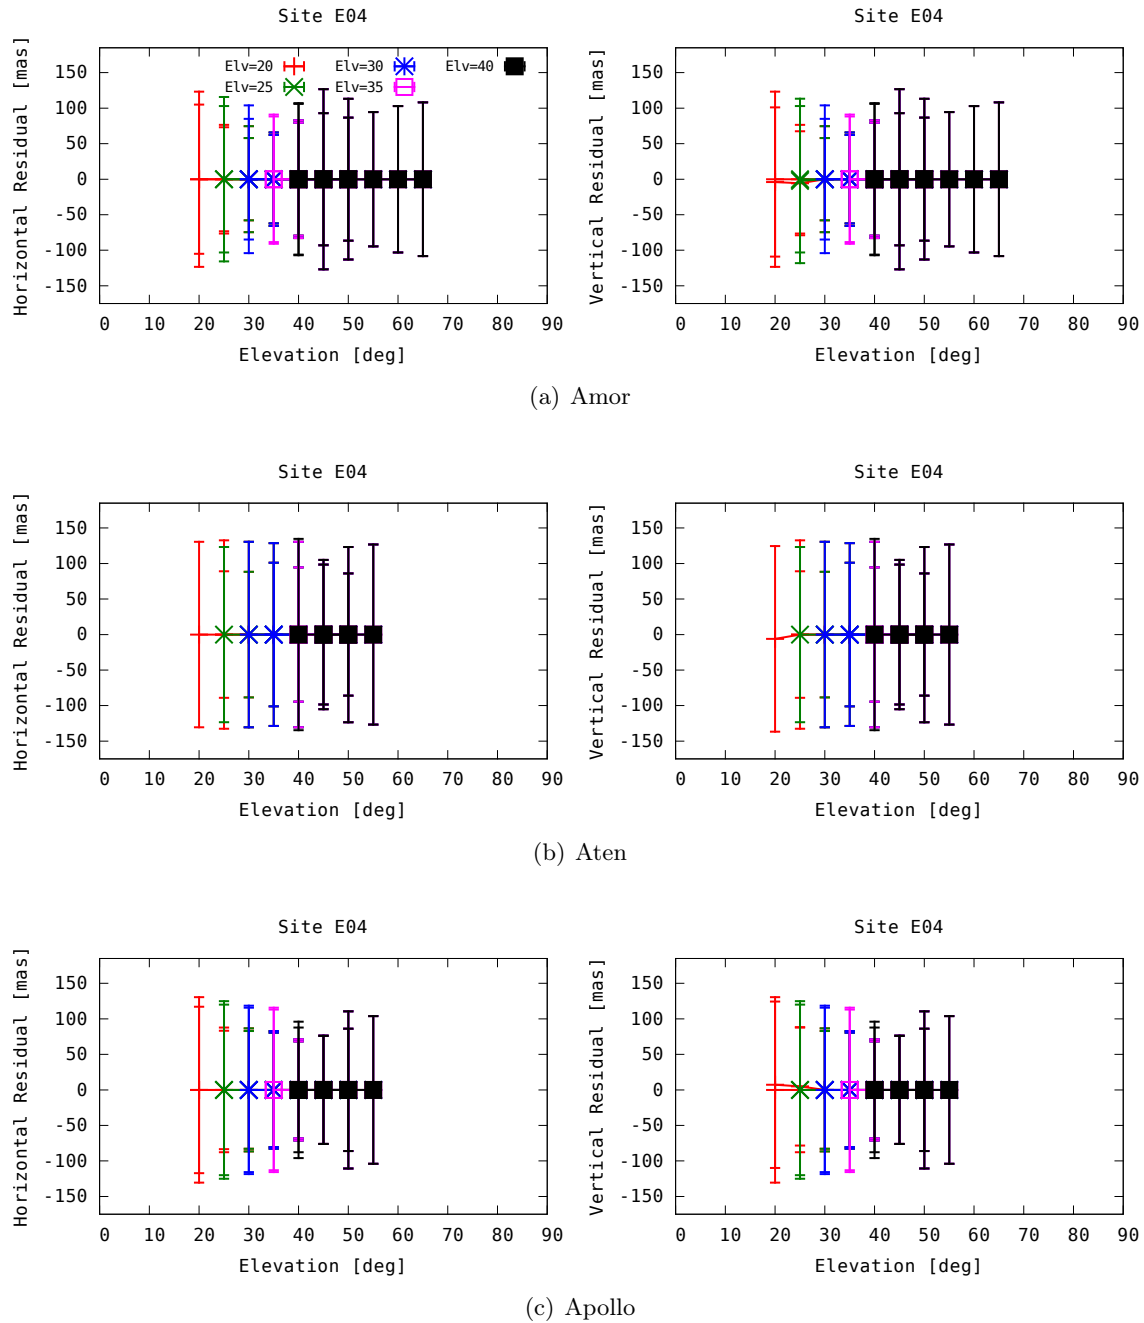
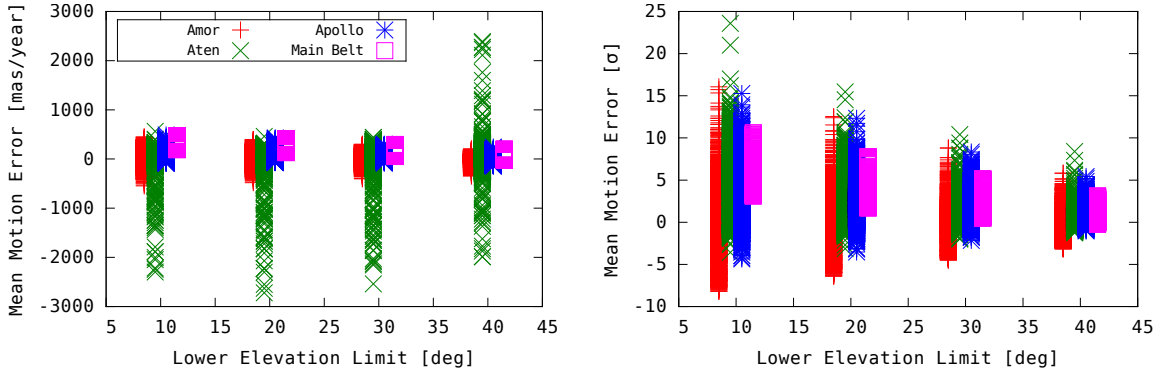
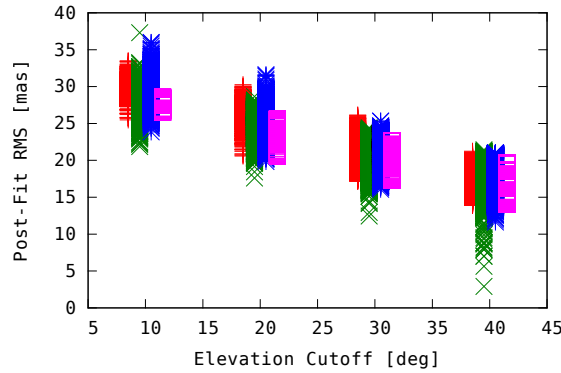


Figure 5-3: Distribution Statistics from a Single Station of Simulated Post-Fit Residuals. Representative mean $\pm 3\sigma$ of distributions of aggregate horizontal and vertical single-apparition post-fit residuals from a individual station of Amor, Aten, and Apollo class asteroid examples with non zero-mean vertical error distributions. While the distributions of post-fit residuals in the aggregate set of all observations are zero mean, distributions from some individual sites are not zero mean in the vertical coordinate. A non-zero mean distribution only occurs for low elevation observations and only for the lowest allowable elevation limit. Given the formal error of the individual observations (250 mas RMS), even with several hundred observations per elevation bin, the mean of the distribution is statistically indistinguishable from zero and DCR bias is not evident in the post-fit residuals from any one station across a single apparition orbit solution, nor is it evident in the aggregate of residuals from all stations (Figure 5-2).



(a) Absolute mean motion bias for all simulated orbits (b) Mean motion bias normalized to its standard error



(c) Single apparition fit RMS to noiseless biased observations. The RSS of this quantity with the per-observation RMS gives the post-fit RMS to real noisy data.

Figure 5-4: Single Apparition Fit Mean Motion Bias and Postfit Residuals for Asteroid Orbits with DCR Bias. Progressively lower allowable elevation angles simultaneously subject the mean motion to greater DCR-induced bias while lowering the formal error by allowing more of the arc to be observed and by supplying more data to constrain the fit. The absolute values of the DCR-induced mean motion bias are less than an arcsecond per year in the worst case at low elevation cutoff, but this quantity is nevertheless statistically significant. It is not possible to deduce the presence of such a bias from the post-fit residuals across only a single apparition. (c) For a single-apparition fit, the contribution of DCR bias to the fit residual across a single apparition is not expected to be detectable above the background per-observation RMS of 250 mas.

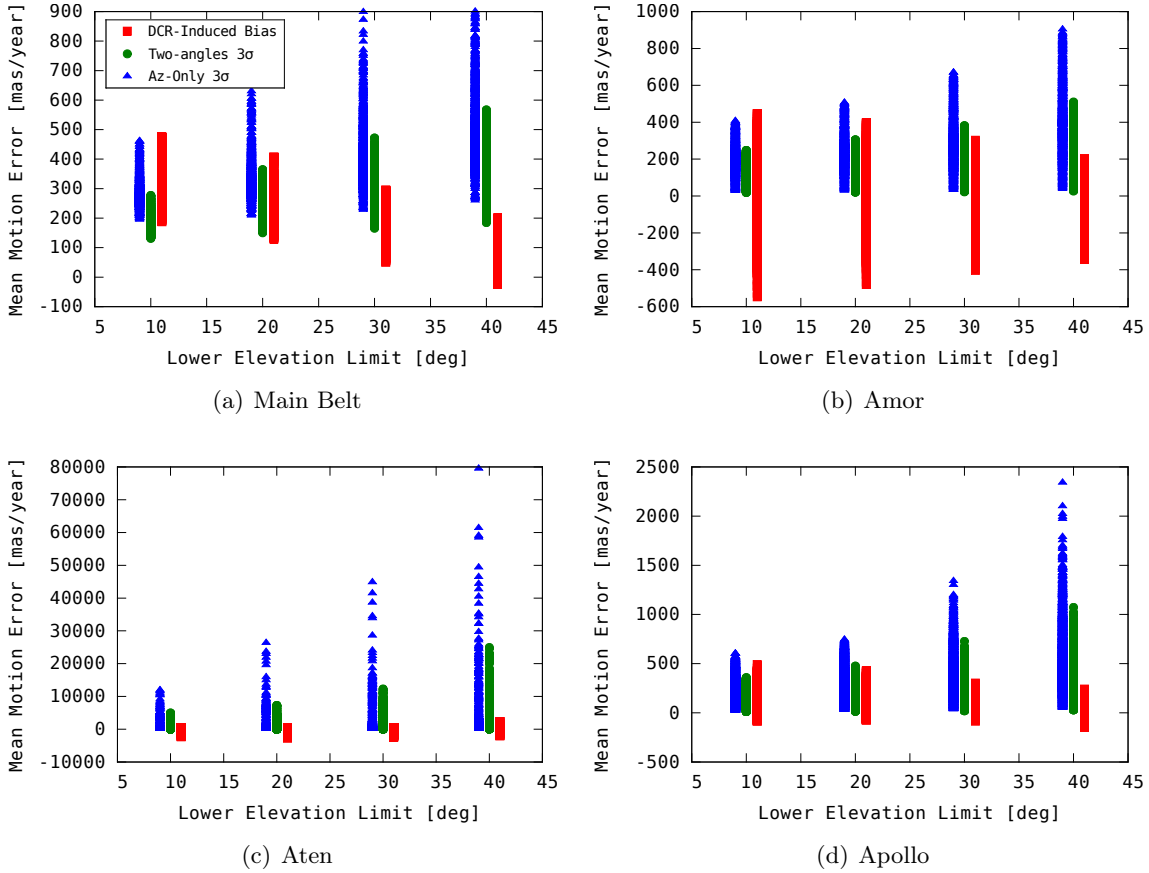


Figure 5-5: Trade-off in Accuracy and Precision of Mean Motion Estimate in a Single-Apparition Fit. Raising the lower limit on the allowable elevation of input observations reduces the expected error in mean motion caused by DCR bias at the expense of reduced precision in the mean motion coordinate. The trade space can be improved by using only the bias-free azimuth component of each astrometric observation at a lower elevation cutoff for certain orbits, but the feasibility of an azimuth-only fit depends on having a sufficient geographic diversity of observing stations. Figure 5-6 more clearly shows individual scenarios where better estimates of mean motion are obtained by ignoring elevation and performing an azimuth-only fit.

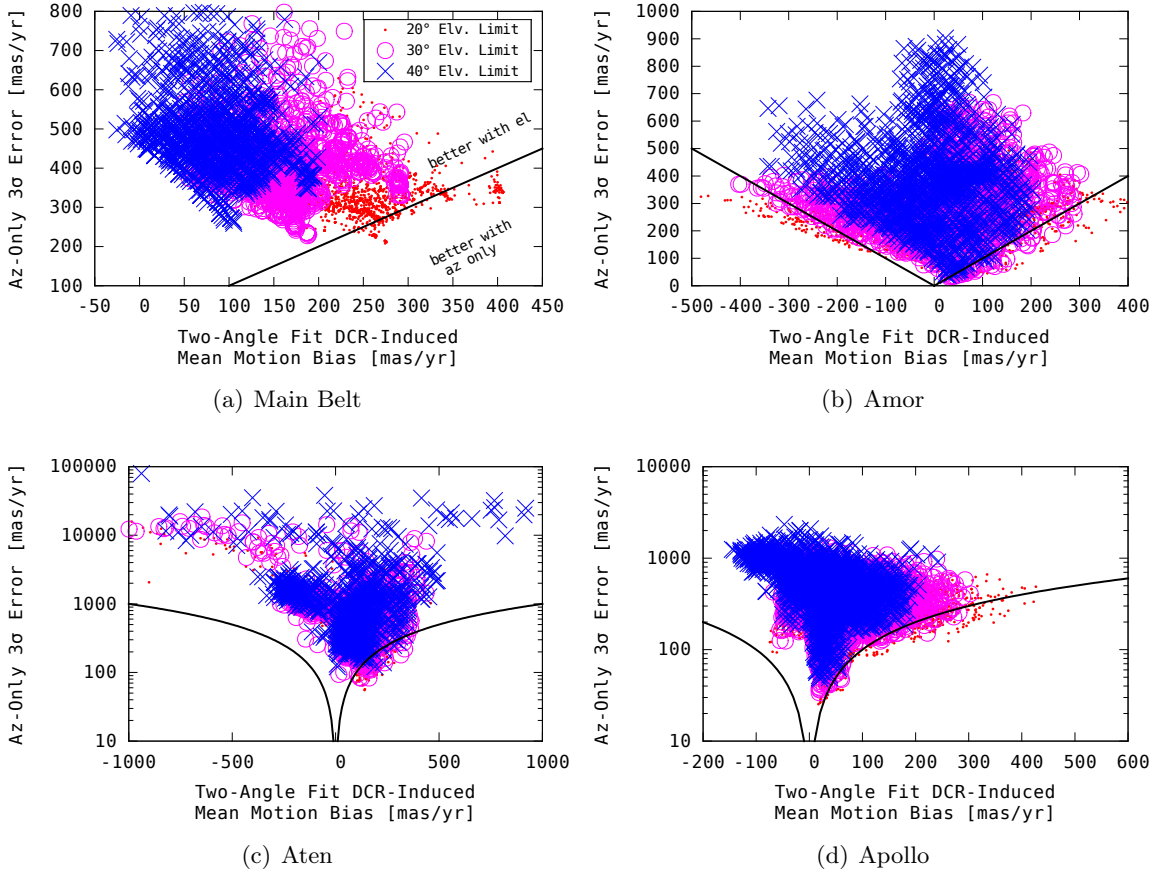


Figure 5-6: Trade-off in Accuracy of Mean Motion Estimate by Discarding Elevation Data. For certain orbits, the effect of DCR, even with a conservative lower elevation limit of 40 degrees, can be such that it pays to discard the vertical component of astrometry and compute fits using the horizontal measurement only.

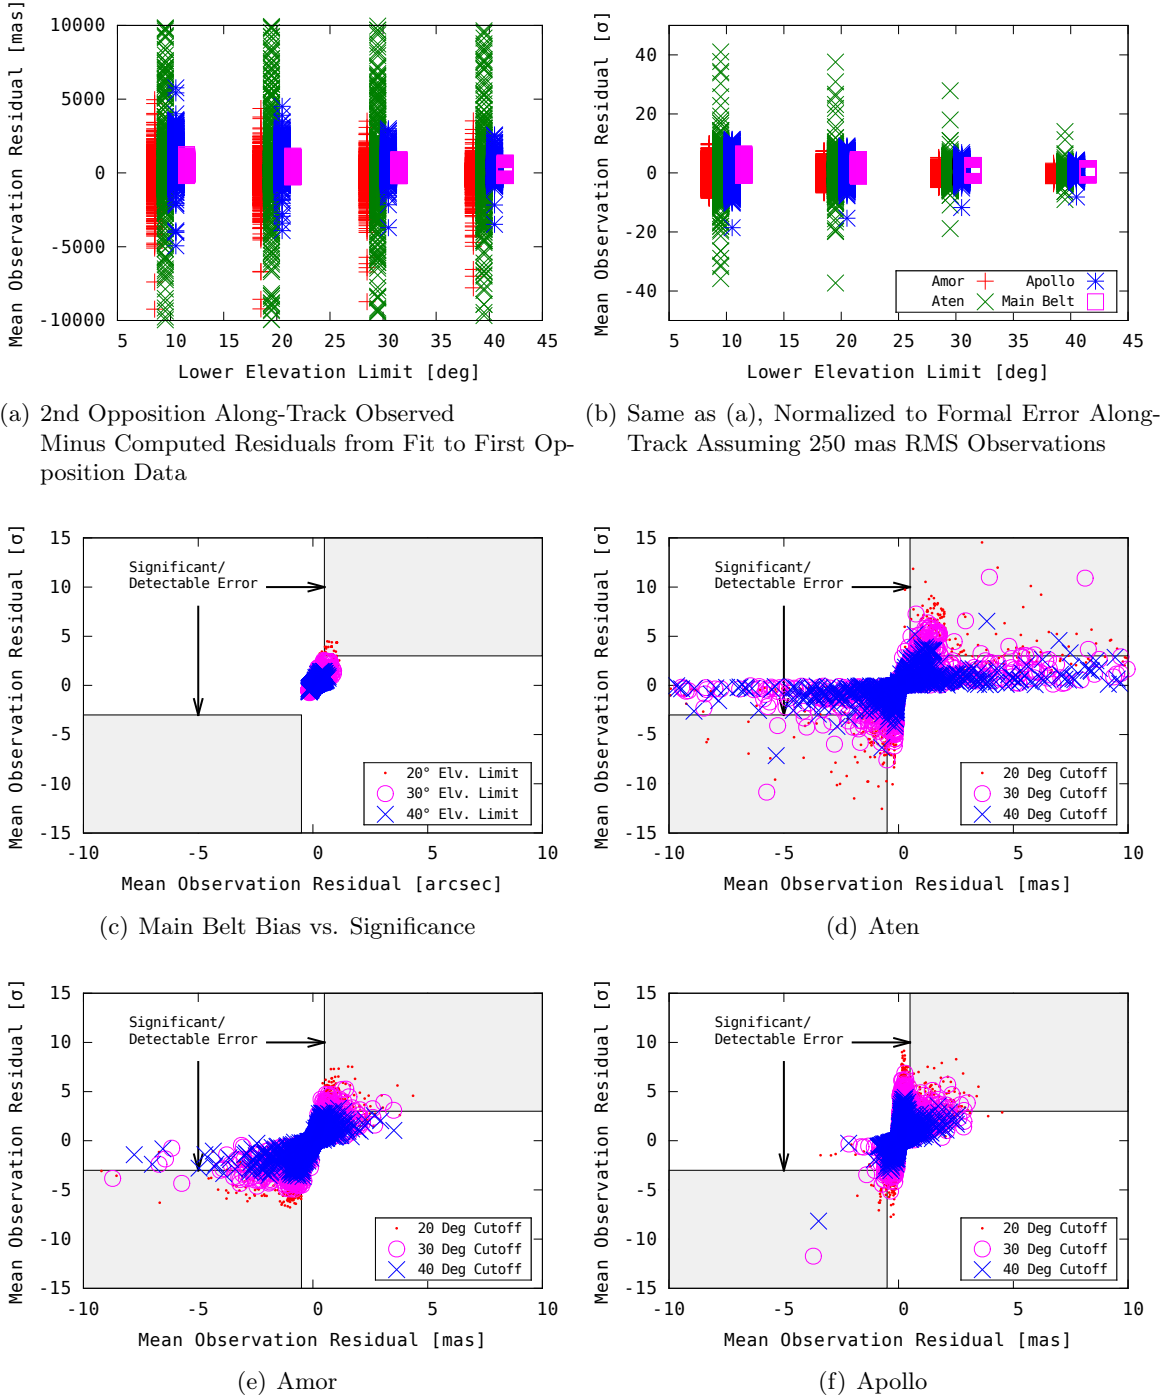


Figure 5-7: DCR-Induced Bias in In-Track Component of Recovery/Precovery Residual from a Single-Apparition Optics-Only Fit vs. SNR of Bias. For all orbit classes, a non-negligible bias of several tenths of an arcsecond per year can occur and exceed the formal error bounds on mean motion by more than 5σ , even when observations are constrained by the conservative elevation limit of 40 degrees. This builds into a statistically significant recovery residual aligned with the direction of the orbital velocity of the target. The single-apparition orbit solution is sufficiently well-constrained that this bias never makes the difference between recovery and loss of target, but in many instances in our sample of simulated orbits, the mean motion bias builds into a sufficiently large value to rise above the per-observation noise floor on the next apparition.




Figure 5-8: Orbits with Statistically Significant Mean Motion Bias vs. Inclination, Semi-Major Axis, and Observation Lower Elevation Limit

5.1.4 Exploration of the Orbital Element Space

I search the orbital element space to find regions where the DCR-induced mean motion bias exceeds the formal uncertainty in mean motion by a large amount. The search is parametrized over elevation cutoffs. Results in semi-major axis/inclination space are shown in Figure 5-8.

For the realistic but optimistic elevation cutoff of 20° , almost all examined orbits showed a 5σ bias in mean motion and several showed a 10 or 15σ bias. Raising the elevation cutoff to a more conservative 30° , nearly all orbits with inclination above 10° retain at least at 5σ mean motion bias while some orbits with a semi-major axis close to 1 AU have at least a 10σ mean motion bias.

The ground network used in these simulations is biased slightly toward having more Northern Hemisphere sites. The effect of this geographic bias was explored by simulating the Aten and Apollo orbits with initial opposition centered at the equinoxes and solstices. The effect of season on the 15° inclined asteroids is shown in Figure 5-9, where significant biases are parametrized by RAAN offset from opposition for each season. While there is little seasonal variation at a 20 degree elevation cutoff, there are some seasonal asymmetries evident with a 30 degree elevation cutoff that can be attributed to the effect of the geographic bias of the observing network.

The statistically significant bias points indicate that there is sufficient biased data that the formal error is smaller than the accumulated mean motion bias. Absence of an indication of statistically significant bias on these plots does not indicate that the DCR bias attenuates for that particular orbit. Rather it is an indication that there is insufficient data to constrain mean motion to a level of precision where DCR bias would be evident. Increased data volume for the same orbit arc or




Figure 5-9: Orbits with Statistically Significant Mean Motion Bias vs. Semi-Major Axis, Offset Between Ascending Node and Opposition, Season of Opposition, and Observation Lower Elevation Limit. Instances of Aten and Apollo orbits with 15° inclination where the DCR-induced mean motion bias in a single-opposition fit exceeds 5σ . The slight seasonal asymmetry reflects the Northern Hemisphere bias of the simulated observing network.

improved accuracy of each individual measurement would, for some of these orbits, bring the formal error down to a point where DCR-induced mean motion error would be significant.

5.2 Spurious Detection of Yarkovsky Acceleration

5.2.1 Sensitivity of Yarkovsky Term to Observation Bias

The DCR-induced apparent mean motion bias in the single-apparition orbit fits described in Section 5.1.2 invites the question of how the estimated mean motion of the fit to biased data evolves with the incorporation of observations from subsequent apparitions. Although in a single-apparition fit the signature of DCR bias is not detectable in the post-fit residuals, I have shown examples where on subsequent apparitions the in-track residual of a fit to data from a single apparition is detectable and statistically significant. The alignment of this prediction error with the direction of the asteroid's orbital velocity implies that a fit incorporating data from both apparitions will reduce that error by changing the state estimate in both mean motion (semi-major axis) and mean anomaly at epoch. This effect, caused entirely by observational bias, would look like the mean motion changing over time, which is consistent with a Yarkovsky acceleration acting on the object.

To gauge the magnitude of a potential spurious Yarkovsky detection, a two-apparition case was considered. A zero-inclination Amor asteroid with $q = 1.3$ and $Q = 2.1$ was simulated over two successive apparitions. Observations were simulated for both apparitions but the relative weighting




Figure 5-10: Spurious Yarkovsky Signal in Two-Apparition Optics-Only Fit. $q = 1.3, Q = 2.1, i = 0$ orbit observed from simulated network with 250 mas RMS observations once per hour when above elevation limit and in astronomical twilight conditions.

of the two apparitions was modulated to gauge the sensitivity of best-fit normalized Yarkovsky acceleration parameter A_2 to observation bias.

In Figure 5-10, the results of an orbit fit to a full-fidelity dynamical model that includes a Yarkovsky acceleration as a free parameter is shown for three sets of weightings of the two apparitions:

1. A short 30-day data arc on the first apparition followed by a long (up to 360 day) data arc;
2. A long data arc followed by a short data arc;
3. Two long data arcs.

Although the magnitude of the spurious Yarkovsky acceleration for this scenario is non-physical (see Section 5.2.2), and no one would try to make a Yarkovsky measurement across only two apparitions, this result nevertheless gives a reason to continue this line of investigation. The fact that inverting the weighting for the two apparitions inverts the sign of the best-fit Yarkovsky term for the orbit solution implies there can be a strong sensitivity of the Yarkovsky estimate to observational bias. The fact that the SNR of the spurious Yarkovsky detection remains above unity when fitting across all the data from both apparitions implies that increased arc length across more apparitions may not attenuate out the effect of observation bias, and there may be realistic scenarios where a spurious Yarkovsky detection is possible. In the remainder of this section, I examine whether there exist circumstances under which DCR bias could appear to be a physically-plausible Yarkovsky acceleration acting on the target.




Figure 5-11: Yarkovsky Effect Magnitudes of Asteroids in the JPL Database vs. Data Arc. Data retrieved December 2018. Orbital period expressed in years and in orbits. While many of the asteroids with published Yarkovsky accelerations have radar-constrained orbits or data arcs spanning many decades, several optics-only short-arc objects exist in this dataset, the shortest arc being 2011 PT with a three-apparition optics-only fit spanning six years. While the Yarkovsky acceleration measured on 2011 PT is not an artifact of observation bias, its 1.5 year period orbit forms the basis of the subsequent analysis of the potential worst-case effect of DCR bias on the measurement of Yarkovsky accelerations in asteroids with short-arc optics-only orbit solutions.

5.2.2 Physically-Plausible Yarkovsky Accelerations

While the analysis is purely astrometric, the physical plausibility of a Yarkovsky acceleration depends on non-astrometric quantities. Large targets (as indicated by photometric size estimates) would indicate that a meaningful Yarkovsky effect is unlikely while a noticeable variation in mean motion over time in a small, faint target across more than two apparitions would make the inclusion of a Yarkovsky term plausible. The relationship between physically plausible Yarkovsky accelerations, target detection opportunity, and the degree to which a Yarkovsky acceleration can be constrained with optics-only measurements restricts the possibilities over which a short arc optics-only orbit composed of data from a few apparitions can yield a spurious Yarkovsky detection. Within these constraints, there exist scenarios where along-track observation bias can look like a physically-plausible Yarkovsky effect.

Before presenting such hypothetical scenarios, it is necessary to note that short-arc optics-only detections of Yarkovsky acceleration are realistic. As of December 2018, the JPL Small-Body Database contains 73 objects with defined non-gravitational terms in their force models. Of these, 38 are objects whose orbit solution is derived entirely from optical observations. The A_2 term and its uncertainty for these objects is plotted as a function of data arc length in Figure 5-11. The

largest A_2 for a short-arc optics-only object is 2011 PT, which was observed on three consecutive apparitions over six years, totaling four of its orbits, with an orbit solution derived from 205 distinct observations spread over those three apparitions. Those 205 distinct observations over six years and three apparitions yield a formal uncertainty on A_2 of 3×10^{-14} AU/day².

The orbit of 2011 PT places opposition at a distance of about 0.01 AU. This is a restrictive criterion and it is necessary to bound the volume in space over which larger but more distant objects may have similar detection opportunities while still having plausible Yarkovsky effects. Following the analysis in [23, 132], I assume that the transverse Yarkovsky acceleration follows the inverse square law and that for an object with albedo A :

$$A_2 \approx \frac{4/9}{1-A} \Phi(1\text{AU}) f(\theta) \cos \gamma \quad (5.1)$$

where for a roughly spherical object with density ρ and diameter D ,

$$\Phi(1\text{AU}) = \frac{1367\text{W}}{c} \frac{3}{2\rho D} \quad (5.2)$$

and that $\cos \gamma$ and $f(\theta)$ can take on their maximum plausible values of 1 and approximately 0.2, respectively.

Assuming purely diffuse scattering, the “visual magnitude” M of that same spherical target at a heliocentric distance r_h , geocentric distance r_g and phase angle ϕ can be roughly given as:

$$M \approx 6.8 - 2.5 \log_{10}(A) - 5 \log_{10}\left(\frac{D}{1 \text{ m}}\right) + 5 \log_{10}\left(\frac{r_g}{10^3 \text{ km}}\right) + 5 \log_{10}\left(\frac{r_h}{1 \text{ AU}}\right) - 2.5 \log_{10}\left(\frac{1}{2} + \frac{1}{2} \cos \phi\right) \quad (5.3)$$

which implicitly assumes an approximate bare CCD photometric zero point.

Plotting the maximum plausible A_2 versus brightest plausible magnitude near opposition for a variety of object diameters from 20 meters up to 10 km and a variety of opposition distances from 0.01 to 1.2 AU in Figure 5-12, it is evident that telescopes with realistic limiting magnitudes around $M = 18$ to $M = 20$ can detect objects with physically plausible Yarkovsky accelerations well in excess of the formal uncertainty in A_2 achievable on a three-apparition fit, and that detection distances as far as 0.6 AU near opposition are consistent with A_2 values above the sensitivity floor of an optics-only orbit fit.

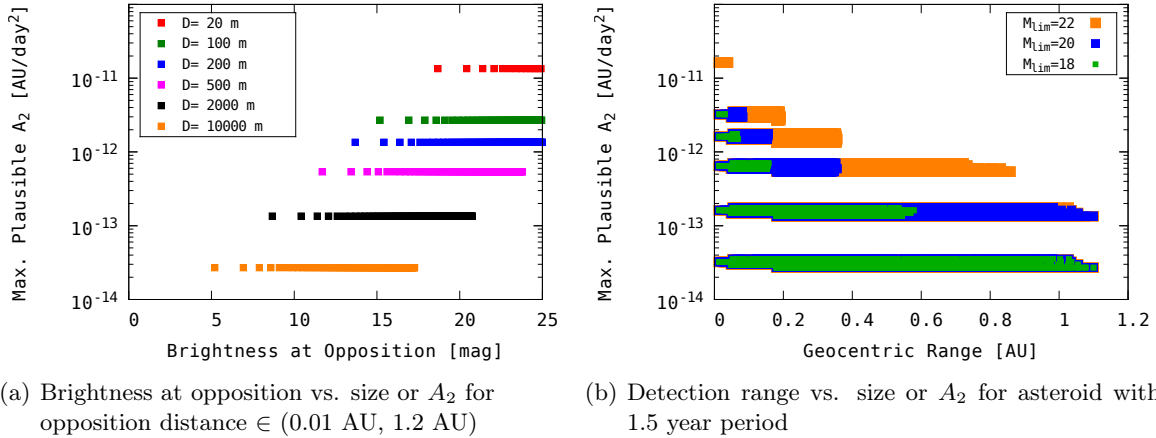


Figure 5-12: Maximum Plausible Yarkovsky Acceleration vs. Brightness and Physical Size. The maximum plausible Yarkovsky acceleration expressed as the normalized A_2 parameter is a function only of the physical size, density, and albedo of an asteroid. Assuming a constant albedo of 0.1 and a density of 2.5 g/cc, the brightness of the target at opposition is plotted in the left panel for geocentric distances ranging from 0.01 to 1.2 AU. Yarkovsky accelerations above the noise floor of a three-apparition optics-only orbit solution like that of 2011 PT are very readily detectable with modest-sized ground-based telescopes. The right panel examines specifically the case of a hypothetical asteroid with the same physical parameters and size ranges as in the left panel and in a 1.5 year period, zero inclination orbit with eccentricity between 0.1 and 0.4. In this case, I plot maximum detection range for different limiting magnitudes accounting for the effects of variable heliocentric and geocentric range and phase. For each asteroid size, orbits with different eccentricities are plotted as adjacent horizontal bands. There is little difference in the maximum detection range among the orbits, and little evidence that phase angle is significant in determining detection range near opposition, implying that detection range is a valid selection criterion for realistic simulations of observations of asteroids in a 1.5 year period orbit. Furthermore, a modest-sized telescope with a limiting magnitude around $M = 20$ is potentially capable of observing objects in this orbit with a large physically-plausible Yarkovsky effect.

5.2.3 Existence of Orbits Susceptible to Spurious Yarkovsky Detections

Having shown that there exists a large volume in space over which a newly detected object can plausibly have a large-enough physically plausible Yarkovsky acceleration to be detected with optics-only measurement over several apparitions, the task is now to demonstrate that uncompensated DCR bias in observations can project into a statistically significant but spurious Yarkovsky acceleration, and that the post-fit residuals of a dynamic model that includes the Yarkovsky term are better than the post fit residuals of a model with only gravitational accelerations. The search for these orbits was performed as follows.

Taking 2011 PT as the model, I considered similar orbits (those with a 1.5 year period and three distinct apparitions over six years) and searched for geometries where the DCR-induced observation bias over three apparitions looks like a Yarkovsky term. Specifically, with the orbital period fixed

to 1.5 years, I varied the eccentricity, argument of perigee, inclination (from zero to 15°), and for inclined cases sampled the RAAN in 60° increments. A three apparition orbit fit was simulated for asteroids in these orbits with observations subjected to uncompensated DCR bias from the observing network in Figure 5-1. DCR bias equivalent to the difference between a C-type and S-type reflectance was injected into the observations. From these synthetic observations, the weighted linear least squares equations for a dynamic model that incorporates a Yarkovsky term were computed. Each observation's coordinates were rotated to be aligned with the local horizon rather than inertial axes, giving the standard weighted least squares formulation:

$$\mathbf{W}^{1/2} \mathbf{A} \begin{bmatrix} \Delta X \\ \Delta Y \\ \Delta Z \\ \Delta V_x \\ \Delta V_y \\ \Delta V_z \\ \Delta A_2 \end{bmatrix} = \mathbf{W}^{1/2} \begin{bmatrix} \Delta A_{z_1} \\ \Delta \text{Elv}_0 \\ \vdots \end{bmatrix} \quad (5.4)$$

The bottom row of the inverse $(\mathbf{A}^T \mathbf{W} \mathbf{A})^{-1} \mathbf{A}^T \mathbf{W}$ is the sensitivity of the computed A_2 term to observation error and alternate entries in the bottom row are the sensitivity to horizontally-aligned and vertically-aligned errors, respectively. If all observation errors are drawn from zero-mean distributions, then the expected error in A_2 (and all other state components) is zero, but if the vertical component of the observation error is not zero mean, then there may exist scenarios where the cumulative effect of that unaccounted-for systematic vertical bias across all observations may yield a nonzero Yarkovsky signal whose formal uncertainty is lower than its amplitude.

To mimic realistic observation volumes for short-arc, optics-only asteroids with known Yarkovsky accelerations, I selected small numbers of observations on each apparition. The sensitivity vector is used as the selection heuristic for observations. First the sensitivity vector in the alternate entries of the bottom row of $(\mathbf{A}^T \mathbf{W} \mathbf{A})^{-1} \mathbf{A}^T \mathbf{W}$ is computed for all possible observation opportunities from all possible sites, simulated at one hour intervals and subject only to minimum elevation, maximum range, and local twilight constraints. The resulting sensitivity vector is then sorted according to the numerical value of the sensitivity of the best-fit A_2 term to vertical bias. The selection criterion is then applied to include the observations that project DCR bias most strongly onto A_2 .

Table 5.3: Simulated Orbits with DCR-Induced Spurious Yarkovsky Detections

Parameter	Case 1	Case 2	Case 3	Case 4
Epoch (JD TDB)	2457233.5003724921982666			
a (AU)	1.3103707192196954			
e	+2.2894285169462208E-01	+2.2894286136384409E-01	+1.5262857228421725E-01	+2.2894285720538432E-01
ω (deg)	+1.8633353675420889E+02	+1.8633353859638953E+02	+1.2633353888447590E+02	+3.0633353003450515E+02
i (deg)	+4.9999936443468860E+00	+9.9999988964569884E+00	+1.5000002982413168E+01	+1.5000002357916664E+01
Ω (deg)	+1.1999998425981468E+02	+1.1999999741181772E+02	+1.1999999947616297E+02	+6.0000001352585386E+01
m_0 (deg)	+2.5349776904368387E-04	+2.4474150550803629E-04	+4.5719879697966910E+01	+3.2077194090149857E+02
R_{\max}	0.3 AU			
Elv_{\min}	20 deg			
A_2 (AU/day ²)	2.22×10^{-13}	9.36×10^{-14}	-5.48×10^{-14}	-5.39×10^{-14}
A_2 1- σ	5.68×10^{-14}	2.86×10^{-14}	1.40×10^{-14}	1.50×10^{-14}
A_2 SNR	3.91	3.27	3.91	3.60

The contribution of each individual observation to the magnitude of the spurious A_2 depends not just on the geometry of that one observation but also on the entire ensemble of observations used in the fit. That is, by virtue of selecting the observations with greatest effect on A_2 when the sensitivity is computed with respect to all possible observation opportunities, the value of that sensitivity in the reduced set of observations changes. As such, the following selection rule was used: 10 observations per apparition with the most positive or most negative influence on A_2 and 630 observations with influence closest to zero, for a total observation volume of 1920 distinct observations over three apparitions. There is no particular justification for this selection strategy and it is not claimed to be optimal; it was a reasonable heuristic that yielded the necessary results with an observation volume that is within an order of magnitude of observation volumes on real asteroids. The entire ensemble of sensitivities from which the selection criterion is applied is computed across all three apparitions and constrained by a 20° lower elevation limit and maximum detection range of 0.3, 0.6, and 1.2 AU. I make no claim of the optimality of this strategy, only that it yields instances where a physically-plausible A_2 exists in the fit to biased data. Orbit fits were cross-checked against OrbFit 5.0.5 [133], and four of the orbits with Yarkovsky SNR > 3 are listed in Table 5.3.

The spurious Yarkovsky detections are not entirely dependent on the inclusion of low elevation observations in the data arc. Figure 5-13 shows the effect on best-fit A_2 and its formal error of an increased elevation cutoff for the four cases. While the magnitude of the best-fit Yarkovsky acceleration and its SNR decreases with the reduced data volumes associated with higher elevation cutoffs, the SNR (assuming 250 mas RMS input observations) remains above unity for all four scenarios up to a 40° elevation limit. For Case 1, in particular, the best-fit value of A_2 is identical at both 20° and 30° elevation cutoffs, and a detection is still present at 30° with a reduced SNR of 3.05. This implies that for higher-quality astrometry, it is possible for DCR-induced spurious detections





Figure 5-13: Evolution of Best-Fit Yarkovsky Accelerations and Their 1σ Uncertainties with Increased Lower Elevation Limit for the Example Cases in Table 5.3. Physically plausible values of A_2 exist at elevation limits up to 40° . This implies that it is not always just low-elevation observations where systematic bias projects into apparent Yarkovsky acceleration and that even with conservative elevation limits, a spurious Yarkovsky detection may occur in a three-opposition optics-only fit. While the statistical significance of the Yarkovsky signal at 40° elevation is below an SNR of 3, the formal errors are based on the assumption of 250 mas RMS formal error on input astrometry. Thus, improved astrometric accuracy or increased data volume may result in statistically significant detections even at the 40° limit.

of non-gravitational acceleration to exist even with seemingly conservative elevation limits.


While the best-fit Yarkovsky terms for Cases 2 through 4 are small in magnitude relative to the short-arc optics-only Yarkovsky measurements in the JPL database, the magnitude of the best-fit Yarkovsky term for Case 1 is comparable to that of 2011 PT and can be considered an existence proof for DCR-induced spurious Yarkovsky detections. The marginal SNR of Case 1 is predicated on the assumption that underlying noise in the optical observations is drawn from a zero-mean Gaussian error distribution with $\sigma = 250$ mas. Higher-quality observations, such as those realized in the reprocessing of LINEAR data, would increase this SNR. Furthermore, higher quality observations would raise the apparent SNR on a large number of marginal spurious detections found in the search.

5.2.4 Post-Fit Residuals of Orbit Fits with (Spurious) Yarkovsky Acceleration

Part of the argument for the plausibility of inclusion of a non-gravitational term in the force model for the asteroid orbits simulated here is that the orbit and the maximum detection range of the simulated object place physical limits on A_2 that are consistent with the numerical values of the DCR-induced spurious Yarkovsky effect. The second half of the argument in favor of plausibility is that the inclusion of the Yarkovsky term reduces the magnitude of the post-fit residuals relative to the gravity-only dynamics model. In Figure 5-14, the mean post-fit residual and its standard



(a) RA/Dec Residuals



(b) Horizontal/Vertical Residuals

Figure 5-14: Mean and Standard Deviation of Mean Post-Fit Residual of Three-Apparition Optics-Only Asteroids in Table 5.3 vs. Lower Elevation Limit on Input Observations and Inclusion or Exclusion of Yarkovsky Acceleration, in Inertial and Station-Local Coordinates. Note that unlike other figures in this thesis, the abscissa represents the elevation cutoff for the entire set of observations rather than the center of a bin. When aggregate residuals are analyzed in inertially-aligned RA/Dec coordinates, an improvement in post-fit residuals for Cases 1 and 2 is evident in the RA coordinate but is not evident in Cases 3 and 4. When analyzed in station-local horizontal/vertical coordinates, all three cases show a zero-mean horizontal post-fit residual both with and without the inclusion of a Yarkovsky acceleration in the dynamic model, while the mean vertical residual improves by between 1 and 2 standard deviations in all four scenarios over a fit that excludes the Yarkovsky acceleration. If the plausible range of physical sizes for these hypothetical objects implies a plausible Yarkovsky effect, then a three-apparition optics-only fit to observations with uncompensated DCR bias would yield better mean residuals with the inclusion of a Yarkovsky term, thereby creating a spurious detection of Yarkovsky acceleration. Note also that the post-fit vertical residuals across three apparitions do not match the input “S-C” bias applied to the original observations. While the post-fit bias for the fit without Yarkovsky acceleration manifests in the vertical coordinate, it is still attenuated by at least half when compared with the input bias.

deviation (under the assumption of 250 mas RMS observations) with and without non-gravitational acceleration for Cases 1-4 is shown versus the elevation cutoff. If residuals are analyzed in RA/Dec coordinates, then the inclusion of non-gravitational acceleration slightly reduces the RA residual for Cases 1 and 2 but has no discernible effect on Cases 3 and 4. Resolving the residuals into station-local horizontal/vertical coordinates, on the other hand, yields an unambiguous reduction by at least 1σ in the vertical residual for all four cases and as much as 2σ in Cases 2-4. This also shows how DCR bias can hide if residuals are analyzed in inertial, rather than station-local, coordinates.

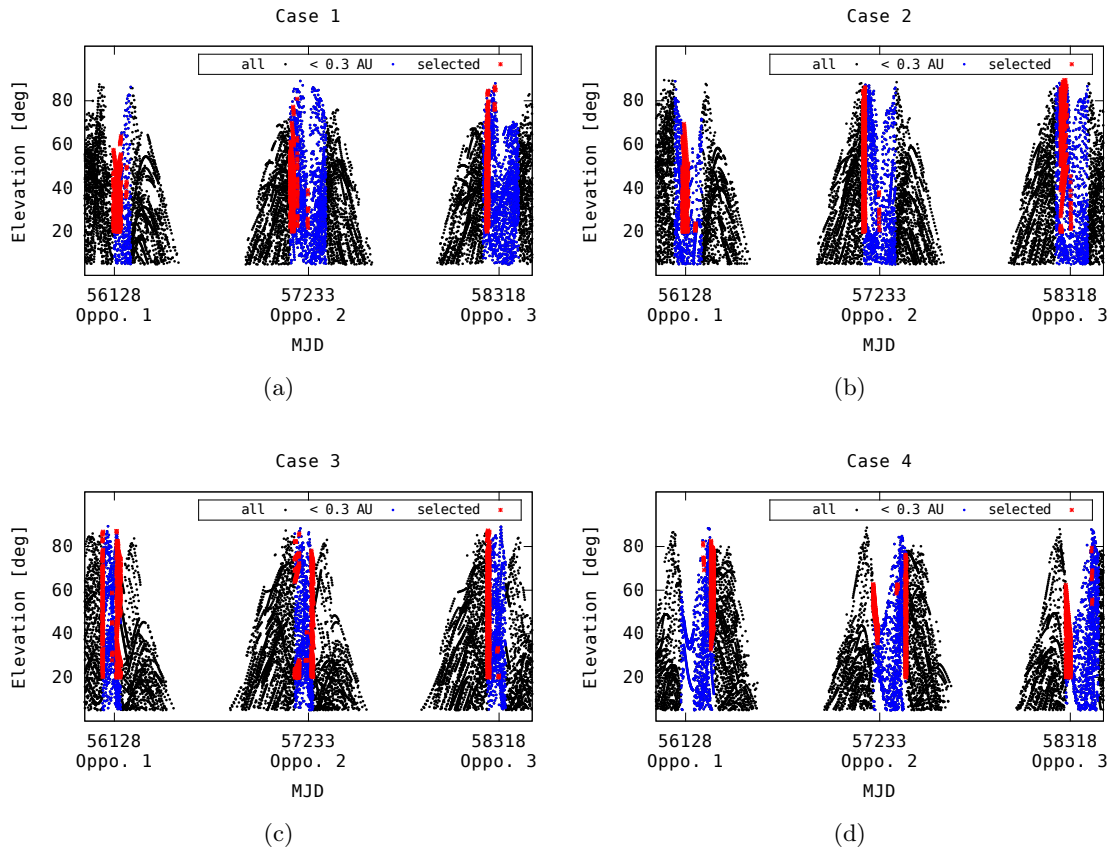


Figure 5-15: Distribution in Time and Elevation of Observations Selected for Maximal Spurious Yarkovsky Detection Relative to Distribution of All Possible Observations

5.2.5 Likelihood of a Scenario Susceptible to Spurious Yarkovsky Acceleration

Although I identified scenarios where spurious detections of Yarkovsky acceleration are possible, I have examined only a small slice of possible orbits (those with 1.5 year period) and have not attempted an optimal search for such scenarios in that orbit. This leads to the conclusion that when DCR dominates systematic error, spurious Yarkovsky detections are possible over a large portion of the orbital element space. It is, however, important to temper that statement with a recognition that the likelihood of observations occurring at times that yield spurious detections is low. Of the hundreds of distinct orbits with 1.5 year period that I analyzed, and for which I applied the same observation selection criterion, only a few dozen cases resulted in spurious Yarkovsky detections. This implies a low probability of drawing by accident a set of observations on a real asteroid that would yield a spurious detection. Furthermore, the distribution of selected observations for Cases 1-4 is not uniform over all possible opportunities. No attempt is made to rigorously quantify the

latter statement, but Figure 5-15 shows that of all possible observation opportunities, the set of observations that yield the spurious detection are clustered.

Lastly, I have checked the 205 optical observations from which the orbit of 2011 PT is derived and found negligible contribution of any vertically-aligned systematic bias to the detected Yarkovsky acceleration. The overall conclusion that to draw from this analysis is that while it is improbable that randomly selected observing opportunities would coalesce into a spurious Yarkovsky detection, the fact that it is possible implies that it is important to check any set of observations for a short-arc object for susceptibility to spurious Yarkovsky detection. I have also checked the observations of 11 ‘Oumuamua and can state that its non-gravitational acceleration is not caused by DCR bias in the ground-based observations used for its orbit solution.

5.3 Impact Monitoring of Potentially Hazardous Asteroids

Having demonstrated, using full-fidelity force models in three dimensions, that, just as in the two-dimensional motivating example in Chapter 1, orbit estimates of inner Solar System asteroids can be biased by unaccounted-for DCR bias, this section will develop the scenario presented in that motivating example: the effect of DCR bias on orbit estimates of hazardous asteroids.

5.3.1 False Confidence in Optics-Only Predictions of Near Misses: An Example


The dominance of mean motion bias in the DCR-induced error in orbit solutions for Main Belt and near-Earth asteroids implies that monitoring and collision predictions for hazardous asteroids may also be significantly affected by DCR bias, especially if orbit predictions are dominated by optical, rather than radar, observations. Inspired by the case of 99942 Apophis, I simulate an Apophis-like orbit that makes a near-Earth pass in 2029 and then impacts in 2036.

Simulated observations subjected to the DCR bias associated with the positive and negative difference between S and C type reflectance spectra spanning the thirteen years before the 2029 flyby were generated and used to form long-arc fits that were then propagated forward to 2036. Simulated observations from the observing network in Figure 5-1 were assumed to occur once per 90 minutes and have an accuracy of 0.25 arcsec. Observations were simulated at all times the target was at least 20 degrees above the horizon with the Sun at least 15 degrees below the horizon. A total of 14662 distinct simulated observations over thirteen years from all of the observing stations were used in the fit.


5.3.1.1 Bias in the B-Plane

The parameters of the impacting orbit (selected to have zero in-track and cross-track b-plane offset at the 2036 impact) and the orbit estimates derived from optics-only observations subjected to “S-C” and “C-S” bias are given in Table 5.4. The fits to biased observations miss the Earth in 2036 by approximately 500,000 km with a 1σ bound of about 173,000 km. Monte-Carlo simulations of the orbit fits were computed to numerically estimate the reduction in apparent Earth impact probability in 2036 when estimates are made using uncompensated DCR biased observations rather than bias-free observations. Monte-Carlo fitting was performed in observation space, with observation noise drawn from circular-symmetric Gaussian distributions centered around the biased and unbiased observations at each observation opportunity. The Gaussian parameters $\sigma = 250$ mas and $\sigma = 100$ mas were used to verify the analytical result from the noise-free simulation and to estimate the probability of impact that would be computed in the presence of higher-precision astrometric observations subject to the same uncompensated DCR bias.


Figure 5-16 shows the projection of the Monte-Carlo fit results onto the 2036 b-plane for the $\sigma = 250$ mas case. This represents the best-available optics-only knowledge of the 2036 flyby geometry available before the 2029 flyby. The baseline impact probability computable from unbiased data is approximately 10% with 250 mas RMS observations and approximately 25% with 100 mas observations. With observations subjected to uncompensated DCR bias, however, the probability of impact is below the 1% sensitivity of the 2000 Monte Carlo runs. Using the analytical approximation, the orbit fit derived from the 14662 observations with 250 mas formal error yields an impact probability of about 6×10^{-4} in 2036, and if those observations had 100 mas formal error (or alternatively if there were roughly 60,000 observations with 250 mas formal error), the impact probability would be below 10^{-6} . Although the difference between the impact probabilities computable from biased versus unbiased data are not the difference between near zero and near unity, the point is that in this underlying scenario, the true probability of impact is exactly 100% but the probability computable from optics-only orbit fits without accounting for DCR is near zero. This situation, while obviously not common in the true distribution of potentially hazardous asteroids (PHAs), can exist for any PHA that encounters the Earth regularly. If, on a given encounter with the Earth, that PHA happens to actually pass through a gravitational keyhole for an impact on a subsequent encounter, the DCR bias in optics-only data is sufficient to move the prediction of the first encounter out of that keyhole with seemingly high confidence.



(a) Distribution in Along-Track / Cross-Track B-Plane Coordinates



(b) Histogram along ζ coordinate



(c) Histogram of perigee heights

Figure 5-16: Monte Carlo Simulations of Optics-Only Predictions of 2036 B-Plane of Apophis-Like Impactor Derived from 13 Year Data Arc Prior to 2029. The Monte Carlo simulation consisted of 2000 random draws sampled in observation space. The Gaussian approximation derived from a fit to noiseless data and its formal covariance is a good approximation for the distribution of samples in the b-plane.

Table 5.4: Simulated Hazardous Asteroid Orbital Parameters

Parameter	Impactor	Observed with S-C bias	Observed with C-S bias
Epoch (JD TDB)	2462219.5003725188478256		
a (AU)	+9.1434884792794213E-01	+9.1434884957529661E-01	+9.1434884628055058E-01
e	+1.9313756713845182E-01	+1.9313754210271095E-01	+1.9313759217550019E-01
ω (deg)	+2.3083111328772463E+02	+2.3083110819910704E+02	+2.3083111837640683E+02
i (deg)	+3.2643398722924717E+00	+3.2643390531825585E+00	+3.2643406914290125E+00
Ω (deg)	+2.0318174546242173E+02	+2.0318174943271859E+02	+2.0318174149206752E+02
m_0 (deg)	+8.7536343096987736E+01	+8.7536339887054595E+01	+8.7536346307448838E+01




Figure 5-17: Observation Volume and Impact Estimate Probability vs. Elevation Limit for Hypothetical Apophis-Like Impactor. For this Apophis-like scenario, the majority of the observing opportunities occur at low elevation angles. Raising the elevation limit above 20° for data collected between 2016 and 2029 begins to quickly degrade the precision of the prediction for 2036, and while the probability of impact computed from DCR-biased and unbiased data are equivalent at an elevation limit of 40°, that probability is much smaller than what can be computed from the incorporation of low-elevation data, whether the DCR bias is corrected or the vertical component of the data is discarded.

5.3.1.2 Estimated Impact Probability vs. Observation Elevation Limit

The large difference between the impact probability estimate made from unbiased versus biased observations is driven, in part, by the choice to include observations down to an elevation of 20°. It is reasonable to ask whether, in the case of an identified potentially hazardous object, it would make more sense to include observations down to a lower elevation in order to increase the precision of the trajectory prediction at the potential impact time or whether, in light of the potential presence of an unmodeled DCR bias in the observations, it would be more sensible to exclude observations at low elevations in favor of higher weights on observations made at high elevations.

The trade-off between observation volume and the resulting b-plane uncertainty and impact probability estimate is in Figure 5-17. For this scenario, the majority of observations of the target are only possible at elevation angles below 30°. Raising the elevation limit to 40° eliminates DCR-

induced 2036 b-plane bias relative to the size of the b-plane uncertainty, but does so at the cost of a far larger uncertainty. At a 40° elevation limit, the estimate of probability of impact made with biased observations is indistinguishable from that made with unbiased observations, but that probability is approximately roughly 50% lower than that which could be computed from DCR-corrected observations down to an elevation limit of 20° .

In this multi-apparition scenario, an azimuth-only fit yields b-plane uncertainty bounds and impact probability estimates that are comparable with what is obtained from including both horizontal and vertical information. Therefore, despite the larger effect of DCR bias at low elevations, in this scenario it pays to observe at low elevation. The low-elevation observations must either be corrected for DCR bias or the model of per-observation uncertainty at low elevations must reflect the true information content of that data. Whether, at one extreme, DCR is completely corrected, or at the other extreme, the vertical component of the observations is completely discarded, a more accurate impact probability may be obtained over what is possible from raising the elevation limit on observations included in the fit.

5.3.1.3 Impact Probabilities with Albedo-Only DCR Estimation

This analysis is a comparison between total spectral agnosticism and perfect spectral knowledge. As discussed in Section 4.2.1, there exists a middle ground where an object's spectrum may be unknown but its albedo may be known. An Earth-crossing asteroid is likely to belong to the S-complex, so in this scenario this albedo-only middle ground would look like the DCR difference between intermediate-to-high albedo objects. The mean-value miss distance and probability of collision computable from biased observations corrected with imperfect albedo-based spectral estimates is shown for these objects in Table 5.5.

As implied by the discussion in Section 4.2.1, the improvement in the accuracy of the collision probability estimate gained from using albedo-only information can be significant. If the object in this hypothetical scenario were to have a moderate-to-high albedo and, in fact, be a member of the S-complex, then the assumption that it has a mean S-type reflectance spectrum would lead to an impact probability estimate of no less than 4%, in contrast with a 10% – 25% estimate possible with perfect spectral knowledge. If, however, the object happened to be Q-type, then the error in computing the 2036 impact point would be comparable to that of the completely color-agnostic estimate.

Table 5.5: 2036 B-Plane Miss Distance Estimate vs. Fidelity of DCR Correction and Observation Quality

DCR Fidelity	Miss Distance (250 mas RMS)	Prob. Impact (250 mas RMS)	Prob. Impact (100 mas RMS)
Perfect	$0 \pm 170 \times 10^3$ km	$\approx 10\%$	$\approx 25\%$
None (S vs. C)	$530 \times 10^3 \pm 170 \times 10^3$ km	0.06%	$\ll 10^{-10}\%$
Albedo-only (S vs. Sa)	$120 \times 10^3 \pm 170 \times 10^3$ km	5.8%	4.0%
Albedo-only (S vs. Sr)	$98 \times 10^3 \pm 170 \times 10^3$ km	6.3%	6.7%
Albedo-only (S vs. Sv)	$77 \times 10^3 \pm 170 \times 10^3$ km	6.8%	9.9%
Albedo-only (S vs. Q)	$506 \times 10^3 \pm 170 \times 10^3$ km	0.09%	$\ll 10^{-10}\%$

Table 5.6: Simulated Mumbai Impactor Orbital Parameters

Parameter	Impactor
Epoch (JD TDB)	2462209.5930692358
a (AU)	+9.1484779250612003E-01
e	+1.9243178538344743E-01
ω (deg)	+2.3095603237802106E+02
i (deg)	+3.3326502659653530E+00
Ω (deg)	+2.0318072484092306E+02
m_0 (deg)	+7.6320249150262526E+01

5.3.2 Impact Point Bias of Known Impactors

5.3.2.1 Evolution of the Impact Point Estimate of the Apophis-Like Example After 2029

Moving to a scenario closer to that of the two-dimensional example in Section 1.3.2.1, a further motivation for the careful modeling of DCR bias in orbit solutions is that once the hypothetical Apophis-like asteroid has passed its first encounter and ground-based observations begin to indicate a probable impact, the prediction of the impact point can also be sensitive to DCR bias. The impact point of the trajectory described in the Section 5.3.1 happens to be in the middle of the Indian Ocean and a slight modification to the initial trajectory (Table 5.6) places the impact in Mumbai, India, a city of over 10 million people.

Before examining the degree to which the impact point estimate is biased relative to its formal error, it is necessary to specify the calculation of that formal error. The projection of an inertial 6-DOF covariance into a two-dimensional error distribution on the Earth’s surface is a three-step procedure, illustrated in Figure 5-18.

As before, let $\Phi(t)$ be the state transition matrix from orbit fit epoch to time t and let $\Sigma_{\text{fit,epoch}}$ be the formal covariance of the six-dimensional orbital state at the fit epoch. First, the inertial

covariance at impact time is converted to Earth-fixed coordinates:

$$\begin{aligned}
\mathbf{x}_{\text{impact,gcrf}} &= \Phi(t_{\text{impact}})\mathbf{x}_{\text{fit,epoch}} \\
\mathbf{x}_{\text{impact,itrf}} &= \begin{bmatrix} \mathbf{R}_{\text{gcrf2itrf}} & \mathbf{0} \\ \dot{\mathbf{R}}_{\text{gcrf2itrf}} & \mathbf{R}_{\text{gcrf2itrf}} \end{bmatrix} \mathbf{x}_{\text{impact,gcrf}} \\
&\Downarrow \\
\Sigma_{\text{itrf,impact}} &= \begin{bmatrix} \mathbf{R}_{\text{gcrf2itrf}} & \mathbf{0} \\ \dot{\mathbf{R}}_{\text{gcrf2itrf}} & \mathbf{R}_{\text{gcrf2itrf}} \end{bmatrix} \Phi(t_{\text{impact}})\Sigma_{\text{fit,epoch}}\Phi(t_{\text{impact}})^T \begin{bmatrix} \mathbf{R}_{\text{gcrf2itrf}} & \mathbf{0} \\ \dot{\mathbf{R}}_{\text{gcrf2itrf}} & \mathbf{R}_{\text{gcrf2itrf}} \end{bmatrix}^T
\end{aligned}$$

Next, the Earth-fixed covariance must be collapsed to two dimensions along the inbound velocity vector. Figure 5-18(b) provides some graphical intuition for why that collapse operation is required. The two-dimensional plane perpendicular to the velocity vector is defined by along-track \mathbf{a} and cross-track \mathbf{c} directions as shown in Figure 5-18(b) and the collapse operation consists of:

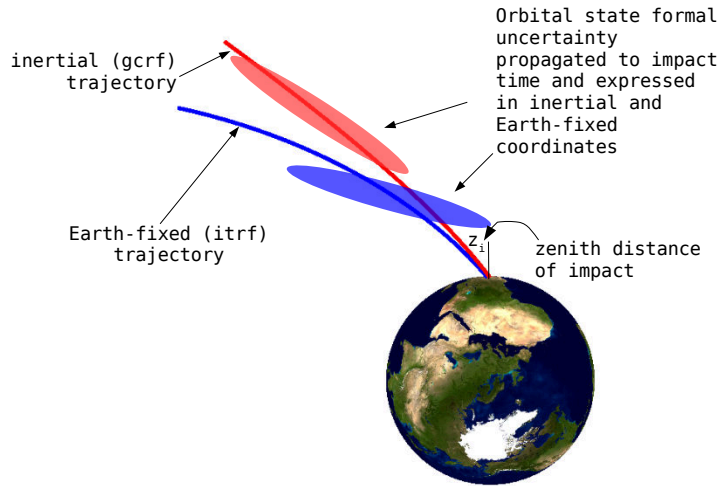
$$\Sigma_{\text{AC,space}} = \begin{bmatrix} \mathbf{a}^T & \mathbf{0} \\ \mathbf{c}^T & \mathbf{0} \end{bmatrix} \Sigma_{\text{itrf,impact}} \begin{bmatrix} \mathbf{a} & \mathbf{c} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \quad (5.5)$$

Lastly, the collapsed covariance must be “laid down” on the surface of the Earth by a cosine projection:


$$\Sigma_{\text{AC,ground}} = \begin{bmatrix} \frac{1}{\cos z_i} & 0 \\ 0 & 1 \end{bmatrix} \Sigma_{\text{AC,space}} \begin{bmatrix} \frac{1}{\cos z_i} & 0 \\ 0 & 1 \end{bmatrix} \quad (5.6)$$

The square root of the largest eigenvalue of $\Sigma_{\text{AC,ground}}$ is quoted as the ground uncertainty of the impact point. This implicitly approximates the surface of the Earth as flat and neglects the curvature of the inbound trajectory near impact time. Selection of impact scenarios where the inbound velocity comes from low zenith distances makes these approximations valid for error analysis along that dominant eigenvector, though in practice the nonlinearities matter for cross-track errors. Cross-track errors are orders of magnitude smaller than the in-track errors and are not considered in this analysis.

Returning to the Mumbai impact scenario, and assuming the same constraints on the simulated observing network and an optics-only impact point prediction, a DCR bias in the “C-S” sense (that is, the target is actually C-type but observations are corrected assuming it to be S-type) places the best-estimate predicted impact point several hundred kilometers out to sea. The time-evolution of



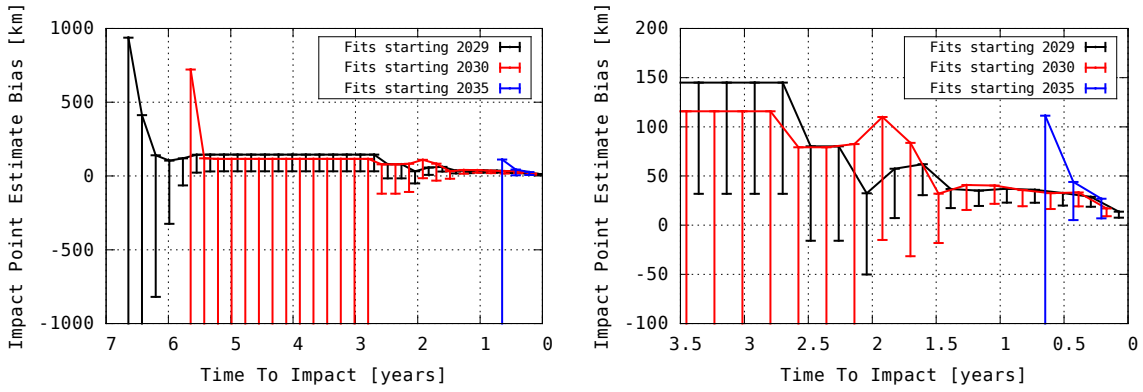
(a) Cartesian Error at Impact Time



(b) Collapsing Along Velocity

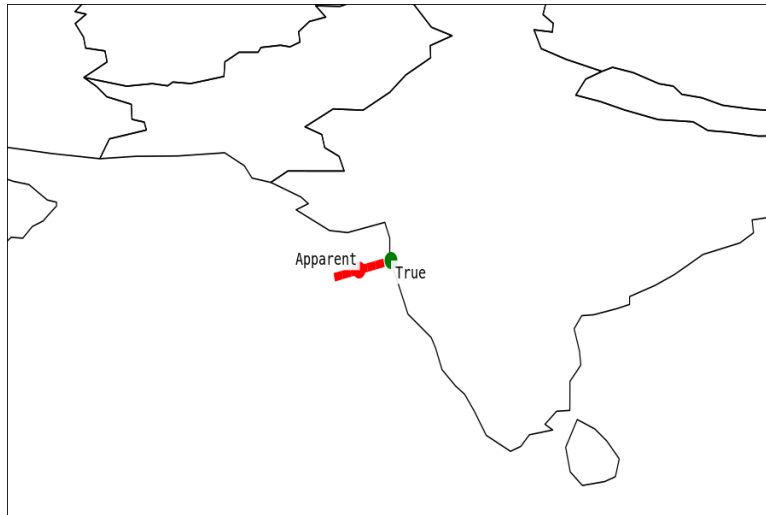
Figure 5-18: Asteroid Impact Point Uncertainty Calculation

the impact point prediction and its uncertainty are shown in Figure 5-19. The observation bias pushes the predicted impact point out to sea with at least a 3σ confidence. The dearth of observing opportunities between 2030 and 2033 (when Earth and the asteroid are on opposite sides of the Sun) means that the best-available ground-based optics-only impact point estimate remains wrong by a significant amount for a long time and does not begin to drift back toward its true location in the middle of a populated area until a few years before impact. In these scenarios, linearity holds and the $\pm 3\sigma$ formal error bounds are excellent approximations for the distribution resulting from Monte-Carlo analysis. As such, I present only the formal error bars in the dominant uncertainty direction and not the results of Monte-Carlo simulations.



(a) Mean and Mean- 3σ Impact Point Bias

(b) In Years Leading up to Impact



(c) Impact Point Estimate Bias $\pm 3\sigma$ Three Years Before Impact

Figure 5-19: Simulated Drift of Apparent Impact Point with True Impact of Apophis-Like Asteroid at Mumbai, India. A “C-S” bias in optical observations beginning after the 2029 near-Earth flyby pushes the apparent impact point out to sea by a seemingly safe distance with seemingly high confidence. An observation outage between about 3 and 6 years before impact results in the false estimate persisting until the asteroid can again be observed three years before impact.

5.3.2.2 Short-Timescale Evolution of Impact Point for Newly Discovered Impactors


In the Apophis-like scenario, one may argue that a 3σ confidence interval is insufficient and that in the event of such a potential catastrophic impact, larger bounds would be used. There do, however, exist scenarios where the bias distances are measured in thousands of kilometers and correspond to apparent confidence regions spanning over 10σ . To illustrate a pathological case, I consider newly-discovered impacting asteroids and the quality of the impact point estimate possible from a thirty-day rolling window of high-cadence observations at elevations below 20 degrees.

The sensitivity of the impact point to DCR is especially high if the impact occurs as the asteroid approaches the Earth from inside 1 AU, forcing ground-based observations at low elongation and high zenith distances. This scenario (Figure 1-6) is exactly the one analyzed in two dimensions in Section 1.3.2.1 and the result holds up when considered rigorously in three dimensions.

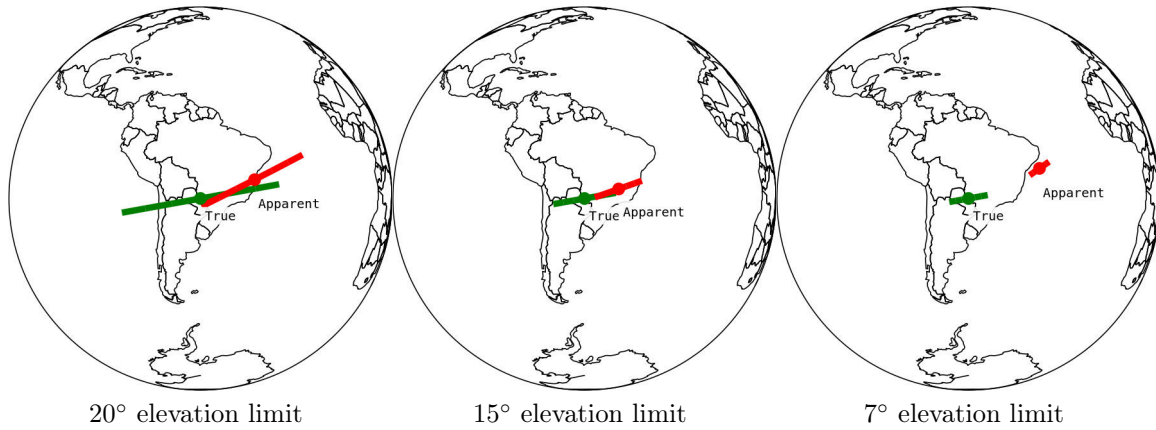
The three-dimensional analysis is performed as follows. The orbits that impact the Earth at a given point in time are constrained by the three-dimensional location of the Earth at the impact time, leaving three free parameters for the orbit: perihelion q , aphelion Q , and inclination i . The free parameters are sampled to search for cases where DCR-induced observation error introduces a statistically significant bias in the impact point prediction. The sampling is performed such that $q \in [0.5, 0.9]$, $Q \in [1.0, 3.5]$ and $i \in [-40, 40]$.

Given desired values of q , Q , and i , an impact date is selected and a heliocentric Keplerian orbital state coincident with the Earth at the chosen impact time is computed. Using two-body heliocentric physics, this initial state is propagated backward six months before impact. This becomes the seed state. The velocity component of the seed state is differentially corrected such that when propagated forward six months using the full JPL Solar System force model, an Earth impact is achieved. The only criterion for terminating the differential correction is that the incoming velocity vector of the impact have an elevation angle above 45 degrees. This is to avoid nonlinearities from the curvature of the Earth when using linear transformations to project state covariances of orbit estimates onto the surface of the Earth. Specific aiming points on the Earth's surface are not chosen for this analysis.

Once the orbit is generated, observations are simulated over the six month period prior to impact from all stations in the observing network in Figure 5-1. The stations are assumed to have 100% weather availability past astronomical twilight, be able to detect the target down to an elevation angle of 7 degrees, and to generate astrometric observations at a cadence of once every two seconds



(a) Bias and uncertainty vs time and elevation limit



(b) True and biased impact points $\pm 3\sigma$ uncertainty 105 days before impact

Figure 5-20: Evolution of Impact Bias and Uncertainty with Time for 30 Day Running Window of Impact Point Estimates for a Hypothetical Newly-Discovered Asteroid. Asteroid orbital parameters $Q = 1$, $q = 0.7$ and $i = 0$. Uncertainty estimates assume observations with ground-based telescopes at a cadence of 0.5 Hz. The true impact point bias from unaccounted-for DCR increases with progressively lower elevation cut-off. However, progressively lower elevation cutoffs increase the data volume, thereby also lowering the formal error of the impact point estimate if DCR bias is not accounted for.

with a formal accuracy of 200 mas RMS. To save computation time, the observations are simulated as normal points at half-hour intervals with formal covariance of $200 \text{ mas}/\sqrt{1800} = 4 \text{ mas}$, and S-type vs. C-type DCR bias is injected into the observations along the station-local vertical direction.

It is important to acknowledge that these assumptions on station capability and availability are not necessarily realistic given the way real telescopes operate and the impact trajectories and timelines are specifically chosen to highlight a worst-case scenario. However, it is fair to assume that in the event of a real-world scenario such as the one described here, ground-based telescopes observing at low elevation angles would dominate the astrometric observations of such a hazardous object in the months leading up to impact.




Figure 5-21: Evolution of Impact Bias and Uncertainty with Time for Florence-Like Impactor. Orbital parameters: $Q = 2.5$, $q = 1.0$ and $i = \pm 20^\circ$. Calculations are for a rolling 30-day observation window with observation cadence 0.5 Hz from all stations in view down to a 7° elevation limit.

Considering the case of an Aten asteroid with $Q = 1$, $q = 0.7$, and $i = 0$, Figure 5-20 demonstrates the potential existence of worst-case scenarios where the surprise comes with as little as two months to go. The existence of conditions where the true bias exceeds the formal error are not confined to the particular case of $Q = 1$, $q = 0.7$ and $i = 0$. Figure 5-21 shows the same scenario for an orbit like that of 3122 Florence with $Q = 2.5$, $q = 1.0$ and $i = \pm 20^\circ$, where again there exist periods where the DCR bias in the observations does not attenuate and contributes to a statistically significant bias in the impact point. These cases exist for nearly all choices of free parameters at some point within the six month window in the simulation, meaning that the original motivating analysis holds up when performed with high fidelity force models in three-dimensions.

5.4 Prediction of Stellar Occultations by Kuiper Belt Objects

DCR bias is also significant in orbit determination for Kuiper belt objects. The long orbital periods and relatively recent discovery of many members of this family of objects makes nearly all orbit determination a short-arc problem. As such, it is not appropriate to isolate the analysis of DCR-induced bias to any one orbital element like mean motion. Rather, the analysis will focus on the effect of DCR-induced orbit estimate bias on the prediction of stellar occultations by KBOs, specifically how DCR bias in observations of KBOs projects into cross-track bias of the ground track of occultation predictions.

5.4.1 Scenario

Similar to the analysis of asteroids, DCR-biased observations of a KBO are simulated and the result of interest is how the observation bias projects into relevant prediction errors. The DCR bias for a Quaoar-like object, the worst-case bias for an unfiltered silicon sensor, is used in the simulation. Realistic observation volumes for KBOs are on the order of between several hundred to several thousand distinct observations over many decades, with total observing arcs ranging from years to decades. In this thesis, I specifically examine the scenario where a new object is discovered and occultation predictions are made from observations with a small number of dedicated telescopes capable of generating a single astrometric observation per hour. Rather than the full simulated observing network used for mean motion and bias and apparent Yarkovsky analysis, the simulated observations only originate from the locations of MPC site 689 at 35 degrees north latitude and site I02 30 degrees south latitude, both at approximately 2 km altitude. This corresponds to approximately one thousand distinct observations per opposition and is a somewhat optimistic, but not unrealistic, scenario in terms of data volume.

An $a = 40$ AU circular orbit with inclination ranging from zero to 40° is considered. Successive oppositions are approximately 366 days apart. For nonzero inclination the RAAN offset relative to opposition is sampled in 60 degree increments. To capture any geography-induced bias given the chosen position of the simulated observing sites, I simulate both a Northern Hemisphere winter-time opposition at MJD=57012 and an equinox opposition at MJD=57103. Highly inclined KBOs are not observable at high elevation angles in equal volumes from both hemispheres. The disparity is greatest in winter/summer, while the observation volumes from the two sites are less lopsided when the opposition occurs at an equinox. For the case where similar observation volumes are possible from both sites, the ground track bias resulting from the incorporation of data from both sites and the ground track bias resulting from incorporating only the Northern Hemisphere site's data are examined separately.

The DCR-induced cross-track bias is a function of KBO orbital parameters and a function of the length of the available data arc. The simulations cover the following data arc length scenarios. On the low end, the fit span in the simulation runs approximately 300 days across a single opposition with cross-track error projection evaluated at the subsequent opposition. Then, the fit incorporates data from the second opposition with cross-track bias evaluated over the weeks after the end of the data arc in the second opposition. This analysis is repeated up to four oppositions after simulated

discovery, showing both the cross-track bias and its $\pm 3\sigma$ bounds that come from projecting an orbit fit made to all data from prior oppositions excluding the current opposition and from an orbit fit made to all data up to and including the current opposition.

5.4.2 Simulation Results

Figure 5-22 shows the evolution of the cross-track prediction error of the occultation path for a zero inclination KBO as a function of available data volume. Operating under the assumption that observations are generated during astronomical twilight when the target is at least 20 degrees above the horizon and that those observations are each drawn from a zero-mean circular symmetric Gaussian distribution with $\sigma = 0.25$ arcsec, the $\pm 3\sigma$ width of the occultation ground track one year after discovery and derived from only the data available on the discovery opposition is on the order of an Earth radius for a hypothetical occultation on the second opposition. The error bounds shrink down to under several thousand kilometers when data from the second opposition is incorporated into the orbit prediction. By the third and fourth opposition after discovery, the $\pm 3\sigma$ precision of the occultation prediction made at any time during the opposition for the coming several weeks using all available data up to that time reaches an approximate steady state value of about 1000 km wide when data from both observing sites is used and approximately 1500 km (that is, about $\sqrt{2}$ worse) when only the Northern Hemisphere site is used.

The reduction in the formal uncertainty with incorporation of more recent data is most dramatic during the first opposition after discovery, where there is a clear diminishing stair-step pattern to the widths of the uncertainty bounds with incorporation of additional data. As soon as the second opposition after discovery, the information added over the course of the opposition no longer reduces the uncertainty of the occultation path by more than a few percent. By the time of the fourth opposition, the orbit uncertainty has diminished to the point where there is no discernible benefit to incorporating an additional month of fresh data to the fit and the formal uncertainty in the cross-track direction of the occultation ground track is roughly the same regardless of whether the fit is derived from data that is fresh or data that is up to one year old.

As with objects orbiting in the inner Solar System, the systematic observation bias does not attenuate when averaged over many observations and builds into a statistically significant bias in the orbit prediction derived from those biased observations. In the zero-inclination case, the systematic bias in the prediction of the occultation path on the first opposition after discovery using only data from the initial discovery opposition is on the order of an Earth radius for the springtime opposition

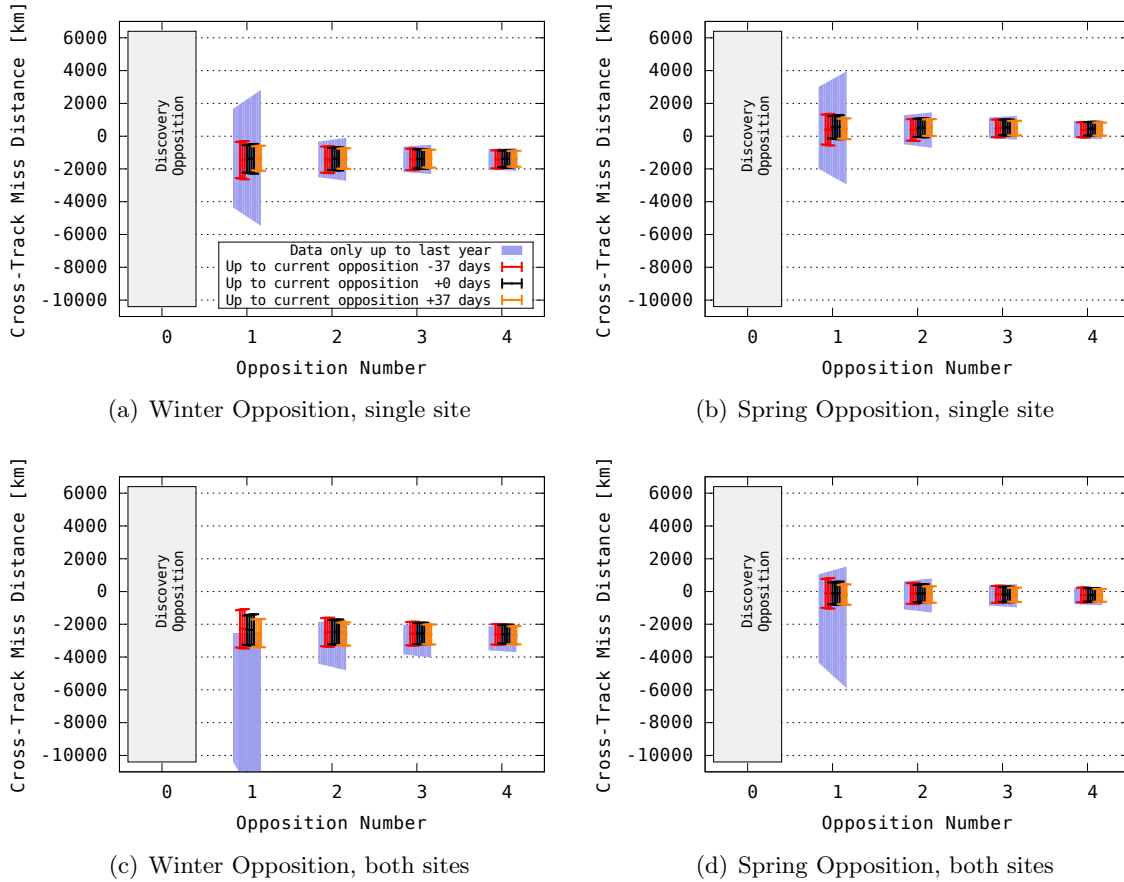


Figure 5-22: Time-Evolution of Bias of Cross-Track Prediction and Formal its 3σ Bounds for Zero-Inclination KBO. $a = 40$ AU KBO observed near opposition from single Northern Hemisphere site or jointly by Northern and Southern Hemisphere sites with opposition on the Spring Equinox or Winter Solstice.

case, leading to roughly a 50% probability of predicting an occultation path that is suitable for terrestrial observation while the true occultation path misses the Earth entirely.

When fresh data is incorporated into the orbit fit, the discrepancy between prediction and reality reduces significantly, and there are no scenarios where a predicted occultation misses the Earth entirely. Nevertheless there remain instances where the DCR-induced cross-track bias in a prediction that incorporates all available data exceeds the 3σ formal error associated with the prediction. This is particularly true for the zero inclination KBO with a Spring-time opposition when incorporating data from only the Northern Hemisphere site. In that scenario, the DCR observation bias projects into an over 500 km cross-track bias in the occultation ground track prediction even when incorporating data from five oppositions and making an immediate prediction. Incorporating data from both the Northern and Southern Hemisphere stations reduces the ground track bias to

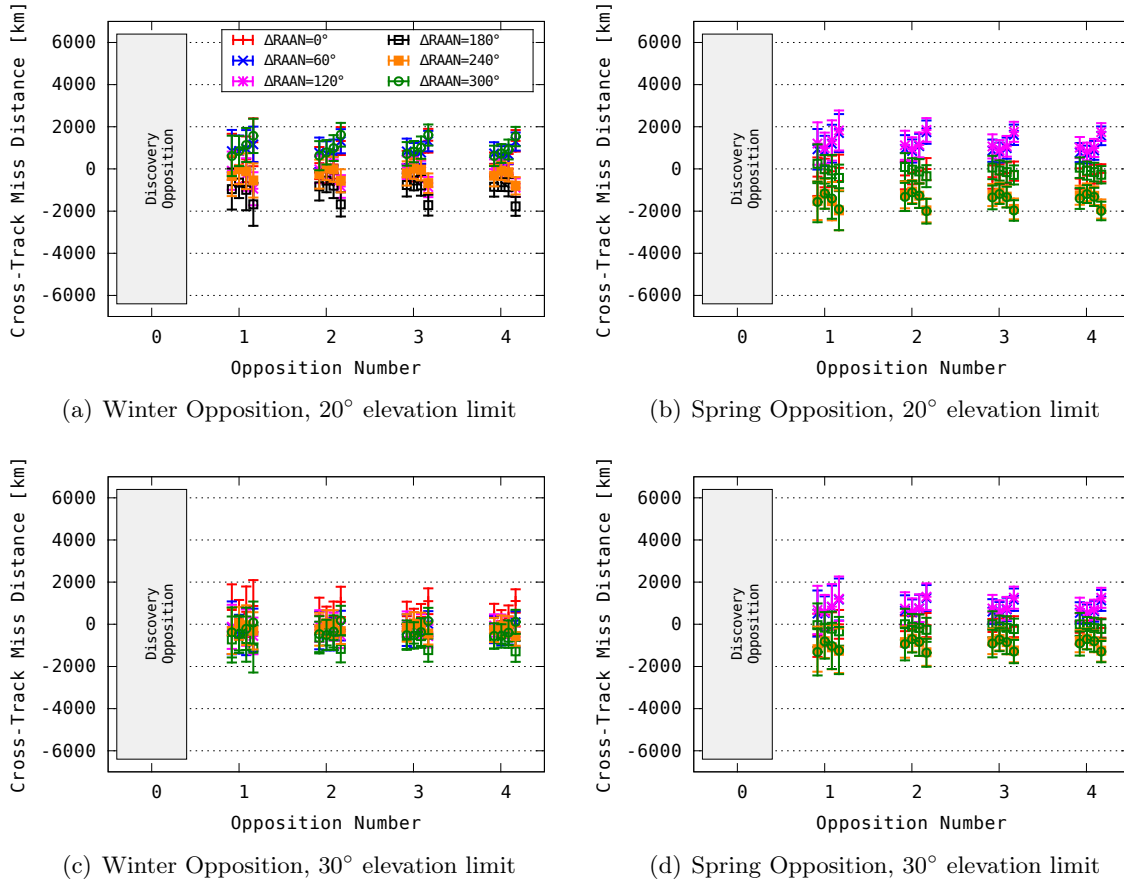


Figure 5-23: Time-Evolution of Bias of Cross-Track Prediction and Formal its 3σ Bounds for 20° Inclined KBOs. All possible $i = 20^\circ, a = 40$ AU orbits observed from stations in both hemispheres.

near-zero, but this is a case of two observation biases canceling each other out by chance. An imbalance in the number of observations from the two sites, or different weightings for data from the two stations may result in a mean bias anywhere between the -500 km in this case up to an opposite bias of around $+500$ km in the case that data from the Southern Hemisphere site dominates the orbit fit.

For KBOs with nonzero inclination, the bias in the best-available occultation path estimate depends on the offset of the KBO's RAAN relative to opposition as well as the elevation cutoff and relative data volumes between northern and Southern Hemisphere sites. Figure 5-23 shows an example of this dependence for a KBO inclined at 20 degrees to the ecliptic.

While it does appear that increasing elevation cutoff reduces the magnitude of the bias, it also reduces the available data volume and increases the formal error bounds of the occultation path. It is difficult to analyze the entire parameter space with charts like Figure 5-23, and a better performance

metric is the probability that a shadow center will be intercepted by a line of telescopes that are spread over a $\pm 3\sigma$ line perpendicular to the occultation ground track. This performance metric implicitly captures the variation in shadow path uncertainty that depends on the particular data volumes. Recall that the $\pm 3\sigma$ uncertainty width is between 1000 km and 2000 km when stations are assumed to generate one observation per hour with a formal error of 250 mas RMS. Figure 5-24 shows the median probability over the time period around opposition of shadow intercept across the third opposition past discovery using all available data up to the prediction time as a function of orbital inclination, RAAN offset from opposition, and elevation cutoffs from 20 to 40 degrees. Assuming the telescope spacing is commensurate with the expected diameter of the object¹, the baseline for probability of shadow interception for a $\pm 3\sigma$ wide line of telescopes is over 99 percent.

While the overall trend of these results is that increasing the elevation limit on observations and favoring a two-site over a single-site solution improves prediction performance, there are a few pathological geometries where this is either not the case or is not possible. Figure 5-25 shows the total available observation volume available for a third-opposition prediction from the Northern Hemisphere and Southern Hemisphere site for the several orbits with $i = 30^\circ$ as a function of the allowable lower elevation limit as well as the probabilities of successfully measuring the occultation.

For $i = 30^\circ, \Delta\Omega = 0^\circ$ shown in Figures 5-25(a) and 5-25(b), high elevation observations are not available from the Southern Hemisphere station for a Wintertime opposition, and in this scenario, the increase in observation arc length that comes with incorporating more data from lower elevations from the single Northern Hemisphere site beats out the effect of more DCR bias in observations made at low elevation in terms of increasing the shadow path prediction accuracy. Geographic diversity of observing stations is also not always a guarantee of improved prediction performance. For the case of $i = 30^\circ, \Delta\Omega = 300^\circ$ shown in Figures 5-25(c) and 5-25(d), a single-site fit gives better prediction accuracy for elevation limits of 20° or 30° for both a Springtime and Wintertime opposition. In the Winter case, the data set is dominated by the Northern site, but the effect of DCR bias is not significant within that data set. In this scenario, incorporation of data from the Southern Hemisphere site is a case of fixing what isn't broken, and the overall prediction performance degrades. For the Winter case in Figures 5-25(c) and 5-25(d), the improvement in prediction performance from only a single site is spurious, as the low observation volume from the Northern site in this scenario implies that the $\pm 3\sigma$ uncertainty is larger than the Earth. Figures

¹The observation cadences and orbital parameters used in this simulation imply an object on the order of 100 km in diameter.

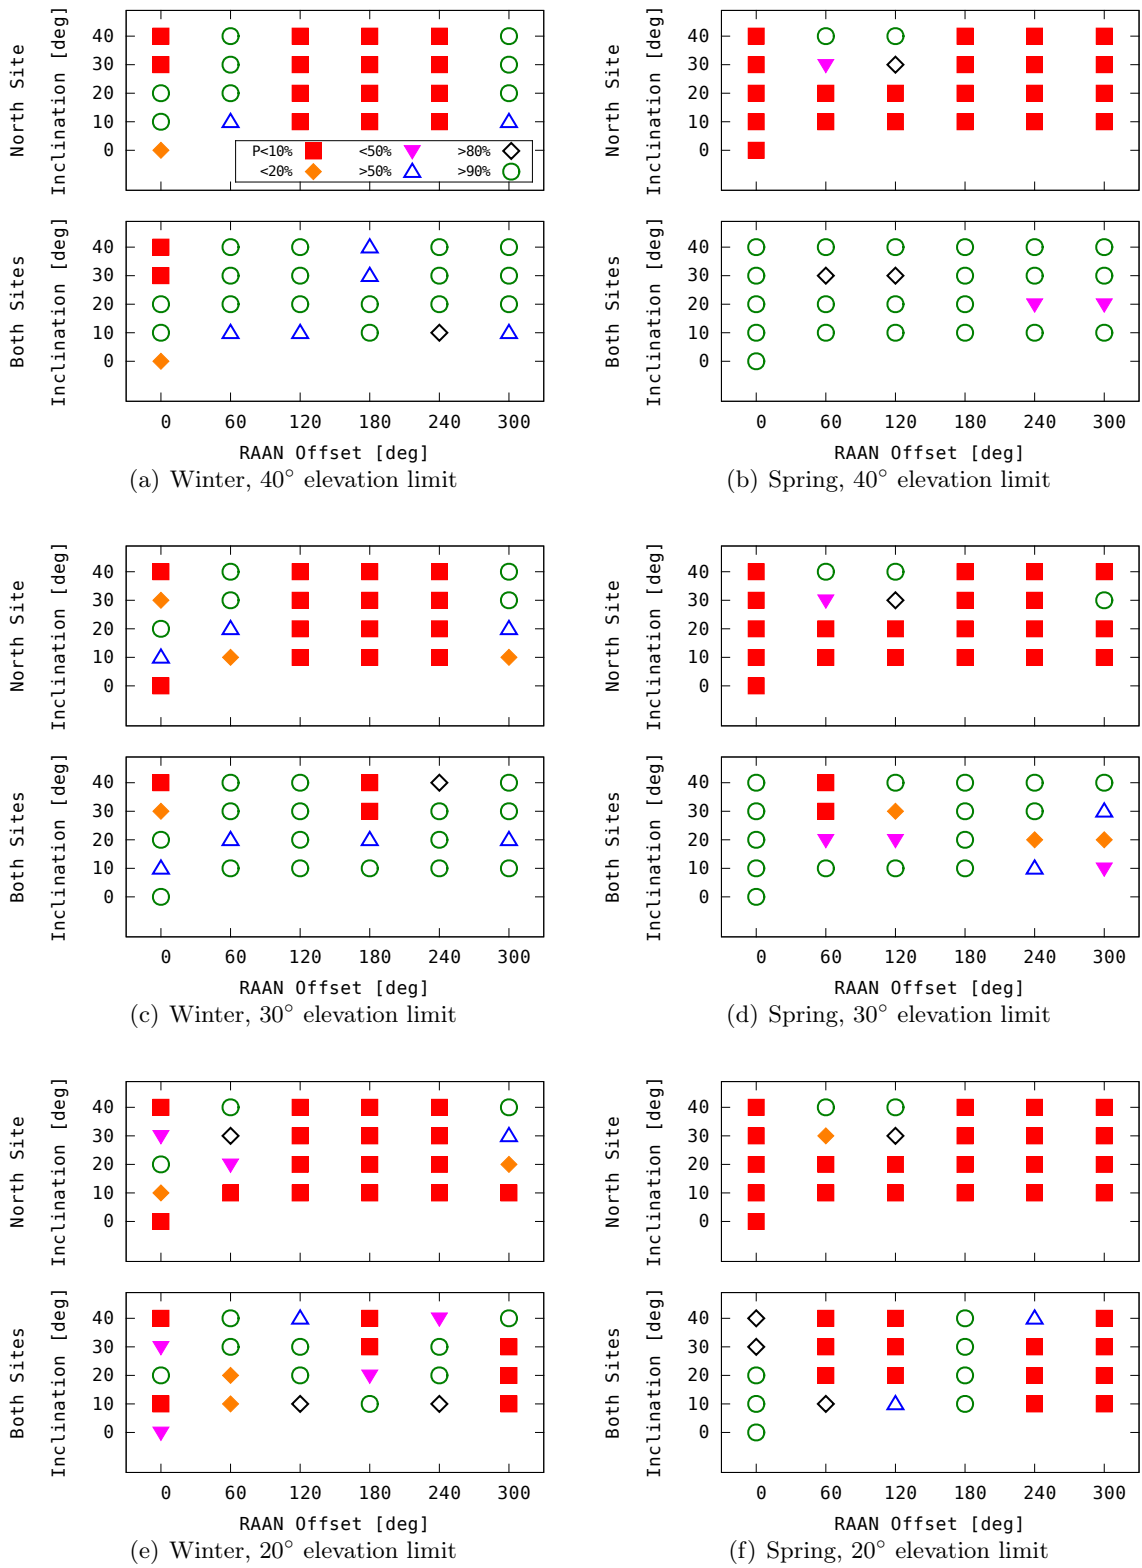


Figure 5-24: Probability of Shadow Path Interception by $\pm 3\sigma$ Wide Line of Telescopes vs. KBO Orbital Parameters and Observation Lower Elevation Limit. The probability quoted is taken as the median over time around the third opposition after discovery, incorporating into the shadow path estimate all available data up to the time of the potential occultation.

5-25(e) and 5-25(f) demonstrate another takeaway from this analysis: throwing out low elevation data is not always desirable even when data volumes are not critically low. In the Winter opposition $i = 30^\circ, \Delta\Omega = 120^\circ$ scenario in Figure 5-25(e) and 5-25(f), there is no appreciable sensitivity to elevation limit in the probability of successfully measuring an occultation predicted with a two-site fit. Thus, it would pay to incorporate all available low-elevation data into the prediction in order to narrow the uncertainty interval for the shadow path prediction.

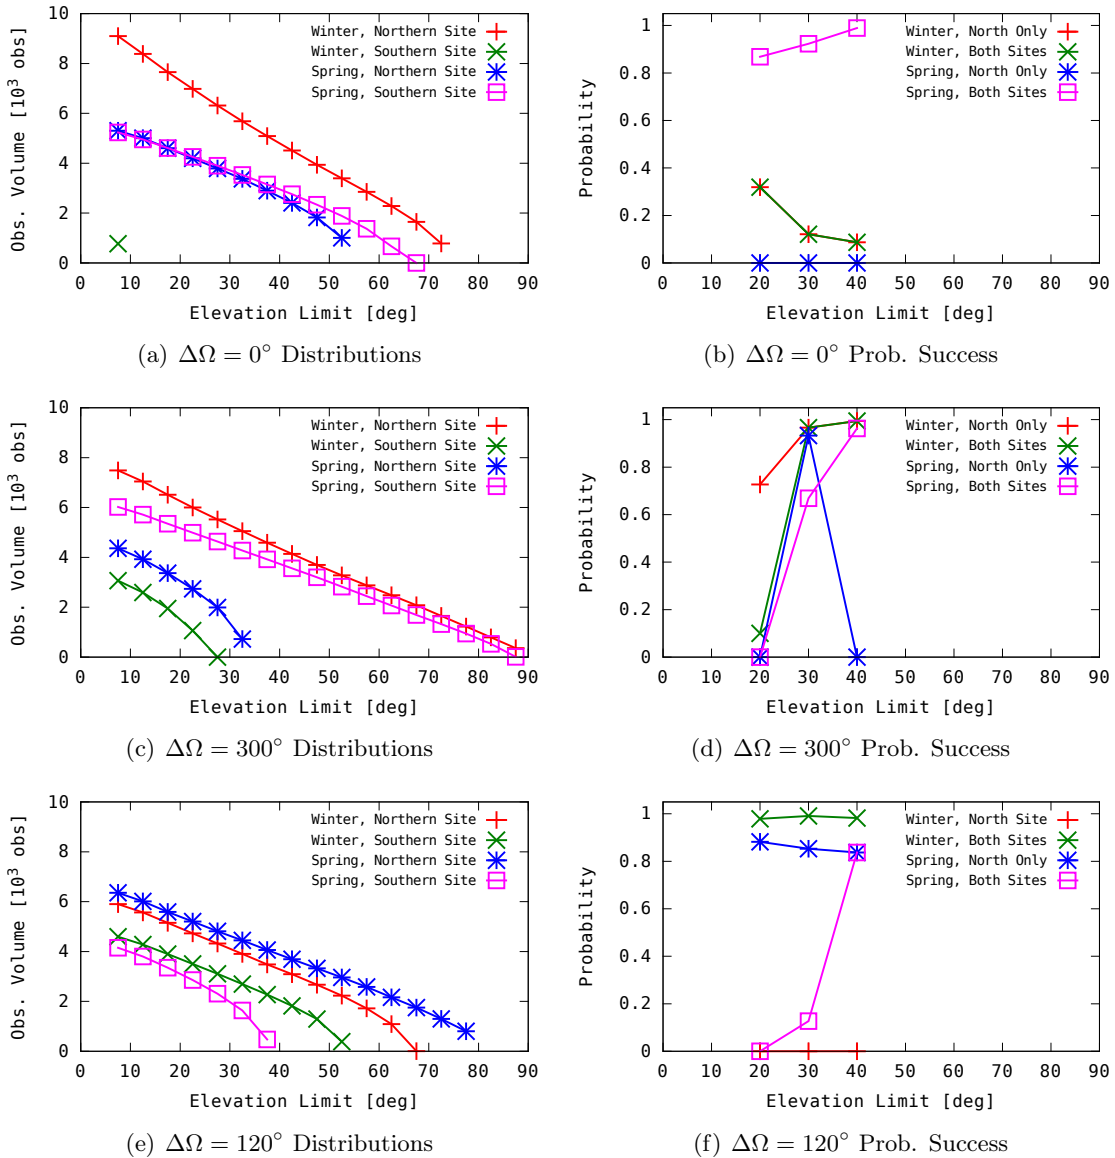


Figure 5-25: Observation Volume and Occultation Measurement Success Probability vs. Elevation Limit for selected $a = 40$ AU, $i = 30^\circ$ KBO Orbits. Total observation volume vs. elevation limit available to make a prediction for a stellar occultation on the third opposition after discovery, and probability of successfully observing that occultation using a line of telescopes $\pm 3\sigma$ wide centered on and oriented perpendicular to the predicted shadow path.

5.4.3 Conclusions

As the simulations show, there is a significant reduction in the probability of successful measurement of the occultation for nearly all sampled orbits when only a single observing site is used. When two observing sites are used, the reduction is less severe, although there are still pathological cases where the bias in the occultation path exceeds the formal error bounds by an amount sufficient to reduce the the intercept probability to below 50%. Raising the allowable lower elevation limit on observations generally improves prediction performance when two observing sites are used, but does so at the expense of requiring a wider baseline of telescopes for observing the occultation and does not improve performance appreciably when only a single site is available from which to make occultation predictions. There exist certain pathological geometries, mostly for higher inclination orbits, for which raising the elevation limit is not possible or actually reduces the prediction performance.

5.5 DCR Bias in Geostationary Orbits

An observation bias on the order of a few hundred milliarcseconds from the ground corresponds to a Cartesian bias of several tens of meters at geosynchronous altitude. The orbit prediction error that results from this observation bias would also be on the order of a few tens of meters or a few hundred milliarcseconds in angle and would grow with time. Any mission requiring orbit solutions and predictions that are accurate to that level is significantly impacted by DCR bias. Three use cases with such requirements are:

1. Solar radiation pressure (SRP) Area-to-mass ratio (AMR) estimation for an individual active RSOs to refine the force model for that object;
2. Prediction of collisions with uncontrolled debris (whose trajectory predictions are by definition derived from non-cooperative measurements and whose spectra are variable);
3. Prediction of appulses against natural guide stars for the purpose of resolved imaging of GEO satellites with ground-based sensors.

The primary focus of this section is motivated by the later case, and in particular on the IARPA Amon-Hen program [30] to construct a ground-based instrument to image geosynchronous satellites using natural guide stars. The requirement for this application is to predict an appulse between the selected target and a sufficiently-bright guide star. The orbit prediction requirement is formulated

in angle space. The in-track component of the orbit prediction error corresponds to the timing of the predicted appulse and the cross-track component of the error corresponds to the appulse distance. The latter is the more stringent requirement for this application. For interferometric or adaptive optics imaging, the appulse distance must be within the isopistonc or isoplanatic angle, which is on the order of single-digit arcseconds [31]. This analysis examines the relationship between the amount of telescope time necessary to devote to any one SSA target in support of appulse prediction and the resulting evolution of the worst-case cross-track error of a short-arc orbit fit and examines the effect of systematic DCR bias on this prediction error and the relation of AMR estimation to prediction accuracy.

5.5.1 Scenario

Suppose that appulse predictions are made from observations collected by a telescope co-located with the resolved imaging instrument and that four hours of telescope time are dedicated to a round-robin observation of several geosynchronous satellite targets. The sensor is assumed to generate astrometric observations with 250 mas RMS at a cadence of up to 8 Hz. The first thing to consider is the time evolution of the cross-track component of the formal covariance of the orbit prediction for a target as a function of observation volume obtained in that four hour observing period. After determining the data volume necessary to obtain the required cross-track precision in the orbit prediction in the absence of systematic observation bias, the same case is simulated with observations subjected to a DCR bias of a 50% solar panel and 50% Kapton reflectance spectrum.

The target's orbital inclination, and longitude and RAAN offsets from the observing telescope are varied to sample the free parameters of near-circular geosynchronous orbits. Orbit fits are simulated at the Winter and Summer Solstices and Spring Equinox to gauge the effect of solar declination angle on the cross-track prediction accuracy. An Earth gravity model up to degree and order 8, solar, lunar, and planetary third body gravitational perturbers, relativistic correction, and solar radiation pressure is used for the simulation. An area-to-mass ratio of 0.03 m²/kg, typical of many geostationary satellites, is used in the simulation. The simulations are made for a site at 30° north latitude and 1.6 km altitude.

5.5.2 Cross-Track Prediction Accuracy

Figure 5-26 shows an example of the evolution of in-track and cross-track prediction $\pm 3\sigma$ error bounds for a fit to unbiased observations of a geostationary and an inclined satellite observed over

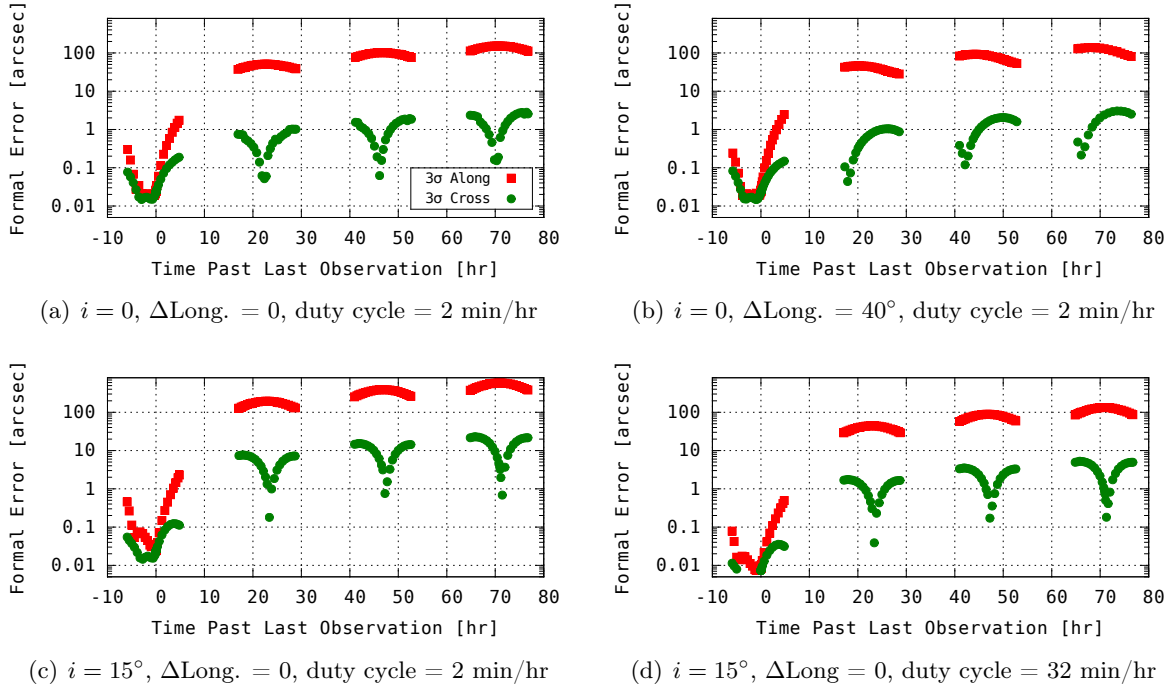


Figure 5-26: Time-Evolution of $\pm 3\sigma$ Along-Track and Cross-Track Bounds of GEO Orbit Predictions from Four-Hour Arc vs. Observing Duty Factor and Orbit Geometry. Simulated observations at 8 Hz cadence and 250 mas RMS noise.

a four-hour arc with a fit epoch around the Winter Solstice.

The formal $\pm 3\sigma$ cross-track error bounds on a zero-inclination target anywhere on the GEO belt (Figures 5-26(a) and 5-26(b)) observed for as little as 2 minutes per hour stay within 3 arcseconds for up to three days past the last observation, potentially enabling a single telescope to issue predictions for up to thirty zero-inclination targets. Inclined targets, however are more difficult to pin down in both the along and cross-track coordinate as shown in Figure 5-26(c). Increased data volume is required to achieve comparable prediction performance for inclined targets. In the short-arc regime, the data volume increase is roughly proportional to the square of the desired improvement, as shown in nearly four-fold improvement over Figure 5-26(c) that comes from increasing data volume by a factor of 16 in Figure 5-26(d).

Considering only the next-night prediction accuracy against geostationary and 15° inclined orbits, the worst-case $\pm 3\sigma$ cross-track bounds are shown in the left column of Figure 5-27 as a function of free parameters of the orbit and the season. Predictions are notably more difficult during eclipse season because of reduced data volume. Still, for many orbits it is possible to achieve bounds under 2 arcsec for the next night with a four hour observing arc and a 2 min/hour duty cycle, with

only certain orbit geometries necessitating more data volume. These geometries are predictable in advance, meaning overall throughput for next-night appulse predictions from a single telescope can be between two and thirty satellites depending on the particular targets being observed.


The previous statement only holds when astrometric measurements are free of systematic bias. The effect of DCR bias can be significant. As shown in the right column of Figure 5-27, for certain orbit geometries the DCR-induced cross-track bias in the orbit prediction can be as large as, or even exceed, the $\pm 3\sigma$ bounds of the prediction. This problem cannot be relieved by increasing collection cadence or duty cycle, or arc length within a single night. As shown in Figure 5-28, a factor of 16 increase in data volume uniformly tightens the predicted $\pm 3\sigma$ bounds by a factor of four but has nearly no effect on the systematic DCR-induced error.

5.5.3 Area-to-Mass Ratio Bias


Analysis of the behavior of the fits in the previous section reveals that the DCR bias is absorbed by the SRP AMR. The best-fit values of AMR can be highly non-physical. Figure 5-29 shows the best-fit AMR and its formal uncertainty as a function of season, longitude offset from the telescope and RAAN offset at the time of observation. The AMR bias is consistently positive in the summer and negative in the winter.

As with the cross-track error performance metric, raising the duty factor or observation cadence does not improve AMR estimation accuracy. As shown for the case of the Winter Solstice fits in Figures 5-29(b) and 5-29(e), with more data volume or arc length within a single night, the formal error in AMR goes down to under 10% of the true AMR value used in the simulation while the best-fit estimate remains biased. This gives false confidence in the AMR estimate and, in turn, false confidence in orbit predictions derived from this orbit solution for the purposes of conjunction assessment, maneuver detection, debris identification, and appulse prediction. This is a particular hazard for lower inclinations or smaller DCR biases which may result in an AMR estimate that is incorrect but not obviously non-physical.


It is important to note that AMR is measurable to a high relative precision over a short arc in the absence of DCR bias. Recall that from Section 3.3.6 that while the geocentric force model is specified to a precision of 10^{-20} m/s², the uncertainty in the actual force on a satellite above LEO is driven by uncertainty in non-isotropic thermal radiation, antenna thrust, and the dynamic component of the gravity field. At GEO, the uncertainty in instantaneous acceleration is about 10^{-9} m/s². Thus, the AMR of 0.03 m²/kg used in this simulation, typical of GEOs and corresponding to an acceleration




(a) Zero inclination $\pm 3\sigma$ Bounds




(b) Zero inclination DCR-induced Bias




(c) $i = 15^\circ$, Winter Solstice $\pm 3\sigma$ Bounds




(d) $i = 15^\circ$, Winter Solstice DCR-induced Bias




(e) $i = 15^\circ$, Spring Equinox $\pm 3\sigma$ Bounds



(f) $i = 15^\circ$, Spring Equinox DCR-induced Bias



(g) $i = 15^\circ$, Summer Solstice $\pm 3\sigma$ Bounds



(h) $i = 15^\circ$, Summer Solstice DCR-induced Bias

Figure 5-27: GEO Next-Night Cross-Track Prediction Error for 2 Minutes Per Hour Observation Duty Cycle Over Four-Hour Arc. $\pm 3\sigma$ bounds and DCR-induced systematic bias vs. inclination, longitude and RAAN offset, and season. 8 Hz data rate, 250 mas RMS noise.

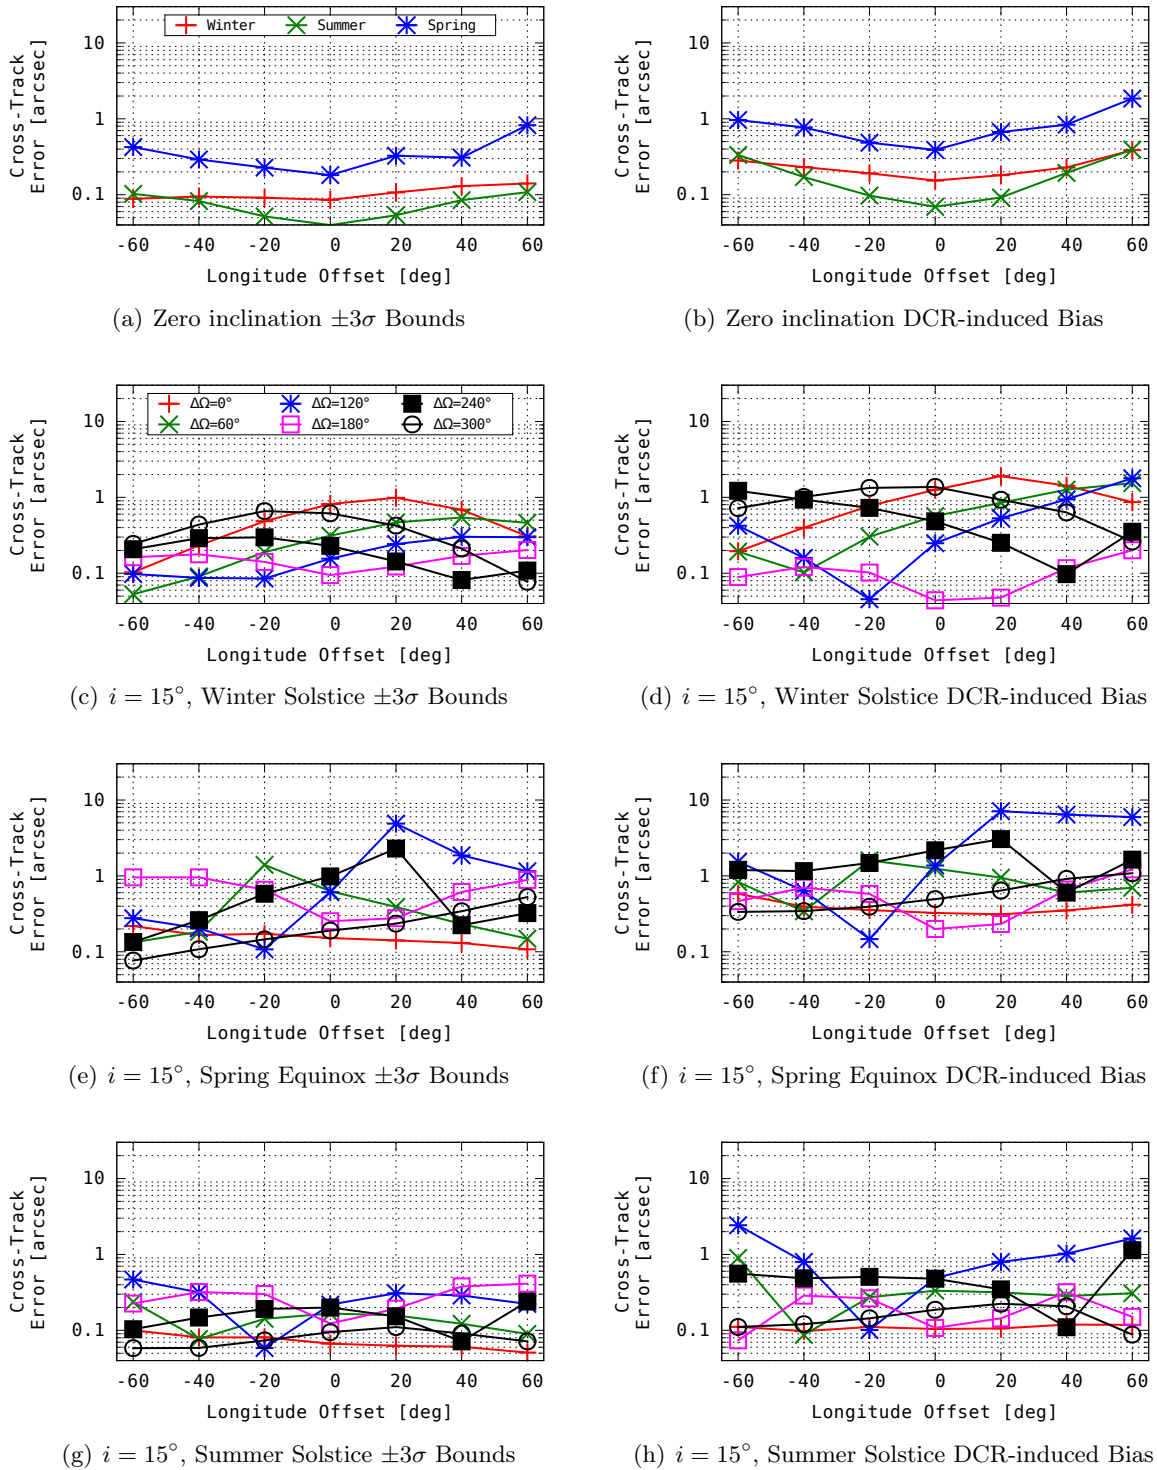





Figure 5-28: GEO Next-Night Cross-Track Prediction Error for 32 Minutes Per Hour Observation Duty Cycle Over Four-Hour Arc. $\pm 3\sigma$ bounds and DCR-induced systematic bias vs. inclination, longitude and RAAN offset, and season. 8 Hz data rate and 250 mas RMS noise. While the formal error bounds contract by a factor of approximately four with respect to Figure 5-27, the DCR-induced bias is not attenuated by the incorporation of additional data.




(a) $i = 15^\circ$, Winter Solstice, 4 hr arc, 2 min/hr duty cycle





(b) $i = 15^\circ$, Winter Solstice, 4 hr arc, 32 min/hr duty cycle



(c) $i = 15^\circ$ Spring Equinox, 4 hr arc, 2 min/hr duty cycle



(d) $i = 15^\circ$ Summer Solstice, 4 hr arc, 2 min/hr duty cycle



(e) Zero Inclination, all night arc, 16 min/hr duty cycle

Figure 5-29: Bias and $\pm 3\sigma$ Formal Uncertainty in AMR With 8 Hz, 250 mas RMS Astrometry

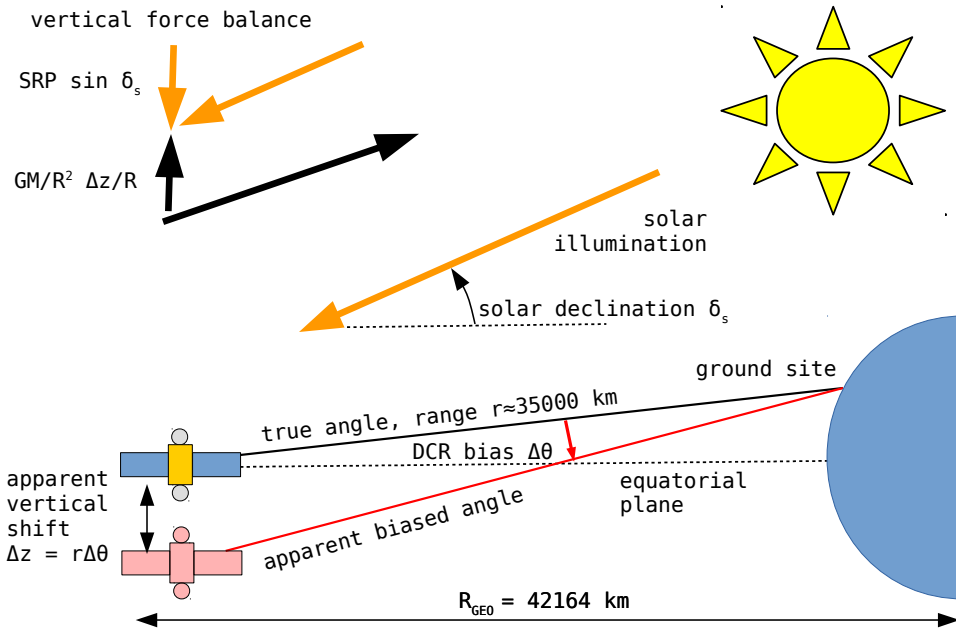


Figure 5-30: Geometric Intuition for Bias in Geostationary Satellite Solar Radiation Pressure Coefficient Estimate

of $1.3 \times 10^{-7} \text{ m/s}^2$, ought to be constrainable to a precision of 1% or better².

Figure 5-30 provides intuition for the behavior of DCR-induced AMR bias for a geosynchronous satellite. An apparent downward bias from a Northern Hemisphere ground station looks like an apparent downward shift of the satellite. This shift is small, $\Delta z = r \times \delta\theta \approx 16$ meters for a bias of 100 mas. Ordinarily, this would indicate a tiny inclination offset $\sin \delta i \approx \Delta z/R$ of the orbit with respect to the equatorial plane. However, a four hour fit span is sufficient to constrain the semi-major axis and inclination to their true values. The result is that the AMR must make up the difference by the force balance in the z direction. At the solstices, the Sun is at its maximum declination extremes of ± 23.5 degrees, meaning $\sin 23.5^\circ \approx 40\%$ of the solar radiation pressure is in the vertical direction.

The force balance requires that this force cancel out the $\frac{GM}{R_{\text{GEO}}^2} \sin \delta i$ gravitational force if the satellite were to remain at its apparent z offset. Assuming the force balance holds exactly and the DCR bias is on the order of 100 mas, the residual acceleration necessary to keep the satellite at its z offset is about $9 \times 10^{-8} \text{ m/sec}^2$. This corresponds to the vertical component of solar radiation pressure, and at the extreme solar declinations of summer and winter, the total acceleration from radiation pressure would be $(9 \times 10^{-8} \text{ m/sec}^2) / \sin 23.5^\circ \approx 2.25 \times 10^{-7} \text{ m/sec}^2$. For a solar constant

²Antenna thrust, while real and measurable over the long term, is negligible with respect to SRP for a satellite weighing several tons and radiating kilowatts in the radial direction.

of 1367 W/m^2 , this acceleration corresponds to an AMR on the order of $0.05 \text{ m}^2/\text{kg}$, which agrees with the AMR bias in the numerical results in Figure 5-29(e).

In the summer, the Sun is at a positive declination with respect to the satellite at midnight and the AMR bias is positive. In the winter the Sun is at a negative declination and this effect causes the AMR estimate to be negative. In both cases, the true AMR is masked by a spurious signal that is larger in amplitude than the true value of AMR while the formal error in AMR is small. Without correcting for DCR bias, several nights' worth observations are necessary to constrain AMR to its true physical value.

5.5.4 Sub-20 Meter Conjunction Predictions

It is important to emphasize that the DCR-induced AMR bias is a primary driver of prediction error for GEO Cartesian states and ground-based angle measurements. While a force balance perpendicular to the orbital plane sets the level of excess radiation pressure acceleration in the orbit fit, solar radiation pressure acts mostly within the plane of the orbit. An excess solar radiation pressure acceleration of $2.25 \times 10^{-7} \text{ m/sec}^2$ builds into a Cartesian error on the order of 20 meters in the orbital plane over twelve hours. The exact coupling of that error into in-track and cross-track angular prediction error on the next night does not have a simple graphical explanation like Figure 5-30, but the order of magnitude of the imparted error would be at least 100 mas by this reasoning and three-dimensional conjunction predictions in GEO would have a growing error floor of tens of meters per day, driven by DCR-induced AMR estimation bias.

5.5.5 Empirical Validation of Predictability of GEO Appulses

A multi-night measurement was conducted to validate the intuition that zero-inclination targets' orbits are easier to predict in cross-track than inclined targets and that the AMR estimate absorbs much of the DCR bias from real observations in short-arc orbit fits. Geosynchronous satellites were observed over three consecutive nights from January 30, 2018 (2018:030) to February 1, 2018 (2018:032) using the 31" telescope at the MIT Lincoln Laboratory Experimental Test Site (ETS) in Socorro, NM.

Six targets were split into two groups of three, with each target in the group being observed round-robin in ten minute bursts for a total arc-length of 3.5 hours for each group. Five satellites were successfully observed for all three nights. Observations from all five satellites with three nights of astrometry were reduced using UCAC5 with stellar DCR correction applied.

5.5.5.1 Ballistic Target Selection

To determine whether the targets remained ballistic for the entire observing period, orbit fits were computed to the entire three-night data span and to two-night data spans excluding the first or the last night. The ballistic model included Earth gravity up to degree and order eight, all gravitational perturbers, relativity, and solar radiation pressure AMR as a free parameter of the fit. The variation in best-fit values of AMR over the three-night arc and the two two-night arcs were used as an indicator for a maneuver having occurred. The best-fit AMR estimates and the truth (where available) are listed in Table 5.7. Truth values for the AMR were obtained by fitting the same dynamical model to satellite operator published ephemerides from Intelsat for Galaxy-25 and from the Goddard Space Flight Center (GSFC) Flight Dynamics Facility for TDRS-3 and TDRS-6. Formal errors for the three-night fit are on the order of $0.006 \times 10^{-2} \text{ m}^2/\text{kg}$ and formal errors for the two-night fits are on the order of $0.02 \times 10^{-2} \text{ m}^2/\text{kg}$.

Table 5.7: Best-Fit Multi-Night Area-to-Mass Ratio

SCC	Name	Inclination	AMR [m^2/kg]			Truth	Source
			Nights 0-2	Nights 0-1	Nights 1-2		
19548	TDRS-3	14°	3.48×10^{-2}	3.35×10^{-2}	3.72×10^{-2}	3.54×10^{-2}	GSFC
22314	TDRS-6	14°	3.66×10^{-2}	3.66×10^{-2}	3.28×10^{-2}	3.43×10^{-2}	GSFC
24812	Galaxy-25	0°	6.12×10^{-2}	6.13×10^{-2}	6.12×10^{-2}	5.80×10^{-2}	Intelsat
36411	GOES-15	0°	1.56×10^{-2}	8.4×10^{-3}	1.15×10^{-2}	N/A	
41866	GOES-16	0°	-7.6×10^{-2}	6.6×10^{-3}	-1.94×10^{-1}	N/A	

The GOES satellites all appeared to execute a non-negligible station-keeping maneuver or momentum dump at some point on the second night as indicated by their inconsistent or nonphysical AMR estimates. The consistency of the AMR estimates for all combinations of optical data and the operator-published truth for TDRS-3, TDRS-6, and Galaxy-25 indicate that these targets remained purely ballistic over the observing span, allowing for a meaningful examination of the effect of DCR bias on one-night AMR estimate and cross-track prediction accuracy for inclined and geostationary orbits.

5.5.5.2 One-Night AMR Estimation

Single-night AMR estimates were computed for all three ballistic satellites. The resulting (biased) AMRs are listed in Table 5.8. With the caveat that quantitative AMR bias predictions cannot be made without having spectroscopic measurements coincident with every astrometric observation, several conclusions can be drawn from the results. As expected the AMR bias for inclined satellites

Table 5.8: One-Night AMR for Ballistic Targets

SCC	Name	AMR [m ² /kg]			Truth
		Night 0 (2018:030)	Night 1 (2018:031)	Night 2 (2018:032)	
19548	TDRS-3	-0.42 ± 0.38	-13.7 ± 2.50	-1.71 ± 1.2	3.54 × 10 ⁻²
22314	TDRS-6	+2.00 ± 0.70	+13.2 ± 3.11	-17.4 ± 8.7	3.43 × 10 ⁻²
24812	Galaxy-25	+0.26 ± 0.03	+0.02 ± 0.03	+0.19 ± 0.05	5.80 × 10 ⁻²

is significantly higher than for the zero-inclination Galaxy-25. Further, a mostly positive AMR bias on Galaxy-25 observed from the Northern Hemisphere in the winter implies that its DCR bias should also be positive, which would be consistent with a satellite presenting an optical cross-section composed of mostly solar panel signature.

5.5.5.3 Cross-Track Predictability

The predictability of appulses of geosynchronous targets using bare-silicon observations without DCR compensation was evaluated for the three purely ballistic targets. Performance was measured by the observation error from a one-night fit from the first night propagated forward into the second and third nights. The error in cross-track and time is shown in Figure 5-31.

As predicted for the zero-inclination target, a one-night fit is sufficient to make cross-track appulse predictions accurate to at least ±2 arcsec, even in the presence of uncompensated DCR bias. However, for the inclined satellites the one-night fit is insufficient to make predictions to the same level of accuracy. For the inclined case, the location of the target is pinned around the observation time, as shown by the zero-crossing in the cross-track error at one-day intervals, but the prediction error away from that short time period will quickly exceed the isoplanatic/isopistonc angles on the second night.




Figure 5-31: Ballistic Targets Time and Cross-Track Errors for One-Night Orbit Fit

5.6 Chapter Summary

This chapter has addressed the second of the three original contributions of this thesis:

2. The consequences of leaving DCR bias uncorrected in orbit solutions are examined for Earth satellites, inner Solar System asteroids, and Kuiper Belt Objects. It is shown that there exist certain use-cases where DCR bias in observations can significantly affect the resulting orbit estimate in all of these orbital regimes.

It has been shown that DCR bias can introduce statistically-significant biases into orbit solution and prediction products for targets in all of the analyzed orbit regimes.

For inner Solar System asteroids, DCR can project into a statistically-significant bias in the Keplerian mean motion estimate of asteroids in Aten, Apollo, Amor, and Main Belt orbits in a single-apparition fit. There exist geometries where a lower elevation limit on input observations as high as 40° still allows much as 5σ biases in Keplerian mean motion. For certain orbits, this mean motion bias builds into a detectable and significant bias in the recovery (or precovery) residual on the next opposition. This residual can look like a non-gravitational perturbation acting on the asteroid and a scenario was developed where DCR bias can look like a physically-plausible and statistically-significant Yarkovsky acceleration in an optics-only orbit fit over three successive apparitions.

Failing to account for the effect of DCR bias in orbit predictions for near-Earth asteroids can significantly skew the calculation of impact probabilities and bias the geographic location of estimates of known future impacts. In a scenario modeled on the orbit of 99942 Apophis, an optics-only orbit solution formed from observations between 2016 and 2029 would predict at least a 10% probability of impact in 2036 if DCR bias were fully corrected out of the input astrometry, but only a 0.06% chance of impact if DCR bias were not accounted-for. DCR corrections based only on albedo information and not spectroscopic measurements can in some circumstances raise that probability back above the 5% level, but in the case of an ambiguity between an S-type and Q-type object with comparable albedo, the best-case impact probability estimate is 0.14%.

Raising the elevation limit on allowable observations lowers the difference between impact probability estimates computed from biased versus unbiased data, but at the expense of overall accuracy. For near-Earth asteroids which are more readily observable at low elevation angles, the gain in observation volume and precision of the orbit estimate achievable from incorporating low-elevation data can be significant, even when only the azimuth component of low-elevation data is used in the

orbit solution.

For the Apophis-like scenario, an impact point estimate derived from measurements between 2029 and 2036 remains biased by over 100 km away from its true location with high confidence until only a few years before impact. Similar results were found in any scenario where the impact point estimate is derived from observations made at low elevation angles in the months leading up to a known impact.

In scenarios demonstrating the theme of DCR bias interfering with the planning of a scientific measurement, unaccounted-for DCR was shown to reduce the probability of successfully predicting a stellar occultation by a Kuiper Belt Object, to near zero for some orbits, and limiting the ability to predict a stellar appulse by a geosynchronous satellite to an accuracy of 2 – 3 arcseconds.

In geosynchronous orbit, DCR bias introduces a bias into the solar radiation pressure area-to-mass ratio estimate of an object observed over a single night. While this bias is radically non-physical for large satellites, it is on par with the AMR of debris, and can limit the time-horizon of accurate predictions of the orbits of this class of objects.

Chapter 6

DCR Compensation of Observations of GPS and GLONASS Satellites with Concurrent Spectroscopy at the Firepond Optical Facility

The final contribution of this thesis are the results of a spectroscopic measurement campaign of Earth satellites observed under a variety of illumination and atmospheric conditions and a demonstration of DCR compensation using those measurements. Over the several years leading up to the measurement campaign, a number of concepts for the measurement were considered before settling on an easy-to-build slitless instrument built around commercial off-the-shelf (COTS) equipment. This chapter will discuss the selection process and the design considerations for the spectrograph, present the results of preliminary measurements on satellites used to motivate the design decisions, and describe in detail the measurement campaign conducted at the Firepond Optical Facility to demonstrate DCR compensation of observations of GPS and GLONASS satellites using concurrent spectroscopy.




Figure 6-1: Dispersive Optical Elements

6.1 Slitless Spectroscopy

6.1.1 Spectrograph Design Considerations

6.1.1.1 Mechanisms to Divide Light

The spectroscopic resolution requirement for DCR compensation derived in Section 4.6 calls for a resolution of 40 nm over a passband of 800 nm in order to achieve the design goal of DCR estimation to an accuracy of 20 mas. This is a 20-channel measurement and there are two ways to make it: by multiplexing between 20 blocking filters or by dividing the light. While for slow-moving and slowly-varying astronomical targets, it may be appropriate to multiplex blocking filters, it is impractical to multiplex between 20 blocking filters when observing satellite targets and when using equipment readily obtainable and usable at the Firepond facility. This restricts consideration to the use of dispersive elements: prisms, reflective gratings, and transmission gratings. The basic functionality of these elements is shown in Figure 6-1.

Each of these elements has their respective advantages and disadvantages. A prism achieves high efficiency in terms of not throwing away light, but this comes at the expense of requiring a large spatial volume. Additionally, a system built around a prism requires tighter dimensional tolerancing than gratings because a prism completely divides incoming light into its spectral components without preserving any fiducial ray paths that localize the origin of the spectrum. While this is not a problem in designs with in-built kinematic stability between the input light source and the dispersive element, for an instrument observing a moving target from a moving platform, this is not the best option.

Diffraction gratings, both reflective and transmissive, provide an alternative to the prism that reduces the volume requirement. A diffraction grating operates by introducing a regular spatial modulation to an optical element like a window or a mirror at a length scale of the modulation on the order of magnitude of a wavelength. The modulation breaks up the incoming wavefront, and by

the Huygens principle, generates outgoing wavefronts centered at the grating features. The spatial regularity of the grating features causes constructive interference to occur in certain preferential directions, termed “spectral orders.” The presence of a 0th order along the nominal beam path that would exist in the absence of the modulation preserves the fiducial beam path and allows in-situ registration of the output spectrum’s origin point, thereby overcoming a deficiency of the prism.

6.1.1.2 Grating Spectrographs

A fundamental property of the grating is that if constructive interference occurs for any wavelength at a particular output angle, then integer multiples of that wavelength will also constructively interfere in at an integer multiple of that output angle. Thus, for every wavelength, energy will necessarily be deposited into all orders. The geometry of the modulating pattern on the grating surface can be selected to vary the amplitude of the signal sent into the distinct orders, but it is not possible to select a grating where all of the energy from the entire silicon passband is deposited into a single spectral order. There will always be some overlap between energy from adjacent spectral orders at wavelengths that are integer multiples.

For this reason, spectrographic instruments that need to sample a passband wider than a multiple of its lower edge wavelength must further divide the light. In an echelle spectrograph, this division can occur after the grating, where a prism oriented at right angles to the grating’s dispersion separates the overlapping orders [134, 135]. The echelle design is more appropriate for applications in astronomy requiring high spectral resolution and is not necessary for the coarse resolution requirement of DCR compensation. Alternatively, the division can occur before the grating, with multiple spectrographs observing different parts of the target passband [136, 137].

This latter approach, shown in Figure 6-2(a), was considered for this thesis. This concept consisted of two COTS slit spectrographs, one for measuring the blue end of the silicon band from 300 nm to 550 nm and one for measuring the remainder of the band from 550 nm to 1100 nm, without spectral order overlap within either band. A single Oriel MS125 spectrograph and detector was already available and it would have been necessary to purchase a second unit and detector for a minimum cost in excess of \$17,000.

The purpose of a slit spectrograph is to spatially filter out other sources in the field of view so that only light from the target source is divided by the grating and measured by the detector. A further benefit of the slit is suppression of sky background, which, as discussed in Section 4.6.3, has the effect of boosting the SNR of the measured spectrum and reducing the time-aperture requirement

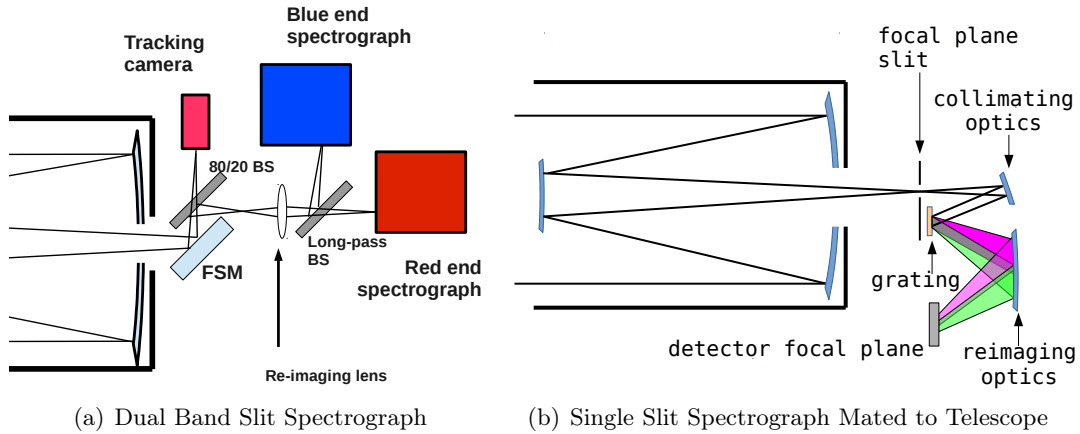
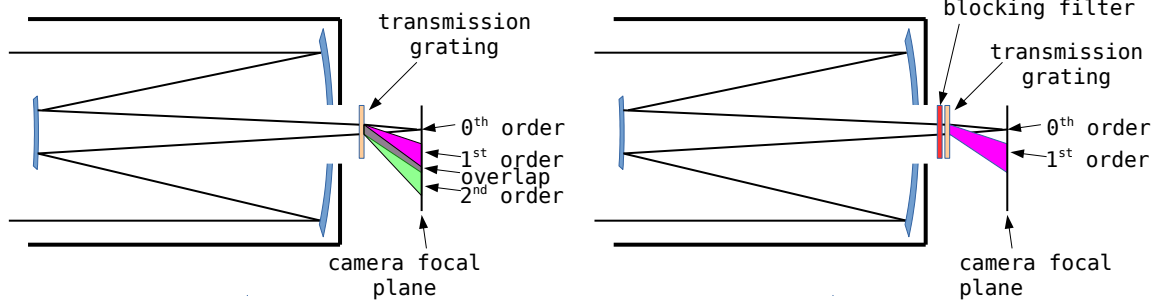


Figure 6-2: Slit Spectrograph

on a measurement. A complication with slit spectrographs is the necessity to deterministically locate the target over the slit at the focal plane of the foreoptics. The width of the slit is usually selected to match the seeing achievable at the site, which for the Firepond facility is on the order of 5 arcsec full-width half-max (FWHM). For slow-moving astronomical targets, it is not challenging to locate and keep the target over the slit to that level of accuracy, but for faster-moving satellites, transient mechanical disturbances present a greater challenge to keeping the instrument on target continuously. A fast steering mirror with a guiding camera would be necessary in such a design, at a cost of several thousand dollars more.

With a single steering mirror, the mechanical tolerances on the allowable deformation in the overall optical backend would be dictated by the allowable misalignment between the split blue and red beams and the two slits of the two spectrographs. With foreoptic focal lengths on the order of 2 – 3 meters, and slits measuring a few arcseconds across, this dimensional tolerance works out to a few tens of microns across a distance of a few tens of centimeters between the two slits. The instrument would need to be placed on a moving telescope exposed to ambient outside conditions. This means the tolerance would need to be met across mechanical orientations from near-vertical to near-horizontal and across a temperature range of several tens of degrees Celsius corresponding to the difference between an indoor assembly conditions and a typical cold winter night at the Firepond facility. For reference, the thermal expansion coefficient of aluminum is on the order of 20 microns per meter-Kelvin and for steel it is on the order of 10 microns per meter-Kelvin, meaning the worst-case thermal drift between the two slits could be on the order of several tens of microns.

While it is certainly possible to meet the mechanical alignment and stability requirements for



(a) Slitless Spectrograph Mated to Telescope

(b) Placement of Blocking Filter

Figure 6-3: Slitless Spectrograph. The final design of the slitless spectrograph for this thesis consists of two pairs of matched gratings and blocking filters selected to individually measure the blue and red halves of the silicon passband (Section 6.4).

a dual-band slit spectrograph design, the design effort necessary to ensure that the requirement was met was deemed excessive given the time, funding, and engineering resources available for the execution of this thesis. For this reason (and from experience with data from NOFS, as described in Section 6.2), the slitless design (Figure 6-3) was chosen for the measurements campaign at the Firepond facility.

6.1.2 Transmission Grating Geometry

COTS transmission gratings for spectrographs are blazed to optimize the placement of the energy from a specified passband into the first spectral order. This is characterized by a blaze wavelength, for which all of the energy goes into the first spectral order. The geometry of the grating and the ray paths is shown in Figure 6-4(a).

Assuming normal incidence on the flat side of the grating, the grating equation relating the wavelength and the output angle is:

$$m\lambda = \Delta_g \sin \theta_b \quad (6.1)$$

where m is the spectral order.

To relate the blaze and groove angles, note that the blaze angle θ_b is related to the groove angle θ_g via:

$$\theta_b = \theta_o - \theta_g \quad (6.2)$$

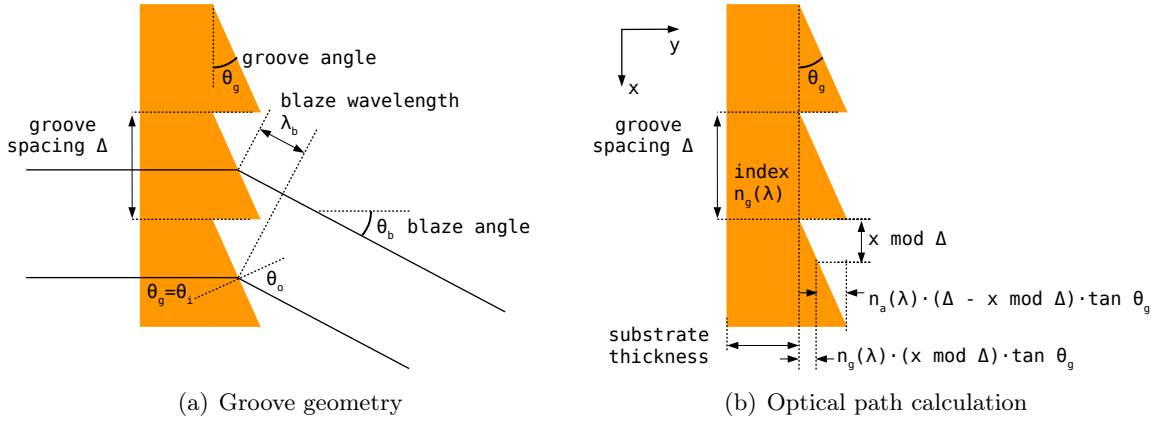


Figure 6-4: Transmission Grating Geometry

Noting that:

$$\theta_i = \theta_g \quad (6.3)$$

and applying Snell's law at the output surface,

$$n_g(\lambda) \sin \theta_g = n_a(\lambda) \sin \theta_o \quad (6.4)$$

and putting everything together, the geometry of the grating is fully characterized by the grating equation and:

$$\theta_g = \tan^{-1} \frac{n_a(\lambda_b) \sin \theta_b}{n_g(\lambda_b) - n_a(\lambda_b) \cos \theta_b} \quad (6.5)$$

The procedure to analyze any grating is to compute the blaze angle from the specified groove spacing and blaze wavelength, and to then compute the groove angle using that blaze angle and the specification of the grating material's index of refraction at its blaze wavelength.

6.1.3 Grating Efficiency Calculation

An optical element inserted into a converging beam of an imaging system gives an output at the focal plane that is a spatially scaled Fourier Transform of the optical element's transmission function [138]. For a purely transmissive element, the transmission function is an optical path difference, which is a pure phase modulation. In reference to Figure 6-4(b), the optical path of a transmission grating is computed as follows.

The calculation of total optical path is made with reference to an input and output plane orthogonal to the optical axis. The input plane is coincident with the grating's flat surface and

the output plane is coincident with the tops of the grooves. Let the refractive index of the grating substrate be denoted $n_g(\lambda)$, the substrate thickness t , the groove angle θ_g , and the groove spacing Δ . Upon entry into the grating, the optical path through the flat part of the substrate is given by $t \cdot n_g(\lambda)$. The path through the groove is given by $(x \bmod \Delta) \tan \theta_g \cdot n_g(\lambda)$. The remaining path through free space (air, index $n_a(\lambda)$) from the end of the grating material to the output plane is given by $[\Delta - (x \bmod \Delta)] \tan \theta_g \cdot n_a$.

Summing these components gives:

$$\text{path} = n_g(\lambda) \cdot (t + (x \bmod \Delta) \tan \theta_g) + n_a(\lambda) \cdot [\Delta - (x \bmod \Delta)] \tan \theta_g \quad (6.6)$$

$$\text{path} = \underbrace{n_g(\lambda) \cdot t + n_a(\lambda) \cdot \Delta \tan(\theta_g)}_{\text{spatially invariant}} + \tan(\theta_g)(n_g(\lambda) - n_a(\lambda)) (x \bmod \Delta) \quad (6.7)$$

The spatially invariant component of the optical path is a constant uniform phase offset and has no effect on the form of the output energy distribution at the imaging plane. Phase scales at $2\pi/\lambda$, so total phase delay function for the grating is given by

$$T(x) = e^{j \tan(\theta_g)(n_g(\lambda) - 1) (x \bmod \Delta) \frac{2\pi}{\lambda}} \quad (6.8)$$

The Fourier transform of a periodic function is given by the Fourier transform of an individual periodic element multiplied by an array factor. The array factor is a high-frequency periodic comb function, whose maxima for any single wavelength correspond to the output angles of the individual spectral orders, while the Fourier transform of the individual array element is a slowly-varying amplitude envelope that modulates the strength of the respective spectral orders at each wavelength [139]. The magnitude-squared of this envelope, evaluated at all the spectral orders for any wavelength, is the grating efficiency of that wavelength at that order.

The Fourier transform of the phase ramp function $T(x)$ is given as follows:

$$\mathcal{F}\{T(x)\} = \int_{-\frac{\Delta}{2}}^{\frac{\Delta}{2}} e^{-j2\pi x f_x} e^{j \tan \theta_g (n_g - n_a) x \frac{2\pi}{\lambda}} dx \quad (6.9)$$

$$= \frac{\Delta \left(e^{j2\pi \left(f_x - \tan \theta_g \frac{(n_g - n_a)}{\lambda} \right) \frac{\Delta}{2}} - e^{-j2\pi \left(f_x - \tan \theta_g \frac{(n_g - n_a)}{\lambda} \right) \frac{\Delta}{2}} \right)}{2j2\pi \left(f_x - \tan \theta_g \frac{(n_g - n_a)}{\lambda} \right) \frac{\Delta}{2}} \quad (6.10)$$

$$= \Delta \cdot \text{sinc} \left(\pi \Delta \left(f_x - \tan \theta_g (n_g - n_a) \frac{1}{\lambda} \right) \right) \quad (6.11)$$

The spatial frequency f_x corresponds to a spatial coordinate at the imaging plane. If the grating is at a distance d from focal plane, then the focal plane u, v coordinate for that frequency is given by $f_x \rightarrow \frac{u}{\lambda d}$, where u is evaluated at the center of the 0th, 1st, 2nd order for a given wavelength. That is,

$$u = d \tan \theta_m = d \tan \left(\sin^{-1} \left(\frac{m\lambda}{\Delta} \right) \right) \quad (6.12)$$

In the small angle approximation, $\tan x \approx x$ and $\sin^{-1} x \approx x \implies f_x = m/\Delta$.

Putting it all together, the grating efficiency for the m th order at wavelength λ is the amplitude-squared of the Fourier Transform of the phase ramp evaluated at the center of the order for that wavelength:

$$I(\lambda) \propto \left(\text{sinc} \left(\pi \left(m - \frac{\Delta}{\lambda} \tan \theta_g (n_g(\lambda) - n_a(\lambda)) \right) \right) \right)^2 \quad (6.13)$$

Equation 6.13 refers to an intensity per unit wavelength. A spectrum measured from gratings with different amounts of dispersion (i.e. grooves per mm) will have also have a scaling factor to reflect that difference in dispersion such that conservation of energy holds.

6.2 Slitless Spectroscopy of Geostationary Satellites at the US Naval Observatory Flagstaff Station

Part of the motivation in choosing a slitless design for the Firepond facility was experience with the results of slitless spectroscopy of geostationary satellites made at NOFS. In concert with the observations made at ETS in January/February 2018 (Section 5.5.5), the 1.3 m slitless spectrograph instrument at NOFS operated over two nights observing the same parts of the GEO belt (though not simultaneously). The NOFS slitless spectrograph instrument is used for RSO identification, rather than astrometric compensation, and consists of a Richardson 906R transmission grating on uncoated BK7 glass located approximately 55 mm in front an e2v Corporation silicon detector. Satellite exposures of one minute are used and the resulting low-resolution spectra are valid redward of 450 nm and contain mixtures of the $m = 1$ and $m = 2$ spectral orders.

The approximate sampling resolution is approximately 9 nm per pixel. The angular pixel scale is approximately 0.6 arcsec/pixel [140]. The NOFS 1.3 m telescope is an equatorial telescope and the grating is oriented to align the spectrum with the North/South axis so that when geostationary satellites are observed, the star streaks cross perpendicular to the spectrum (Figure 6-5). This means that the spectrum is subjected to DCR itself: shorter wavelengths will fall into spectral bins

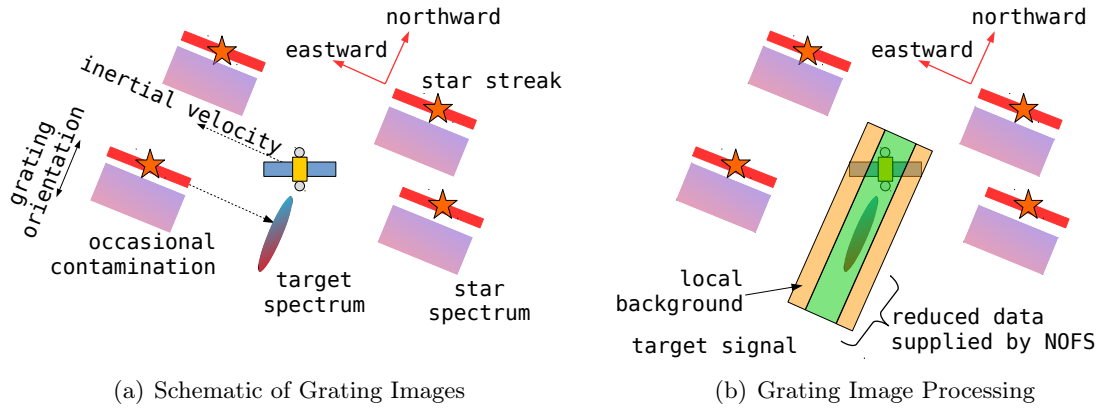


Figure 6-5: Schematic of NOFS Spectrograph Data Product. In slitless spectroscopy, the target and its spectrum are simultaneously present in the image along with the centroids and spectra of all other sources in the image. Extracting the spectrum requires identifying the desired target centroid to find the origin of the spectrum and reading out the difference between the pixels containing the spectrum and their local background. For a moving target especially, the computation of a local background for each spectrum pixel is necessary to mitigate the effect of contamination by star and star spectrum streaks through the desired target’s spectrum.

Table 6.1: Geostationary Satellites Observed by NOFS Spectrograph

SCC	Name	Elevation [deg]
24812	Galaxy-25	44
28154	SES	42
29155	GOES-13	22
36411	GOES-15	42
41866	GOES-16	34

higher than nominal while longer wavelengths will fall into spectral bins lower than nominal.

For the four satellites observed above 30 degrees elevation, the worst-case refraction difference between the 400 nm and 1000 nm ends of the spectrum is about 3 arcsec, or five pixels, giving an error of about 2.5 pixels of 22 nm of wavelength in the identification of input wavelengths. For the one GOES satellite observed at low elevation, the worst-case error is greater at 4 arcsec, but at that low elevation, the declination direction is more aligned with the horizontal, attenuating the geometric error from DCR.

6.2.1 Slitless Spectroscopy Processing

Equation 6.13 was used to compute the nominal grating efficiency for the overlapping $m = 1$ and $m = 2$ orders for the NOFS configuration. Without blocking filters present, and with only one grating, there is potentially significant overlap around of wavelengths between 300 nm and 400 nm

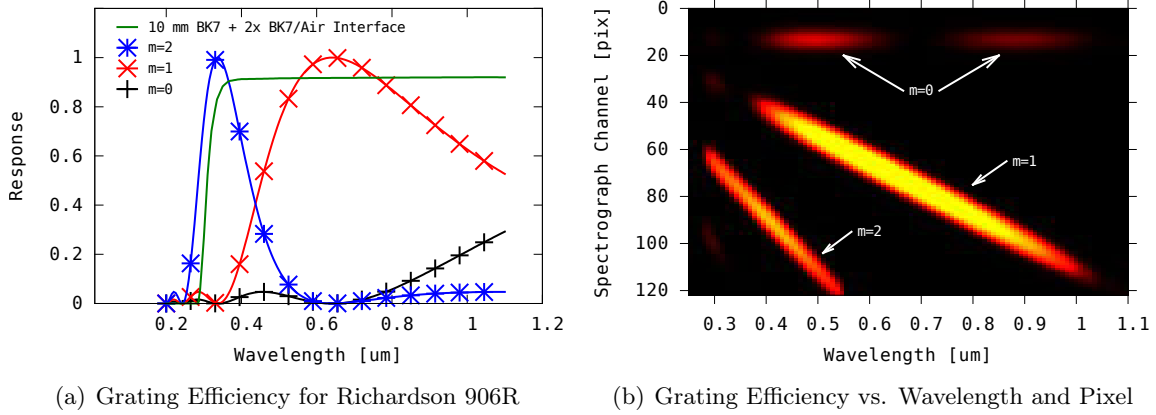


Figure 6-6: Grating Efficiency for NOFS Slitless Spectrograph

with signal from 600 nm to 800 nm. The degree to which such overlap actually exists depends on the energy content of the satellite signature between 300 nm and 400 nm.

The data from NOFS is presented as a single linear vector of measured fluxes. The distribution of true flux into idealized 9 nm wavelength bins may also be thought of as a linear vector and the transformation of that true flux into measured signal by the action of the grating may be thought of as multiplication by a matrix whose columns correspond to the grating efficiencies in the relevant spectral orders for each wavelength. Figure 6-6(b) shows this matrix for $m = 0, 1, 2$, which captures the flux distribution as measured by the NOFS system.

This matrix is not directly invertible into a spectrum of the target flux over the full silicon band but does contain sufficient information to resolve satellite signatures into a mixture of basis materials. Kapton, aluminum, and solar panel basis functions were chosen to represent the satellite spectra. The transfer function from material ratios to spectrograph data was computed with MODTRAN [123] estimates of atmospheric transmittance, handbook values for spectral reflectivity of aluminum-coated mirrors and transmissivity of the dewar window and grating substrate, and focal plane manufacturer datasheet spectral quantum efficiency.

Figure 6-7 shows a sample reconstruction of the measured signal based on the assumption that the signature is composed entirely of these three basis functions. The reconstructions are cosmetically good, but they are imperfect. The measurement is not an ideal one as there are several potential sources of systematic error:

1. Grating non-orthogonality or a change in the actual distance between the grating and focal plane will change the grating efficiencies used to compute the nominal transfer function matrix;




Figure 6-7: Reconstruction of NOFS Spectra from Basis Functions for Two Satellites


2. Unaccounted-for DCR combined with an unknown orientation of the grating (spectrum above or below the target) introduces a possible systematic shift of a few tens of nm in the assumed transfer function along the spectrum;
3. DCR in the direction perpendicular to the spectrum may cause energy to leak out away from where the NOFS software expects it to be;
4. Deviation of atmospheric conditions from the nominal model generated by MODTRAN may introduce a systematic bias to the transfer function and resulting reconstruction;
5. Unmodeled optical coatings present anywhere in the optical system can also introduce a systematic bias to the reconstruction;
6. The model basis functions assumed for the material properties of the satellites may deviate from the spectral reflectance of the actual materials in the particular satellites observed and the satellites may contain non-negligible fractions of materials that are not one of the three chosen basis functions.

6.2.2 Results from NOFS Measurements


Despite the imperfect reconstruction, the fact that the satellites are geostationary and observed through the same atmospheric conditions means that it is sensible to compare the reconstructed material ratios among the satellites against each other. As suggested by its mostly-positive one-night AMR bias (Table 5.8), Galaxy-25 (SCC 24812) resolves into a mostly solar panel signature (Figure

6-8(a)). While many satellites' relative composition of the three materials remained constant over the observing period, GOES-15 (SCC 36411) showed rapid variation (Figure 6-8(b)). This variation is attributed to the strong phase angle dependence of solar panel spectral reflectance.


Equation 1.2 was used to compute a nominal DCR bias from this mixture of materials. Figure 6-8(c) shows the computed DCR bias for all observed satellites as a function of RA offset from the anti-solar direction, highlighting the strong time and phase angle dependence of DCR for active satellites with solar panels that track the Sun.



(a) 2018:032 Galaxy-25 Material Ratio



(b) 2018:032 GOES-15 Material Ratio



(c) DCR From Material Ratios

Figure 6-8: NOFS Spectroscopy of Selected Geosynchronous Satellites

6.2.3 Conclusions

This sampling shows that on the GEO belt, DCR bias can vary by over 100 mas for different satellites and in the case of SCC 28154 it can vary by 100 mas for the same satellite observed on two subsequent nights at the same phase angle. As shown in the case of GOES-15 in the time leading up to opposition, DCR bias can change by 100 mas on the timescale of ten minutes. This demonstrates that DCR bias for GEO RSOs is not predictable without concurrent spectroscopic measurements and gives a notional requirement for the cadence of the spectroscopic measurement.

An important caveat is that because spectra were not measured blueward of 450 nm, the material ratios and the quantities derived from those material ratios should not be considered estimates of absolutes. The relative variation of material ratio estimates and computed DCR bias, however, may be interpreted as a measure of DCR variation among the GEO population with respect to a fixed (but unmeasured) zero point.

6.3 Multi-band Photometry of GPS and GLONASS Satellites

The partial coverage of the silicon band of the NOFS measurement gives notional values for the variability of DCR bias in the GEO population. However that measurement lacks coverage in the blue end of the silicon band where refraction and DCR contributions are greatest. As such, the NOFS measurement was only suitable for gauging the variability in DCR among the GEO population and for individual satellites over time and not for computing absolute DCR with respect to solar-analog. Prior to commencing assembly of the slitless spectrograph for the Firepond facility, a conceptual follow-on experiment to the NOFS measurement, and a path-finding measurement for the Firepond spectrograph, was made against GPS and GLONASS satellites using multiband photometry covering the full silicon passband.

6.3.1 Scenario

On the night of May 24, 2018, the Firepond Facility observed two GLONASS satellites nearing opposition and entering eclipse¹. Bare silicon CCD imagery collected at 8 Hz was reduced using the Gaia DR2 star catalog and solutions containing at least 15 stars were retained for analysis.

The residuals against the IGS final orbits, resolved into station-aligned horizontal and vertical coordinates are shown versus elevation in Figure 6-9. The RMS residual of individual observations

¹See Figure 6-48 in Section 6.8.1 for an illustration of the satellite illumination conditions in this measurement.




Figure 6-9: Mean $\pm 3\sigma$ Astrometric Residuals of Two GLONASS Satellites vs. Elevation Angle. N=14758 distinct observations of SCC 32276 and N=33642 distinct observations of SCC 40001.


due to atmospheric seeing was about 380 mas in this dataset. As residuals are averaged by elevation angle bins, the random error quickly attenuates, leaving a near zero-mean horizontal residual. This measurement was made shortly after the release of Gaia DR2 and the near zero-mean horizontal residual is indicative of a high-quality astrometric solution with the Gaia DR2 catalog.

The structure in the vertical residual is indicative of the variability of DCR bias as a function of target spectrum. As the satellites neared opposition, the blue glints from their solar panels manifested in a positive elevation residual. As they entered eclipse, the shorter wavelengths of solar illumination were filtered out by the atmosphere, resulting in a progressively redder target spectrum. This is evident in both satellites' negative elevation residual before loss of signal.


The collection on GLONASS satellites was accompanied by a second co-boresited 6-inch telescope collecting multicolor photometry by cycling through a full set of Sloan filters every 75 seconds. The five Sloan filters cover the full silicon band and are about 160 nm wide [128]. This is insufficient for reconstruction of DCR to 25 mas accuracy, but it was a measurement that could be made quickly with equipment already on-hand to demonstrate the utility of real-time DCR compensation with measurements over the full silicon band.

6.3.2 Results from Multiband Photometry Experiment

Figure 6-10(a) shows the normalized flux in the five Sloan bands as a function of time for the two satellites as they enter eclipse. The majority of the satellites' flux prior to eclipse falls within the redder bands with little signal measured in the shorter wavelengths. This may be interpreted as a spectrum containing a mix of Kapton signature and non-glinting solar panel. As the satellites



(a) Multiband Photometry





(b) DCR Prediction and Actual Astrometric Bias vs. Time

Figure 6-10: Multiband Photometry and DCR Estimation of Two GLONASS Satellites

near opposition just before entering eclipse, they both exhibit blue glints from their solar panels as indicated by an increased fraction of their flux that is present in the u' and g' bands. This transition occurs over a period of about ten minutes. Once the satellites enter eclipse, their absolute brightness rapidly drops and a majority of the remaining signal is made up of long wavelength photons captured in the i' and z' passbands.

DCR bias was estimated from the multiband photometry data using a discretized version of Equation 1.2. Figure 6-10(b) shows that the trend in the predicted DCR bias and the actual DCR bias correspond closely. The discrepancy in the DCR prediction for 32276 and the excess positive bias prediction for 40001 at the nominal time of entry into eclipse are attributed to the time lag of the photometric measurement. Cycling through all filters took 75 seconds to complete while the timescale for entry into eclipse was also on the order of a minute. Nevertheless, agreement between the trends of the DCR prediction and the observed vertical residual in the time leading up to eclipse demonstrates that DCR bias may be corrected with simultaneous spectroscopic measurement.



(c) Bracketry Built at Firepond. Dimensions in inches. Credit: Paul Pepin / MIT LL

Figure 6-11: Equipment Available at the Firepond Facility

6.4 Design of the Firepond Slitless Spectrograph

With the experience with the NOFS slitless data and the determination that a two-channel slit spectrograph was not feasible for the completion of this thesis, a slitless design based on COTS equipment was chosen. The Firepond facility had as available equipment a 16-inch Meade LX200 telescope with a Finger Lakes Instruments Centerline CL-1-10 dual filter wheel and bracketry to accept one of the several Andor iXon3 DU888 cameras with 1024×1024 13 micron pixels also available for use at the site. This equipment (Figure 6-11) had been purchased and used in 2012, and was suitable for conversion into a slitless spectrograph with the purchase of transmission gratings and blocking filters.

The limitations imposed by the design of that equipment place the focal plane approximately

10 cm behind the gratings. This fixed distance and the 13 mm size of the camera focal plane necessitated a low-dispersion grating. Following the example of the NOFS system, Richardson gratings were considered for this design, and the lowest-dispersion stock items in their catalog were the 906R grating with 35 grooves per mm blazed for 640 nm and the 606R grating with 30 grooves per mm blazed for 405 nm.

With the nominal separation of 10 cm between the gratings and the focal plane, the 906 grating would give a wavelength scale of 3.7 nm / pixel on the detector and the 606 grating would give 4.3 nm / pixel. At this scaling, a full spectrum from $m = 0$ to 1100 nm would fit into approximately 300 pixels. The target spectral resolution for this measurement was 40 nm and this grating selection resulted in a dispersion roughly ten times larger than necessary relative to the pixel grid. Nevertheless, the angular size of each pixel in this telescope is about 0.68 arcsec, and typical seeing at the Firepond facility (optimistically estimated as about 5 arcsec FWHM, or eight pixels), meant that the target resolution wavelength bins of about ten pixels across would be just big enough so that the measurement would not be degraded by seeing.

Four 50 mm square gratings on 3 mm BK7 substrate were ordered from Richardson: two uncoated 906 gratings and two AR-coated 606 gratings. The total cost of the gratings was \$2,090. One set of gratings would be placed to orient the spectrum horizontally and one set would be placed to orient the spectrum vertically. The vertically-oriented spectrum would be placed with blue on top, so that DCR stretches the spectrum rather than collapses it. The two orientations are meant to track satellite targets that move mostly perpendicular to the direction of the spectrum in order to minimize the effect of star streaks cluttering up the data with spurious signals near-parallel to the direction of the spectra. The vertical spectra were intended to make measurements of geostationary satellites while the horizontal spectra were intended for use with GPS and GLONASS satellites, the majority of whose apparent motion relative to the stars is in the vertical direction.

With gratings selected, colored glass blocking filters were selected from the Newport Corporation catalog. An FSQ-BG40 filter was selected to measure the shorter wavelengths in concert with the 606 grating while for the longer wavelengths with the 906 grating, both an FSQ-OG515 and an FSQ-OG570 filter were purchased. All three blocking filters cost approximately \$75 each.

The experience with the 906 grating NOFS data showed that the overlap between the second order of shorter wavelengths and the first order of the band from 600 nm to 1100 nm can be significant, necessitating strong suppression of short wavelengths when using the 906 grating. Uncertainty in the true versus the advertised shape of the cut-on transmittance of the orange glass filters motivated

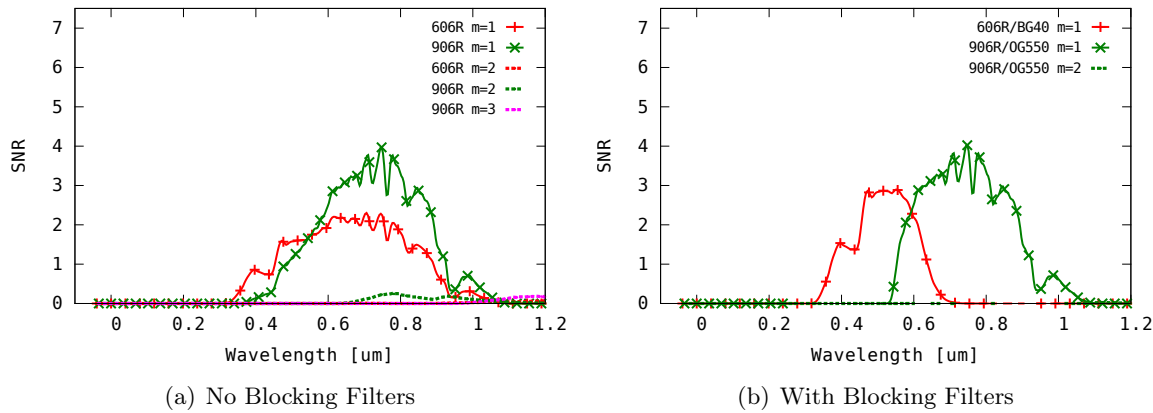


Figure 6-12: SNR Along Spectrum for 12.5 mag 80% Solar Panel Target

the purchase of both the OG515 and OG570 filters. Upon receipt and bench testing of the filters, it was determined that the OG570 filter was the more appropriate to use in this instrument and all data collected with the fielded spectrograph utilized the BG40 and OG570 filters only.

6.5 Performance Predictions

6.5.1 SNR and Suppression of Contamination by Higher Spectral Orders

The purpose of this instrument is to make measurements on GPS and GLONASS satellites. Past experience observing these calibration targets indicated that their expected brightness should be between 10.5 mag and 12.5 mag. From the experience with the NOFS data, geostationary satellites are between 60% and 80% solar panel. In terms of predicting the ability of the spectrograph to accurately reconstruct DCR, the most stressing condition would be a satellite with the larger fraction of solar panel. Kapton is highly reflective over the red portion of the silicon passband, but the solar panel blue reflectance peak coincides with the Rayleigh-scattering roll-off in atmospheric transmittance. The interplay between this roll-off, the amount of energy in the blue, and the fact that atmospheric refraction is greatest in magnitude for short wavelengths is what determines the behavior of the total DCR bias for a satellite signature composed of mostly solar panel. Taking this all together, a spot-check performance prediction for the instrument was computed for a 80% solar panel, 20% Kapton target with an overall brightness of 12.5 mag observed at 30 degrees elevation.

Figure 6-12 shows the per-pixel SNR along the spectrum for this target assuming a one second exposure of the 16-inch telescope, a sky background of 17.5 mag/arcsec², and a (very optimistic)

seeing disk that is three pixels FWHM. The calculation is made for both the 606 and 906 gratings with and without blocking filters. Given the initial uncertainty in which orange glass filter (OG515 or OG570) was going to be appropriate, the calculations were made using an OG550 blocking filter for measurements over the longer wavelengths in the silicon passband.

The purpose of the spot-check calculation was to determine whether the selected blocking filters provide sufficient suppression of contamination between $m = 1$ and $m = 2$ spectral orders for a target with a lot of energy in the blue half of its detectable spectrum, and the necessity of the blocking filters is evident especially with the 906 grating. Without a blocking filter, a significant amount of signal in the $m = 2$ and $m = 3$ spectral order contaminates the measurement redward of 700 nm. With the orange glass blocking filter, it is possible to obtain an uncontaminated measurement of the red part of the spectrum at high SNR with the 906 grating.

The inclusion of the blue glass blocking filter also boosts the SNR of the spectral measurement in the blue part of the passband with the 606 grating. Because a flat background spectrum had been assumed in the calculation, this is not caused by any spectral effects in the sky background or transmittance; rather it is simply a result of the fact that the BG40 passband is only about 300 nm wide and cuts out the sky background over the region where the silicon detector's quantum efficiency is high. The improvement in SNR for the 906 grating with the orange glass blocking filter is only marginal, as the amount of detected sky background photons cut out by the blocking filter is smaller.

While there is slight contamination of higher spectral orders with the 606 grating, it is confined mostly redward of 1000 nm. On the one hand, this implies that it ought to be possible to make the measurement with the 606 grating only, at the expense of approximately a factor of 2–3 loss in SNR across the spectrum. Alternatively, the lack of contamination around the cut-off and cut-on wavelengths of the blocking filters around 600 nm implies that the 606 may be used for an in-situ calibration of the filter transmission functions in that region, obviating the need for laboratory testing of those optical components.

6.5.2 SNR of $m = 0$ PSF

A practical consideration in slitless spectroscopy is the ability to locate the origin of the target spectrum within the focal plane. The measurements at NOFS on geostationary satellites were made with a 1.3 meter diameter aperture and exposures of about 1 minute long. In those images, there was always sufficient SNR on those targets to enable automatic localization of the target




Figure 6-13: Comparison of SNR Along Spectrum for 12.5 mag 20% and 80% Solar Panel Target

spectrum origins. This may not be the case with the smaller aperture telescope tracking GPS and GLONASS satellites. While geostationary satellites require minimal mount motion, GPS and GLONASS satellites move at about 35 – 40 arcsec/sec. The mechanical perturbations from that motion will likely be too high to allow exposures as high as one minute. For this reason, it is necessary to predict whether the target spectrum origin will be detectable in short-exposure images.

In Figure 6-13, the SNR calculation is made for the $m = 0$ mode for both an 80% solar panel and a 20% solar panel target, with blocking filters in place. With a one-second exposure and an optimistic seeing of 2 arcsec FWHM, the SNR of the brightest pixel in the $m = 0$ centroid ranges between 5 – 6.5 depending on the grating and blocking filter used. Although this is technically detectable, the assumptions in this calculation are optimistic. A degradation in seeing from 2 arcsec to 4 arcsec will reduce the pixel peak SNR to below the commonly-adopted minimum detection threshold of 3. A reduction in exposure time from 1 second that might be necessitated by mechanical jitter would also reduce this SNR.

Because for targets at the faint limit of 12.5 magnitude it may not be possible to locate the $m = 0$ PSF in the focal plane to act as a fiducial point for extracting the spectra, an alternate strategy will be necessary for automatically extracting the spectra. The strategy planned for this instrument will be to multiplex image frames collected without gratings or filters between spectrographic data collections. These “clear/clear” astrometry frames enable detection and measurement of the satellite targets’ focal plane positions over time, and the locations of the spectra during spectrographic frames can be inferred from the time-history measured during the clear/clear frames.

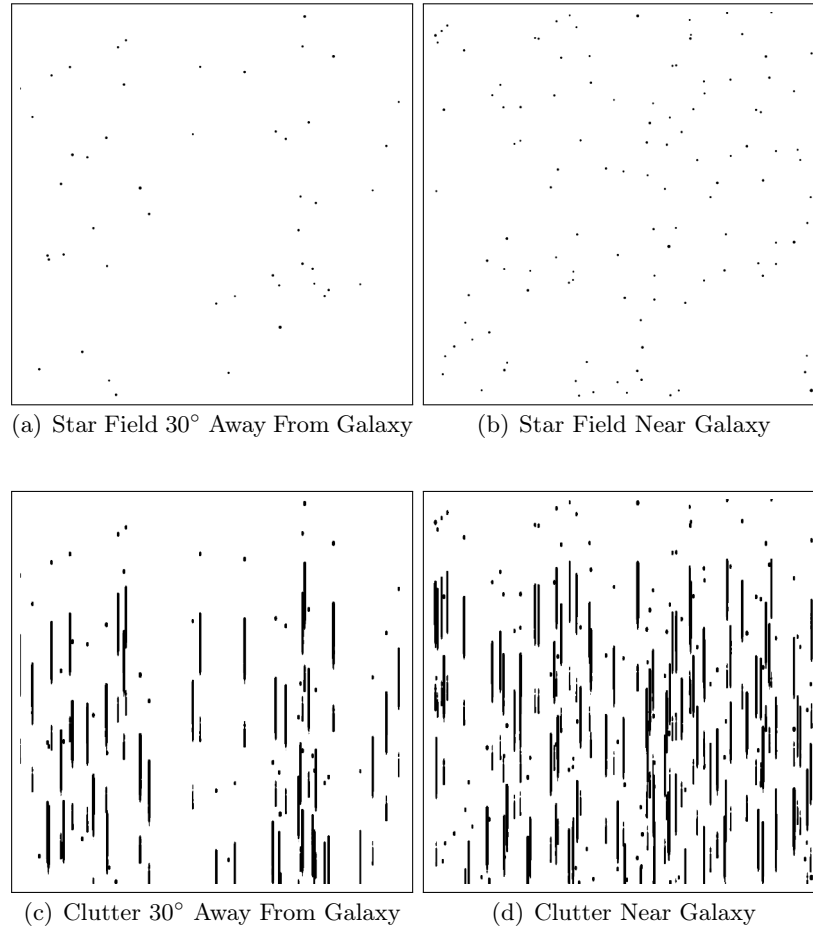


Figure 6-14: Simulated Effect of Star Clutter on Slitless Spectroscopy of a Satellite.

6.5.3 Clutter

The last consideration in slitless spectroscopy is the effect of star clutter. To address the degree to which star clutter will contaminate the measured spectra, a simulated star field was convolved with a linear mask representing the regions where the spectrum generated by the grating would have an SNR above 3 for a bright star. The star field was drawn from the UCAC5 star catalog along the path of a geostationary satellite at two points in its trajectory: 30 degrees away from the galactic plane (Figure 6-14(a)) and crossing the galactic plane (Figure 6-14(b)). The entire catalog has a faint cutoff around 17th magnitude, but one would expect only stars that are within a few visual magnitudes of the target (say 15 or brighter) to contribute a meaningful contamination. Thus the clutter images in Figure 6-14(c) represents a worst-case scenario representative of the majority of the time away from the galactic plane while Figure 6-14(d) represents a worst-case scenario for crossing the galactic plane.

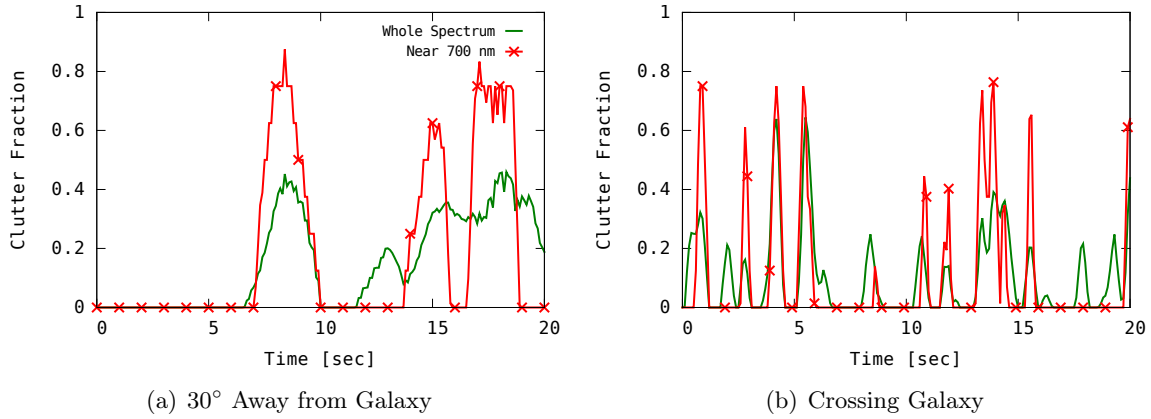


Figure 6-15: Fraction of Spectrum Contaminated by Star Clutter

This scenario was simulated for a twenty-second interval and the fraction of the simulated satellite spectrum contaminated by either a star’s $m = 0$ centroid or its extended spectrum was computed for every simulated frame. The resulting clutter fraction versus time is shown in Figure 6-15. Though there are instances where a large fraction of the entire spectrum of the satellite is potentially contaminated by star clutter, these cases do not always imply that every part of the spectrum is contaminated. The region around 700 nm is illustrated as a sample. About 1/4 of the time when there is some contamination, this region is free of contamination. Further, there are also long stretches lasting a second or more where there is zero contamination of the entire spectrum, even near the galactic plane.

The first implication of this simulation is that it should be possible to compute an uncluttered target spectrum, even in the presence of dense star fields, by taking the median value over time of pixels where the satellite spectrum is expected. As shown in Figure 6-5, the spectrum must be measured with respect to the local background perpendicular to the direction of the spectrum. The percentage of the clutter-free gap times in Figure 6-15 provides a reasonable estimate of the size for the region perpendicular to the spectrum over which the background level can be computed.

Away from the galactic plane, the gaps between contamination can be over five seconds but even near the galactic plane, there are consistent contamination-free gaps of about 2 seconds for a geostationary satellite moving at 15 arcsec/sec across the sky. For the 0.7 arcsec pixels in this instrument, this means that the total gap is about 40 pixels wide, implying that the average could be computed over a ± 21 pixel region, even near the galactic plane. The fact that the overall downtime from clutter over these simulated 20 second intervals is less than half means that a median over at

least 20 seconds should be sufficient to collect a good spectrum under most conditions.

The second implication is that the exposures used to collect the spectrum should be shorter than the characteristic gaps between fully-contaminated regions seen in the simulation. This means that the one-second exposures used to make the performance predictions are potentially too optimistic since the gaps between clutter events can be as short as two seconds. Out of this reasoning, a 500 msec exposure was used in on-sky measurements of satellites.

6.6 Assembly and Laboratory Testing

6.6.1 Test Overview

The purpose of laboratory testing of the spectrograph after assembly is to verify proper functioning of all electronic, mechanical and optical components and to verify the performance predictions generated during the design phase. To accomplish the latter objective, the testing setup depicted in Figure 6-16 was envisioned: a collimated pinhole illuminated by a broadband spectrum would be placed at the front aperture of the spectrograph to simulate a distant point source. The spectrum output from the pinhole would be measured against a calibration standard lamp with a fiber-coupled spectrograph in order to enable validation of the instrument's spectral throughput and grating efficiency predictions.

A free-space segment in the path between the broad-band source and the spot projector would be used to insert narrow band interference filters in order to measure the wavelength scale on the spectrograph focal plane. The use of a free-space segment for the filter, rather than placing the filter between the spot projector and the spectrograph or between the pinhole and the collimator, was necessary to ensure that the introduction of the interference filter (or any other filter) would not turn the beam. Placing the filter at the output of the spot projector at anything other than a right angle to its optical axis would introduce a nonzero shift to the $m = 0$ origin of the spectrum measured in the device under test, thereby removing the utility of introducing the filter for geometric calibration, especially if the filter isolated a wavelength that was readily detectable in the $m = +1$ order but not detectable in the $m = 0$ order.

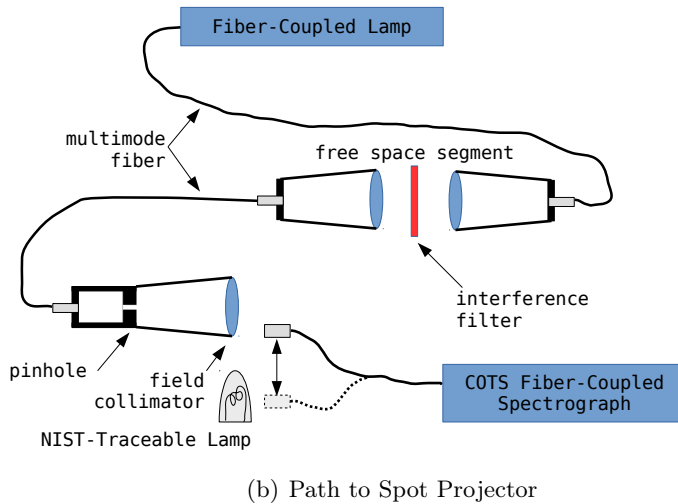
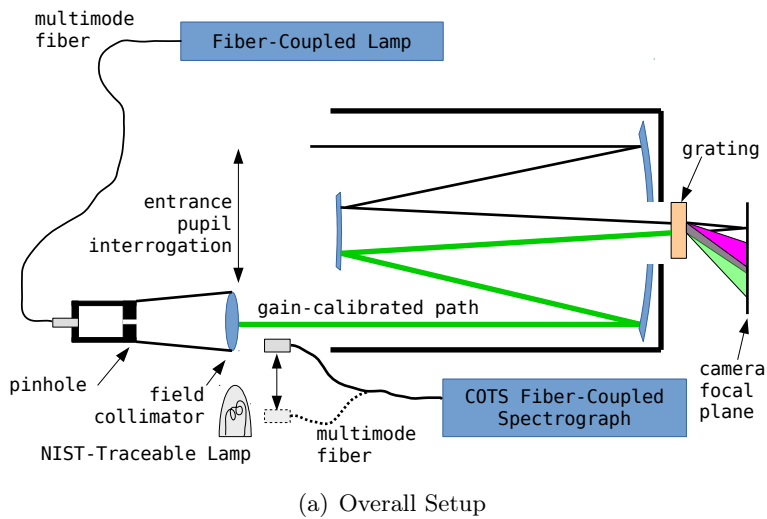


Figure 6-16: Bench Test Setup Overview. The purpose of the testing is to validate the assumptions used to compute the transformation from the measured amplitude along the spectral orders in grating images to the inferred amplitude of the spectrum of the input illumination. The setup in panel (a) consists of a spot projector made with a multimode fiber behind a pinhole placed at the focal plane of a collimator. The fiber is connected to a broad-band source. The shape of the spectrum at the output of the spot projector must be measured against a calibration standard in order to measure the shape of the combined spectral throughput of the telescope optics and detector quantum efficiency. This is accomplished with a laboratory spectrograph that measures the output of the spot projector and a standard radiometric calibration lamp. The radiometric standard lamp is too bright to be directly measured with the 16-inch telescope, necessitating indirect calibration. To establish the wavelength scale on the spectrograph focal plane, the spot projector is fed with a narrow-band light source to establish fiducial points for different wavelengths on the focal plane. This is shown in panel (b) where a free-space segment allows the insertion of a suitable interference filter into the path between the lamp and the projector. A free space segment is preferred over inserting a filter between the spot projector and the telescope because the free-space segment in the fiber does not introduce tilt into the grating images, thereby enabling a measurement that does not depend on accurate placement of the interference filter.

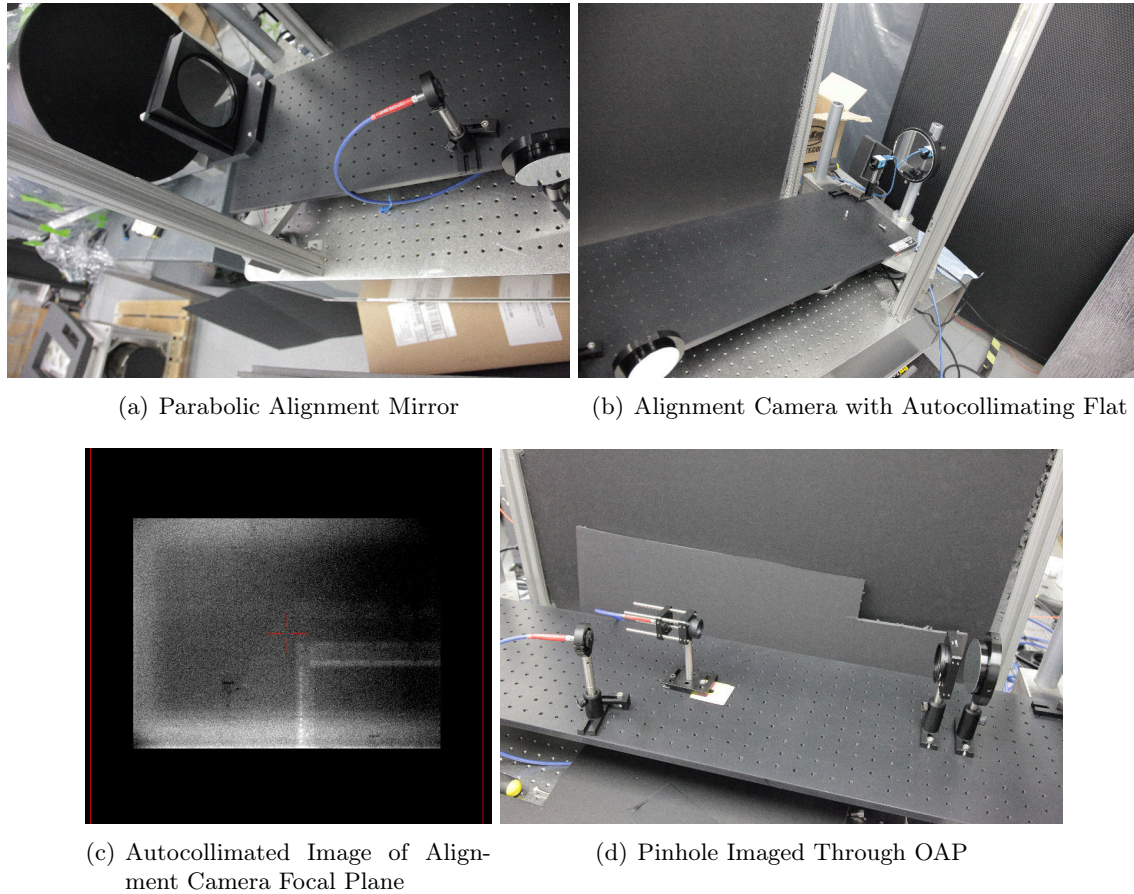


Figure 6-17: Photographs of First Attempted Spot Projector

6.6.2 Spot Projector

6.6.2.1 Breadboard Free-Space OAP Spot Projector

The first attempt to construct the spot projector consisted of using entirely breadboard optical components and free-space optics. A six inch glass substrate parabolic mirror (Figure 6-17(a)) was used to establish an optical axis parallel to that of the 16-inch telescope. An alignment camera imaged itself through the paraboloid via an autocollimating flat (Figure 6-17(b)) to establish infinity focus and alignment to the optical axis of the six inch mirror by imaging the edges of its own focal plane (Figure 6-17(c)).

Having established this optical axis and infinity focus in the placement of the alignment camera, the autocollimating flat was removed and a metal-substrate off-axis paraboloid (OAP) with a 60 cm focal length was placed on the optical bench. A 5 micron pinhole illuminated by an optical fiber was placed at the focus of the OAP (Figure 6-17(d)) giving a nominal $10 \mu\text{rad}$ image. The pinhole




Figure 6-18: Free-Space OAP Spot Projector Imaged Through 16-inch Telescope. The shape of the PSF in panel (a) is more diffuse than one would like in order to make laboratory tests of optical quality. In panel (b), there is evident banding radiating away from the $m = 1$ spectral order when the broadband source is imaged through the grating. This was seen in all bench test images of the collimated point source and is caused by the fact that a multimode fiber was used in the setup, meaning the spectral content of the signal coming out of the fiber-coupled pinhole is not uniform across the extent of the output fiber. The bands emanating from the spectral order correspond to the different modes (with different spectral content) coming out of the fiber.

was illuminated with an Ocean Optics HL-2000-HP-FHSA fiber-coupled white light source through an Ocean Optics QP400-2-VIS-NIR optical fiber. The position and angle of the pinhole and the OAP was adjusted until focus was achieved in the alignment camera. Once alignment and focus through the 6-inch paraboloid was achieved, the paraboloid was removed and the infinity-focused collimated spot from the OAP went into the 16-inch spectrograph.

An one-inch aperture iris had to be placed in front of the OAP to improve the optical quality of the spot by reducing the influence of the edges of the metal-substrate OAP. Figure 6-18(a) shows the spot from this first attempt at a projector imaged through the spectrograph OTA. The spot diameter is approximately 25 pixels across, which is significantly worse than the expected seeing.

Images taken through gratings show distinct wavelength-dependent banding (Figure 6-18(b)). Some sort of banding was observed in all subsequent attempts to build a spectrograph and may be a result of the wavelength-dependence in the spatial distribution of the power output from the fiber illuminating the pinhole from behind.





Figure 6-19: Samyang 500 mm Mirror Lens Spot Projector




(a) Focused with 550 nm Illumination (b) Same Focus with 750 nm Illumination

Figure 6-20: Chromatic Aberration in Samyang Mirror Lens

6.6.2.2 Commercial Mirror Lens Spot Projector

The spot size realized in the first iteration of the spot projector was unacceptably large to make meaningful tests with grating images. An attempt was made to improve on this optical quality by using a commercial photography lens. A 500 mm focal length Samyang mirror lens was available at the site, and the second attempt consisted of mating the fiber-illuminated pinhole to this lens (Figure 6-19). Although the best-achievable spot size with this lens was substantially better than with the free-space OAP, the lens suffered from substantial chromatic aberrations.

Figure 6-20 shows this chromatic aberration in action. A 550 nm interference filter was inserted into the free space gap in the pinhole illumination fiber and the mirror lens was focused. When the 550 nm filter was replaced with a 750 nm filter, optical quality was significantly degraded.



(a) Telescope With Illuminated Pinhole

(b) Sheetmetal Aperture Mask

Figure 6-21: Spot Projector Made with Orion Telescope

6.6.2.3 All-Reflecting Amateur Telescope Spot Projector

Figure 6-21 shows the final attempt to construct the spot projector. This consisted of mounting the fiber-illuminated pinhole to the focal plane of an all-reflective Orion Optics Starblaster 6 Newtonian telescope. Unfortunately, the large aperture diameter of this telescope caused the 16-inch telescope to saturate when imaging the collimated pinhole. Although it would have been possible to attenuate the signal in the pinhole by inserting a neutral color filter into the free-space path, an aperture mask was a simpler color-neutral method to attenuate the signal. A $1\frac{3}{8}$ inch diameter hole punched into sheetmetal and placed between the spot projector and the 16-inch telescope provided sufficient attenuation.

6.6.3 Mechanical Stability Validation

A key question stemming from the multiplexing operation of the spectrograph is the degree to which the cycling filter wheel will maintain repeatability of the placement of gratings and blocking filters. While a lateral displacement of either grating or filter is inconsequential to the recovery of the spectrum, a nonrepeatability in tilt of either the grating or the blocking filter will introduce a shift in the origin of the spectrum relative to the PSF center measured during a clear/clear frame.

Mechanical repeatability does not require a tight spot size to check, and the data from the first attempted spot projector is adequate for making this comparison. With the spot projector on and no mechanical or optical perturbations applied to the bench, the gratings and blocking filters were cycled repeatedly in the filter wheel. The resulting shift in the $m = 0$ centroid in the image is shown in Figure 6-22.




Figure 6-22: Repeatability of $m = 0$ Centroid Location Under Cycling of Filter Wheel

Perturbations to the centroid are most noticeable when the filter wheel cycles between different gratings. The maximum amount of perturbation between a clear and grating image appears to be about five pixels when the gratings are the fastest to cycle, and almost ten pixels if the gratings cycle slowly. Fortunately, the largest nonrepeatability is confined to the vertical coordinate while the worst-case horizontal shift is under five pixels. When taking data with spectra aligned with the horizontal, this will introduce a possible bias of no more than 15–20 nm into the spectral measurement. The mechanical performance of this instrument is acceptable if the target spectral resolution is 40 nm.

6.6.4 Absolute Photometric Calibration

With the spot projector constructed, an attempt was begun at a full radiometric calibration. The fundamental task necessary toward this objective is to make a calibrated measurement of the spectrum coming out of the spot projector. Toward that end, a Newport Model 63358 45 Watt Quartz-Tungsten-Halogen lamp with a NIST-traceable Newport Model OPS-Q250 power supply were purchased and set up in the optical test lab at the Firepond Facility (Figure 6-23), and an Ocean Optics USB4000-VIS-NIR-ES fiber-coupled spectrograph, capable of measuring from 340 – 1047 nm was borrowed to make the calibration.





Figure 6-23: Radiometric Calibration Source Set Up at Firepond

6.6.4.1 Measurement Through Achromat Lens

The first idea to measure the spectrum output by the spot projector was to simply place an optical fiber in front of the aperture as implied by Figure 6-16, but the area of the fiber was too small to capture enough energy from the spot projector with the 5 micron pinhole in place. This led to an attempt to measure the spectrum out of the spot projector by coupling the spot projector output into the measurement fiber with an $f/2$ achromat lens (Figure 6-24(b)).

Because the spot projector spectrum was measured with a lens to couple the light into the fiber, the calibration source spectrum also had to be measured with the same lens. Although measuring the calibration lamp with the lens enabled easy positioning of the image of the lamp onto the fiber input, the lens collimated too much light onto the fiber and the Ocean Optics spectrograph saturated at its minimum exposure of 4 msec when the calibration lamp was turned on to its full power in order to generate the calibrated spectrum.


At this point, an attempt was made to place an aperture stop onto the front of the lens. After some trial and error, a 0.041 inch diameter aperture mask was found to give an adequate SNR in the fiber-coupled spectrograph at a reasonable exposure. This aperture mask was placed on the lens



(a) Measurement With Bare Fiber

(b) Measurement With Achromat

Figure 6-24: Setup for Measurement of Spot Projector Spectrum



(a) Drilled Washer Aperture Mask

(b) Aperture Mask on Cage

Figure 6-25: Aperture Mask for Measuring Calibration Lamp

by drilling a hole into a washer and taping the washer to the front of the lens mount with aluminum tape.

The need to mask out everything but a small portion of that lens had the following implication for the radiometric calibration. The lens, though an achromat, still has chromatic aberration over its full aperture over the full 340 – 1047 nm band of the calibration measurement. This chromatic aberration will cause wavelength-dependent defocusing at the fiber input, on top of any wavelength-dependent coupling efficiency variation at the fiber as a function of incident angle over the full $f/2$ light cone. The use of the aperture mask introduces a complication because any chromatic aberration that is present over the full aperture of the lens for the spot projector measurement will be nulled out by selecting only the inner 0.041 inch diameter section of the lens for the calibration measurement.

This problem is caused by the dynamic range difference between the two sources meant to be measured: the fiber-coupled light source behind a 5 μm pinhole and the 45 W calibration lamp.

The small aperture required to measure the lamp is insufficient to measure the output of the spot projector. The dynamic range of the fiber-coupled spectrograph is insufficient to bridge the gap. It would be theoretically possible to make the measurement by inserting calibrated attenuating elements into the path between the calibration lamp and the spectrograph, but in addition to the added complexity of such a measurement, the verification of the spectral transmittance of such elements itself requires a high dynamic range measurement.

6.6.4.2 Measurement of Spot Projector Transmittance and Pinhole Spectrum

Given the difficulty in making a direct calibrated measurement of the output of the spot projector, a second approach was tried: combining a calibrated measurement of the spectrum of the output of the pinhole at the focal plane with a measurement of the spectral transmittance of the optics of the spot projector. In this approach, the only measurement required to be referenced to the calibration lamp is that of the spectrum of the output of the pinhole.

The measurement of the pinhole spectrum is straightforward: the spectrum of the calibration lamp is measured with a bare fiber to calibrate and then that same measurement fiber is placed into the pinhole assembly opposite the pinhole from the illuminating fiber (Figure 6-26(a)). The distance between the pinhole and the measurement fiber is dialed in to match the $f/5$ focal ratio of the Orion telescope. Both of these measurements are easy to make and are within the dynamic range of the fiber-coupled spectrograph.

The measurement of the spectral transmittance of the Orion telescope was attempted as follows. A flat mirror was placed at the focal plane instead of the pinhole. An illuminating fiber was coupled to the achromat lens with the cage assembly and a second cage assembly was connected to the fiber-coupled spectrograph. The two cages were pointed into the aperture of the Orion telescope and aligned such that the measurement fiber caught the retroreflection of the input illumination through the focal plane mirror. This is shown in Figure 6-26(b). This measures two trips through the spot projector and one reflection off of the flat at the focal plane.

The reflectance of the flat placed at the focal plane had to be measured with the same achromat-coupled fiber pair. As such, a normal incidence measurement was not possible, and the measurement was made at an incidence of about 45 degrees as shown in Figure 6-26(c). Finally, this measurement must be referenced against the transmittance through the two achromat-coupled fibers (Figure 6-26(d)).

These two latter steps is where the calibration attempt failed. While the measurement of the




Figure 6-26: Setup for Piecewise Measurement of Spot Projector Spectrum

round trip through the spot projector was within the dynamic range of the fiber-coupled spectrograph, the strength of the received signal in the face-on spectrum of one achromat aimed at the other was beyond the saturation point of the spectrograph. Furthermore, the face-on spectral reflectance off of the retroreflective flat as used in the round-trip measurement may not be identical with the spectral reflectance of the flat by itself at 45° incidence angle, since it was not known if the flat had a protective coating that would introduce angle-dependence into its spectral reflectivity.

6.6.5 Relative Photometric Calibrations

In lieu of an absolute radiometric calibration, several relative measurements were possible that did not require knowledge of the shape of the input illumination spectrum. The two that were of particular relevance to the spectrograph data processing are the verification of blocking filter transmittance and the grating efficiency.




Figure 6-27: Blocking Filter Bench Calibration

6.6.5.1 Blocking Filter Transmittance Validation

The first relative measurement that can be performed without knowing the exact shape of the input spectrum is a measurement of the roll-on and roll-off of the transmittance of the blocking filters. This measurement is made by multiplexing between a clear slot and the blocking filter while imaging through the same grating.

The result of the blocking filter calibration made with the Orion telescope spot projector is shown in Figure 6-27. Although measurement of the short wavelength rolloff of the BG40 filter is not possible given the low detector quantum efficiency in that region, the quality of the rest of the measurement is sufficient to measure the difference between the filter datasheet and the actual transmittance.

The measurement of the short-wavelength rollon of the OG515 made here is what determined that the OG515 does not provide enough suppression of wavelengths below its nominal cutoff. As such, the OG570 filter was used for all on-sky measurements given that its suppression of wavelengths below 500 nm was confirmed.

6.6.5.2 Grating Efficiency Validation

The second relative measurement is a validation of grating efficiency. The ratio of spectral grating efficiencies is invariant to the shape of the input spectrum. The ratio of signal strengths at common wavelengths may be compared against predictions from Equation 6.13, properly scaled by the ratio of the dispersions of the two gratings. For the red end of the silicon passband, the orange glass blocking filter is necessary to eliminate contamination by $m = 2$ spectral order of short wavelengths. While




Figure 6-28: Ratio of Responses of 606 and 906 Grating vs. Prediction

for short wavelengths, no blocking filters are necessary, the blue glass filter was used to maintain focus at the same position.

A sample of this ratio of responses is shown in Figure 6-28, with a nominal AR coating response added to the 606 grating, as ordered from Richardson. Both gratings are made of the same 3 mm BK7 and have an otherwise common transmittance. The measured ratio does not match the nominal prediction. While a potential explanation would be that the order was processed incorrectly and the AR coating was applied to the 906 grating instead of the 606 grating or not applied at all, there is no plausible combination of AR coating mismatch that recreates the observed data.

A more plausible explanation is that the gratings are not orthogonal to the optical path. Figure 6-29 shows the effect on the optical path difference function for a tilted grating. The action of a tilt by an angle θ_t transforms the total phase shift in Equation 6.8 to:

$$T(x) = e^{j \tan(\theta_g)(n_g(\lambda)(\tan(\theta_g + \theta_t) - \tan \theta_t) - n_a(\lambda) \tan(\theta_g + \theta_t)) (x \bmod (\Delta \cos \theta_t)) \frac{2\pi}{\lambda}} \quad (6.14)$$

which in turn transforms the grating response for mode m from Equation 6.13 to:

$$I(\lambda) \propto \left[\text{sinc} \left(\pi \frac{\Delta}{\lambda} \cos \theta_t \left(\tan \left(\sin^{-1} \left(m \frac{\lambda}{\Delta} \right) - \theta_t \right) - (n_g - n_a) [\tan(\theta_g + \theta_t) - \tan \theta_t] \right) \right) \right]^2 \quad (6.15)$$

The influence of a tilt perturbation to the ratio of grating responses can be significant. Figure 6-30 shows that a tilt of just 0.02 degrees, or about 17 microns over a 50 mm grating, can change

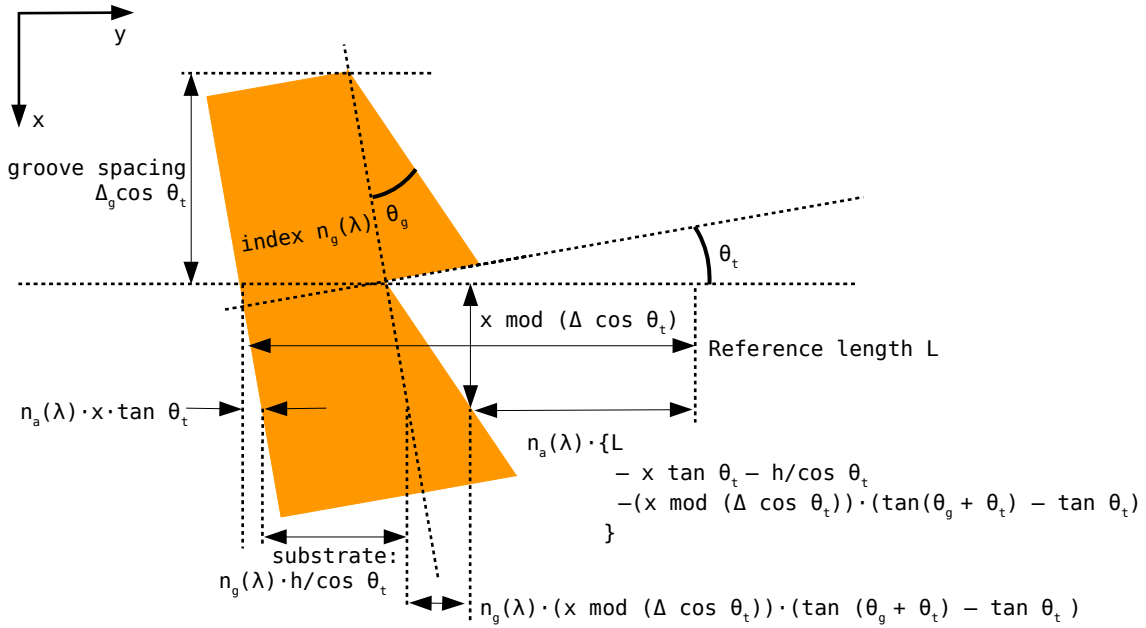


Figure 6-29: Optical Path Difference Calculation for Titled Grating

the response ratio by several percent relative to an orthogonal assumption.

It is important to understand the degree to which the tilts of both gratings may be estimated from the grating images. The most straightforward way to do this is by brute-force computation of all reasonable tilts for both gratings and the selection of the best-fit response ratio that matches the data. Furthermore, since we are talking about a few tens of microns worth of shift over the 50 mm span of the filter, which corresponds to perhaps no more than a tenth of a millimeter of shift across the diameter of the grating carousel, this measurement cannot be confined to the laboratory bench and must be performed on-sky as the telescope changes orientation. The on-sky stability of the grating tilt in the installed system will be discussed in Section 6.7.4.

As a preview of on-sky measurements, Figure 6-31 shows the predicted zero-tilt grating response ratio, the measured grating response ratio, and the predicted best-fit tilt response ratio as measured on a single bright star, and the log squared error of the reconstruction for all possible grating tilts. The best-fit tilt angles give a cosmetically good reconstruction of measured data. The shape of the error function implies that the tilt of the 606 grating is more tightly constrained by the response ratio than the tilt of the 906 grating.

Tilt reconstruction requires measurements with both gratings, which necessarily requires move-




Figure 6-30: Influence of Grating Tilt on Ratio of Grating Efficiencies




Figure 6-31: Estimation of Grating Tilt Using Response Ratio

ment of the filter wheel between measurements. The fact that the grating tilt may be unstable with cycling of the filter wheel means that a laboratory bench absolute photometric calibration is not meaningful. Thus, although the attempt to make the absolute photometric calibration did not succeed due to a mismatch between the dynamic ranges of available test equipment, the consequences of that failure are obviated by the fact that the measurement would not have been useful in light of the mechanical nonrepeatability present in the spectrograph instrument itself.




Figure 6-32: Installation of the Spectrograph at Firepond, August 16, 2018

6.7 Installation, Checkouts, and On-Sky Calibration

The work of assembling the spectrograph hardware and modifying the Firepond telescope data acquisition software to enable remote control of the dual filter wheel occurred during June and July 2018 and the bench testing took place over the first half of August 2018. At that point, the instrument was ready for deployment and on August 16, 2018, it was lifted into the dome and installed underneath the Firepond 48" telescope. Figure 6-32 shows highlights from the installation day.

The Firepond facility was tasked with higher-priority operations until the end of September and first light for the spectrograph was on the night of September 28, 2018. The primary tasks for first light were focus and alignment on a terrestrial reference and on-sky calibration. The purpose of the on-sky calibration was to measure the kinematic stability of the spectrograph optical tube assembly (OTA) with respect to the orientation of the 48" telescope and the stability of the installed geometry of the gratings in the spectrograph with cycling of the filters and changing orientation of the OTA. This assessment of the geometric stability entailed the measurement of the stability of the wavelength scale (that is, distance between camera and grating) and of grating tilt with both cycling of the filter wheel and the changing orientation of the OTA.

These initial characterizations were necessary for the basic operation of the spectrograph as they informed the development of algorithms necessary to extract spectra from the images for targets whose $m = 0$ PSFs were too faint to detect the origin points of their spectra.

During the first night of checkouts, it was discovered that the primary mirror mounting of the

16-inch telescope was insufficient to maintain focus above 60 degrees elevation. The amount of travel in the instrument's focus slide was not sufficient to compensate for the amount of shift in the primary mirror. This is a known problem with 16-inch Meade telescopes. Although this reduced the potential diversity of atmospheric conditions over which measurements on satellites could be made, the lost observing opportunities occur where little DCR bias is expected in the first place.

Once the geometric alignment and stability characterizations were complete, further measurements were made to validate the models used in the analysis in Chapters 4 and 5, which in turn inform the algorithms necessary for spectroscopic compensation of DCR bias. These validation measurements entailed

1. Verification of the spectral transmittance of the blocking filters. The stability of optical filters over time and over changing temperature cannot be assumed. Although there was no particular reason to expect the filters purchased for this program to degrade, the unwritten tribal knowledge passed down from observer to observer at Lincoln Laboratory dictated that this was a good measurement to repeat.
2. Observation of stars with known spectra to validate the assumptions made on the spectral transmittance and quantum efficiency of the atmosphere and the spectrograph optical elements that were used to design the spectrograph and perform the analyses described in Chapters 4 and 5. This measurement is partially a repeat of the full radiometric calibration attempted during laboratory testing, but also incorporates validation of the MODTRAN atmospheric transmittance models used for all analyses in this thesis and also used in image reductions to compute the DCR bias of stars.
3. Observation of bright stars, not necessarily with known spectra, to validate the atmospheric refraction model used for the analyses in this thesis.

Table 6.2 summarizes the calibration and satellite measurements undertaken between September 2018 and April 2019. The calibration procedures described in this chapter evolved over the Fall of 2018 as informed by experience with the system. In particular, the geometric calibration procedures on bright stars were not fully fleshed out or tested before the fifth operational night in November 2018. For this reason, the satellite collections on the third and fourth nights (2018:311-313) did not include calibration measurements.

Table 6.2: Summary of Spectrograph Operational Nights

Date	Collection Objectives	Weather	Result/Comment
2018:272	Mount bias model, stars and satellites	Partly cloudy, then clear, humid	First light, focus check (limited to below 60 deg)
2018:279	Stars and satellites	Partly cloudy all night	
2018:311/2	Mount bias model, satellites only	Clouds early, then clear	
2018:312/3	Satellites only	Clouds early, then clear	Focus quality degraded by end of night
2018:341/2	Stars and satellites	Light clouds early, then clear	Refocused spectrograph to middle of focuser travel range
2018:346/7	Stars and 1 satellite	Light haze on satellite track	Confirmed good focus
2019:094	Stars and satellites	Windy and cloudy early, then clear and calm	Replaced malfunctioning grating/filter wheel. Measured a blue glint

6.7.1 Kinematic Stability Relative to 48" Telescope

6.7.1.1 The Necessity of Mount Bias Models

Azimuth and elevation or hour angle and declination are mathematical abstractions that are measured by physical angle encoders in real telescopes. From these measurements, the pointing of the telescope's OTA and the orientation of the telescope's focal plane within well-defined terrestrial and inertial coordinates are inferred. Rigid-body assumptions that relate the measured shaft angle of the telescope to the orientation of the focal plane are accurate to a few tenths of a degree for even the most well-designed telescopes in the presence of nonorthogonalities between mechanical axes and deflections under load. Modeling of deviations from the idealized assumptions is necessary to achieve pointing accuracy at the arcsecond level. Standard physical effects that are included in telescope pointing models account for nonorthogonality between the mechanical axes, misalignment of the OTA with respect to the telescope mechanism, OTA tube flexure with changing orientation, and tilt of the telescope bearing [141].

The Firepond Telescope Control System (TCS) includes an integrated model with terms for these physical effects that allows the 48" telescope to achieve all-sky pointing accuracy below 1.5 arcsecond RMS on a good night. The procedure to compute the coefficients for this pointing model involves pointing the telescope to a large number of star fields sampling the full range of motion of both axes (81 points are used operationally) and computing the offsets between the observed astrometric solution (as expressed in idealized telescope axes coordinate) and the measured telescope encoder values. The model is then computed from these offsets.

While the Firepond mount is made of welded and bolted steel, the mounting between the spectrograph OTA and the Firepond telescope mount is constructed of aluminum and is clamped onto the 48" telescope truss. Although the spectrograph was co-aligned with the 48" telescope at




Figure 6-33: Post-Fit Pointing Residuals of Mount Bias Models Computed for the 48" Telescope and 16" Spectrograph Telescope from Data Collected on the Night of Nov 7-8, 2018. The worst-case corrected pointing error for the 48" telescope is under a few arcseconds, but the worst-case corrected error for the 16" telescope is a few tens of arcseconds. There is hysteresis in the kinematic mounting of the telescope that is most evident in the separation of traces in the vertical component of the error. One branch of the curve corresponds to moving downward from high elevation while the other manifests when moving up from low elevation. The particular shape of the curve is a function of time-history of the mount motions used to collect the dataset from which the mount bias model coefficients were computed.

installation, the kinematic mounting between the two telescopes is not rigid, and a separate set of mount model coefficients is necessary for the spectrograph. The inability of the spectrograph to maintain focus above 60 degrees elevation limits the range of axes motion over which the model needs to be applied.

A representative measurement of the post-fit mount bias modeling errors for both the 48" telescope and the 16" spectrograph is shown in Figure 6-33. While the 48" telescope pointing accuracy is sufficient to blindly place a slowly-moving target to within a few pixels of its desired location on the 48" telescope focal plane, the performance of the 16" telescope is significantly worse, as manifested in the obvious hysteresis in the vertical component of the error.

6.7.1.2 Empirical Bias Model for the Spectrograph

The implementation of the mount bias model software within the Firepond TCS is memoryless, and it is not possible to include hysteresis effects in the existing software that tags recorded images with their nominal pointing. As such, it was necessary to post-process imagery from the 16" telescope to compute the pixel locations of satellite targets in spectrograph frames where the $m = 0$ PSF may not be detectable. There are two possible approaches to this: a physical modeling approach that attempts to model the hysteresis and compute the pointing for each image from the time-history

of prior mount motions, and an empirical approach that interpolates the alignment between focal plane and inertial coordinates from nearby fiducial clear/clear frames where astrometric solutions are possible.

The efficacy of the physical modeling approach is questionable given that one would expect the shape of the hysteresis plots to depend on the time-history of mount motions. The plot in Figure 6-33(b) is a product of step-and-stare motions up and then down over a finite set of points, whereas the motion history of tracking a MEO satellite going up, going down, and changing tracks under different circumstances may result in a different error signal. Thus, in order to ensure that the model is valid over all possible satellite tracks, it would be necessary to make fiducial astrometric measurements during the satellite tracks in order to build the model.

The empirical approach entails the same data collection without attempting to construct the physical pointing model. When data is collected on satellite tracks, the spectrograph collection frames are interleaved with 30-second windows for astrometric calibration without gratings or blocking filters in the optical path. Figure 6-34 shows the time-history of the pointing difference for the 16" telescope between the astrometric solutions during satellite tracks and the nominal pointing computed from encoder angles supplemented with a mount bias model. Although the time-history of the difference is discontinuous across mount slews between different targets, during periods of smooth tracking of GPS and GLONASS satellites, the time-history of the deviation between computed and measured pointing is also smooth and continuous.

It is therefore sufficient to use the empirical approach to infer the location of target spectra within the focal plane. The telescope pointing for spectroscopic collection frames is computed by applying the interpolated difference between adjacent astrometry frames. The location of the $m = 0$ PSF in the focal plane is then inferred from the IGS truth ephemeris and this empirically-corrected 16-inch OTA pointing. For nearly all satellite data collected during the seven nights of system operation, this technique enabled the automatic extraction of satellite spectra by locating the $m = 0$ PSF of satellites to within three pixels of their true location as gauged by instances where the $m = 0$ PSF was detectable on bright targets.

6.7.2 Wavelength Scale

Although the nominal distance between the grating and the focal plane was approximately 10 cm, the process of installing the filter wheel into the optical backend and aligning and focusing the system placed the focal plane at an unknown distance from the back face of the gratings. Furthermore, the

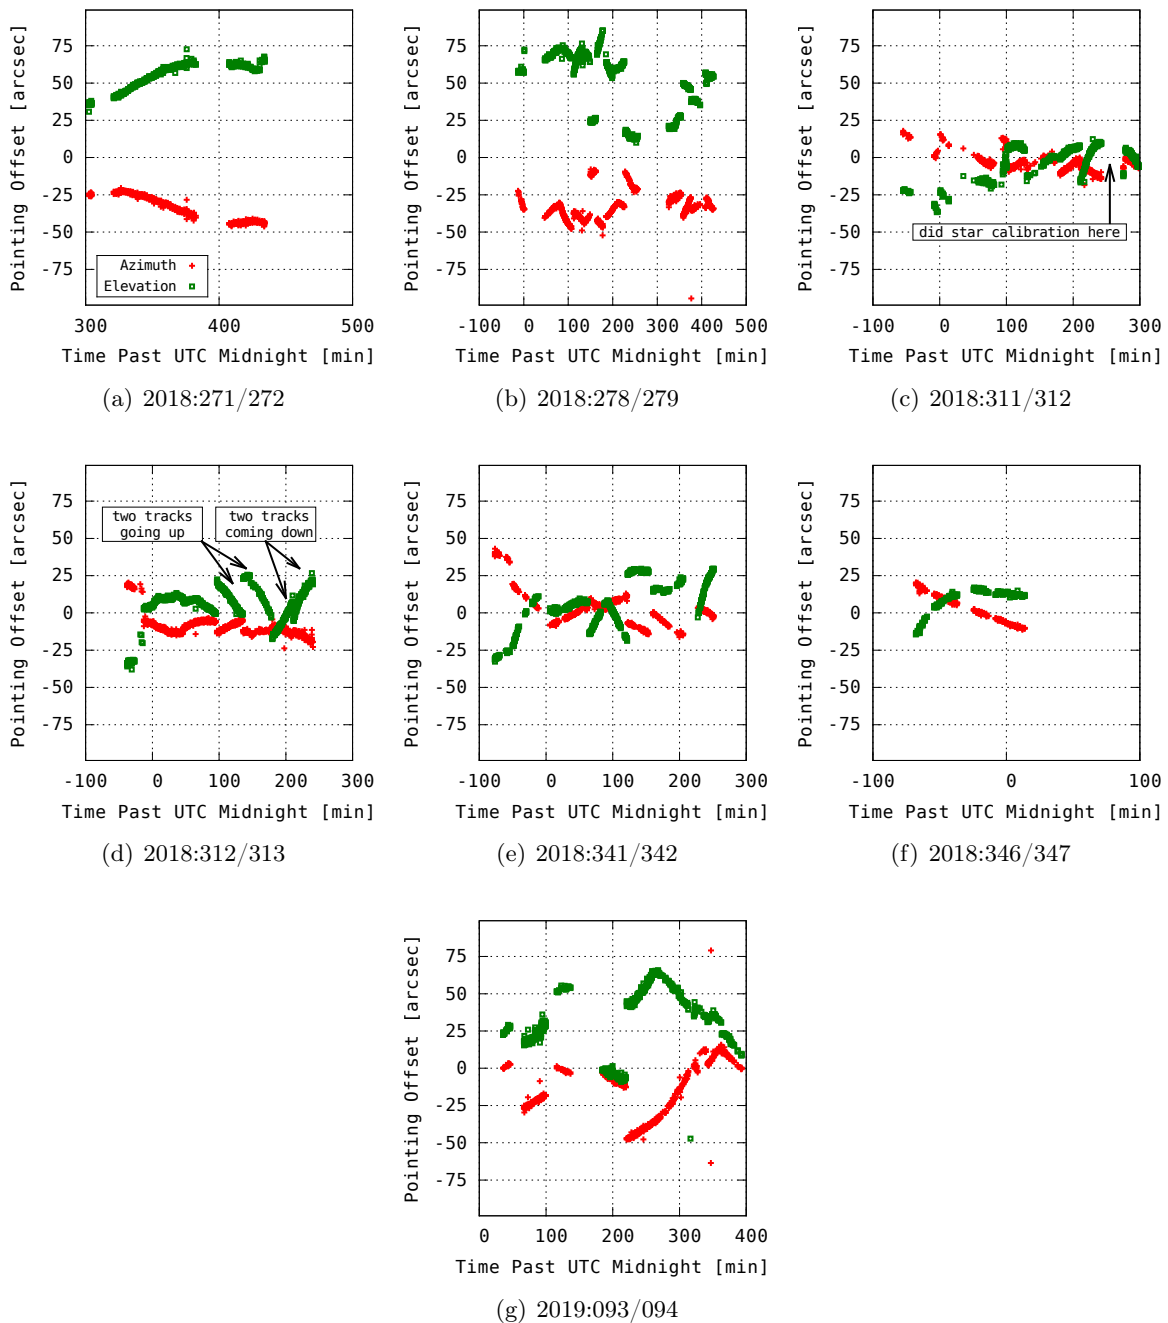


Figure 6-34: Time-History of Difference Between 16" Telescope Pointing Measured from Bare-CCD Astrometry Frames and Pointing Computed from Mount Bias Model. The mount bias model was collected on the night of 2018:311/312 and used to seed the astrometric solution for all images collected over the seven operational nights. Bare-CCD astrometry frames interleaved between grating images of stars and satellites were then used for image registration in order to locate the origin of the desired target's spectrum in the grating images. The smoothness of the offsets of the telescope pointing as measured from the astrometry against the pointing predicted from the mount bias model indicates that interpolation of focal plane orientation between successive astrometric measurements is a valid operation that will not result in failure to locate the target spectrum in grating images.

slop in the Meade telescope’s primary mirror requires moving the focal plane via the focus stage within a range of motion of up to ± 12.5 mm from center to keep up with focus. It is therefore necessary to measure the actual wavelength scale of the spectra on the focal plane.

The most straightforward way to do this is to use atmospheric absorption bands at known locations to establish the wavelength scaling. Figure 6-35 shows the predicted atmospheric transmittance for extremes of temperature and humidity encountered during the spectroscopic measurement campaign, along with measured spectra of sampled from those two nights. The broad water vapor line at 940 nm is a tempting target for calibration, but on particularly dry nights and for stars with relatively little energy around 940 nm, it is not the strongest feature in the measured spectrum. The molecular oxygen line at 765 nm, on the other hand, is always present.

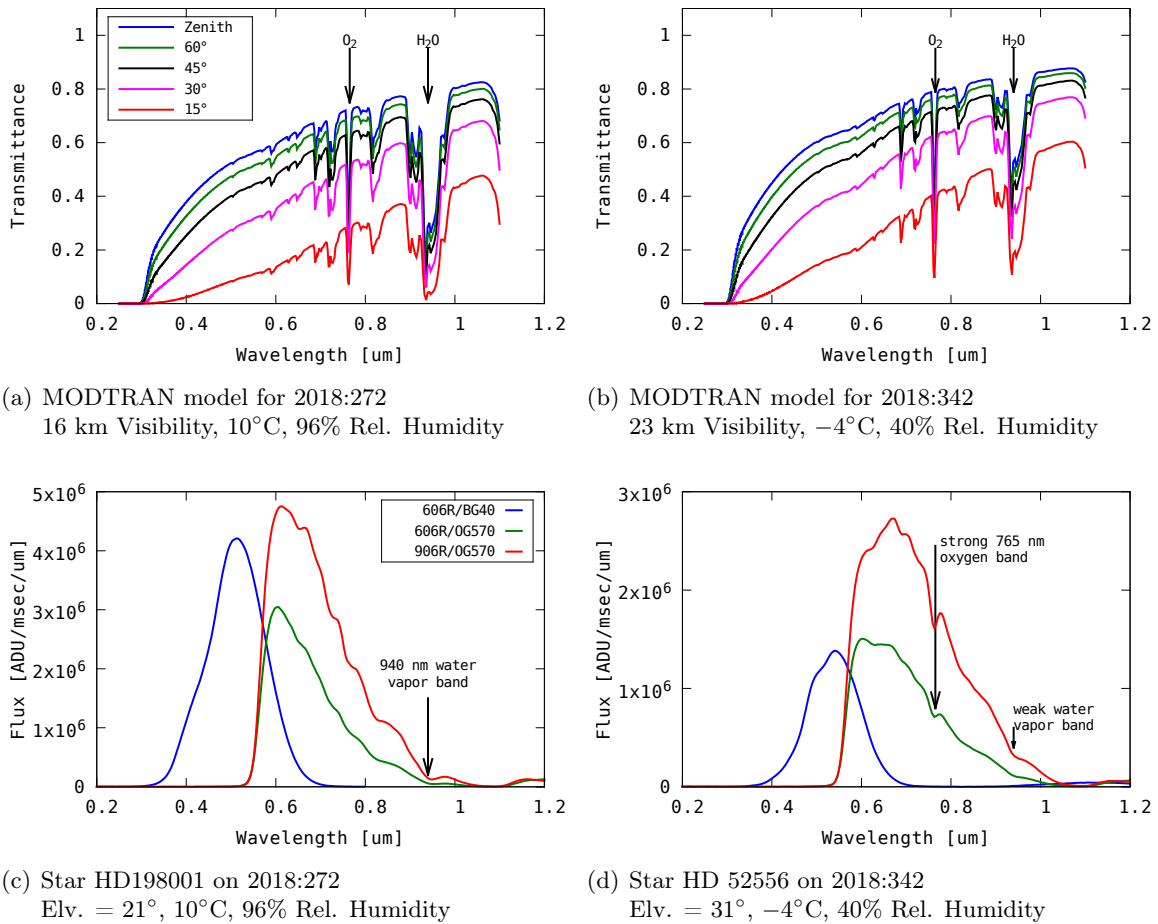


Figure 6-35: MODTRAN Atmospheric Spectral Transmittances and Sample Stellar Spectra for a Warm Wet Night and a Cold Dry Night. The water vapor band at 940 nm can, in practice, be washed out on dry nights. The molecular oxygen band at 765 nm is present at all times with constant strength and can be used as a fallback.

The process for automatically performing the calibration is to manually dial in an assumed wavelength scale based on visual inspection of the measured spectral components and then to find the maximum of the spectrum convolved with a double-difference kernel. The width of the double-difference kernel is computed from the width of a Gaussian fit directly to the spectrum in the vicinity of the band. It is necessary to perform this fit for both gratings, and for that reason it is necessary to include observations made with the blue-optimized 606 grating with the OG570 blocking filter in the data collection sequence. The stability of the wavelength scaling for the two gratings measured against both absorption bands is shown in Figure 6-36.




Figure 6-36: Stability of Wavelength Scale Over Operational Nights with Star Measurements. The large discontinuity between 2018:279 and 2018:342 corresponds to a refocusing of the telescope to place neutral focus closer to the center of the range of travel for the focus slide and the slight discontinuity between 2018:347 and 2019:094 corresponds to disassembly of the optical back end to swap out a broken filter wheel, without deliberately modifying focus.

The geometric scaling calibration can only be performed on stars brighter than about 7th magnitude. Satellites are insufficiently bright to perform this calibration in a timely fashion and the mount jitter and star clutter encountered during the track further confounds the measurement. Thus, the mean-value wavelength scale calibration for a full night measured against stars is used to process satellite data as well. For the nights of 2018:311-312 and 2018:312-313, stars were not observed, and the geometric scaling measured on 2018:272 and 2018:279 was used instead. This is justified based on the fact that for the five of seven operational nights when calibration stars were observed, geometric stability over the course of the operating period is within 2%.

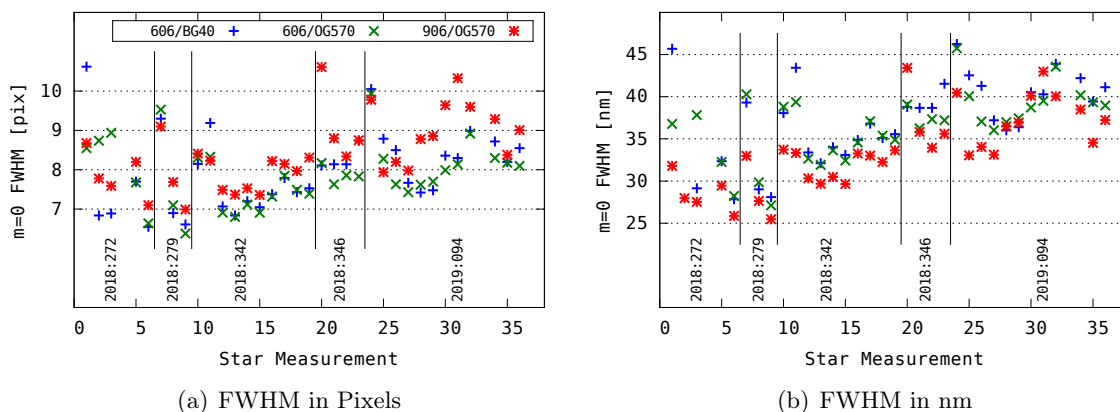


Figure 6-37: Achieved Optical Quality and Spectral Resolution

6.7.3 Seeing and Achieved Resolution

The horizontal width of the $m = 0$ PSF indicates the achievable spectral resolution of the instrument as limited by seeing. Figure 6-37 shows the measurement of the FWHM of the $m = 0$ PSF using the bright stars observed during five out of the seven operational nights. Overall the seeing disk had a FWHM diameter of about 8 pixels. One would theoretically expect seeing to be worse at short wavelengths, but there is no appreciable difference between measurements with the OG570 filter, where $m = 0$ would mostly consist of wavelengths redward of 800 nm, and measurements made with the BG40 filter, where $m = 0$ would consist mostly of wavelengths between 500 nm and 550 nm.

The achieved angular resolution translates into a wavelength resolution of just under 40 nm. The original design called for a spectral resolution of 40 nm, and the achieved on-sky performance just meets this goal. It is not possible to improve upon this resolution with this telescope. The way to improve resolution in the presence of fixed seeing is to increase the dispersion on the focal plane.

Dispersion can be increased with a higher dispersion grating or it can be done by moving the grating further away from the camera. Higher dispersion gratings that would still make the spectrum fit onto the 13 mm focal plane of the Andor cameras available at the site are not stocked items, rendering this impractical. Increasing the distance between the grating and the focal plane is also not possible as the grating is already as far forward as it can mechanically go. Moving the camera back is also not an option as it would degrade overall optical quality to refocus the telescope that far away from its intended focal length. In either case, increasing the dispersion would also lower the per-wavelength-bin SNR as each wavelength bin would take up a commensurately larger amount of




Figure 6-38: On-Sky Measurements of Grating Tilt

sky background noise.

The only reasonable alternative for slitless spectroscopy is to use a shorter focal length OTA, while keeping the gratings roughly the same distance away from the focal plane. This would tighten up the seeing-limited spot size but increase the background noise commensurately, making this more of a break-even proposal. Unfortunately, a shorter focal length telescope with the same 16-inch aperture and bracketry already built up was not available for this experiment.

6.7.4 Grating Tilt Calibration

The grating tilt as inferred from the ratio of responses of the 906 and the 606 grating measured with the OG570 blocking filter for the five nights when calibration stars were observed is plotted in Figure 6-38. The 606 grating's tilt estimate is more stable than the 906 grating's tilt estimate, as one would expect from the shape of the error function in Figure 6-31.

Allowing for outliers in the measurements caused by observing spectra in a cluttered environment or during periods of photometric instability caused by clouds blowing through, the stability for all days except 2019:094 seems good and it seems reasonable to use the median measured grating tilt over the whole night to process individual satellite spectra during that night. Between 2018:346 and 2018:094, the filter wheel malfunctioned and was replaced with a spare. It is possible that the apparent instability on 2019:094 is the result of the replacement having sloppier mechanical tolerances than the original unit, but it is also possible that the jitter at the beginning of the night

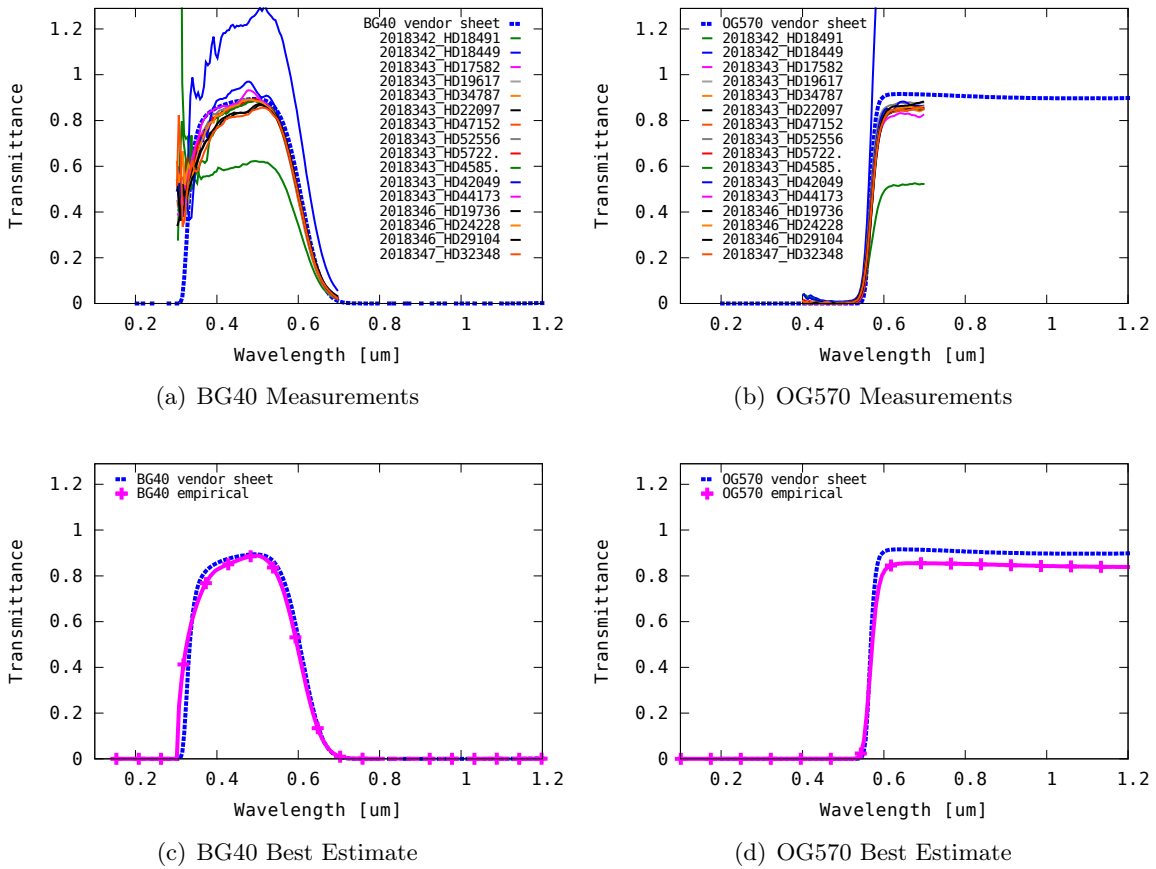


Figure 6-39: On-Sky Measurements of Blocking Filter Transmittance

is just an unlucky set of bad measurements given that the tilt estimates at the start and end of the night agree with both each other and with the measurements from the other nights.

These calibration measurements were made on bright stars (about 6 – 7 mag). An attempt was made to measure grating tilt from satellite measurements, but at 10th magnitude or fainter, the satellite spectra are insufficiently bright.

6.7.5 Photometry Validation

6.7.5.1 Filter Transmittance Check

During the last two operational nights of 2018, the ambient temperature was -4°C , and the blocking filter spectral transmittance was spot-checked using bright stars. The comparison of the resulting transmittance estimate in Figure 6-39 with the laboratory measurements at room temperature shows a slight change in the peak transmittance. The difference is most likely attributable to dust and/or condensation, as a light coating of grime was discovered on the blocking filters when


the filter wheel was swapped out for the spare in February 2019. The fact that the shape of the transmittance cut-off and cut-on agrees cosmetically with the laboratory measurement implies that whatever surface contamination there may have been was spectrally flat and thus inconsequential to the data processing.

6.7.5.2 Total Spectral Transmittance Validation


The estimation of DCR requires only knowledge of the detected photoelectron spectrum on the ground. In principle, this means that the only quantity that needs to be grounded in absolute calibration is the grating efficiency and blocking filter transmittance. Knowledge of the spectral transmittance of the atmosphere and the optical system, and the spectral quantum efficiency of the detector are not relevant as they are implicitly part of the measurement. If spectra of fiducial stars and the target were measured simultaneously, this is all that would be necessary. However since stars' spectra are not simultaneously measured by this instrument, their DCR bias with respect to the solar analog zero point, and the solar analog zero point against which the DCR of the target is tabulated, must be computed using a model of atmospheric and telescope optics transmittance and the detector spectral quantum efficiency. The purpose of this calibration is to observe stars with known exoatmospheric spectra and to use the measured spectra of those stars in the spectrograph to validate that overall spectral transmittance model.

The spectral typing and parallax for the Tycho-2 catalog published by Pickles in 2010 [142] was used as the source of reference spectra for the bright stars necessary for this measurement. This catalog is not a direct exoatmospheric measurement of stellar spectra or interstellar reddening for these stars but rather uses multicolor photometry to estimate the best-fit Pickles Atlas spectrum and unreddened free-space loss for each star. This model was implemented in OrbitTool and the observed bright stars' spectral type was looked up in the VizieR system [143].

The absolute calibration of some of these stars is better than others. Figure 6-40 shows two examples of stars with an underestimate and an overestimate of total detectable photoelectron flux. The model spectra need to be normalized to the measurement, and the region around 700 nm was chosen as the normalization point. Once normalized, the shape of the reference model spectra mostly matches the shape of the measured reconstructed spectrum, but as Figure 6-41 shows, sometimes the spectral type predicted by the catalog is incorrect. Part of this comes from the fact that the calibration must be made against bright stars. Bright stars must be used in order to be able to take the data in a timely fashion so as to still have time to measure satellites. This means most of the




(a) Star HD198001, 2018:272: Model Too Bright




(b) Star HD183913, 2018:272: Model Too Faint

Figure 6-40: Absolute Photometry Mismatch



(a) Star HD52556, 2018:343: Good Match




(b) Star HD42049, 2018:343: Model Too Red

Figure 6-41: Spectral Type Mismatch


stars measured brush up against the bright cutoff of the spectral type catalog where the errors are somewhat higher.

6.7.5.3 Results of On-Sky Photometric Validation


Despite the fact that the Pickles Tycho2 spectral type catalog does not appear to be suitable for a rigorous photometric calibration, for most stars the spectrum appears to match, and it is possible to make inferences about the performance of the system and the data processing necessary to extract meaningful DCR corrections from it. As Figure 6-42 shows, whenever the shape of the spectrum matches, a better match is obtained by assuming that the 16-inch telescope includes two anti-reflective coatings in its optics rather than only one. Plotting the percentage point difference between the observed spectra and the modeled spectra for all stars (including outliers), one does




(a) HD183913, One AR Coat



(b) HD183913, Two AR Coats



(c) All Reconstruction Errors, One AR Coat




(d) All Reconstruction Errors, Two AR Coats

Figure 6-42: One vs. Two AR Coatings

indeed see that the match appears better if two AR coatings are included in the model rather than one.

Furthermore, outside of a few bad stars whose measurements were likely perturbed by clutter or sky variability during the time of measurement, the model used to compute the reference spectra agrees with measurement to under 5 percentage points in the region away from where the red and blue ends of the spectrum are stitched together. This validates all of the spectral transmittance assumptions made in the analysis presented in this document and implemented in the software used to process the spectrograph data.



(a) Blue Portion Observed with 606/BG40

(b) Red Portion Observed with 906/OG570

Figure 6-43: Grating Images of HD198001 Observed at 21° Elevation During Checkouts

6.7.6 Refraction Model Validation

The plan to observe GPS and GLONASS satellites with vertical gratings giving horizontal spectra was motivated by the fact that the majority of those satellites' motion relative to the stars was in the vertical direction and by the fact that the wavelength scale on the focal plane of a horizontal spectrum is constant and is not perturbed by DCR. This is illustrated in Figure 6-43 which shows an image of HD198001 observed with the spectrograph during the first night of operations in 2018:272. The image was taken at 21 degrees elevation, and in the blue portion of the spectrum, the presence of DCR is especially evident in the upward curl of the shortest wavelengths.

Because that curl is directly observable, horizontally-aligned spectra also present an opportunity to validate the atmospheric spectral refraction model used in the analysis and data processing for this thesis. Using images of bright stars, the location of the spectrum's vertical centroid versus wavelength is a direct measurement of DCR. By observing bright stars over different nights and at different elevation angles, the validity of the refraction model can be checked over the same environmental conditions through which satellites are observed and under which astrometric reductions and DCR corrections are computed. As there is no way to directly measure refraction given the mechanical nonrepeatability and imprecision of the 16-inch telescope, the measurement must be a relative one with respect to some reference wavelength.




Figure 6-44: Grating Efficiency of Spectral Orders Suitable for Horizon Reference

6.7.6.1 Horizon Reference

A horizontal baseline must be established within the image in order to make the measurement. The pixel grid of the camera is not perfectly aligned to the horizon and the gratings are not perfectly aligned to the pixel grid. In general there are rarely enough bright stars in the field to compute astrometric solutions for grating images made with short exposures. Regions of common wavelengths (subject to the same wavelength-dependent refraction) among the spectral orders must be used to establish the local horizon.

Figure 6-44 shows the combined effect of blocking filters, quantum efficiency, and grating efficiency for the $m = -1$, $m = +1$, and $m = +2$ spectral orders for all three combinations of gratings and blocking filters used to observe bright stars. In general the $m = -1$ signal is higher than the $m = +2$ signal, and its symmetry with the $m = +1$ signal and compactness² relative to the more spread-out higher orders makes it the most attractive to use as the horizon reference.

The region from 500–600 nm has the highest signal for the 606/BG40 combination, 700–800 nm for 606/OG570, and 800–900 nm for the 906/OG570 data. The shape of the $m = -1$ responses in those regions are also similar to the shape of the $m = +1$ region, meaning the centroid averaged over those regions would not be expected to have a bias caused by one end of the region having a higher signal strength than the other since both are illuminated by the same star spectrum. That said, with the 606 grating, the peak signal strength in the $m = -1$ order is between 2–3% relative to the $m = 1$ signal and for the 906 grating the peak strength is under 1% relative to $m = 1$. This means the measurement must be made on bright stars, in practice around 6th magnitude.

²Compactness is desirable to avoid clutter from other nearby sources.

6.7.6.2 Results of Spectral Refraction Validation

Figure 6-46 shows a sample of the measured vertical centroid versus wavelength from grating images, corrected to the horizon reference, subtracted from the expected spectral refraction for that wavelength, at that elevation, computed using the mean atmospheric conditions for that night. The reference wavelength for comparison is 550 nm, where the blue and red ends of the spectrum are stitched together. The measurement is made separately for all three combinations of grating and blocking filter used to observe bright stars. Overplotted on top of the measured error in both $m = -1$ and $m = +1$ is a straight-line fit to the error computed from the high-SNR $m = 1$ centroid.

For most stars, the measurement of error with respect to the DCR model is of adequate quality, as indicated by agreement between the linear fit computed from the $m = +1$ data and the $m = -1$ data. Further, agreement between the slopes of the linear fits among the measurements with the three combinations of gratings and filters also indicates that the measured error is real, and the model of atmospheric dispersion can differ from observation by as much as a few hundred milliarcseconds between the red and blue ends of the silicon passband. This is a relative error of a few percent with respect to the actual amount of dispersion.

The roll-off of signal strength at the ends of the $m = +1$ is indicated by an increase in the centroiding noise but especially for the blue end data by a downward rolloff in the centroid. This is not likely to be a real DCR error, but rather an effect of computing a centroid over a diminishing signal. While the $m = +1$ spectral order curls up in the blue, when the signal cuts out the centroid calculation is done over empty sky. The average centroid location of a column of sky noise is closer to zero than any particular extreme. The same phenomenon occurs at the red ends of the $m = -1$ mode, especially in Figure 6-46(c) and 6-46(d). The expected centroid location there is more negative than the reference wavelength and as the signal cuts out the average of empty sky noise comes back to zero.

The entire set of deviations from the nominal refraction model versus elevation angle is plotted in Figure 6-47. There does not appear to be any large-scale systematic trend to these deviations as a function of elevation or any systematic difference between different nights. There are two instances of large deviations from the overall trend: one at high elevation on 2019:094 and one at low elevation on 2018:272. The lack of agreement in the slope among the measurements with the three grating and filter combinations on 2019:094, combined with the large number of nearby measurements without any noticeable deviation from the model, implies that this is simply a bad

measurement in a cluttered environment or on a low SNR star, which can be confirmed in Figure 6-46(f). The case of the spike at low elevation on 2018:272 is less clear-cut. Although it is possible this is a bad measurement, there is agreement among the three measurements and inspection of Figure 6-46(b) indicates that this observation may indeed be real. The weather on 2018:272 was not photometric, with occasional pockets of clouds and water vapor going through the field of view, implying a weather front. If an observation at low elevation happened to go through a patch of warm air, the overall dispersion in the blue may be lower than expected for that one observation, resulting in a larger positive slope.

The amount of deviation from the model is a small percentage of the absolute amount of dispersion. In Figure 6-45, the average error in the slope of a few hundred mas per micron is under 5% below 40 degrees elevation and under 30% at 60 degrees elevation. Although the relative error would seem to be large at high elevation, the implications for DCR prediction accuracy are minor. Also indicated on Figure 6-45 are the elevation angle shifts equivalent to the relative error. That is, below 50 degrees elevation, the amount of DCR error observed is equivalent to a shift of less than 3 degrees along the DCR slope vs. elevation curve, and above 50 degrees, the shift is under 7 degrees. From Figures 4-16, 4-15, and 4-11, it is clear that these shifts would introduce no more than a few tens of milliarcseconds of DCR bias to the average fiducial star or target, and the effect on the measured difference between a fiducial star and a target would be less than that as the bias in observed DCR relative to the modeled DCR occurs in the same direction for both targets and stars.

6.7.7 Discussion of On-Sky Calibration

While mechanical stability of the instrument internals is acceptable, the kinematic alignment of the spectrograph OTA to the telescope mount can be off by many tens of arcseconds. The amount of misalignment depends on the dynamics and the time-history of the satellite tracking motion, but the short-timescale stability is under a few arcseconds. This is acceptable for measurements of satellites with well-defined orbits but necessitates the interleaving of fiducial astrometry frames between grating images in order to infer the pixel location of the spectrum origin for targets too faint for reliable localization of the $m = 0$ PSF directly in the grating images.

The measurement of on-sky seeing of 8 pixels (or about 5 arcsec FWHM) is worse than the design assumption by a factor of about two. The first implication of this is that the $m = 0$ PSF of faint targets will, in fact be at too low of an SNR to detect, necessitating the use of the empirical




Figure 6-45: Magnitude of Observed DCR Error Compared to Absolute DCR

model. The second implication is that the spectral resolution achieved by the instrument is right at the design goal of 40 nm.

The most important conclusion from the checkout measurements is assurance that the assumptions that were used in the analysis and data processing are valid to within the ability of the spectrograph to measure them. These validated assumptions include both the model of spectral refraction and the spectral transmittance of the atmosphere.

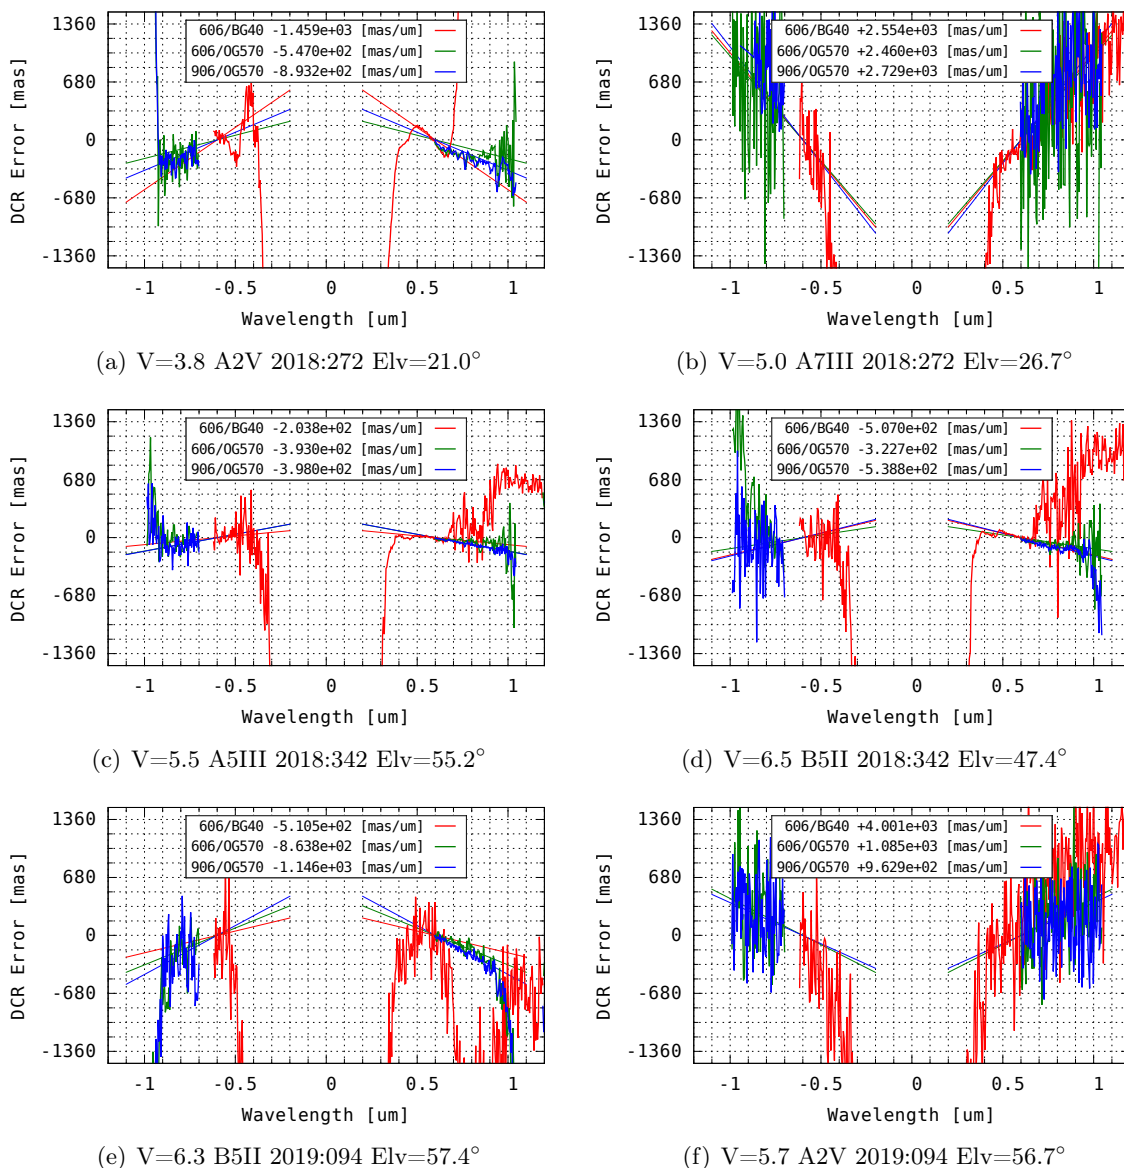





Figure 6-46: Example of DCR Verification Data. Stellar classification and magnitude from Pickles classification [142]. The sense of the error is observed minus modeled. Positive error means the observed centroid appears higher than expected implying more refraction at that wavelength than expected. Negative error implies less refraction at that wavelength than expected. A positive slope means less refraction in the blue and more in the red. A negative slope means more refraction in the blue than in the red. Vertical axis tick marks are placed approximately one CCD pixel apart to demonstrate the scale of the measurements relative to the pixel grid. Note also that seeing for most of the observing conditions was on the order of several pixels, and thus the DCR error inferred from the measurement here is a small fraction of the PSF.




(a) Everything




(b) 2018:272




(c) 2018:279



(d) 2018:342



(e) 2018:346



(f) 2019:094

Figure 6-47: Measured Deviation from Refraction Model vs. Elevation. Measurements from individual nights do not appear to deviate from the aggregate average over all nights. Large deviations from individual stars on any given day sometimes represent bad measurements but can occasionally represent a real, but transient, deviation, such as on 2018:272.

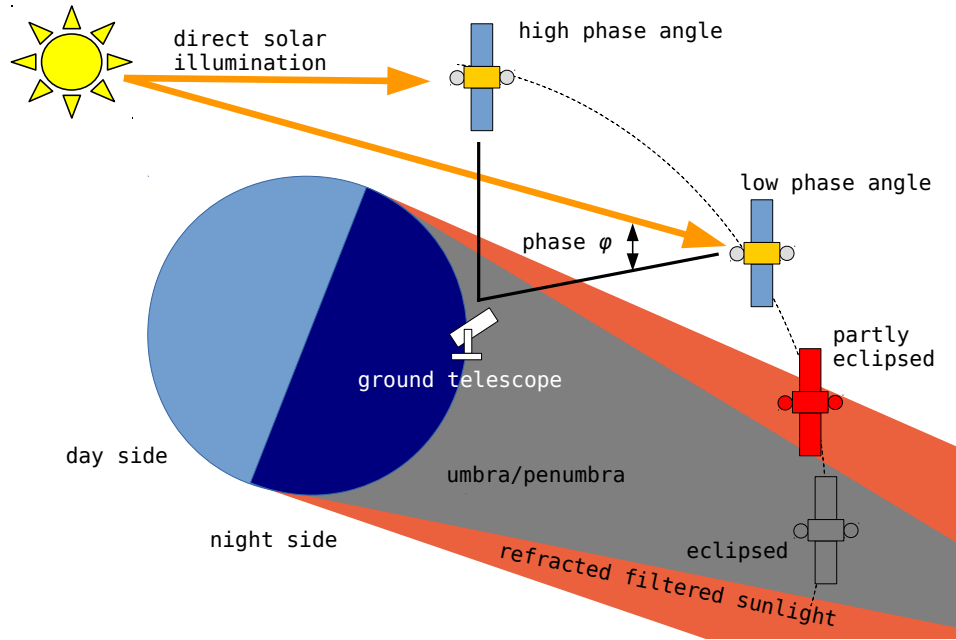


Figure 6-48: Satellite Illumination Conditions. Ground-based telescopes must observe during night time while the satellite is illuminated by sunlight. To first order, the illumination conditions may be characterized by the phase angle ϕ . Nearly all active satellites have movable solar panels that track the Sun. At high phase angles, the Sun is behind the satellite and the front side of the solar panel is not visible to the ground station. At low phase angles, the front of the solar panels are seen face-on from the ground site. When observing from the ground, low phase angles necessarily occur near eclipse. Not all satellite orbits go into eclipse, but if they do, shortly before entering eclipse, the satellite passes through a region of illumination by light that has been refracted and filtered by the Earth's atmosphere. The short wavelength content of this light has been scattered away, leaving very red illumination. This region extends about one degree out from the limb of the Earth. At GPS/GLONASS altitude, this region spans about 450 km and a satellite at that altitude moving directly into the center of the umbra at 4 km/sec will take just under two minutes to traverse this region.

6.8 Satellite Measurements at the Firepond Optical Facility

6.8.1 Planning Observations

Figure 6-48 shows the illumination conditions of a satellite in Earth orbit observed by a ground-based telescope. The purpose of this experiment was to measure a variety of targets through a variety of atmospheric conditions. The variety in atmospheric conditions comes from observing the same satellite over different elevation angles and the variety in targets comes from the fact that the GPS and GLONASS satellites observed during the measurement campaign comprise four different vehicle designs with presumably different optical cross-sections, but mainly the target diversity comes from the variation in illumination conditions.

At low phase angles, one would expect the satellite to present a cross-section dominated by




Figure 6-49: Elevation vs. Phase Angle of Satellites Observed with Spectrograph

solar panels and at higher phase angles one would expect the signature to be dominated by other spacecraft materials as the solar panels track the Sun and present a progressively smaller cross-section to the ground telescope. During each operational night, observations were planned to try to observe the same satellite at a variety of phase angles at roughly the same range of elevation angles. Where possible, instances where the satellites entered or grazed eclipse were prioritized for observation to try to capture a solar panel glint³.

6.8.2 Summary of Data Collections

Satellites were observed over each of the seven operational nights between September 2018 and April 2019. Including calibration measurements and astrometric fiducial frames, a total of 920 GB of data was collected with the 16-inch spectrograph. A total of 1600 GB of bare-CCD astrometric images were collected with the 48" telescope from which 580731 distinct astrometric observations of GPS and GLONASS were formed. The distribution of GPS and GLONASS targets over elevation and phase angle is shown in Figure 6-49.

6.8.2.1 Brightness of Targets

Figure 6-50 shows the average measured brightness versus phase angle of the satellites that were observed during the measurement campaign. The observations are grouped by the type of satellite.

³Recall from Section 2.4.1 that solar panel glints are only a few degrees in angular extent.




Figure 6-50: Measured Brightness of Satellites Observed During Spectroscopy Measurements

Three types of GPS satellite were observed: Block IIR, Block IIR-M, and Block IIF, while all GLONASS satellites observed were of the GLONASS-M type.

The first conclusion that can be drawn from this measurement is that the targets are fainter than expected, thereby lowering useful data volume. The design of the spectrograph assumed a worst-case brightness of 12.5 mag for all targets, and while this is mostly true of GLONASS satellites, the GPS Block IIR and IIR-M satellites are fainter than that at phase angles above 60° . This means that the measurements for which target SNR is high will be biased toward low phase angles, thereby reducing the diversity of target illumination conditions over which the DCR reconstruction technique may be exercised.

The second conclusion pertains to the error budget on the astrometric measurement of these satellites. The GPS Block IIR and IIR-M satellites have a solar panel wingspan of 11.4 meters while the Block IIF satellites have a wingspan of 35.5 meters [144]. The GLONASS satellites are about half a visual magnitude brighter than the GPS Block IIF satellites implying their wingspan is between 30% and 60% larger.

At a typical range of 20,000 km, a one meter offset between the center of mass and center of illumination corresponds to an angular measurement bias of 10 mas. While it is very unlikely that the center of illumination of a satellite is at the very tip of its solar panels, if the center of illumination is somewhere in the middle of one solar panel, then the worst-case astrometric bias for a GPS Block IIR and IIR-M satellite would be 25 mas, but the worst case bias for a GPS Block




Figure 6-51: Measured Sky Brightness at Firepond During Spectroscopy Measurements

IIF or a GLONASS-M would be between 87 mas and 140 mas at near zero phase angle and as high as 100 mas at 45 degrees phase angle⁴. This is a paradox: the target with the brightest signal and therefore the best DCR reconstruction is also subject to the greatest astrometric bias from defect of illumination.

6.8.2.2 Sky Brightness

The design of the spectrograph assumed a flat spectrum sky brightness of 17.5 mag/arcsec². The sky is generally brighter farther away from zenith and the computation of a rigorous SNR threshold to filter out observations that are too faint for good DCR estimation merits a measurement of actual sky conditions. The 48" telescope was used to compute equivalent sky background based on the observed noise strength in the images. The 48" telescope was used rather than the 16" telescope's astrometry frames because the smaller aperture and lower field of view of the 16" telescope resulted in fewer frames with large numbers of stars suitable for a good calibration.

The electron gain of the camera in the 48" telescope is about 0.44 ADU / photoelectron and the per-pixel read noise is about 2 ADU in the operating mode used for collecting the bare CCD astrometry frames. Let Z_0 be the photometric zero point in the image computed from the stars,

⁴If the solar panels track the Sun, the maximum extent is the satellite's wingspan times the cosine of solar phase.

such that for every star with a detected photoelectron signal ϕ_i , and CCD magnitude⁵ M_i ,

$$M_i = -2.5\log_{10}(\phi_i) + Z_0 \quad (6.16)$$

where Z_0 is recomputed for every image. The Z_0 computed from stars (used to transform CCD amplitudes of targets into magnitudes in Figure 6-50) was used to transform the measured sky background noise in the images into an equivalent background level expressed as magnitudes per square arcsecond.

The background noise level was computed as the standard deviation of all pixels not belonging to either the satellite target or a detected star. This noise level is the RSS of shot noise from sky background and read noise. If all noise levels are expressed in photoelectron units rather than ADU's, then the variance of background signal (measured background σ^2 minus separately measured read noise strength N^2) is identical with the sky background photoelectron count per pixel. Thus, the background β term in Equation 4.11 may be computed as

$$\beta = -2.5\log_{10}(\sigma^2 - N^2) + Z_0 \quad (6.17)$$

with a scale correction for the solid angle subtended by the 48" telescope's 0.71 arcsec pixels.

The background level as measured from the photometric zero point in all 48" astrometry frames over the seven operational spectrograph nights is plotted in Figure 6-51. Aside from occasional patches of bright sky from clouds blowing through the field, the sky was mostly fainter than 17.0 visual magnitudes per square arcsecond, but overall was often brighter than the assumed value of 17.5. At elevations above 40 degrees, the 17.5 visual magnitudes per square arcsecond assumption used in the design is valid, but at lower elevations there is some variability among the nights. A more conservative bound from 16.5 visual magnitudes per square arcsecond at 20° elevation down to 18.5 visual magnitudes per square arcsecond at 60° elevation was used to set thresholds for target SNR to compute DCR corrections to a desired accuracy.

⁵The conversion between catalog magnitudes in the Vega system and CCD magnitudes in the Vega system can be defined in any number of ways. The convention I use is that a star's CCD magnitude is the same as that of a star with a Vega spectrum if their photoelectron counts are identical. This is an elevation-dependent conversion. The conversion for Gaia DR2 stars is precomputed with synthetic photometry of Pickles stars similarly to the way the conversion from G , BP , RP and π to DCR estimate is computed.




Figure 6-52: SNR vs. Predicted DCR Estimation Accuracy for Spectroscopy Measurements at Firepond for 40 nm Spectral Resolution

6.8.2.3 SNR Requirement vs. Data Availability

In the analysis of the slit spectrograph presented in Section 4.6.3, it was assumed that the noise in the measured spectrum comes from both shot noise in the divided light from the target and attenuated shot noise from the averaged in-band sky background, where the averaging was defined by the height of the slit relative to its width. For the slitless case, each measured segment of target spectrum is subject to the full sky background. Equation 4.11 may be modified to reflect this difference by assuming that the slit height H , rather than a multiple of its width, is a fraction equal to the ratio of its width relative to the full measured passband. This effectively transforms $1/\sqrt{H}$ attenuation of background shot noise by averaging into a \sqrt{H} increase in background shot noise by addition over the full spectrum.

Making this modification and using a worst-case sky background measurement from the operational nights at Firepond, the SNR (as measured in the 48" telescope with 120 msec exposure) may be plotted against the predicted DCR estimation accuracy from the 16-inch telescope (for spectra measured over 24 second exposures). This is shown in Figure 6-52, where it is evident that the 48" SNR threshold to achieve 100 mas or tighter⁶ reconstruction may be accurately modeled by a quadratic exponential function of elevation angle. Applying these thresholds to the observed SNR

⁶Recall that even though stellar DCR can be reconstructed to better than 20 mas and the design goal of the spectrograph was 25 mas accuracy in the DCR reconstruction, the astrometry error budget due to unmodeled defect of illumination can be as high as 100 mas.

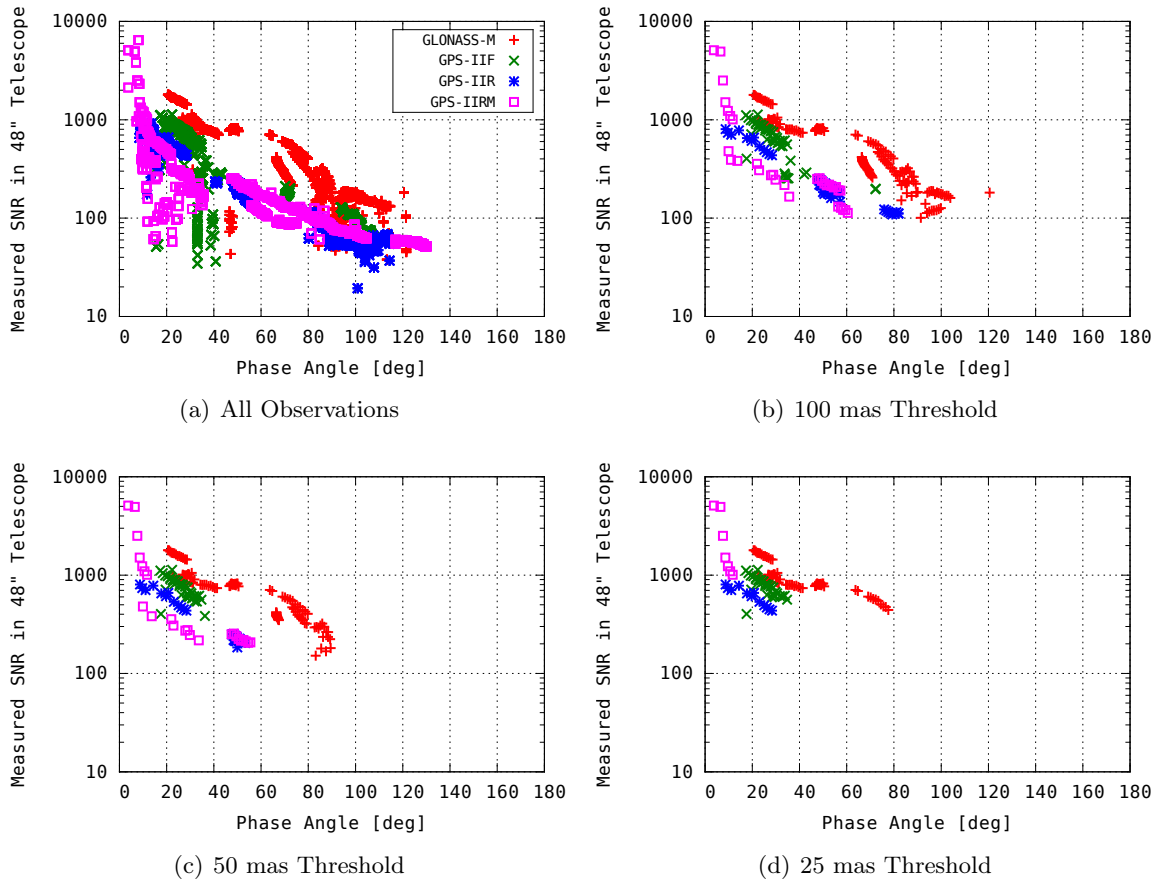


Figure 6-53: Data Volume vs. Desired DCR Estimation Accuracy

of our satellites (Figure 6-53) shows that with a 100 mas desired accuracy, data is present up to a phase angle of 100° . Tightening the desired accuracy down to 50 mas removes all instances of phase angles greater than 90° removes all observations of smaller GPS satellites above 60° phase. Lowering the desired accuracy further still down to the original design goal of 25 mas leaves data only at very low phase angles and thereby removes a large amount of target diversity over which the technique may be exercised. Therefore, the remaining analysis here will be restricted to data that meets the 50 mas and 100 mas accuracy requirement.

The original design analysis assumed that spectra could be averaged over several cycles of the gratings to yield a tighter reconstruction. In Figure 6-53(a), the bulk of the data at high phase angles requires an SNR boost by a factor of two in order to be included. The SNR requirement is inversely proportional desired accuracy and without that SNR boost, the predicted DCR estimation error on those observations is in excess of 200 mas. The effective SNR boost gained from averaging is roughly \sqrt{N} , thus doubling the SNR would require averaging over four measurement intervals.

In practice this proves difficult for GPS and GLONASS satellites. These targets move at about 40 arcsec/sec, meaning that over the roughly two minutes of the full measurement cycle, they move about 1.3 degrees, and an average spectrum measured over more than four cycles can span a phase angle difference of up to 6 degrees. Between 20° and 60° phase, the targets all appear to gain or lose signal strength at a rate of about 0.05 magnitudes/degree phase, meaning that the 6 degree span of the measurement could see the target changing brightness by as much as 30%.

If one assumes that for a box/wings satellite model, most of that brightness change comes from the solar panels' apparent cross-section changing, then this growth in total brightness could look like as much as 30% increase in solar panels, which for a target that is already mostly solar panel can be a more than ten percentage point change in the ratio of solar panel in its total signature. From Figure 4-16, we see that a ten percentage point change in solar panel composition for a target that is mostly solar panel can cause a DCR difference of more than 50 mas below 50 degrees elevation. Thus, this attempt at averaging would gain 100 mas of accuracy but would lose up to 50 mas.

Although this is a worst-case analysis, the point is that for a 16" spectrograph observing moving targets, averaging over longer intervals would be pushing further into the point of diminishing returns. The impracticality is the result of the fact that the GPS and GLONASS satellites observed for this experiment move at up to 40 arcsec/sec, changing phase angle quickly. Geostationary satellites are farther out and move more slowly relative to the sky, and at 15 arcsec/sec they would experience at worst a 20 mas change in DCR over the same averaging interval.

6.8.3 Spectroscopy Results

6.8.3.1 DCR Compensation with Spectroscopy

Figure 6-54 shows the predicted vs. observed DCR bias for all GPS and GLONASS satellites observed over the seven operational nights. The astrometric observations are generated using the Gaia DR2 star catalog, processed with per-star DCR correction, and residuals are computed with parallactic refraction included in the observation model. For clarity, each point represents the mean of residuals over a 75-second averaging period for each spectroscopic measurement. Astrometric observations are included in the average if they meet the relevant SNR requirement, if they are the result of an astrometric plate solution with either 20 or more stars in star-dense regions or five or more stars in sparse regions, and if the horizontal residual of each individual observation is under 300 mas. Only measurement periods with at least 30 distinct observations are included.

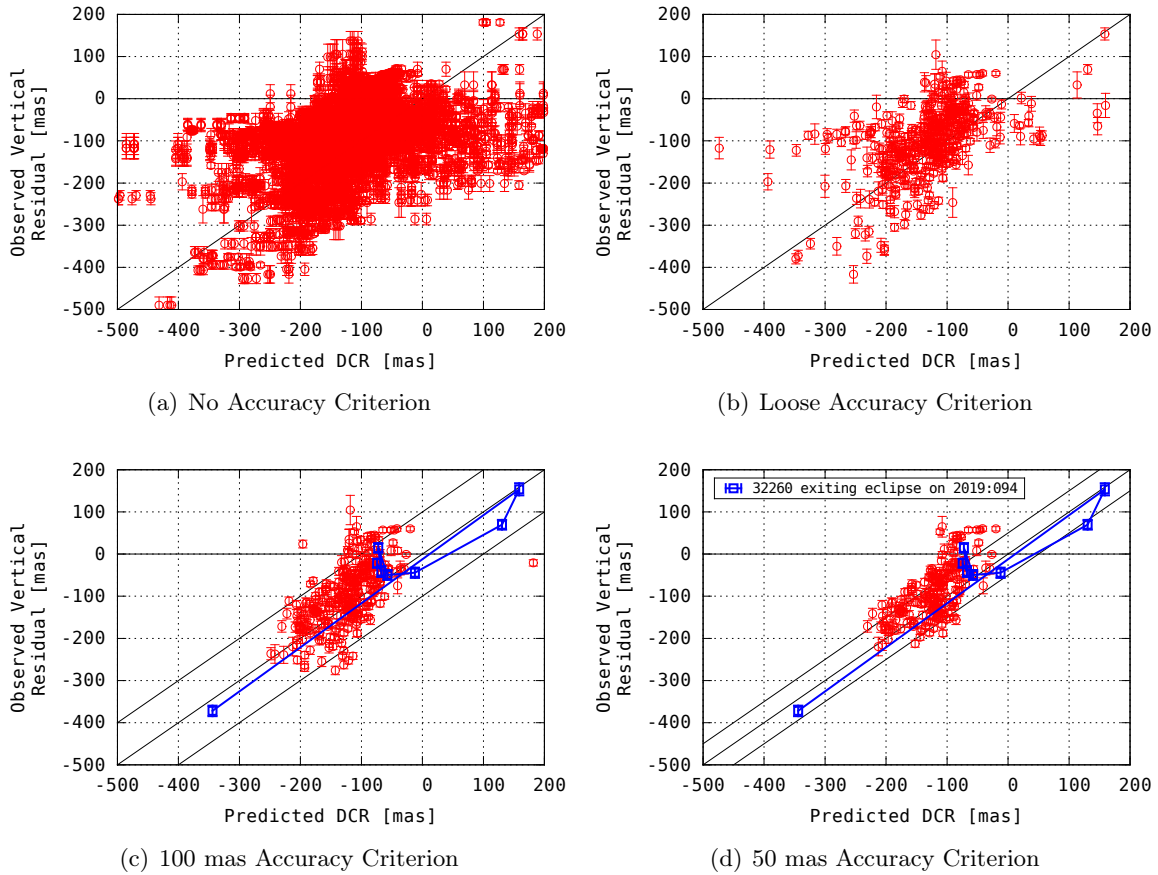


Figure 6-54: Predicted vs. Observed Vertical Astrometric Residual of Satellites with Measured Spectra. Input astrometry is processed using the Gaia DR2 star catalog with per-star DCR corrections in each image frame and parallactic refraction is accounted for in the observation model. Datapoints in the plot are selecting according to the minimum SNR criterion over 75-second averaging periods. (a) With no accuracy criterion applied, the approximately 12,000 datapoints show only a weak correlation between observed and predicted DCR bias. (b) When the accuracy criterion is tightened, but still left loose enough to allow at least one datapoint per observed satellite, the correlation becomes more evident. When the minimum SNR criterion is strictly applied at the expense of excluding observations of some satellites in panels (c) and (d), the agreement between prediction and observation is quite good and deviations from predictions are within an order of magnitude of the error bounds used to generate the SNR threshold. The highlighted blue trace in panels (c) and (d) represents the single instance of a satellite observed emerging from eclipse and glinting blue (see Section 6.8.3.4).

There is overall good agreement between the DCR bias as predicted from spectroscopic measurements and the DCR bias as observed in the astrometry from the 48" telescope. Furthermore, the deviations between prediction and observation generally stay within the error bounds used to generate the SNR criterion. It should be noted, however, that the distributions are not Gaussian, thus the relation between the $1\text{-}\sigma$ specification of the accuracy criterion and the fraction of residuals outside the $1\text{-}\sigma$ bounds is not expected to follow Gaussian statistics.




Figure 6-55: Distributions of Residuals Under Different Processing

6.8.3.2 Analysis of DCR Prediction Error Statistics

Recall that for a parallactic refraction correction to be computed, the range to the target must be known. Strictly speaking, this is not possible using only two-dimensional image data. For this reason, this correction must be implemented as part of the observation model in the orbit determination software, where a range estimate is by definition accessible. This logical split between parallactic refraction correction and image processing has implications for how data are likely to be analyzed. The present state-of-the-art in several operational orbit determination software packages is to assume that input observations are fully corrected for all physical biases and to make atmosphere-agnostic orbit fits. Similarly, the present state of the art is to make atmosphere-agnostic reductions when processing imagery.

This means that no stellar DCR compensation is made in the data reduction, no target DCR correction is applied to observations, and in many cases it means no parallactic refraction is included in the observation model. Because all of these corrections would logically be applied at different stages of the whole processing chain, it is illustrative to measure the effect of stellar DCR, target DCR, and parallactic refraction together and separately.

Figure 6-55 shows the error distributions of vertical residuals under the conditions of:

1. Raw astrometry: no DCR correction and no parallactic refraction in observation model;
2. No DCR correction but parallactic refraction in observation model;

Table 6.3: Distribution Statistics of Residuals Under Different Processing

Criterion	Processing	Mean [mas]	Std. Dev [mas]	N
All	raw horizontal	-17	31	12508
All	raw vertical	-62	97	12508
100 mas	raw vertical	-61	100	312
All	parallax only	-19	87	12508
100 mas	parallax only	-29	94	312
All	no parallax, stellar DCR	-141	93	12468
100 mas	no parallax, stellar DCR	-131	102	312
All	parallax, stellar DCR	-96	83	12476
100 mas	parallax, stellar DCR	-94	77	312
100 mas	parallax, stellar DCR, target correction	24	68	312
100 mas	parallax only, target correction	90	105	312
50 mas	parallax, stellar DCR, target correction	27	61	223

3. Stellar DCR correction but no parallactic refraction;
4. Stellar DCR correction and parallactic refraction;
5. Stellar DCR correction, parallactic refraction, and per-target DCR correction, both with no minimum SNR and with a 100 mas accuracy criterion.

Table 6.3 summarizes the statistics of those distributions, with a distinction between all input data points and just those with a high-enough SNR to enable 100 mas correction. Note that different input constraints on the astrometric processing will result in slightly different ensembles of stars in the different frames, resulting in more or fewer astrometric solutions that pass the goodness criteria for inclusion into the averaging periods. This is why sample sizes for each distribution vary.

The horizontal residuals are essentially zero-mean given the error budget from defect of illumination, and their 31 mas standard error over the 75 second averaging period sets the order of magnitude of the error floor for the rest of the other measurements. It is evident that the astrometry in the vertical direction is both noisier and more biased than in the horizontal direction. With the atmosphere-agnostic data reduction and observation model, the mean bias is 62 mas, or 2σ with respect to the horizontal baseline, and the noise is more than three times worse. Including the effect of parallactic refraction reduces the bias but does not make an appreciable dent in the noise. Applying stellar DCR compensation but neglecting to include parallactic refraction in the observation model makes the bias considerably worse.

Only when the parallax correction is included in the observation model and the astrometric reductions are corrected for stellar DCR does the astrometric noise start to come down from about 90 mas to 77 mas. But again, if no per-target DCR correction is made, the bias is still large relative

to the accuracy baseline. Once data reductions are made with stellar DCR corrections, parallactic refraction is included in the observation model, and per-target DCR corrections are applied, then the mean residual comes down to under 30 mas (comparable with the horizontal residual), and the noise level is reduced from near 100 mas to below 70 mas, a more than 30% performance improvement over atmosphere-agnostic processing.

The most important takeaway from these distributions is that everything has to be in place for per-target stellar DCR correction to make a quantifiable impact. If stellar DCR, target DCR, and parallactic refraction corrections are made in isolation, no performance improvement will be seen and in some cases performance will visibly degrade⁷.

6.8.3.3 Difference in Atmosphere vs Difference in Targets

An important question to ask here is whether the variation in DCR bias spanning Figure 6-54 comes from a diversity of targets or a diversity of atmospheric conditions. This question is answered by plotting the observed vertical residual attributable to DCR bias against elevation. Figure 6-56 breaks down Figure 6-54 by elevation range. This shows that while there is some variation of target DCR within narrow elevation bands, it is the atmosphere that is responsible for most of the variation in the satellites' DCR bias. The one exception is SCC 32260 exiting eclipse (Section 6.8.3.4).

Plotting the entire set of vertical residuals versus elevation in Figure 6-57 shows again that most of the bias variation for the targets with high SNR appears to come from variation in elevation rather than variation in target type. Furthermore, for most of the data for which there are good spectra, and for much of the data where there are not, the DCR bias is well-approximated by assuming that the target is a mixture of 60% solar panel / 40% Kapton.

The statistics of the error distributions for astrometry corrected with spectroscopy and astrometry corrected with the constant mixture model are shown in Figure 6-58. The correction made with the constant 60% solar panel / 40% Kapton mixture model is marginally better than the correction made with concurrent spectroscopy, with the mean error reduced from 24 mas to 4 mas and the standard deviation reduced from 68 mas to 61 mas.

While this would seem to call into question the utility of spectroscopy when a blind correction

⁷In practice, such a scenario can discourage further exploration in the presence of operational requirements and deadlines. When he was working on the metric accuracy assessments for the SST operational software in the early 2010s, Alex Szabo tried to find evidence of DCR bias and parallactic refraction in residuals of GPS satellites reduced with the early UCAC catalogs. He was unable to find any definitive evidence of these phenomena, and in retrospect it was never going to work without simultaneously accounting for parallactic refraction, stellar DCR, and the DCR of individual satellites.

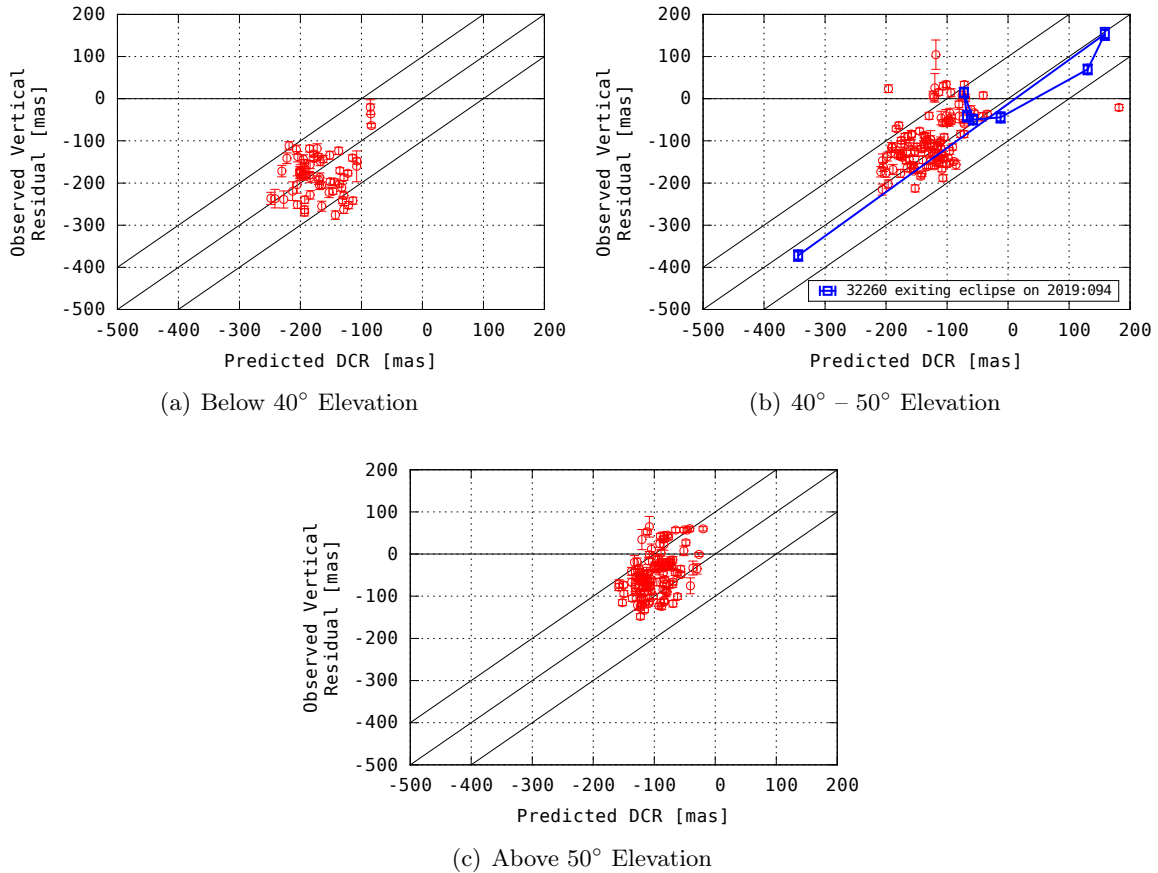


Figure 6-56: Observed DCR vs. Predicted Separated by Elevation Range. Same as Figure 6-54(c) (100 mas accuracy criterion) except broken down by elevation range to examine whether any target spectrum variation was observed through the same atmosphere. Although there is some variation within the bands, most of the predicted and observed DCR within each elevation band forms a single cluster about 100 mas across in both axes. The exception is the instance of 32260 emerging from eclipse, whose spectrum was known to undergo a drastic change while being observed through the same atmosphere.

seems better, it is important to recall that the spectra measured in this experiment are mostly confined to phase angles below 90° , with the majority of observations of GPS satellites occurring at phase angles below 60° . It is not surprising that in these restricted conditions, the satellites present what is essentially a constant mixture of materials to the observer. One would expect that at higher phase angles, the targets would look different than at low phase angle. Further, the case of SCC 32260 exiting eclipse demonstrates why spectroscopy is necessary: while most of the time “satellites are spectroscopically boring⁸,” sometimes they are not. Occasionally the satellite will, either through chance geometric alignment or as part of its normal operations, present a radically different cross-section to the observer.

⁸As Dave Monet likes to say.




Figure 6-57: Observed DCR vs. Elevation. The DCR bias of most satellites is well-approximated by that of a spectrum consisting of a 60% solar panel / 40% Kapton mixture.




Figure 6-58: Distributions of Observations Corrected with Spectroscopic Measurements and with Solar Panel/Kapton Mixture Model. Correcting with spectroscopy, the distribution has a mean of 24 mas and a $1\text{-}\sigma$ spread of 68 mas. Correcting with a constant 60% solar panel / 40% Kapton mixture model yields a mean residual of 4 mas and a $1\text{-}\sigma$ spread of 61 mas.




Figure 6-59: Measured Brightness of SCC 32660 Exiting Eclipse.




Figure 6-60: Reconstructed Spectra of SCC 32660 Exiting Eclipse.

6.8.3.4 Blue Glint

During the spectroscopic measurement campaign, instances where the satellites entered or grazed eclipse were prioritized for observation to try to capture a solar panel glint. Of the 14 distinct satellites observed over the seven operational nights from September 2018 to April 2019, only one instance of a blue glint was observed, that of SCC 32260 exiting eclipse on 2019:094. Interestingly, SCC 32384 was observed entering eclipse shortly afterward, but no glint was seen. The rarity of the blue glint merits further examination of the one that was observed.

Figure 6-59 shows the measured brightness of the satellite as it transitioned the partial eclipse region and entered full sunlight. Underlaid beneath the signature is the timing of the multiplexed

gratings and blocking filters used to make the spectrographic measurement. The time between detection of the satellite signal and its peak brightness is on the order of two minutes and it is worth noting that during partial eclipse, the target brightness ranges over five visual magnitudes between the center of the blue end measurement with the 606/BG40 grating/filter combination and the measurement of the red end of the spectrum with the 906/OG570 grating/filter combination. The target saturated the 48" telescope at peak brightness, so it is likely that the blue glint peaked brighter than 5th magnitude, but the point is that during the glint, in the approximately 45 seconds between the center of the blue and the red measurement on the first filter wheel cycle to capture the glint, the target brightness dropped by at least 20%, and it dropped by about 40% during the second cycle to capture the glint (Figure 6-59(b)).

The effect of the rapidly changing signature between measurement of the two ends of the target spectrum is visible in the reconstructed spectra in Figure 6-60. It is almost miraculous that the effect of the rapidly changing signature did not spoil the measured spectra to the point where DCR predictions derived from them were not usable.

Cycle 1 measures the target as it emerges from eclipse and undergoes a 7 visual magnitude change in brightness between the start of the blue measurement and the end of the red measurement. The spectra for each combination of grating and filter are derived from the median of the grating images, but that is not very satisfying as that still leaves about 5 visual magnitudes worth of signal growth between the red and the blue. What makes the reconstruction valid is that despite the nominal change between the blue and red ends of the measurement, while the satellite is in partial eclipse, almost all blue light in the illumination spectrum is filtered out by the atmosphere (Figure 6-61) and there is no blue signal to measure during Cycle 1.

The effect of the changing signature is more apparent in the spectra reconstructed during Cycle 2 and Cycle 3. During both of these measurement cycles, the solar panels were glinting blue and the overall target brightness during the blue measurement was higher than during the red measurement. This is reflected in the straight line stitching together the red and blue ends of the reconstructed spectrum in Figure 6-60. Despite this artifact caused by the need to multiplex passbands and gratings during the measurement of the full spectrum, the cosmetic match between the reconstructed spectrum and a 100% solar panel model is quite good, validating the intuition that the signature of a satellite near zero phase is dominated by a solar panel blue glint.

The agreement between observed signatures and the solar panel model further validates the assumptions used in Chapter 3 to predict the range of possible DCR bias of satellite targets, and




Figure 6-61: Estimate of Atmospheric Transmission into Partial Eclipse Region

the processing of the NOFS slitless spectroscopy using a model of satellite signatures as a mixture of nominal material reflectances derived from laboratory measurements.

6.8.4 Experiment Summary

For this thesis, I constructed a slitless spectrograph with COTS equipment and used it to measure the silicon passband spectra of GPS and GLONASS satellites under a diversity of atmospheric and illumination conditions. The instrument multiplexes between two transmission gratings and suitable colored glass blocking filters to collect a full spectrum every 75 seconds on targets brighter than 12.5 magnitude. These spectra were used to infer a DCR bias for the target for that time period.

Simultaneously, the Firepond 48" telescope collected high frame-rate bare silicon images of the same targets. The images were processed into observations with color-aware astrometric reductions using the Gaia DR2 star catalog. The bare-silicon astrometry was compared with the final orbit solutions published by the International GNSS Service and the inferred DCR bias was compared against the astrometric residuals of observations. 312 unique spectroscopic measurements on 14 distinct satellites with SNR sufficient for DCR compensation were collected between 20 and 60 degrees elevation over seven observing nights from September 2018 to April 2019.

Good agreement was achieved between the predicted DCR bias using the spectra and the observed DCR bias from the bare CCD astrometry. Further, the majority of the measurements indicate that the DCR bias of both GPS and GLONASS satellites may be adequately modeled by assuming the satellites are composed of 60% solar panel and 40% Kapton. However, unusual illumination conditions can perturb the DCR bias away from this nominal value by several hundred milliarcseconds. Partial eclipse during times where the satellite grazes the shadow cone cut out all short-wavelength illumination and shift the DCR bias downward by several hundred milliarcseconds. Alternatively,

glints off of the solar panels at low phase angles can significantly increase the proportion of blue light in the target spectrum and shift the DCR bias upward by several hundred milliarcseconds. While these unusual illumination conditions are rare and short-lived for GPS and GLONASS satellites, and only one such case was observed over seven nights, they may be encountered more frequently in observations of geostationary satellites and may last tens of minutes during certain times of year, thereby introducing bias into orbit solutions generated from uncompensated observations that overlap with these events.

In concert with the satellite measurements, bright stars were observed in order to validate the models of atmospheric spectral transmittance and spectral refraction that were used to compute stellar DCR and infer targets' DCR bias from their spectra. Although there was a single possible instance where, during non-photometric high humidity conditions, the spectral refraction model deviated significantly from predictions, the measurements indicate that on average, the assumptions used in the analysis portion of this thesis are accurate to a level necessary to compensate for target DCR to an accuracy of 25 mas on sufficiently bright targets above 20 degrees elevation.

6.9 Chapter Summary

This chapter has addressed the final original contribution of this thesis:

3. A DCR mitigation technique based on concurrent spectroscopic measurements is developed. The efficacy of the technique is demonstrated on measurements of GPS and GLONASS satellites at the MIT Lincoln Laboratory Firepond Optical Facility in Westford, MA. In data collected over seven nights, a reduction in the magnitude of residual DCR bias by 60% with respect to color-agnostic processing is shown.

Measurements of a subset of geosynchronous satellites from NOFS demonstrated that the DCR bias for active satellites can be highly variable, with some targets' DCR bias changing by 100 mas over the course of a few tens of minutes and varying by as much as 100 mas night-to-night. Having provided the necessary motivation for spectroscopic measurements of satellites for DCR compensation, a proof-of-concept measurement using multicolor photometry of GLONASS satellites entering eclipse demonstrated that real-time DCR compensation with concurrent measurements is feasible.

Toward that end, a slitless spectrograph was constructed with COTS equipment and deployed at the MIT LL Firepond Optical Facility in Westford, MA. This instrument was operated over seven

nights between September 2018 and April 2019 and was used to measure the full silicon passband spectra of bright stars and GPS and GLONASS satellites.

Measurements of stars were used to validate the assumptions used to make predictions in this thesis. Measurements of the stars' spectra validated the atmospheric spectral transmittance and instrument overall spectral efficiency models to within 5 – 10 percentage points and measurements of the morphology of the spectral orders were used to validate the spectral refraction model to a comparable level of accuracy.

Measurements of the silicon passband spectra of GPS and GLONASS satellites were made under a diversity of atmospheric and illumination conditions yielding 312 unique spectroscopic measurements on 14 distinct satellites with SNR sufficient for DCR compensation over the seven nights. Using the current state-of-the-art analysis of those observations, that is: no parallactic refraction correction, no stellar DCR correction, and no target DCR correction, the distribution of vertical error in this data set is -61 ± 100 mas. Applying all three of these corrections improves this to 24 ± 68 mas. Constraining the data set to 223 higher-SNR observations yields an improvement in the spread of the distribution to 27 ± 61 mas. The overall improvement in the magnitude of the bias is approximately 60% and the reduction in the spread is approximately 30%. Failing to account for any of these three effects simultaneously yields to either a larger mean bias or spread.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 7

Conclusion

7.1 Summary of Original Contributions

7.1.1 Dominance of DCR Bias in Ground-Based Optical Astrometry

This thesis demonstrated that DCR bias is a dominant long-timescale source of systematic error in ground-based astrometric observations made with telescopes using silicon sensors. The bias is detectable and significant in observations made using both unfiltered detectors and narrow passband filters.

Analysis of archived Minor Planet Center observations of asteroids with known spectral types and well-constrained orbits made with bare silicon detectors reveals the unambiguous signature of DCR bias. The systematic error in these observations is confined to the station-local vertical coordinate and is near zero-mean in the horizontal coordinate. A separation in the vertical, but not horizontal, mean residual of redder S-type versus neutral C-type targets demonstrates that it is DCR bias that is responsible. Analysis of over 165,000 distinct bare-silicon astrometric observations of GPS and GLONASS satellites made as part of regular calibration checks at the Firepond Optical Facility between 2012 and 2016 also exhibit a near zero-mean error in the horizontal residuals and a significant non-zero mean vertical residual against IGS reference ephemerides derived from multifrequency RF measurements insensitive to DCR bias.

Archived observations of asteroids of known spectral types made with narrow passbands exhibit the same separation in the vertical residuals of S-type versus C-type targets as the bare-silicon observations. A high-volume observing campaign made with both bare-silicon and selected Sloan passbands during the 2017 near-Earth flyby of the asteroid 3122 Florence shows a larger residual

in the narrow passband data than in the bare silicon data, demonstrating conclusively that narrow passbands alone do not always mitigate DCR bias.

7.1.2 Consequences of Unaccounted-for DCR Bias in Orbit Solutions

This thesis explored the consequences of failing to account for DCR bias in orbit solutions for Solar System bodies and Earth satellites. Several scenarios were presented where the bias in the orbit solution resulting from the bias in the input observations can be of academic or practical significance.

In the Kuiper Belt, DCR bias can be as high as 200 milliarcseconds for Barucci Group IV objects such as Quaoar, and this bias can prevent successful planning for stellar occultations, even with several years' worth of observation arc.

For inner Solar System asteroids, DCR bias in input observations results in short-arc orbit solutions biased in Keplerian mean motion with respect to ground truth. This bias is small, on the order of an arcsecond per year or less, but is sometimes more than 5σ with respect to its formal error bounds. Though this will not make the difference between a successfully recovered and a lost asteroid, there exist orbits for which over several apparitions this mean motion bias can appear to be a physically-plausible Yarkovsky acceleration in an optics-only orbit solution.

In the particular case of hazardous asteroids, there exist orbits where failing to account for DCR bias can make the difference between predicting a near miss and an impact. A hypothetical Apophis-like asteroid is one such example, where in the counter-factual scenario of a later discovery and absence of radar observing opportunities, an optics-only orbit solution from unbiased 0.25 arcsec RMS observations between 2016 and 2029 would predict a true impact in 2036 with a $\approx 10\%$ probability of detection while an orbit solution derived from biased observations would predict a 99.94% probability of a miss.

The prediction of impact points of hazardous asteroids that are known impactors is also sensitive to DCR bias. Late discoveries of asteroids that emerge from inside 1 AU in the months leading up to an impact are particularly susceptible as most of the ground-based observations from which the impact point estimate is computed must be made at low elevation. For both Aten and Apollo class asteroids, there exist scenarios where in the months leading up to impact, the running estimate of the impact point can be biased away from its true location by many hundreds of kilometers to a high degree of false confidence.

Most observations used to track RSOs in the geostationary belt come from optical telescopes with unfiltered silicon detectors. Unaccounted-for DCR bias in observations from those sensors

results in a bias in the SRP AMR estimate for orbit solutions spanning only a single night. This has implications for measuring the properties of HAMR objects and for predicting the trajectories of debris objects for the purpose of collision avoidance. The DCR-induced AMR bias introduces a Cartesian error of tens of meters per day into short-arc orbit solutions. It is also significant in applications requiring very precise orbit solutions for applications such as resolved imaging of geostationary satellites using natural guidestars.

7.1.3 Compensation of DCR Bias with Concurrent Spectroscopy

Stellar DCR is correctable on a per-star basis during astrometric reductions using the multicolor photometry and parallax measurements in the Gaia DR2 catalog. Compensation of DCR bias for SSA targets requires per-target spectroscopic characterization.

A slitless spectrograph designed to measure the spectra of GPS and GLONASS satellites at a resolution of 40 nm over the entire silicon passband was built and deployed at the Firepond Optical Facility. Over seven operational nights between September 28, 2018 and April 4, 2019, measurements of bright stars were used to validate atmospheric refraction and spectral transmittance models used in the analysis in this thesis. DCR compensation was demonstrated with 312 unique spectral measurements of 14 distinct GPS and GLONASS satellites at phase angles between 0° and 100° , improving the mean vertical residual from 61 mas to 24 mas and improving the 1σ spread in vertical residuals from 97 mas down to 68 mas with respect to atmosphere-agnostic processing.

Most of the GPS and GLONASS satellites maintained a constant spectral reflectance equivalent to that of a target composed of 60% solar panel and 40% Kapton, and much of the variation in DCR bias observed over that target set was caused by variation in atmospheric conditions over a range of elevation angles rather than by a change in target spectrum at a constant elevation angle. For the case of the GPS and GLONASS targets observed in the experiment, using the assumption of constant reflectance spectrum results in a marginally tighter error distribution than that achieved with real-time compensation with measured spectra.

A notable exception to the constant signature assumption was observed on the last operational night when SCC 32260 was observed exiting eclipse and presented a blue solar panel glint to the sensor. Though dramatic changes in spectral signature are rare and short-lived for GPS satellites, spectroscopic measurements of the geostationary belt made from the 1.3 meter telescope at the United States Naval Observatory Flagstaff Station indicate that these events may occur more frequently and with longer duration for RSOs in GEO, thereby necessitating spectroscopic measure-

ment rather than the use of a constant reflectance signature assumption.

7.2 Recommendations

The significance of DCR bias and the proof-of-concept demonstration of a compensation technique motivates several recommendations for data collection, data reduction, and data exploitation in orbit prediction software.

7.2.1 Astrometric Data Collection

Local meteorological conditions, especially meteorological visibility, should be measured and recorded at the time of observation to enable DCR compensation during the data reduction phase. The configuration of the observing instrument, including the usage of optical filters, the existence of any coatings on other optical elements, and the focal plane quantum efficiency, should be recorded and that record maintained along with raw imagery. This especially applies to observations of asteroids, whose faint signal implies that spectra for DCR compensation are unlikely to be available at the time of astrometric measurement.

7.2.2 Astrometric Data Reduction

Operational data reduction software should be upgraded to use the Gaia DR2 star catalog and to apply per-star DCR compensation using Gaia DR2 photometry and stellar parallax, the telescope configuration, and the meteorological conditions at the time of observation. Pending the release of better-calibrated per-star photometry in future Gaia data releases, meta-information on the stars used for each observation and the per-star stellar DCR compensation used in the solution should be retained to enable reversibility and future improvement.

Validation measurements conducted at Firepond indicate that the standard models included in MODTRAN predict atmospheric spectral transmittance to an accuracy sufficient for DCR compensation. The atmospheric spectral refraction models published by either Filippenko [1] or in the Explanatory Supplement [67] are also sufficient to compute DCR.

Unless spectroscopic characterization of the target being observed is available at the time of data reduction, reduced observations should not be corrected using any nominal models. Further, any DCR correction that is applied to the target should be clearly indicated in the observation so that the process is reversible in the event an improved spectroscopic measurement becomes available after

the fact. While this may not be relevant for a sensor and software package designed to make DCR-compensated measurements of Earth satellites, it is relevant for asteroid and KBO observations where astrometry will be available before spectra are.

7.2.3 Orbit Fitting

Parallactic refraction should be implemented in the optical observation model in any piece of orbit determination software used for operational purposes. Because the computation of parallactic refraction requires a range estimate, it is logical to compute it in the observation model of an orbit determination routine rather than to attempt to apply it during the measurement or data reduction phase where nominally it is only angles, and not range, that are immediately available.

Above GPS altitude, parallax is insensitive to the actual amount of refraction with its range of variability with meteorological conditions, and thus the implementation within orbit determination software need not be conditioned on knowledge of local meteorological conditions at the time of observation.

Given that some measurements will occur without per-target DCR compensation, orbit determination software should be modified to allow injection of DCR compensation on a per-target basis along the station-local vertical direction to allow after-the-fact corrections to be made. In a related vein, the software should also be modified to enable independent weighting of residuals along station-local horizontal/vertical coordinates rather than sky-aligned RA/Dec coordinates. Analysis of residuals should also take place in station-aligned rather than sky-aligned coordinates.

Lastly, the software should also allow completely discarding elevation data for certain situations where accepting more biased data results in a worse orbit solution than what results from accepting less unbiased data. Though this may be impractical for single-site SSA of the geostationary belt, it can be practical for long data arc observations of Solar System bodies.

The reason to recommend that all of these modifications be made to the orbit determination routine rather than as external tools is to reduce the complexity by placing all of the sky-aligned-to-station-local coordinate transformations into one place, where they may need not even be immediately accessible to a user.

7.3 Future Work

Aside from the immediate recommendations for operational software presented in the previous section, the results of this thesis motivate several open-ended measurement campaigns and avenues of research.

7.3.1 Further Spectroscopic Measurements

7.3.1.1 Asteroids and KBOs

Of the several hundred thousand numbered asteroids in the Minor Planet Center database, only a few thousand have been spectroscopically characterized. Of the asteroids that have been characterized by the Tholen or SMASS surveys, few have had their spectra measured blueward of 450 nm. Outside of a few measurement campaigns such as that of Thomas et al. [51], there are few instances of the same object being deliberately measured at different parts of its orbit to capture any phase-angle or temperature dependence of its reflectance spectrum.

This leaves a potentially large gap in the ability to predict the DCR bias for many asteroids, including many asteroids for which radar observations are not possible. Given that ground-based sensors will always be less expensive and thus more available than space-based sensors, it is reasonable to assume that the majority of astrometric observations of these objects will continue to be subject to DCR bias.

The large time-aperture products required to achieve corrections to DCR bias commensurate with the accuracy of modern star catalogs will require prioritization of observations. Roughly speaking, the value gained from DCR compensation will be greatest for relatively short-arc near-Earth asteroids with few radar observing opportunities, Kuiper Belt Objects of interest in occultation measurements, and non-hazardous asteroids with potential Yarkovsky accelerations that can only be measured from optical astrometry.

The time-aperture requirement for such a measurement campaign presents a use case for existing and forthcoming large-aperture telescopes. The cost of observing time on these instruments for spectroscopy of asteroids and KBOs must be weighed against the scientific or practical benefit gained from refined astrometry, the cost of space-based observing platforms, observing asteroids and KBOs at low zenith distances from more sites, observing to low elevations from more sites while discarding azimuth in order to retain accuracy while improving precision, and incorporating higher-fidelity uncertainty modeling in orbit determination routines.

7.3.1.2 Earth Satellites

The majority of the spectroscopic measurements on GPS and GLONASS satellites made for this thesis were restricted to low phase angles. All of the measurements were on satellites belonging to one of four common spacecraft designs. The uniformity of illumination conditions and target type is part of the reason that the majority of the observed DCR bias was consistent with a target that was 60% solar panel and 40% Kapton.

If it is indeed possible to fully capture satellite DCR as a combination of a small number of basis materials, then it would be possible to make a spectroscopic measurement sufficient to compensate for target DCR at a resolution coarser than 40 nm and/or at a lower SNR than what is necessary for a measurement predicated on a direct measurement of the entire spectrum. The only way to determine whether this is true is to conduct spectroscopic measurements on a larger variety of satellites and at a wider variety of illumination conditions.

The wide variety of geostationary satellite types launched since the 1960s and the prominence of the geostationary belt as an orbital regime where the majority of SSA observations come from ground-based optics, combined with the way DCR bias projects into AMR bias in that orbit, suggests that measurements of geostationary satellites are the most interesting given that the efficacy of DCR compensation has been demonstrated. The 16-inch spectrograph constructed for this thesis will be used to make preliminary measurements of this orbital regime at the Firepond Facility toward that objective.

7.3.2 Handling Existing Data

7.3.2.1 Reprocessing Raw Imagery of Solar System Bodies

In Table 6.3 of Chapter 6 describing the measurement campaign on GPS and GLONASS satellites, the error distribution statistics of astrometric reductions with and without stellar DCR compensation are compared showing a significant reduction in the mean and the spread of the error with the application of per-star DCR corrections in the data reduction. This motivates the re-reduction of existing astrometric observation of solar system bodies, independently of measurements of the spectra of the individual targets of those observations. The LINEAR program, SST, and the Catalina Sky Survey, to name a few large contributors to the Minor Planet Center, have retained a portion of their CCD images, and this thesis serves to motivate an effort to re-reduce those datasets using Gaia DR2 with stellar DCR correction.

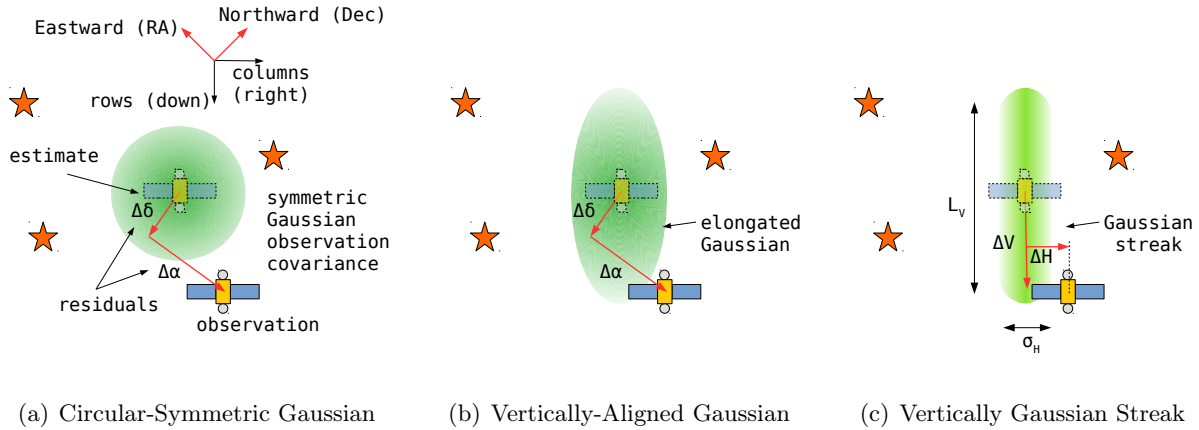


Figure 7-1: Error Distributions in Orbit Determination Routines

7.3.2.2 Color-Aware Farnocchia-Style Corrections

In the case where original imagery has not been retained, the evident improvement in systematic errors realized from the application of the Farnocchia catalog corrections to observations of long-arc asteroids with known spectral types motivates a similar approach. If the optical configuration of a sensor and meteorological conditions at the time of observations are known, then it is possible to guess which stars were used in the solution and to formulate color-aware corrections for that sensor to remove the effect of stellar DCR using Gaia DR2 photometry and parallax. Rather than a static correction to the observed RA and Dec, the correction would be along the station-local vertical and would be elevation-dependent.

7.3.2.3 Maximum-Likelihood Orbit Determination

It is reasonable to expect that even if stellar DCR compensation were to become a universally adopted practice, the large time-aperture requirement for spectroscopic DCR compensation will mean that it will still be necessary to handle observations with an unknown DCR bias. Given that DCR bias has been shown to be significant in certain applications, it is prudent to address how orbit solutions may be formed using observations with uncertain DCR.

Although this is not a universal statement, many operational and academic orbit determination software packages have a baked-in assumption that the error statistics of input optical observations are circularly-symmetric Gaussians. Some packages, such as GTDS and OrbFit, do not assume circular symmetry, but nevertheless assume that the observation covariance matrix is diagonal. That is, that the error distributions of observations supplied in sky-aligned RA/Dec coordinates

may be non-circular Gaussians, but if there is any elongation then it is aligned with the RA/Dec axes.

The Gaussian assumption is attractive because least squares differential correction is a maximum-likelihood estimator in the case where the observation noise is Gaussian. The circular-symmetric or diagonal assumption is there to save time or memory. The computation of the $\mathbf{W}^{\frac{1}{2}}$ weighting matrix in Equations 3.21 and 3.22 is faster if the matrix is diagonal. If all the weights are identical (i.e. circular symmetric Gaussians) then it factors out entirely from both equations.

Some of the operational software, like GTDS, has heritage back to the early days of operational SSA when computer time and memory were more precious than they are now. This motivates an examination of what would happen to orbit solutions if the underlying assumption were that the error distribution on optical observations were an elongated Gaussian aligned with the local vertical for each station. Such a distribution is shown schematically in Figure 7-1(b) and within a software routine formulated in terms of residuals in sky-aligned inertial angles $(\Delta\alpha, \Delta\delta)$ would have the form,

$$P(\Delta\alpha, \Delta\delta) = \frac{1}{2\pi\sigma_H\sigma_V} e^{-\frac{1}{2} \begin{bmatrix} \Delta\alpha & \Delta\delta \end{bmatrix} \mathbf{R}_{\text{rdo2eno}} \begin{bmatrix} \sigma_H^2 & 0 \\ 0 & \sigma_V^2 \end{bmatrix}^{-1} \mathbf{R}_{\text{rdo2eno}}^T \begin{bmatrix} \Delta\alpha \\ \Delta\delta \end{bmatrix}} \quad (7.1)$$

at a nominal RA/Dec of (α, δ)

From an information-theoretic perspective, the distribution that best matches the information content of an observation with a completely unknown, but bounded, amount of target DCR bias is uniform distribution along the vertical direction [145], where the extent of the vertical direction is the maximum plausible amount of DCR. The convolution of this distribution with the seeing-limited circularly-symmetric Gaussian distribution matches the information content of seeing-limited Gaussian horizontal distribution. This distribution is shown schematically in Figure 7-1(c) and its form is identical to that of the focal plane charge distribution for a streaking Gaussian source. In station-local horizontal/vertical coordinates $(\Delta H, \Delta V)$, the error distribution is

$$P(\Delta H, \Delta V) = \frac{1}{\sqrt{2\pi}2L_V\sigma_H} e^{-\frac{\Delta H^2}{\sigma_H^2}} \left(\text{erf} \left(\frac{\Delta V + L_V/2}{\sigma_H\sqrt{2}} \right) - \text{erf} \left(\frac{\Delta V - L_V/2}{\sigma_H\sqrt{2}} \right) \right) \quad (7.2)$$

This distribution is not usable in a least-squares formulation and an orbit determination routine that employs it would need to be re-written from the start to use gradient ascent within a maximum-likelihood framework. This may be a large undertaking or may not be feasible at all for software

where the least-squares formulation is deeply ingrained in the structure of the code.

Nevertheless the analysis of the effect of DCR bias on orbit determination of solar system bodies shows that for some cases, a better orbit solution is achieved by completely ignoring elevation and fitting to azimuth-only data. This is equivalent to an error distribution that is an infinitely long Gaussian streak in the vertical direction, and logically one would expect an improvement by constraining its maximum vertical extent.

References

- [1] A. V. Filippenko. The importance of atmospheric differential refraction in spectrophotometry. *PASP*, 94:715–721, Aug 1982.
<https://doi.org/10.1086/131052>.
- [2] L. H. Auer and E. M. Standish. Astronomical Refraction: Computational Method for All Zenith Angles. *AJ*, 119:2472–2474, May 2000.
<https://doi.org/10.1086/301325>.
- [3] W. J. Faccenda. GEODSS: Past and future upgrades. Technical report, MITRE Corporation, 2000.
- [4] W. J. Faccenda, C. M. Williams, and D. Brisnehan. Deep Stare technical advancements and status. Technical report, MITRE Corporation, 2003.
- [5] G. H. Stokes, J. B. Evans, H. E. M. Viggh, F. C. Shelly, and E. C. Pearce. Lincoln Near-Earth Asteroid Program (LINEAR). *Icarus*, 148:21–28, Nov 2000.
<https://doi.org/10.1006/icar.2000.6493>.
- [6] D. F. Woods, R. Shah, J. A. Johnson, et al. Space Surveillance Telescope: focus and alignment of a three mirror telescope. *Optical Engineering*, 52(5):053604, 2013.
<https://doi.org/10.1117/1.OE.52.5.053604>.
- [7] E. Pearce, F. Shelly, and J. Johnson. High precision real time metric processing for the MOSS and LINEAR systems. In *Proceedings of the 2000 Space Control Conference*, pages 145–157.
<https://apps.dtic.mil/docs/citations/ADA377625>.
- [8] MIT Lincoln Laboratory. *Annual Report*. 2010.
<https://apps.dtic.mil/docs/citations/ADA596371>.
- [9] D. G. Monet, S. E. Levine, B. Canzian, et al. The USNO-B catalog. *AJ*, 125(2):984–993, 2003.
<https://doi.org/10.1086/345888>.
- [10] C. T. Finch, N. Zacharias, and G. L. Wycoff. UCAC3: Astrometric reductions. *AJ*, 139(6):2200–2207, June 2010.
<https://doi.org/10.1088/0004-6256/139/6/2200>.
- [11] N. Zacharias, C. Finch, J. Subasavage, et al. The first U.S. Naval Observatory robotic astrometric telescope catalog. *AJ*, 150(4):101, 2015.
<https://doi.org/10.1088/0004-6256/150/4/101>.

- [12] R. P. Binzel, A. S. Rivkin, J. Stuart, et al. Observed spectral properties of near-earth objects: results for population distribution, source regions, and space weathering processes. *Icarus*, 170(2):259 – 294, 2004.
<https://doi.org/10.1016/j.icarus.2004.04.004>.
- [13] D. Monet. The general release of USNO-A 1.0. In *Bulletin of the American Astronomical Society*, volume 28 of *BAAS*, page 1282, Dec 1996.
<https://ui.adsabs.harvard.edu/abs/1996BAAS...28.1282M>.
- [14] D. Farnocchia, S. Chesley, A. Chamberlin, and D. Tholen. Star catalog position and proper motion corrections in asteroid astrometry. *Icarus*, 245:94 – 111, 2015.
<https://doi.org/10.1016/j.icarus.2014.07.033>.
- [15] JPL Horizons.
<https://ssd.jpl.nasa.gov/horizons.cgi>.
- [16] IAU Minor Planet Center.
<https://www.minorplanetcenter.net>.
- [17] National Research Council. *Continuing Kepler’s Quest: Assessing Air Force Space Command’s Astrodynamics Standards*. National Academies Press, 2012. ISBN 9780309261432.
- [18] M. Shao, B. Nemati, C. Zhai, et al. Finding very small near-earth asteroids using synthetic tracking. *ApJ*, 782(1):1, 2014.
<https://doi.org/10.1088/0004-637X/782/1/1>.
- [19] C. Zhai, M. Shao, B. Nemati, et al. Detection of a faint fast-moving near-earth asteroid using the synthetic tracking technique. *ApJ*, 792(1):60, 2014.
<https://doi.org/10.1088/0004-637X/792/1/60>.
- [20] R. Schweickart, C. Chapman, D. Durda, et al. Threat Characterization: Trajectory Dynamics. *ArXiv Physics e-prints*, Aug 2006.
<https://arxiv.org/abs/physics/0608155>.
- [21] D. Farnocchia, S. R. Chesley, P. G. Brown, and P. W. Chodas. The trajectory and atmospheric impact of asteroid 2014 AA. *Icarus*, 274:327 – 333, 2016.
<https://doi.org/10.1016/j.icarus.2016.02.056>.
- [22] D. Farnocchia, S. R. Chesley, D. J. Tholen, and M. Micheli. High precision predictions for near-Earth asteroids: the strange case of (3908) Nyx. *Celestial Mechanics and Dynamical Astronomy*, 119:301–312, Aug 2014.
<https://doi.org/10.1007/s10569-014-9536-9>.
- [23] D. Farnocchia, S. Chesley, D. Vokrouhlický, et al. Near earth asteroids with measurable Yarkovsky effect. *Icarus*, 224(1):1 – 13, 2013.
<https://doi.org/10.1016/j.icarus.2013.02.004>.
- [24] D. Jewitt and J. Luu. Discovery of the candidate Kuiper belt object 1992 QB1. *Nature*, 362: 730–732, Apr 1993.
<https://doi.org/10.1038/362730a0>.

- [25] J. L. Elliot, E. W. Dunham, and C. B. Olkin. Exploring small bodies in the outer solar system with stellar occultations. In M. R. Haas, J. A. Davidson, and E. F. Erickson, editors, *From Gas to Stars to Dust*, volume 73 of *Astronomical Society of the Pacific Conference Series*, 1995.
<https://ui.adsabs.harvard.edu/abs/1995ASPC...73..285E>.
- [26] D. J. Trump. Space Policy Directive 3: National Space Traffic Management Policy, 2018.
- [27] P. M. Cunio, M. Bantel, B. R. Flewelling, et al. Photometric and other analyses of energetic events related to 2017 GEO RSO anomalies. In *Proceedings of the AMOS Technical Conference*, 2017.
<https://ui.adsabs.harvard.edu/abs/2017amos.confE..75C>.
- [28] T. Schildknecht, R. Musci, and T. Flohrer. Properties of the high area-to-mass ratio space debris population at high altitudes. *Advances in Space Research*, 41(7):1039 – 1045, 2008. ISSN 0273-1177.
<https://doi.org/10.1016/j.asr.2007.01.045>.
- [29] T. Kelecy and M. Jah. Analysis of high area-to-mass ratio (hamr) geo space object orbit determination and prediction performance: Initial strategies to recover and predict hamr geo trajectories with no a priori information. *Acta Astronautica*, 69(7):551 – 558, 2011. ISSN 0094-5765.
<https://doi.org/10.1016/j.actaastro.2011.04.019>.
- [30] Amon-Hen Program. IARPA-BAA-17-02, June 2017.
<https://www.iarpa.gov/index.php/research-programs/amon-hen/amon-hen-baa>.
- [31] R. Teoste, J. Daley, R. N. Capes, J. J. Alves, and M. D. Zimmerman. Measurements of tilt anisoplanatism at the firepond facility. Technical Report 815, MIT Lincoln Laboratory, 1988.
<https://apps.dtic.mil/docs/citations/ADA202681>.
- [32] F. Schlesinger. Photographic Determinations of Stellar Parallax Made with the Yerkes Refractor. *ApJ*, 32:372, Dec 1910.
<https://doi.org/10.1086/141819>.
- [33] F. Schlesinger. Photographic Determinations of Stellar Parallax Made with the Yerkes Refractor II. *ApJ*, 33:8, Jan 1911.
<https://doi.org/10.1086/141826>.
- [34] F. Schlesinger. Photographic determinations of stellar parallax made with the Yerkes refractor. III. *ApJ*, 33, Mar 1911.
<https://doi.org/10.1086/141844>.
- [35] E. W. Rork. Private communication, 2007.
- [36] K. G. Henize. The Baker-Nunn Satellite-Tracking Camera. *S&T*, 16, Jan 1957.
<https://ui.adsabs.harvard.edu/abs/1957S%26T...16..108H>.
- [37] California Institute of Technology. The 48-inch (1.2-meter) Samuel Oschin Telescope, Accessed 2019-05-28.
<http://www.astro.caltech.edu/palomar/about/telescopes/oschin.html>.

- [38] W. F. van Altena and E. B. Folamont. The Earth’s atmosphere: refraction, turbulence, delays and limitations to astrometric precision. In W. F. van Altena, editor, *Astrometry for Astrophysics*. Cambridge University Press, 2013.
<https://ui.adsabs.harvard.edu/abs/2013asas.book.....V>.
- [39] N. Zacharias, S. E. Urban, M. I. Zacharias, et al. The first US Naval Observatory CCD astrograph catalog. *AJ*, 120(4):2131–2147, Oct 2000.
<https://doi.org/10.1086/301563>.
- [40] M. Carpino, A. Milani, and S. R. Chesley. Error statistics of asteroid optical astrometric observations. *Icarus*, 166(2):248 – 270, 2003.
[https://doi.org/10.1016/S0019-1035\(03\)00051-4](https://doi.org/10.1016/S0019-1035(03)00051-4).
- [41] N. Zacharias, C. T. Finch, T. M. Girard, et al. The Fourth US Naval Observatory CCD Astrograph Catalog (UCAC4). *AJ*, 145:44, Feb 2013.
<https://doi.org/10.1088/0004-6256/145/2/44>.
- [42] The HIPPARCOS and TYCHO catalogues. Astrometric and photometric star catalogues derived from the ESA HIPPARCOS Space Astrometry Mission. In *ESA Special Publication*, volume 1200, Jan 1997.
<https://ui.adsabs.harvard.edu/abs/1997ESASP1200.....E>.
- [43] L. Lindegren, U. Lammers, U. Bastian, et al. Gaia Data Release 1. Astrometry: one billion positions, two million proper motions and parallaxes. *A&A*, 595:A4, Nov 2016.
<https://doi.org/10.1051/0004-6361/201628714>.
- [44] Gaia Collaboration, A. G. A. Brown, A. Vallenari, et al. Gaia Data Release 2. Summary of the contents and survey properties. *A&A*, 616:A1, Aug 2018.
<https://doi.org/10.1051/0004-6361/201833051>.
- [45] N. Zacharias, C. Finch, T. Girard, et al. The Third US Naval Observatory CCD Astrograph Catalog (UCAC3). *AJ*, 139:2184–2199, Jun 2010.
<https://doi.org/10.1088/0004-6256/139/6/2184>.
- [46] N. Zacharias, C. Finch, and J. Frouard. UCAC5: New proper motions using Gaia DR1. *AJ*, 153(4):166, 2017.
<https://doi.org/10.3847/1538-3881/aa6196>.
- [47] C. R. Chapman and M. J. Gaffey. Reflectance spectra for 277 asteroids. In T. Gehrels, editor, *Asteroids*, pages 655–687. University of Arizona Press, 1979.
<https://ui.adsabs.harvard.edu/abs/1979aste.book..655C>.
- [48] D. J. Tholen and M. A. Barucci. Asteroid taxonomy. In R. P. Binzel, T. Gehrels, and M. S. Matthews, editors, *Asteroids II*, pages 298–315. University of Arizona Press, 1989.
<https://ui.adsabs.harvard.edu/abs/1989aste.conf..298T>.
- [49] S. J. Bus and R. P. Binzel. Phase II of the Small Main-Belt Asteroid Spectroscopic Survey. A Feature-Based Taxonomy. *Icarus*, 158:146–177, Jul 2002.
<https://doi.org/10.1006/icar.2002.6856>.
- [50] F. E. DeMeo, R. P. Binzel, S. M. Slivan, and S. J. Bus. An extension of the Bus asteroid taxonomy into the near-infrared. *Icarus*, 202:160–180, Jul 2009.
<https://doi.org/10.1016/j.icarus.2009.02.005>.

- [51] C. A. Thomas, J. P. Emery, D. E. Trilling, et al. Physical characterization of warm Spitzer-observed near-Earth objects. *Icarus*, 228(Supplement C):217 – 246, 2014.
<https://doi.org/10.1016/j.icarus.2013.10.004>.
- [52] J. A. Sanchez, V. Reddy, A. Nathues, et al. Phase reddening on near-Earth asteroids: Implications for mineralogical analysis, space weathering and taxonomic classification. *Icarus*, 220(1):36 – 50, 2012.
<https://doi.org/10.1016/j.icarus.2012.04.008>.
- [53] J. R. Masiero, A. K. Mainzer, T. Grav, et al. Main Belt Asteroids with WISE/NEOWISE. I. Preliminary Albedos and Diameters. *ApJ*, 741:68, Nov 2011.
<https://doi.org/10.1088/0004-637X/741/2/68>.
- [54] B. D. Warner, A. W. Harris, and P. Pravec. The asteroid lightcurve database. *Icarus*, 202:134–146, Jul 2009.
<https://doi.org/10.1016/j.icarus.2009.02.003>.
- [55] S. Marchi, M. Lazzarin, S. Magrin, and C. Barbieri. Visible spectroscopy of the two largest known trans-Neptunian objects: Ixion and Quaoar. *A&A*, 408(3):L17–L19, 2003.
<https://doi.org/10.1051/0004-6361:20031142>.
- [56] D. P. Cruikshank and C. M. Dalle Ore. Spectral Models of Kuiper Belt Objects and Centaurs. In *The First Decadal Review of the Edgeworth-Kuiper Belt*, pages 315–330. Springer, 2004.
<https://doi.org/10.1023/B:M00N.0000031948.39136.7d>.
- [57] D. C. Jewitt and J. Luu. Crystalline water ice on the Kuiper belt object (50000) Quaoar. *Nature*, 432(7018):731, 2004.
<https://doi.org/10.1038/nature03111>.
- [58] M. A. Barucci, M. Fulchignoni, M. Birlan, et al. Analysis of Trans-Neptunian and Centaur colours: continuous trend or grouping? *A&A*, 371:1150–1154, Jun 2001.
<https://doi.org/10.1051/0004-6361:20010501>.
- [59] K. K. Luu, C. L. Matson, J. Snodgrass, S. M. Giffin, and K. Hamada. Object characterization from spectra data. In *Proceedings of the AMOS Technical Conference*, 2003.
- [60] R. H. French, J. Rodriguez-Parada, M. Yang, et al. Optical properties of materials for concentrator photovoltaic systems. In *Photovoltaic Specialists Conference (PVSC), 2009 34th IEEE*, pages 000394–000399. IEEE, 2009.
<https://doi.org/10.1109/PVSC.2009.5411657>.
- [61] D. L. Bowers, L. D. Wellems, M. J. Duggin, W. Glass, and L. G. Vaughan. Broadband spectral-polarimetric BRDF scan system and data for spacecraft materials. Technical report, Air Force Research Laboratory, Kirtland AFB, NM, Space Vehicles Directorate, 2011.
<https://apps.dtic.mil/docs/citations/ADA550687>.
- [62] K. Jorgensen, J. Okada, L. Bradford, et al. Obtaining material type of orbiting objects through reflectance spectroscopy measurements. In *Proceedings of the AMOS Technical Conference*, 2003.

- [63] K. Abercromby, J. Okada, M. Guyote, K. Hamada, and E. Barker. Comparisons of ground truth and remote spectral measurements of the FORMOSAT and ANDE spacecraft. In *Proceedings of the AMOS Technical Conference*, 2006.
<https://ui.adsabs.harvard.edu/abs/2006amos.confE..79A>.
- [64] K. Abercromby, B. Buckalew, P. Abell, and H. Cowardin. Infrared telescope facility’s spectrograph observations of human-made space objects. In *Proceedings of the AMOS Technical Conference*, 2015.
<https://ui.adsabs.harvard.edu/abs/2015amos.confE..19B>.
- [65] A. Vananti, T. Schildknecht, and H. Krag. Reflectance spectroscopy characterization of space debris. *Advances in Space Research*, 59(10):2488–2500, 2017.
<https://doi.org/10.1016/j.asr.2017.02.033>.
- [66] D. Monet. Private communication, July 2017.
- [67] S. Urban and P. Seidelmann. *Explanatory Supplement to the Astronomical Almanac*. University Science Books, 2013. ISBN 9781891389856.
- [68] B. E. Schaefer and W. Liller. Refraction near the horizon. *PASP*, 102:796–805, Jul 1990.
<https://doi.org/10.1086/132705>.
- [69] A. W. Harris. Physical properties of Neptune and Triton inferred from the orbit of Triton. In J. T. Bergstrahl, editor, *NASA Conference Publication*, volume 2330, pages 357–373, Oct 1984.
<https://ui.adsabs.harvard.edu/abs/1984NASCP2330..357H>.
- [70] A.W. Harris. Private communication, 2018.
- [71] R. C. Stone. The effect of differential color refraction on declinations determined in meridian circle programs. *A&A*, 138:275–284, Sep 1984.
<https://ui.adsabs.harvard.edu/abs/1984A&A...138..275S>.
- [72] J. Gubler and D. Tytler. Differential Atmospheric Refraction and Limitations on the Relative Astrometric Accuracy of Large Telescopes. *PASP*, 110:738–746, Jun 1998.
<https://doi.org/10.1086/316172>.
- [73] M. Altmann, S. Bouquillon, F. Taris, et al. GBOT: ground based optical tracking of the Gaia satellite. In *Observatory Operations: Strategies, Processes, and Systems V*, volume 9149 of *Proc. SPIE*, page 91490P, Aug 2014.
<https://doi.org/10.1117/12.2056571>.
- [74] C. Briceño, S. Heathcote, G. Cecil, et al. The SOAR telescope atmospheric dispersion corrector. In *Proc.SPIE*, volume 10700, page 107003Z, 2018.
<https://doi.org/10.1117/12.2314026>.
- [75] A. J. Drake, S. G. Djorgovski, A. Mahabal, et al. The Catalina Real-time Transient Survey. In E. Griffin, R. Hanisch, and R. Seaman, editors, *New Horizons in Time Domain Astronomy*, volume 285 of *IAU Symposium*, pages 306–308, Apr 2012.
<https://doi.org/10.1017/S1743921312000889>.

- [76] J. Ruprecht, G. Ushomirsky, D. Woods, et al. Asteroid Detection Results Using the Space Surveillance Telescope. In *Proceedings of the AMOS Technical Conference*, page 92, 2015. <https://ui.adsabs.harvard.edu/abs/2015amos.confE..92R>.
- [77] W. M. Folkner, J. G. Williams, D. H. Boggs, R. S. Park, and P. Kuchynka. The Planetary and Lunar Ephemerides DE430 and DE431. *Interplanetary Network Progress Report*, 42-196: 1–81, Feb 2014. <https://ui.adsabs.harvard.edu/abs/2014IPNPR.196C...1F>.
- [78] N. K. Pavlis, S. A. Holmes, S. C. Kenyon, and J. K. Factor. An Earth gravitational model to degree 2160: EGM2008. *EGU General Assembly*, 2008(4):4–2, 2008.
- [79] G. Petit and B. Luzum. IERS Conventions (2010). *IERS Technical Note*, 36:1, Jan 2010. <https://ui.adsabs.harvard.edu/abs/2010ITN....36....1P>.
- [80] O. Montenbruck and E. Gill. *Satellite Orbits: Models, Methods, and Applications*. Physics and astronomy online library. Springer Berlin Heidelberg, 2000. ISBN 9783540672807. <https://doi.org/10.1007/978-3-642-58351-3>.
- [81] J. M. Dow, R. E. Neilan, and C. Rizos. The International GNSS Service in a changing landscape of Global Navigation Satellite Systems. *Journal of Geodesy*, 83(3):191–198, 2009. <https://doi.org/10.1007/s00190-008-0300-3>.
- [82] O. L. Colombo. The dynamics of global positioning system orbits and the determination of precise ephemerides. *Journal of Geophysical Research: Solid Earth*, 94(B7):9167–9182, 1989. <https://doi.org/10.1029/JB094iB07p09167>.
- [83] K. Sośnica. LAGEOS sensitivity to ocean tides. *Acta Geophysica*, 63(4):1181–1203, Aug 2015. <https://doi.org/10.1515/acgeo-2015-0032>.
- [84] J. D. Giorgini, D. K. Yeomans, A. B. Chamberlin, et al. JPL’s On-Line Solar System Data Service. In *AAS/Division for Planetary Sciences Meeting Abstracts #28*, volume 28 of *Bulletin of the American Astronomical Society*, page 1158, Sep 1996. <https://ui.adsabs.harvard.edu/abs/1996DPS....28.2504G>.
- [85] C. H. Acton Jr. Ancillary data services of NASA’s navigation and ancillary information facility. *Planetary and Space Science*, 44(1):65–70, 1996. <https://ui.adsabs.harvard.edu/abs/1996P&SS...44...65A>.
- [86] C. Acton, N. Bachman, B. Semenov, and E. Wright. A look towards the future in the handling of space science mission geometry. *Planetary and Space Science*, 150:9–12, 2018. <https://ui.adsabs.harvard.edu/abs/2018P&SS..150...9A>.
- [87] M. Schenewerk. A brief review of basic GPS orbit interpolation strategies. *GPS solutions*, 6(4):265–267, 2003. <https://doi.org/10.1007/s10291-002-0036-0>.
- [88] L. Smullin and G. Fiocco. Project luna see. *Proceedings of the Institute of Radio Engineers*, 50(7):1703, 1962.
- [89] S. F. Catalano. The Firepond Kalman filter. Technical report, MIT Lincoln Laboratory, 1977. <https://apps.dtic.mil/docs/citations/ADA040131>.

- [90] A. B. Gschwendtner and W. E. Keicher. Development of coherent laser radar at Lincoln Laboratory. *Lincoln Laboratory Journal*, 12(2):383–396, 2000.
- [91] R. Tyson. Innovative wavefront estimators for zonal adaptive optics systems. In *Adaptive Optics Engineering Handbook*, pages 203–232. CRC Press, 1999. ISBN 9780824782757.
- [92] Z. Folcik. Private communication, 2016.
- [93] IAU SOFA Board, IAU SOFA Software Collection.
<http://www.iausofa.org>.
- [94] J. Bangert, W. Puatua, G. Kaplan, et al. User’s Guide to NOVAS Version C3.1, 2011.
<https://ui.adsabs.harvard.edu/abs/2012ascl.soft02003K>.
- [95] T. Moyer. *Formulation for Observed and Computed Values of Deep Space Network Data Types for Navigation*. JPL Deep-Space Communications and Navigation Series. Wiley, 2005. ISBN 9780471726173.
<https://doi.org/10.1002/0471728470>.
- [96] M. Galassi, J. Davies, J. Theiler, et al. GNU Scientific Library - Reference Manual.
<https://www.gnu.org/software/gsl/>.
- [97] G. D. Byrne and A. C. Hindmarsh. A polyalgorithm for the numerical solution of ordinary differential equations. *ACM Trans. Math. Softw.*, 1(1):71–96, Mar 1975.
<https://doi.org/10.1145/355626.355636>.
- [98] P. N. Brown, G. D. Byrne, and A. C. Hindmarsh. VODE: A variable-coefficient ODE solver. *SIAM Journal on Scientific and Statistical Computing*, 10(5):1038–1051, 1989.
<https://doi.org/10.1137/0910062>.
- [99] J. Kabelac. Atmospheric models and astronomical and parallactic refraction. *Studia Geophysica et Geodaetica*, 11:1–22, 1967.
<http://www.dtic.mil/docs/citations/AD0850231>.
- [100] P. C. Hansen, V. Pereyra, and G. Scherer. *Least squares data fitting with applications*. JHU Press, 2013.
<https://doi.org/10.1353/book.21076>.
- [101] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. OUP Oxford, 2001. ISBN 9780198572220.
- [102] M. Pearlman, J. Degnan, and J. Bosworth. The international laser ranging service. *Advances in Space Research*, 30(2):135 – 143, 2002.
[https://doi.org/10.1016/S0273-1177\(02\)00277-6](https://doi.org/10.1016/S0273-1177(02)00277-6).
- [103] J. D. Giorgini. Private communication, 2017.
- [104] F. Lyard, F. Lefevre, T. Letellier, and O. Francis. Modelling the global ocean tides: modern insights from fes2004. *Ocean Dynamics*, 56(5):394–415, Dec 2006. ISSN 1616-7228.
<https://doi.org/10.1007/s10236-006-0086-x>.

- [105] P. Knocke, J. C. Ries, and B. D. Tapley. Earth radiation pressure effects on satellites. In *Proceedings of the AIAA-/AAS Astrodynamics Conference, Inst. of Aeronaut. Astronaut., Washington, D.C.*, 1988.
<https://doi.org/10.2514/6.1988-4292>.
- [106] S. B. Luthcke, N. P. Zelensky, D. D. Rowlands, F. G. Lemoine, and T. A. Williams. The 1-centimeter orbit: Jason-1 precision orbit determination using GPS, SLR, DORIS, and altimeter data special issue: Jason-1 calibration/validation. *Marine Geodesy*, 26(3-4):399–421, 2003.
<https://doi.org/10.1080/714044529>.
- [107] M. Ziebart, S. Edwards, S. Adhya, et al. High precision GPS IIR orbit prediction using analytical non-conservative force models. In *ION GNSS: Long Beach, California, USA.*, pages 1764–1770, 2004.
- [108] E. Bertin and S. Arnouts. SExtractor: Software for source extraction. *A&AS*, 117:393–404, Jun 1996.
<https://doi.org/10.1051/aas:1996164>.
- [109] V. Kouprianov. Distinguishing features of CCD astrometry of faint GEO objects. *Advances in Space Research*, 41:1029–1038, 2008.
<https://doi.org/10.1016/j.asr.2007.04.033>.
- [110] D. B. Owen. A table of normal integrals. *Communications in Statistics-Simulation and Computation*, 9(4):389–419, 1980.
<https://doi.org/10.1080/03610918008812164>.
- [111] F. G. Valdes, L. E. Campusano, J. D. Velasquez, and P. B. Stetson. FOCAS Automatic Catalog Matching Algorithms. *PASP*, 107:1119, Nov 1995.
<https://doi.org/10.1086/133667>.
- [112] J. S. Heyl. k-d Match: A Fast Matching Algorithm for Sheared Stellar Samples. *arXiv e-prints*, art. arXiv:1304.0838, Apr 2013.
<https://arxiv.org/abs/1304.0838>.
- [113] J. A. Cardelli, G. C. Clayton, and J. S. Mathis. The relationship between infrared, optical, and ultraviolet extinction. *ApJ*, 345:245–256, Oct 1989.
<https://doi.org/10.1086/167900>.
- [114] A. J. Pickles. A Stellar Spectral Flux Library: 1150-25000 Å. *PASP*, 110:863–878, Jul 1998.
<https://doi.org/10.1086/316197>.
- [115] M. Sekiguchi and M. Fukugita. A Study of the B-V Color-Temperature Relation. *AJ*, 120:1072–1084, Aug 2000.
<https://doi.org/10.1086/301490>.
- [116] A. N. Cox. *Allen's Astrophysical Quantities*. Springer, 2000.
<https://ui.adsabs.harvard.edu/abs/2000asqu.book.....C>.
- [117] L. He, D. C. B. Whittet, D. Kilkenny, and J. H. Spencer Jones. Interstellar Extinction from 0.35 to 2.2 Microns: A Study Based on Luminous Southern Stars. *ApJS*, 101:335, Dec 1995.
<https://doi.org/10.1086/192243>.

- [118] D. Whittet. *Dust in the Galactic Environment, 2nd Edition*. Series in Astronomy and Astrophysics. CRC Press, 2002. ISBN 9780750306249.
<https://doi.org/10.1201/9781315274645>.
- [119] G. Catanzaro. How to determine surface gravity from stellar spectra. In E. Niemczura, B. Smalley, and W. Pych, editors, *Determination of Atmospheric Parameters of B-, A-, F- and G-Type Stars: Lectures from the School of Spectroscopic Data Analyses*, pages 97–109. Springer International Publishing, Cham, 2014. ISBN 978-3-319-06956-2.
https://doi.org/10.1007/978-3-319-06956-2_9.
- [120] B. Smalley. Stellar parameters from photometry. In E. Niemczura, B. Smalley, and W. Pych, editors, *Determination of Atmospheric Parameters of B-, A-, F- and G-Type Stars: Lectures from the School of Spectroscopic Data Analyses*, pages 111–120. Springer International Publishing, Cham, 2014. ISBN 978-3-319-06956-2.
https://doi.org/10.1007/978-3-319-06956-2_10.
- [121] Gaia data release scenario.
<https://www.cosmos.esa.int/web/gaia/release>.
- [122] R. Andrae, M. Fouesneau, O. Creevey, et al. Gaia Data Release 2. First stellar parameters from Apsis. *A&A*, 616:A8, Aug 2018.
<https://doi.org/10.1051/0004-6361/201732516>.
- [123] A. Berk, G. P. Anderson, P. K. Acharya, et al. MODTRAN5: a reformulated atmospheric band model with auxiliary species and practical multiple scattering options. In S. S. Shen and P. E. Lewis, editors, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery X*, volume 5425 of *Proc. SPIE*, pages 341–347, Aug 2004.
<https://doi.org/10.1117/12.546782>.
- [124] M. S. Bessell. UBVRI passbands. *PASP*, 102:1181–1199, Oct 1990.
<https://doi.org/10.1086/132749>.
- [125] P. R. Jorden, D. Jordan, P. A. Jerram, J. Pratlong, and I. Swindells. e2v new CCD and CMOS technology developments for astronomical sensors. In *High Energy, Optical, and Infrared Detectors for Astronomy VI*, volume 9154, page 91540M. International Society for Optics and Photonics, 2014.
<https://doi.org/10.1117/12.2069423>.
- [126] D. Farnocchia, S. R. Chesley, A. Milani, G. F. Gronchi, and P. W. Chodas. Orbits, Long-Term Predictions, Impact Monitoring. In P. Michel, F. E. DeMeo, and W. F. Bottke, editors, *Asteroids IV*, pages 815–834. University of Arizona Press, 2015.
https://doi.org/10.2458/azu_uapress_9780816532131-ch041.
- [127] B. Sesar, J. S. Stuart, Ž. Ivezić, et al. Exploring the Variable Sky with LINEAR. I. Photometric Recalibration with the Sloan Digital Sky Survey. *AJ*, 142:190, Dec 2011.
<https://doi.org/10.1088/0004-6256/142/6/190>.
- [128] M. Fukugita, T. Ichikawa, J. E. Gunn, et al. The Sloan Digital Sky Survey Photometric System. *AJ*, 111:1748, Apr 1996.
<https://doi.org/10.1086/117915>.

- [129] F. Vilas and M. J. Gaffey. Phyllosilicate absorption features in main-belt and outer-belt asteroid reflectance spectra. *Science*, 246:790–792, Nov 1989.
<https://doi.org/10.1126/science.246.4931.790>.
- [130] T. Burbine. *Asteroids: Astronomical and Geological Bodies*. Cambridge University Press, 2016. ISBN 9781316867396.
<https://doi.org/10.1017/9781316156582>.
- [131] I. Molotov, V. Agapov, V. Titenko, et al. International scientific optical network for space debris research. *Advances in Space Research*, 41(7):1022–1028, 2008.
<https://doi.org/10.1016/j.asr.2007.04.048>.
- [132] D. Vokrouhlický, W. F. Bottke, S. R. Chesley, D. J. Scheeres, and T. S. Statler. The Yarkovsky and YORP Effects. In P. Michel, F. E. DeMeo, and W. F. Bottke, editors, *Asteroids IV*, pages 509–531. University of Arizona Press, 2015.
https://doi.org/10.2458/azu_uapress_9780816532131-ch027.
- [133] Orbfite Consortium. OrbFit: Software to Determine Orbits of Asteroids. Astrophysics Source Code Library, Jun 2011.
<https://ui.adsabs.harvard.edu/abs/2011ascl.soft060150>.
- [134] K. G. Libbrecht and M. L. Peri. A fiber-fed echelle spectrograph for the Hale 5-m telescope. *Publications of the Astronomical Society of the Pacific*, 107:62, Jan 1995.
<https://doi.org/10.1086/133516>.
- [135] J. D. Eastman, T. M. Brown, J. Hygelund, et al. NRES: the network of robotic echelle spectrographs. In *Proc.SPIE*, volume 9147, 2014.
<https://doi.org/10.1117/12.2054699>.
- [136] H. Sugai, N. Tamura, H. Karoji, et al. Prime focus spectrograph for the Subaru telescope: massively multiplexed optical and near-infrared fiber spectrograph. *Journal of Astronomical Telescopes, Instruments, and Systems*, 1(3):1 – 11 – 11, 2015.
<https://doi.org/10.1117/1.JATIS.1.3.035001>.
- [137] H. Jiang, Z. Hu, M. Xu, et al. The preliminary design of the next generation Palomar spectrograph for 200-inch Hale telescope. In *Proc.SPIE*, volume 10702, 2018.
<https://doi.org/10.1117/12.2312550>.
- [138] J. Goodman. *Introduction to Fourier Optics*. W. H. Freeman, 2005. ISBN 9780974707723.
- [139] M. Born and E. Wolf. *Principles of Optics*. Cambridge University Press, Oct 1999. ISBN 0521642221.
<https://ui.adsabs.harvard.edu/abs/1999prop.book.....B>.
- [140] J. A. Munn, H. C. Harris, T. von Hippel, et al. A Deep Proper Motion Catalog Within the Sloan Digital Sky Survey Footprint. *AJ*, 148:132, Dec 2014.
<https://doi.org/10.1088/0004-6256/148/6/132>.
- [141] P. Wallace. Pointing and tracking algorithms for the Keck 10-meter telescope. In L. B. Robinson, editor, *Instrumentation for Ground-Based Optical Astronomy*, pages 691–706, New York, NY, 1988. Springer New York. ISBN 978-1-4612-3880-5.
https://doi.org/10.1007/978-1-4612-3880-5_72.

- [142] A. Pickles and É. Depagne. All-Sky Spectrally Matched UBVRI - ZY and u' g' r' i' z' Magnitudes for Stars in the Tycho2 Catalog. *PASP*, 122:1437, Dec 2010.
<https://doi.org/10.1086/657947>.
- [143] F. Ochsenbein, P. Bauer, and J. Marcout. The VizieR database of astronomical catalogues. *A&AS*, 143:23–32, Apr 2000.
<https://doi.org/10.1051/aas:2000169>.
- [144] Air Force Space Command. Global positioning system fact sheet.
<https://www.af.mil/About-Us/Fact-Sheets/Display/Article/104610/global-positioning-system>.
- [145] T. M. Cover and J. A. Thomas. *Elements of Information Theory, 2nd Edition*. John Wiley & Sons, 2006.