Characterization and Mitigation of GNSS Carrier Phase Cycle Slips

by

B. W. Breitsch

B.S., Miami University, 2014

M.S., Colorado State University, 2017

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Aerospace Engineering and Sciences

2021

Committee Members:

Prof. Y. Jade Morton

Dr. Gary McGraw

Prof. Penina Axelrad

Prof. Tomoko Matsuo

Dr. Attilla Komjathy

Breitsch, B. W. (Ph.D., Aerospace)

Characterization and Mitigation of GNSS Carrier Phase Cycle Slips

Thesis directed by Prof. Y. Jade Morton

Originally designed for navigation, signals from Global Navigation Satellite Systems (GNSS) are now being used for remote-sensing of Earth's ionosphere, atmospheres, and surfaces. Since their earliest application, the carrier phase of these signals have been afflicted by cycle slips, which manifest as rapid and discrete changes in the phase measurement bias. They can occur due to the effects of noise, multipath interference, receiver processing, or a combination of all these factors. In both navigation and remote-sensing contexts, GNSS signals are subject to a variety of harsh conditions, such as propagation through ionospheric and atmospheric structures, arrival from low elevation angles, or reflection off Earth's surfaces. Under such conditions, cycle slips become a major error source for remote sensing and high accuracy navigation algorithms relying on continuous phase measurements, and therefore must be mitigated.

This work addresses the characterization and mitigation of GNSS carrier cycle slips that occur under these harsh conditions. In particular, we examine the origins of cycle slips due to the effects of both noise and phase transitions, which are slip-inducing phase fluctuations related to signal multipath interference. We discuss how phase transitions are related to canonical fading, which is where a signal exhibits a fade in amplitude coinciding with a rapid half-cycle phase change. Using phase screen simulations of ionosphere diffraction, we are able to characterize the cycle slips rate-of-occurrence during ionosphere scintillation. We also show how many slips can occur due to phase transitions where the signal fading is shallow, which has implications for algorithms that try to make use the signal amplitude measurements when mitigating cycle slips. We then assess actual phase transitions and cycle slips in real multi-frequency ionosphere scintillation data. We discuss two different approaches to cycle slip mitigation that have been successful in the past, but which fail to adequately address the slips in the ionosphere scintillation datasets.

Our assessment of cycle slips in real scintillation data motivates our development of a general probabilistic model for assessing cycle slip occurrences given an arbitrary set of GNSS carrier phase measurements. In our model, we use Gaussian processes to represent signal phase components and discuss tuning of the covariance parameters to deal with different levels of uncertainty or variation in these components. We also discuss the roles of thermal noise and noise due to unmodeled errors such as ionosphere diffraction. We consider the model performance for estimating a single cycle slip under a variety of hypothetical conditions, and show that generally we require an extended window ($\geq$16 seconds) of high-rate ($\geq$20 Hz) measurements in order to reliably estimate cycle slip occurrences. However, as we show in our assessment and characterization of real cycle slip occurrences, several slips are prone to occur over such a window under harsh signal conditions. Therefore, cycle slip amplitude estimates for a given window of harsh signal carrier phase measurements will be highly interdependent. To address this, we develop a batch cycle slip detection and estimation method that can reliably estimate cycle slips under harsh conditions. Our approach makes use of a variety of novel techniques including sparse estimation of slip occurrences and an adaptation of the search-and-shrink algorithm traditionally used to find the solutions of integer-least-squares problems. We assess the algorithm performance on simulated and real datasets. We demonstrate its effectiveness when applied during ionosphere scintillation, weak ocean reflections, or radio occultations through the lower troposphere, and we show that it can work with triple-, dual-, and single-frequency signal measurements.

Dedication

To my parents and family, with love.

Acknowledgements

First and foremost, I am grateful to my parents for their unending encouragement, support, and love. I am also grateful to all my family, friends, and mentors who have encouraged me and helped me along the way. Special thanks to Dr. Chuck (Charles) Rino, who has been an invaluable mentor and role model during my graduate career and whose work on radio scintillation has been a foundation for some of the work in this dissertation. My utmost gratitude also extends to the members of the SeNSe Lab, who through their camaraderie and collaboration have helped me to complete this work, and whose own accomplishments inspire me to become a better researcher. Thank you to all my committee members – Gary McGraw, Penny (Penina) Axelrad, Tomoko Matsuo, and Attilla Komjathy – who are my academic role models and whose knowledge and achievements inspire me. Their mentoring, insight, and feedback have substantially improved the quality of this dissertation. Finally, I am forever indebted to my advisor, Jade Morton, whose patience, wisdom, and assiduousness as a mentor have been instrumental in shaping me into the person and researcher I am today. Thank you for leading me through this journey.

This work was completed under support from the GAANN (Graduate Assistance in Areas of National Need) fellowship.

Contents

Chapter

1 An Overview of GNSS Cycle Slips **1**

1.1 Introduction . . . 1

1.2 Background . . . 3

1.2.1 Origin of Cycle Slips . . . 3

1.2.2 Phase Measurement Model . . . 4

1.2.3 Phase Combinations for Cycle Slip Mitigation . . . 8

1.2.4 Cycle Slip Mitigation Using Phase Combinations . . . 13

1.3 State-Space Approaches to Cycle Slip Mitigation . . . 15

1.3.1 Modern Approaches and Challenges in Cycle Slip Mitigation . . . 17

1.4 GNSS Cycle Slip Datasets . . . 18

1.4.1 LEO Ocean Reflection . . . 20

1.4.2 Mountaintop RO . . . 20

1.4.3 Datasets: Summary . . . 23

1.5 Summary . . . 23

2 Cycle Slip Simulation and Characterization **25**

2.1 Introduction . . . 25

2.1.1 Phase Transitions . . . 26

2.1.2 Previous Work . . . 29

2.2 Background . . . 33

2.2.1 Phase Screen Scintillation Model . . . 33

2.2.2 Simulating the Impact of Noise . . . 38

2.3 Methodology . . . 39

2.3.1 Simulation . . . 39

2.3.2 Signal Fade Definition . . . 39

2.3.3 Identifying Cycle Slips . . . 40

2.4 Results and Analysis . . . 41

2.4.1 Cycle Slip Occurrence Versus Fade Depth and Duration . . . 42

2.4.2 Cumulative Cycle Slip Occurrence . . . 42

2.5 Conclusion . . . 46

3 Occurrence of Challenging Cycle Slips in Real-World Data **50**

3.1 Introduction . . . 50

3.2 Background . . . 51

3.2.1 Geodetic Detrending . . . 51

3.2.2 Linear combinations . . . 53

3.3 Analysis of challenging cycle slips . . . 55

3.3.1 Ascension Island Scintillation . . . 56

3.3.2 Hong Kong Scintillation Example . . . 58

3.4 Cycle Slip Mitigation Performance: De Lacy 2012 . . . 61

3.4.1 Algorithm Description . . . 62

3.4.2 Analysis . . . 64

3.5 Cycle Slip Mitigation Performance: Filtering Algorithm . . . 69

3.5.1 Algorithm Description . . . 72

3.5.2 Assessment . . . 74

3.6 Summary and Discussion . . . 79

4 Probabilistic Modeling of Cycle Slip Detection and Estimation **81**

4.1 Introduction 81

4.2 System Model 82

4.3 Distributions of $\mathcal{G}$, $\mathcal{I}$, and ϵ 85

4.3.1 Noise Models for Harsh Conditions 87

4.3.2 Fixing Model Hyperparameters 90

4.4 Probabilistic Estimation 92

4.4.1 Integer Least Squares 94

4.5 Simulated Cycle Slip Estimation Performance 96

4.6 Summary 98

5 Novel Batch Algorithm for Cycle Slip Detection and Estimation **101**

5.1 Hyperparameter Tuning: Hong Kong Dataset 104

5.2 Sparse Detection 104

5.2.1 MM Algorithm 106

5.3 Detection and Tuning Parameters 108

5.4 Reduced Model 110

5.5 Search for ILS Solution 112

5.5.1 LAMBDA 112

5.5.2 Cliques 116

5.5.3 Method 1: Search and Shrink Over Cliques 118

5.5.4 Method 2: Approximate Support Using Marginals 120

5.6 Method Summary and Computational Considerations 124

5.7 Results 127

5.7.1 Simulated Data 127

5.7.2 Real Data 132

5.8 Summary and Discussion 141

6 Summary and Conclusion **144**

Bibliography **148**

Appendix

A Linear Combinations **156**

B Hyperparameter Estimation **158**

C Algebra for MAP and ILS Equivalence **160**

Tables

Table

1.1 Approximate range and variation in the components of L-band measurements for a typical stationary ground receiver. 8

2.1 Phase Screen Parameters 36

2.2 Scintillation Parameters 41

2.3 Cycle Slips Per Minute 48

2.4 Cumulative Cycle Slip Error Correlation Coefficient 49

3.1 Various linear combinations of observables useful for triple-frequency cycle slip detection. 54

3.2 Shows the sensitivity of each combination in Table 3.1 to different cycle slip amplitudes. Different colors are shown to help indicate the sensitivity, with yellow being largest and purple being smallest. 55

3.3 Manually detected cycle slip occurrences. Occurrence times are measured in seconds from 11:00 UT. 60

4.1 Characteristics and hyperparameter values for different scenarios corresponding to a variety of signal conditions. 97

5.1 Scintillation window characteristics and hyperparameter values 104

5.2 Comparison of predicted number of slips based on results from Chapter 3 and the actual number of estimated slips on each signal. 135

5.3 Estimation window C/N_0 and hyperparameter values for ocean reflection and mountaintop RO examples. 139

Figures

Figure

1.1 Shows two prototypical examples of cycle slip occurrences in GNSS phase measurements, both taken from the 2013-10-05 Hong Kong dataset, which we introduce in Section 1.4. The left panel shows a more easily identifiable slip in the presence of noise and a background trend. The right panel shows simultaneous phase measurements from two signals, with a less-obvious slip occurring on the L5 signal. 2

1.2 Hypothetical example of cycle slips occurring during phase unwrapping. 5

1.3 Illustrates the process of ionosphere scintillation, in which irregular structures in the ionosphere plasma density induce fluctuations received GNSS signals. 7

1.4 Examples of detrended carrier phase, code and carrier GF combinations, and the HWM combination in 30-second measurements for a receiver located in Brazil. Cycle slips occur due to effects caused by the low satellite elevation, indicated in the dashed line in the third panel. This data was obtained from the CDDIS (Crustal Dynamics Data Information System) website [55]. 11

1.5 Up-close example of the carrier GF and HWM combinations from Figure 1.4. In the top panel, the running estimates for the HWM mean and $\pm 4\sigma$ bounds are shown along with the raw HWM samples. Both panels indicate the magnitude of the jump in the combinations due to a cycle slip occurrence around 72 minutes into the window. 14

1.6 Triple-frequency C/N_0, phase, and phase combinations corresponding to a signal from GPS PRN 24 that was measured during a scintillation event on 2013-03-10 by a receiver on Ascension Island. 19

1.7 Triple-frequency C/N_0, phase, and phase combinations corresponding to a signal from GPS PRN 24 that was measured during a scintillation event on 2013-10-05 by a receiver near Hong Kong. 21

1.8 Measurements from L1 and L2 GPS signals that were reflected off the ocean surface and received by the low-Earth orbiting Spire radio occultation satellite. Top panel shows C/N_0 fluctuations, middle panel shows the detrended phase, and bottom panel shows the presence of cycle slips using the GF combination of L1 and L2 carrier phases. 22

1.9 Example of L1 signal C/N_0 and carrier phase measurements collected using a high-gain dish antenna on top of Mount Haleakala, Hawaii. 24

2.1 The first panel depicts how the trajectories of the complex modulation term in Equation 2.2 for two hypothetical fade scenarios. The second panel shows the corresponding time series in resulting amplitude and phase modulation of the signal. . . 27

2.2 Shows simulated phase and amplitude time series with the same parameters as example A in Figure 2.1, but with added noise assuming a C/N_0 of 40 dBHz. 28

2.3 Demonstrates two examples of canonical fades. The first two panels shows the diffractive field phasor trajectories in the complex plane. The bottom two panels show the signal amplitude and phase. Each canonical fade corresponds to a decrease of more than 10 dB in signal amplitude along with a rapid half-cycle phase change. The blue and red colors distinguish the noise-free and noisy versions of the signal. The flat blue line in the second panel represents the true phase offset due to a phase transition that occurs in the true phase during the second fade only. Meanwhile, the noisy signal shows how the impact of noise can mask or induce cycle slips that are or are not present in the noise-free signal. The alternate shaded portions of signal amplitude in the second panel demonstrates the definition of a single fade that is used for analysis in this paper. 29

2.4 The density of the random walk of the diffraction field phasor in the complex plane. The point at 1 (marked by a white dot) corresponds to zero diffractive perturbations, while the origin corresponds to complete destructive interference. The consecutive panels illustrate how, as scintillation strength increases, it may be more probable for the random walk to wrap around the origin and cause cycle slips. 33

2.5 An example simulation of scintillation amplitude (first panel) and phase (second panel) corresponding to the L1, L2, and L5 frequencies for a strong scintillation scenario S3T2. Note that phase has been scaled to TEC units. 37

2.6 Diffraction phase residual (obtained as the full phase observation minus the phase screen component) for the scenario shown in Figure 2.5. Phase is plotted in cycles, and the occurrence of diffraction-induced phase transitions clearly leads to integer-valued biases. The results of applying the TVD fit algorithm are shown with the dashed black line. 40

2.7 Shows the rate of cycle slips per fade as a function of fade amplitude (i.e. fading depth) for each baseline C/N_0. The scintillation parameters demonstrate increasing scintillation strength with $S_4 = 0.6, 0.8, 0.9, 1.0$ and $\tau = 0.6, 0.5, 0.5, 0.4$ for the respective panels (which corresponds to L5 signal for scenarios S1T2, S2T2, S3T2, and S4T2). A histogram estimate of the fade amplitude pdf is also shown in the dashed line. In all cases, lower baseline C/N_0 leads to an increase in cycle slips at deeper fades, while increase in scintillation strength leads to an increase in cycle slips for all $\overline{C/N_0}$ values and especially at shallower fade amplitudes. 43

2.8 Shows the rate of cycle slips per fade as a function of fade duration for each baseline C/N_0. The scintillation parameters are $S_4 \approx 0.95$ and $\tau \in \{1.0, 0.5, 0.2\}$ for each of the respective panels (which corresponds to L5 signal for scenarios S3T1, S3T2, and S3T3). A histogram estimate of the fade duration pdf is also shown in the dashed line. We generally see that lower baseline C/N_0 leads to an increase of cycle slip rates, but especially so for low fade rates ($\tau = 1$) and long fade durations. At higher fade rates (lower τ) all $\overline{C/N_0}$ values show similar fade rates at all durations. 43

2.9 Examples of pmf estimates obtained from histograms of the cumulative cycle slip error over 5-minute windows. This particular case shows the three frequencies for scenario S4T2 given in Table 2.2, and for a baseline C/N_0 of 36 dB-Hz. The black lines show the results of fitting a symmetric Skellam distribution to the pmfs. 45

2.10 Cycle slip rate dependence on baseline C/N$_0$, decorrelation time, and S_4 index. The cycle slip rates increase as baseline C/N$_0$ decreases, as decorrelation time (τ) decreases, and as scintillation strength (S_4) increases. The rates are derived from fitting the rate parameter of symmetric Skellam distributions to empirical distributions of cumulative cycle slip error over 5 minutes. The scintillation parameters increase in strength with $S_4 = 0.6, 0.8, 0.9, 1.0$ for the respective panels, while the different colored lines indicate different decorrelation times $\tau \approx 0.2 - 1.2$. These scintillation parameters correspond to the different scenarios given for the L5 signal in Table 2.2, and demonstrate the general relationship between cycle slip rate and different factors given in Table 2.3. 46

3.1 Shows the overlay of derivative of L1/L2 IF combinations of carrier phase. An overall linear trend has been removed for clarity. Each color corresponds to a different PRN and the white line shows the derivative of receiver clock variation that is obtained by averaging the IF derivatives. 52

3.2 Examples of L1/L2 and L1/L5 IF combinations and the L1/L2/L5 GIF combination using 100 Hz carrier phase measurements from the Ascension Island dataset. Bottom panel shows a zoomed-in portion of the most disturbed period. 56

3.3 Example of canonical fade behavior occurring on L2 and L5 signals at around 76311 seconds. L2 shows a sharp half-cycle phase change, corresponding to simulation scenario B, whereas L5 shows a slightly more gradual transition similar to scenario A. The GIF combination bias stays the same before and after the fade, indicating no cycle slip (phase transition) occurred. 58

3.4 Example of a canonical fade on the L5 signal. The behavior is similar to that of simulation scenario B. Similar to the example in Figure 3.3, there were no full-cycle transitions during this deep fade. 59

3.5 Examples of canonical fades on the L5 signal. The first canonical fade at 76485 seconds does not show a corresponding change in the GIF bias, and presumably no cycle slip has occurred. The second fade appears to show a noise-induced cycle-slip, corresponding to simulation scenario D. We see a corresponding drop in the GIF combination bias. 61

3.6 Using detrended phase and signal intensity, shows an interesting example of triple-frequency carrier behavior where it it is not clear from the individual phase records which signals (if any) accrue full cycle transitions and which just show canonical fade behavior. Careful measuring of the change in the average GIF combination value before and after the fade event reveals a jump of around 0.29, suggesting that a full-cycle transition occurred in the L2 signal only. 62

3.7 Using detrended phase and signal intensity, shows moderate canonical fades on the L1, L2 and L signals. There is a post-fade discrepancy in the L2 and L5 signals suggesting a possible full-cycle transition. A change in average GIF combination value of around 0.25 before versus after the fades suggests that the L5 signal underwent a full-cycle transition. 63

3.8 Using detrended phase and signal intensity, shows consecutive canonical fades on L2 and L5 signals, as well as a canonical fade on the L1 signal. The first fade in L5 coincides with a jump in GIF combination value of 0.25, indicating a positive full-cycle transition, while its second fade shows a downward half-cycle transition. The L2 signal shows two canonical fades in the same downward direction. 64

3.9 Using detrended phase and signal intensity, shows consecutive fades for L2 and L5 signals. The GIF combination show a jump in average value of around 0.50 before versus after the fade, indicating two full-cycle transitions in the L5 signal. The half-cycle transition behavior on the L2 signal is consistently downward while the full-cycle transition behavior of the L5 signal is consistently upward. 65

3.10 Using detrended phase and signal intensity, shows consecutive deep fades mainly for the L2 and L5 signals. The first fade coincides with canonical upward half-cycle transition of the L2 signal. After accounting for its noise-induced cycle-slip, the L5 signal also shows a canonical downward half-cycle transition. For the second fade, the L2 signal shows a downward full-cycle transition with a corresponding change in average GIF combination value, while the L5 signal shows a canonical upward half-cycle transition. 66

3.11 Shows examples of L1/L2 and L1/L5 ionosphere-free combinations and the L1/L2/L5 geometry-ionosphere free combination using 100 Hz carrier phase measurements. Top and bottom panels show the combinations before and after manual cycle slip correction. 67

3.12 The residuals of ΔY_i, $i = 1, 2, 3$, along with circles indicating the occurrence of an actual cycle slip in the manual truth reference for the corresponding signal (L1= 1, L2= 2, L5= 3). Dashed lines indicate slip detection thresholds, and overall we see many missed detections and false alarms for these combination residuals. 68

3.13 Detrended GF combinations of the L1/L2, L1/L5, and L2/L5 code phase observables, revealing a substantial increase in code phase noise and fluctuations during the scintillation event. 69

3.14 Residuals $\Delta\mathrm{GIF}_{\mathrm{L1,L2,L5}}$ plotted against $\Delta\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}$ along with their $\pm 4\sigma$ thresholds (dashed lines). Residuals corresponding to actual slips in the manual truth reference are encircled, similar to Figure 3.12 70

3.15 Estimated cycle slip amplitudes at each epoch for the method outlined in [20] as well as the amplitudes from the truth reference (shown with circles). 71

3.16 Comparison of truth reference of cycle slip bias sequence versus sequence estimated during the filtering process for L1, L2, and L5 signals (offset for clarity). There are 15, 63, and 78 slips in the truth reference and 26, 89, and 90 jumps in the filter estimate for the L1, L2, and L5 signals respectively. Of these, the filter estimated 8, 35, and 36 of the true jumps correctly. 75

3.17 Comparison of IF and GIF combinations of the carrier phase estimates from the Hong Kong dataset before and after applying the filtering method. 76

4.1 Illustrates the hypothetical structure of the various matrices relevant to a system model for triple-frequency signals. Panels a and b show the **A** and **B** matrices, respectively, from Equation 4.8. Panels c and d show the covariance matrices $\mathbf{Q_x}$ from Equation 4.15 and $\mathbf{Q_y}$ from Equation 4.23. 85

4.2 PSDs for ionosphere diffraction (top two panels) and corresponding phase noise variances for different bandwidth (bottom panel). Bandwidth is assumed equal to $1/T$ for sampling interval T. 91

4.3 Pull-in regions and corresponding probabilities for quiet (left panel) and disturbed (right panel) signal conditions. Each region is colored according to the probability of the float estimate $\mu_{\Delta\hat{\mathbf{z}}}$ lying in that region, given that no slip has actually occurred. 98

4.4 Probability of false slip amplitude identification under different signal conditions and for different model structures. Each row corresponds to signal conditions with model hyperparameters chosen according to Table 4.1. The columns show results for models using single-, dual- and triple-frequency measurements. 99

5.1 Example of cycle slips occurring in the Hong Kong scintillation dataset for GPS PRN 24 on 2013-10-03. The top panel shows signal C/N_0, the middle panel shows detrended carrier phase measurements, and the bottom panel shows the triple-frequency IF and GIF phase combinations. The shaded regions indicate the occurrences of cycle slips. 102

5.2 Hong Kong scintillation dataset that was introduced in Section 1.4.0.2. The first and second panels show the C/N_0 and geodetic-detrended phase for the three signals. The third and fourth panels show the scintillation indices S_4 and decorrelation times τ, respectively, for each of the signals. The shaded areas indicate the two time windows of interest. 105

5.3 Illustrates the MM process for a hypothetical univariate minimization problem. The black curve is the objective function, the gray curves are the series of quadratic majorizers, and the blue lines illustrate how the values that minimize the majorizers progress towards finding the minimum of the objective function. 108

5.4 Shows the rate (in slips per minute) of false alarms and missed detections for a range of values for the detection threshold a_{det} and tuning parameter λ. The colored points indicate crossover points of the false alarm and missed detection curves for the various values of λ. The black lines indicate curves for the value of λ that minimizes the crossover point. 111

5.5 Examples of the covariance $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}$ of the reduced model float estimate corresponding to the real data from Window 2. The second and third panels show the covariance and precision (i.e. inverse of the covariance) matrices for the permuted reduced model, where parameters are ordered by slip epoch. Different cliques over the model parameters are indicated in the third panel. 111

5.6 Illustrates the process of search and shrink, which seeks the best integer candidates within a given ellipsoid of particular float solution. 117

5.7 Illustrates the probabilistic graph structure of the cycle slip problem before and after marginalization of the phase components $\mathbf{x}$. Nodes indicate sets of variables in the model and edges imply conditional dependence between nodes. In this case, we show a complete graph for variables $\mathbf{z}$, indicating that all variables are interdependent. . . 119

5.8 Illustrates the concept of cliques in a PGM. 119

5.9 Block diagram describing the inputs and processes involved in the cycle slip estimation method developed in this paper. 128

5.10 Example of simulated triple-frequency scintillation data containing cycle slips. Panels a and b show the C/N_0 and phase for the three signals. Panel c shows the diffraction error containing cycle slips along with a truth estimate of slip occurrence. Panel d shows the float estimates of the slip amplitude sequences $\Delta\tilde{\mathbf{z}}$ obtained through the sparse estimation method. Panel e shows the MAP estimate of slip bias sequences. Panels f, g, and h show the IF, GIF, and GF phase combinations respectively, with gray and colored lines corresponding to the combinations obtained before (raw) and after (corrected) subtracting off the MAP slip bias estimate $\check{\mathbf{z}}$. 133

5.11 Distribution of phase error due to cycle slips before (raw) and after (corrected) subtracting MAP cycle slip estimate. The distribution is generated from histograms of the phase error over 100 simulations of scintillation measurements over 5-minute windows, with simulation parameters corresponding to scenario 1. 134

5.12 Detection and estimation results for Window 1 of the real scintillation data set. Panels a and b show the triple-frequency C/N_0 and phase, respectively. Panel c shows the float estimates of the slip amplitude sequences $\Delta\tilde{\mathbf{z}}$ obtained through the sparse estimation method. Panel d shows the MAP estimate of slip bias sequences. Panels e, f, and g show the IF, GIF, and GF phase combinations respectively, with gray and colored lines corresponding to the combinations obtained before (raw) and after (corrected) subtracting off the MAP slip bias estimate $\check{\mathbf{z}}$. 136

5.13 Detection and estimation results for Window 2 of the real scintillation data set. Panels and layout are the same as in Figure 5.12. 137

5.14 Dual-frequency detection and estimation results for Window 2 of the Hong Kong scintillation data set. Similar to Figure 5.12, panels a and b show the C/N_0 and phase, respectively. Panel c shows the float estimates of the slip amplitude sequences $\Delta\tilde{\mathbf{z}}$ obtained through the sparse estimation method. Panel d shows the MAP estimate of slip bias sequences. Panels e and f show the IF and GF phase combinations respectively. 138

5.15 Dual-frequency detection and estimation results for the GNSS-R ocean reflection dataset. Similar to Figure 5.14, panels a and b show the C/N_0 and phase, respectively. Panel c shows the float estimates of the slip amplitude sequences $\Delta\tilde{\mathbf{z}}$ obtained through the sparse estimation method. Panel d shows the MAP estimate of slip bias sequences. Panels e and f show the IF and GF phase combinations respectively. . . . 140

5.16 Single-frequency detection and estimation results for the Hawaii mountaintop RO dataset. Similar to Figure 5.14, panels a and b show the C/N_0 and phase, respectively. Panel c shows the float estimates of the slip amplitude sequences $\Delta\tilde{\mathbf{z}}$ obtained through the sparse estimation method. Panel d shows the MAP estimate of slip bias sequences. 142

B.1 Grid of average log-likelihood for corrected simulated scintillation measurements corresponding to Window 1 parameters. The white dot indicates the location for parameter values that minimize this average log-likelihood. 159

Chapter 1

An Overview of GNSS Cycle Slips

1.1 Introduction

Cycle slips are discrete and rapid changes in a signal carrier phase measurement bias. A wide range of factors can contribute to the occurrence of cycle slips in measurements from Global Navigation Satellite Systems (GNSS), including signal blockage, low signal-to-noise ratio, high platform dynamics, and effects of propagation through a turbulent ionosphere or troposphere. In short, harsh signal propagation conditions create a ripe environment for cycle slips to afflict carrier phase measurements. When left uncorrected, cycle slips introduce persistent biases into the measurement model that is used for a wide range of GNSS applications. As such, it is important to understand how they arise, how they behave, and how well we can mitigate their impacts.

The problem of cycle slip mitigation can be stated as follows: given a set of carrier phase measurements $\phi(t)$, determine whether a cycle slip has occurred and, if so, when did it occur and what is its amplitude. Figure 1.1 illustrates two prototypical examples of GNSS phase measurements containing cycle slips. In both panels, a cycle slip occurs at about halfway through on the L5 signal phase (red line) in the presence of a larger overall phase trend (black dashed line) and noise. The slip in the left panel occurs quickly relative to the variation in the trend, making it easier to identify. On the other hand, the slip in the right panel does not happen as quickly and is harder to identify. Its occurrence only becomes clear after comparison with measurements from the L1 signal transmitted by the same satellite, since in this particular case we know both signals should show similar phase trends.

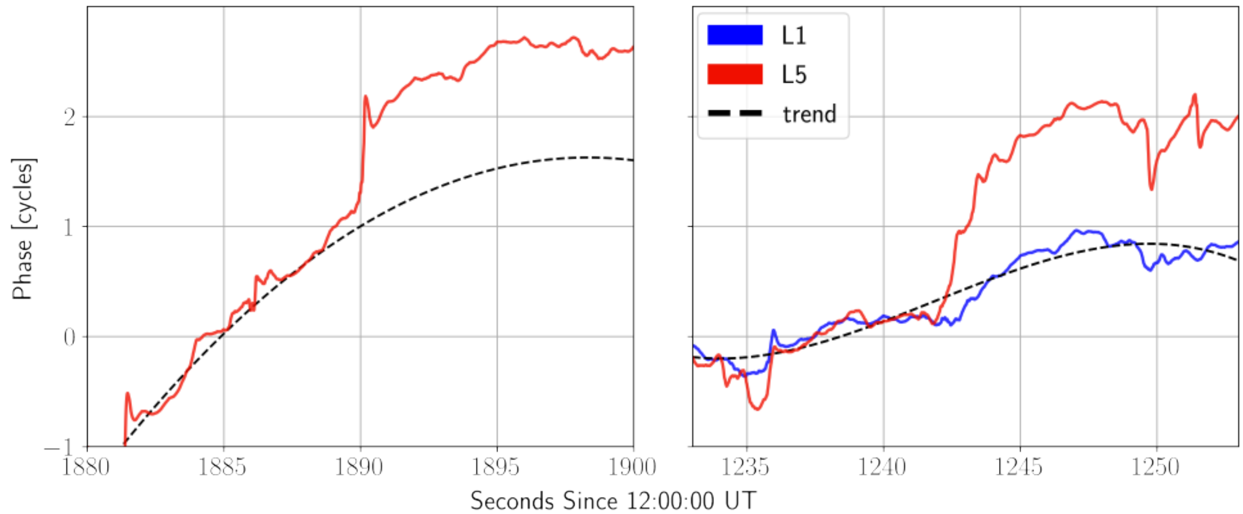

Figure 1.1: Shows two prototypical examples of cycle slip occurrences in GNSS phase measurements, both taken from the 2013-10-05 Hong Kong dataset, which we introduce in Section 1.4. The left panel shows a more easily identifiable slip in the presence of noise and a background trend. The right panel shows simultaneous phase measurements from two signals, with a less-obvious slip occurring on the L5 signal.

In general, our ability to visually identify these slips is dependent on the relationship between the slip characteristics, noise amplitude, and our knowledge about the background trend of $\phi(t)$. The presence of excessive noise or uncertain background trends can mask cycle slip occurrences and decrease the probability of their identification. The potential for multiple consecutive cycle slips exacerbates this issue. Some or all of these factors – i.e. large noise amplitude, uncertain background trends, and consecutive slip occurrences – are present in GNSS measurements collected under the conditions of multipath, weak signal power, strong atmospheric or ionospheric disturbances, and highly dynamic receiver platforms. These harsh conditions make cycle slip mitigation difficult and are the motivation for this dissertation. With such challenging data, it will sometimes be impossible to correctly identify cycle slips with high confidence. Therefore, in order to enable the effective use of GNSS phase measurements under such conditions, it is important to understand not only how to deal with cycle slips, but also what are the limits for how well they can be mitigated.

This dissertation includes six chapters. In the remainder of this chapter, we briefly discuss the relevant history and background of GNSS and techniques that have been applied to the cycle slip problem. We also show motivating examples of cycle slips in the context of different GNSS remote sensing applications. In Chapter 2, we use simulations to assess the interplay between signal diffraction and noise when it comes to cycle slip occurrence. We also characterize the cumulative error that cycle slips can cause in the case of diffractive ionosphere scintillation. In Chapter 3, we will take a closer look at cycle slip occurrence in real data containing deep signal fading and ionosphere diffraction and assess the effectiveness of two different cycle slip mitigation techniques under such conditions. The performance of these methods establish a need for an improved algorithm. Subsequently, in Chapter 4, we lay out a framework for cycle slip estimation using an arbitrary window multi-frequency phase measurements. We establish parameters for describing phase components smoothness, establish approaches for determining phase noise, and use these to quantify the probabilities of cycle slip identification under a variety of signal conditions. Then, in Chapter 5, we use these models to develop a technique for detecting sparse slip occurrences in a window of high-rate measurements and finding integer-least squares estimates in high dimensions. We demonstrate application of this framework to both simulated and real data sets. In Chapter 6, we provide a brief summary and discussion of the topics presented in this dissertation.

1.2 Background

1.2.1 Origin of Cycle Slips

A received baseband GNSS signal after demodulation can be modeled as:

$$s(t) = A(t) \exp(i\phi(t)) + \eta(t) \tag{1.1}$$

where A is the signal amplitude, ϕ is its phase, and $\eta(t)$ is assumed to be circular complex-valued noise. During tracking, a reference signal with phase $\phi_{\mathrm{ref}}(t)$ is correlated against the baseband signal to produce noisy observations of the received amplitude and residual phase:

$$\Delta s(t) = A(t) \exp(i\Delta\phi(t)) + \eta'(t) \tag{1.2}$$

where $\Delta\phi(t) = \phi(t) - \phi_{\text{ref}}(t)$. The raw phase measurement can then be reconstructed by adding the unwrapped phase residual to the reference phase:

$$\hat{\phi}(t) = \phi_{\text{ref}}(t) + \text{unwrap}\left(\angle \Delta s(t)\right) \tag{1.3}$$

In this process, there are generally three things that can cause cycle slips to arise. 1) If the difference between the reference phase and true phase change too much from epoch to epoch, i.e. the receiver loses phase lock, then the residual measurement will not reflect their actual phase difference and cycle slips will occur until lock is reestablished. 2) If noise dominates the signal, the residual phase measurements jump around and the unwrap operation will erroneously introduce cycle slips. As an example, Figure 1.2 illustrates how cycle slips can arise due to noise during the unwrapping process. 3) If the signal experiences interference due to multipath, the residual baseband signal Δs can wrap around the complex origin. We call this phenomenon a phase transition, and it introduces a cycle slip when the complex signal is unwrapped.

The first cause we mentioned, loss-of-lock, is one that we consider to be due to receiver malfunction. In other words, adjusting the PLL bandwidth or otherwise ensuring an adequate reference phase model can avoid such slips. There is also extensive literature concerning the second cause of a noise-driven PLL (c.f. [67], [1]). Work in [37] considers ionosphere-scintillation-induced cycle slips from this perspective, but does not specifically differentiate the impact of phase transitions. In Chapter 2 and Chapter 3, we discuss phase transitions in more detail and examine how the interplay between phase transitions and noise affects cycle slip occurrence.

1.2.2 Phase Measurement Model

So far we have introduced the origin of cycle slips, but their impact as a source of error is only meaningful within the context of an actual phase measurement model. In this work, we consider the following models for GNSS code phase and carrier phase observables:

$$\rho_k(t) = \mathcal{G}(t) + \beta_k \mathcal{I}(t) + \mathcal{B}_k + n_k(t) \tag{1.4}$$

$$\Phi_k(t) = \frac{\lambda_k}{2\pi}\phi_k(t) = \mathcal{G}(t) - \beta_k \mathcal{I}(t) + \mathcal{B}_k + \lambda_k z_k(t) + \epsilon_k(t) \tag{1.5}$$

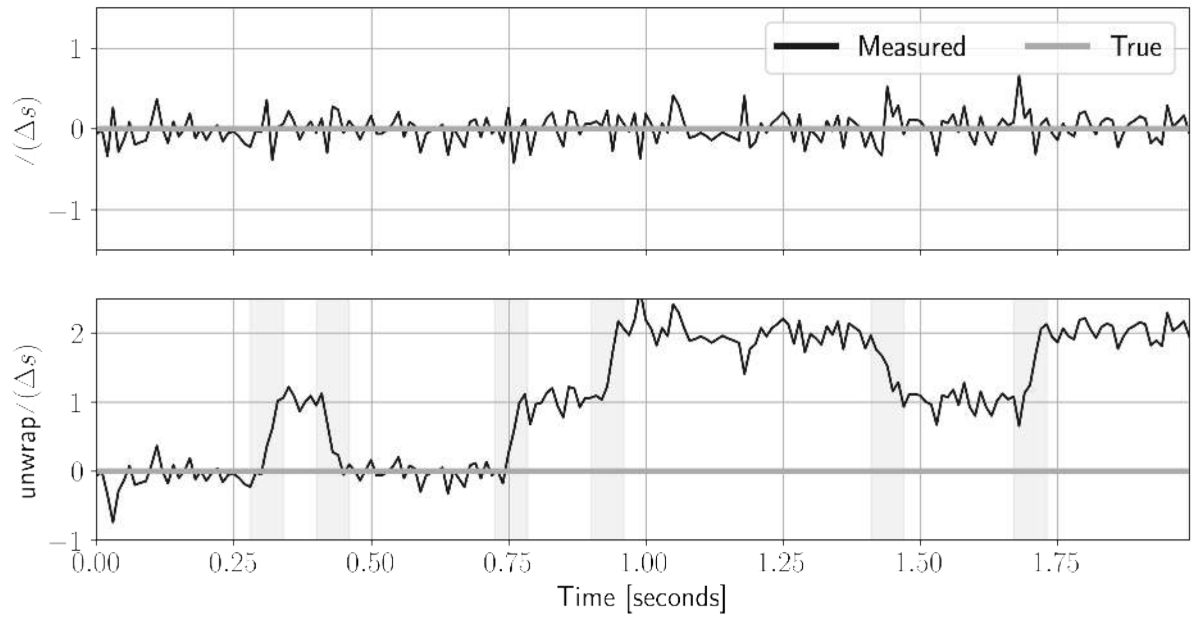

Figure 1.2: Hypothetical example of cycle slips occurring during phase unwrapping.

where

- k denotes a quantity corresponding to the k^{th} signal with carrier frequency f_k and wavelength λ_k
- $\mathcal{G}$ is the non-dispersive phase component, including any geometric range, clock, and tropospheric effects, etc.
- $\mathcal{I}$ is the first-order ionosphere delay on the first signal frequency
- $\beta_k = f_1^2/f_k^2$ defines the proportionality of the ionosphere delay among the different signals
- $\mathcal{B}_k$ is a fractional bias term, including the effects of hardware delays, etc.
- z_k is the term for the integer-valued bias in the phase measurement due to integer ambiguity and cycle slips
- n and ϵ account for noise and unmodeled effects, including multipath and scintillation-induced fluctuations

The $\mathcal{G}$ term can be broken down into its main components as:

$$\mathcal{G} = r + c\,(\delta t_{\mathrm{tx}} - \delta t_{\mathrm{rx}}) + D_{\mathrm{tropo}} \tag{1.6}$$

where

- r is the transmitter-to-receiver antenna range
- δt_{tx} and δt_{rx} are the transmitter and receiver clock errors
- D_{tropo} is the tropospheric range delay

Additionally, the $\mathcal{I}$ term can be related to the underlying physical parameter of ionosphere total electron content (TEC):

$$\mathcal{I}(t) \approx \frac{\kappa}{f^2}\mathrm{TEC}(t) \tag{1.7}$$

where

- TEC measured in TEC units (1 TECu $= 1 \times 10^{16}$ electrons/meter^2) is an integrated measure of the ionosphere plasma content along the signal ray path
- $\kappa \approx 40.308 \times 10^{16}\mathrm{m} \cdot \mathrm{s}^2/\mathrm{TECu}$ is a constant

This ionospheric term only accounts for the first-order refractive ionosphere effect, though higher-order effects are known to be comparatively small [38].

There are various effects that contribute to the noise terms ϵ and n. In particular, the ionosphere is known for causing phase and amplitude disturbances when signals diffract through plasma irregularities. This effect is called ionosphere scintillation, and in addition to larger variations in $\mathcal{I}$ it also causes increased measurement noise levels and carrier phase cycle slips. Figure 1.3 illustrates the process of ionosphere scintillation. A similar phenomenon occurs when occulting signals diffract through irregularities in the lower troposphere, which is known as tropospheric scintillation. More generally, any form of multipath or signal interference will lead to increased code and carrier phase noise levels. Additionally, there are the effects due to thermal noise, which will be especially important for signals with very low $\mathrm{C/N_0}$ or signals experiencing amplitude fading.

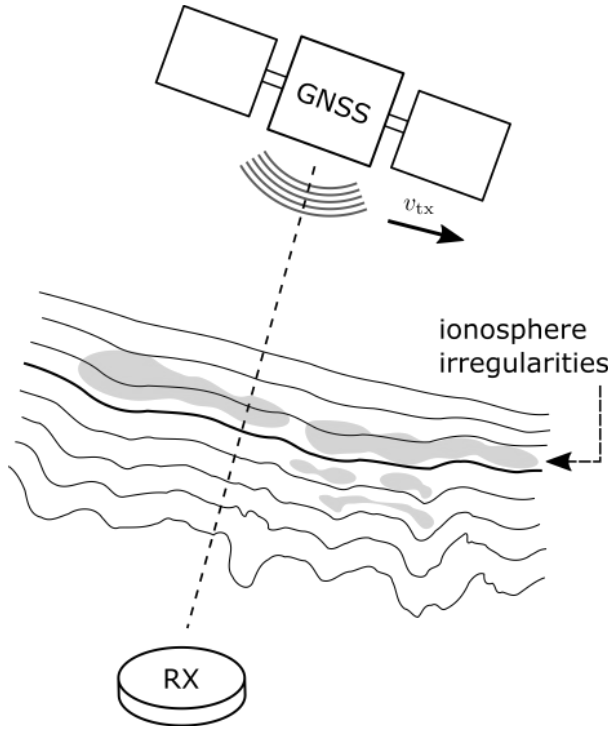

Figure 1.3: Illustrates the process of ionosphere scintillation, in which irregular structures in the ionosphere plasma density induce fluctuations received GNSS signals.

When it comes to cycle slip mitigation, it is important to understand how the variation in these components compares to potential changes in the integer-bias due to cycle slips. Table 1.1 lists the size and variation of the different phase components one might see in measurements from a stationary receiver on the ground. The largest and most highly varying components are the satellite-receiver range r, receiver clock error δt_{rx}, and refractive ionosphere effect $\mathcal{I}$. Any effective approach to cycle slip mitigation should make some effort to model and estimate these phase components. The troposphere variation is smaller or less variable, and so it is usually not as critical to model. Also note that the code phase noise n_k is generally 1-2 orders of magnitude larger than the carrier phase noise ϵ_k. This has important implications with regards to the ineffectiveness of code phase observables when it comes to detecting or estimating cycle slips with small amplitudes. Note that multipath or diffraction effects will lead to carrier phase errors significantly larger than the typical values listed in the table. Overall, this model will be sufficient for discussing previous cycle slip

	range	variation
r	20,000 - 26,000 km	0-4000 m/s
$c\delta t_{\text{rx}}$	$\pm 1 - 2$ m	0-300 m/s
$\mathcal{I}$	1-50 m	0-0.8 m/min
D_{tropo}	2.5-25	
n_k	0.05-0.5 m @ 1 Hz	
ϵ_k	±4mm @ 1 Hz	

Table 1.1: Approximate range and variation in the components of L-band measurements for a typical stationary ground receiver.

mitigation techniques in this chapter as well as for introducing our own methods in Chapter 4.

1.2.3 Phase Combinations for Cycle Slip Mitigation

There is extensive literature on cycle slip mitigation for GNSS signals, documenting a progression of techniques since the launch of the Global Positioning System in the 1980s. The earliest algorithms developed to work with single or dual-frequency GPS L1 and L2 signals all looked for change outliers in specific combinations of measurements. For instance, [32] applied a Kalman filter to the so-called geometry-free carrier phase combination and flagged outliers in its variation as cycle slips to be fixed. The work in [7] extended this approach to use both code-minus-carrier and the geometry-free carrier combinations. Both authors acknowledged how an active ionosphere and/or receiver motion and clock dynamics associated with large measurement fluctuations can present challenges when detecting outliers in phase time series. As such, linear combinations of measurements that can isolate or remove these signal components have become a common thread in much of the literature on cycle slip mitigation. Their use is also motivated by practical and computational concerns in various other GNSS applications and processing steps, including standard positioning and ambiguity resolution [60]. We use linear combinations of phase observations throughout this work in order to demonstrate the effect of cycle slips in real data, to explain previous mitigation approaches, and to assess improvement in mitigation outcomes.

For K signals at different carrier frequencies, an arbitrary linear combination of code and carrier phase measurements can be expressed as:

$$Y = \sum_k c_{\rho_k} \rho_k + c_{\Phi_k} \Phi_k \tag{1.8}$$

where c_{ρ_k} and c_{Φ_k} are the linear combination coefficients for code phase and carrier phase respectively. Also note that if the noise terms in Equation 1.4 and Equation 1.5 are described by covariances $\mathbf{Q}_\rho$ and $\mathbf{Q}_\Phi$, then the variance of the resulting combination is:

$$\sigma^2 = \mathbf{c}_\rho^T \mathbf{Q}_\rho \mathbf{c}_\rho + \mathbf{c}_\Phi^T \mathbf{Q}_\Phi \mathbf{c}_\Phi \tag{1.9}$$

where $\mathbf{c}_P$ and $\mathbf{c}_L$ are the coefficient vectors for code and carrier phase measurements.

Most often, linear combinations are chosen such that they remove contributions from the non-dispersive effects (modeled by $\mathcal{G}$) and/or the refractive ionosphere effect (modeled by $\mathcal{I}$) in the resulting combination. These are referred to as geometry-free (GF) and ionosphere-free (IF) combinations, respectively, and their coefficients satisfy:

$$\sum_k c_{\rho_k} + c_{\Phi_k} = 0 \qquad \text{GF} \tag{1.10}$$

$$\sum_k \beta_k c_{\rho_k} - \beta_k c_{\Phi_k} = 0 \qquad \text{IF} \tag{1.11}$$

1.2.3.1 Dual-Frequency Combinations

Among the various geometry-free and ionosphere-free combinations, the dual-frequency combinations using only code phase or only carrier phase have possibly the widest application. For the first two decades of GNSS, satellites in the GPS and GLONASS constellations only transmitted signals in the L1 and L2 frequency bands, which are centered at 1575.42 MHz (L1) and 1227.60 MHz (L2) for GPS specifically. The L1/L2 dual-frequency ionosphere-free combination is widely used when correcting for the ionosphere effect during positioning [27]. Meanwhile, the dual-frequency geometry-free combination is used for measuring ionosphere TEC [68]. We express the GF and IF combinations (applied to either code or carrier phase) as:

$$\mathrm{GF}_{j,k} = \rho_j - \rho_k \qquad \text{or} \qquad \Phi_j - \Phi_k \tag{1.12}$$

$$\mathrm{IF}_{j,k} = \frac{f_j^2 \rho_j - f_k^2 \rho_k}{f_j^2 - f_k^2} \qquad \text{or} \qquad \frac{f_j^2 \Phi_j - f_k^2 \Phi_k}{f_j^2 - f_k^2} \tag{1.13}$$

The dual-frequency carrier phase GF combination has been widely applied for cycle slip detection and estimation (c.f. [6], [8], [50]). This is because the non-dispersive phase components in $\mathcal{G}$ are often the largest and most uncertain, and so their presence makes cycle slip estimation more difficult. As an example, we consider the dual-frequency phase measurements in the top panel of Figure 1.4. The data was recorded at 30-second intervals by a receiver located in Brazil that is part of the International GNSS Service (IGS) network. We have subtracted off the largest phase variations, which are due to the satellite-receiver range, so that we can more easily see the effect of cycle slips and other phase variations that are present in the data. In this case, the cycle slips arise due to the low satellite elevation (indicated by the dashed line in the third panel) which can often cause excessive noise, multipath, and loss-of-lock when below $\sim 15 - 20°$. These slips are notably different from the earlier examples shown in Figure 1.1; they are larger and effectively instantaneous due to the low sampling rate. Even still, they are somewhat difficult to discern in the top panel due to residual phase variations from the receiver oscillator, troposphere, and refractive ionosphere effects. Rather, they are more easily identified in the geometry-free phase combination shown in the second panel, which still contains the ionosphere variation but is overall much smoother due to being rid of the non-dispersive components. The noisier GF code measurement combination is also displayed to indicate the approximate GF phase trend without cycle slips.

When the ionosphere variation is relatively smooth, as is the case for our example in Figure 1.4, the GF carrier phase combination is clearly a good option for detecting cycle slips. However, this combination alone does not contain enough information to estimate the amplitude of cycle slips on the two different carriers. If we want to estimate cycle slip amplitudes using phase combinations, we need another independent measurement combination that is also effective for observing cycle

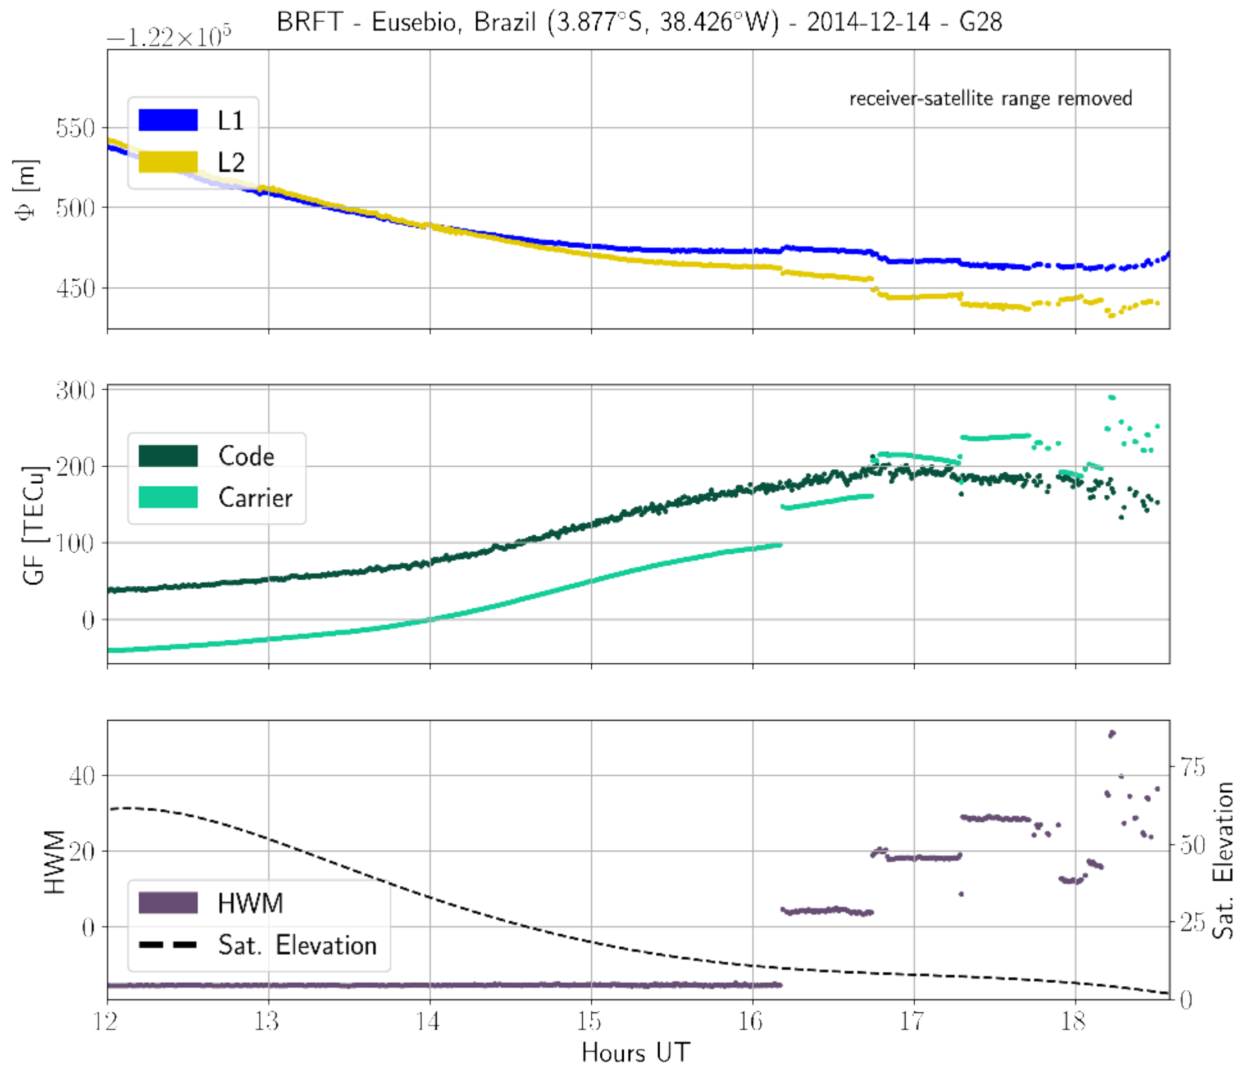

Figure 1.4: Examples of detrended carrier phase, code and carrier GF combinations, and the HWM combination in 30-second measurements for a receiver located in Brazil. Cycle slips occur due to effects caused by the low satellite elevation, indicated in the dashed line in the third panel. This data was obtained from the CDDIS (Crustal Dynamics Data Information System) website [55].

slips. We consider the widelane and narrowlane combinations, which are given as:

$$\mathrm{WL}_{j,k} = \frac{f_j \rho_j - f_k \rho_k}{f_j - f_k} \quad \text{or} \quad \frac{f_j \Phi_j - f_k \Phi_k}{f_j - f_k} \tag{1.14}$$

$$\mathrm{NL}_{j,k} = \frac{f_j \rho_j + f_k \rho_k}{f_j + f_k} \quad \text{or} \quad \frac{f_j \Phi_j + f_k \Phi_k}{f_j + f_k} \tag{1.15}$$

Their names stem from how the difference or summation of phase (in cycles) can be interpreted as a measurement corresponding to a signal with a wider or narrower (shorter) wavelength. For example, the WL combination can be considered to have a wavelength of $\lambda_{\mathrm{WL}} = c/(f_j - f_k)$, which is clearly larger than either λ_j or λ_k. The widelane and narrowlane combinations are not typically applied on their own, but rather are combined to form the Hatch-Melbourne-Wübenna (HMW) combination, which is given by:

$$\mathrm{HMW}_{j,k} = \frac{f_j \Phi_j - f_k \Phi_k}{f_j - f_k} - \frac{f_j \rho_j + f_k \rho_k}{f_j + f_k} \tag{1.16}$$

Note that this combination uses both carrier and code observations. Consider the result of substituting Equation 1.4 and Equation 1.5 into Equation 1.16, which after some simplification yields:

$$\mathrm{HMW}_{1,2} = \lambda_{\mathrm{WL}}(z_1 - z_2) + \frac{\lambda_{\mathrm{WL}}}{\lambda_1}\epsilon_{\Phi_1} + \frac{\lambda_{\mathrm{WL}}}{\lambda_2}\epsilon_{\Phi_2} + \frac{\lambda_{\mathrm{NL}}}{\lambda_1}\epsilon_{\rho_1} + \frac{\lambda_{\mathrm{NL}}}{\lambda_2}\epsilon_{\rho_2} \tag{1.17}$$

From this expression, it is clear why this combination is used for cycle slip detection. In addition to it being both ionosphere- and geometry-free, its use of the widelane carrier phase combination amplifies the integer ambiguity terms while its use of the narrowlane code phase combination suppresses the large code phase noise (since $\lambda_{\mathrm{WL}} > \lambda_{\mathrm{NL}}$). The third panel in Figure 1.4 shows the HMW combination, and the slips that are obvious in the GF combination are also apparent in the HMW combination. Moreover, the ionosphere variation is eliminated in this combination, although there is larger noise presence than in the GF carrier phase combination.

1.2.3.2 Triple-Frequency Combinations

Since the launch of the first GPS Block-IIF and QZSS satellites in 2010, more and more GNSS satellites broadcasting at three or more frequencies have come online. Even before their launch, the various linear combinations involving triple-frequency signals were investigated [18]. In particular, the GPS L5 signal, centered at 1,176.45 MHz, allows for computing a combination that is both ionosphere-free and geometry-free and that uses only carrier phase measurements. The so-called geometry-ionosphere-free (GIF) combination has been used in a variety of applications, from identification of subtle variations in the satellite hardware biases [53] to fixing of integer ambiguities [84] and cycle slips [20]. We choose to express the carrier-phase-only GIF combination as:

$$\mathrm{GIF}_{i,j,k} = (\beta_j - \beta_k)\Phi_i + (\beta_k - \beta_i)\Phi_j + (\beta_i - \beta_j)\Phi_k \tag{1.18}$$

for β as defined just after Equation 1.5. Derivation of these and other combination coefficients is provided in Appendix A. Examples of the cycle slips in the triple-frequency GIF phase combination are included in Section 1.4 and in Chapters 3 and 4.

1.2.4 Cycle Slip Mitigation Using Phase Combinations

So far, we have seen how certain measurement combinations can be useful for discerning the effect of cycle slips. One of the most widely cited approaches to cycle slip correction, the TurboEdit algorithm from [8], makes use of the HMW and GF carrier combinations to mitigate cycle slips in dual-frequency measurements. To deal with the effects of code measurement noise, it computes the running mean and variance of the HWM combination over phase-connected segments. If a subsequent sample of the HWM combination lies outside of a specified threshold, say ± 4 standard deviations from the mean, then a slip is detected. The top panel of Figure 1.5 illustrates this process for one of the slips that occurs in the data segment from Figure 1.4. The running average of HWM is shown in a solid line with $\pm 4\sigma$ bounds indicated by the dashed gray lines. Once a slip is detected, the mean and variance estimates are reset and the slip amplitudes can be estimated using the changes in the HWM and GF combinations caused by the slip. For the HWM combination, we compute ΔHWM as the difference between first sample after a slip and the last mean estimate before the slip. For the GF carrier phase combination, in order to account for the ionosphere trend, we use a polynomial extrapolation to predict the next GF sample after the slip. We take ΔGF to

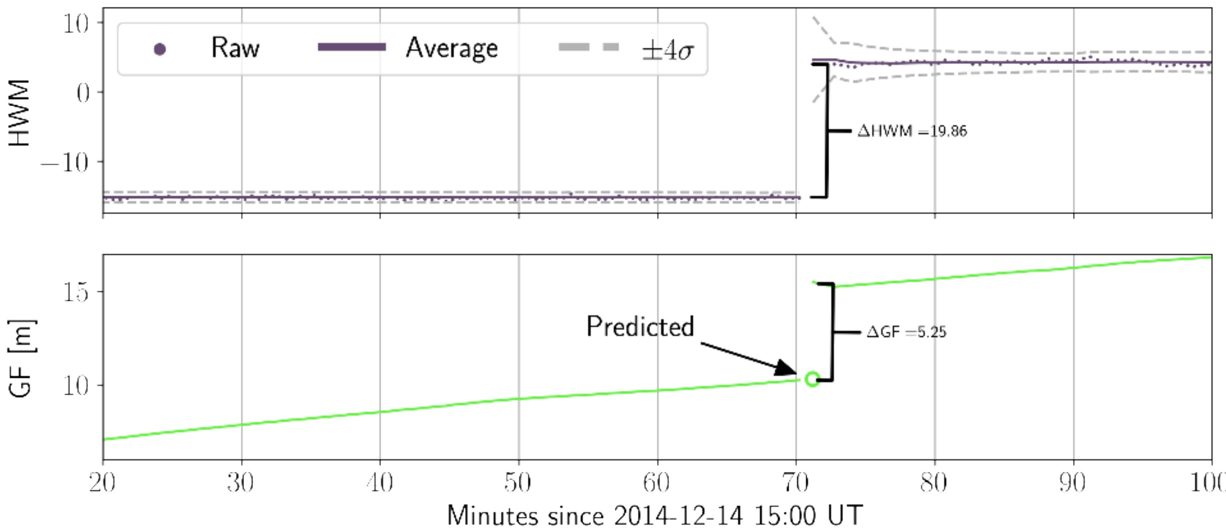

Figure 1.5: Up-close example of the carrier GF and HWM combinations from Figure 1.4. In the top panel, the running estimates for the HWM mean and $\pm 4\sigma$ bounds are shown along with the raw HWM samples. Both panels indicate the magnitude of the jump in the combinations due to a cycle slip occurrence around 72 minutes into the window.

be the difference between the actual and predicted GF samples. This process is also illustrated in second panel of Figure 1.5. With ΔHWM and ΔGF we can estimate the slip amplitudes by solving the system:

$$\begin{bmatrix} \Delta \text{HWM} \\ \Delta \text{GF} \end{bmatrix} = \begin{bmatrix} \lambda_{WL} & -\lambda_{WL} \\ \lambda_1 & -\lambda_2 \end{bmatrix} \begin{bmatrix} \Delta z_1 \\ \Delta z_2 \end{bmatrix} \tag{1.19}$$

where Δz_k are the cycle slip amplitudes on each signal.

After solving this system, the values for Δz_k need to be fixed to integers. Cycle slip mitigation algorithms tend to do one of two things: 1) they round them to the nearest integer, or 2) they search for the best integer solution taking into account correlation between the estimates of Δz_1 and Δz_2. For this second approach, many algorithms implement the least-squares ambiguity decorrelation and adjustment (LAMBDA), which is an established method for fixing integer-valued parameters that is widely used in GNSS ambiguity resolution [77]. We discuss this algorithm in more detail when we introduce our approach to cycle slip estimation in Chapter 4.

Overall, the essential steps of the TurboEdit algorithm can be described as detection, regression, and fixing. To elaborate, we first detect slips as outlier in the changes of some. Then, we use the deviation in these combinations from their nominal or predicted values in order to regress on the cycle slip amplitude parameters Δz. Finally, Δz are fixed to integer values. These are the prototypical steps for most cycle slip algorithms, and many of the subsequent mitigation approaches can be seen as extensions of the Turboedit algorithm. For example, [52] introduced weighting factors for the HWM coefficients based on satellite elevation, and the authors from [103] and [102] extended the HMW to analogous triple-frequency combinations. Other algorithms try to make use of other measurements combinations. [94] included use of the code-carrier widelane (WL) combination to help address difficult-to-detect slips in dual-frequency measurements. [20] suggested a series of combinations to address slips in triple-frequency signals with an emphasis on the GPS and Galileo systems, and later [35] presented a similar approach designed for the Beidou system. Additionally, the work in [26] offers some improvement to these methods by introducing dynamic detection thresholds based on the posterior distribution of the various phase combinations. However, when it comes to small cycle slips (on the order of 1-2 cycles), it has been repeatedly shown that code phase noise in linear combinations overwhelms the ability of any combination using code phase observables [20]. Therefore, many authors focused their efforts on tracking the ionosphere variation in order to make optimal use of the precise information in the GF phase combination [50], [14], [101]. Alternatively, IF combinations can be used when triple-frequency signals are available and/or the non-dispersive signal component can be estimated [20], [54]. Either way, the fact remains that correct identification of small amplitude slips (1-2 cycles) relies heavily on the use of multi-frequency carrier phase measurements and constraints on uncertainty of dispersive and non-dispersive signal components.

1.3 State-Space Approaches to Cycle Slip Mitigation

While the use of specific phase combinations provides an intuitive and computationally convenient approach to identifying cycle slips, performance of such methods cannot outperform those that use a full set of unaltered phase measurements. This fact was pointed out in [3], where the authors considered a state-space approach for cycle slip estimation. By "state-space" approach, we simply mean using models of the form:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{B}\Delta\mathbf{z} + \epsilon \tag{1.20}$$

where

- $\mathbf{y}$ is a full set of raw measurements
- $\mathbf{x}$ is some set of float-valued parameters
- $\Delta\mathbf{z}$ is some set of integer-valued cycle slip amplitudes
- $\mathbf{A}$ and $\mathbf{B}$ represent the system relating these parameters to our measurements
- ϵ is noise

The key aspect of these approaches is that they use uncombined measurements in $\mathbf{y}$, in contrast to the approaches we discussed in the last section that lose information when combining measurements. However, this does not mean that the use of measurement combinations is necessarily inferior. Methods using carefully chosen combinations can be viewed as efficient approximations of state-space approaches. More specifically, their combination coefficients arise as approximations to the eigenvectors of the matrix $\mathbf{A}\mathbf{Q}_{\mathbf{x}}\mathbf{A}^T + \mathbf{Q}_\epsilon$, where $\mathbf{Q}_\epsilon$ and $\mathbf{Q}_{\mathbf{x}}$ are covariances quantifying the uncertainty in the noise and non-integer parameter model terms. This matrix is a key part of our estimation approach that we develop in Chapter 4.

In models of this form, the terms $\mathcal{G}$ and $\mathcal{I}$ would be included in $\mathbf{x}$. Clearly, their behavior and how they are modeled has a direct impact on the estimation of $\Delta\mathbf{z}$, and vice versa. [3] quantified how uncertainty in these dispersive and non-dispersive signal components leads to incorrect estimation of slip amplitudes in dual-frequency signals. Subsequent work in [99] assessed how a third signal frequency would benefit estimation performance. The work in [90] demonstrates the performance of models using dual- and triple-frequency measurements from multiple satellites simultaneously. One of the main benefits of state-space approaches is that they offer greater flexibility in how we model the different phase components and their uncertainty. For example, with the approach we introduce in Chapter 4 we assume stationary Gaussian process models for the non-dispersive and refractive ionosphere phase components.

1.3.1 Modern Approaches and Challenges in Cycle Slip Mitigation

Most cycle slip approaches try to detect slip occurrences using just two measurement epochs, or alternatively, a single epoch of time-differenced measurements. These types of algorithms have been successful under normal signal conditions for 1 Hz or slower measurements. For instance, the authors in [20] found zero faulty cycle slip estimates when applying their algorithm to 1 Hz GPS data, even with phase noise standard deviations reaching up to 1/2 cycle. However, when it comes to more challenging signal conditions, uncertainty in the variation of signal phase components can cause most cycle slip algorithms to fail. The authors of [2] first discussed the particularly difficult problem of dealing with cycle slips during ionosphere plasma bubble events when phase measurements contain multiple cycle slips and unpredictable variations due to changes in ionosphere total electron content (TEC). Along with the authors in [101] and [65], they emphasize the utility of IF phase combinations and careful estimation of non-dispersive phase components when it comes to detecting cycle slips in these scenarios. However, as noted in [3], even with careful estimation of non-dispersive components it may still be impossible to effectively estimate slip occurrences in time-differenced measurements when the ionosphere variability is severe.

Several authors noted the greatly improved estimation capability when using multiple measurement epochs. For instance, the authors in [21] provided a probabilistic approach to detecting slip occurrences in a window of measurements by deriving the posterior for slip occurrence in a polynomial regression. The authors in [14] considered a forward-backward moving window algorithm specifically targeting dual-frequency signals during high ionosphere activity, and [102] presented a similar algorithm for triple-frequency signals. Also, while a majority of cycle slip mitigation research has centred around ionosphere effects, it is not the only source of challenging signals containing cycle slips. One example comes from the work in [85], which presents an algorithm to detect and remove cycle slips in weak signals reflected off the ocean surface. Other challenging sources of cycle slip occurrence include troposphere scintillation [96], urban multipath [28], and high-dynamic receiver platforms or oscillators. For this latter case, [51] proposed detecting slips as outliers in the singular spectrum over a window of measurements, and [49] proposed multi-epoch detection in the context of sequential filtering. Overall, the common thread in the most effective algorithms in each of these domains is that they use as much information as possible: i.e. extended windows of high-rate raw phase measurements.

1.4 GNSS Cycle Slip Datasets

Here we introduce datasets of signals under harsh conditions that contain multiple cycle slips. We will use these datasets in chapters 3, 4, and 5 to assess and validate different mitigation algorithms. For the Ascension Island and Hong Kong datasets introduced below, we use phase measurements where a majority of non-dispersive phase components have been removed using a detrending technique that we describe in Section 3.2.1 of Chapter 3.

1.4.0.1 Ascension Island

The first dataset consists of commercial receiver output (Septentrio PolaRxS) along with raw intermediate-frequency (IF) data for triple-frequency GPS signals from a receiver near Ascension Island (7.95°, 14.36° W) on 2013-03-10. Raw data from GPS L1, L2, and L5 bands was collected and recorded by the SeNSe Lab, and was tracked using the robust tracking algorithm from [92]. As shown in the first panel of Figure 1.6, the signal from GPS PRN 24 contained significant fluctuations in amplitude beginning at around 20:00 UT (also local) and lasting until around 21:30 UT. During this time, the satellite rose from around 13° to 27° elevation. The middle and bottom panels of Figure 1.6 show the detrended carrier phase and GIF carrier phase combination for the L1, L2, and L5 signals. Many discrete jumps occur, indicating the frequent occurrence of diffraction-induced cycle slips. These jumps coincide with fading of the signal amplitude. We take a closer look at the cycle slips that occur in this dataset in Section 3.3.1 of Chapter 3.

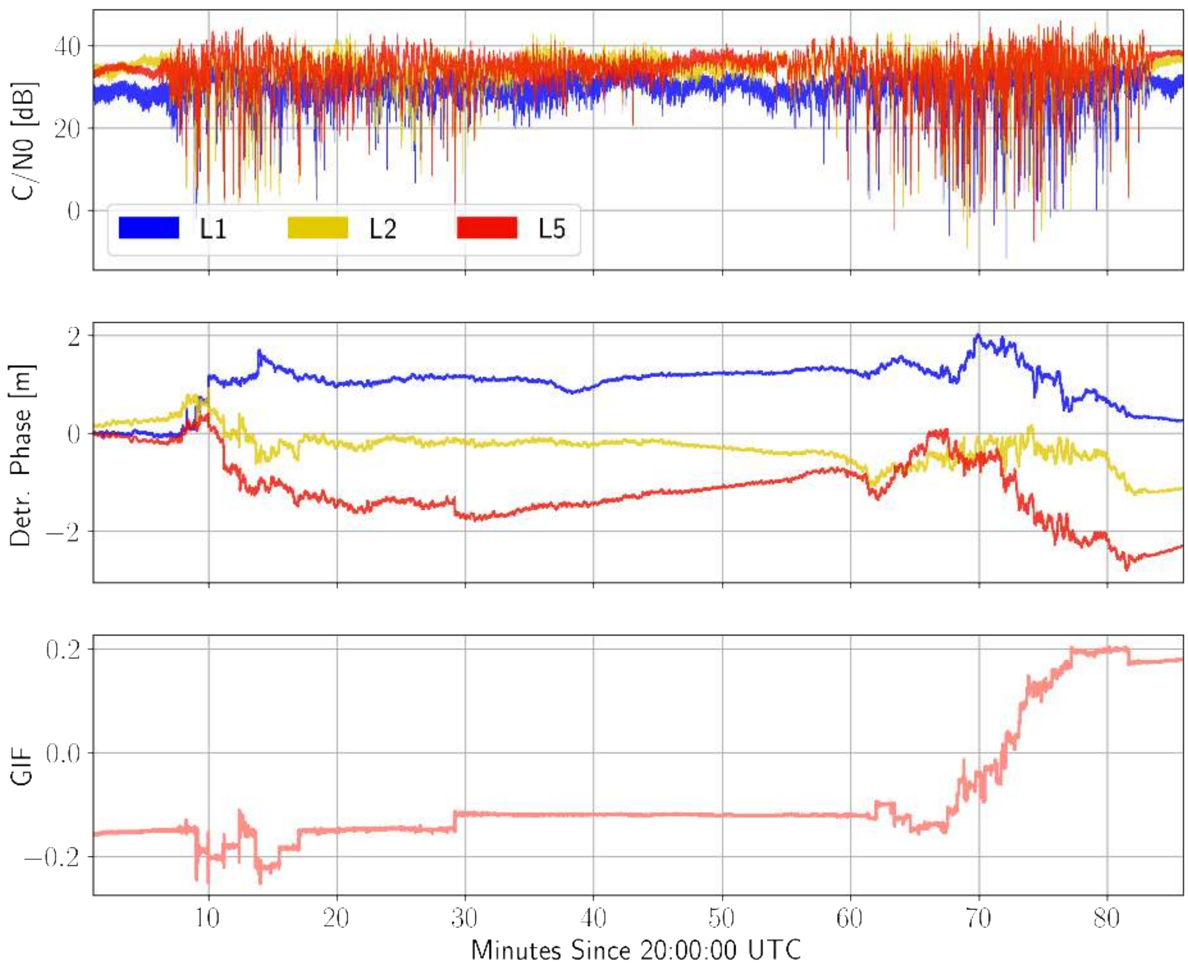

Figure 1.6: Triple-frequency C/N_0, phase, and phase combinations corresponding to a signal from GPS PRN 24 that was measured during a scintillation event on 2013-03-10 by a receiver on Ascension Island.

1.4.0.2 Hong Kong

The second dataset comes from a commercial GNSS receiver (Septentrio PolaRxS) located just outside of Hong Kong (22° N, 115° E). GNSS measurements exhibited effects due to a strong ionosphere scintillation event on 2013-10-05. We again consider signals from GPS PRN 24, which showed particularly strong deep fades lasting from around 12:00 to 13:00 UT or around 20:00 to 21:00 LT. During this time, the satellite rose from around 31° to 61° elevation. The top panel of Figure 1.7 shows the signal C/N_0 obtained for L1, L2, and L5 signals, demonstrating the presence of amplitude fluctuations associated with diffractive scintillation. The second panel shows the detrended phase measurements and the bottom two panels show the IF and GIF carrier phase combinations, respectively. Again, jumps in the phase combinations indicate the occurrence of cycle slips.

1.4.1 LEO Ocean Reflection

In addition to cycle slips caused by ionosphere scintillation, we also consider slips in the L1/L2 GPS signals that were reflected off the ocean and collected by a receiver on board a low Earth orbiting satellite from Spire Inc.. Figure 1.8 shows the signal amplitudes in the top panel, the detrended phases in the middle panel, and the GF combination of L1 and L2 carrier phases in the bottom panel. The occurrence of discrete jumps in the GF signal indicate cycle slips. The chaotic phase behavior at the beginning of the plot likely corresponds to non-coherent signal measurements, similar to the previous example. It is important to acknowledge that as signal conditions degrade we will reach a point where cycle slip mitigation is no longer feasible or even meaningful.

1.4.2 Mountaintop RO

Our final example of cycle slips comes from mountaintop radio occultation (RO) measurements. A high-gain dish antenna located on top of Mount Haleakala, Hawaii (3 km altitude) collected the signals from very-low-elevation and over-the-horizon GNSS satellites. GPS signals were acquired for a rising event with PRN 03 on 8 May 2017. Figure 1.9 shows the L1 signal C/N_0

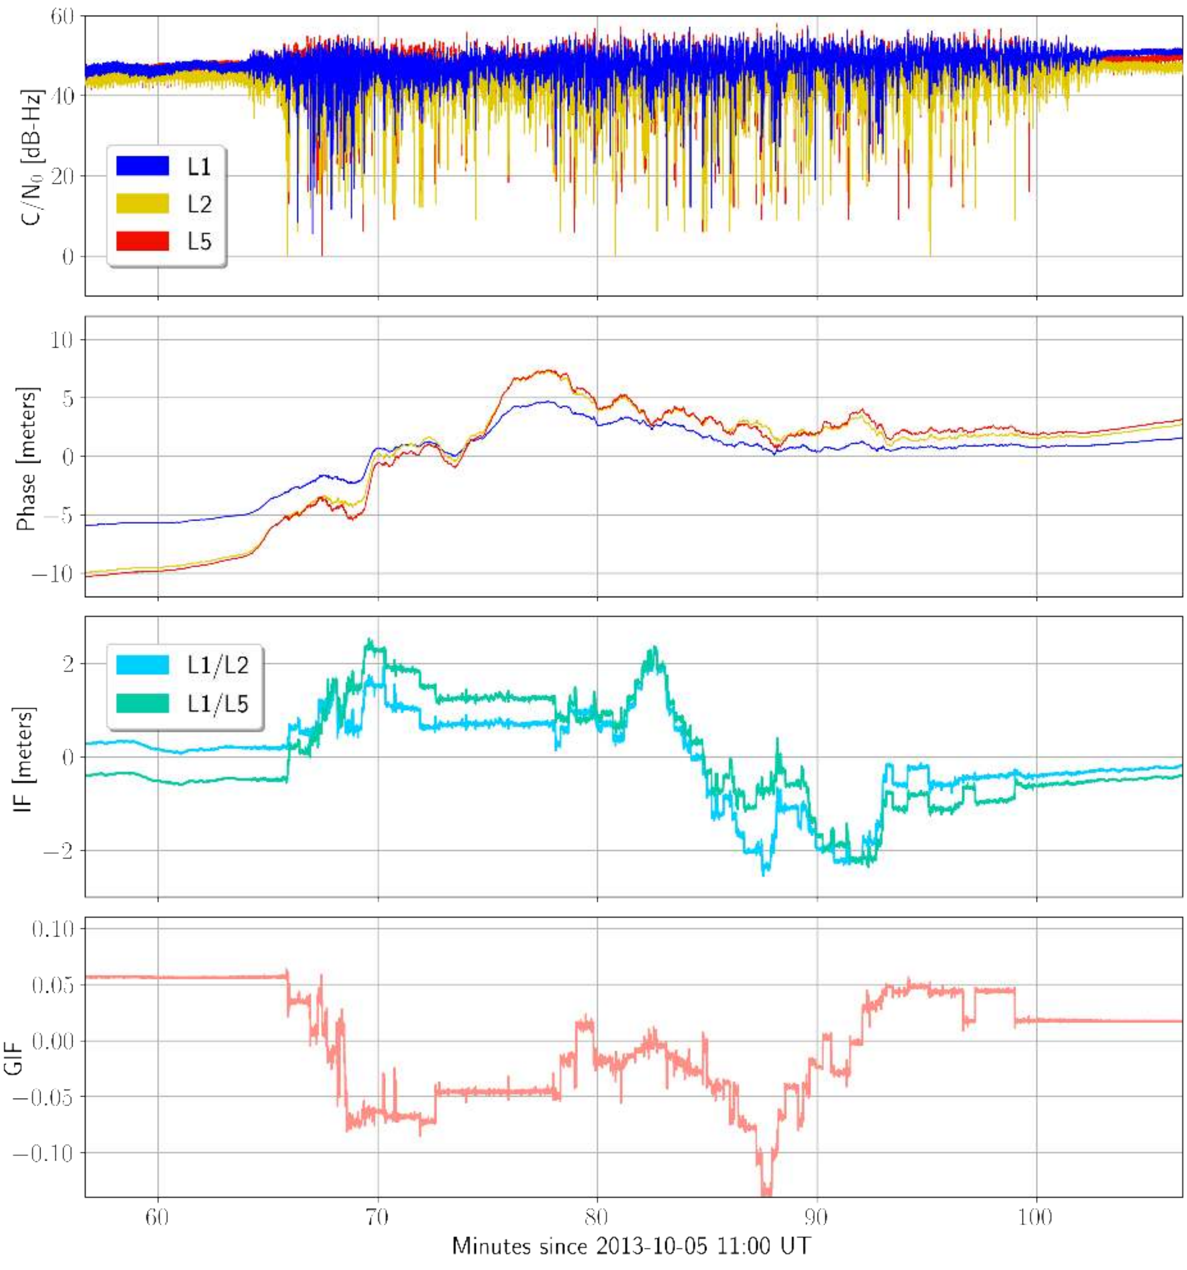

Figure 1.7: Triple-frequency C/N$_0$, phase, and phase combinations corresponding to a signal from GPS PRN 24 that was measured during a scintillation event on 2013-10-05 by a receiver near Hong Kong.

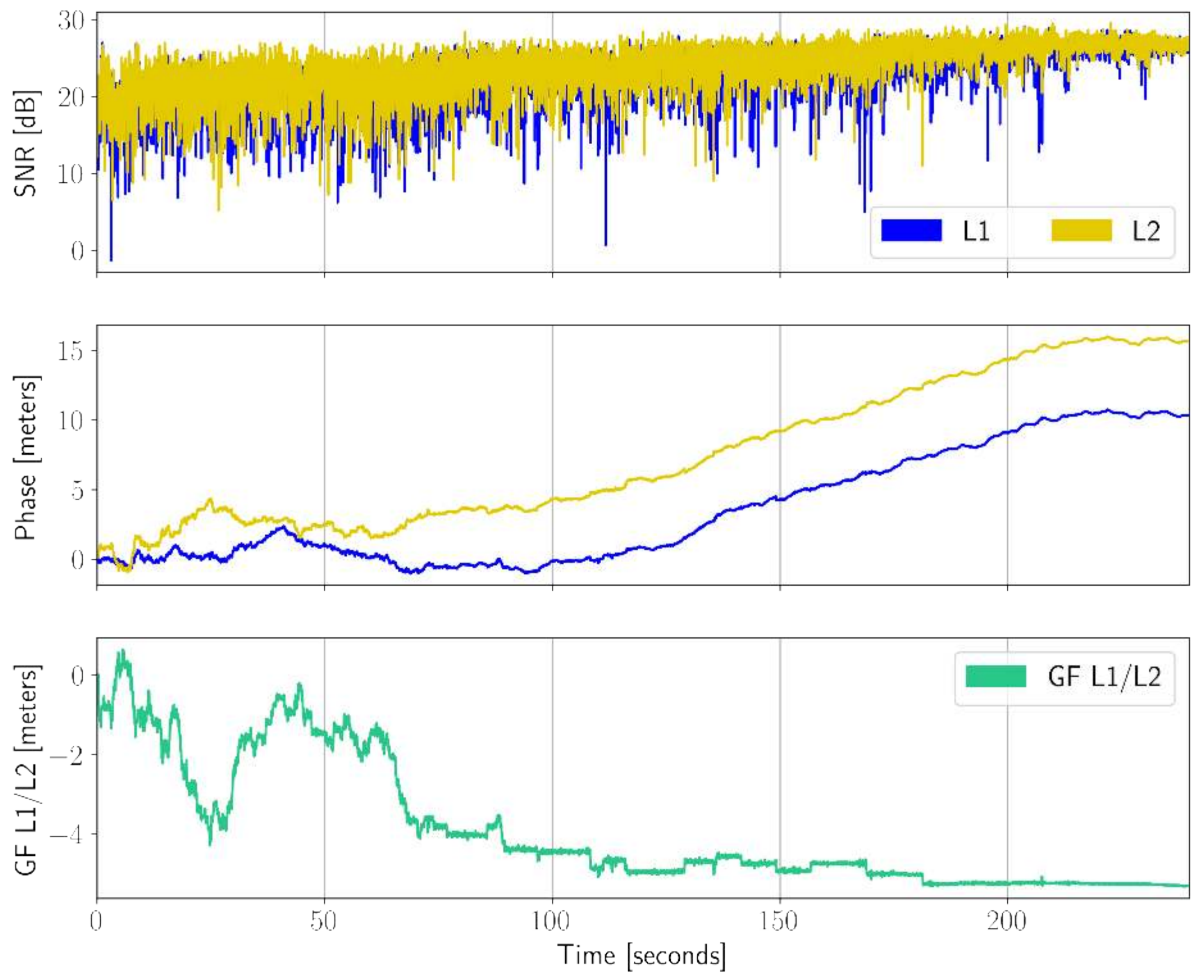

Figure 1.8: Measurements from L1 and L2 GPS signals that were reflected off the ocean surface and received by the low-Earth orbiting Spire radio occultation satellite. Top panel shows C/N_0 fluctuations, middle panel shows the detrended phase, and bottom panel shows the presence of cycle slips using the GF combination of L1 and L2 carrier phases.

and carrier phase measurements. In this case, the navigation data bits were not removed and so the carrier phase measurements are afflicted by half-cycle (instead of full-cycle) slips. In addition to already very low C/N_0 the signal experiences several deep fades during which the signal phase rapidly accrues a bias due to cycle slips.

1.4.3 Datasets: Summary

Looking at each of these examples, a common aspect is the occurrence of many cycle slip over a periods less than a minute. One of the reasons we are able to see these slips in this data is our use of high-rate (100 Hz) carrier phase measurements. Without such measurements, many of the slips in these examples may not be identifiable. Another pattern is the coincidence of cycle slips and signal fading. We explore this relationship more in chapters 2 and 3, and we make explicit use of signal amplitude when performing mitigation based on the backpropagation filter in Chapter 5.

1.5 Summary

In this chapter, we introduced the cycle slip problem along with a model for GNSS measurements. We discussed linear combinations of these measurements and how they have been used for cycle slip mitigation, along with a broader discussion about approaches to cycle slip mitigation. We then introduced some GNSS datasets with harsh signal conditions containing phase fluctuations, noise, and many cycle slips. In the next chapter, we will characterize the occurrence of cycle slips in simulations. In the subsequent chapter, we will take a closer look at the occurrence of cycle slips in the real datasets we just considered in Section 1.4. All of this will be in preparation for our new approach to cycle slip mitigation that we present in Chapter 4 and Chapter 5.

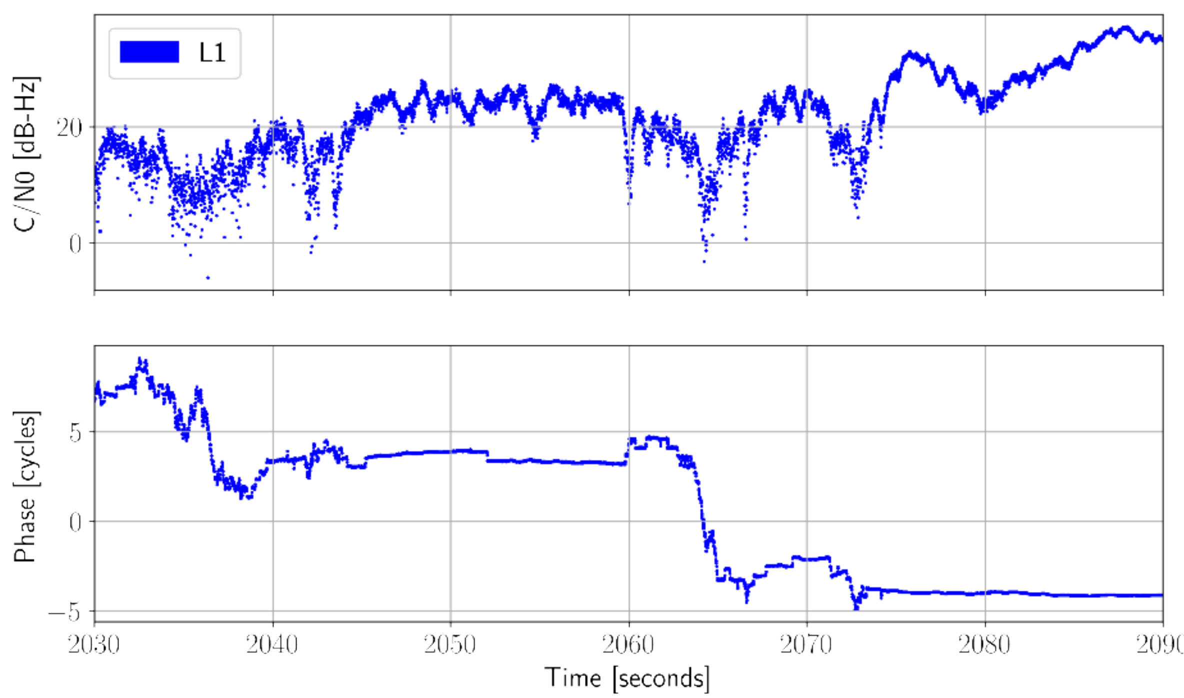

Figure 1.9: Example of L1 signal C/N_0 and carrier phase measurements collected using a high-gain dish antenna on top of Mount Haleakala, Hawaii.

Chapter 2

Cycle Slip Simulation and Characterization

2.1 Introduction

In Section 1.2.1 in Chapter 1, we reviewed the different processes that cause of cycle slips to occur. Two of these causes are the unwrapping of noisy phase measurements and the occurrence of *phase transitions*, both of which occur for signals under harsh conditions. In particular, when a radio signal experiences deep amplitude fading, its phase will likely undergo a simultaneous rapid half-cycle change. In the context of GNSS signals, these events are known as canonical fades and commonly occur during strong scintillation or multipath reflection [58]. Occasionally, the rapid half-cycle phase changes are part of actual full-cycle phase transitions, or cycle slips, which occur over the duration of the fade. It is important to stress that these cycle slips are a product of the propagation environment and are not necessarily due to receiver processing errors. However, regardless of their origins, all cycle slips cause a persistent change in carrier ambiguity. During conditions of strong scintillation or multipath associated with frequent and intense fading, these diffraction- or multipath-induced cycle slips can be the largest and most challenging source of error that corrupts GNSS phase measurements. In order to advance navigation and remote-sensing applications under such harsh conditions, it is important to analyze and characterize this error source.

[0] This chapter is adapted from the paper *GNSS Carrier Phase Cycle Slips Due to Diffractive Ionosphere Scintillation: Simulation and Characterization* published in IEEE Transactions on Aerospace and Electronic Systems [10].

2.1.1 Phase Transitions

As we introduced in Chapter 1, the received and demodulated baseband GNSS signal can be modeled as a complex exponential:

$$s(t) = A(t)\exp(i\phi(t)) + \eta(t) \tag{2.1}$$

When a signal refracts through atmospheric structures or reflects off of rough surfaces, the received signal essentially becomes a superposition of multiple copies of the "direct" signal. The result is called interference or multipath, and it leads to fluctuations in the phase and amplitude of the received signal. In the simplest case, we consider the scenario of interference between two continuous-wave signals:

$$A\exp\left(i\phi\right) + \tilde{A}\exp\left(i\tilde{\phi}\right) = A\exp\left(i\phi\right)\left[1 + \alpha\exp\left(i\Delta\phi\right)\right], \tag{2.2}$$

where $\alpha = \tilde{A}/A$ is the amplitude ratio and $\Delta\phi = \tilde{\phi} - \phi$ is the difference in phase between the two signals. The resulting signal can be interpreted as a nominal signal $A\exp\left(i\phi\right)$ modulated by the complex number $1 + \alpha\exp\left(i\Delta\phi\right)$ [1] . Looking at Equation 2.2, we see that whenever $\Delta\phi$ approaches 180°, and α is large enough, the resulting signal experiences a deep fade. It is illustrative to consider the signal behavior in time-series that are modeled by this equation. The top image in Figure 2.1 depicts two trajectories of the modulation term in the complex plane, with corresponding time-series of the resulting amplitude and phase shown beneath. [17] provides a similar visual descriptions of phase transition occurrence for the case of ionosphere scintillation. In the context of GNSS signals, this behavior is called *canonical fading*.

Scenarios A and B both show the typical canonical fade behavior that occurs when the phasor passes close to zero. A crucial difference is that the signal phasor in scenario B wraps around the origin thereby causing a cycle slip relative to the nominal signal phase. A third example is provided by scenario C in Figure 2.2, where the simulated parameters are the same as in scenario A, but with

[1] Note that, during diffractive scintillation, it is usually possible to define a direct signal phase that is consistent with the signal phase measured in absence of diffraction. For instance, this would be the unperturbed field that is considered in the Rytov approximation of the electromagnetic field propagating through a random structure (c.f. [88]).

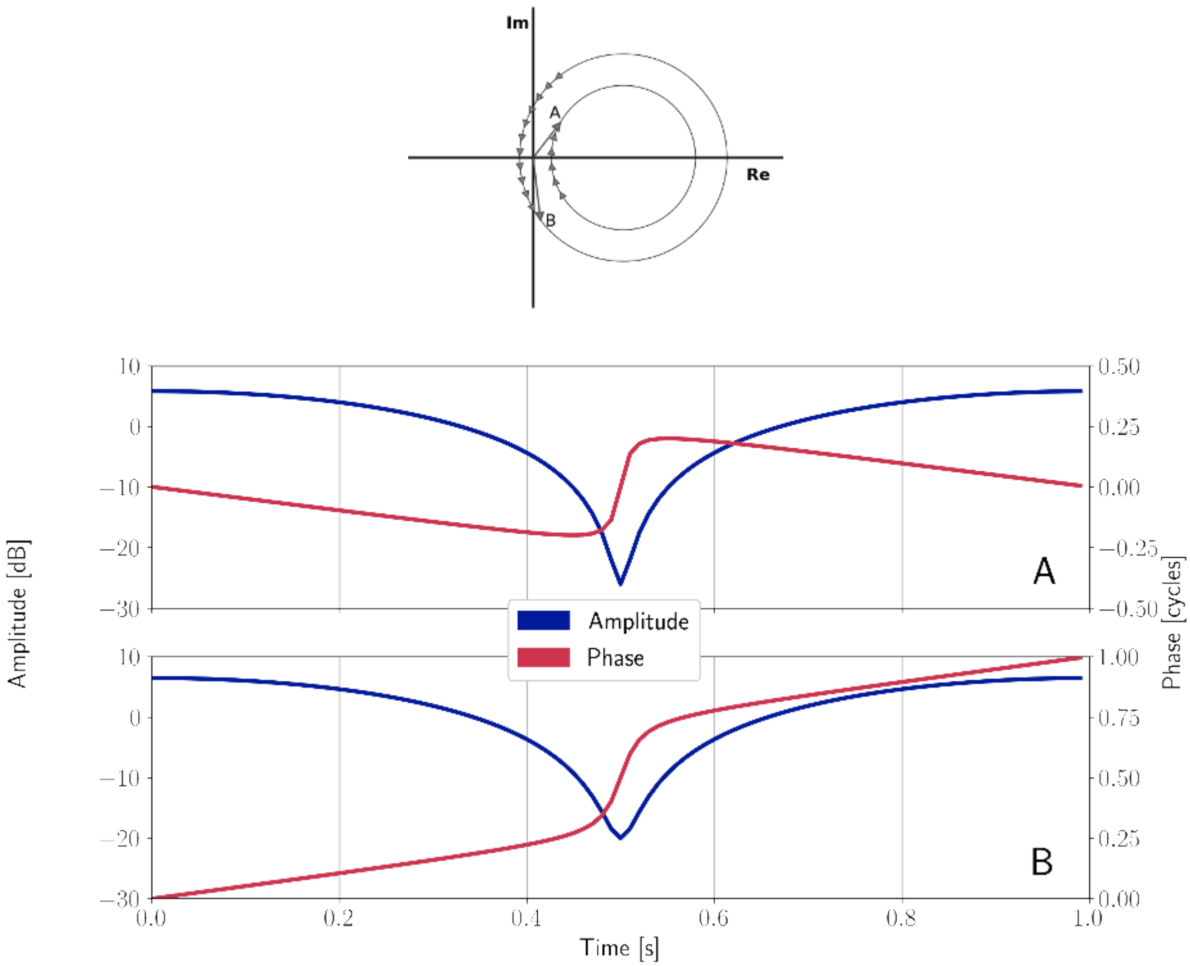

Figure 2.1: The first panel depicts how the trajectories of the complex modulation term in Equation 2.2 for two hypothetical fade scenarios. The second panel shows the corresponding time series in resulting amplitude and phase modulation of the signal.

added noise assuming a C/N_0 of 40 dB-Hz. In this case, the noise causes a cycle slip to occur even though there was no slip in the original noise-free phase trajectory. This simple example illustrates two key points. First, canonical fading will only sometimes correspond to cycle slip occurrence, and whether it does is not clear just from looking at the phase time series from one signal. The phase in scenarios A and B actually look very similar when reflected vertically, but only scenario B contains a cycle slip. Secondly, the effects of noise can induce cycle slips in the measured phase, and its effects will compound with scintillation diffraction. Deeper fading and lower baseline C/N_0 leads to a greater chance for cycle slips.

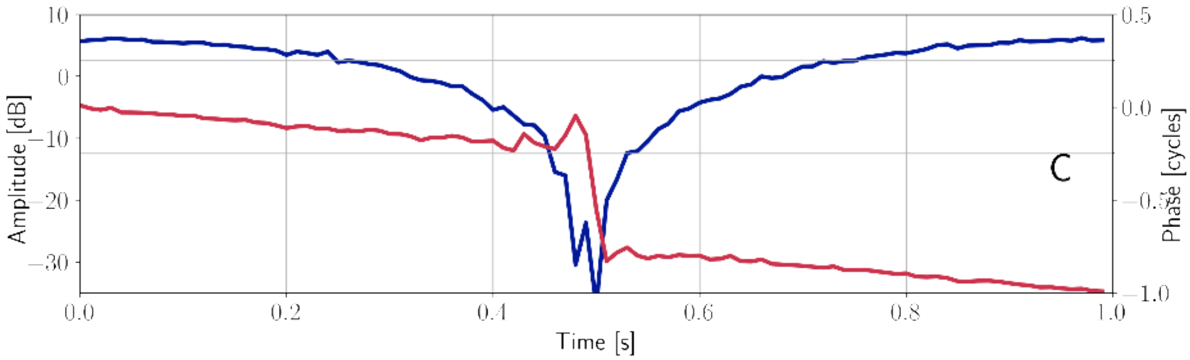

Figure 2.2: Shows simulated phase and amplitude time series with the same parameters as example A in Figure 2.1, but with added noise assuming a C/N_0 of 40 dBHz.

Admittedly, real-world diffraction is quite a bit more complicated than the idealized interference model we just considered. More generally, the motion of the phasor in Figure 2.1 corresponds to a sort of random walk [31], [97]. For example, Figure 2.3 shows the complex phasor along with amplitude and phase measurements corresponding to a more realistic scintillation signal, which we simulate using the method outlined in Section 2.2.1. The figure shows two cases: noise-free and noisy. Two canonical fades occur as the phasor passes close to the origin. For the noise-free case, only the second fade corresponds to a phase transitions, while in the noisy version the first fade corresponds to a cycle slip and the phase transition in the second fade no longer occurs. This more realistic example further illustrates the intricate relationship between phase transitions and noise in these types of signals and how their effects compound to either create or mask the occurrence of cycle slips. This is one of the main topics we investigate in the remainder of this chapter. Moreover, as we discuss towards the end of this chapter and in Section 3.3 of Chapter 3, signal fading in real data does not always exhibit sharp drops in intensity like we see in the simulation examples. Sometimes deep fading is extended or followed by lingering regions of moderate or weak fading [47], [42]. For these scenarios it can be especially unclear whether or not a cycle slip has occurred.

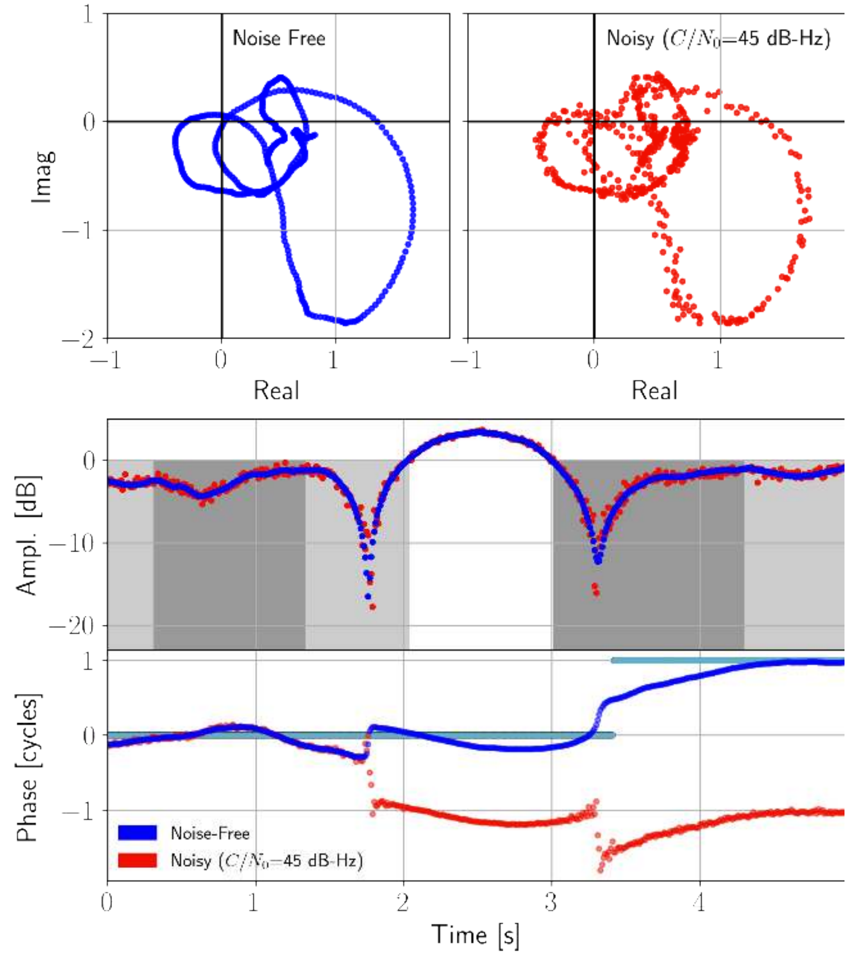

Figure 2.3: Demonstrates two examples of canonical fades. The first two panels shows the diffractive field phasor trajectories in the complex plane. The bottom two panels show the signal amplitude and phase. Each canonical fade corresponds to a decrease of more than 10 dB in signal amplitude along with a rapid half-cycle phase change. The blue and red colors distinguish the noise-free and noisy versions of the signal. The flat blue line in the second panel represents the true phase offset due to a phase transition that occurs in the true phase during the second fade only. Meanwhile, the noisy signal shows how the impact of noise can mask or induce cycle slips that are or are not present in the noise-free signal. The alternate shaded portions of signal amplitude in the second panel demonstrates the definition of a single fade that is used for analysis in this paper.

2.1.2 Previous Work

In this section, we review previous work that has been done to characterize cycle slip occurrence. Traditionally, cycle slips are considered to be the result of a non-linear interaction that occurs in a PLL driven by noisy phase measurements, and there are established results about the mean time to slip occurrence in phase-lock loops (PLLs) [76], [1]. The authors in [36] extend this analysis to cycle slips that occur for GNSS PLLs during strong ionosphere scintillation. They assumed a squaring PLL that is insensitive to half-cycle phase changes (e.g. the Costas PLL) and loosely identified cycle slips as belonging to one of two classes: half-cycle slips due to canonical fading, and half-cycle or greater sized slips due to extended deep fading plus noise. For the latter case, the mean time between slips was assumed to occur at the rate consistent with original analyses for a noise-driven PLL. The authors in [24] applied this analysis to real data from a scintillation campaign and provided real-world context for cycle slip measurement results. They concluded that signal diffraction and associated cycle slips are a dominant error source during strong scintillation. One issue with these analyses is that C/N_0 estimation is poor at low SNR, so modeling cycle slip occurrence based on fading depth does not accurately capture the true probability. More importantly, these analyses did not account for the contribution of phase transitions to cycle slip occurrence during strong scintillation.

The concept of phase transitions in the context of GNSS scintillation was originally identified in phase screen models used to simulate strong scattering from ionosphere irregularities. These models typically use the parabolic approximation to the wave equation to propagate the effects of one or more phase screens representing the cumulative impact of ionosphere irregularities, and have been successfully used to describe multi-frequency observations of real diffractive scintillation [57], [43]. The hybrid technique introduced in [30] is one such model that was later used in [31] and [97] to assess GNSS phase errors. Similarly, [17] used the model introduced in [16] to analyze the impact of diffractive fluctuations on GNSS phase. These authors identify how phase transitions occur when the random walk of the diffractive field phasor (representing the complex multipath modulation imparted on the GNSS signal) wraps around the complex origin. As we saw in the last section, due to the nature of these random walks, phase transitions almost always coincide with deep signal fading and rapid half-cycle phase changes (i.e. canonical fades).

Authors from both [31] and [17] investigated errors in the dual-frequency ionosphere-free combination, specifically acknowledging the impact of phase transitions at stronger scintillation strengths. In [17] they noted that their cumulative impact on phase error could be up to several meters. In [97] they used the random phase screen model to study the occurrence of canonical fading and its relationship to scintillation intensity as characterized by the S_4 index. They used the term "phase transition" to refer to phase variation during canonical fades, regardless of whether and actual full-cycle phase transition occurred. Note, we choose to use the term only to describe the full-cycle transition event (e.g. the second fade shown in blue for Figure 2.3) that introduces biases into the true unwrapped phase. The authors also investigate cycle slip occurrence, but consider any canonical fade (with field amplitude below -10 dB and a rapid half-cycle phase change) as a proxy for the occurrence of cycle slips. This may be applicable for PLLs that are insensitive to half-cycle phase changes, but is not true for general four-quadrant phase tracking. Overall, both sets of authors agree that the effect of phase transitions is an important error to contend with in strong diffractive scintillation.

The terminology surrounding phase transitions thus far has been inconsistent and deserves clarification. While in [31] and [97] they essentially equate phase transitions to cycle slips, the authors in [17] argue that because phase transitions occur "gradually over many samples. . . [they are] distinct from cycle slips, which are abrupt phase changes of a cycle (or more). . . ." In our opinion, both authors are correct. The conventional usage of the term cycle slip would seem to indicate that phase transitions are an error in the receiver measurement and that they occur instantaneously between two measurement epochs, neither of which is true. On the other hand, phase transitions introduce an integer-cycle bias into the phase measurement, just like cycle slips. To geodetic users who use 1 Hz data or slower there may be no way to differentiate between phase transitions and cycle slips caused by other factors, such as receiver artifacts. Moreover, even at high sampling rates there is not necessarily a clear distinction between phase transitions and the noise-induced cycle slips that occur during deep fading, as studied in [36]. Introducing noise can cause new cycle slips to occur or mask the occurrence of cycle slips due to phase transitions. As an example, Figure 2.3 provides a noisy version (in red) of the true phase (in blue) during canonical fading. In this case, the noisy phase ends up slipping during the first fade, for which there was no slip in the true phase, and not slipping during the second fade, for which there was a slip in the true phase. The purpose here is to recognize how the impacts of signal fading and noise can compound on one another to produce many cycle slips. In the end, due to their common origins, we suggest the terminology "diffraction-induced cycle slips" to encompass both phase transitions and noise-induced cycle slips that occur during strong scintillation.

In the remainder of this chapter, we aim to clarify the contribution of phase transitions towards cycle slip occurrence during strong scintillation. To do this, we look at simulations of ionosphere scintillation at three L-band frequencies in order to analyze the impact of diffraction-induced cycle slips on GNSS phase measurements. We use the ionosphere phase screen model from [62], which has been validated using real multi-frequency GNSS data in [43] and [91]. A caveat to this model is that it only really applies to equatorial scintillation scenarios, and while diffractive scintillation is mostly associated with low-latitudes, it can also occur on radio occultation links [56] or occasionally even at high latitudes [71]. Despite this caveat, the methodology applied in this work can be extended to these other domains where signal diffraction occurs, given adequate models. Using these simulations, we look at how often diffraction-induced cycle slips occur for given fading depth and duration, as well as look at the statistics of their cumulative impact on measured phase. The main contributions of this work are the following. First, we quantify the cycle slip occurrence rate dependence on the signal baseline C/N0 and the scintillation intensity. Second, we confirm that the Poisson process is an appropriate model to describe the cycle slip occurrence and identify the Skellam distribution (Equation 2.12) as an accurate representation of the cumulative impact cycle clips on phase measurement errors. Finally, correlations between the cumulative impact of cycle slips for the three carrier frequencies are established for varying levels of scintillation levels and baseline C/N_0. As a whole, the results serve to partially distinguish cycle slips associated with phase transitions and those occurring due to deep fading plus noise. This has implications for previous characterizations of cycle slip occurrence that only consider fading parameters. Moreover, these results can provide a useful baseline for assessing or predicting cumulative phase errors during strong diffractive scintillation.

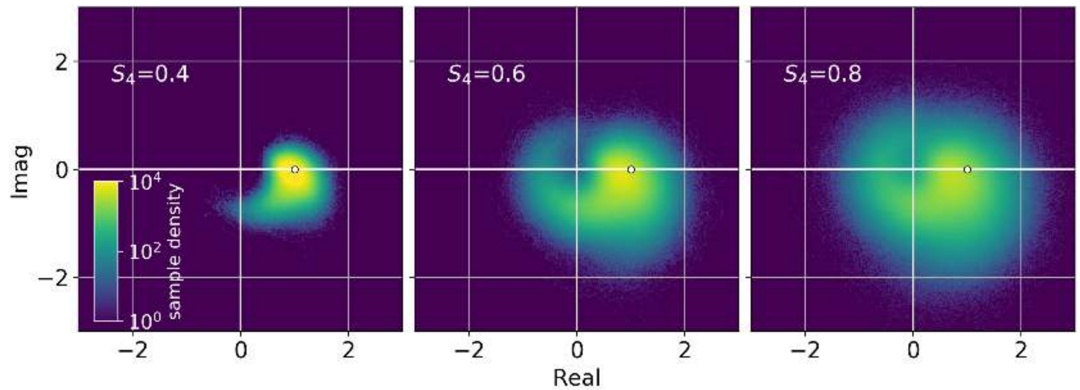

Figure 2.4: The density of the random walk of the diffraction field phasor in the complex plane. The point at 1 (marked by a white dot) corresponds to zero diffractive perturbations, while the origin corresponds to complete destructive interference. The consecutive panels illustrate how, as scintillation strength increases, it may be more probable for the random walk to wrap around the origin and cause cycle slips.

The chapter is structured as follows. In the first section, we introduce the scintillation simulation model and present a model for the impact of noise on cycle slip occurrence. In the subsequent section, we introduce definitions of fading depth and duration and describe a technique for extracting the occurrences of and cumulative effect of cycle slips from the simulated phase. Then we present and interpret results for the dependence of cycle slips on scintillation and fading parameters, as well as on the cumulative error due to cycle slips and its correlation across frequencies. Finally, we summarize the essential results and discuss the challenges and potential solutions associated with diffraction-induced cycle slips.

2.2 Background

2.2.1 Phase Screen Scintillation Model

For this chapter, we consider the following model for the received signal phase, which ignores any initial ambiguities and biases from Equation 1.5 in Chapter 1:

$$\phi_k(t) = 2\pi \frac{\mathcal{G}(t)}{\lambda_k} - \frac{\kappa}{f_k^2}\mathrm{TEC}(t) - \phi_{s,k}(t) \tag{2.3}$$

We introduce the term $\phi_{s,k}$ to specifically represent the diffractive scintillation phase, including phase transitions. Note that we have substituted the term $\beta_k\mathcal{I}$ from our original phase model in Equation 1.5 with the definition of refractive ionosphere phase from Equation 1.7. Simulating scintillation phase amounts to creating realistic time series for the last two terms in this model. In order to simulate ϕ during strong scintillation, we use the equivalent phase screen model from [62]. Given a set of phase screen parameters and a set of signal frequencies, the model generates a consistent set of phase screens $\bar{\phi}_k(t)$ and realizations Ψ_k of the complex field at the receiving antenna. In our simulations, we consider the true (noise-free) phase and amplitude of the signal to be:

$$\phi_k(t) = \text{unwrap}\left(\angle\Psi_k(t)\right) \tag{2.4}$$

$$A_k(t) = \left|\Psi_k(t)\right| \tag{2.5}$$

There is one small caveat to Equation 2.4, since extremely deep fades can lead to incorrect phase unwrapping in the discrete signal. However, this problem can be solved by applying Fourier interpolation as described in [61]. Then, the time series $\phi_k(t)$ effectively simulates Equation 2.3, assuming the term $G(t)$ is zero.

Figure 2.5 shows examples of noise-free simulated scintillation intensity and phase time series for a case of moderately strong equatorial scintillation for all three signal frequencies. Note that the phase is converted to TEC units via:

$$\text{TEC}_k(t) = -\frac{\phi_k(t)}{\lambda_k r_e} \tag{2.6}$$

The similarity between all three signals' phases is consistent with the majority of variation being due to the phase screens, which are approximately identical among the three frequencies. The persistent discrepancies between the phases are due to the diffraction-induced phase transitions.

We can deconstruct the simulator output $\Psi(t)$ into:

$$\Psi = \Psi_s \exp(i\bar{\phi}(t)) \tag{2.7}$$

where

$$\Psi_s = |\Psi| \exp(i\phi_s(t)) = |\Psi| \exp(i(\phi(t) - \bar{\phi}(t))) \tag{2.8}$$

is interpreted as the diffractive perturbation to a normalized, unperturbed field $\Psi_0 = \exp(i\bar{\phi}(t))$. In other words, we treat the phase screen as the refractive phase (i.e. the TEC term from Equation 2.3). Figure 2.4 shows histograms of Ψ_s for various scintillation conditions, revealing how as scintillation strength increases it becomes more probable for the diffraction perturbation to wrap around the origin and cause a phase transition. The colored lines in Figure 2.6 show the diffractive phases $\phi_{s,k}(t)$ corresponding to the simulated example in Figure 2.5. It reveals how the fluctuations and phase transitions due to strong scattering are sometimes correlated and sometimes uncorrelated among signals at different frequencies.

The model specifying the equivalent phase screens requires five parameters: U, ρ/v_{eff}, μ_0, p_1, and p_2. We briefly describe them here, and for an in-depth discussion readers are referred to [62]. The universal strength parameter U determines the magnitude of plasma irregularities relative to the background density, and generally describes the strength of amplitude and phase fluctuations in the resulting scintillation. The parameter ρ/v_{eff} comprises the ratio of the first Fresnel radius ρ to the effective scan velocity v_{eff} through the ionosphere structure, and is correlated to the time between diffractive fluctuations. The remaining parameters μ_0, p_1, and p_2 respectively define the break scale and slopes of a two-component inverse power law spectrum derived from the statistical structure of the ionosphere irregularities. This set of parameters is defined for a given signal frequency and then mapped to physically consistent values for other frequencies. The stochastic phase screen structure is initialized using the same random seed so that signal fluctuations across different frequencies are consistent with propagation through the same random structure.

Our goal is to provide an intuitive analysis of phase transition or cycle slip occurrence under a variety of realistic equatorial scintillation scenarios, which is difficult to do using five model-specific variables. The authors in [91] suggest a reduction of the model to just two parameters – the scintillation index S_4 and decorrelation time τ – that are commonly used in characterizing equatorial scintillation (c.f. [37], [24]). As such, we use S_4 and τ to characterize the different scintillation scenarios that we consider in this study, where S_4 is defined in the usual way as the normalized deviation of signal intensity and decorrelation time is defined as when the intensity autocorrelation drops to $1/e$ of its peak value. Work from [93] shows that for ground stations at low latitudes, μ_0, p_1, and p_2 remain close to nominal values of 0.8, 2.7, and 3.6, respectively. Meanwhile, the parameters U and ρ/v_{eff} are able to capture almost all of the observed variation in scintillation characteristics. They empirically derive the mappings between parameters by simulating scintillation for various values of U and ρ/v_{eff} and then computing the corresponding S_4 and τ from the resulting simulated intensity. The results show that the value of S_4 can be directly related to U, while τ and ρ/v_{eff} have a linear dependence that varies with U. We did this in order to determine the phase screen parameters U and ρ/v_{eff} provided in Table 2.1 that correspond to the approximate range of S_4 and τ values in Table 2.2. The decrease in slope relating τ to ρ/v_{eff} as U increases is in agreement with the analysis by [16] and [24] that showed how higher S_4 generally corresponds to lower τ for realistic equatorial scintillation.

Table 2.1: Phase Screen Parameters

		S1T1	S1T2	S1T3	S2T1	S2T2	S2T3	S3T1	S3T2	S3T3	S4T1	S4T2	S4T3
U	L1	0.30	0.30	0.30	0.60	0.60	0.60	1.20	1.20	1.20	2.00	2.00	2.00
	L2	0.68	0.68	0.68	1.37	1.37	1.37	2.73	2.73	2.73	4.56	4.56	4.56
	L5	0.79	0.79	0.79	1.57	1.57	1.57	3.15	3.15	3.15	5.24	5.24	5.24
$\frac{\rho}{v_{\text{eff}}}$	L1	1.50	0.70	0.30	1.50	0.70	0.30	1.50	0.70	0.30	1.50	0.70	0.30
	L2	1.70	0.79	0.34	1.70	0.79	0.34	1.70	0.79	0.34	1.70	0.79	0.34
	L5	1.74	0.81	0.35	1.74	0.81	0.35	1.74	0.81	0.35	1.74	0.81	0.35

In our analysis, scintillation strength and scintillation fading rate are two of the primary factors associated with cycle slip occurrence due to scintillation, with higher occurrence rates corresponding to stronger scintillation (high S_4) and faster fading (low τ). It is known that phase transitions begin to occur in scintillation signals when S_4 passes around 0.6 [31], which also tends to be the point at which receivers start experiencing cycle slips and loss-of-lock [23]. During solar maximum, receivers at low latitudes, especially near the equatorial ionization anomaly, experience

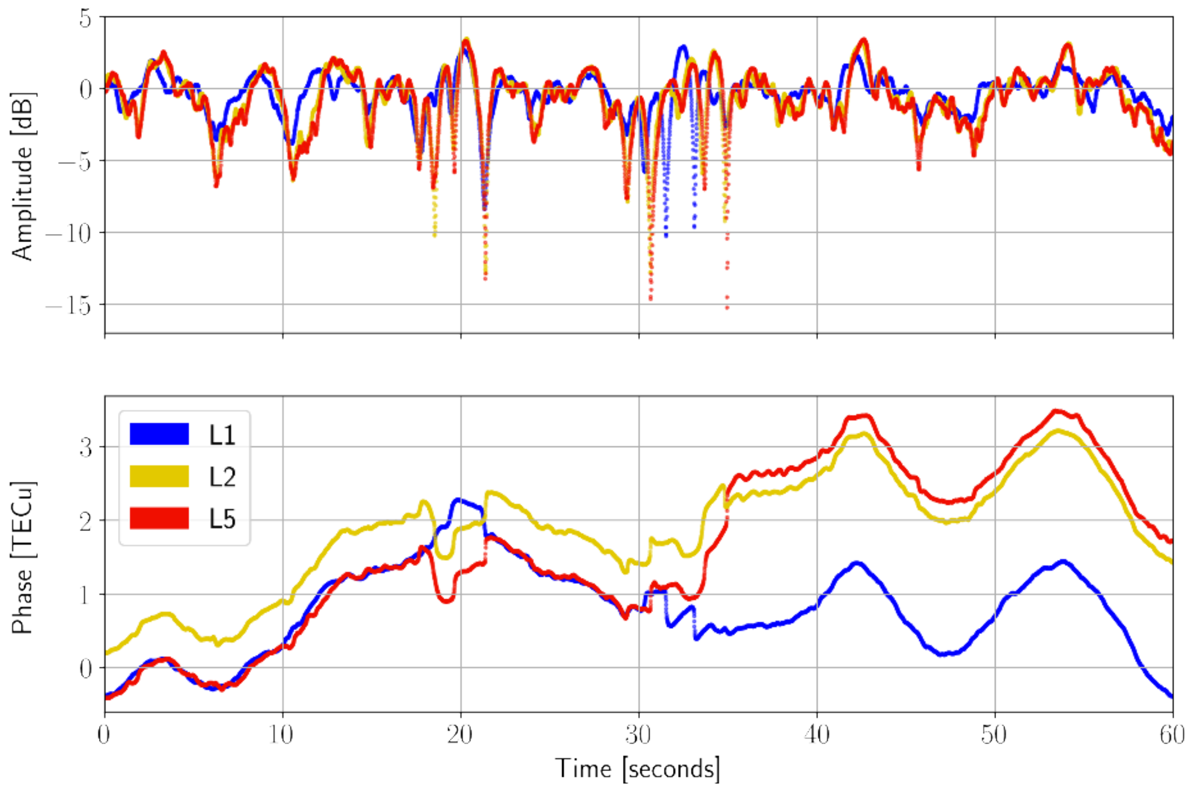

Figure 2.5: An example simulation of scintillation amplitude (first panel) and phase (second panel) corresponding to the L1, L2, and L5 frequencies for a strong scintillation scenario S3T2. Note that phase has been scaled to TEC units.

S_4 values reaching 1 or higher, while τ ranges between $\approx 0.2 - 1.3$ depending on scintillation strength [24], [23]. Based on these analyses, we chose pairs of S_4 and τ parameters for each frequency (L1, L2, L5) corresponding to a set of 12 scintillation scenarios that we consider in this study. Their values are provided in Table 2.2. For convenience, we enumerate these scenarios using the template "S#T#" to indicate the different strengths and decorrelation times. Increasing numbers correspond to an increase in S_4 and decrease in τ, which translates overall to an increase in cycle slips. As an example, "S4T3" corresponds to the scenario with the highest S_4 value and shortest τ value, which should produce the most cycle slips.

2.2.2 Simulating the Impact of Noise

Scintillation strength and fading rate both impact the occurrence rates of cycle slips in strong GNSS scintillation. However, in real-world scenarios the impact of noise cannot be ignored. We consider a baseline carrier-to-noise density ratio $\overline{C/N_0}$ to be the relevant parameter determining noise impact on cycle slip occurrence. $\overline{C/N_0}$ is defined as the carrier-to-noise density ratio that would be measured in the absence of any scintillation diffraction. In order to simulate the impact of different $\overline{C/N_0}$, we generate noise according to:

$$\eta_k(t) \sim \mathcal{N}(0, \sigma^2) \qquad \sigma = \frac{B}{\overline{C/N_0}} \tag{2.9}$$

i.e. $\eta_k(t)$ are assumed to be independent random variables drawn from the zero-mean circular complex normal distribution with variance σ^2. We assume the noise bandwidth B is equal to $1/T$, where $T = 0.01$ seconds is the measurement integration time and sampling interval used in this study. The noisy simulations of phase and amplitude are then given by:

$$\phi_k(t) = \text{unwrap}\left(\angle\left(\Psi_k(t) + \epsilon_k(t)\right)\right) \tag{2.10}$$

$$A_k(t) = \left|\Psi_k(t) + \epsilon_k(t)\right| \tag{2.11}$$

Note that this model for introducing noise into the phase measurements accounts for thermal and processing noise only, and specifically does not account for phase jitter due to the receiver oscillator. This is generally not an important factor for ionosphere monitoring receivers, which are usually designed to have low phase noise [82], but may be an important consideration for other types of receivers and should therefore be studied further. The approach we take should still provide a good diagnostic assessment on the relative contributions of noise to cycle slip occurrence. In our analysis, we consider baseline C/N$_0$ values from 36 to 50 dB-Hz, which covers the range of observed values for low-latitude scintillation events studied in [23].

2.3 Methodology

To reiterate, our aim is to highlight the impact on cycle slip occurrence due to scintillation and fading characteristics (similar to previous work in [24]), as well as the importance of C/N_0 and noise. Additionally, we look to analyze the cumulative error distribution due to cycle slips. To do this requires simulation of a large amount of phase and amplitude time series for each of the different scintillation scenarios. We also need to establish the definition for a signal fade and a method to determine cycle slip occurrence in the simulated output. In this section, we outline our approach to these tasks.

2.3.1 Simulation

For each scintillation scenario in Table 2.2, we generate over 1000 hours of scintillation. A large volume of simulations is necessary since phase transitions or cycle slips can be a sparse occurrence, especially at lower scintillation strengths. We then add noise and compute the measured signal for baseline carrier-to-noise density ratios from 36 to 50 dB-Hz according to Equation 2.10. We used these simulated amplitudes and phases to obtain fade statistics and cycle slip occurrences as outlined in the next two subsections.

2.3.2 Signal Fade Definition

We define a fade segment as a contiguous interval between two local maxima in signal amplitude where the normalized amplitude is also less than 0 dB. Before finding its local maxima, it is useful to remove smaller oscillations in the signal amplitude by applying a low-pass filter with a cutoff around 4 Hz. This helps ensure fade segments are not too short. The alternating shaded portions of the amplitude time series in the first panel of Figure 2.3 illustrate what we consider to be fade segments. Using this definition, we define the fade duration as the segment duration and the fade amplitude as the minimum amplitude over the segment.

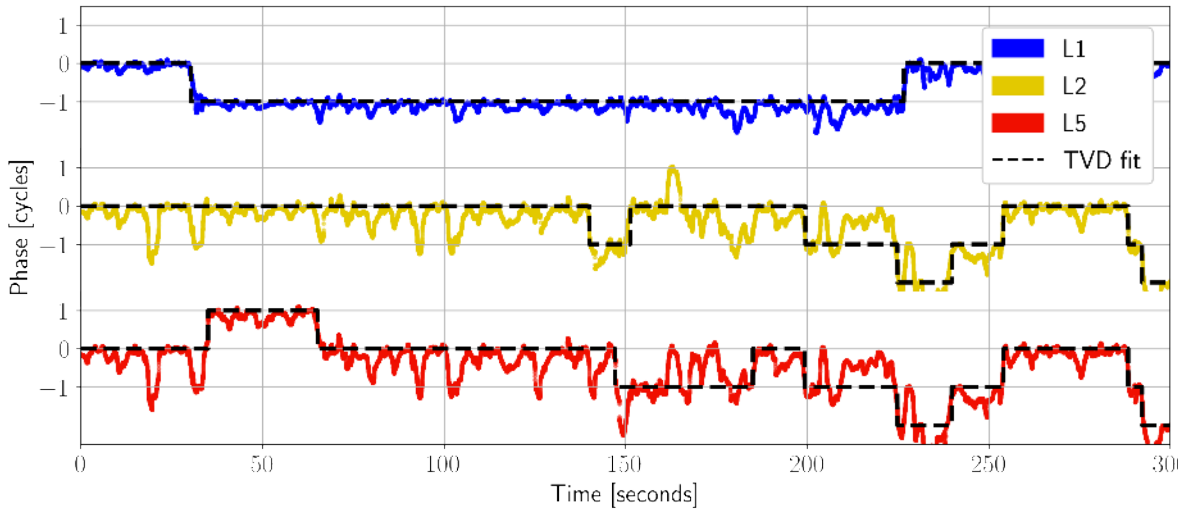

Figure 2.6: Diffraction phase residual (obtained as the full phase observation minus the phase screen component) for the scenario shown in Figure 2.5. Phase is plotted in cycles, and the occurrence of diffraction-induced phase transitions clearly leads to integer-valued biases. The results of applying the TVD fit algorithm are shown with the dashed black line.

2.3.3 Identifying Cycle Slips

Here we introduce a technique that uses total variation denoising (TVD) to determine the occurrence of phase transitions or cycle slips in simulated phase time series. While cycle slip detection in real-world GNSS data is complicated by the presence of unknown phase components, in simulations these other phase components are known and can be removed in order to isolate the diffractive phase and cycle slips. Consider the example of diffractive phase that is illustrated in Figure 2.6. In general, its variations are comprised of constant-mean fluctuations that are highly correlated across frequencies and contain sparse integer jumps in signal mean due to phase transitions or cycle slips. TVD is an optimization technique commonly used to estimate signals with sparse derivative components in the presence of noise [66]. Such is the case for diffractive phase, when we treat the cumulative bias due to cycle slips as the desired signal and the remaining diffractive fluctuations as noise.

We apply a weighted TVD fit (c.f. [4], [5]) to the diffractive phase ϕ_s, where the weights are chosen in such a way that ensures only one "jump" occurs per fade. This allows us to associate each phase transition or cycle slip with a fading amplitude and duration. We achieve this using a penalty weight that is 1 at the minima of fade segments and sufficiently large (e.g. 10,000) at all other times. This penalty applies to the magnitude of the derivative of the TVD fit so that the algorithm only estimates a jump at times corresponding to minima in signal amplitude. We remove an overall negative bias in the diffractive phase (modulo 1 cycle), which can be understood as the "tail" of the histograms shown in Figure 2.4. Finally, we quantize the fit by rounding to the nearest integer. An example of the results of this process are shown with dashed black lines in Figure 2.6. The occurrences of individual phase transitions or cycle slips are obtained by finding times where the difference of this signal is non-zero. Since each jump is constrained to occur at a fade minimum, we can associate each phase transition or cycle slip with a single fade and its corresponding amplitude and duration traits.

Table 2.2: Scintillation Parameters

		S1T1	S1T2	S1T3	S2T1	S2T2	S2T3	S3T1	S3T2	S3T3	S4T1	S4T2	S4T3
	L1	0.36	0.36	0.36	0.50	0.51	0.51	0.70	0.70	0.70	0.85	0.86	0.86
S_4	L2	0.52	0.53	0.53	0.73	0.73	0.72	0.93	0.92	0.92	1.00	1.02	1.02
	L5	0.56	0.56	0.56	0.77	0.77	0.76	0.96	0.95	0.94	1.02	1.04	1.04
	L1	1.29	0.61	0.26	1.20	0.56	0.24	1.07	0.49	0.21	0.94	0.44	0.19
τ	L2	1.32	0.62	0.27	1.18	0.55	0.24	1.00	0.47	0.20	0.86	0.40	0.17
	L5	1.32	0.62	0.27	1.17	0.54	0.24	0.99	0.46	0.20	0.84	0.39	0.17

2.4 Results and Analysis

Here we analyze and discuss the fade and cycle slip occurrences obtained using the methods outlined in the last section. We first describe cycle slip rates with respect to fading amplitude and duration, showing how these rates change for relevant S_4, τ, and $\overline{C/N_0}$ parameters. We then shift our focus to the distribution of cumulative cycle slip error, providing tables of cycle slip rates and correlation coefficients of error among two different signal frequencies for the different simulation scenarios.

2.4.1 Cycle Slip Occurrence Versus Fade Depth and Duration

Figure 2.7 shows the fraction of fades at different amplitudes that contain cycle slips. The scintillation strength increases for each panel from left to right, and we can see corresponding changes in the dashed lines that show the pdf of fade amplitudes. As is well-established, the stronger scintillation leads to deeper fades [42]. The solid lines show what fraction of fades with a given amplitude experience cycle slips, with each line corresponding to a different $\overline{C/N_0}$ value. Here we see how a lower $\overline{C/N_0}$ leads to an increase in the number of fades with cycle slips at deeper fading amplitudes. This corresponds with the larger influence of noise at lower signal amplitudes. We also observe how as scintillation strength increases the cycle slip occurrence per fade increases at all fading depths and for all $\overline{C/N_0}$ values. However, at shallower fading depths this increase is proportionally greater and is present in all $\overline{C/N_0}$ values equally. Since noise cannot induce cycle slips as easily at low fading amplitudes, this implies that phase transitions must be the cause.

Figure 2.8 shows the fraction of fades of different durations that contain cycle slips. Similar to Figure 2.7, the pdf of fade durations is provided with dashed lines. The decorrelation time decreases in each panel from left to right, highlighting how longer fades are less frequent for shorter decorrelation times. Again, the solid lines show what fraction of fades contain cycle slips. There is a general trend of increasing cycle slip fraction for longer fade durations, reflecting how when a signal spends more time at low amplitude, there is a greater chance that random perturbation due to noise will cause cycle slips. We see this in the leftmost panel, where a larger τ value leads to longer fades, and increased noise leads to increased cycle slips at these longer fade durations. However, as fading rate increases (with a decrease in τ) the fade durations become shorter and the dependence of cycle slip occurrence on $\overline{C/N_0}$ diminishes.

2.4.2 Cumulative Cycle Slip Occurrence

Let χ denote the cumulative occurrence of cycle slips measured over 5-minute intervals. We compute χ for each independent simulation interval, for each signal frequency, and for all

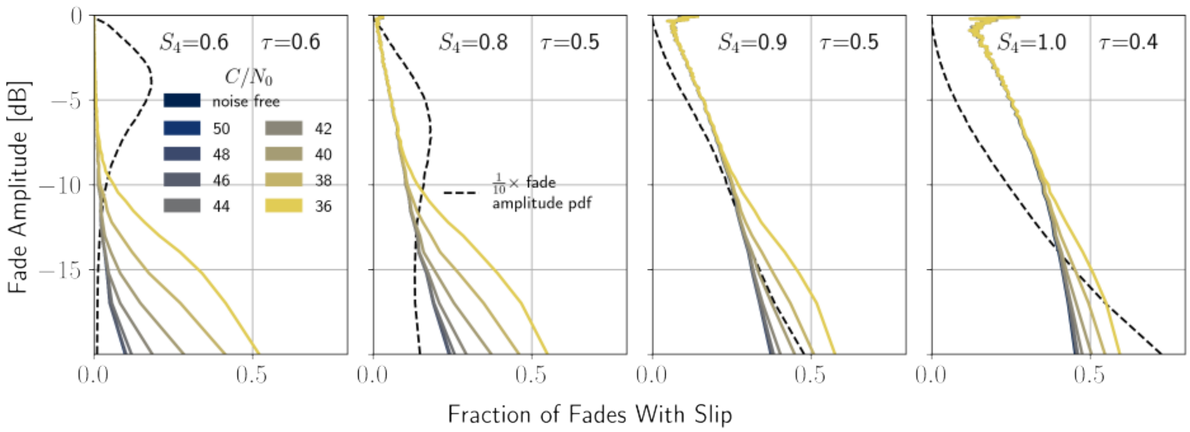

Figure 2.7: Shows the rate of cycle slips per fade as a function of fade amplitude (i.e. fading depth) for each baseline C/N$_0$. The scintillation parameters demonstrate increasing scintillation strength with $S_4 = 0.6, 0.8, 0.9, 1.0$ and $\tau = 0.6, 0.5, 0.5, 0.4$ for the respective panels (which corresponds to L5 signal for scenarios S1T2, S2T2, S3T2, and S4T2). A histogram estimate of the fade amplitude pdf is also shown in the dashed line. In all cases, lower baseline C/N$_0$ leads to an increase in cycle slips at deeper fades, while increase in scintillation strength leads to an increase in cycle slips for all $\overline{C/N_0}$ values and especially at shallower fade amplitudes.

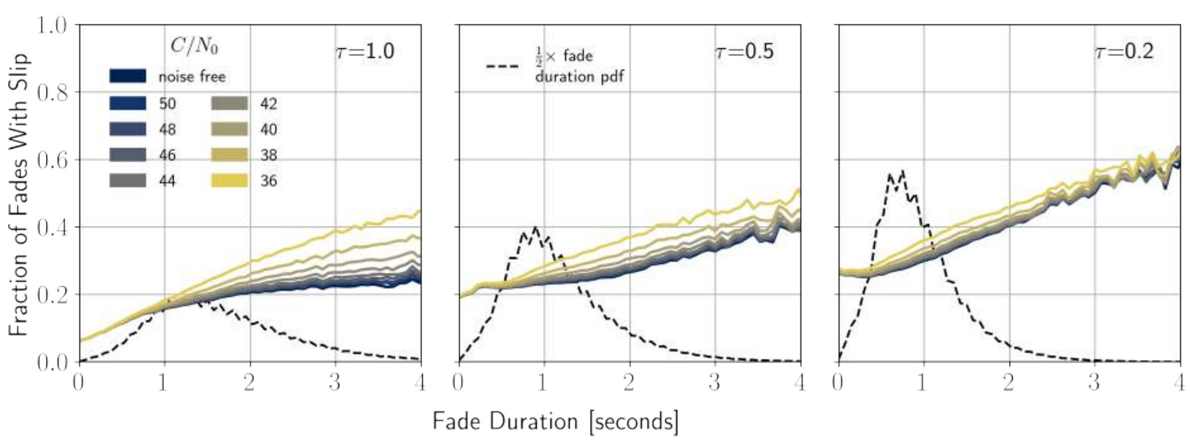

Figure 2.8: Shows the rate of cycle slips per fade as a function of fade duration for each baseline C/N$_0$. The scintillation parameters are $S_4 \approx 0.95$ and $\tau \in \{1.0, 0.5, 0.2\}$ for each of the respective panels (which corresponds to L5 signal for scenarios S3T1, S3T2, and S3T3). A histogram estimate of the fade duration pdf is also shown in the dashed line. We generally see that lower baseline C/N$_0$ leads to an increase of cycle slip rates, but especially so for low fade rates ($\tau = 1$) and long fade durations. At higher fade rates (lower τ) all $\overline{C/N_0}$ values show similar fade rates at all durations.

scintillation scenarios (including different $\overline{C/N_0}$ values). For the volume of simulations used in this study, this results in around 12,500 samples for each combination of signal, $\overline{C/N_0}$, and scintillation scenario. We histogram these samples in order to estimate the probability mass function (pmf) of cumulative cycle slip occurrence over 5 minutes. Figure 2.9 shows one scenario of the resulting pmf on three frequencies. These pmfs follow the Skellam distribution, which is expected when we assume cycle slip occurrence is a Poisson process [70]. The authors in [24] also consider a Poisson process model for cycle slip occurrence, however they do not consider the cumulative effect of positive and negative cycle slips. The Skellam distribution describes the difference between two Poisson processes, and in this case allows us to interpret positive and negative cycle slips as individual Poisson processes with rate parameters μ_+ and μ_- [70]. We found the mean of the distribution to be essentially zero for all scintillation parameter choices, indicating that the rates for positive and negative cycle slip occurrence are equal, i.e. $\mu_+ = \mu_-$. The resulting symmetric Skellam distribution can be written:

$$p(n; \mu) = e^{-\mu} I_{|n|}(\mu) \tag{2.12}$$

where n is the independent variable reflecting cumulative cycle slips, $I_{|n|}$ is the modified Bessel function of the first kind, and $\mu = 2\mu_- = 2\mu_+$ is interpreted as the overall rate of cycle slip occurrence. We fit symmetric Skellam distributions to cumulative cycle slip histograms for each signal and scintillation scenario considered in this study. Table 2.3 shows the cycle slip rate values obtained via this process. Figure 2.10 visually summarizes the cycle slip rate results for the L5 frequency. Plugging the appropriate rate parameter into the Skellam distribution provides the distribution of cumulative phase error due to cycle slips over arbitrary time periods.

As an attempt to characterize the joint distribution of cumulative cycle slip occurrence, we can compute the Pearson correlation coefficient for cumulative cycle slip error among each frequency pair. Since the variance of a symmetric Skellam distribution is equal to its overall rate parameter μ, we compute the correlation coefficient as:

$$\frac{\mathrm{cov}(\chi_k, \chi_l)}{\sqrt{\mu_k \mu_l}} \tag{2.13}$$

where k and l correspond to the two different frequencies. Table 2.4 provides the Pearson correlation coefficients for each frequency pair derived from the same χ samples used above. For lower S_4 values,

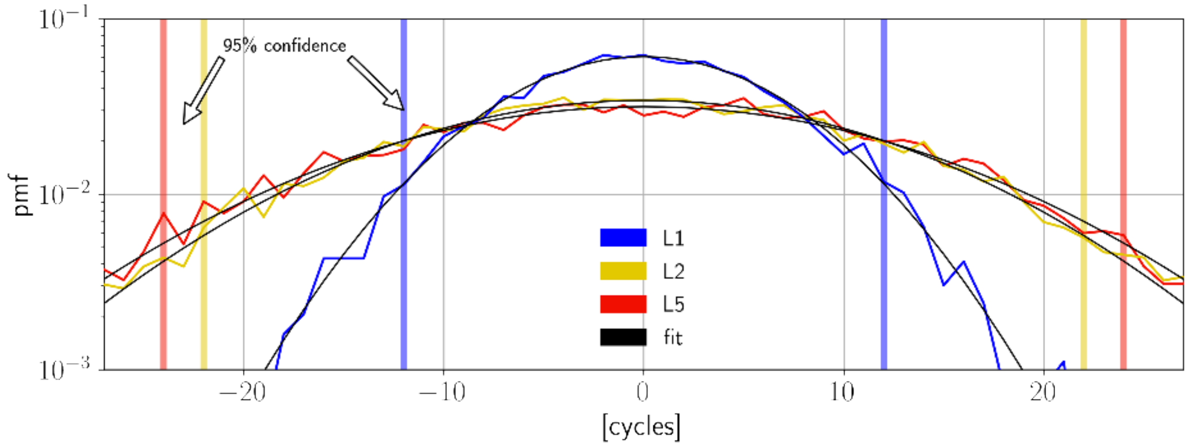

Figure 2.9: Examples of pmf estimates obtained from histograms of the cumulative cycle slip error over 5-minute windows. This particular case shows the three frequencies for scenario S4T2 given in Table 2.2, and for a baseline C/N_0 of 36 dB-Hz. The black lines show the results of fitting a symmetric Skellam distribution to the pmfs.

there are some instances where cycle slip occurrence was too low to compute valid correlation coefficients, and those entries are marked with dashes. In general, the correlation coefficients show larger values for less noise (higher $\overline{C/N_0}$) and stronger scintillation. The L2/L5 signal pair shows the highest correlation, followed by L1/L2 then L1/L5, which is expected behavior according to their frequency ratios. At low $\overline{C/N_0}$, the correlation is practically zero for moderate scintillation, where phase transition behavior mostly coincides with deep fades and is susceptible to the random influence of noise. However, for stronger scintillation, the correlation increases to similar values as obtained for the noise-free case. This is in agreement with our earlier observations that more phase transitions occur with shallow fading amplitudes, and so noise has less influence to alter phase transition behavior. While the correlation coefficient provides some useful information about how similar the cumulative errors will be between two frequencies, it is important to stress that it does not necessarily describe the full joint distribution of errors for the two frequencies.

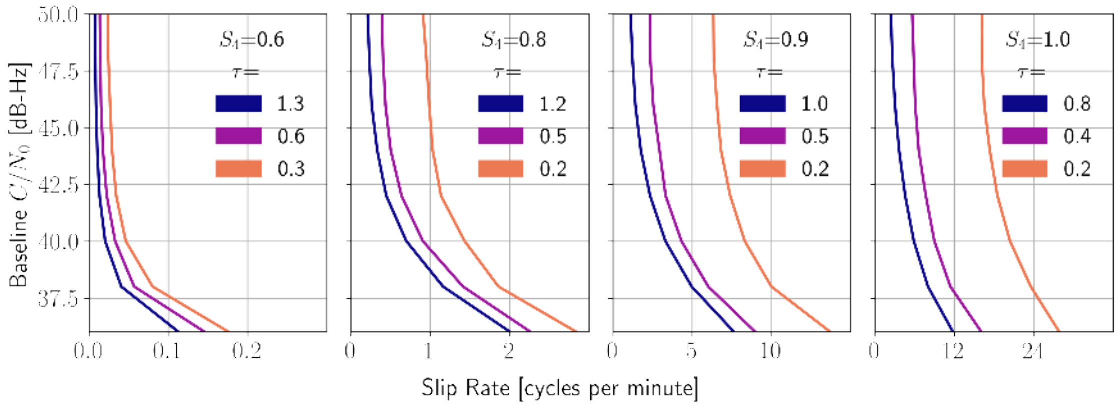

Figure 2.10: Cycle slip rate dependence on baseline C/N_0, decorrelation time, and S_4 index. The cycle slip rates increase as baseline C/N_0 decreases, as decorrelation time (τ) decreases, and as scintillation strength (S_4) increases. The rates are derived from fitting the rate parameter of symmetric Skellam distributions to empirical distributions of cumulative cycle slip error over 5 minutes. The scintillation parameters increase in strength with $S_4 = 0.6, 0.8, 0.9, 1.0$ for the respective panels, while the different colored lines indicate different decorrelation times $\tau \approx 0.2 - 1.2$. These scintillation parameters correspond to the different scenarios given for the L5 signal in Table 2.2, and demonstrate the general relationship between cycle slip rate and different factors given in Table 2.3.

2.5 Conclusion

In this chapter, we used simulations of strong equatorial scintillation to analyze the compounding effects of diffraction and noise on cycle slip occurrence. The key confirmations and findings from the results of our simulations are as follows. 1) As scintillation strength increases there is a notable increase in cycle slip occurrence rates for all baseline C/N_0 values at shallower fade amplitudes. This suggests that phase transitions are a dominant mechanism for cycle slips in this regime (as opposed to noise-induced), and that fading depth is not a consistent indicator for cycle slip probability as scintillation strength increases. 2) The Skellam distribution does a good job of describing the cumulative impact of cycle slips on phase measurement errors, suggesting that cycle slip occurrence can indeed be modeled as a Poisson process. The distribution was symmetric, supporting the observation that positive and negative phase transitions or cycle slips are equally likely. We provided cycle slip rate parameters that can be used to describe error distributions for a wide variety of equatorial scintillation conditions. 3) We provided correlation coefficients for the cumulative impact of cycle slips on different frequency pairs. The correlation increases as scintillation strength increases, which also suggests that correlated diffractive phase transitions (and not just canonical fading plus noise) play an increasingly important role in cycle slip occurrence for stronger scintillation. We want to emphasize that the cycle slip rate and correlation results presented here do not consider the interaction between noise and tracking loop implementations, but they should provide a good baseline for assessing or predicting cycle slip error in unfiltered phase measurements.

Table 2.3: Cycle Slips Per Minute

$\overline{C/N_0}$	[dB-Hz]	S1T1	S1T2	S1T3	S2T1	S2T2	S2T3	S3T1	S3T2	S3T3	S4T1	S4T2	S4T3
Noise Free	L1	0.00	0.00	0.00	0.00	0.00	0.00	0.08	0.17	0.40	0.48	0.98	2.55
	L2	0.00	0.00	0.01	0.11	0.22	0.53	0.73	1.69	4.32	1.80	4.48	12.09
	L5	0.01	0.01	0.02	0.18	0.34	0.83	0.95	2.13	5.95	2.20	5.15	14.76
50	L1	0.00	0.00	0.00	0.00	0.00	0.00	0.10	0.19	0.43	0.54	1.10	2.70
	L2	0.00	0.01	0.01	0.13	0.25	0.57	0.90	1.83	4.60	2.09	4.88	13.67
	L5	0.01	0.01	0.02	0.21	0.39	0.91	1.12	2.34	6.34	2.44	5.66	16.24
48	L1	0.00	0.00	0.00	0.00	0.00	0.00	0.10	0.20	0.45	0.58	1.14	2.79
	L2	0.00	0.01	0.01	0.14	0.26	0.58	1.00	1.88	4.63	2.24	4.96	13.85
	L5	0.01	0.01	0.02	0.24	0.41	0.95	1.25	2.37	6.36	2.67	5.90	16.19
46	L1	0.00	0.00	0.00	0.00	0.00	0.00	0.12	0.20	0.46	0.68	1.20	2.84
	L2	0.00	0.01	0.01	0.16	0.28	0.60	1.15	2.07	4.79	2.51	5.16	14.36
	L5	0.01	0.02	0.03	0.26	0.44	0.99	1.43	2.55	6.57	3.06	6.11	16.55
44	L1	0.00	0.00	0.00	0.00	0.00	0.01	0.14	0.23	0.49	0.83	1.35	2.94
	L2	0.00	0.01	0.01	0.21	0.32	0.64	1.43	2.32	5.20	3.04	5.47	15.32
	L5	0.01	0.02	0.03	0.33	0.50	1.03	1.77	2.95	6.82	3.68	6.65	17.38
42	L1	0.00	0.00	0.00	0.00	0.00	0.01	0.17	0.30	0.54	1.06	1.57	3.16
	L2	0.01	0.01	0.01	0.29	0.41	0.72	1.88	2.77	5.52	3.78	6.44	16.41
	L5	0.01	0.02	0.03	0.45	0.64	1.14	2.35	3.33	7.40	4.60	7.72	18.54
40	L1	0.00	0.00	0.00	0.01	0.01	0.01	0.26	0.40	0.64	1.53	1.93	3.60
	L2	0.01	0.01	0.02	0.44	0.56	0.89	2.64	3.55	6.38	5.09	7.69	18.64
	L5	0.02	0.03	0.05	0.70	0.91	1.43	3.32	4.33	8.31	5.91	9.01	20.45
38	L1	0.00	0.00	0.00	0.01	0.01	0.02	0.47	0.62	0.88	2.28	2.81	4.56
	L2	0.02	0.03	0.03	0.74	0.86	1.22	3.81	4.78	7.80	7.42	9.72	21.18
	L5	0.04	0.06	0.08	1.16	1.41	1.86	4.97	6.02	9.99	8.05	11.43	23.58
36	L1	0.00	0.00	0.00	0.03	0.04	0.05	0.95	1.07	1.36	3.60	4.35	6.30
	L2	0.05	0.07	0.09	1.39	1.51	1.91	6.28	7.28	10.89	11.65	13.68	25.61
	L5	0.11	0.15	0.18	2.01	2.27	2.84	7.64	8.98	13.70	11.82	16.12	27.90

Table 2.4: Cumulative Cycle Slip Error Correlation Coefficient

$\overline{C/N_0}$	[dB-Hz]	S1T1	S1T2	S1T3	S2T1	S2T2	S2T3	S3T1	S3T2	S3T3	S4T1	S4T2	S4T3
Noise Free	L1, L2	-	-	-	0.04	0.03	0.02	0.04	0.08	0.16	0.08	0.11	0.26
	L1, L5	-	-	-	0.02	0.00	0.00	0.05	0.07	0.14	0.05	0.12	0.25
	L2, L5	0.40	0.45	0.43	0.55	0.56	0.58	0.59	0.61	0.68	0.59	0.65	0.72
50	L1, L2	-	-	-	0.04	0.05	0.01	0.05	0.10	0.17	0.07	0.14	0.28
	L1, L5	-	-	-	0.01	0.01	0.02	0.05	0.07	0.16	0.05	0.12	0.27
	L2, L5	0.32	0.40	0.40	0.42	0.48	0.54	0.49	0.57	0.66	0.51	0.62	0.71
48	L1, L2	-	-	-	0.04	0.04	0.02	0.05	0.09	0.17	0.06	0.14	0.28
	L1, L5	-	-	-	0.02	0.01	0.02	0.04	0.07	0.15	0.06	0.11	0.27
	L2, L5	0.30	0.37	0.40	0.37	0.45	0.52	0.45	0.54	0.65	0.47	0.60	0.70
46	L1, L2	-	-	-	0.03	0.06	0.03	0.05	0.09	0.16	0.06	0.13	0.29
	L1, L5	-	-	-	0.00	0.02	0.02	0.05	0.06	0.16	0.05	0.10	0.27
	L2, L5	0.27	0.33	0.37	0.31	0.40	0.50	0.38	0.50	0.64	0.41	0.56	0.70
44	L1, L2	-	-	0.04	0.04	0.06	0.03	0.04	0.07	0.16	0.05	0.13	0.29
	L1, L5	-	-	0.00	0.01	0.03	0.01	0.06	0.06	0.15	0.06	0.11	0.27
	L2, L5	0.19	0.27	0.33	0.23	0.34	0.46	0.30	0.44	0.61	0.34	0.52	0.68
42	L1, L2	-	-	0.03	0.03	0.06	0.03	0.04	0.07	0.16	0.06	0.12	0.29
	L1, L5	-	-	0.00	0.00	0.00	0.03	0.04	0.05	0.15	0.04	0.09	0.28
	L2, L5	0.17	0.21	0.28	0.17	0.27	0.40	0.25	0.38	0.58	0.27	0.46	0.66
40	L1, L2	-	-	0.03	0.02	0.04	0.03	0.03	0.06	0.16	0.04	0.13	0.28
	L1, L5	-	-	0.02	0.01	0.03	0.03	0.03	0.05	0.16	0.03	0.08	0.27
	L2, L5	0.12	0.13	0.18	0.10	0.19	0.34	0.17	0.29	0.52	0.18	0.39	0.64
38	L1, L2	0.00	0.00	0.01	0.03	0.01	0.03	0.02	0.05	0.15	0.04	0.13	0.29
	L1, L5	0.00	0.00	0.03	0.02	0.00	0.03	0.01	0.03	0.15	0.02	0.09	0.28
	L2, L5	0.07	0.06	0.11	0.07	0.12	0.25	0.09	0.19	0.45	0.14	0.32	0.60
36	L1, L2	0.00	0.03	0.01	0.00	0.01	0.01	0.01	0.05	0.14	0.04	0.10	0.29
	L1, L5	0.00	0.00	0.02	0.03	0.03	0.02	0.00	0.05	0.16	0.01	0.11	0.29
	L2, L5	0.06	0.05	0.04	0.02	0.05	0.17	0.07	0.14	0.39	0.09	0.27	0.56

Chapter 3

Occurrence of Challenging Cycle Slips in Real-World Data

3.1 Introduction

In Chapter 2, we provided a characterization of cycle slips that occur due to the combined effects of phase transitions and noise. In this chapter, we will take a closer look at cycle slips occurrence due to these effects in actual GPS scintillation data sets. In particular, we look at the Ascension Island and Hong Kong scintillation datasets containing fade-induced cycle slips that challenge typical detection and repair algorithms. For the first dataset, we observe carefully detrended phase records from multi-frequency GNSS signals and attempt to manually interpret the occurrence of cycle slips. This process provides valuable insight into the challenges and limitations of cycle slip detection. In our work, we utilize high-rate (100 Hz) phase and intensity measurements to interpret the many and frequent cycle slip occurrences due to diffractive ionosphere scintillation. We show how these slips can occur over the course of 0.5-1 seconds or longer and how sometimes several cycle slips can occur in the span of less than a minute. With the second dataset, we apply and assess the performance of two cycle slip mitigation techniques from [20] and from [12], which is an adapted version of the cycle slip filtering algorithm originally presented in [85]. Although we are able to manually identify many cycle slips using these high-rate measurements, as we will show, these techniques struggle to estimate slips correctly.

In the next section, we discuss the geodetic detrending technique we apply to the real-world scintillation datasets in order to reduce the effect of non-dispersive phase components. We also

[0] The content of this chapter is largely adapted from the conference papers in [11], [9], and [12].

discuss additional background on different phase combinations and their varying sensitivities to cycle slips. In Section 3.3, we introduce motivating examples of deep-fade-induced phase transitions during strong scintillation and multipath. In the Section 3.4, we introduce a triple-frequency cycle slip detection algorithm from [20] and assess its performance on the Hong Kong scintillation dataset. In the Section 3.5, we introduce the algorithm from [12]. Finally, we provide concluding remarks and summarize the main challenges associated with cycle slips under harsh signal conditions in Section 3.6. Overall, the purpose of this chapter is to serve as motivation for our development of yet another approach to cycle slip mitigation in chapters 4 and 5.

3.2 Background

3.2.1 Geodetic Detrending

When introducing the GNSS code and carrier phase measurements models from Section 1.2.2 in Chapter 1, we discussed the relative contributions of various phase components. In order to better see diffraction-induced cycle slips in the phase measurements, it is important to remove the large contributions from the geometry term $\mathcal{G}$. That is why many cycle slip algorithms adopt the geometry-based approaches that utilize information from other received satellite signals in order to resolve the non-dispersive signal phase component ([15], [59]). Indeed, optimal cycle slip mitigation techniques should use all available information, including knowledge of non-dispersive phase components. Therefore, we use the relatively straightforward geodetic detrending technique to remove a majority of non-dispersive phase components from our signals. Originally applied in [44] in order to isolate scintillation effects in 1 Hz data, this method is explained in detail in [54]. Here we outline the essential details.

The first step is to remove all signal components that can be estimated a-priori. Since we are dealing with a stationary antenna, the satellite range can be obtained by using the known receiver antenna position and precise orbits supplied by the International Geodetic Service (IGS) [48]. We also estimate the satellite clock and antenna phase variations using IGS products. We use the

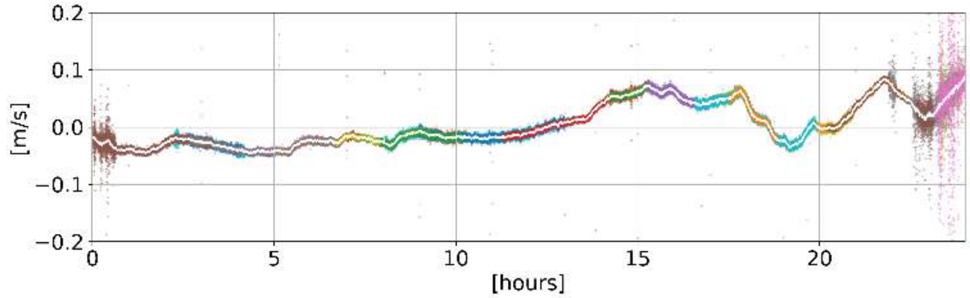

Figure 3.1: Shows the overlay of derivative of L1/L2 IF combinations of carrier phase. An overall linear trend has been removed for clarity. Each color corresponds to a different PRN and the white line shows the derivative of receiver clock variation that is obtained by averaging the IF derivatives.

simple Hopfield model for troposphere delay estimation, which should be accurate to within several centimeters [39]. Combining these factors gives an estimate for the non-dispersive component (G in eqs. 1.4 and 1.5) except for the receiver clock term.

The next step is to estimate the receiver clock variations. Following the approach in [44], we first remove the non-dispersive components estimated in the first step from each of the code and carrier phase measurements and for each satellite (in our case, we considered only GPS satellites). Then we take the IF combination of the L1 and L2 carrier phase measurements, which at this point should only contain the receiver clock variations, carrier bias, cycle slips, and noise. Taking the discrete derivative of the IF combination eliminates the carrier bias and turns cycle slips into outliers. By averaging the IF combinations across all satellites we obtain an estimate of the derivative of receiver clock variations. Figure 3.1 illustrates this process applied to the Ascension Island scintillation dataset that was introduced in Section 1.4.0.1, with the white line showing the estimated clock variation derivative. We then integrate this result to obtain the final receiver clock variations and remove them from each of the code and carrier phases to obtain our final detrended observations. Note that these observations will still contain all variations due to the ionosphere, but the non-dispersive components have been removed at sub meter-level accuracy.

3.2.2 Linear combinations

As we saw in Chapter 1, given a set of code and carrier phase observations we can produce various linear combinations with desirable properties, which have been a popular tool in the detection of cycle slips. In particular, we discussed how geometry- and ionosphere-free combinations are useful when it comes to reliable cycle slip detection. However, it is also important that linear combinations used for cycle slip detection are actually sensitive to cycle slips that occur. Table 3.1 lists the coefficients for various measurement combinations for GPS L1, L2, and L5 signals, including 4 new combinations that we did not introduce in Chapter 1 but which we will use for the cycle slip detection algorithm in Section 3.4.

The first three new combinations, which we denote Y_{L1}, Y_{L2}, and Y_{L5}, are code and single-carrier GF combinations. In each combination, the coefficient for the respective carrier is equal to one while the other two carrier coefficients are zero. Meanwhile the code phase coefficients are optimized to produce a combination that cancels the carrier phase geometry component while minimizing the overall combination variance given by Equation 1.9. The combinations allow for detection of large cycle slip amplitudes on individual carriers, and play a role analogous to the HWM combination that was used in the TurboEdit algorithm from Chapter 1 Section 1.2.4. The fourth new combination is a GF carrier-phase-only combination denoted $\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}$. Its coefficients are chosen such that they maximize the ratio between a 1-cycle slip on the L1 signal and the noise variance corresponding to the time-difference observation $\Delta\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}} = \mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}(t_i) - \mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}(t_{i-1})$, which is assumed to be twice the variance of the undifferenced observation. This combination is designed to be sensitive to slips in the L1 signal that are not easily identified in the first signal.

In [20], the authors assume independent noise with code noise standard deviations $\sigma_{P_1} = \sigma_{P_2} = 0.15$m and $\sigma_{P_5} = 0.1$m (where the lower code noise on L5 corresponds to its higher chipping rate) and carrier noise standard deviations $\sigma_{L_1} = \sigma_{L_2} = \sigma_{L_5} = 0.002$m. While signals under harsh conditions, and under ionosphere diffraction specifically, are subject to higher levels of both

Name	Type	Coefficients					
Y_{L1}	GF	-0.235	-0.235	-0.529	1.000	0.000	0.000
Y_{L2}	GF	-0.235	-0.235	-0.529	0.000	1.000	0.000
Y_{L5}	GF	-0.235	-0.235	-0.529	0.000	0.000	1.000
$\mathrm{GIF}_{\mathrm{L1,L2,L5}}$	GIF				0.215	-1.215	1.000
$\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}$	GF				1.000	-0.370	-0.630
$\mathrm{GF}^{(2)}_{\mathrm{L1,L2,L5}}$	GF				1.000	-2.780	-1.780
$\mathrm{IF}_{\mathrm{L1,L2}}$	IF				-2.546	1.546	0.000
$\mathrm{IF}_{\mathrm{L1,L5}}$	IF				-2.261	0.000	1.261
$\mathrm{GF}_{\mathrm{L1,L2}}$	GF				1.000	-1.000	0.000
$\mathrm{GF}_{\mathrm{L1,L5}}$	GF				1.000	0.000	-1.000

Table 3.1: Various linear combinations of observables useful for triple-frequency cycle slip detection.

code and carrier phase noise, for consistency we will use their values in when implementing their algorithm in Section 3.4. We also note that since the proportion of these noise amplitudes will be similar under diffraction, the coefficients in Table 3.1 would be approximately the same for combinations optimized for ionosphere scintillation anyways.

In order to assess the sensitivities of these combinations to different cycle slips, a useful metric is the ratio of the magnitude of the cycle slip effect to the noise standard deviation for a particular combination. This is essentially the slip signal-to-noise ratio (SNR). Let $\Delta\mathbf{z} \in \mathbb{Z}^3$ be the vector representing the cycle slip amplitude at a given epoch for the three carriers, and let $\mathbf{b} = \mathrm{diag}(\lambda_1, \lambda_2, \lambda_3)\Delta\mathbf{z}$ be the bias vector introduced into the carrier measurements after a slip occurrence. Then $|\mathbf{c}_\Phi^T \mathbf{b}|$ is the magnitude of the cycle slip effect on the linear combination with carrier coefficients $\mathbf{c}_\Phi$ and the standard deviation of noise is provided by Equation 1.9 from Chapter 1. We can then compute the SNR metric as the ratio of these values. To do this, let us consider potential slips at a given epoch of up to one cycle in amplitude. This allows for 3 possible outcomes (-1, 0, +1 cycle) on each of the 3 carriers, for a total of $3^3 = 27$ possible combinations. Out of these, one is the no-slip case (0, 0, 0) and half the remaining cases are redundant since they are additive inverses. This leaves 13 one-cycle combinations.

Table 3.2 lists each combination's SNR measure for each of the 13 possible single-slip scenarios. Blank cells indicate when a combination is entirely insensitive to a slip. While the assumed

	(1,1,1)	(1,0,0)	(1,1,0)	(0,1,0)	(-1,0,-1)	(0,0,-1)	(0,-1,-1)	(1,-1,0)	(1,0,-1)	(-1,1,1)	(0,1,-1)	(1,-1,1)	(1,1,-1)
Y_{L1}	2.6	2.6	2.6		2.6			2.6	2.6	2.6		2.6	2.6
Y_{L2}	3.4		3.4	3.4			3.4	3.4		3.4	3.4	3.4	3.4
Y_{L5}	3.5				3.5	3.5	3.5		3.5	3.5	3.5	3.5	3.5
$\mathrm{GIF}_{\mathrm{L1,L2,L5}}$	0.3	12.9	80.5	93.4	93.1	80.2	13.2	106.3	67.3	26.1	173.6	186.5	160.7
$\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}$	24.5	76.8	40.3	36.5	12.0	64.8	101.3	113.3	141.6	178.1	28.3	48.5	105.2
$\mathrm{GF}^{(2)}_{\mathrm{L1,L2,L5}}$	136.6	27.6	70.8	98.4	38.2	65.8	164.2	126.0	93.3	191.8	32.7	60.2	5.1
$\mathrm{IF}_{\mathrm{L1,L2}}$	18.0	81.3	18.0	63.4	81.3		63.4	144.7	81.3	144.7	63.4	144.7	18.0
$\mathrm{IF}_{\mathrm{L1,L5}}$	21.0	83.1	83.1		21.0	62.1	62.1	83.1	145.2	145.2	62.1	21.0	145.2
$\mathrm{GF}_{\mathrm{L1,L2}}$	19.1	67.3	19.1	86.3	67.3		86.3	153.6	67.3	153.6	86.3	153.6	19.1
$\mathrm{GF}_{\mathrm{L1,L5}}$	22.8	67.3	67.3		22.8	90.1	90.1	67.3	157.4	157.4	90.1	22.8	157.4

Table 3.2: Shows the sensitivity of each combination in Table 3.1 to different cycle slip amplitudes. Different colors are shown to help indicate the sensitivity, with yellow being largest and purple being smallest.

noise levels are too low to correspond to severe scintillation conditions, the values in Table 3.2 still provide some important insights into the limitations of cycle slip detection. For instance, we see that all combinations using code phase (Y_{L1}, Y_{L2}, and Y_{L5}) have fairly low SNR metrics compared to combinations only using carrier phase. Also, the SNR metric for L1 slips is generally much smaller than those for L2 and L5. This is due to the L1 signal's shorter wavelength and the wider spectral gap between L1 and the lower-frequency bands. In the next section, we will use the GIF combination (listed as $\mathrm{GIF}_{\mathrm{L1,L2,L5}}$ in tables 3.1 and 3.2) to assess slip occurrences in the scintillation time series.

3.3 Analysis of challenging cycle slips

We want to assess the occurrence of this slips, and in doing so relate them back to our prototypical examples of canonical fades in Figure 2.5 from Chapter 2. Recall that there were 3 cases: A) a canonical fade occurs but there is no cycle slip, B) a canonical fade occurs and there is a phase transition that leads to a cycle slip, C) a canonical fade occurs during which noise appears to induce a cycle slip. We refer to these three prototypical examples when analyzing the phase transitions in the real data.

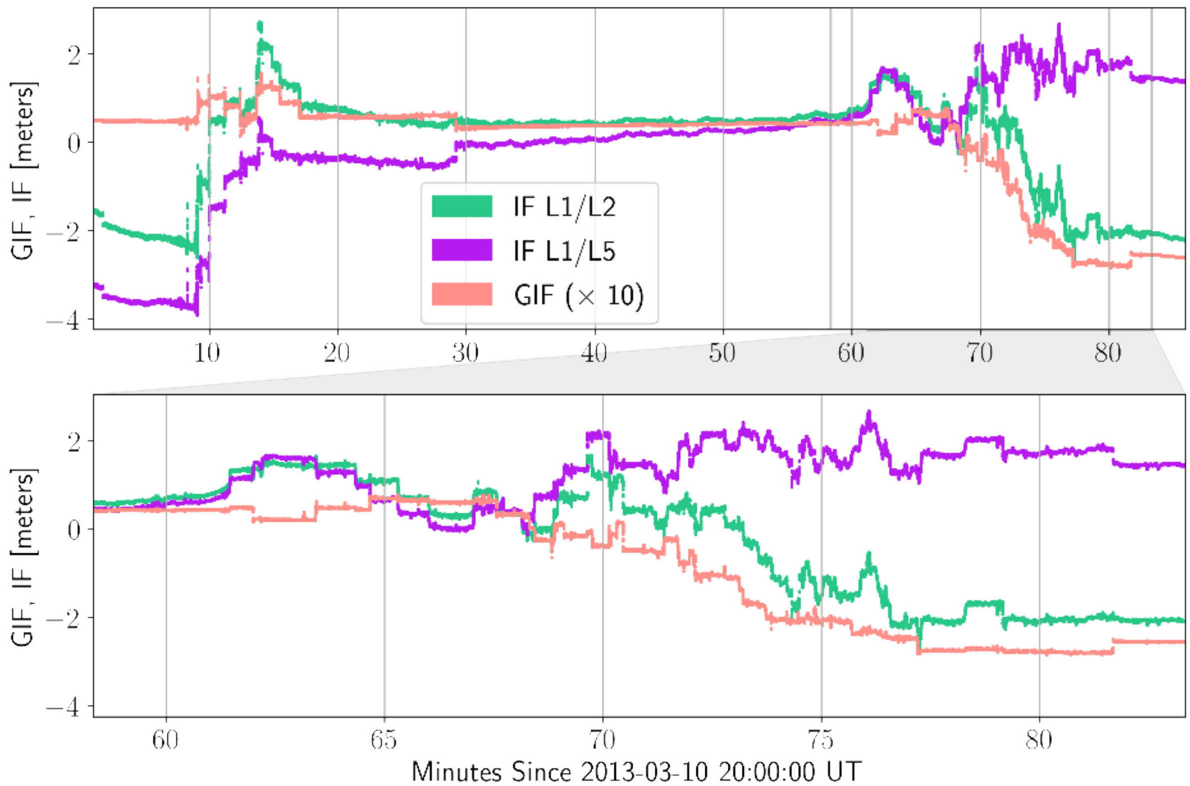

Figure 3.2: Examples of L1/L2 and L1/L5 IF combinations and the L1/L2/L5 GIF combination using 100 Hz carrier phase measurements from the Ascension Island dataset. Bottom panel shows a zoomed-in portion of the most disturbed period.

3.3.1 Ascension Island Scintillation

We first consider the 2013-03-10 Ascension Island scintillation event for GPS PRN 24, which we introduced in Section 1.4.0.1. Figure 3.2 shows the IF and GIF phase combinations computed using the detrended 100 Hz phase measurements. The discrete jumps in these combinations indicate the presence of diffraction-induced cycle slips. The bottom panel of Figure 3.2 shows a zoomed-in portion of the most intense diffractive effects around minutes 75 to 80. The jumps in these combinations show multiple cycle slips occurring each minute.

Figures 3.3 and 3.4 show examples of canonical fades that occur on the L2 and L5 signals. Based on analysis of the GIF combination, there are no cycle slips caused by ionosphere diffraction in any of these examples, and for the most part the phase behavior corresponds well with that shown in simulation scenarios A and B from Section 2.1.1. Figure 3.5 shows a (presumably) noise-induced cycle slip occurring in the middle of an otherwise canonical fade. This behavior is very similar to that of simulation scenario C. The instantaneous nature of the phase transition during this noise-induced cycle-slip helps distinguish it from full-cycle transitions corresponding to simulation scenario B. The GIF combination during these events remains relatively constant before and after the fades, indicating the absence of full-cycle phase transitions.

Figure 3.6 show an example of a full-cycle transition that occurs during deep fades in the L2 and L5 signals. Though there is clearly a bias discrepancy among the detrended phases, it is not immediately clear which of the fading signals underwent a full-cycle transition. The average value of the GIF combination before versus after the jump event indicates a jump of about 0.29. Referencing Table 3.1, we deduce that this jump corresponds to a full-cycle transition in the L2 signal since $1.215 \times \lambda_{\mathrm{L2}} = 0.297$. Figure 3.7 shows another example of this behavior. Beginning with weak and canonical fades on all three signals, the L2 and L5 signals experience additional weak fading 1.25 seconds after their initial fades. Again, the detrended phase biases imply a full-cycle phase transition. There is a jump in the GIF combination signal value from before and after the fade event of around 0.25, corresponding to a full-cycle transition in the L5 signal (since $1 \times \lambda_{\mathrm{L5}} = 0.255$). The detrended signal intensity for L5 only shows around -10 dB fading when this transition must have occurred. This corroborates our finding from Chapter 2 that weak to moderate fades may be as important as very deep fades when it comes to the occurrence of cycle slips in very strong scintillation.

Figures 3.8, 3.9, and 3.10 all show examples of consecutive deep fades, mostly in the L2 and L5 signals. Just as in each of the previous examples, whenever moderately sharp fades occur, the carrier phase exhibits an approximate half-cycle phase transition. For each of these examples, parsing out the true phase behavior is an exercise in distinguishing canonical half-cycle transitions from those that are actual full-cycle transitions. The phase behavior from the examples in Figures 3.10 and 3.9 are particularly messy in this regard. The GIF combination proves to be a useful tool in identifying the occurrence and culprit of some full-cycle phase transitions. However, these

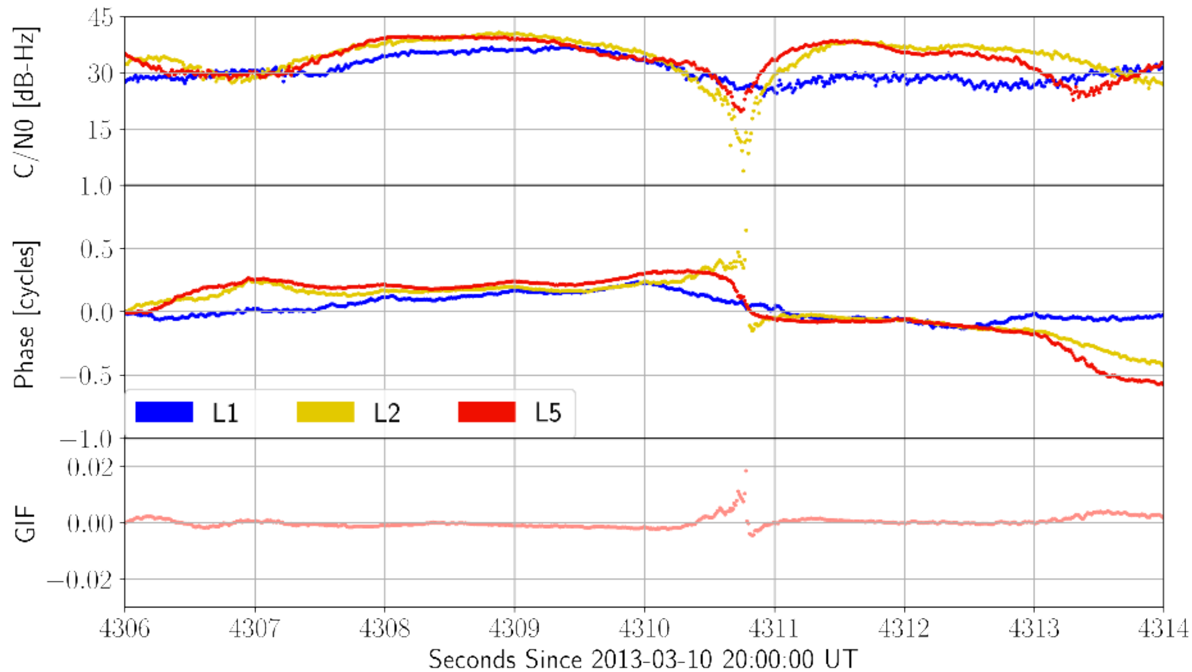

Figure 3.3: Example of canonical fade behavior occurring on L2 and L5 signals at around 76311 seconds. L2 shows a sharp half-cycle phase change, corresponding to simulation scenario B, whereas L5 shows a slightly more gradual transition similar to scenario A. The GIF combination bias stays the same before and after the fade, indicating no cycle slip (phase transition) occurred.

examples also show how the jumps in average GIF combination values are neither immediate nor easy to measure due to persistent corruption by diffractive fluctuations.

3.3.2 Hong Kong Scintillation Example

We now consider Hong Kong scintillation dataset that we introduced in Section 1.4.0.2. After applying the geodetic detrending technique from Section 3.2.1, we computed the GIF phase combination and the dual-frequency IF combinations of phase. In [63], the authors show how the diffraction effects are somewhat suppressed in ionosphere-free combinations of carrier phase. This makes identification of cycle slips easier, especially when using very-high-rate (100 Hz) carrier phase measurements. We performed a manual cycle slip detection and correction procedure by iteratively identifying and removing cycle slips from the three 100 Hz phase measurements. Identification of slips was done by observing and interpreting the behavior of jumps in the GIF, IF, and GF

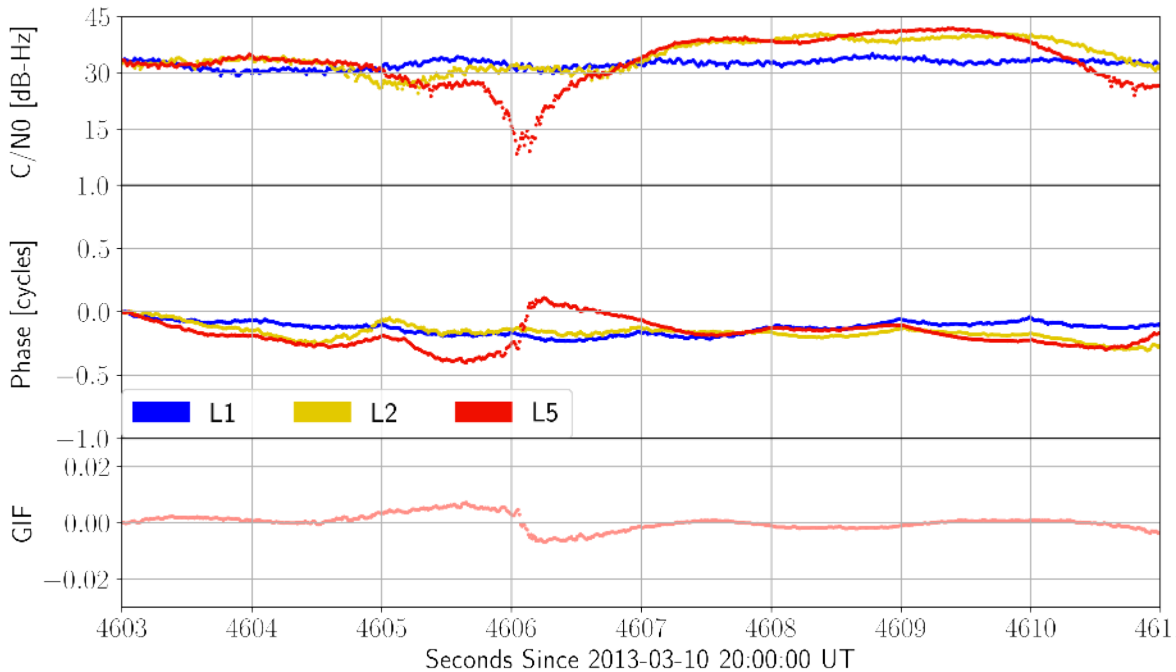

Figure 3.4: Example of a canonical fade on the L5 signal. The behavior is similar to that of simulation scenario B. Similar to the example in Figure 3.3, there were no full-cycle transitions during this deep fade.

combinations during fades in signal intensity, as discussed in the last section. The result of this procedure is summarized in the bottom panel of Figure 3.11, which shows the now relatively smooth IF combinations that have been rid of large jumps. We found 18 slips on L1, 53 on L2, and 78 on L5. These results are roughly consistent with the cycle slip rates we found in Chapter 2. We acknowledge that there are likely errors in this procedure, and in particular cycle slips occurring on all three signals simultaneously are extremely difficult to detect in these IF combinations (e.g. see Table 3.2) or during ionosphere diffraction in general. Nevertheless, the residuals of the IF and GIF combinations in Figure 3.11 give us some confidence that most of the cycle slips are correctly identified, and that the results of this procedure are adequate for assessing the performance of the different cycle slip mitigation methods in the rest of this chapter.

Time	Amplitude														
7554	0	0	-1	7849	0	1	0	8610	0	1	1	8893	0	0	1
7558	0	-1	-1	7915	0	1	1	8634	0	1	1	8898	1	0	0
7582	-1	-1	-1	7958	0	0	1	8653	0	1	1	8901	-1	0	0
7616	0	0	-1	8285	0	1	1	8688	0	0	1	8913	0	0	1
7636	0	0	1	8299	0	-1	0	8694	0	1	1	8948	0	1	0
7638	1	0	0	8321	0	-1	-1	8699	0	1	0	8957	0	-1	-1
7640	1	1	1	8337	0	0	-1	8712	-1	0	0	8962	0	0	1
7647	0	0	-1	8341	0	0	1	8723	0	-1	-1	8975	0	1	1
7662	0	0	-1	8342	0	0	1	8732	0	1	1	8979	0	0	1
7673	0	-1	-1	8389	0	0	-1	8743	0	-1	-1	8995	0	1	1
7689	-1	1	0	8391	-1	0	0	8761	0	0	-1	9016	0	0	1
7695	0	-1	1	8442	0	1	1	8766	0	1	1	9038	0	0	-1
7708	0	0	-1	8463	0	0	-1	8773	0	0	1	9047	0	1	1
7710	0	1	0	8466	0	0	1	8780	0	1	1	9090	0	0	1
7715	0	0	-1	8472	0	-1	-1	8784	0	0	-1	9123	0	-1	0
7730	1	0	0	8481	0	-1	-1	8800	0	1	1	9161	0	-1	-1
7732	0	1	0	8503	0	-1	-1	8834	0	0	-1	9177	0	-1	-1
7733	0	0	1	8521	1	0	0	8848	0	0	-1	9179	1	0	0
7757	0	-1	-1	8538	0	-1	-1	8850	-1	0	0	9189	0	-1	-1
7764	1	0	0	8540	0	1	1	8872	0	-1	-1	9207	0	1	1
7815	0	-1	0	8554	0	-1	-1	8877	0	0	1	9247	0	-1	-1
7817	0	1	1	8563	0	1	1	8883	1	0	0	9305	0	1	1
7818	0	0	-1	8585	-1	0	0	8889	0	-1	-1	9398	0	0	-1
7821	0	1	1	8589	-1	0	-1	8890	0	-1	-1	9432	0	0	1
7845	0	-1	0	8592	0	0	1	8892	0	1	1	9540	0	0	-1

Table 3.3: Manually detected cycle slip occurrences. Occurrence times are measured in seconds from 11:00 UT.

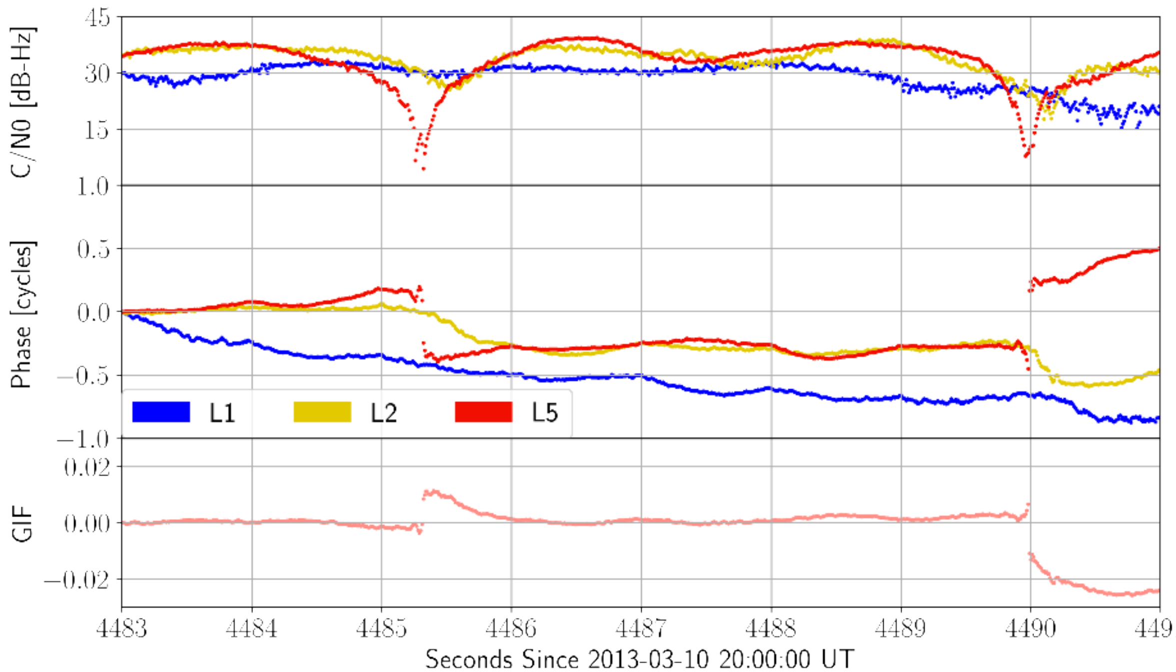

Figure 3.5: Examples of canonical fades on the L5 signal. The first canonical fade at 76485 seconds does not show a corresponding change in the GIF bias, and presumably no cycle slip has occurred. The second fade appears to show a noise-induced cycle-slip, corresponding to simulation scenario D. We see a corresponding drop in the GIF combination bias.

3.4 Cycle Slip Mitigation Performance: De Lacy 2012

Here we assess the effectiveness of cycle slip mitigation using the algorithm presented in [20]. We test the algorithm using 1 Hz data, for which it was designed. It makes use of several of the linear combinations of observables from Table 3.1. Just like we saw for the TurboEdit algorithm in Chapter 1, a common approach in many cycle slip algorithms is to first use both code and carrier phase measurements together in order to identify large cycle slips. Any larger cycle slips (more than 2 cycles) are most likely caused by receiver processing errors or loss-of-lock and can be easily detected in such combinations. Once these large jumps are identified, carrier phase combinations are then used to detect and correct smaller slips. This technique is known as cascading detection and is used by many other authors including [50], [102], [2], [14]. In this section, we will first briefly describe the algorithm steps, then we will show and assess its performance in detecting and

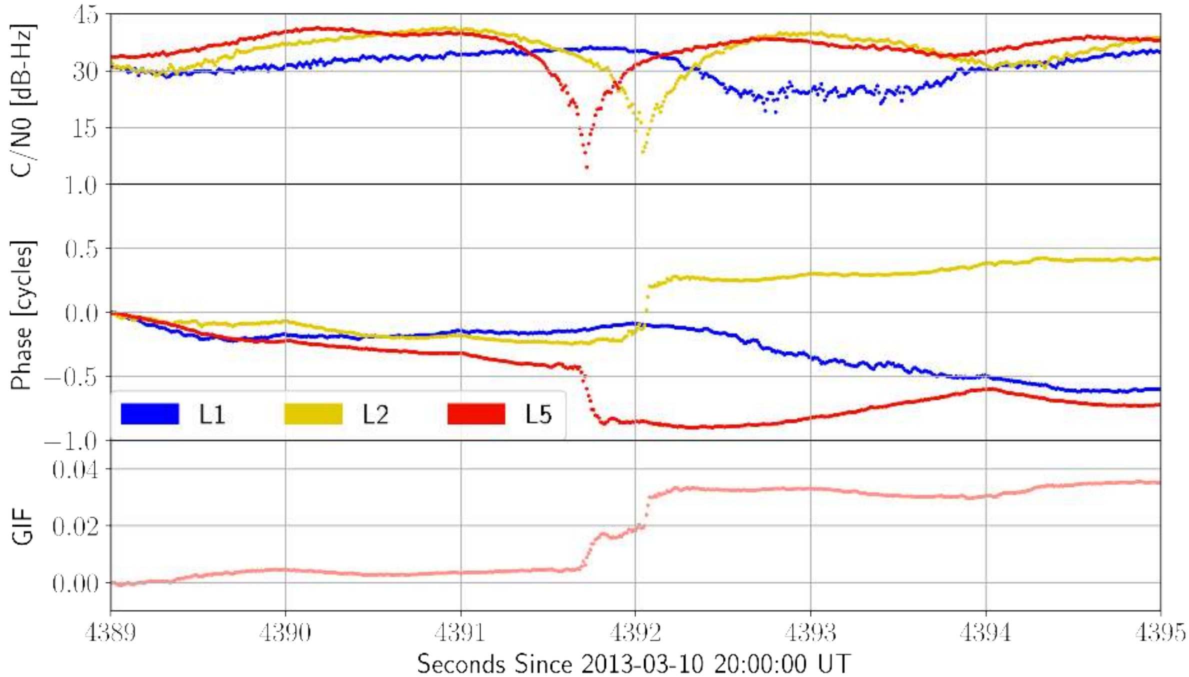

Figure 3.6: Using detrended phase and signal intensity, shows an interesting example of triple-frequency carrier behavior where it it is not clear from the individual phase records which signals (if any) accrue full cycle transitions and which just show canonical fade behavior. Careful measuring of the change in the average GIF combination value before and after the fade event reveals a jump of around 0.29, suggesting that a full-cycle transition occurred in the L2 signal only.

estimating cycle slips in the Hong Kong dataset.

3.4.1 Algorithm Description

The De Lacy algorithm uses the Y_{L1}, Y_{L2}, Y_{L5}, $\mathrm{GIF}_{\mathrm{L1,L2,L5}}$, $\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}$, and $\mathrm{GF}^{(2)}_{\mathrm{L1,L2,L5}}$ combinations from Table 3.1 and compares the magnitudes of their epoch-wise differences (e.g.$\Delta Y_i = Y_i(t) - Y_i(t-1)$) to a threshold in order to detect slips. Similar to the Turboedit algorithm we assessed in Chapter 1, for each detection the authors suggest a 4σ threshold (i.e. 4 times the standard deviation of ΔY_i). The steps to detect and estimate cycle slips are performed at each epoch of time-differenced combinations:

(1) The time-differenced combinations ΔY_{L1}, ΔY_{L2}, ΔY_{L5} are each tested for large slip occurrences using a threshold of 0.41 meters for $i = 1, 2, 3$. If a slip is detected in any of the

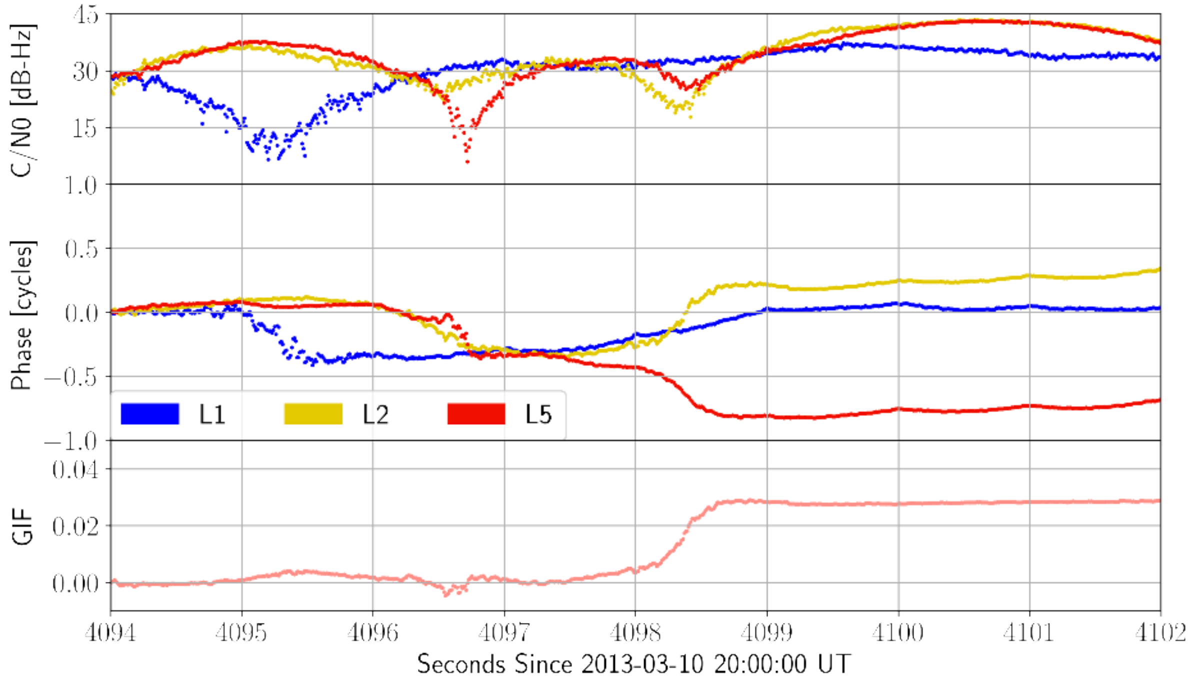

Figure 3.7: Using detrended phase and signal intensity, shows moderate canonical fades on the L1, L2 and L signals. There is a post-fade discrepancy in the L2 and L5 signals suggesting a possible full-cycle transition. A change in average GIF combination value of around 0.25 before versus after the fades suggests that the L5 signal underwent a full-cycle transition.

combinations, move to step 4.

(2) The time-differenced combination $\Delta\mathrm{GIF}_{\mathrm{L1,L2,L5}}$ is tested against a threshold of 1.8 cm. If a slip is detected, move to step 4.

(3) The time-differenced combination $\Delta\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}$ is tested against a threshold of 1.7 cm. If a slip is detected, move to step 4.

(4) When a slip is detected, first its approximate magnitude is estimated as $\Delta z_k = \mathrm{round}(Y_k/\lambda_k)$. At this point, the residual slip magnitude is assumed to be between ± 2 cycles.

(5) Similar to the TurboEdit algorithm, the remaining joint cycle slip amplitude $\Delta\mathbf{z}$ is then chosen to minimize the residual in a set of time-differenced phase combinations. In particular, the algorithm searches over the joint slip amplitudes for the one that minimizes $|\Delta\mathrm{GF}_{\mathrm{L1,L2}}| + |\Delta\mathrm{GF}_{\mathrm{L1,L5}}|$.

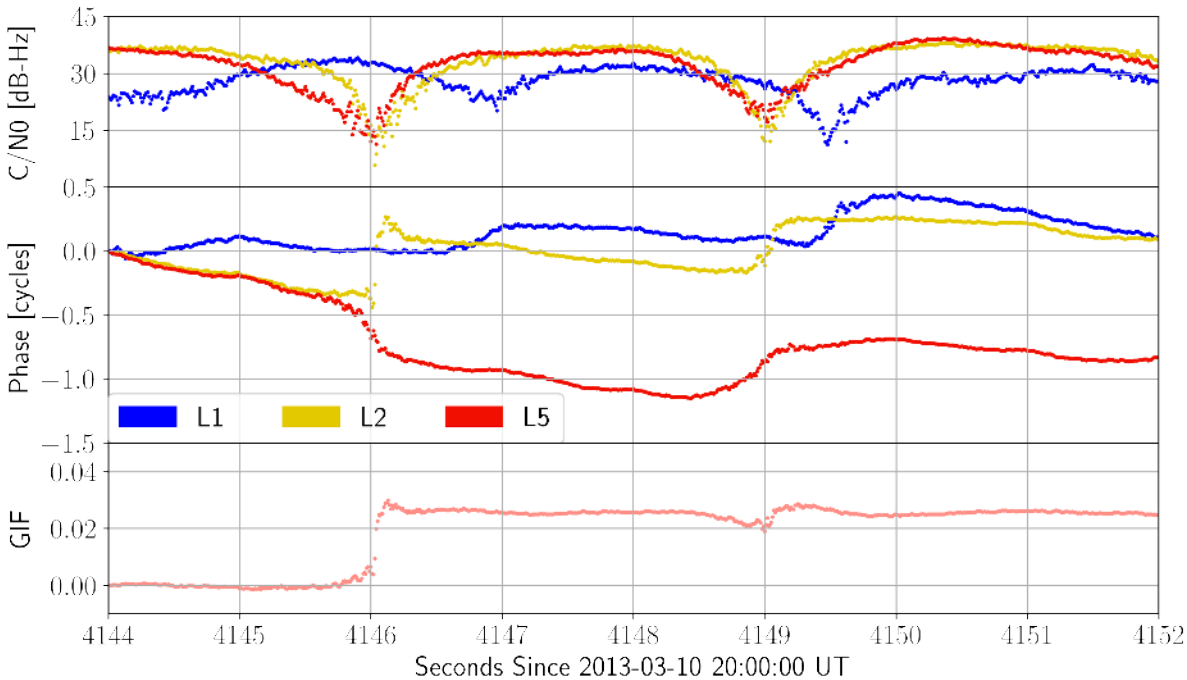

Figure 3.8: Using detrended phase and signal intensity, shows consecutive canonical fades on L2 and L5 signals, as well as a canonical fade on the L1 signal. The first fade in L5 coincides with a jump in GIF combination value of 0.25, indicating a positive full-cycle transition, while its second fade shows a downward half-cycle transition. The L2 signal shows two canonical fades in the same downward direction.

(6) The following pairs of joint amplitudes cannot be resolved by the two combinations in Step 5: $\{(-2,-2,-2),(-2,-1,1)\}$, $\{(-2,-1,-2),(2,2,1)\}$, $\{(2,1,2),(-2,-2,-1)\}$, $\{(2,2,2),(-2,-1,-1)\}$. Therefore, if the estimated amplitude lies in one of these pairs, the sixth and final combination $\mathrm{GF}^{(2)}_{\mathrm{L1,L2,L5}}$ is used to distinguish between the two outcomes. Specifically, the joint amplitude from the corresponding pair that minimizes $|\Delta\mathrm{GF}^{(2)}_{\mathrm{L1,L2,L5}}|$ is chosen.

3.4.2 Analysis

Figure 3.12 shows the residuals ΔY_i along with the $\pm 4\sigma$ thresholds. Additionally, residuals corresponding to cycle slips in the truth reference are encircled. There is a high false-detection rate for the test on ΔY_i, $i = 1, 2, 3$. Notably, during manual cycle slip adjustment presented in the last section, all identified slips were of 1 cycle in amplitude, though there were consecutive slips over a

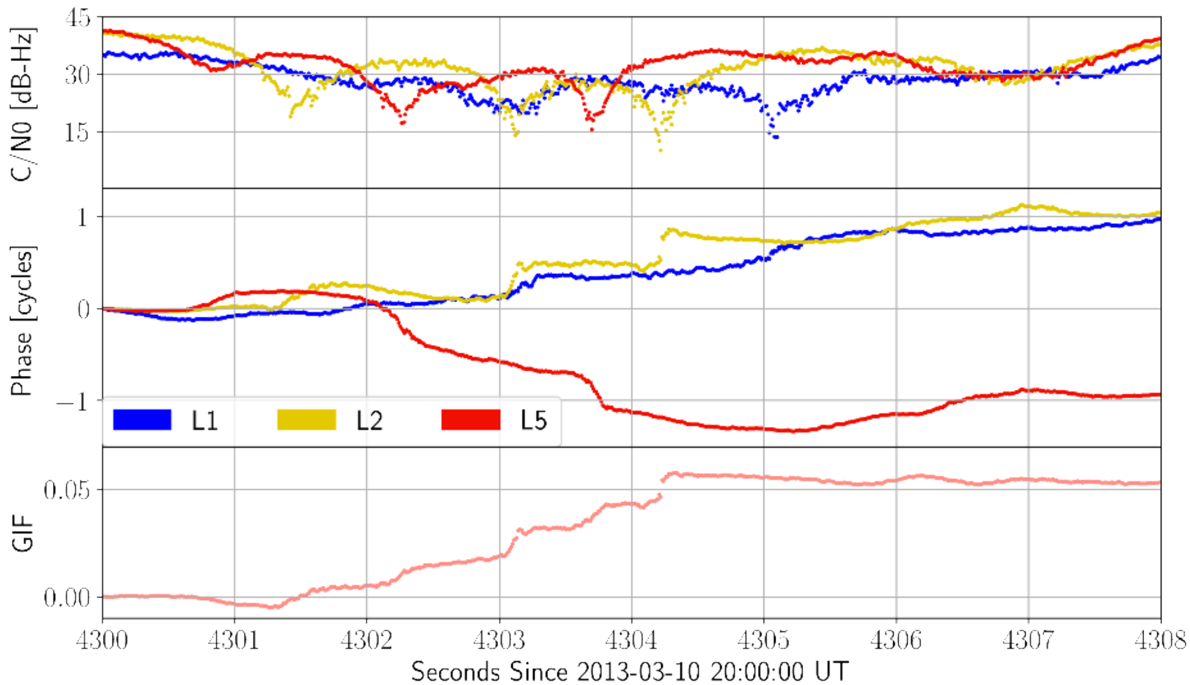

Figure 3.9: Using detrended phase and signal intensity, shows consecutive fades for L2 and L5 signals. The GIF combination show a jump in average value of around 0.50 before versus after the fade, indicating two full-cycle transitions in the L5 signal. The half-cycle transition behavior on the L2 signal is consistently downward while the full-cycle transition behavior of the L5 signal is consistently upward.

few seconds that would appear as multiple-cycle slips in lower-rate data. However, for 1 Hz data, this means that all detected cycle slips in our manual truth reference are considered as small slips. This further implies that no cycle slips should have been detected when using these combinations that include code phase observables, and yet 112 slips for L1, 143 for L2, and 153 for L5 were detected using the standard thresholds. After accounting for the cycle slip occurrence from the truth reference, this corresponds to false detection rates of 91, 79, and 60 percent, respectively.

There is clearly increased noise due to ionosphere variability and diffraction. While the thresholds could be adjusted to account for ionosphere variability and diffraction fluctuations, this would mean that if there were any larger slips some of them may be missed during detection using code phase. Additionally, it is interesting to note that it is not just diffraction effects on carrier phase that are causing increased noise in these observables, but that there is also a large increase

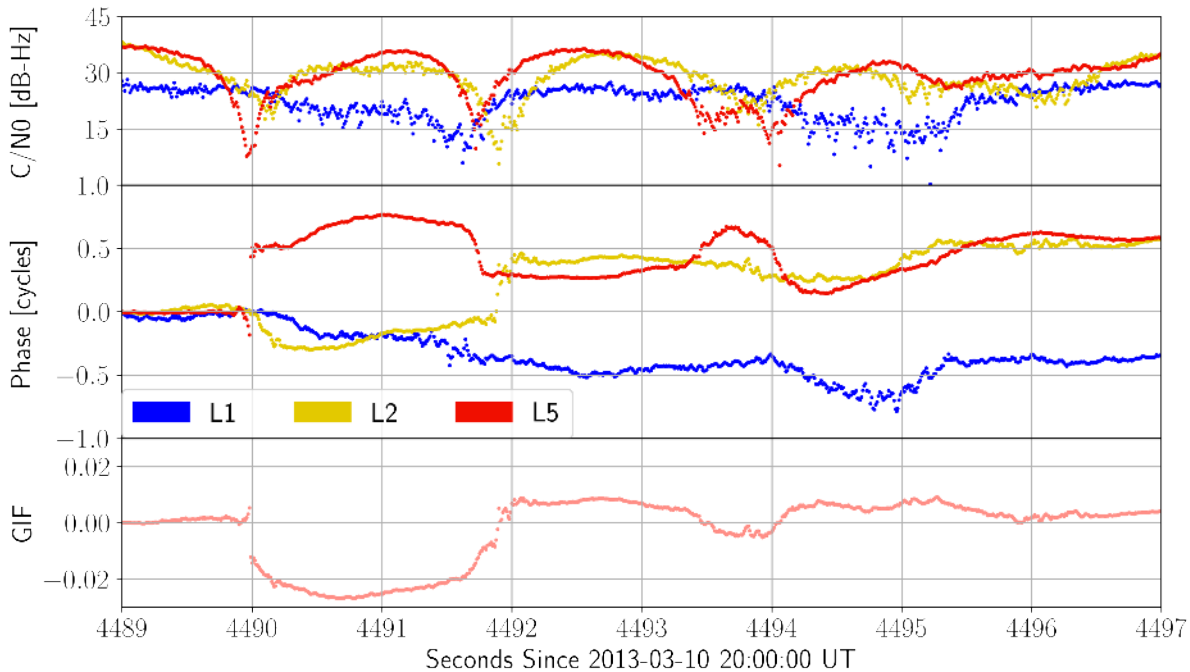

Figure 3.10: Using detrended phase and signal intensity, shows consecutive deep fades mainly for the L2 and L5 signals. The first fade coincides with canonical upward half-cycle transition of the L2 signal. After accounting for its noise-induced cycle-slip, the L5 signal also shows a canonical downward half-cycle transition. For the second fade, the L2 signal shows a downward full-cycle transition with a corresponding change in average GIF combination value, while the L5 signal shows a canonical upward half-cycle transition.

in fluctuations and errors in the code phase measurements. This can be seen more clearly in Figure 3.13, which shows detrended GF combinations of dual-frequency code phase pairs. In this case, we are unsure of the actual origin of increased code phase fluctuations and whether they are directly a result of diffraction effects on the signal or some receiver processing artifact such as carrier-based smoothing of the code phase measurements. In either case, the increased code phase noise is clearly associated with the ionosphere diffraction.

The next stage of detection considers $\Delta\mathrm{GIF}_{\mathrm{L1,L2,L5}}$ and $\Delta\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}$, which are shown in Figure 3.14 along with their $\pm 4\sigma$ thresholds. Additionally, residuals corresponding to jumps in the manual truth reference are encircled, similar to Figure 3.12. It is interesting to note how different slip combinations tend to lie along certain subspaces of the joint $\Delta\mathrm{GIF}_{\mathrm{L1,L2,L5}}, \Delta\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}$ residual. This hints at more optimal testing procedures for these different slip combinations, as

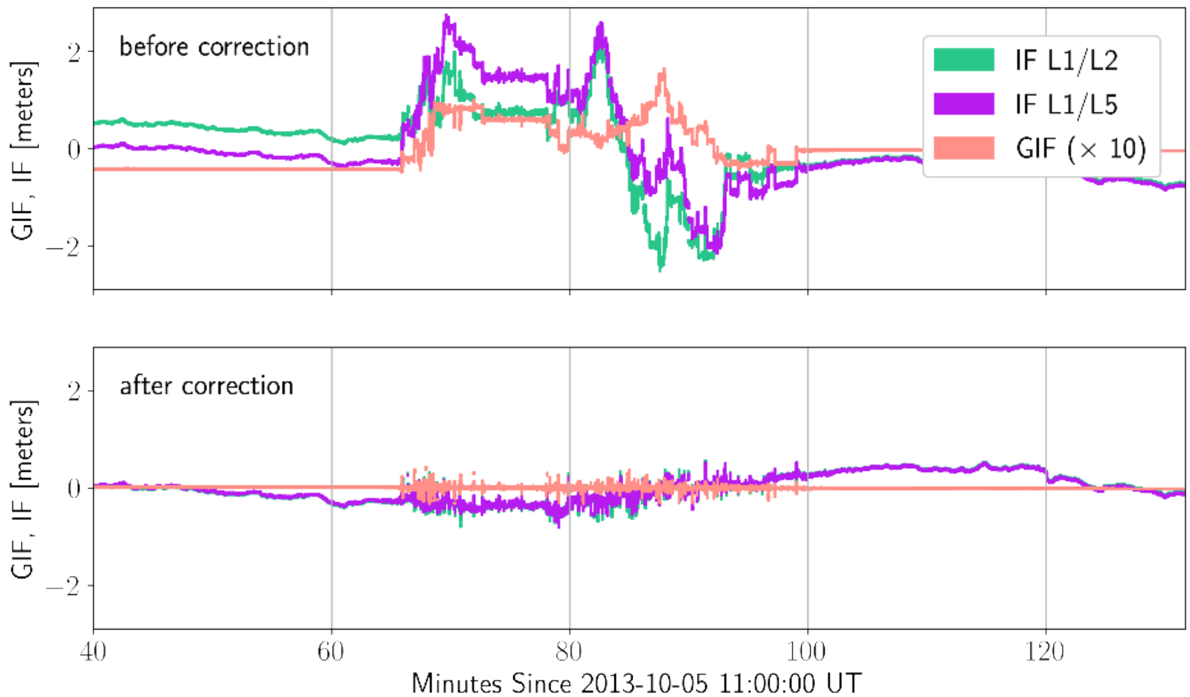

Figure 3.11: Shows examples of L1/L2 and L1/L5 ionosphere-free combinations and the L1/L2/L5 geometry-ionosphere free combination using 100 Hz carrier phase measurements. Top and bottom panels show the combinations before and after manual cycle slip correction.

discussed in [81] or [3]. We will explore this idea more in Chapter 4. For this method, a cycle slip is detected if the threshold is exceeded for either $\Delta\text{GIF}_{\text{L1,L2,L5}}$ or $\Delta\text{GF}^{(1)}_{\text{L1,L2,L5}}$. This detection scheme achieved 100 percent detection rate for epochs containing cycle slips in the truth reference. However, similar to the issues faced during large slip detection, we again see large false-detection rates. This method determined that there were cycle slips at 1216 epochs, thus yielding a false-detection rate of 91 percent. This false detection rate is reduced when neglecting epochs where estimated slip amplitudes are all zero, which we discuss next. Even still, such high rates of false detection are inefficient, especially for when it comes to algorithms that use detections to reset carrier ambiguities.

In the last step, slip amplitudes are estimated using the three different GF phase combinations. Because these combinations use carrier phase only, they should normally be precise enough to reliably estimate the cycle slip amplitudes on all three signals with high confidence. For the case

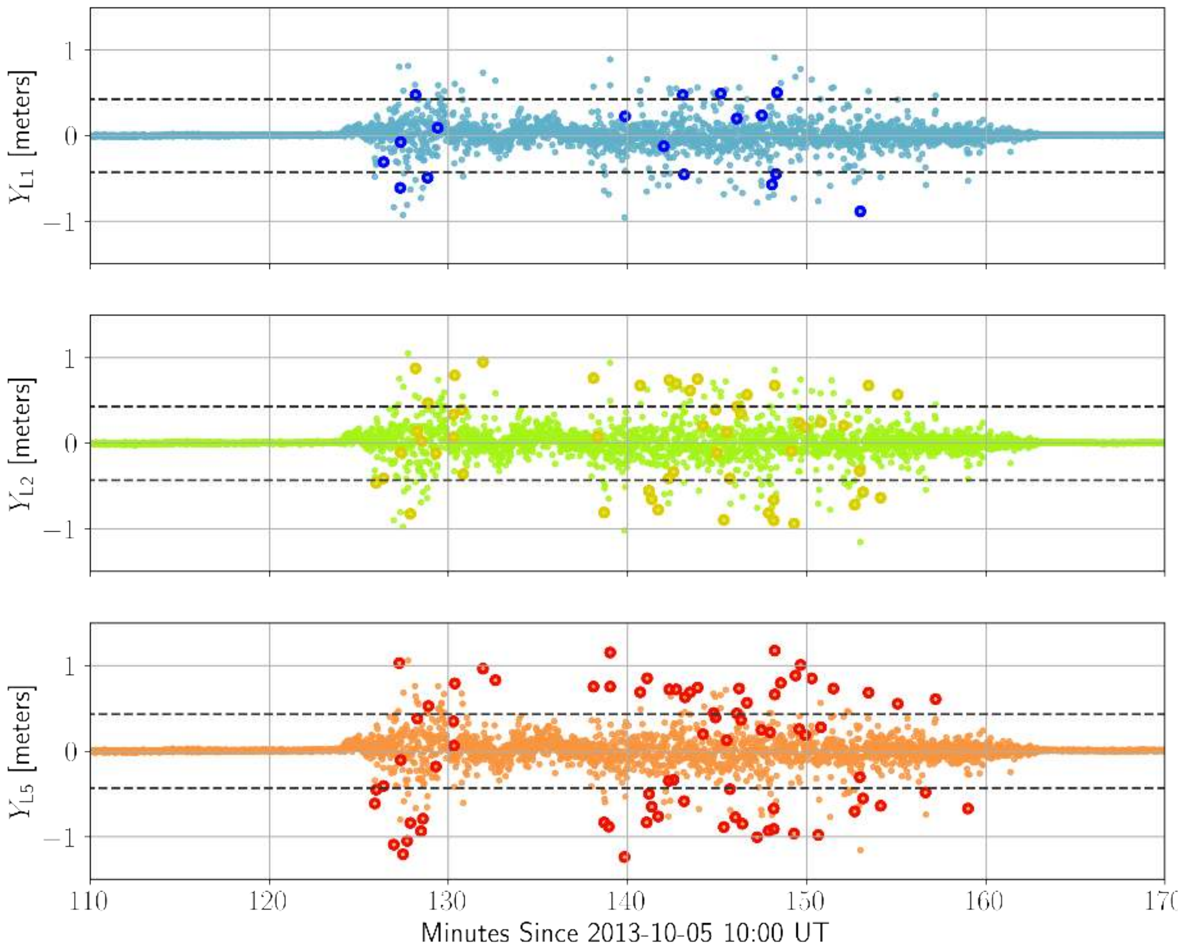

Figure 3.12: The residuals of ΔY_i, $i = 1, 2, 3$, along with circles indicating the occurrence of an actual cycle slip in the manual truth reference for the corresponding signal (L1= 1, L2= 2, L5= 3). Dashed lines indicate slip detection thresholds, and overall we see many missed detections and false alarms for these combination residuals.

of ionosphere scintillation, however, the ionosphere variation and diffraction noise will especially impact these combinations, and so we expect poor estimation results. Figure 3.15 shows the slip estimated along with slip amplitudes from the truth reference. Overall, 594 of the epochs that tested positive for cycle slips resulted in an estimate of no slip on all three signals. The false-detection rate for epochs containing cycle slips goes down to 44 percent if we discount their occurrences. While there are many epochs with ± 2 cycle amplitude slip estimates, they are much less numerous than the ± 1 cycle slip amplitudes. Their occurrence also correlates fairly well with the slip occurrence from the truth reference. Careful inspection of the slip amplitudes shows that often $+1$ or

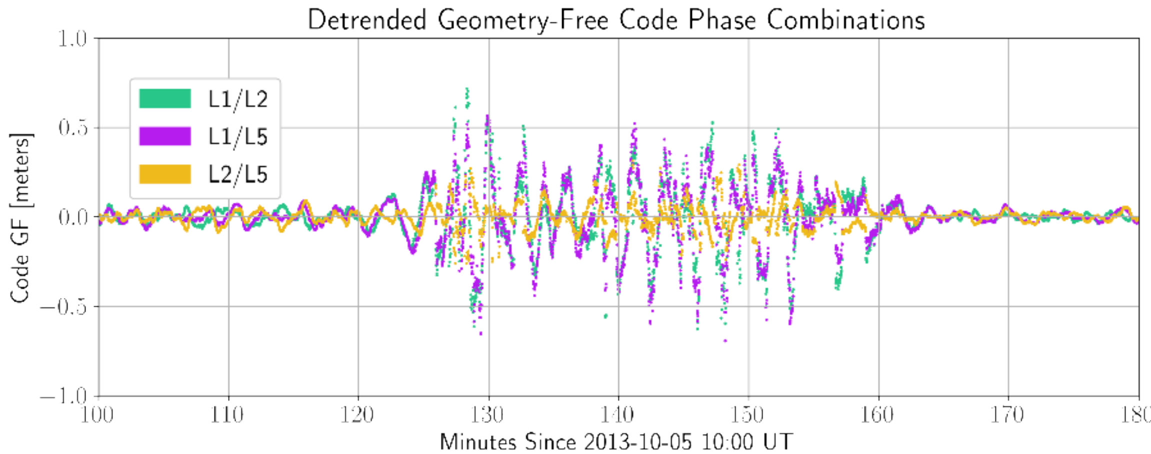

Figure 3.13: Detrended GF combinations of the L1/L2, L1/L5, and L2/L5 code phase observables, revealing a substantial increase in code phase noise and fluctuations during the scintillation event.

−1 cycle slips are estimated simultaneously on all three frequencies. In fact, this was the case for 499 epochs, or nearly 89 percent of the detected jump epochs with non-zero slip amplitudes. The authors in [20] acknowledge the relative insensitivity of GF combinations to simultaneous slips in the same direction on all three signals. Indeed, this may be the most challenging aspect of small slip detection during ionosphere diffraction.

3.5 Cycle Slip Mitigation Performance: Filtering Algorithm

Work from [85] presents a Kalman filter algorithm to mitigate cycle slip occurrences for signals reflected off water surfaces. The state-space algorithm adapts its gain based on fluctuations in signal carrier-to-noise density ratio C/N_0 while estimating the occurrence of integer-cycle jumps due to cycle slips. Both reflected and scintillating signals exhibit interference and deep fading due to scattered signals, and as such it is natural to apply this algorithm (with some modification) to scintillation signals and consider its performance. Here we develop and assess an adaptation of the algorithm from [85], which we originally presented in [12], to apply to triple-frequency scintillation data. In the next section, we describe the algorithm. Then we assess its performance when applied to the Hong Kong dataset. Overall, we will find this algorithm has better performance than the

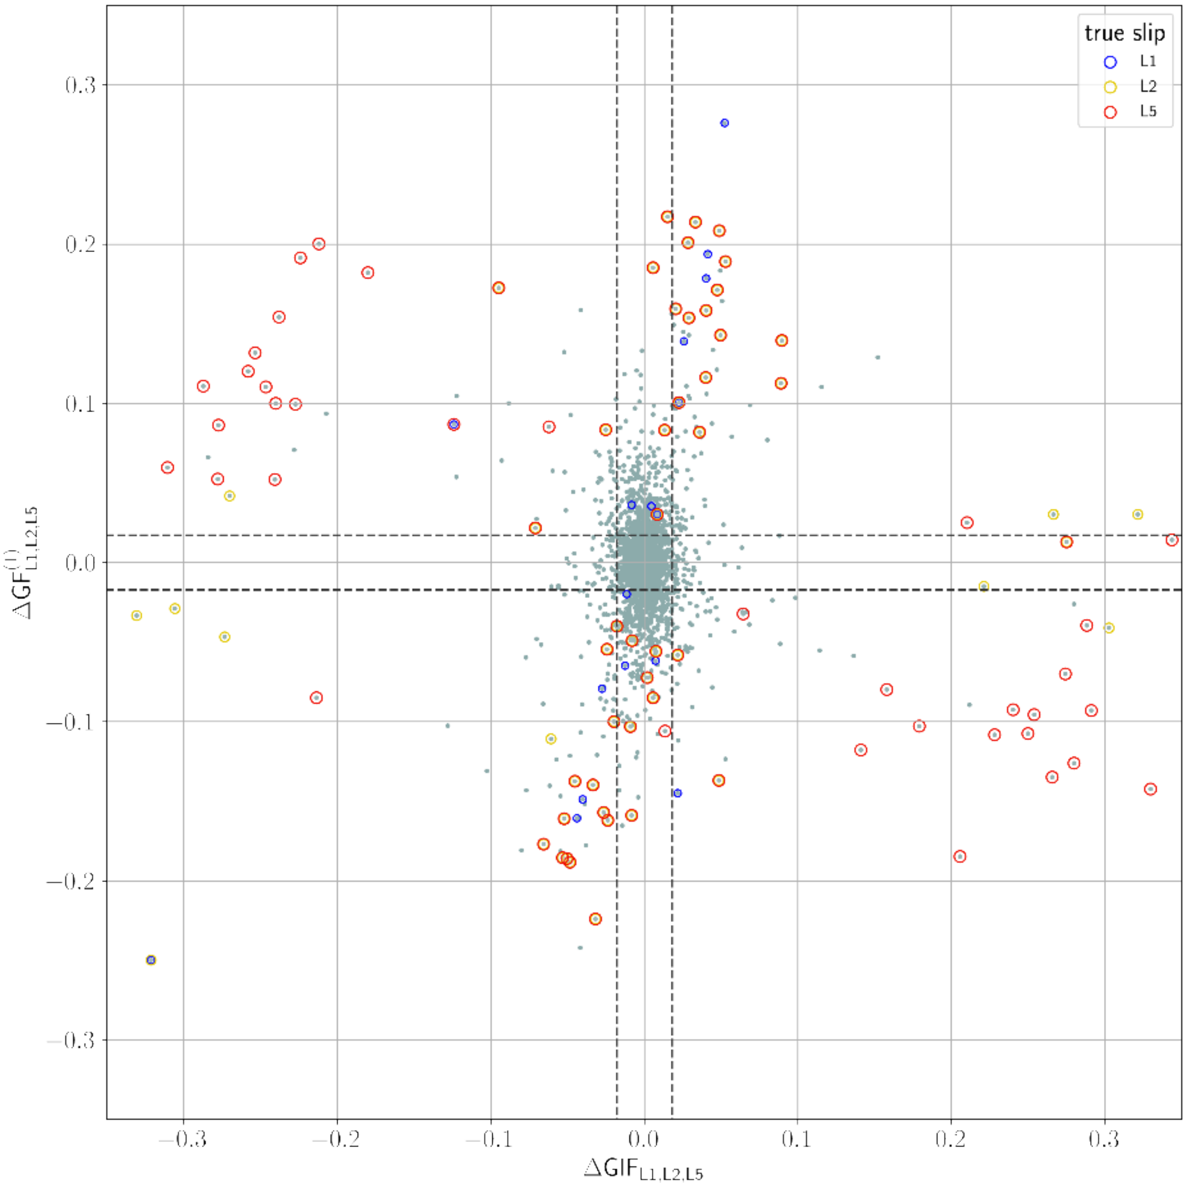

Figure 3.14: Residuals $\Delta\mathrm{GIF}_{\mathrm{L1,L2,L5}}$ plotted against $\Delta\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}$ along with their $\pm 4\sigma$ thresholds (dashed lines). Residuals corresponding to actual slips in the manual truth reference are encircled, similar to Figure 3.12

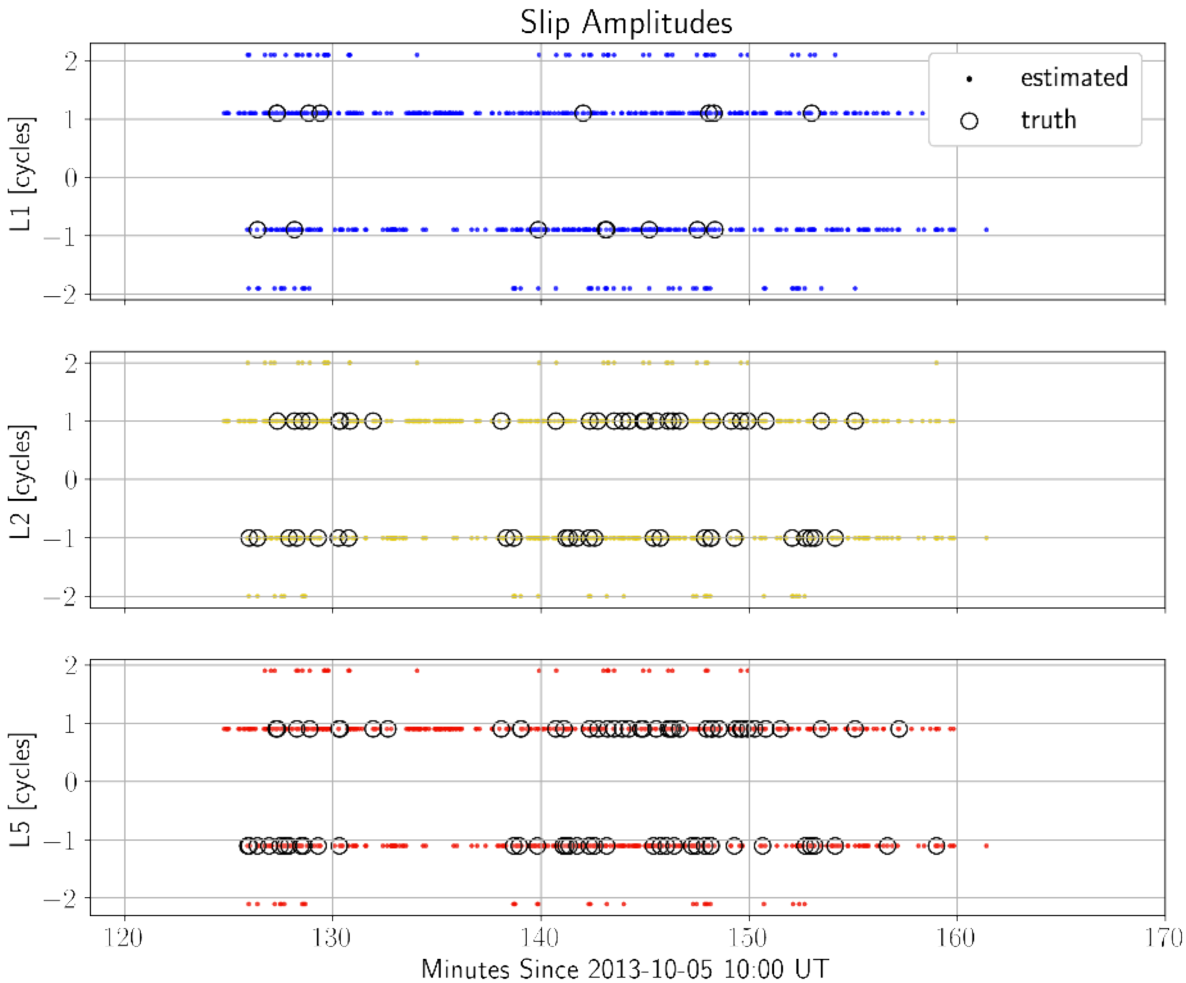

Figure 3.15: Estimated cycle slip amplitudes at each epoch for the method outlined in [20] as well as the amplitudes from the truth reference (shown with circles).

De Lacy method presented in the last section, but it still produces many incorrect slip amplitude estimates.

3.5.1 Algorithm Description

Following the development in [85], we consider a state-space formulation for filtering the phase measurements in order to mitigate the impact of cycle slips. We collect the geodetically detrended phase for various signals into a measurement vector $\mathbf{y}$:

$$\mathbf{y}(t_i) = \begin{bmatrix} \Phi_1(t_i) & \dots & \Phi_k(t_i) \end{bmatrix}^{\mathrm{T}} \tag{3.1}$$

We consider a state vector $\mathbf{x}$ consisting of the filtered phase residuals $\bar{\Phi}_{f_i}$ along with their rate estimates $\nabla\bar{\Phi}_{f_i}$:

$$\mathbf{x}(t_i) = \begin{bmatrix} \bar{\Phi}_1(t_i) & \dots & \bar{\Phi}_K(t_i) & \nabla\bar{\Phi}_1(t_i) & \dots & \nabla\bar{\Phi}_K(t_i) \end{bmatrix}^{\mathrm{T}} \tag{3.2}$$

With these definitions for $\mathbf{y}$ and $\mathbf{x}$, we construct the measurement model:

$$\mathbf{y}(t_i) = \mathbf{H}\mathbf{x}(t_i) + \mathbf{b}(t_i) + \mathbf{v}(t_i) \tag{3.3}$$

where $\mathbf{H}$ is given by

$$\mathbf{H} = \begin{bmatrix} 1 & \dots & 0 & 0 & \dots & 0 \\ & \ddots & & \vdots & & \vdots \\ 0 & \dots & 1 & 0 & \dots & 0 \end{bmatrix} \tag{3.4}$$

and $\mathbf{b}(t_i) = \begin{bmatrix} \lambda_1 z_1(t_i) & \dots & \lambda_K z_K(t_i) \end{bmatrix}^T$ is an estimate of the bias due to cycle slips with λ_{f_i} as the wavelength corresponding to the i-th signal. We also assume the measurement noise $\mathbf{v}(t_i) \sim \mathcal{N}(\mathbf{0}, \mathbf{R}(t_i))$ with covariance matrix $\mathbf{R}(t_i) = \mathrm{diag}(\sigma_1^2(t_i), \dots, \sigma_K^2(t_i))$. Here the measurement variances are adjusted based on signal $\mathrm{C/N_0}$:

$$\sigma_i^2(t_i) = \frac{1}{2T \cdot \mathrm{C/N}_{0,i}(t_i)} \left(1 + \frac{1}{2T \cdot \mathrm{C/N}_{0,i}(t_i)}\right) \tag{3.5}$$

where $\mathrm{C/N_0}$ is estimated as

$$\mathrm{C/N_0}(t_i) = \frac{|\delta s(t_i)|^2}{\sigma_{N,f}^2 T} \tag{3.6}$$

This is assuming we have access to the complex correlation outputs δs (as introduced in Section 1.2.1 of Chapter 1) from our receiver. The nominal noise variance σ_N^2 can be estimated from the amplitude variance of $\delta s(t_i)$ during ideal conditions that occur either before or after scintillation or multipath occurrence.

The discrete-time state dynamics model is given by:

$$\mathbf{x}(t_{k+1}) = \mathbf{F}\mathbf{x}(t_i) + \mathbf{w}(t_i) \tag{3.7}$$

where we have the discrete-time transition matrix:

$$\mathbf{F} = \begin{bmatrix} 1 & \dots & 0 & T & \dots & 0 \\ & \ddots & & \vdots & \ddots & \\ 0 & \dots & 1 & 0 & \dots & T \\ 0 & \dots & 0 & 1 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & 0 & \dots & 1 \end{bmatrix} \tag{3.8}$$

and where $\mathbf{w}(t_i) \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$. Our adapted dynamics model leads to a slightly different discrete-time process noise covariance $\mathbf{Q}$ than that in [85]. We express it using auxiliary matrices $\mathbf{Q}_\Phi$ and $\mathbf{Q}_{\mathcal{G},\mathcal{I}}$:

$$\mathbf{Q}_\Phi = \mathrm{diag}\left(\sigma^2_{\Phi_1}, \dots \sigma^2_{\Phi_M}\right) \tag{3.9}$$

$$\mathbf{Q}_{\mathcal{G},\mathcal{I}} = \begin{bmatrix} \sigma^2_{\mathcal{G}} + \beta_1\beta_1\sigma^2_{\mathcal{I}} & \dots & \sigma^2_{\mathcal{G}} + \beta_1\beta_M\sigma^2_{\mathcal{I}} \\ \vdots & \ddots & \vdots \\ \sigma^2_{\mathcal{G}} + \beta_M\beta_1\sigma^2_{\mathcal{I}} & \dots & \sigma^2_{\mathcal{G}} + \beta_M\beta_M\sigma^2_{\mathcal{I}} \end{bmatrix} \tag{3.10}$$

$$\mathbf{Q} = \begin{bmatrix} T\mathbf{Q}_\Phi + \frac{T^3}{3}\mathbf{Q}_{\mathcal{G},\mathcal{I}} & \frac{T^2}{2}\mathbf{Q}_{\mathcal{G},\mathcal{I}} \\ \frac{T^2}{2}\mathbf{Q}_{\mathcal{G},\mathcal{I}} & T\mathbf{Q}_{\mathcal{G},\mathcal{I}} \end{bmatrix} \tag{3.11}$$

The parameters $\sigma^2_{\Phi_i}$, $\sigma^2_{\mathcal{G}}$, and $\sigma^2_{\mathcal{I}}$ are the values for the power spectral densities of the respective white-noise processes. The noise associated with $\sigma^2_{\Phi_i}$ accounts for unmodeled and uncorrelated errors amongst the different signals, while the noise associated with $\sigma^2_{\mathcal{G}}$, and $\sigma^2_{\mathcal{I}}$ come from random variations in the unmodeled non-dispersive and refractive ionosphere components that are present in all signals. The covariance matrix $\mathbf{Q}$ must be tuned to achieve the best possible cycle slip mitigation. If its values are too large, the filter will be too forgiving to large phase variations and will not be able to distinguish cycle slips. If its values are too small, the filter may become overconfident and diverge. We performed a grid search of various combinations of values between 0.00001 to 2 and selected those that produce the best performance for the Hong Kong dataset in the sense that it yielded the smallest number of slips in the resulting measurements (i.e. smallest number of missed slips plus added slips). For the results presented in this the next section, we use values of $\sigma^2_{\Phi_i} = 0.0001$, $\sigma^2_{\mathcal{G}} = 0.00001$, and $\sigma^2_{\mathcal{I}} = 0.00005$.

Having established the measurement and state dynamics models for our system, the remainder of the algorithm follows mostly standard Kalman filter procedure. The one exception is the estimation of the bias sequence $\mathbf{b}(t_i)$, which is done using the propagated state $\mathbf{x}^-(t_i)$ and measurements $\mathbf{y}(t_i)$ as:

$$\mathbf{b}(t_i) = \text{argmin}_{\mathbf{b}'} ||\mathbf{y}(t_i) - \mathbf{H}\mathbf{x}^-(t_i) - \mathbf{b}'|| \tag{3.12}$$

Again, the values of $\mathbf{b}$ are restricted to an integer number of wavelengths. The estimated bias $\mathbf{b}(t_i)$ is then subtracted from the measurement residual when updating $\mathbf{x}$.

Otherwise, a state estimation covariance $\mathbf{P}(t)$ is propagated via:

$$\mathbf{P}^-(t_{i+1}) = \mathbf{F}\mathbf{P}^+(t_i)\mathbf{F}^T + \mathbf{Q} \tag{3.13}$$

and the remaining update steps follow:

$$\mathbf{K}(t_i) = \mathbf{P}^-(t_i)\mathbf{H}^T \left(\mathbf{H}\mathbf{P}^-(t_i)\mathbf{H}^T + \mathbf{R}(t_i)\right)^{-1} \tag{3.14}$$

$$\mathbf{x}^+(t_i) = \mathbf{x}^-(t_i) + \mathbf{K}(t_i)(\mathbf{y}(t_i) - \mathbf{H}\mathbf{x}^-(t_i) - \mathbf{b}(t_i)) \tag{3.15}$$

$$\mathbf{P}^+(t_i) = (\mathbf{I} - \mathbf{K}(t_i)\mathbf{H})\,\mathbf{P}^-(t_i) \tag{3.16}$$

3.5.2 Assessment

We apply the filtering algorithm to the Hong Kong dataset after decimating our measurements to 20 Hz. We obtain estimates of the slip bias sequences $z(t_k)$ for each signal. These are plotted

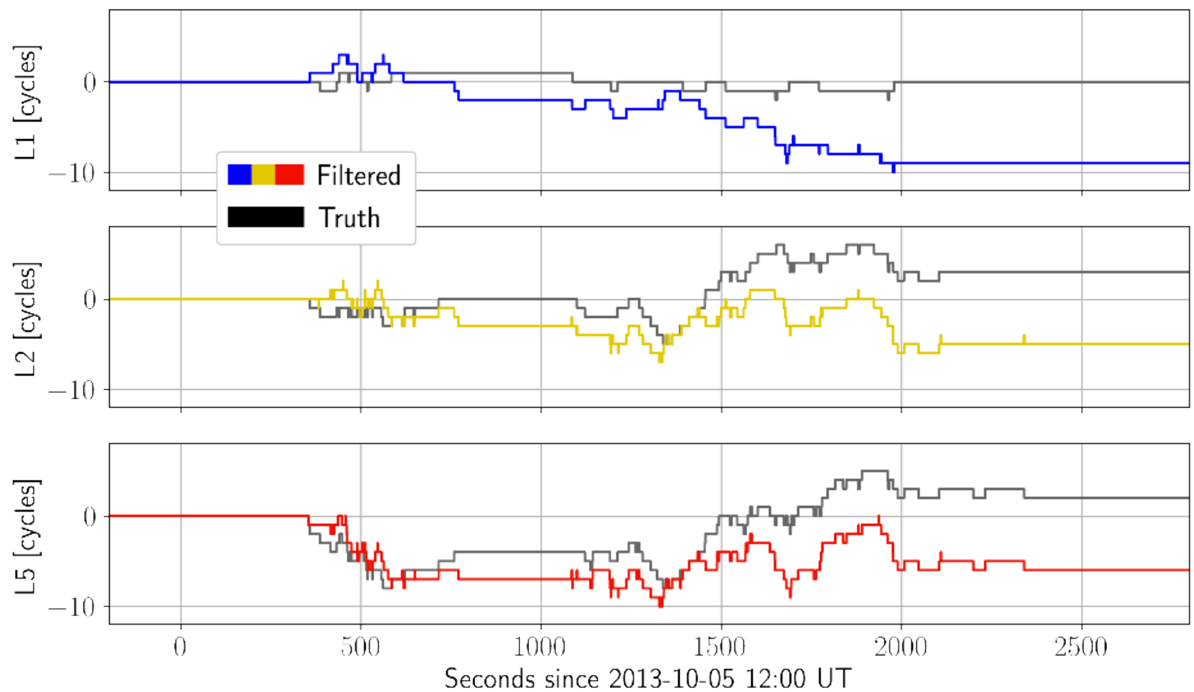

Figure 3.16: Comparison of truth reference of cycle slip bias sequence versus sequence estimated during the filtering process for L1, L2, and L5 signals (offset for clarity). There are 15, 63, and 78 slips in the truth reference and 26, 89, and 90 jumps in the filter estimate for the L1, L2, and L5 signals respectively. Of these, the filter estimated 8, 35, and 36 of the true jumps correctly.

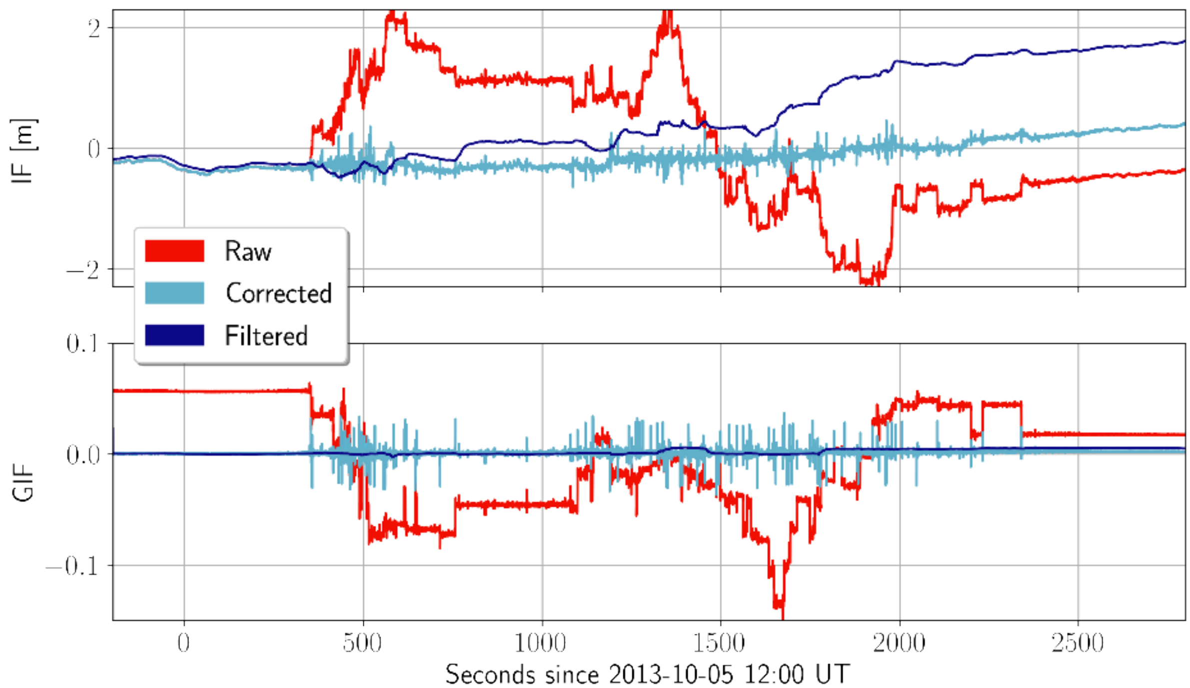

Figure 3.17: Comparison of IF and GIF combinations of the carrier phase estimates from the Hong Kong dataset before and after applying the filtering method.

in Figure 3.16 along with the bias sequences for the truth reference. The discrepancy between the true and estimated bias sequences can be due to cycle slips that the filter fails to remove or that the filter erroneously adds into its output. In this case, for the L1, L2, and L5 signals there are 15, 63, and 78 slips in the truth reference and 22, 89, and 90 slips in the filter estimate, respectively. Of these, the filter estimated 8, 35, and 42 of the true slips correctly and erroneously added 14, 54, and 48slips. Note then that for this dataset the filter adds more slips than it effectively mitigates, which raises some doubt as to its usefulness. However, some of these added slips may be due to rapid consecutive slips that were not included in the truth reference. For instance, between 43500 to 44000 seconds in Figure 3.16, all three phase bias estimates show occurrences where the estimated bias jumps and then almost immediately returns to its previous value. Due to the way we manually identified slips over a window of observations, rapid consecutive slips that did not result in a net bias were not always identified in the truth reference. Therefore, we argue that the number of slips added by the filter is moderately less than the reported number. After counting up the number of rapid consecutive (i.e. less than 1 second apart) slips and removing them from the total, we arrive at 8, 40, and 21 added slips.

An additional way to assess the removal of cycle slips is by observing the IF and GIF carrier phase combinations. Note again that since the non-dispersive phase components were removed, the IF combination should be mostly flat, as should the GIF combination. Figure 3.17 shows the IF (top panel) and GIF (bottom panel) combinations before and after applying the filtering algorithm. While there are some new jumps in the filtered IF and GIF combinations when compared to the raw versions due to added cycle slips, our qualitative assessment is that the phase combinations after filtering appear more flat overall. In particular, jumps in the GIF combination between 44500 and 45000 seconds have been significantly reduced. There are a couple beneficial qualities of the filtered output that we observe through these linear combinations: 1) the filtered phase is less noisy, and 2) many consecutive cycle slip occurrences have been removed by the filter. This second point can be seen particularly by observing the GIF combination, e.g. around 43800 seconds. The removal of these rapid consecutive cycle slips may be beneficial to window-based cycle slip detection methods that may otherwise struggle to deal with multiple slips occurring within one detection window. Otherwise, the fact that falsely added more slips than it correctly mitigated clearly indicates the need for a better algorithm.

Overall, when we consider the performance of this algorithm, it makes sense that it is only able to mitigate some of the slip occurrences in the scintillation measurements. The filter is designed to "filter out" cycle slips by reducing its gain when the signal amplitude fades. As such, any cycle slips due to phase transitions that occur during deep signal fades are likely to be correctly filtered out and estimated by the algorithm. On the other hand, in the last chapter we saw how strong scintillation contains many phase transitions with shallow fading amplitudes. Cycle slips due to these phase transitions will not be as effectively filtered out by the algorithm. In actuality, the performance of this algorithm on a given dataset is highly dependent upon the tuning of its process noise. Smaller process noise generally leads to filtering of more slips, but if it is too small then the algorithm will start introducing slips in order to accommodate variations in the actual signal phase components. Recall that we optimized the tuning to achieve the best performance for this particular dataset in terms of total cycle slips present in the filtered measurements. Even with the optimal tuning we saw this effect of added slips in the filtering results. Moreover, because the algorithm is sequential, there is no way for it to recognize if the phase deviations it observes are actually part of the phase component dynamics (e.g. the refractive ionosphere phase) or if they correspond to actual phase transitions. This is particularly the case when phase transitions occur over a longer duration or has a shallow fading amplitude, as we just discussed. Our main takeaway from this assessment is this: while the principles of filtering out cycle slips using an adaptive algorithm like the one described here make sense, ultimately one needs a full window of measurements surrounding a slip to adequately assess its occurrence. This will be a key point that motivates our work in Chapter 4 and Chapter 5.

3.6 Summary and Discussion

In this work, we analyzed the occurrence of GNSS cycle slips associated with signal fading. In particular, we looked at two datasets containing strong diffractive ionosphere scintillation. In both datasets, jumps in the IF and GIF combinations clearly indicated the presence of cycle slips. For the first dataset from Ascension Island, we analyzed various examples of canonical fades. We showed how signal intensity and jumps in various linear combinations of phase observables allow for identifying cycle slip occurrence. We applied this analysis to the second dataset from Hong Kong in order to estimate a truth reference for cycle slip occurrence. A representative result of our correction is shown in Figure 3.11, where any jumps in the IF and GIF combinations are essentially eliminated. Then we used this truth reference to assess the performance of two cycle slip mitigation algorithms. The first was the De Lacy algorithm: the single-epoch time-differenced cycle slip mitigation algorithm for triple-frequency signals that is described in [20]. Although the algorithm showed a 100 percent detection rate, large false-detection rates seem to be a consistent problem at each stage of the algorithm. Based on the detection results from Figure 3.14, we cannot reduce this false detection rate very much without introducing missed detections. The second algorithm was an adapted version of the filtering method from [85]. This algorithm showed better results than those from [20], but still ended up introducing more cycle slips than it accurately corrected, despite having chosen an optimal tuning of the algorithm's process noise covariance. We argue that the added slips in this case are due in large part to phase transitions with extended or shallow fading that can occur during strong scintillation or multipath, as we discussed in the last chapter.

Observing the difference residuals of the various linear combinations used in the De Lacy algorithm provided some insight into the several challenges that ionosphere diffraction presents to traditional methods of cycle slip mitigation: 1) increased code phase errors/noise, 2) increased carrier phase noise/modelling errors due to diffractive fluctuations, and 3) the rapid succession of multiple cycle slips. The first of these issues, increased code phase noise, is problematic because it somewhat hinders the computational benefit of addressing large cycle slips using code phase measurements. In this case, increased code phase noise contributes to large false-detection rates that reduces the utility of code phase for large slip detection. The second issue is perhaps an even larger contributor to the standard algorithm's false-detection rates. Ionosphere diffraction is associated with canonical fading and rapid phase fluctuations that can behave similarly regardless of whether a phase transition (cycle slip) actually occurs. In turn, it becomes uncertain whether these fast phase changes that show up in the residuals $\Delta\mathrm{GIF}_{\mathrm{L1,L2,L5}}$ and $\Delta\mathrm{GF}^{(1)}_{\mathrm{L1,L2,L5}}$ correspond to actual cycle slips. Similarly, we discussed how the adaptive filtering algorithm struggles to discriminate between actual phase component variations and the variations associated with phase transitions or cycle slips, as evidenced by its addition of multiple slips in the filtered results. Indeed, some sort of averaging is needed in order to reduce the impact of diffractive fluctuations and other unmodeled effects. Toward the end of Chapter 1, we discussed how window-based detection as a promising technique. However, the occurrence of several cycle slips in rapid succession can be a problem for window-based methods that usually assume the occurrence of only one cycle slip in a given window. It at least requires more careful consideration before being applied to mitigation of diffraction-induced cycle slips, including careful assessment of what measurement duration and sampling rate are necessary to correctly estimate slips. This will be a topic we explore in Chapter 4.

Chapter 4

Probabilistic Modeling of Cycle Slip Detection and Estimation

4.1 Introduction

In Chapter 1, we discussed several approaches to cycle slip mitigation. These types of cycle slip mitigation algorithms that test for outliers in time-differenced phase combinations have been successful under normal signal conditions. However, as we discussed in the last chapter, when it comes to more challenging conditions, uncertainty in the variation of signal phase components can cause most cycle slip algorithms to fail. The authors of [2] first discussed the particularly difficult problem of dealing with cycle slips during ionosphere plasma bubble events when phase measurements contain multiple cycle slips and unpredictable variations due to changes in ionosphere total electron content (TEC). Along with the authors in [45], [101], and [65], they emphasize the utility of IF phase combinations and careful estimation of non-dispersive phase components when it comes to detecting cycle slips in these scenarios. However, as noted in [3], even with careful estimation of non-dispersive components it may still be impossible to effectively estimate slip occurrences in 1 Hz dual-frequency measurements when the ionosphere variability is severe. Our analysis of cycle slips associated with phase transitions in the last chapter suggests that high-rate measurements over an extended window are actually necessary to reliably correct for cycle slips under the most challenging conditions.

In this chapter, we aim to bring together all our available carrier phase measurements into a general probabilistic framework for estimating cycle slips. In doing so we highlight the fundamental similarities between cycle slip estimation and ambiguity resolution as mixed-integer inference problems, and we frame the cycle slip problem as a trade-off between estimation fidelity and computational burden. An important aspect of this work is the quantification of how much information is actually necessary to adequately resolve cycle slips. We show that high-rate (>20 Hz) measurements over a window of at least 16 seconds is necessary for estimating slips during ionosphere scintillation. Although false estimations can still occur, this resolution of data substantially reduces their probability.

This chapter is divided into five remaining sections. In Section 4.2 we introduce a state-space model for cycle slip estimation for an arbitrary number of signals, measurement sampling rate, and window duration. In Section 4.3, we discuss modeling of the dispersive and non-dipsersive phase components as Gaussian processes, as well as ways for modeling the impact of noise. In Section 4.4, we discuss the cycle slip problem from a probabilistic perspective and derive the expression for the posterior distribution of cycle slip amplitudes given our measurements. In Section 4.5 we present results of characterizing the failure rates for slip amplitude estimation under a variety of harsh signal conditions and for different measurement rates and window durations. Finally, in Section 4.6 we summarize the results from this chapter.

4.2 System Model

In this chapter, we use the carrier phase measurement model that we introduced in Equation 1.5 from Chapter 1. Many approaches to cycle slip mitigation include code phase measurements in their formulation, e.g. the method we analyzed at the end of Chapter 2. However, because they are orders of magnitude more noisy than carrier phase, code phase measurements make little contribution towards estimation of small slip amplitudes, especially when there is any sort of uncertainty in non-dispersive or ionosphere phase components. Therefore, to simplify our analysis, we only consider carrier phase measurements in this work. However, in principle, the system model we present here can incorporate code phase measurements.

Here we describe the full set of measurement and state variables relevant to the cycle slip problem. We can vectorize our phase measurements for K signals transmitted from one satellite at a time epoch t:

$$\mathbf{y}(t) = \begin{bmatrix} \Phi_1(t) & \cdots & \Phi_K(t) \end{bmatrix}^T$$

where each Φ_k is modeled according to Equation 1.5. We denote a state vector consisting of the non-dispersive and refractive ionosphere phase components along with the fractional phase bias for each signal:

$$\mathbf{x}(t) = \begin{bmatrix} \mathcal{G}(t) & \mathcal{I}(t) & \mathcal{B}_1 & \dots & \mathcal{B}_K \end{bmatrix}^T$$

We collect the integer-cycle bias components into a separate vector:

$$\mathbf{z}(t) = \begin{bmatrix} z_1(t) & \cdots & z_K(t) \end{bmatrix}^T$$

Then we can write the following equation modeling measurements at a single time epoch:

$$\mathbf{y}(t) = \mathbf{A}_1\mathbf{x}(t) + \mathbf{B}_1\mathbf{z}(t) + \epsilon(t) \tag{4.1}$$

where

$$\mathbf{A}_1 = \begin{bmatrix} 1 & \beta_1 & 1 & & \\ \vdots & \vdots & & \ddots & \\ 1 & \beta_K & & & 1 \end{bmatrix}, \qquad \mathbf{B}_1 = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_K \end{bmatrix} \tag{4.2}$$

and where $\epsilon(t) = [\epsilon_1(t), \ \cdots \ , \ \epsilon_k(t)]^T$ is comprised of noise terms that we assume to be independent and normally distributed with zero-mean and variances σ_k^2. Note that if we want to estimate half-cycle instead of integer-cycle slips, we can change the entries in $\mathbf{B}_1$ to be half-wavelengths.

We extend this model to consider a window of measurement and state terms at a series of N_t times $\mathbf{t} = [t_1, \ \cdots \ , \ t_{N_t}]^T$, and adopt the notation $\Phi_k(\mathbf{t}) = [\Phi_k(t_1), \cdots, \Phi_k(t_{N_t})]^T$ (likewise for $\mathcal{G}$, $\mathcal{I}$, etc.). We express our full system model over multiple time epochs as:

$$\mathbf{y}(\mathbf{t}) = \mathbf{A}\mathbf{x}(\mathbf{t}) + \bar{\mathbf{B}}\mathbf{z}(\mathbf{t}) + \epsilon(\mathbf{t}) \tag{4.3}$$

where we have

$$\mathbf{A} = \left[\begin{bmatrix} 1 & \beta_1 \\ \vdots & \vdots \\ 1 & \beta_K \end{bmatrix} \otimes \mathbf{I}_{N_t} \;\middle|\; \mathbf{I}_K \otimes \mathbf{1}_{N_t} \right], \qquad \bar{\mathbf{B}} = \mathbf{B}_1 \otimes \mathbf{I}_{N_t}, \tag{4.4}$$

and where $\mathbf{I}_N$ is the $N \times N$ identity matrix, $\mathbf{1}_N$ is a length N vector of all ones, and $\otimes$ denotes the Kronecker product. As an example, the structure of the matrix $\mathbf{A}$ is illustrated in Figure 4.1a.

When it comes to addressing cycle slips, what we are actually interested in is the *change* in integer phase bias $z(t)$. To accommodate this, let $\mathbf{t'} = [t'_1, \cdots, t'_M]^T$ be a vector of times at which slips, with amplitudes given by $\Delta z_k(t')$, may potentially occur in our model. We define:

$$z_k(\mathbf{t}) = \mathbf{S}\Delta z_k(\mathbf{t'}) \tag{4.5}$$

where $\Delta z_k(\mathbf{t'})$ (which we call the *cycle slip amplitude sequence*) is analogous to the discrete derivative of $z_k(\mathbf{t})$ and $\mathbf{S}$ is analogous to a discrete integral, having entries

$$S_{i,j} = \begin{cases} 1 & \text{if } t_i \geq t'_j \\ 0 & \text{otherwise} \end{cases} \tag{4.6}$$

We then use this matrix to construct:

$$\mathbf{B} = \mathbf{B}_1 \otimes \mathbf{S} \tag{4.7}$$

As an example, Figure 4.1b depicts a hypothetical structure of $\mathbf{B}$ for triple-frequency measurements. Incorporating $\mathbf{B}$ and $\Delta\mathbf{z}$ into Equation 4.3 gives us our final model equation:

$$\mathbf{y}(\mathbf{t}) = \mathbf{A}\mathbf{x}(\mathbf{t}) + \mathbf{B}\Delta\mathbf{z}(\mathbf{t'}) + \epsilon(\mathbf{t}) \tag{4.8}$$

Up until this point we have explicitly referred to the times at which a quantity is evaluated using the vectors $\mathbf{t}$ and $\mathbf{t'}$. The purpose of this is to emphasize how $\mathbf{y}$ and $\Delta\mathbf{z}$ are generally evaluated at different sets of times. For simplicity of notation, in the remainder of the chapter we will often drop the $(\mathbf{t})$ and instead use just a single symbol to refer to a quantity evaluated at all applicable time epochs; e.g. we use $\mathbf{y}$ instead of $\mathbf{y}(\mathbf{t})$. The intent is to make things easier to read, and the actual dimensions of the objects should hopefully be clear from context.

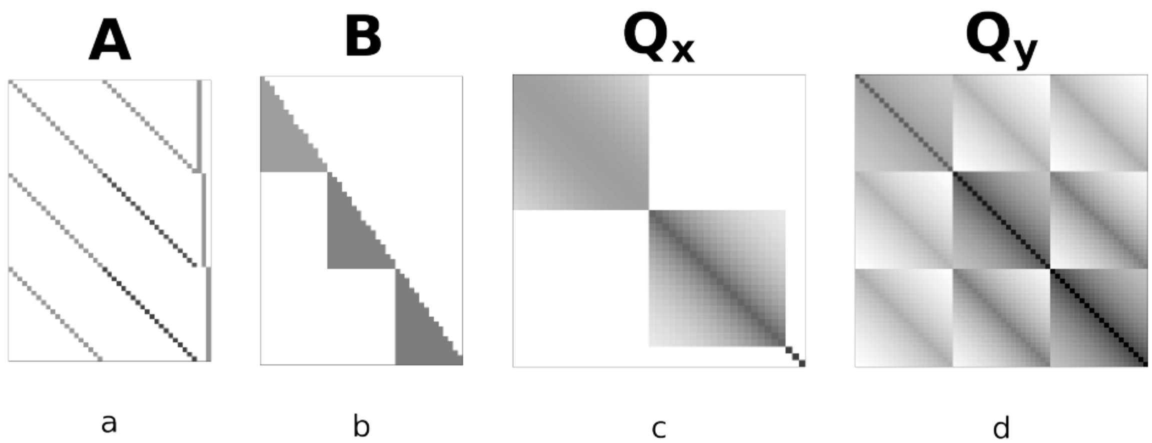

Figure 4.1: Illustrates the hypothetical structure of the various matrices relevant to a system model for triple-frequency signals. Panels a and b show the $\mathbf{A}$ and $\mathbf{B}$ matrices, respectively, from Equation 4.8. Panels c and d show the covariance matrices $\mathbf{Q_x}$ from Equation 4.15 and $\mathbf{Q_y}$ from Equation 4.23.

4.3 Distributions of $\mathcal{G}$, $\mathcal{I}$, and ϵ

In order to effectively estimate the occurrence of cycle slips, it is important to carefully model the behavior of the non-dispersive and ionosphere phase components. In this work, we assume the time series for $\mathcal{G}(t)$ and $\mathcal{I}(t)$ can be adequately modeled as zero-mean Gaussian processes (GP). That is, any sampling of their time series at a discrete set of points $\mathbf{t}$ is assumed to be drawn from an appropriate joint normal distribution. So, for the vectors representing non-dispersive and ionosphere phase time series, we have:

$$p(\mathcal{G}) = \mathcal{N}\left(\mathcal{G}; \mathbf{0}, \mathbf{Q}_{\mathcal{G}}\right) \tag{4.9}$$

$$p(\mathcal{I}) = \mathcal{N}\left(\mathcal{I}; \mathbf{0}, \mathbf{Q}_{\mathcal{I}}\right) \tag{4.10}$$

Here we use $p(\cdot)$ to denote the probability distribution for the random variable corresponding to its argument. We also denote the expression for the normal density with mean μ and covariance $\mathbf{Q}$ evaluated at $\mathbf{v} \in \mathbb{R}^N$ as:

$$\mathcal{N}\left(\mathbf{v}; \mu, \mathbf{Q}\right) = \left|2\pi\mathbf{Q}\right|^{-\frac{1}{2}} \exp\left[-\frac{1}{2}\left(\mathbf{v} - \mu\right)^T \mathbf{Q}^{-1}\left(\mathbf{v} - \mu\right)\right] \tag{4.11}$$

In addition to being GPs, we assume the phase components are time-stationary processes defined by their autocovariance kernels, which we express as:

$$q_{\mathcal{G}}(\tau) = E[\mathcal{G}(t)\mathcal{G}(t+\tau)] \tag{4.12}$$

$$q_{\mathcal{I}}(\tau) = E[\mathcal{I}(t)\mathcal{I}(t+\tau)] \tag{4.13}$$

Here E denotes expectation.

In other works on cycle slip mitigation, trends in these components have been modeled as linear [50], polynomial [21], or as autoregressive processes [98]. Compared to these approaches, GPs offer more flexible modeling of phase component behavior. Depending on the choice of autocovariance kernel, they can represent a variety of interpolation approaches, including those we just mentioned as well as splines and Fourier modes [46], [34]. Moreover, GP models can deal with missing or irregularly-sampled data [74], which can be a common occurrence for GNSS signals under harsh conditions. However, this flexibility tends to come at a computational cost. Fortunately, this cost can be reduced by taking advantage of the Toeplitz structure (i.e. constant along diagonal) of time-stationary covariance matrices, which allows us to efficiently compute matrix-vector products using the Fast Fourier Transform [75].

There are several standard options for parameterized autocovariance kernels. The exponential kernel corresponds to continuous but non-differentiable processes while the squared exponential kernel corresponds to processes that are infinitely differentiable. Realizations of these types of processes are often either too rough or too smooth to represent realistic data. The Matérn kernel with parameter ν generalizes these two kernels and provides a middle ground. It is equal to the exponential kernel at $\nu = 1/2$ and converges to the squared-exponential kernel as $\nu \to \infty$. It is given by:

$$k(\tau) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu}\frac{\tau}{\rho}\right)^{\nu} K_\nu \left(\sqrt{2\nu}\frac{\tau}{\rho}\right) \tag{4.14}$$

where Γ denotes the Gamma function, K_ν is the modified Bessel function of the second kind, σ^2 controls variance, and ρ is a scale parameter that controls the rate of probable fluctuations. For this work, we have heuristically chosen to use the Matérn kernel with parameter $\nu = 3/2$ to represent both non-dispersive and refractive ionosphere phase components. Typically the autocovariances can be obtained from their process' corresponding power spectral densities (PSDs). However, we want to show that our model is generally applicable, and we do not have good estimates for the PSDs of the phase components for all the scenarios we consider in Chapter 5. From our experience, the Matérn ($\nu = 3/2$) kernel achieves a good balance between long-term "memory" of the process trends while still allowing for small short-term fluctuations.

In addition to these models for $\mathcal{G}(t)$ and $\mathcal{I}(t)$, we assume that the fractional biases are drawn from independent Gaussian distributions with variance given by $\sigma_{\mathcal{B}}^2$, i.e. $\mathcal{B}_k \sim \mathcal{N}(0, \sigma_{\mathcal{B}}^2)$. Then the combined state $\mathbf{x}$ can be considered as one big zero-mean Gaussian process:

$$p(\mathbf{x}) = \mathcal{N}\left(\mathbf{x}; \mathbf{0}, \mathbf{Q}_{\mathbf{x}}\right) \tag{4.15}$$

where $\mathbf{Q}_{\mathbf{x}} = \text{blockdiag}\left(\mathbf{Q}_{\mathcal{G}},\ \mathbf{Q}_{\mathcal{I}}, \sigma_{\mathcal{B}}^2, \ldots, \sigma_{\mathcal{B}}^2\right)$. As an example, Figure 4.1c illustrates the hypothetical structure of $\mathbf{Q}_{\mathbf{x}}$.

4.3.1 Noise Models for Harsh Conditions

In addition to the state variables in $\mathbf{x}$, we take the measurement errors $\epsilon(t)$ to be zero-mean and normally distributed:

$$p(\epsilon) = \mathcal{N}\left(\epsilon; \mathbf{0}, \mathbf{Q}_{\epsilon}\right) \tag{4.16}$$

There are various types of noise that can impact our signal including thermal noise, oscillator jitter, receiver platform vibrations, and unmodeled errors such as ionosphere diffraction fluctuations. In this work, we will only consider the impacts of thermal noise and unmodeled errors (specifically those due to ionosphere diffraction), although the impacts of these other types of noise can be inferred from our results. In general, thermal noise will be important to consider for weak signal reflections or low-elevation satellites, whereas the impact of unmodeled random phase fluctuations will be most important for signals experiencing any amount of scintillation or scattering from structures in the atmosphere. Other types of noise, like jitter or vibrations, likely have a comparatively small impact (c.f. [86]). How these different types of noise manifest in our measurement model depends on the receiver system response, its output bandwidth, and how measurement downsampling or compression (if any) is performed. We will again make a simplifying assumption that the noise bandwidth is equal to $1/T$ where T is the sampling interval of our measurements.

In Section 2.2.2, we introduced the thermal noise process that defines η as circular complex white noise and how the power of this noise with respect to our signal is measured by $\mathrm{C/N_0}$. Usually we can approximate the resulting noise process in our phase measurements as real-valued discrete white noise and with variance (in units of cycles-squared) given by:

$$\sigma^2_{\text{thermal}} = \frac{1}{2\mathrm{C/N_0}T}\left(1 + \frac{1}{2\mathrm{C/N_0}T}\right) \tag{4.17}$$

However, this approximation does not work at low $\mathrm{C/N_0}$ or fast sampling rates. In particular, it breaks down when the signal-to-noise ratio is roughly less than 1. At a bandwidth of 10 Hz, this corresponds to a $\mathrm{C/N_0}$ of $\mathrm{SNR} \times \mathrm{BW} = 10$ dB-Hz, which is very low. Meanwhile, cycle slips tend to arise due to noise in phase unwrapping when the SNR reaches below 2.5, which corresponds to a $\mathrm{C/N_0}$ of 25 dB-Hz at 10 Hz bandwidth. This is also very low $\mathrm{C/N_0}$, but we still see it for weak reflection (like the GNSS-R example from Section 1.4.1 in Chapter 1), very low elevation or partially obstructed satellites, and during deep signal fading due to multipath or interference.

When it comes to diffraction-based noise processes, we can no longer necessarily consider the noise process to be white. Consider the top two panels of Figure 4.2, which show power spectral densities (PSD) for ionosphere diffraction error, which we obtain by averaging periodograms of the diffraction error from our ionosphere scintillation simulations. The left panel shows spectra for different scintillation strengths (as indicated by S_4) while the right panels shows the spectra for the same scintillation strength but with different decorrelation times (indicated by τ). There is a break in the spectrum just below 1 Hz, corresponding to the characteristic fluctuations of scintillation at around this frequency. Notice that there is a strong dependence of spectrum amplitude on scintillation strength. Meanwhile, the value of τ is related to the frequency of the break in the spectrum, but has little effect on the overall amplitude of the noise. Because of this, we consider scintillation strength as measured by S_4 to be the primary factor that determines the level of diffraction-induced phase noise.

While we saw in Chapter 2 how both thermal noise and diffraction compound to create cycle slips, when it comes to modeling phase noise usually one or the other is dominant. To show this, in the bottom panel of Figure 4.2 we plot the phase noise variance versus measurement bandwidth due to ionosphere diffraction for 3 different scintillation strengths, each of which we obtain by integrating the corresponding PSDs from zero up to the measurement sampling frequency. We also plot the phase variance due to thermal noise (dashed black line) for a C/N$_0$ of 20 and 25 dB-Hz according to Equation 4.17. Under our assumptions, we see that for strong scintillation the thermal noise variance will not exceed the diffraction noise variance, even for measurements/noise bandwidth at 100 Hz. For moderately strong scintillation ($S_4 = 0.6$) thermal noise can surpass diffraction noise variance at around 25 Hz, while for weak scintillation the diffraction-induced phase noise is weak and the thermal noise becomes dominant at only 5 Hz. Additionally, note how the diffraction noise variance is approximately constant for measurement bandwidth past the break frequency (around 0.5-1 Hz). Meanwhile, the thermal noise variance is highly dependent upon the measurement bandwidth. This means that when we go to model phase noise, the measurement sampling rate / bandwidth is very important when dealing with low-diffraction low-C/N$_0$ scenarios where thermal noise dominates. On the other hand, when diffraction noise is dominant, sampling rate will not be an important factor in determining phase noise variance above 1 Hz.

Based on this assessment, we can model σ_ϵ^2 as:

$$\sigma_\epsilon^2 = \sigma_{\mathrm{diffr}}^2 + \sigma_{\mathrm{thermal}}^2 \tag{4.18}$$

where $\sigma_{\mathrm{thermal}}^2$ is obtained as in Equation 4.17 and $\sigma_{\mathrm{diffr}}^2$ is obtained by integrating the appropriate PSD estimates as in Figure 4.2. The related work in [73] considers a similar but more intricate model for phase noise variance that accounts for PLL bandwidth and oscillator jitter. They also adjust the thermal noise variance to account for scintillation-induced fading. These more realistic noise models could prove useful when it comes to improving application to real data or if noise mis-modeling turns out to be a limiting factor. We will stick to our simple model in Equation 4.18 for the analysis we perform in this chapter.

One caveat worth mentioning is the correlation between diffraction fluctuations for signals at multiple frequencies. This model assumes uncorrelated noise on different signals, which is only somewhat true for the case for ionosphere scintillation [72]. A large portion of diffractive ionosphere fluctuations are correlated in a way that resembles the refractive ionosphere effect [17]. Similar correlations are likely present in the diffraction due to troposphere scintillation. Ultimately, how we model noise is highly application specific. Our model aims to be both powerful and general, and we will only consider the simple noise model introduced in Equation 4.16. This model will also be sufficient and effective for the detection and estimation approach we introduce in Chapter 5, although some benefit may come from more specific noise modeling under different scenarios, which is a promising topic for future investigation.

4.3.2 Fixing Model Hyperparameters

The phase component covariances $\mathbf{Q}_\mathbf{x}$ will be determined by the parameters $\sigma_\mathcal{G}^2$, $\rho_\mathcal{G}$, $\sigma_\mathcal{I}^2$, $\rho_\mathcal{I}$, and $\sigma_\mathcal{B}^2$ while $\mathbf{Q}_\epsilon$ is determined by the parameter σ_ϵ^2. Finding an appropriate covariance parameterization is an important aspect of the proposed cycle slip mitigation technique. If the covariance is too small then we underfit the data and risk many false alarms; if the covariance size is too big then we will overfit the data and risk missed detections. [1] In general, the covariances $\mathbf{Q}_\mathcal{G}$, $\mathbf{Q}_\mathcal{I}$, and $\mathbf{Q}_\epsilon$ will depend on the data at hand, and there are various approaches to model selection and hyperparameter optimization given that data [89].

One way to choose the model parameters is to simply hand-tune them. This is the approach we take when applying our algorithm to the ocean reflection and mountaintop RO datasets in Chapter 5 Section 5.7.2, since for those cases we do not have a larger set of data with which to validate our choices. For cases where the noise variance is mostly due to thermal noise, we can set σ_ϵ^2

[1] The one exception is $\sigma_\mathcal{B}^2$; this parameter is relatively less important, and simply needs to be chosen large enough to account for the initial bias in the phase measurements.

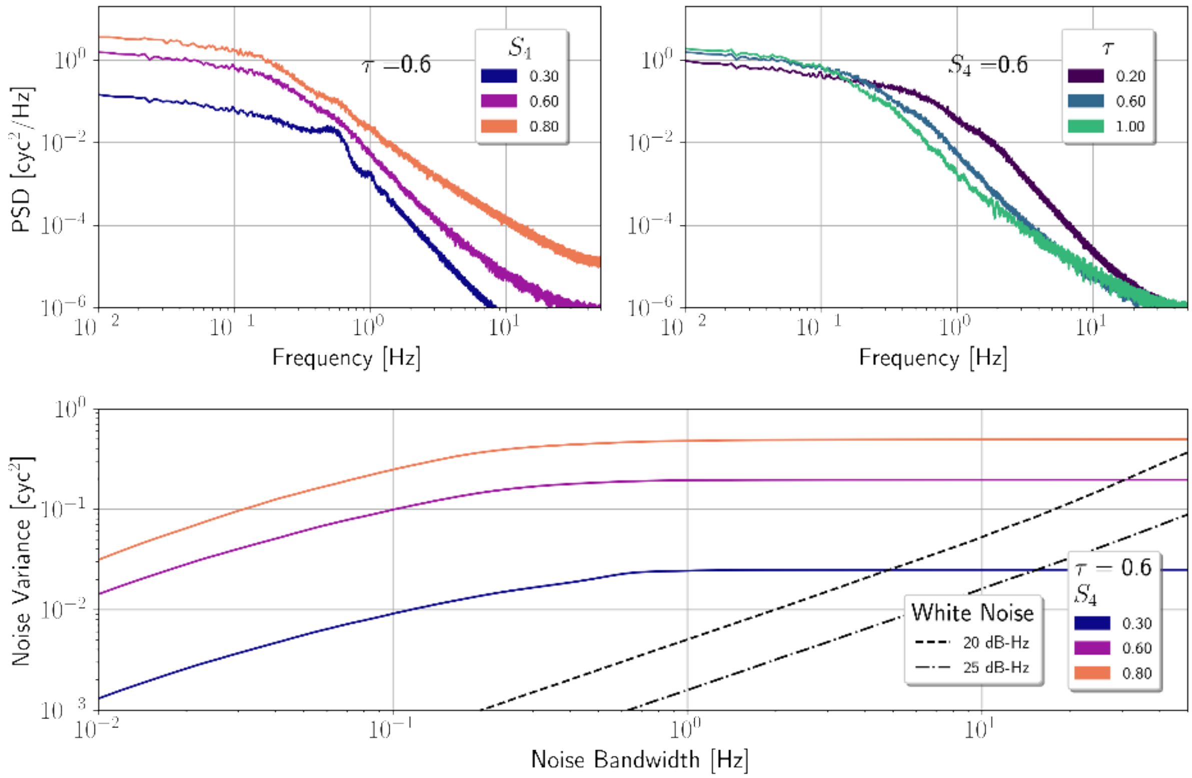

Figure 4.2: PSDs for ionosphere diffraction (top two panels) and corresponding phase noise variances for different bandwidth (bottom panel). Bandwidth is assumed equal to $1/T$ for sampling interval T.

using Equation 4.17. We can look at the behavior of "cycle-slip-free" parts of the signal phase and phase combinations in order to determine reasonable values for $\sigma^2_{\mathcal{G}}$, $\rho_{\mathcal{G}}$, $\sigma^2_{\mathcal{I}}$, and $\rho_{\mathcal{I}}$. Alternatively, when choosing model hyperparameters for scintillation measurements, we have at our disposal realistic simulations using the phase screen model that we introduced in Section 2.2.1 of Chapter 2. In this case, we can obtain our parameter values by performing a grid search over $\sigma^2_{\mathcal{I}}$, $\rho_{\mathcal{I}}$, and σ^2_ϵ in order to find the set of values that best explains the simulated measurements. In Appendix B, we discuss this process in more detail. Diffraction will be the dominant contribution to ϵ, and so another option for setting σ^2_ϵ is to base its value on the diffraction noise PSDs that we discussed in the last section. In actual signals, it is difficult to completely separate the diffractive and refractive phase fluctuations. This means there is some flexibility in how the model hyperparameters are tuned. If the refractive component is specified to be more smooth, we need to account for more noise by increasing σ_ϵ^2. On the contrary, we can allow the refractive component to fluctuate more by decreasing $\rho_\mathcal{I}$, and this in turn allows us to reduce the size of σ_ϵ^2. That is why, in general, it is best to choose these parameters simultaneously using our realistic scintillation simulations. Meanwhile, we choose $\sigma_\mathcal{G}^2$, $\rho_\mathcal{G}$, and $\sigma_\mathcal{B}^2$ heuristically. For instance, parameters for $\mathcal{G}$ can be chosen by observing the amplitude and scale of fluctuations in the IF combinations for real data, or they can be tuned to optimize some performance outcomes.

4.4 Probabilistic Estimation

Now we use our established system model from previous sections to present cycle slip estimation from a probabilistic perspective. Using the system model from Equation 4.8 and noise distribution from Equation 4.16, we obtain the following likelihood distribution for $\mathbf{y}$ given $\mathbf{x}$ and $\Delta\mathbf{z}$:

$$p(\mathbf{y}|\mathbf{x}, \Delta\mathbf{z}) = \mathcal{N}\left(\mathbf{y}; \mathbf{A}\mathbf{x} + \mathbf{B}\Delta\mathbf{z}, \mathbf{Q}_\epsilon\right) \tag{4.19}$$

We also have the prior distribution on $\mathbf{x}$ from Equation 4.15. As for the discrete-valued $\Delta\mathbf{z}$, we can express an arbitrary prior as:

$$p(\Delta\mathbf{z}) = \sum_{i=1}^{N_\chi} w_i^- \delta(\Delta\mathbf{z} - \Delta\mathbf{z}_i) \tag{4.20}$$

Here, each w_i^- is a weight corresponding to a particular cycle slip sequence $\Delta\mathbf{z}_i$ that can be thought of as a point on the integer lattice: $\mathbb{Z}^{K \cdot N_{t'}}$. For practical reasons, we assume $p(\Delta\mathbf{z})$ is only nonzero on some admissible subset this lattice, which we denote χ. This set, which has N_χ elements, should initially be large enough to contain any possible cycle slip sequence. Similar to [22] and [29], we can assume independence between $\mathbf{x}$ and $\Delta\mathbf{z}$ so that:

$$p(\mathbf{x}, \Delta\mathbf{z}) = p(\mathbf{x})p(\Delta\mathbf{z}) \tag{4.21}$$

Using these likelihood and prior distributions we can express the joint distribution of $p(\mathbf{y}, \mathbf{x}, \Delta\mathbf{z})$, and through orthogonal decomposition we can derive the expression for $p(\mathbf{y}|\Delta\mathbf{z})$. Similar derivations in the context of GNSS ambiguity resolution are mentioned in [80] and [29]. The resulting steps can be summarized as:

$$
\begin{aligned}
p(\mathbf{y},\mathbf{x},\Delta\mathbf{z}) &= p(\mathbf{y}|\mathbf{x},\Delta\mathbf{z})p(\mathbf{x})p(\Delta\mathbf{z}) \\
&= \mathcal{N}\left(\mathbf{y};\mathbf{A}\mathbf{x}+\mathbf{B}\Delta\mathbf{z},\mathbf{Q}_\epsilon\right) \\
&\qquad \cdot \mathcal{N}\left(\mathbf{x};\mathbf{0},\mathbf{Q}_\mathbf{x}\right)p(\Delta\mathbf{z}) \\
&= \mathcal{N}\left(\mathbf{x};\mu_{\mathbf{x}|\mathbf{y},\Delta\mathbf{z}},\mathbf{Q}_{\mathbf{x}|\mathbf{y}}\right) \\
&\qquad \cdot \mathcal{N}\left(\mathbf{y};\mu_{\mathbf{y}|\Delta\mathbf{z}},\mathbf{Q}_\mathbf{y}\right)p(\Delta\mathbf{z}) \\
&= p(\mathbf{x}|\mathbf{y},\Delta\mathbf{z})p(\mathbf{y}|\Delta\mathbf{z})p(\Delta\mathbf{z})
\end{aligned}
\tag{4.22}
$$

where

$$\mathbf{Q}_\mathbf{y} = \mathbf{A}\mathbf{Q}_\mathbf{x}\mathbf{A}^T + \mathbf{Q}_\epsilon \tag{4.23}$$

$$\mathbf{Q}_{\mathbf{x}|\mathbf{y}} = \mathbf{Q}_\mathbf{x} - \mathbf{Q}_\mathbf{x}\mathbf{A}^T\mathbf{Q}_\mathbf{y}^{-1}\mathbf{A}\mathbf{Q}_\mathbf{x} \tag{4.24}$$

$$\mu_{\mathbf{x}|\mathbf{y},\Delta\mathbf{z}} = \mathbf{Q}_{\mathbf{x}|\mathbf{y}}\mathbf{A}^T\mathbf{Q}_\epsilon^{-1}(\mathbf{y}-\mathbf{B}\Delta\mathbf{z}) \tag{4.25}$$

$$\mu_{\mathbf{y}|\Delta\mathbf{z}} = \mathbf{B}\Delta\mathbf{z} \tag{4.26}$$

As an example, the structure of the matrix $\mathbf{Q}_\mathbf{y}$ for triple-frequency signals is illustrated in Figure 4.1d.

It follows from Bayes' rule that:

$$
\begin{aligned}
p(\Delta\mathbf{z}|\mathbf{y}) &= \frac{p(\mathbf{y}|\Delta\mathbf{z})p(\Delta\mathbf{z})}{p(\mathbf{y})} \\
&= \frac{p(\mathbf{y}|\Delta\mathbf{z})p(\Delta\mathbf{z})}{\sum_i p(\mathbf{y}|\Delta\mathbf{z}_i)p(\Delta\mathbf{z}_i)}
\end{aligned}
\tag{4.27}
$$

Using Equation 4.20, we can rewrite Equation 4.27 as:

$$p(\Delta\mathbf{z}|\mathbf{y}) = \sum_i w_i^+\delta(\Delta\mathbf{z}-\Delta\mathbf{z}_i) \tag{4.28}$$

where

$$w_i^+ = \frac{w_i^-\mathcal{N}\left(\mathbf{y};\mathbf{B}\Delta\mathbf{z}_i,\mathbf{Q}_\mathbf{y}\right)}{\sum_j w_j^-\mathcal{N}\left(\mathbf{y};\mathbf{B}\Delta\mathbf{z}_j,\mathbf{Q}_\mathbf{y}\right)} \tag{4.29}$$

What Equation 4.29 essentially tells us is that the posterior probability of a particular cycle slip sequence is proportional to its prior time a "modified" likelihood that takes into account the modeled smoothness of phase components through $\mathbf{Q_y}$.

4.4.1 Integer Least Squares

Here we draw a connection between the posterior derived in the previous section and the integer least-squares (ILS) solution, which is an approach to optimizing discrete-valued variables that has commonly been used to resolve GNSS carrier phase ambiguities [78]. This connection has previously been pointed out in [22] and [80]. In ILS, we first relax the integer constraint on $\Delta\mathbf{z}$, and refer to the new float-valued slip amplitude sequence as $\Delta\hat{\mathbf{z}}$. Then, assuming an improper flat prior $p(\Delta\hat{\mathbf{z}}) \propto 1$, we obtain the following posterior distribution:

$$p\left(\Delta\hat{\mathbf{z}} \mid \mathbf{y}\right) = \mathcal{N}\left(\Delta\hat{\mathbf{z}}; \mu_{\Delta\hat{\mathbf{z}}}, \mathbf{Q}_{\Delta\hat{\mathbf{z}}}\right) \tag{4.30}$$

$$\mathbf{Q}_{\Delta\hat{\mathbf{z}}} = \left(\mathbf{B}^T\mathbf{Q}_\mathbf{y}^{-1}\mathbf{B}\right)^{-1} \tag{4.31}$$

$$\mu_{\Delta\hat{\mathbf{z}}} = \mathbf{Q}_{\Delta\hat{\mathbf{z}}}\mathbf{B}^T\mathbf{Q}_\mathbf{y}^{-1}\mathbf{y} \tag{4.32}$$

The approach of ILS is to find the integer-valued argument that maximizes the expression for the density in Equation 4.30. A visual interpretation of this solution is the integer-valued point that lies closest to the float least-squares estimate $\mu_{\Delta\hat{\mathbf{z}}|\mathbf{y}}$ under the metric induced by $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}$. Figure 4.3a illustrates the "pull-in" regions that define which float amplitude estimates are mapped to which integer grid points. For active ionosphere conditions the pull-in regions are stretched out along a particular direction due to the uncertainty in the ionosphere phase components. When there is no correlated uncertainty in the phase components, these pull-in regions will just look like square boxes surrounding each grid point.

The connection between the ILS solution and the posterior from Equation 4.27 stems from the realization that the following expressions are proportional (with respect to argument $\Delta\mathbf{z}_i$):

$$\mathcal{N}(\mathbf{y}; \mathbf{B}\Delta\mathbf{z}_i, \mathbf{Q}_\mathbf{y}) \propto \mathcal{N}(\Delta\mathbf{z}_i; \mu_{\Delta\hat{\mathbf{z}}}, \mathbf{Q}_{\Delta\hat{\mathbf{z}}}) \tag{4.33}$$

so we can rewrite Equation 4.29 as:

$$w_i^+ = \frac{w_i^- \mathcal{N}\left(\Delta \mathbf{z}_i; \mu_{\Delta \hat{\mathbf{z}}}, \mathbf{Q}_{\Delta \hat{\mathbf{z}}}\right)}{\sum_{j=1}^{N_\chi} w_i^- \mathcal{N}\left(\Delta \mathbf{z}_j; \mu_{\Delta \hat{\mathbf{z}}}, \mathbf{Q}_{\Delta \hat{\mathbf{z}}}\right)} \tag{4.34}$$

We expound on the result in Equation 4.33 in Appendix C. The result of Equation 4.34 shows us that when the prior distribution for $\Delta \mathbf{z}$ is flat, i.e. $w_i^- \propto 1$, the argument $\Delta \mathbf{z}_i$ that maximizes the above expression is equal to the ILS estimate. In other words, when the prior on $\Delta \mathbf{z}$ is non-informative, the maximum posterior (MAP) and ILS solutions coincide. The distribution describing the posterior in this case is sometimes called the "discrete Gaussian" or "discrete normal" and can be denoted:

$$\mathcal{DN}\left(\Delta \mathbf{z}_i; \mu_{\Delta \hat{\mathbf{z}}}, \mathbf{Q}_{\Delta \hat{\mathbf{z}}}\right) = \frac{1}{S} \exp\left[\left(\Delta \mathbf{z}_i; \mu_{\Delta \hat{\mathbf{z}}}\right)^T \mathbf{Q}_{\Delta \hat{\mathbf{z}}}^{-1} \left(\Delta \mathbf{z}_i; \mu_{\Delta \hat{\mathbf{z}}}\right)\right] \tag{4.35}$$

where S is the normalizing constant.

ILS has been used to frame the problem of carrier phase ambiguity resolution for a long time, and it is no surprise that estimation of cycle slips shares the same underlying problem structure. Inference on these types of problems is NP-hard in the general case, which means computing or maximizing the posterior involves a brute-force search over the space of $\Delta \mathbf{z}$. There are a number of methods to help make the search process more efficient. As we mentioned in Chapter 1, many ambiguity resolution and cycle slip algorithms apply the least-squares ambiguity decorrelation and adjustment (LAMBDA) when estimating carrier ambiguities or slip amplitudes. The decorrelation aspect in LAMBDA actually refers to a particular implementation of a lattice reduction algorithm. These types of algorithms can make searching over the integer solutions more efficient by finding new bases of the integer lattice that are approximately orthogonal under the metric induced by $\mathbf{Q}_{\Delta \mathbf{z}}$. The second aspect of the LAMBDA approach to ILS involves a search to find the best integer parameter vector in the least-squares sense. We discuss this aspect in more detail in Section 5.5 of Chapter 5 where we introduce a modified search algorithm designed for the windowed cycle slip problem.

4.5 Simulated Cycle Slip Estimation Performance

With our system model established, we now try to answer an elusive but very important question: how well can we actually estimate cycle slip occurrences under harsh conditions? And moreover, how much data is actually necessary to do so reliably? Specifically, it is useful to know for what duration and at what sampling rate we need measurements in order to achieve high probability of correctly estimating cycle slip amplitudes. Similar questions regarding cycle slip estimation performance in the presence of uncertainty in dispersive and non-dispersive phase components or measurement noise have been addressed for single-epoch methods in [3] and [90]. However, no study has addressed this question as it relates to a window of observations at different sampling rates.

We will address four different scenarios corresponding to the model parameters described in Table 4.1. The first three scenarios correspond to mild, medium, and strong scintillation levels at a high $\mathrm{C/N_0}$ of 50 dB-Hz. The fourth scenario is mild scintillation with a low $\mathrm{C/N_0}$ of 25 dB-Hz. Here scintillation strength is completely determined by S_4 and we keep τ constant. Note that the parameters for $\sigma_{\mathcal{G}}^2$ and $\rho_{\mathcal{G}}$ are also kept constant and correspond to very smooth, well-behaved non-dispersive phase components. We obtain the values for σ_ϵ^2 by adding the phase noise variances for thermal and diffraction-induced noises as described in Section 4.3.1. Note in particular that this means the noise variance is bandwidth dependent and will change for different sampling frequencies. The value of σ_ϵ^2 at 10 Hz sampling rate is shown in the table.

In order to demonstrate model strengths under these different scenarios, we first consider the case of estimating the amplitude of a single cycle slip in the middle of a window of observations; i.e. the vector $\mathbf{t}'$ from Equation 4.5 contains a single element, which is the time at the middle of the window. We assume measurements modeled according to Equation 4.19 and for model parameters corresponding to the different scenarios in 4.1. Figure 4.3a depicts pull-in regions that arise when using a 10-second window of 5 Hz dual-frequency measurements assuming model parameters from scenarios 1 and 3. Each region is colored according to the probability of identifying the corre-

	Units	Scenario 1	Scenario 2	Scenario 3	Scenario 4
S_4		0.3	0.6	0.9	0.3
τ	s	0.6	0.6	0.6	0.6
C/N$_0$	dB-Hz	50.0	50.0	50.0	25.0
$\sigma^2_{\mathcal{G}}$	m^2	0.1	0.1	0.1	0.1
$\rho_{\mathcal{G}}$	s	40.0	40.0	40.0	40.0
$\sigma^2_{\mathcal{I}}$	m^2	1.0	1.0	1.0	1.0
$\rho_{\mathcal{I}}$	s	60.0	9.0	5.0	60.0
σ^2_{ϵ}	cyc^2 @ 10 Hz	0.13	0.20	0.47	0.26

Table 4.1: Characteristics and hyperparameter values for different scenarios corresponding to a variety of signal conditions.

sponding slip amplitude, given that no slip has occurred. As expected in this situation, the correct cell (0,0), which is outlined in dark gray, has the highest probability of being identified. However, certain slip amplitudes like (1,0) show a non-negligible probability of identification (denoted P_{id}). Assuming that our measurement model is accurate, summing up the probabilities of these other cells (or equivalently taking $1 - P_{\text{id}}$) produces the probability of false identification (denoted P_{fa}). We consider P_{fa} as a measure for how well or how poorly a particular model can estimate cycle slip occurrence.

For each of the four scenarios, we consider models for single-, dual-, and triple-frequency signals over a range of measurement sampling rates and window durations. Figure 4.4 shows the P_{fa} computed for the different model structures and for each signal scenario. In general, we see decreased P_{fa} for higher sampling rates, longer window durations, and more signal frequencies. It is clear that the biggest factor determining model performance is the presence of diffraction-induced phase noise. For single-frequency signals the situation is particularly dire, with false identification probabilities staying above 30% for the case with the strongest scintillation. For dual- and triple-frequency signals, the false identification probability can still be brought down to roughly 10-20% when using at least a 16-second window of 20-50 Hz measurements. For low-C/N$_0$ conditions, we also note that around 16-20 second windows were necessary to achieve the best performance. Meanwhile, increasing sampling rate improved performance under all scintillation conditions. Under all the scenarios, it was not possible to reliably estimate slips in the 0.5 or 0.2 Hz measurements.

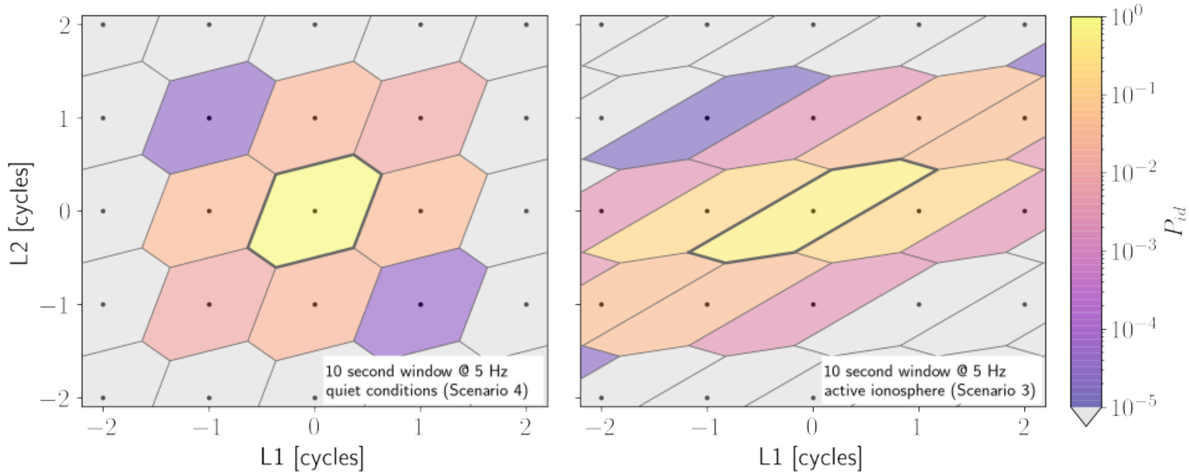

Figure 4.3: Pull-in regions and corresponding probabilities for quiet (left panel) and disturbed (right panel) signal conditions. Each region is colored according to the probability of the float estimate $\mu_{\Delta\hat{\mathbf{z}}}$ lying in that region, given that no slip has actually occurred.

4.6 Summary

In this chapter, we established a general approach to probabilistic inference on the occurrence of cycle slips in GNSS phase measurements. Along the way, we introduced Gaussian process models for phase component time series and noted their improved flexibility when compared to other models. We also discussed the contributions of both thermal and diffraction-induced noise to the phase measurement variance, which we argue are the two most relevant noise sources under harsh conditions. We derived expressions for the posterior distribution of cycle slip amplitudes, and we showed how the MAP estimate is the same as the ILS solution when slip priors are uniform. Finally, we used our model to simulate and assess cycle slip estimation performance under various signal conditions, the results of which were shown in Figure 4.4.

There are many different scenarios for signal conditions that will correspond to varying capabilities when it comes to estimating cycle slip occurrences, and the results from Figure 4.4 only cover a selection of representative examples. Similar analyses can be carried out for different sets of signals and signal conditions. As for the case of diffractive ionosphere scintillation, we see that in order to achieve reasonably low probabilities of false identification we should use at least a 16

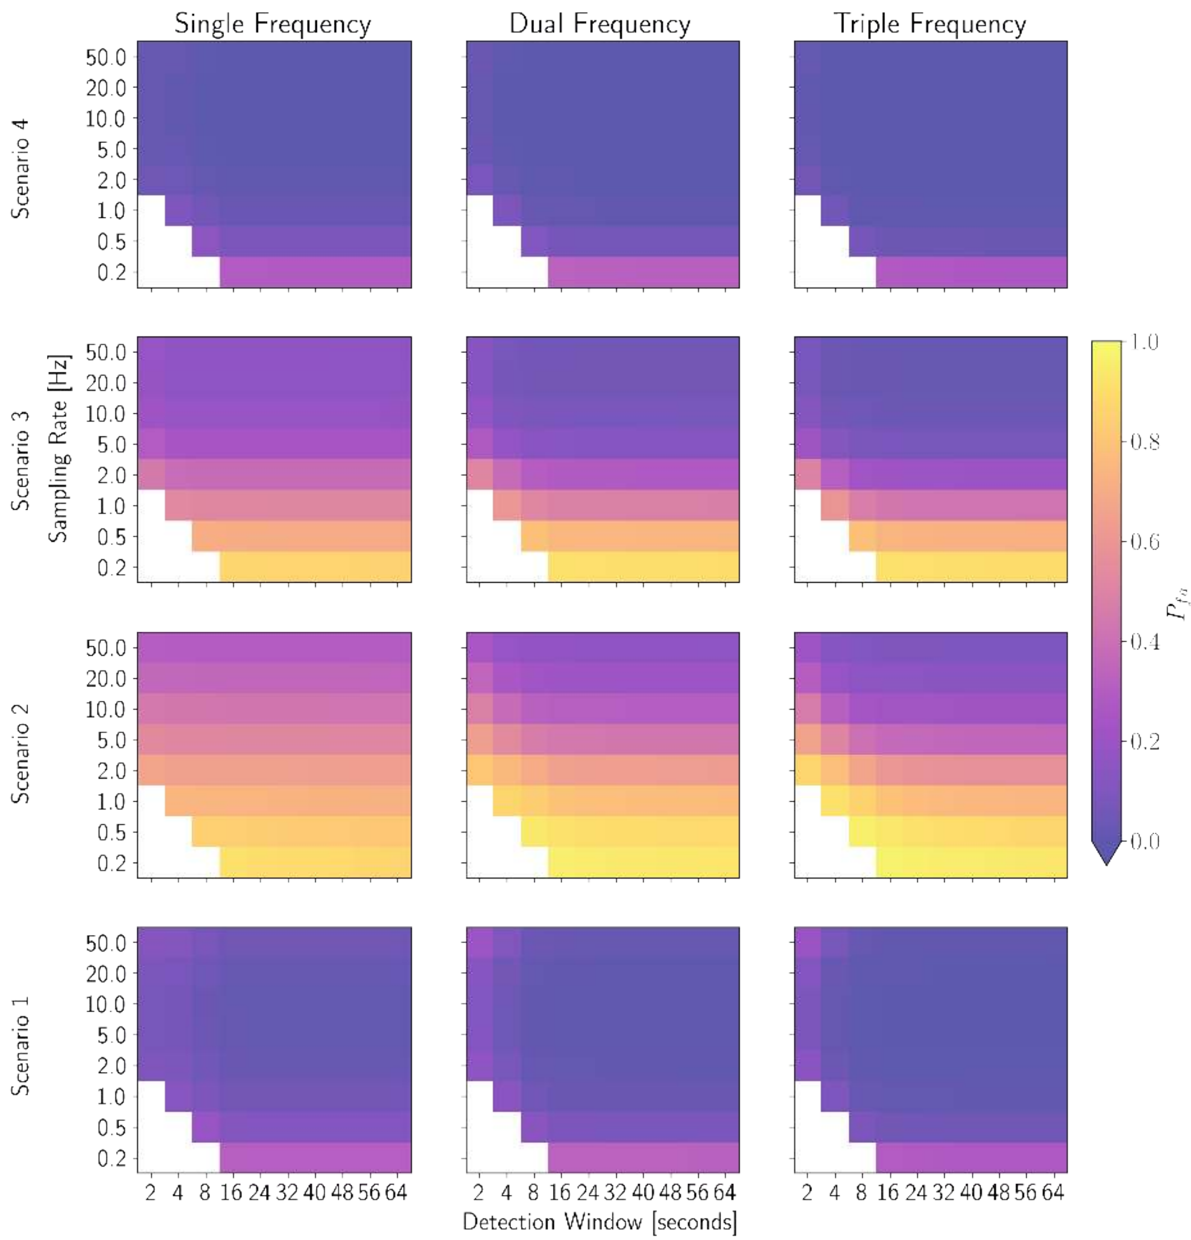

Figure 4.4: Probability of false slip amplitude identification under different signal conditions and for different model structures. Each row corresponds to signal conditions with model hyperparameters chosen according to Table 4.1. The columns show results for models using single-, dual- and triple-frequency measurements.

seconds of measurements sampled at 20 Hz or faster. This will be a motivating premise driving our development of cycle slip detection and estimation methods in the next chapter.

Chapter 5

Novel Batch Algorithm for Cycle Slip Detection and Estimation

In chapters 2 and 3, we saw how harsh signal conditions, ionosphere diffraction specifically, produce cycle slips at a rate of tens of slips per minute and how they are challenging to mitigate. As a concrete example, Figure 5.1 shows a case in the Hong Kong data where nearly a dozen slips occur over the span of 10-20 seconds, with occurrences highlighted in the vertical shaded regions. Meanwhile, our analysis from the last chapter indicates that we need at least that much measurement data to reliably estimate the slip occurrence. This poses an interesting question of how to deal with the multitude of cycle slips that occur under harsh signal conditions. In particular, one might initially consider the potential for cycle slips to occur between any two measurement epochs. In this case, when considering a window of high-rate measurement and slip epochs the problem becomes very high-dimensional and the time-adjacent float estimates will be highly correlated (e.g. consider the difference between estimating a slip occurring at 1.0 seconds versus 1.1 seconds). Even worse, when numerous cycle slips occur in sequence, as is the case under harsh signal conditions, the pairwise dependence between time-adjacent slip amplitudes results in a model that is coupled over the entire estimation window. The various authors that have addressed window-based cycle slip detection (e.g. [21], [14] [49]) all make restrictive assumptions about the number of slips that can occur in a given window. Given the random and chaotic behavior that can occur under harsh signal conditions, such rigid assumptions are ultimately a hindrance to effective mitigation.

In this chapter, we introduce a batch estimation algorithm that uses a window of measurements at arbitrary sampling rate for any number of carrier frequencies. In particular, we propose

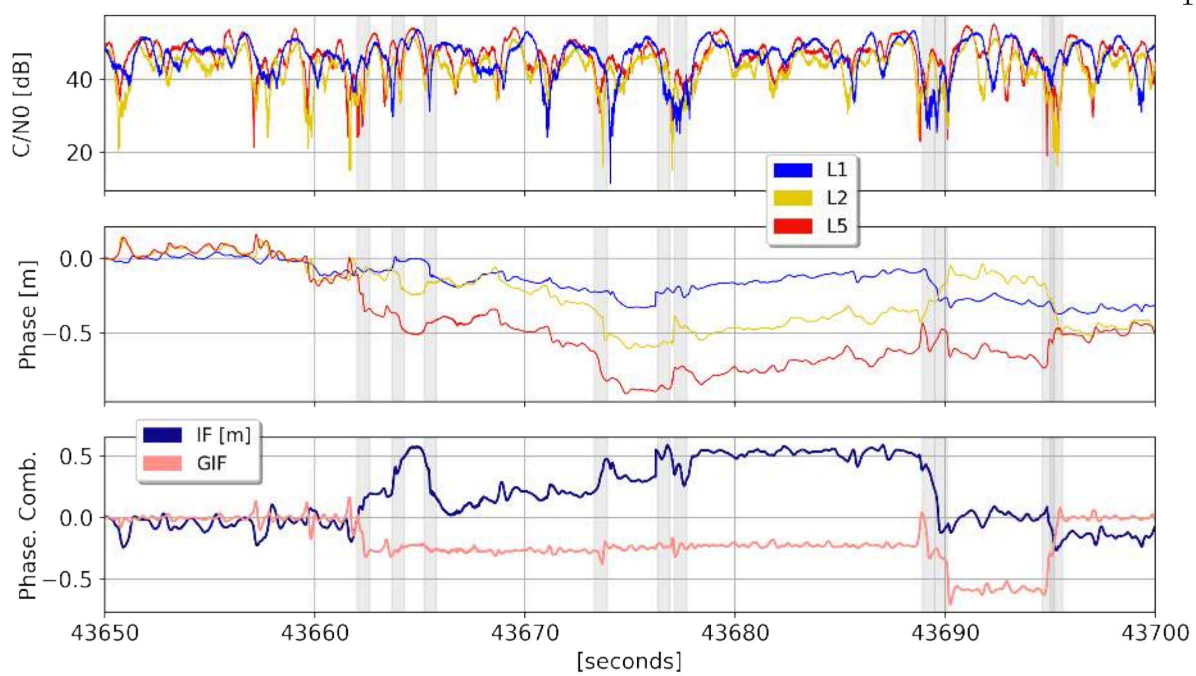

Figure 5.1: Example of cycle slips occurring in the Hong Kong scintillation dataset for GPS PRN 24 on 2013-10-03. The top panel shows signal C/N_0, the middle panel shows detrended carrier phase measurements, and the bottom panel shows the triple-frequency IF and GIF phase combinations. The shaded regions indicate the occurrences of cycle slips.

that a better way to formulate the window-based cycle slip estimation is by using an appropriate prior on $\Delta\mathbf{z}$ that is not actually uniform but instead reflects the sparsity of slip occurrence. Our characterization of slip occurrences as Poisson processes in Chapter 2 provides good empirical justification for this interpretation. As we demonstrate in this chapter, introduction of a sparse prior naturally leads to a way of *detecting* cycle slip occurrences over a window of observations. By first detecting slip occurrences we can effectively restrict the support of the posterior to only the detected epochs, thereby reducing the dimension and making the inference problem more tractable. Once we have our model with reduced dimension, we can search for the most probable cycle slip amplitude sequence. For this last step, we introduce new and adapted approaches to finding the optimal integer parameters in the high-dimensional ILS problem associated with these detected slip occurrences. Overall, this approach can be summarized in these four steps: 1) We compute a float estimate of $\Delta\mathbf{z}$ that incorporates a sparsity-inducing prior. 2) We detect slips as maxima in norm of this sparse float estimate. 3) We compute a reduced system model using the detected slips from the previous step. 4) We search for the best integer candidates in the reduced model to find the MAP cycle slip amplitude sequence. We then use the MAP estimate to correct for cycle slip occurrences.

The results of applying our algorithm to simulated and real scintillation data sets indicate that it is effective, however it also has some downsides. This method is computationally demanding, especially when compared to cycle slip algorithms that are used in practice. It is meant to be a thorough approach to cycle slip inference that is appropriate for post-processing or that can be used as a benchmark for design of other algorithms. Further work will be necessary to see how its principles adapt to sequential and near real-time processing. Also, we focus on use of simulations of ionosphere scintillation in order to tune our algorithm, and the parameters we obtain may or may not be suitable for use in other harsh signal scenarios. Nevertheless, the principles behind the algorithm that we introduce are applicable to a wide variety of contexts, and we demonstrate its effectiveness for different harsh signal conditions in our results.

This remainder of this chapter is divided into 4 sections. Section 5.1 briefly discusses the model we introduced in Chapter 4 and how we can tune its hyperparameters for the Hong Kong scintillation dataset, which is the primary focus when evaluating our results at the end of the chapter. Then, in Section 5.2 we describe our method for obtaining the sparse float estimate of cycle slips, followed by Section 5.3 that describes the detection of cycle slips occurrences using the sparse float estimate. In Section 5.4 we briefly describe the reduced estimation model that is obtained after detecting the slips. In Section 5.5 we provide background on the LAMBDA search algorithm for determining integer ILS parameters, and then we introduce two new approaches to search solutions to our high dimensional ILS problem corresponding to the reduced estimation model. We provide a summary of the algorithm in Section 5.6 and present results of applying it to simulated and real cycle slip datasets in Section 5.7. Finally, we provide a summary and discussion in Section 5.8.

	Units	Window 1	Window 2
S_4		0.86	0.71
τ	s	0.53	0.7 3
C/N$_0$	dB-Hz	48.00	49.00
$\sigma^2_{\mathcal{G}}$	m^2	0.001	0.001
$\rho_{\mathcal{G}}$	s	60.00	60.00
$\sigma^2_{\mathcal{I}}$	m^2	0.80	0.60
$\rho_{\mathcal{I}}$	s	7.00	9.00
σ^2_{ϵ}	cycles^2	0.20	0.16

Table 5.1: Scintillation window characteristics and hyperparameter values

5.1 Hyperparameter Tuning: Hong Kong Dataset

When demonstrating our results in real measurements, we will mainly focus on two windows of the Hong Kong scintillation dataset. The top two panels of Figure 5.2 show the C/N$_0$ and detrended phase measurements for that dataset, with the two windows of interest depicted by the shaded regions. The bottom two panels show the scintillation index S_4 and decorrelation time τ, respectively, which are related to the cycle slip occurrence rates as we discussed in Chapter 2. Table 5.1 lists the average S_4, τ, and baseline C/N$_0$ for the L1 signal that describe the signal conditions for the two windows in Figure 5.2. As discussed in Section 4.3.2 of Chapter 4, we can use measurement simulations in order to tune our model hyperparameters. As such, we use the S_4 and τ values from Table 5.1 when generating the simulations that we use to choose our hyperparameters. We describe this process along with additional considerations in Appendix B. The parameters we use corresponding to the two scintillation scenarios we assess in this paper are also provided in Table 5.1.

5.2 Sparse Detection

Our objective in this section is to detect cycle slips. Put another way, we would like to find a set of times at which there is a non-negligible probability of a slip having occurred. Moreover, we want to do this while making full use of our measurements $\mathbf{y}$ and taking advantage of our knowledge of the sparsity of cycle slip occurrence. To do so, we first introduce a new variable $\Delta\tilde{\mathbf{z}}$, called the

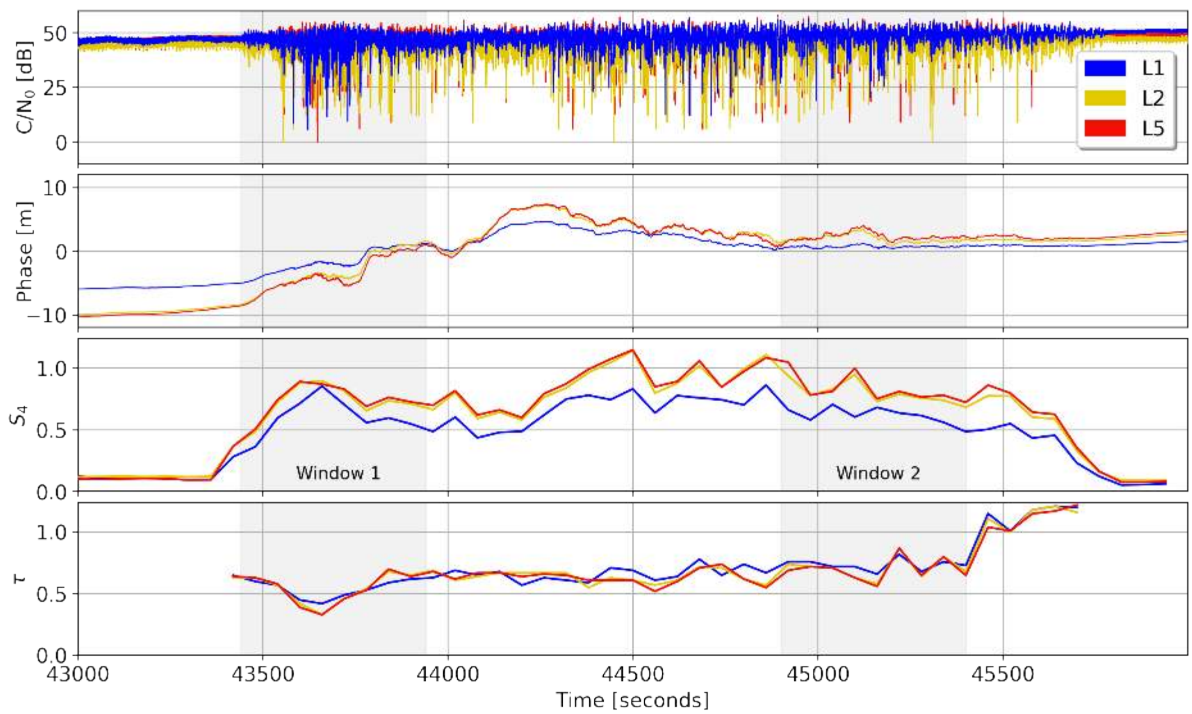

Figure 5.2: Hong Kong scintillation dataset that was introduced in Section 1.4.0.2. The first and second panels show the C/N$_0$ and geodetic-detrended phase for the three signals. The third and fourth panels show the scintillation indices S_4 and decorrelation times τ, respectively, for each of the signals. The shaded areas indicate the two time windows of interest.

sparse float estimate of $\mathbf{\Delta z}$. For this variable, we consider a sparsity-inducing prior that is given by the product of Laplace distributions:

$$p\left(\Delta\tilde{\mathbf{z}}\right) = \prod_{k\in\mathbf{K}}\prod_{t'\in\mathbf{t'}} \frac{\lambda}{2}\exp\left(-\lambda|\Delta\tilde{z}_k(t')|\right) \tag{5.1}$$

The parameter λ (not to be confused with the carrier wavelengths λ_k) is the inverse scale parameter of the Laplace distribution. It plays an important role in the L1-norm optimization that arises when maximizing the posterior corresponding to this prior. We consider how this prior augments the posterior distribution:

$$p\left(\Delta\tilde{\mathbf{z}}|\mathbf{y}\right) \propto p\left(\mathbf{y}|\Delta\tilde{\mathbf{z}}\right)p(\Delta\tilde{\mathbf{z}}) \tag{5.2}$$

From this we see that the log of the posterior can be expressed:

$$\log p(\Delta\tilde{\mathbf{z}}|\mathbf{y}) = -\mathcal{L}(\Delta\tilde{\mathbf{z}}) + C \tag{5.3}$$

where C is constant with respect to $\Delta\tilde{\mathbf{z}}$ and we have the following objective function:

$$\mathcal{L}(\Delta\tilde{\mathbf{z}}) = \frac{1}{2}||\mathbf{y} - \mathbf{B}\Delta\tilde{\mathbf{z}}||^2_{\mathbf{Q_y}} + \lambda||\Delta\tilde{\mathbf{z}}||_1 \tag{5.4}$$

with

$$||\mathbf{y} - \mathbf{B}\Delta\tilde{\mathbf{z}}||^2_{\mathbf{Q_y}} = \left(\mathbf{y} - \mathbf{B}\Delta\tilde{\mathbf{z}}\right)^T\mathbf{Q_y}^{-1}\left(\mathbf{y} - \mathbf{B}\Delta\tilde{\mathbf{z}}\right) \tag{5.5}$$

$$||\Delta\tilde{\mathbf{z}}||_1 = \sum_{k\in\mathbf{K}}\sum_{t'\in\mathbf{t'}}|z_k(t')| \tag{5.6}$$

The maximum a-posteriori (MAP) estimate of the float cycle slip amplitudes is given by:

$$\Delta\tilde{\mathbf{z}} = \underset{\Delta\mathbf{z}\in\mathbb{R}^{KN_{t'}}}{\operatorname{argmin}}\left\{\mathcal{L}(\Delta\mathbf{z})\right\} \tag{5.7}$$

Our detection strategy is to use the MAP estimate of the float slip amplitudes as an indicator of the occurrence of slips.

5.2.1 MM Algorithm

The objective function in Equation 5.4 corresponds to a quadratic optimization problem with L1-norm regularization. This type of problem appears in various contexts, for example in basis pursuit denoising and problems involving LASSO (least absolute shrinkage and selection operator). Our particular formulation is closely related to the problem of total variation denoising with simultaneous low-pass filtering (LPF/TVD), which is introduced in [69]. In that work, the authors use majorizer-minimization (MM) to solve the problem of extracting a sparse signal component from a 1D noisy signal with low-frequency trends. Here, we extend and adapt their MM approach to solve Equation 5.7 for our case of a multidimensional time series containing multiple sparse signal components.

The principle behind MM is to solve a sequence of optimization problems where the objective functions are quadratic majorizers of $\mathcal{L}$. A majorizer of the objective function $\mathcal{L}$ at a reference point $\Delta\tilde{\mathbf{z}}_i$ is a function $\mathcal{M}_i$ such that $\mathcal{M}_i(\Delta\tilde{\mathbf{z}}_i) = \mathcal{L}(\Delta\tilde{\mathbf{z}}_i)$ and $\mathcal{M}_i(\Delta\tilde{\mathbf{z}}) \geq \mathcal{L}(\Delta\tilde{\mathbf{z}})\ \forall\ \Delta\tilde{\mathbf{z}}$. The solution of the optimization problem corresponding to $\mathcal{M}_i$ is used as the reference point for the majorizer $\mathcal{M}_{i+1}$ in the next iteration. The resulting sequence of solutions will converge to the global minimum of $\mathcal{L}$ (so long as $\mathcal{L}$ is convex). This concept is illustrated for a hypothetical 1-dimensional optimization problem in Figure 5.3.

We use the following majorizer of $\mathcal{L}$:

$$\mathcal{M}_i(\Delta\tilde{\mathbf{z}}) = \frac{1}{2}||y - \mathbf{B}\Delta\tilde{\mathbf{z}}||^2_{\mathbf{Q_y}} + \frac{\lambda}{2}\Delta\tilde{\mathbf{z}}^T\Lambda_i^{-1}\Delta\tilde{\mathbf{z}} + \frac{\lambda}{2}||\Delta\tilde{\mathbf{z}}_i||_1 \tag{5.8}$$

where

$$\Lambda_i = \mathrm{diag}(|\Delta\tilde{\mathbf{z}}_i|) \tag{5.9}$$

The minimizer of Equation 5.8 is given by:

$$\Delta\tilde{\mathbf{z}}_{i+1} = \left(\mathbf{B}^T\mathbf{Q}_{\mathbf{y}}^{-1}\mathbf{B} - \lambda\Lambda_i^{-1}\right)^{-1}\mathbf{B}^T\mathbf{Q}_{\mathbf{y}}^{-1}\mathbf{y} \tag{5.10}$$

As iterations progress, many entries in $\Delta\tilde{\mathbf{z}}_i$ approach zero and introduce numerical issues when evaluating expressions containing Λ_i^{-1}. To circumvent the issue, the Woodbury matrix inverse identity can be used to obtain an equivalent expression in terms of Λ_i:

$$\Delta\tilde{\mathbf{z}}_{i+1} = -\frac{1}{\lambda}\left[\Lambda_i - \Lambda_i\mathbf{B}^T\Gamma_i^{-1}\mathbf{B}\Lambda_i\right]\mathbf{B}^T\mathbf{Q}_{\mathbf{y}}^{-1}\mathbf{y} \tag{5.11}$$

$$\Gamma_i = \lambda\mathbf{Q}_{\mathbf{y}} - \mathbf{B}\Lambda_i\mathbf{B}^T$$

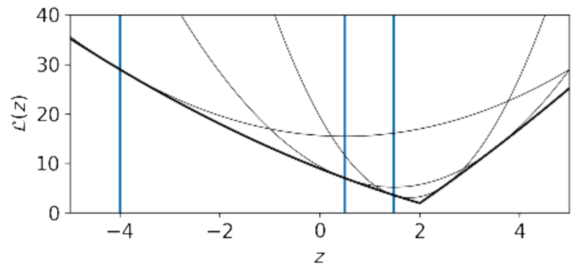

Figure 5.3: Illustrates the MM process for a hypothetical univariate minimization problem. The black curve is the objective function, the gray curves are the series of quadratic majorizers, and the blue lines illustrate how the values that minimize the majorizers progress towards finding the minimum of the objective function.

To perform the algorithm, $\Delta\tilde{\mathbf{z}}_0$ can be initialized as a sequence of all 1s. One then proceeds to minimize the objective function in Equation 5.8 evaluated at $\Delta\tilde{\mathbf{z}}_0$ by applying Equation 5.9 and solving Equation 5.11 to obtain $\Delta\tilde{\mathbf{z}}_1$. This process is repeated until convergence or some set number of iterations. The function MM in Algorithm 1 outlines this process. As a result of this process, we obtain the sparse float amplitude estimate $\Delta\tilde{\mathbf{z}}$, an example of which can be seen in the fourth panel of Figure 5.10 in Section 5.7.

5.3 Detection and Tuning Parameters

If our algorithm is properly tuned, the float estimate $\Delta\tilde{\mathbf{z}}$ at the end of the MM iterations will be close to zero except at epochs where there is a non-negligible probability of cycle slip occurrence. It is important to note that because $\Delta\tilde{\mathbf{z}}$ is a float estimate, it does not necessarily do a good job of indicating which signals contain slips, especially when slip amplitudes are correlated. An example of this can be seen in Figure 5.10 from Section 5.7.1, whose third and fourth panels show the slip bias truth reference and $\Delta\tilde{\mathbf{z}}$, respectively. Any time there is actually a simultaneous slip in one direction on the L2 and L5 signals, $\Delta\tilde{\mathbf{z}}$ shows a spike on the L1 signal in the opposite direction. It makes sense that the optimal argument of a sparsity-inducing cost function would estimate one slip on one signal rather than two slips on the two other signals. Nevertheless, it is clear that any non-zero values in the float estimate correspond to slips in the actual measurements, and so we

Algorithm 1 Majorizor Minimization for computing $\Delta\tilde{\mathbf{z}}$

function $\mathrm{MM}(\mathbf{y}, \mathbf{Q}_\mathbf{y}, \mathbf{B})$
 ▷ For $\mathbf{y} \in \mathbb{R}^{KN_t}$, $\mathbf{Q}_\mathbf{y} \in \mathbb{R}^{KN_t \times KN_t}$, $\mathbf{B} \in \mathbb{R}^{N_t \times N_{t'}}$
 $\Delta\tilde{\mathbf{z}}_0 \leftarrow \mathrm{ones}(N_{t'})$
 for $i \leftarrow 1, \ldots,$ max iterations **do**
  $\Lambda_i \leftarrow \mathrm{diag}(|\Delta\hat{\mathbf{z}}_i|)$
  $\Gamma_i \leftarrow \lambda\mathbf{Q}_\mathbf{y} - \mathbf{B}\Lambda_i\mathbf{B}^T$
  $\Delta\hat{\mathbf{z}}_{i+1} \leftarrow -\frac{1}{\lambda}\left[\Lambda_i - \Lambda_i\mathbf{B}^T\Gamma_i^{-1}\mathbf{B}\Lambda_i\right]\mathbf{B}^T\mathbf{Q}_\mathbf{y}^{-1}\mathbf{y}$
 ▷ estimate is $\Delta\tilde{\mathbf{z}}$ from last iteration
 return $\Delta\tilde{\mathbf{z}}_i$

adopt the approach of detecting cycle slips as relative maxima in the norm of $\Delta\tilde{\mathbf{z}}$ evaluated at each time epoch. We also note that MM is a multiplicative algorithm, and so while most of the values of $\Delta\tilde{\mathbf{z}}$ will be extremely close to zero, they will never equal zero. Thus we also introduce a detection threshold a_{det} to avoid detecting slips from spurious small fluctuations in $\Delta\tilde{\mathbf{z}}$. We denote the times at which we detect slips as $\mathbf{t}_{\mathrm{det}}$, which can be expressed as:

$$\begin{aligned} \mathbf{t}_{\mathrm{det}} = \{t \in \mathbf{t}' \;\mid\; &||\Delta\tilde{\mathbf{z}}(t)||_1 > a_{\mathrm{det}} \text{ and} \\ &||\Delta\tilde{\mathbf{z}}(t)||_1 \text{ is local max.}\} \end{aligned} \tag{5.12}$$

The performance of this detection scheme is dependent upon the choices of both λ and a_{det}. In the MM procedure, increasing the value of λ has the effect of promoting sparsity in the result, i.e. making more entries closer to zero. In theory, the actual value of λ should be related to the rate of cycle slip occurrence among all signals being considered. In the case of ionosphere scintillation, since we have realistic simulations at our disposal, we choose both λ and a_{det} to optimize detection performance on simulated datasets. We performed 50 simulations of 5-minute duration and estimated $\Delta\tilde{\mathbf{z}}$ using values of λ ranging from 15 to 100. We then detected slip occurrences according to the procedures above using different values of a_{det}. We consider a detection to be correct if it falls within 0.2 seconds of an actual slip occurrence in the truth reference. This allows us to calculate the number of missed detections and false alarms for the various values of λ and a_{det}. Figure 5.4 shows the results, with the missed detection and false alarm rates (averaged over all simulations) plotted against the detection threshold. The various lightly-colored lines show

how these curves change for different values of λ, and the colored dots represent the corresponding crossover point of the false alarm and missed detection curves. The dark lines highlight the curves corresponding to the value of λ that minimizes the crossover point. We use this as our criteria for choosing the value of λ, since it optimizes the trade-off between detecting more slips and accruing more false alarms.

With λ established, choosing the value of a_{det} mainly becomes a trade-off between ensuring that we detect all cycle slips and reducing the computational burden of approximating the posterior; the more slips we detect, the more parameters we have to estimate in our reduced model. We found that a threshold of 0.01 works well for all the scenarios we investigated, as it enables detection of all consequential cycle slips while keeping the total algorithm runtime at a reasonable level. One interpretation of this whole process is that the values of λ and a_{det} are hyperparameters for the sparse prior on $\Delta\tilde{\mathbf{z}}$ that is realized through this entire process of detection. Our validation of these parameters is based on simulations, but it is conceivable that their values can be linked to characteristics of the signals themselves. For instance, in Chapter 2 we determined how the rate of cycle slip occurrence relates to signal conditions for the case of diffractive ionosphere scintillation. It is conceivable that other relationships can be developed to quantify the rates of slip occurrence in other scenarios (ocean reflection, troposphere scintillation, etc.), and that these slip rates can be linked to the choice of λ. This will be an interesting topic for future investigation.

5.4 Reduced Model

Having established a method for detecting cycle slip occurrences, we can now form a new system model that only admits slips at the times detected in the previous step. In particular, we construct the matrix $\mathbf{S}$ from Equation 4.5 using $\mathbf{t}' = \mathbf{t}_{\mathrm{det}}$ and create a reduced model matrix $\mathbf{B}_{\mathrm{red}}$ relating slips at these times to the measurements $\mathbf{y}$. We then evaluate the mean and covariance of the float posterior distribution from Equation 4.31. In the new model, $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}$ is of dimensions $KN_d \times KN_d$ where N_d is the number of detected slips. Because of the sparse detection procedure, $N_d \ll N_{t'}$ and so the problem dimension has been significantly reduced. The first panel of Figure

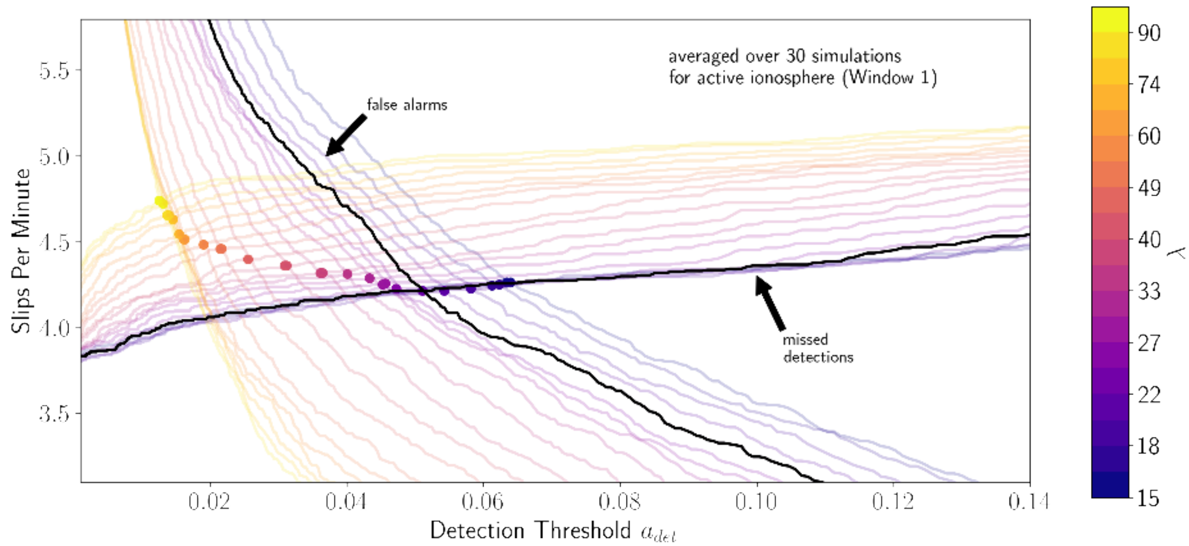

Figure 5.4: Shows the rate (in slips per minute) of false alarms and missed detections for a range of values for the detection threshold a_{det} and tuning parameter λ. The colored points indicate crossover points of the false alarm and missed detection curves for the various values of λ. The black lines indicate curves for the value of λ that minimizes the crossover point.

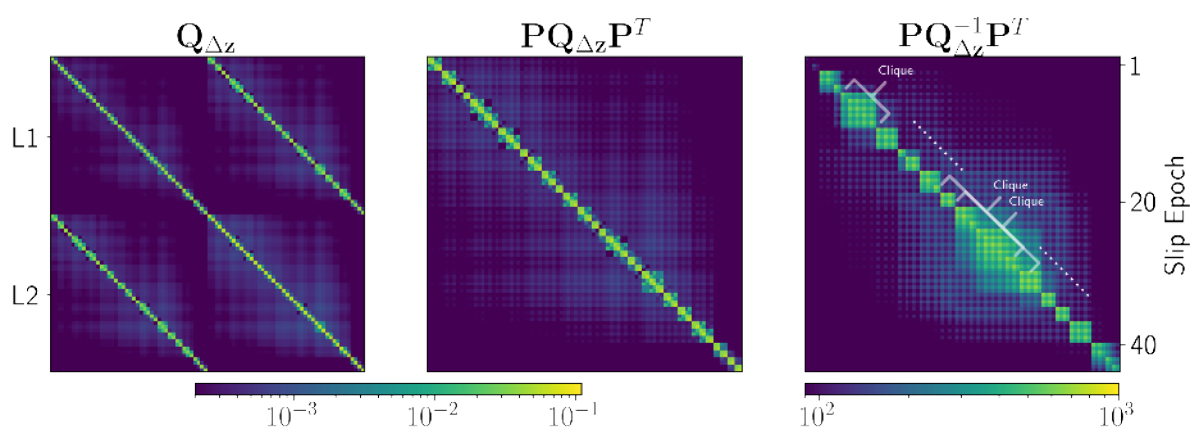

Figure 5.5: Examples of the covariance $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}$ of the reduced model float estimate corresponding to the real data from Window 2. The second and third panels show the covariance and precision (i.e. inverse of the covariance) matrices for the permuted reduced model, where parameters are ordered by slip epoch. Different cliques over the model parameters are indicated in the third panel.

5.5 shows an example of $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}$ in the reduced model for dual-frequency signals. There are four nearly diagonal blocks corresponding to the auto and cross-covariances of the slip amplitudes on the two signals. To make this covariance more diagonal, we apply a permutation $\mathbf{P}$ so that the covariance terms corresponding to a particular epoch are adjacent, i.e. it converts it from a $K \times K$ matrix of $N_d \times N_d$ blocks to a $N_d \times N_d$ matrix of $K \times K$ blocks. Panel b shows the permuted covariance matrix, which is now close to block-diagonal aside from a set of 2×2 blocks just off the main diagonal.

5.5 Search for ILS Solution

With our reduced model established, we are now ready to consider actually finding the best amplitude sequence estimate. As we mentioned, this type of problem involves a discrete search over potential candidates for the maximizer of Equation 4.28. We will assume the prior on cycle slip amplitude sequences for the reduced model is non-informative so that our MAP estimate corresponds to the ILS solution. Recall that, for a float estimate $\hat{\mathbf{a}}$ and covariance $\mathbf{Q}$, the ILS problem can be expressed as:

$$\check{\mathbf{a}} = \text{argmin}_{\mathbf{a} \in \mathbb{Z}^N} \ (\mathbf{a} - \hat{\mathbf{a}})^T \mathbf{Q}^{-1} (\mathbf{a} - \hat{\mathbf{a}}) \tag{5.13}$$

In our case, $\mathbf{a} = \Delta\mathbf{z}$ and $\mathbf{Q} = \mathbf{Q}_{\Delta\mathbf{z}}$, although we will use $\mathbf{a}$ and $\mathbf{Q}$ in this section for notational simplicity and for consistency with other literature.

5.5.1 LAMBDA

To find the solution to Equation 5.13, we search for $\check{\mathbf{a}}$ over a hyper-ellipsoidal region defined by:

$$r(\mathbf{a}) = (\mathbf{a} - \hat{\mathbf{a}})^T \mathbf{Q}^{-1} (\mathbf{a} - \hat{\mathbf{a}}) \leq R_{\text{search}} \tag{5.14}$$

where $r(\mathbf{a})$ is the objective function we are trying to minimize and R_{search} denotes the search radius. We can then compare each $\mathbf{a}$ in this region to find $\check{\mathbf{a}}$. The LAMBDA algorithm prescribes an approach that makes this search as efficient as possible. Normally, it is broken down into two steps: reduction and search. The first step is actually optional, but serves to make the search process more efficient by finding a new basis for the integer lattice which approximately orthogonalizes $\mathbf{Q}^{-1}$. We skip this step for reasons explained in Section 5.5.2. The main way that LAMBDA accomplishes the second step is through the *search-and-shrink* algorithm, which is a method of enumerating points close to $\hat{\mathbf{a}}$ while simultaneously shrinking R_{search}, thereby reducing the number of points we must search over. The algorithm has been carefully explained in [79], [19], and various other publications. Here we go through the essential details, closely following the development provided in the official LAMBDA documentation [25].

5.5.1.1 Sequential Conditional Fixing

To explain the search process, it is helpful to first describe a solution for $\mathbf{a}$ that is obtained by sequentially fixing the individual components of $\mathbf{a}$. This is also called the "bootstrapped" solution. We begin with the $\mathbf{LDL}^T$ decomposition of $\mathbf{Q}$, where $\mathbf{D}$ is diagonal and $\mathbf{L}$ is lower triangular with ones along its diagonal. Inserting this for $\mathbf{Q}$ in Equation 5.14 yields:

$$(\mathbf{a} - \hat{\mathbf{a}})^T \mathbf{L}^{-T}\mathbf{D}^{-1}\mathbf{L}^{-1} (\mathbf{a} - \hat{\mathbf{a}}) \leq R_{\text{search}} \tag{5.15}$$

Next, define $\bar{\mathbf{a}} = \mathbf{a} - \mathbf{L}^{-1} (\mathbf{a} - \hat{\mathbf{a}})$ so that

$$\mathbf{L} (\bar{\mathbf{a}} - \mathbf{a}) = \mathbf{a} - \hat{\mathbf{a}} \tag{5.16}$$

The entries of $\bar{\mathbf{a}}$ are the conditional float estimates of each component of $\mathbf{a}$ given its predecessors, i.e.

$$\bar{a}_i = \hat{a}_{i:1:i-1} = \hat{a}_i - \sum_{j=1}^{i-1}(a_j - \bar{a}_j)L_{i,j}, \ 1 \leq i \leq N \tag{5.17}$$

Here, $\hat{a}_{i:1:i-1}$ is the i-th component of $\hat{\mathbf{a}}$ conditioned on the fixed values of the preceding components, $a_1, \ldots, a_{i-1}$. Note that Equation 5.17 can be obtained using forward substitution applied to the system in Equation 5.16.

Inserting Equation 5.16 into Equation 5.15 yields:

$$(\mathbf{a} - \bar{\mathbf{a}})^T \mathbf{D}^{-1} (\mathbf{a} - \bar{\mathbf{a}}) \leq R_{\text{search}} \tag{5.18}$$

Since $\mathbf{D}^{-1}$ is diagonal, this can be written:

$$\frac{(a_1 - \bar{a}_1)^2}{D_1} + \cdots + \frac{(a_N - \bar{a}_N)^2}{D_N} \leq R_{\text{search}} \tag{5.19}$$

We see that in order to minimize the LHS of Equation 5.19, $\mathbf{a}$ should in some sense be close to $\bar{\mathbf{a}}$. In particular, if we fix the first $i-1$ coordinates of $\mathbf{a}$, then we observe the following bound on a_i:

$$(a_i - \bar{a}_i)^2 \leq D_i \left(R_{\text{search}} - \sum_{j=1}^{i-1} \frac{(a_j - \bar{a}_j)^2}{D_j} \right) \tag{5.20}$$

We can obtain a somewhat decent solution by sequentially fixing a_1 through a_N to the values closest to $\bar{\mathbf{a}}$. That is, we first take $a_1 = \text{round}(\bar{a}_1) = \text{round}(\hat{a}_1)$. Then we compute $\bar{a}_2$ using Equation 5.17 and fix $a_2 = \text{round}(\bar{a}_2)$. This process continues for $a_i,\ 3 \leq i \leq N$. The resulting fixed value of $\mathbf{a}$ is the "bootstrapped" solution (which we denote $\mathbf{a}^{\mathbf{b}}$) that we mentioned earlier. Algorithm 2 shows how to obtain the bootstrapped solution given $\hat{\mathbf{a}}$ and the $\mathbf{L}$ and $\mathbf{D}$ matrices from the $\mathbf{LDL}^T$ decomposition of $\mathbf{Q}$.

Algorithm 2 Integer Bootstrapping

function BOOTSTRAP($\mathbf{L}, \mathbf{D}, \hat{a}$)
    $N \leftarrow \text{length}(\hat{\mathbf{a}})$
    $\bar{\mathbf{a}} \leftarrow \text{zeros}(N)$
    $\mathbf{a}^b \leftarrow \text{zeros}(N)$
    $a_1 \leftarrow \text{round}(\hat{a}_1)$
    for $i \leftarrow 2, \ldots N$ **do**
        $a_i \leftarrow \hat{a}_i - \sum_{j=1}^{i-1} (a_j - \bar{a}_j) L_{i,j}$
    return $\mathbf{a}^b$

5.5.1.2 Search Routine

Now that we have described sequential conditional fixing, we are ready to introduce the search-and-shrink routine for finding the solutions of Equation 5.13. Note the use of the word *solutions*, since technically there can be up to 2^N minimizers of Equation 5.13. However, as noted in [78], this is essentially never the case in practice. Even still, the search method in LAMBDA prescribes a way in which we can obtain the N_{cand} candidate solutions of $\mathbf{a}$ that are closest to $\hat{\mathbf{a}}$. From these candidate solutions, we can evaluate and select our ILS solution.

At the beginning, we consider $R_{\text{search}} = \infty$ and choose N_{cand} initial candidates. For our first candidate we use the bootstrapped solution $\mathbf{a}^{(1)} = \mathbf{a}^b$. We then obtain $N_{\text{cand}} - 1$ other candidates by keeping all their entries the same as $\mathbf{a}^b$ except for the last component, which alternately switches to the next closest integer to $\bar{a}_N^b$, i.e.

$$\begin{aligned}
a_N^{(1)} &= a_N^b \\
a_N^{(2)} &= a_N^b + 1 \\
a_N^{(3)} &= a_N^b - 1 \\
a_N^{(4)} &= a_N^b + 2 \\
&\vdots \\
a_N^{(2p)} &= a_N^b + p \\
a_N^{(2p+1)} &= a_N^b - p \\
&\vdots
\end{aligned}$$

At this point, we have found N_{cand} solutions that are inside our search radius. We shrink the search ellipsoid by setting $R_{\text{search}} = \min\left\{r(\mathbf{a}^{(i)}) \ : \ 1 \le i \le N_{\text{cand}}\right\}$, which for this first case becomes $R_{\text{search}} = r(\mathbf{a}^{(N_{\text{cand}})})$. Next, the search continues starting from coordinate $N-1$, where we take our new $\mathbf{a}$ to be our bootstrapped solution except now we change component $N-1$ to be the next closest integer to $\bar{a}_{N-1}$ and set $a_N = \text{round}(\bar{a}_N)$. At this point, one of two things happens:

(1) If this new candidate $\mathbf{a}^{\text{new}}$ satisfies $r(\mathbf{a}^{\text{new}}) \le R_{\text{search}}$, then we update our list of candidates by replacing the one that has the largest r value with this new candidate. We also set $R_{\text{search}} = \min\left\{r(\mathbf{a}^{(i)}) \ : \ 1 \le i \le N_{\text{cand}}\right\}$. The search then continues by modifying $\mathbf{a}^{\text{new}}$ at component N.

(2) Otherwise, $\mathbf{a}^{\text{new}}$ does not lie within the search ellipsoid, so we continue the search starting

from component $N-2$.

This process continues until all the components of $\mathbf{a}$ have been considered and no new candidates lie within the search ellipsoid.

To illustrate this process, Figure 5.6 shows the candidates and search ellipsoids for the case $N = 2$ dimensions and $N_{\text{cand}} = 6$. It starts with the bootstrapped candidate at (1, 0), which is labeled 1, and then considers the next 5 closest integer coordinates for a_2. Together, these form our initial set of 6 candidates. The dashed blue line shows the R_{search} ellipse corresponding to this initial set of candidates. From here, we take the bootstrapped solution (1, 0) and adjust a_1 to be the next closest integer, which in this case is 0. This point, (0, 0), is labeled 7 in the figure. Since it lies within the search ellipse, we throw out candidate 6 and shrink R_{search} to correspond to candidate 5, which now is the candidate with the largest distance to $\hat{a}$. Since candidate 7 was a success, move the search back up to the 2nd component, with its next best integer value being 1. This new point, (0, 1), is labeled 8 in the figure. Since it also lies within the ellipsoid, we throw out candidate 5, replace it with this new candidate, and shrink R_{search} to correspond to the distance of candidate 4. The same thing happens with the next search candidate, which is labeled 9, and the search ellipse is shrunk further. After adding point 9 to our candidate list, the next point we try is 10, but this point is outside of the search ellipsoid, so it is discarded and the search moves back down to the first component again. This time when we try a new coordinate for a_1, the subsequent conditionally fixed solution, which is labeled 11, lies outside of the search ellipsoid, and so no new candidates can be found starting from the first coordinate. Since there are no new candidates and we have started the search from each component of $\mathbf{a}$, the process is finished and the candidates 1, 2, 3, 7, 8, and 9 form the 6 closest values of $\mathbf{a}$ to $\hat{\mathbf{a}}$.

5.5.2 Cliques

The search-and-shrink procedure is an exhaustive search for the ILS solution that can be fairly time-consuming. For this reason, LAMBDA applies a decorrelating transformation to $\mathbf{a}$ and $\mathbf{Q}$ before performing its search. Various lattice reduction methods have been successful in

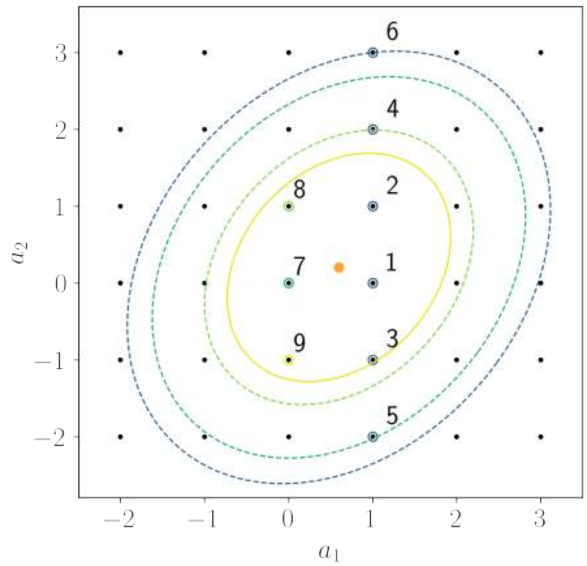

Figure 5.6: Illustrates the process of search and shrink, which seeks the best integer candidates within a given ellipsoid of particular float solution.

handling problems with correlated ambiguities of multi-frequency signals for a fairly high number of dimensions [40]. For an efficient search implementation, dimensions up to 50 can be searched in less than a minute. For problems with small variance, such as ours, even without lattice reduction we find that a search over up to 40 dimensions can be performed in less than a minute (on one 2.4 GHz processor). However any additional increase in problem dimension seems to quickly lead to very long search times. For multi-frequency measurements and long windows of data containing many cycle slips, our problem dimension can easily reach over 150 dimensions, making the standard search-and-shrink routine impractical, even when using a decorrelating transform.

The high dimensionality of our problem would not be an issue if partitioned into completely independent sets of variables, since then we could just split the problem into smaller problems. Unfortunately, the consecutive occurrences of cycle slips that arise during harsh signal conditions leads to widespread interdependence of slip amplitudes across the entire estimation window. This interdependence manifests as the off-diagonal blocks in $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}$, which effectively couple the amplitude estimates between adjacent slip epochs. The widespread interdependence that arises from this coupling can be seen from the large off-diagonal entries of the float solution precision matrix $\mathbf{Q}_{\Delta\mathbf{z}}^{-1}$, an example of which is shown in the third panel of Figure 5.5.

Considering the problem from the perspective of probabilistic graphical models (PGM) offers helpful insight on this problem. Such models represents variables (e.g. cycle slip amplitudes) as nodes and their inter-dependencies (e.g. non-zero values in the float precision matrix) as edges. [33] considers the ILS problem from the perspective of PGMs and suggests that such problems typically have dependency graphs that are complete, i.e. all the variables being estimated are interdependent. This is also somewhat true for our case of inferring the amplitudes of consecutive cycle slips in a time series. Figure 5.7 illustrates the PGM corresponding to the cycle slip problem and how marginalizing the Gaussian process priors ultimately results in an approximately complete dependency graph. We say *approximately* because the interdependence between cycle slip amplitude estimates does diminish over time, as we can see from the tapering of entries in the precision matrix in Figure 5.5. In the context of PGMs, we describe this behavior using the concept of *cliques*. A clique is a maximal set of interdependent nodes; that is, it is a subset of our slip amplitude estimates each of which is interdependent with every other variable in the subset. Figure 5.8 illustrates a hypothetical PGM corresponding to the cycle slip problem where the assumed maximal clique size is 3. In reality, the size of cliques in an ILS model changes depending on the variables and their interdependence. This concept is illustrated in the third panel of Figure 5.5, which shows brackets along the main diagonal of the precision matrix that roughly indicate a few of the subsets of interdependent variables. The concept of cliques is important in our two approaches to finding the ILS solutions to our reduced model, which we discuss next.

5.5.3 Method 1: Search and Shrink Over Cliques

Our first approach to the ILS search for high-dimensional problems is essentially to apply search-and-shrink to subsequent cliques of the ILS model. We call this method search-and-shrink over cliques (SASOC) and the basic steps can be summarized as:

(1) Define the set of cliques over the model.

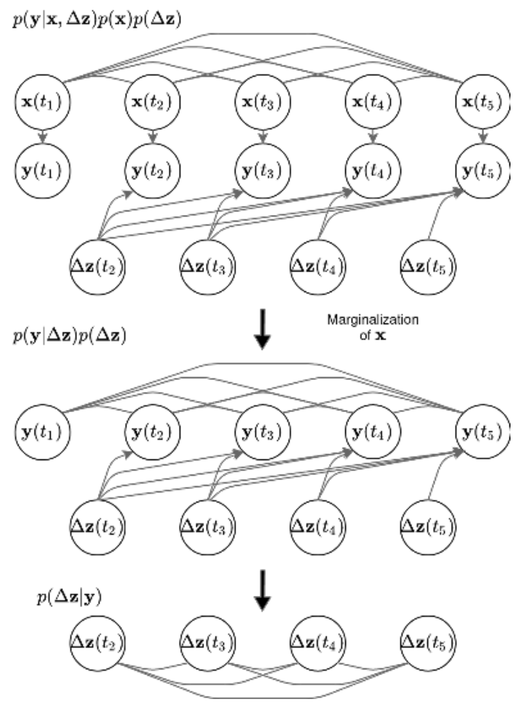

Figure 5.7: Illustrates the probabilistic graph structure of the cycle slip problem before and after marginalization of the phase components $\mathbf{x}$. Nodes indicate sets of variables in the model and edges imply conditional dependence between nodes. In this case, we show a complete graph for variables $\mathbf{z}$, indicating that all variables are interdependent.

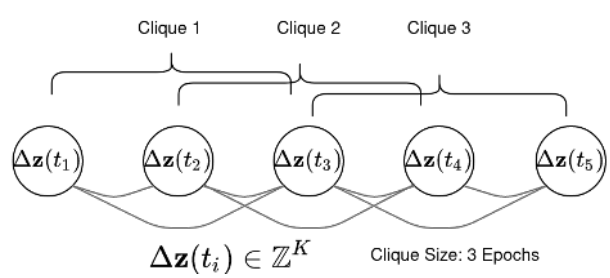

Figure 5.8: Illustrates the concept of cliques in a PGM.

(2) Perform search-and-shrink for N_{cand} candidates over the first clique.

(3) Take the best candidate over the clique $\mathbf{a}^{\text{best}}$ and fix each variables at the beginning of the clique that is not in the subsequent clique: $\check{a}_i = a_i^{\text{best}}$ for each appropriate i.

(4) With those values fixed, continue with search-and-shrink over the subsequent clique.

(5) Repeat this process until we have searched the last clique, at which point we fix all remaining

variables according to the best candidate.

We consider the fixed solution of this process to be our ILS estimate. When defining our set of cliques, we can use the values of the precision matrix to choose maximal sets of interdependent variables. The values of the precision matrix do not end up being exactly zero, and some threshold on their magnitude would have to be used. Alternatively, for the way we actually implement the algorithm, we can set some maximal clique size N_{clique} and assume that no clique for our model exceeds this dimension. Larger N_{clique} captures more of the subtle variable interdependence (as indicated by the values further off the main diagonal of the precision matrix) but also results in a more computationally intensive search. In our implementation, we choose $N_{\text{clique}} = 20$ for a good trade-off between search fidelity and computational speed. Our implementation of SASOC is outlined in Algorithm 3.

5.5.4 Method 2: Approximate Support Using Marginals

While SASOC adapts the search-and-shrink algorithm to be feasible over our high-dimensional problem, it is still somewhat inefficient since it performs a fresh search over an entire clique for almost every variable in the model. As such, we propose another algorithm, which is to bound the support of the joint posterior on cycle slip amplitude sequences using approximations of its marginal distributions at each epoch. We call this algorithm Approximate Support Using Marginals (ASUM). It is inspired by the work from [87]. For a continuous joint-normal random vector $\mathbf{v}$ with mean μ and covariance $\mathbf{Q}$, if we partition $\mathbf{v}$ into $\mathbf{v}_1$ and $\mathbf{v}_2$, then its joint distribution is bounded by the product of its marginals:

$$p(\mathbf{v}) = \mathcal{N}(\mathbf{v}; \mu, \mathbf{Q}) \leq p(\mathbf{v}_1)p(\mathbf{v}_2) = \mathcal{N}(\mathbf{v}_1; \mu_1, \mathbf{Q}_1)\,\mathcal{N}(\mathbf{v}_2; \mu_2, \mathbf{Q}_2) \tag{5.21}$$

From this we can say that the support of $p(\mathbf{v})$ must coincide with the supports of both $p(\mathbf{v}_1)$ and $p(\mathbf{v}_2)$. We can apply this concept when evaluating the posterior for $\Delta\mathbf{z}$, starting first with the marginal distribution of slip amplitudes at individual epochs and then combining successive groups of epochs. Since the variances are small, the marginal support is constrained and remains

Algorithm 3 Find ILS solution using search-and-shrink over cliques

```
function SASOC(L, D, â, N_clique, N_cand)
    N ← length(ẑ)
    R_search ← ∞
    candidates ← zeros(N_cand, N)
    candidates_dist2 ← zeros(N_cand)
    count ← 0
    dist2 ← zeros(N)
    ã ← zeros(N)
    ǎ ← zeros(N)
    step ← zeros(N)
    ã[0] ← â[0]
    ǎ[0] ← round(ã[0])
    i_base ← 0
    done ← False
    while not done do
        if dist2[i] < R_search then
            if i < i_base + N_clique − 1 then
                ▷ if we we have not yet reached the end, keep recursively computing ã
                i ← i + 1
                ã[i] ← â[i] − L[i, 0 : i](ã[0 : i] − ǎ[0 : i])
                ǎ[i] ← round(ã[i])
                dist2[i] ← dist2[i − 1] + (ã[i] − ǎ[i])^2/D[i]
                if ã[i] − ǎ[i] > 0 then
                    step[i] ← 1
                else
                    step[i] ← −1
            else
                ▷ if we reach the end, then store the found candidate and try next valid integer
                candidates[i_max, 0 : N] ← ǎ
                candidates_dist2[i_max] ← dist2[i]
                i_max ← argmax(candidates_dist2)
                R_search ← candidates_dist2[i_max]
                ǎ[i] ← ǎ[i] + step[i]                    ▷ go to next valid integer
                dist2[i] ← dist2[i − 1] + (ã[i] − ǎ[i])^2/D[i]
                step[i] ← −step[i] − sign(step[i])
```

```
    else
        ▷ exit or move down
        if i ← i_base then
            if i_base + N_clique >= N then
                done ← True
            else
                ▷ fix first clique variable using best candidate so far
                ǎ ← candidates[argmin(candidates_dist2), 0 : N]
                ã[i] ← â[i] − L[i, 0 : i](ã[0 : i] − ǎ[0 : i])
                dist2[i] ← (ã[i] − ǎ[i])^2/D[i] + dist2[i − 1]
                ▷ move i_base up and reset candidate search for next clique
                i_base ← i_base + 1
                candidates_dist2[0 : N] ← ∞
                R_search ← ∞
                i_max ← argmax(candidates_dist2)
        else
            ▷ move down
            i ← i − 1
            ǎ[i] ← ǎ[i] + step[i]
            step ← −step[i] − sign(step[i])
            dist2[i] ← (ã[i] − ǎ[i])^2/D[i] + dist2[i − 1]
return candidates, candidates_dist2
```

relatively small even as it exponentiates when we combine multiple slip epochs. We still have to be careful though; the discrete Gaussian is a tricky distribution to work with. While the full discrete distribution is exactly proportional to the full continuous distribution evaluated at the lattice points, its marginals are in general *not* proportional to lattice-samplings of the continuous marginals. Nevertheless, we have found empirically that the sampled continuous marginals still provide a *good enough* approximation of the *support* of the actual discrete marginals. That is, if the continuous marginal evaluated at a point is negligible, the discrete marginal will also be negligible. If the value of the continuous marginal at a point is non-negligible, the discrete marginal at that point *may be* non-negligible.

We must also be aware of how the size of the approximate marginal support grows as we introduce more slip epochs. As we already mentioned, because the variances in the float posterior are small we find the size of the marginal support at each epoch is small enough (around 1-4 points) that the problem remains tractable over several slip epochs. For example, if we expect at most 4 non-negligible slip amplitudes per epoch and our computational resources are limited to evaluating a distribution for around 1000 lattice points at a time, then we can combine around 5 epochs (since $4^5 \approx 1000$). Eventually, depending on the structure of $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}$ and the available computational resources, as we continue to combine marginals evaluating the distribution will become infeasible. At this point, we want different chains of approximate supports to be as independent as possible. For this reason, we determine how to combine marginals by recursively partitioning our variables according to their conditional independence.

To implement the ASUM algorithm, we first recursively partition the float posterior along slip epochs to create a tree of maximally independent subsets of variables, where conditional independence is assessed through the maximum absolute value of entries of the superdiagonal blocks in the precision matrix $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}^{-1}$. This process is outlined in the "partition" function from Algorithm 4. Once we partition the data down to individual slip epochs, we sample the continuous marginal density for each slip epoch as an approximation to the actual discrete marginal. Let χ_0 be the grid of admissible points for a single slip epoch; e.g. if we again consider slip amplitudes between

± 4 cycles then $\chi_0 = \{-4, -3, -2, -1, 0, 1, 2, 3, 4\}^K$. We approximate the integer-valued marginal by evaluating each point in χ_0 according to Equation 4.28 then normalizing by the sum of the results. We select the N_{keep} most probable points from χ_0 according to their continuous marginal probability. We keep track of those points as the approximate support of the marginal. For the next iteration, we consider pairs of epochs with approximate supports χ_1 and χ_2 and we evaluate Equation 4.28 for each point in $\chi_1 \times \chi_2$. This process continues in the next iterations, eventually expanding the marginal to include all slip epochs. At the end, the full distribution is evaluated over the candidates from the full approximate support and the most probable candidate is chosen as our solution. Algorithm 4 shows how this process can be implemented through a recursive function.

5.6 Method Summary and Computational Considerations

In summary, our approach to cycle slip estimation consists of four main steps: 1) Compute the sparse float estimate, 2) detect slips, 3) compute reduced system model, 4) search for the best integer candidates in the reduced ILS model. Under our stated assumptions, this estimate is equivalent to the MAP estimate, and we use it to correct for cycle slip occurrences. These steps are further outlined in Algorithm 5. Also, Figure 5.9 illustrates our entire process for addressing cycle slip occurrence in the real scintillation data that we present in Section 5.7.2. This includes determination of appropriate covariance hyperparameters along with the detection of slip occurrences and bounding/estimation of slip amplitudes. In addition, source code examples for implementing and applying this cycle slip mitigation approach are provided at the SeNSe Lab GitHub page: (https://github.com/cu-sense-lab/cycle-slip-estimation).

Before we dive into the results of applying this algorithm, there are a few remaining computational aspects of the approach that we should discuss. When searching for the ILS slip amplitudes estimate, we proposed two methods. We discussed how the first method, SASOC, is an extension of the traditional ILS search-and-shrink algorithm. In that sense, SASOC may be a more approachable to those familiar with the search-and-shrink approach. The runtime and fidelity of the SASOC estimate depend on the maximum clique size N_{clique}, while for ASUM the performance depends on

Algorithm 4 Recursively approximate support of $\mathcal{ND}(\mu, \mathbf{Q})$ using continuous marginals

function PARTITION(μ, $\mathbf{Q}$)
    ▷ Find separating variable corresponding to the superdiagonal block with the smallest value
    ▷ This can also be done using the precision matrix $\mathbf{Q}^{-1}$ instead of $\mathbf{Q}$
    $N \leftarrow \text{length}(\mu)$
    $i_{\text{min}} \leftarrow 0$
    min_value $\leftarrow \infty$
    for $i \leftarrow 1, \ldots, N$ **do**
        superdiagonal_max $\leftarrow$ max(abs($\mathbf{Q}[0:i, i:N]$))
        if superdiagonal_max $<$ min_value **then**
            min_value $\leftarrow$ superdiagonal_max
            $i_{\text{min}} \leftarrow i$
    return $\mu_1, \mathbf{Q}_1, \mu_2, \mathbf{Q}_2$
function ASUM($\mu, \mathbf{Q}, \chi_0$)
    if size($\mathbf{Q}$) $\leftarrow K$ **then**
        $m(\chi_0) \leftarrow \text{normalize}\left(\mathcal{N}\left(\chi_0; \mu, \mathbf{Q}\right)\right)$
        $\chi_{\text{blk}} \leftarrow \{\Delta\mathbf{z}_i i \in \chi_0 \mid m(\Delta\mathbf{z}_i) > \text{threshold}\}$
    else
        $\mu_1, \mathbf{Q}_1, \mu_2, \mathbf{Q}_2 \leftarrow$ PARTITION(μ, $\mathbf{Q}$)
        $\chi_1 \leftarrow$ ASUM(μ_1, $\mathbf{Q}_1$)
        $\chi_2 \leftarrow$ ASUM(μ_2, $\mathbf{Q}_2$)
        $\chi_3 \leftarrow \chi_1 \times \chi_2$
        $m(\chi_3) \leftarrow \text{normalize}\left(\mathcal{N}\left(\chi_3; \mu, \mathbf{Q}\right)\right)$
        $\chi_{\text{blk}} \leftarrow \{\Delta\mathbf{z}^i \in \chi_3 \mid m(\Delta\mathbf{z}^i) > N_{\text{keep}}^{\text{th}} \text{ largest value of } m(\chi_3)\}$
    return χ_{blk}

the maximum number of points N_{keep} in the approximate marginal supports. For sufficiently large N_{clique} and N_{keep} we found that the SASOC and ASUM solutions were almost always identical. We did find discrepancies when substantially increasing the estimation window size to 40 minutes (as opposed to our the approximately 8-minute windows whose results we actually show) for the Hong Kong scintillation dataset presented in Section 5.7.2. All the results we show in the next section are for 500-second windows of data or less, and the SASOC and ASUM algorithms agree in their found solutions when using $N_{\text{clique}} = 18$ or $N_{\text{keep}} = 100$. Unfortunately, we are not able to provide a rigorous proof regarding the estimate fidelity (i.e. does the found solution actually equal the ILS estimate), but the agreement between the two algorithms suggests that both methods generally work. More substantial proof of their fidelity comes from the results we present in the next section. It is also worth noting that we did find that the ASUM algorithm outperformed SASOC in terms of runtime. As an example, for the 8.5 minutes of 20 Hz triple-frequency scintillation data presented in Section 5.7.2 we found that the ILS search takes between 40-45 seconds for SASOC but only 5-10 seconds for ASUM to run on one 2.4 GHz core.

A very important aspect of the approach, which we only briefly mentioned a few times in this chapter and Chapter 4, is how we implement the mathematical operations in a computationally efficient manner. First, we note that the only densely-evaluated matrix is $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}$. The operators $\mathbf{A}$, $\mathbf{B}$, $\mathbf{B}_{\text{red}}$, and $\mathbf{Q}_{\epsilon}$ can all be efficiently applied in $O(n)$ time and only require $O(n)$ storage or less. The operator $\mathbf{Q}_{\mathbf{x}}$ is block-Toeplitz and can be stored in $O(n)$ space and applied in $O(n\log(n))$ time using FFT-based methods. Therefore, the operator $\mathbf{Q}_{\mathbf{y}}$ can be applied in $O(n\log(n))$ by combining $\mathbf{A}$, $\mathbf{Q}_{\mathbf{x}}$, and $\mathbf{Q}_{\epsilon}$ operators. This is similarly true for the Γ_i operator from Equation 5.11. Inverses of $\mathbf{Q}_{\mathbf{y}}$ and Γ_i can be applied using conjugate gradient descent, and the efficiency of this operation will depend on the number of iterations needed for convergence. These gradient descent operations and computation of $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}$ are the computational bottlenecks for our approach. In general, runtime will depend on the window size, measurement sampling rate, and the number of detected slips. As an example, for the 8.5 minutes of 20 Hz triple-frequency scintillation data presented in Section 5.7.2 we found that the reduced model covariance calculation takes between 3-4 minutes to run on one 2.4 GHz core. Again, we emphasize that this algorithm is designed and intended for batch post-processing. However, this runtime could be substantially improved with more focus on efficient implementation, which will be an important topic of future work.

Algorithm 5 Demonstrates full cycle slip mitigation algorithm

$\triangleright$ Step 1: compute sparse float estimate
$\Delta\tilde{\mathbf{z}} \leftarrow \mathrm{MM}(\mathbf{y}, \mathbf{Q_y}, \mathbf{B})$

$\triangleright$ Step 2: detect slips
$\mathbf{t}_{\mathrm{det}} \leftarrow \{t \in \mathbf{t}' : ||\Delta\tilde{\mathbf{z}}(t)||_1 > a_{\mathrm{det}}\}$
$\triangleright$ Step 3: compute reduced model
$\mu^P_{\Delta\hat{\mathbf{z}}} \leftarrow \mathbf{P}\mu_{\Delta\hat{\mathbf{z}}}$
$\mathbf{Q}^P_{\Delta\hat{\mathbf{z}}} \leftarrow \mathbf{P}\mathbf{Q}_{\Delta\hat{\mathbf{z}}}\mathbf{P}^T$

$\triangleright$ Step 4-a: find best slip amplitude sequence over the approximated support
$\chi_0 \leftarrow \{-N_p, \ldots, N_p\}^K$ $\triangleright$ N_p is the maximum slip amplitude under consideration
$\chi_{\mathrm{approx}} \leftarrow \mathrm{ASUM}(\mu^P_{\Delta\hat{\mathbf{z}}},\ \mathbf{Q}^P_{\Delta\hat{\mathbf{z}}},\ \chi_0)$
$p(\chi_{\mathrm{approx}}) \leftarrow \mathrm{normalize}\left(\mathcal{N}\left(\chi_{\mathrm{approx}}; \mu^P_{\Delta\hat{\mathbf{z}}}, \mathbf{Q}^P_{\Delta\hat{\mathbf{z}}}\right)\right)$
$\Delta\mathbf{z}_{\mathrm{MAP}} \leftarrow \mathrm{argmax}(p(\chi_{\mathrm{approx}}))$

$\triangleright$ Step 4-b: alternatively, find best slip amplitude sequence using search-and-shrink over cliques
$\mathbf{L}, \mathbf{D} \leftarrow \mathrm{ldl}(\mathbf{Q}^P_{\Delta\hat{z}})$
$\mathrm{candidates}, \mathrm{candidates\_dist2} \leftarrow \mathrm{SASOC}(\mathbf{L}, \mathbf{D}, \hat{a}, N_{\mathrm{clique}}, N_{\mathrm{cand}})$
$i_{\mathrm{best}} \leftarrow \mathrm{argmin}(\mathrm{candidates\_dist2})$
$\Delta\mathbf{z}_{\mathrm{MAP}} \leftarrow \mathrm{candidates}[i_{\mathrm{best}}]$

5.7 Results

Having established a methodology for detecting slip occurrences, we now turn to evaluating its performance on both simulated and real GNSS datasets that are impacted by ionosphere scintillation.

5.7.1 Simulated Data

We simulate realistic ionosphere scintillation measurements with parameters corresponding to the S_4, τ, and C/N$_0$ of Window 1 from Table 5.1. Panels a and b from Figure 5.10 show the C/N$_0$ and phase for the simulated triple-frequency measurements. As is to be expected, we see correlated fading of C/N$_0$ and fluctuations in the phase measurements. The overall phase trend is

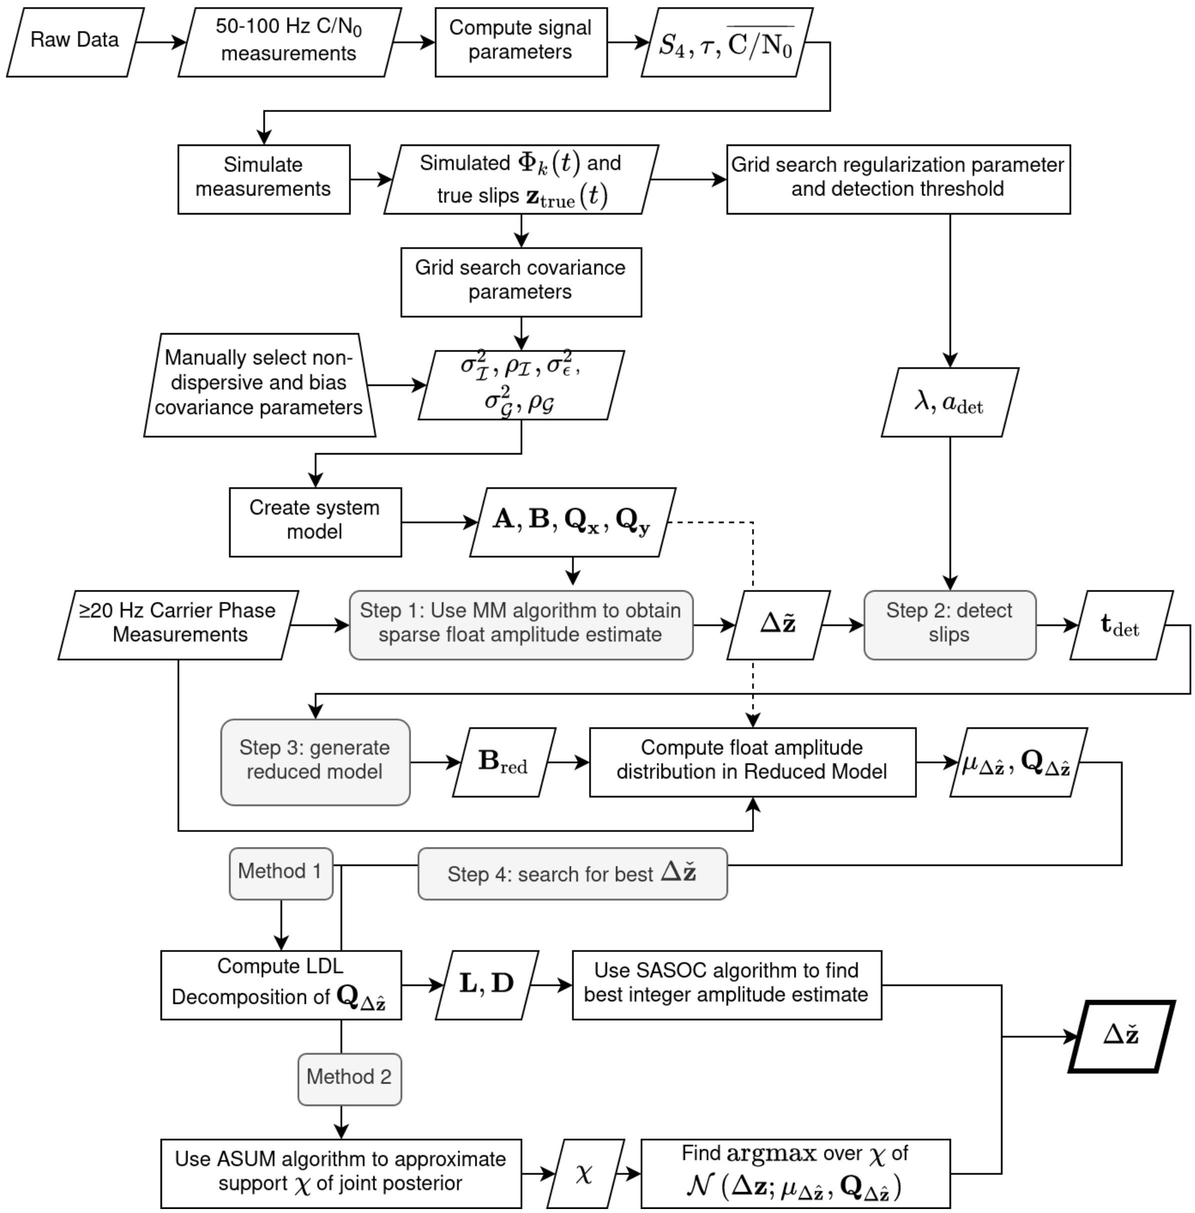

Figure 5.9: Block diagram describing the inputs and processes involved in the cycle slip estimation method developed in this paper.

due to the ionosphere TEC, which for this case of simulated data corresponds to the model phase screen. Panel c shows the diffraction error, which is obtained by subtracting the phase screen from the phase measurements. The diffraction error contains both noise and cycle slips, and the dashed black lines indicate the truth estimate of the slip bias sequence obtained by taking the TVD fit of the diffraction error, which we denote $\mathbf{z}_{\text{true}}$. As is typical for scintillation of this magnitude, the L2 and L5 signals show many simultaneous cycle slips.

Looking closely at the correspondence between the diffraction error, we see it is sometimes not so clear what constitutes a cycle slip in the truth reference. For instance, consider the shaded region between 320-340 seconds (marked Region 6); there is a large fluctuation in the diffraction error for the L5 signal and the truth reference indicates that the phase slips in one direction and then slips back several seconds later. There is a similar but lesser fluctuation in the diffraction error for the L2 signal, but the truth reference does not count them as cycle slips. While some slips have an undeniable and lasting effect on the phase bias, perhaps these instances are better interpreted as large fluctuations in the phase due to diffraction. In fact, these phenomena are the main contribution to the high baseline rate of missed detections that we saw in Figure 5.4. The takeaway here is that it is better to judge our cycle slip estimation results based on their similarity with the truth reference over time rather than evaluating their correspondence for every single slip occurrence.

Panels d and e show $\Delta\tilde{\mathbf{z}}$ and $\check{\mathbf{z}}$, i.e. the sparse float amplitudes and the MAP bias estimates, respectively. Also, the dashed black lines in Panel d indicate the threshold used for detection. We note how the simultaneous slips in the L2 and L5 signals are detected on the L1 signal in the float estimate. The shaded region between 270-300 seconds (marked Region 5) shows one example of this. Fortunately, as we see in the MAP estimate, the fixed amplitudes are correctly estimated as occurring on L2 and L5, and not on L1. Comparing Panels c and e, we observe long-term agreement between the truth reference and MAP estimate of the slip bias sequences; from beginning to end they both show the same overall change in phase bias for each of the three signals. In general, the MAP estimate shares the same slip amplitudes as the truth reference. One important case is between 170-180 seconds (marked Region 4) where the MAP estimate correctly shows a 2-cycle slip in L5, demonstrating the algorithm's ability to correctly deal with different slip magnitudes. In addition to the overall agreement, this example demonstrates successful estimation of two of the most difficult-to-detect slips that can occur for triple-frequency GPS signals. The first is a 1-cycle slip on just the L1 signal, which occurs between 80-100 seconds (marked Region 2), and the second is a simultaneous 1-cycle slip on all three signals, which occurs between 440-450 seconds (marked region 7). These are normally the types of slips that other cycle slip algorithms struggle with, even under moderate conditions, and correct estimation in this instance is a testament to the power of using an extended window of high-rate measurements.

Panels f, g, and h show the IF, GIF, and GF phase combinations. In each panel, the gray line depicts the combination obtained using the raw measurements while the colored line shows the combination after correcting the measurements using the MAP cycle slip estimate. Additionally, for the GF combination (which is scaled to be in TEC units), we show the model phase screen (also scaled to TEC units) in the black dashed line. We see that the IF and GIF combinations are mostly flat after the correction, whereas the raw combinations show many jumps. The GF combination after correction agrees very well with the phase screen trend, except for one instance between 40-70 seconds (marked Region 1). Interestingly, the MAP estimates of slip amplitudes at these epochs are actually opposite of those in the truth reference. It is not clear what exactly causes the algorithm to fail here, but it is worth pointing out that the corrected IF and GIF combinations remain flat during this time, even with the erroneous slip amplitude estimates. More importantly, the algorithm correctly identifies the broader trend in the phase, and the slips revert after a dozen seconds for a net-zero change in cycle bias. There are only a couple other notable discrepancies between the MAP estimate and truth reference. The first occurs with the L5 slips between 320-340 seconds (marked Region 6). As we argued earlier, these missed slips are more of an artifact from generating the truth reference using the phase screen and are not a failure of the algorithm. The other is between 115-130 seconds (marked Region 3) where, according to the truth reference, two consecutive slips occur on the L2 and L5 signals with very short delay. Instead of on L2 and L5, the

MAP estimate interprets these slips as occurring (in the opposite direction) on the L1 signal. This discrepancy is most likely due to the proximity of the detections and the ambiguity of slips on the L1 signal versus simultaneous slips on the L2 and L5 signals. Similar to the other discrepancies, it is of little consequence to the overall agreement in the cycle bias sequences.

To supplement the analysis of this simulation example, we also analyze the cycle slip estimation performance for a batch of simulated measurements. First we ran 100 simulations of scintillation measurements over a 5-minute window using scintillation parameters from Window 1. We then applied the cycle slip estimation algorithm for each simulation run and computed the error due to cycle slips both before and after correction using the MAP cycle slip estimate. For each of these errors, we subtract off the mode of the cycle bias. This is because the overall phase bias is not well-captured by the model we are considering, which is designed only to estimate cycle slips. We also neglect any errors occurring within the first or last 10 seconds of the data window, as faulty estimations during these times are more prevalent due to the lack of measurement context preceding or following the slip detection. Figure 5.11 shows the distributions of these errors for each signal, with the lighter and darker colors indicating the error before and after correction, respectively. The errors before correction reach past ± 5 cycles over the 5-minute window. Meanwhile, the errors after correction are significantly reduced, showing no errors 90% of the time on L1, 85% of the time on L2, and nearly 80% of the time on L5. Moreover, the errors are essentially confined to ± 1 cycle. We identify 3 causes of the errors that persist after correction. 1) The algorithm sometimes simply fails, like in Region 1 of the simulated example in Figure 5.10. Ideally, in these scenarios the errors will not persist, as was the case in the Region 1 example. 2) The algorithm estimates slip correctly, but at slightly different times compared to the truth reference, like in Region 5 from the simulated example. We do not view these errors as a failure of the algorithm, but rather as a byproduct of the way the truth reference is generated and the ambiguity in how scintillation fluctuations are interpreted. 3) Similarly, sometimes phase fluctuations that are interpreted as slips in the truth reference are neglected by the algorithm, like in Region 6 from the simulated example. Since these types of cycle slips are not persistent, we do not consider it a failure if the algorithm does not identify them. It is difficult to discriminate the effects of these latter two causes of errors from actual failures of the algorithm, but we expect that the results summarized in Figure 5.11 underestimate the actual performance of the algorithm. Overall, these results in simulations of scintillation phase measurements suggest that, so long as the algorithm is properly tuned, it should perform well on real scintillation measurements.

5.7.2 Real Data

5.7.2.1 Hong Kong

Here we discuss the results of applying the mitigation algorithm to the real scintillation dataset from Hong Kong. First we assess results for triple-frequency measurements. Figures 5.12 and 5.13 contain panels that illustrate the results for the two windows of the dataset identified in Figure 5.2. Similar to Figure 5.10, panels a and b show the C/N_0 and phase for the three signals. Panels c and d show the sparse float cycle slip amplitude estimates $\Delta\tilde{\mathbf{z}}$ and the the bias sequence estimate $\check{\mathbf{z}}$, respectively. Panels e, f, and g show the IF, GIF, and GF phase combinations before and after cycle slip correction. Just like the simulated example, we see correlated slip estimates in the L2 and L5 signals, and relatively few slips in the L1 signal. Also, for the correlated L2/L5 slips, we see how the float estimate again detects these slips as occurring on the L1 signal. For example, this behavior occurs in Region 3 of Window 1 or Region 5 of Window 2. The similarity in behavior of the float and MAP cycle slip estimates between the simulated and real data is one piece of evidence that the algorithm is behaving properly.

Since we have no truth reference with which to rigorously assess the performance, we turn to phase combinations as another way to evaluate the results. Just like in the simulated example, after correction the IF and GIF combinations remain flat throughout the duration of the window and the GF combination is more smooth. If the algorithm does fail, it is likely to do so in regions where these combinations show noisy or fluctuating behavior, as was the case in the simulated example. For instance, Region 1 marked in Window 1 shows an instance where the IF, GIF, and

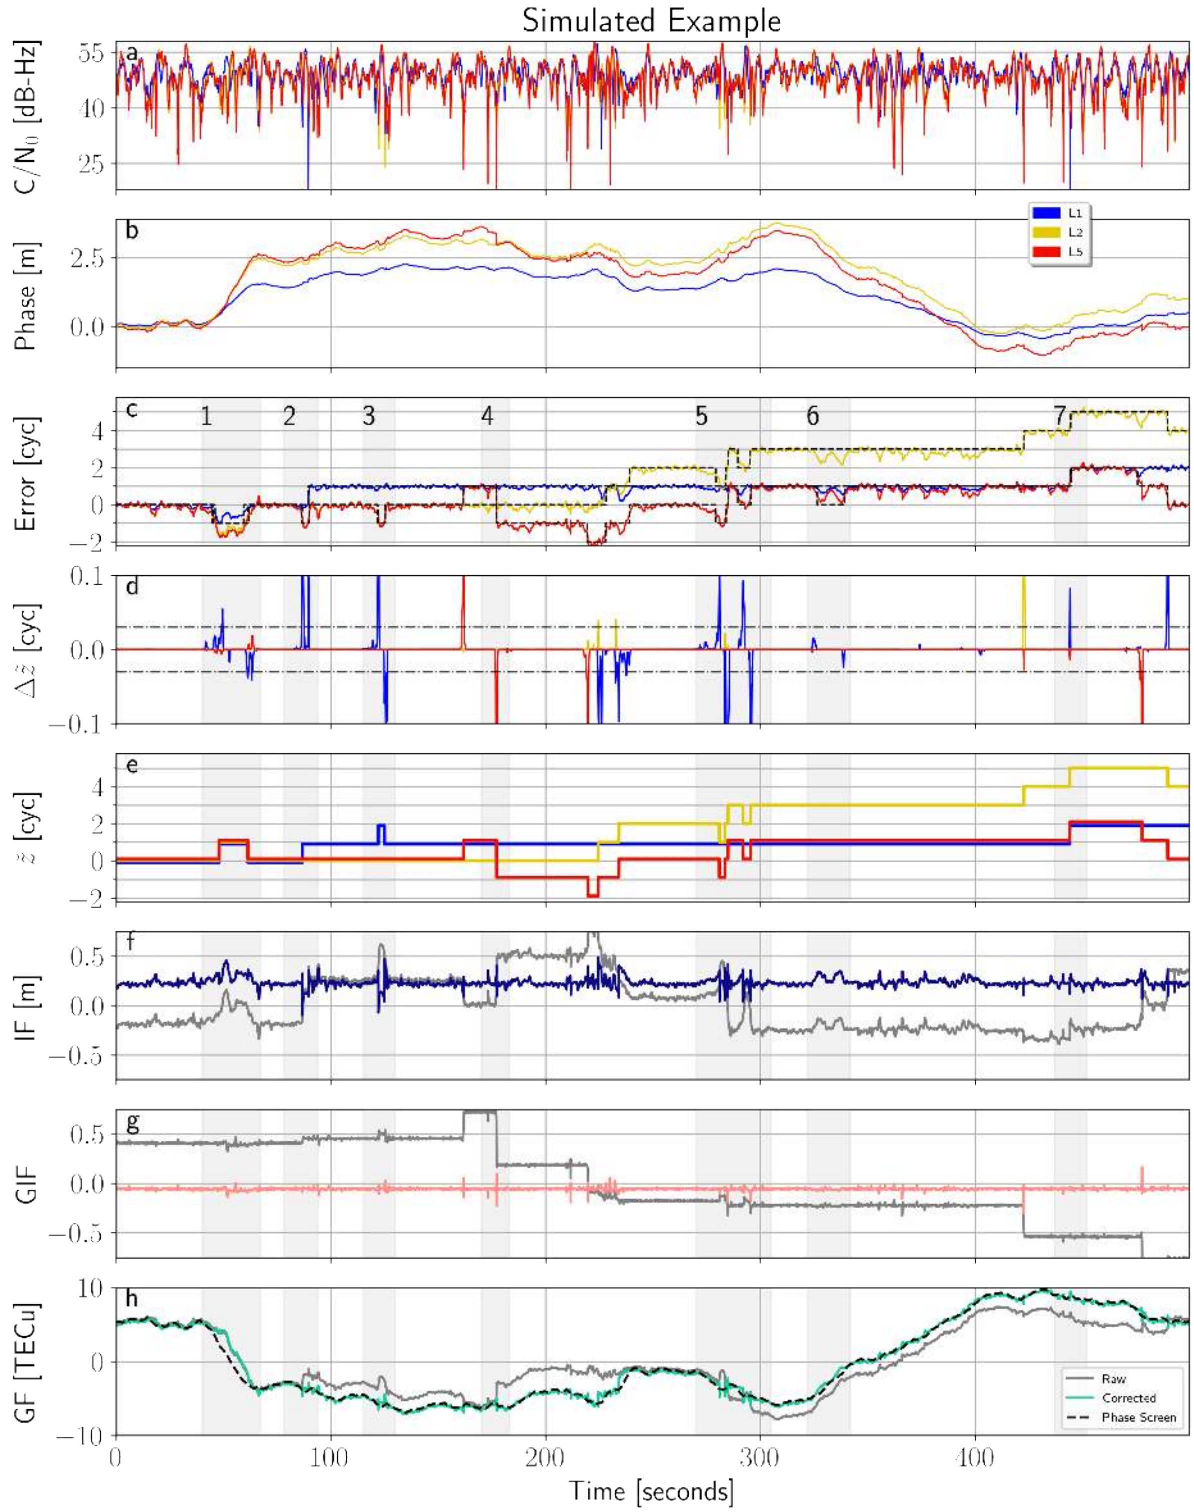

Figure 5.10: Example of simulated triple-frequency scintillation data containing cycle slips. Panels a and b show the C/N$_0$ and phase for the three signals. Panel c shows the diffraction error containing cycle slips along with a truth estimate of slip occurrence. Panel d shows the float estimates of the slip amplitude sequences $\boldsymbol{\Delta\tilde{z}}$ obtained through the sparse estimation method. Panel e shows the MAP estimate of slip bias sequences. Panels f, g, and h show the IF, GIF, and GF phase combinations respectively, with gray and colored lines corresponding to the combinations obtained before (raw) and after (corrected) subtracting off the MAP slip bias estimate $\check{\mathbf{z}}$.

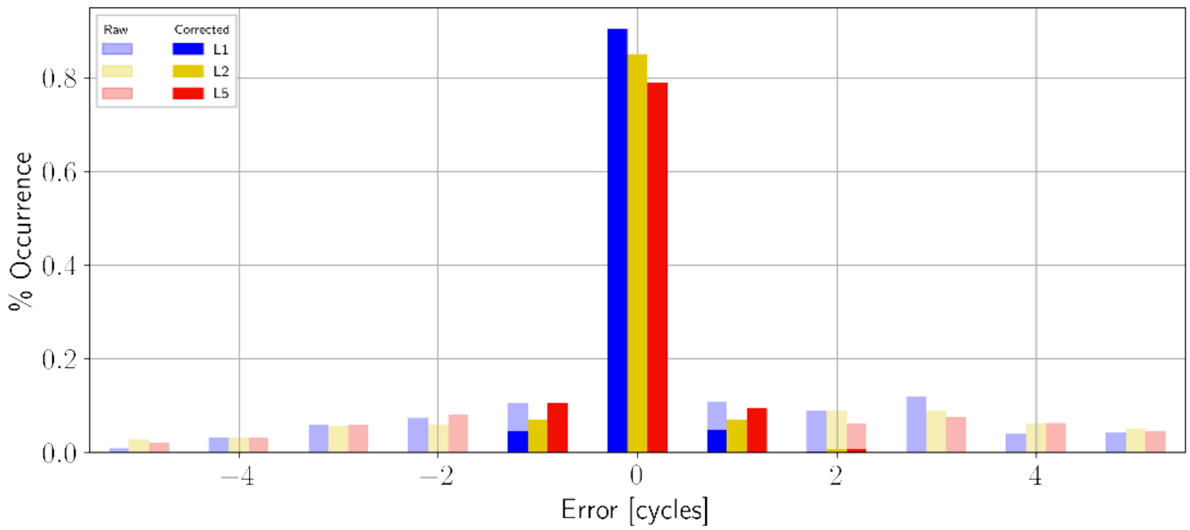

Figure 5.11: Distribution of phase error due to cycle slips before (raw) and after (corrected) subtracting MAP cycle slip estimate. The distribution is generated from histograms of the phase error over 100 simulations of scintillation measurements over 5-minute windows, with simulation parameters corresponding to scenario 1.

GF all show a suspicious bump. During this time, the algorithm estimates a simultaneous 1-cycle slip on all three signals, which is the most difficult cycle slip amplitude to discriminate against the no-slip hypothesis. Another example is Region 4 marked in Window 2, where the GF shows significant variation and the IF shows an increase in noisy fluctuations. The algorithm estimates two slips on the L1 signal during this time, which is generally rare and difficult to identify, so it is quite possible at least one of these was a mis-estimate. On the other hand, the GIF remains quite flat throughout this period and overall the occurrences of slips in the MAP estimate are not suspiciously abundant.

To more rigorously quantify this last statement, we can use the results from Chapter 2 where we modeled cycle slip occurrence as a Poisson process and characterized the slip occurrence rate for different scintillation conditions (using S_4, τ, and C/N$_0$ parameters). For each window of data, Table 5.2 compares the predicted number of slips ($\pm$1-sigma) and the actual number of estimated slips that occur on each signal. We see that for both windows, the number of estimated slips on the L2 and L5 signals is within the 1-sigma bounds of the predicted number of slips. However,

	Window 1			Window 2		
	L1	L2	L5	L1	L2	L5
Predicted	1.6 ± 1.2	14.7 ± 3.8	18.6 ± 4.3	1.2 ± 1.1	11.8 ± 3.4	14.8 ± 3.9
Estimated	10	13	16	6	10	15

Table 5.2: Comparison of predicted number of slips based on results from Chapter 3 and the actual number of estimated slips on each signal.

the number of slips estimated on the L1 signal is larger than the predicted number of slips in both instances. There are a couple instances in Window 1 where the estimated L1 slips are close together and cancel out, just like in the example from Region 3 of the simulated example in Figure 5.10. Aside from this, it is possible that some of the L1 slip estimates are errors, e.g. the L1 slips in Region 4.

It is interesting to compare the results for triple-frequency and dual-frequency estimation. Figure 5.14 shows the results for estimating slips using only the L1 and L2 phase measurements. The panels are similar to those shown in the triple-frequency examples, except for the fact that panel f now shows the GF combination as there is no dual-frequency phase-only GIF combination. Just like for the triple-frequency results, we see much smoother phase combinations after the correction using the estimate slip bias. Comparing the slip bias estimates in panel d, with those from the triple-frequency results in Figure 5.13, we see that the estimates are identical in the latter half of the window. In the first half, however, there are two key discrepancies around 44970 and 45090 seconds, where for the dual-frequency results we see slips occurring on the L2 signal instead of the L1 signal. Given that we saw an overabundance of L1 slips in the triple-frequency estimate, it leads us to believe maybe some of the triple-frequency L1 slip estimates were erroneous. It further suggests that maybe the priors for $\Delta\mathbf{z}$ should not be uniform across different signals but instead should be weighted to reflect the smaller probability of occurrence of L1 slips in comparison to L2 or L5. This will be an interesting topic for future investigation. Overall, despite the extra L1 slip estimates, the algorithm appears to show good performance on the real scintillation datasets in terms of consistency and smoothness in the slip-corrected phase combinations.

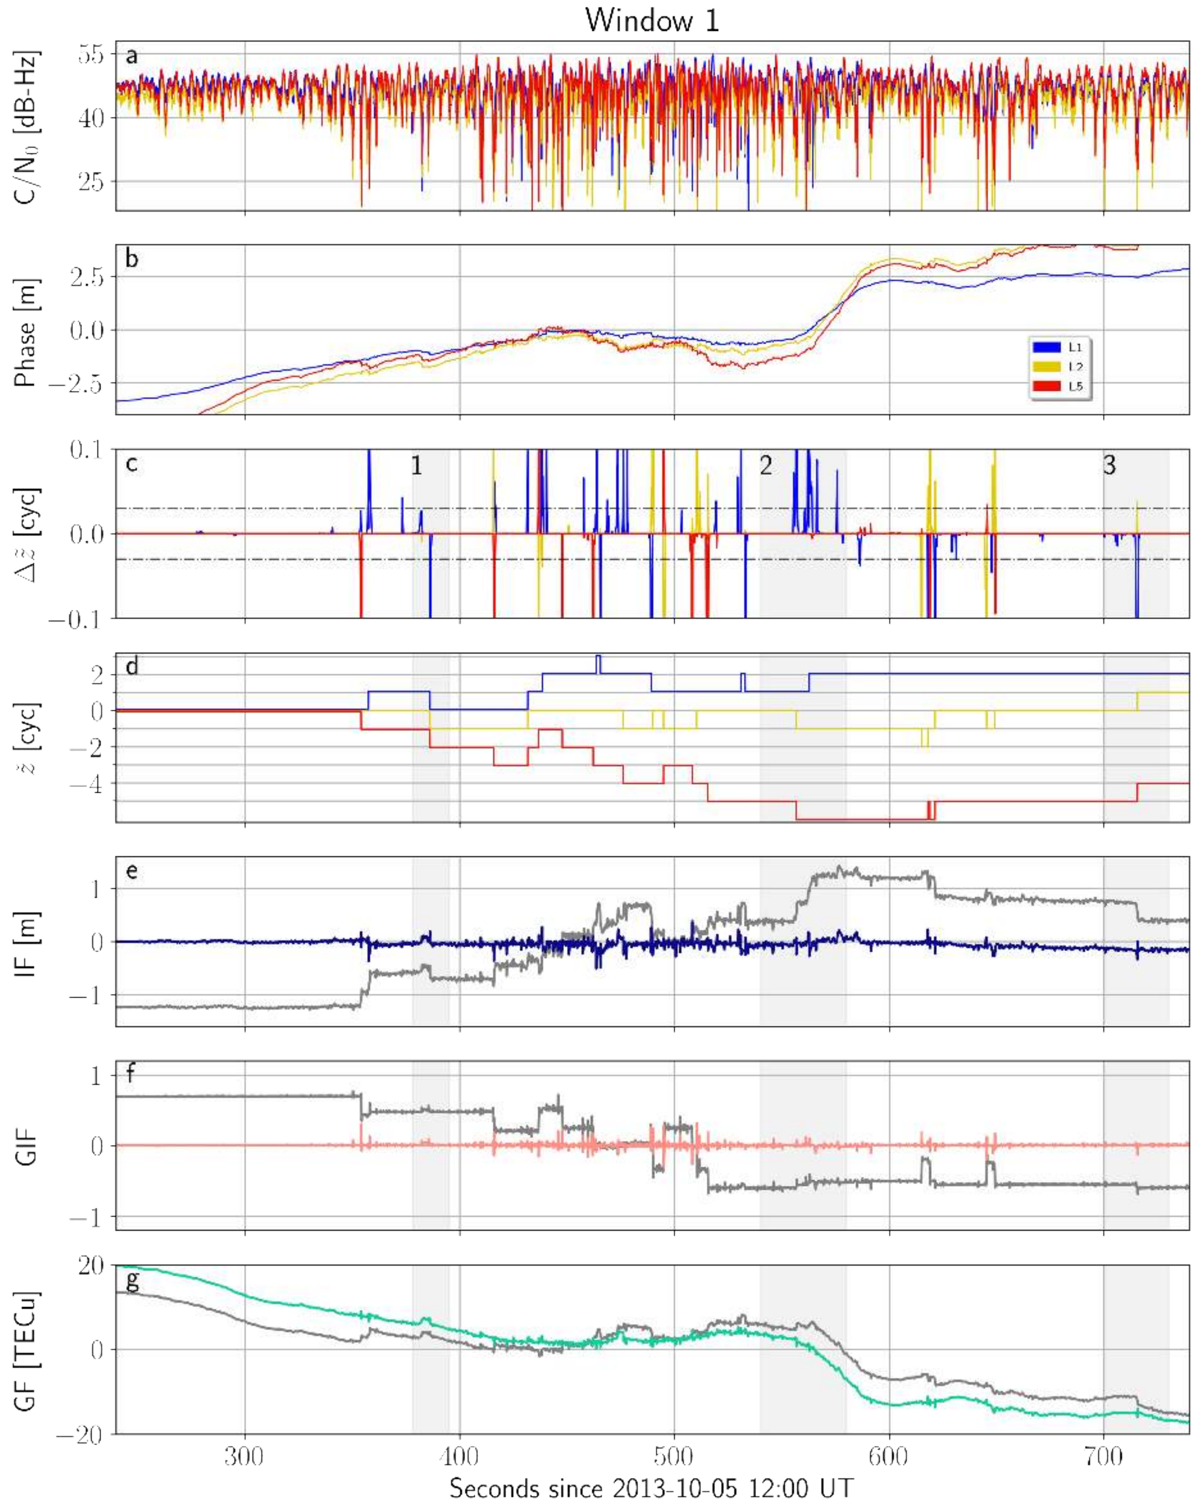

Figure 5.12: Detection and estimation results for Window 1 of the real scintillation data set. Panels a and b show the triple-frequency C/N_0 and phase, respectively. Panel c shows the float estimates of the slip amplitude sequences $\Delta\tilde{\mathbf{z}}$ obtained through the sparse estimation method. Panel d shows the MAP estimate of slip bias sequences. Panels e, f, and g show the IF, GIF, and GF phase combinations respectively, with gray and colored lines corresponding to the combinations obtained before (raw) and after (corrected) subtracting off the MAP slip bias estimate $\check{\mathbf{z}}$.

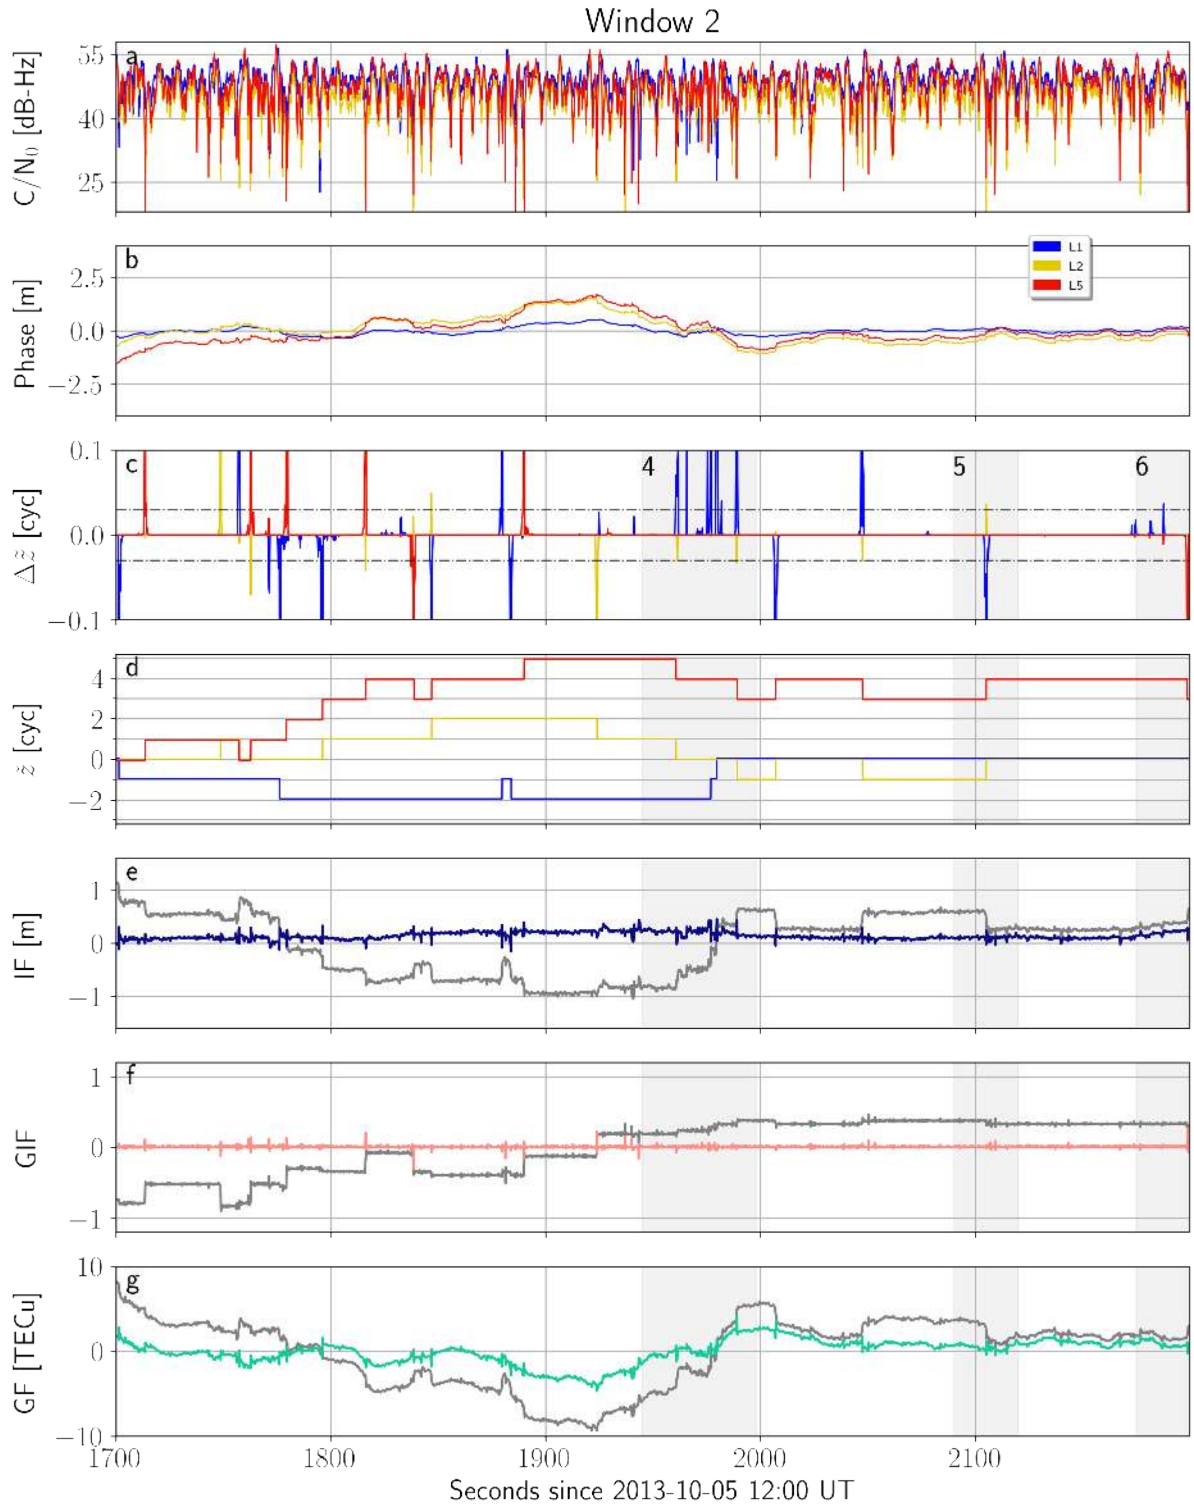

Figure 5.13: Detection and estimation results for Window 2 of the real scintillation data set. Panels and layout are the same as in Figure 5.12.

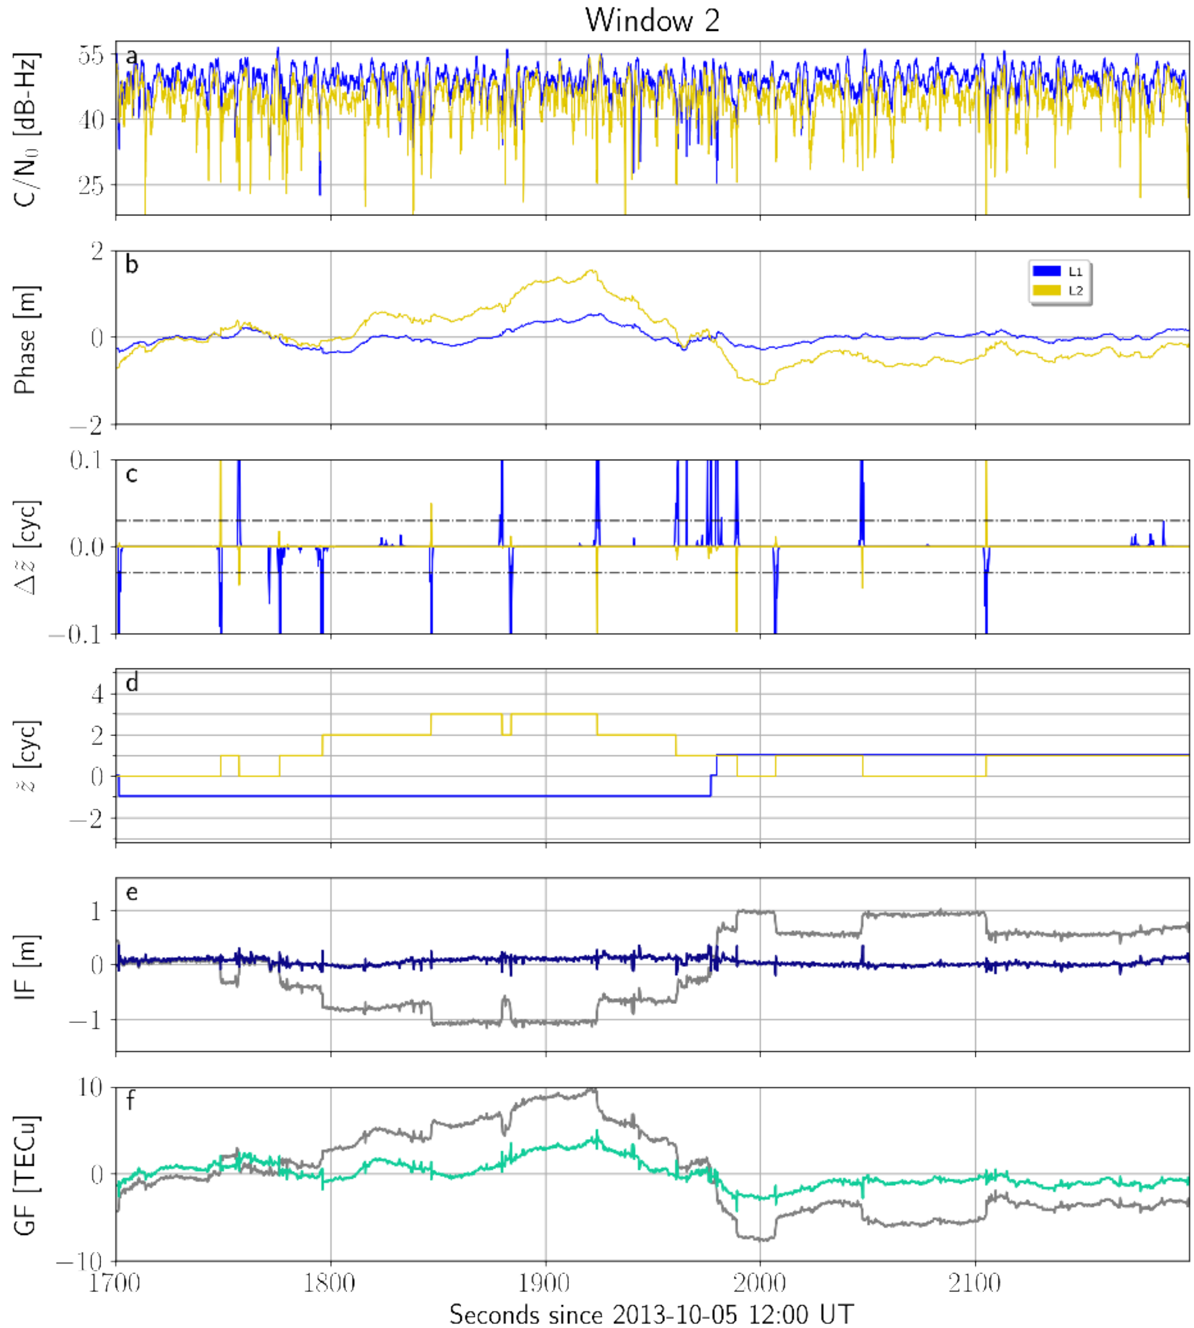

Figure 5.14: Dual-frequency detection and estimation results for Window 2 of the Hong Kong scintillation data set. Similar to Figure 5.12, panels a and b show the C/N_0 and phase, respectively. Panel c shows the float estimates of the slip amplitude sequences $\Delta\tilde{\mathbf{z}}$ obtained through the sparse estimation method. Panel d shows the MAP estimate of slip bias sequences. Panels e and f show the IF and GF phase combinations respectively.

	Units	Ocean Reflection	Mountaintop RO
C/N_0	dB-Hz	22.00	20.00
$\sigma^2_{\mathcal{G}}$	m^2	20	.05
$\rho_{\mathcal{G}}$	s	20	12
$\sigma^2_{\mathcal{I}}$	m^2	1	0.001
$\rho_{\mathcal{I}}$	s	90	90
σ^2_{ϵ}	cycles^2	0.75	0.2
λ		20	15

Table 5.3: Estimation window C/N_0 and hyperparameter values for ocean reflection and mountaintop RO examples.

5.7.2.2 Ocean Reflection

Given the topic that we set in Chapter 1, it is important that we demonstrate the algorithm's effectiveness for other types of harsh signal conditions than just diffractive scintillation. As such, we applied the cycle slip mitigation algorithm to the ocean reflection dataset. When doing so, we need to adjust the model hyperparameters for the new scenario to achieve the best performance. In this case we manually tuned each parameter to achieve adequate performance and runtime. The values we used are listed in Table 5.3. Figure 5.15 shows the results of applying the algorithm for the ocean reflection dataset. Since we do not have as good a model for the non-dispersive component, this time we see more variation in the IF phase combination and nearly completely flat behavior in the GF phase combination. After around 65 seconds, the algorithm appears to perform quite well based on the smooth results in the IF and GF combinations. Before that epoch, the algorithm behavior appears to fail, however this period of the data most likely corresponds to a non-coherent signal [64], and so we do not consider it to be an actual failure of the algorithm.

5.7.2.3 Mountaintop RO

Our final results are for the the mountaintop RO data that we introduced in Chapter 1 Section 1.4.2. In this case we consider measurements on only the L1 signal, which showed several deep fades and had overall quite low C/N_0 during the estimation window we chose. In this dataset, the navigation bits have not been removed, and so the signal is actually afflicted with half-cycle

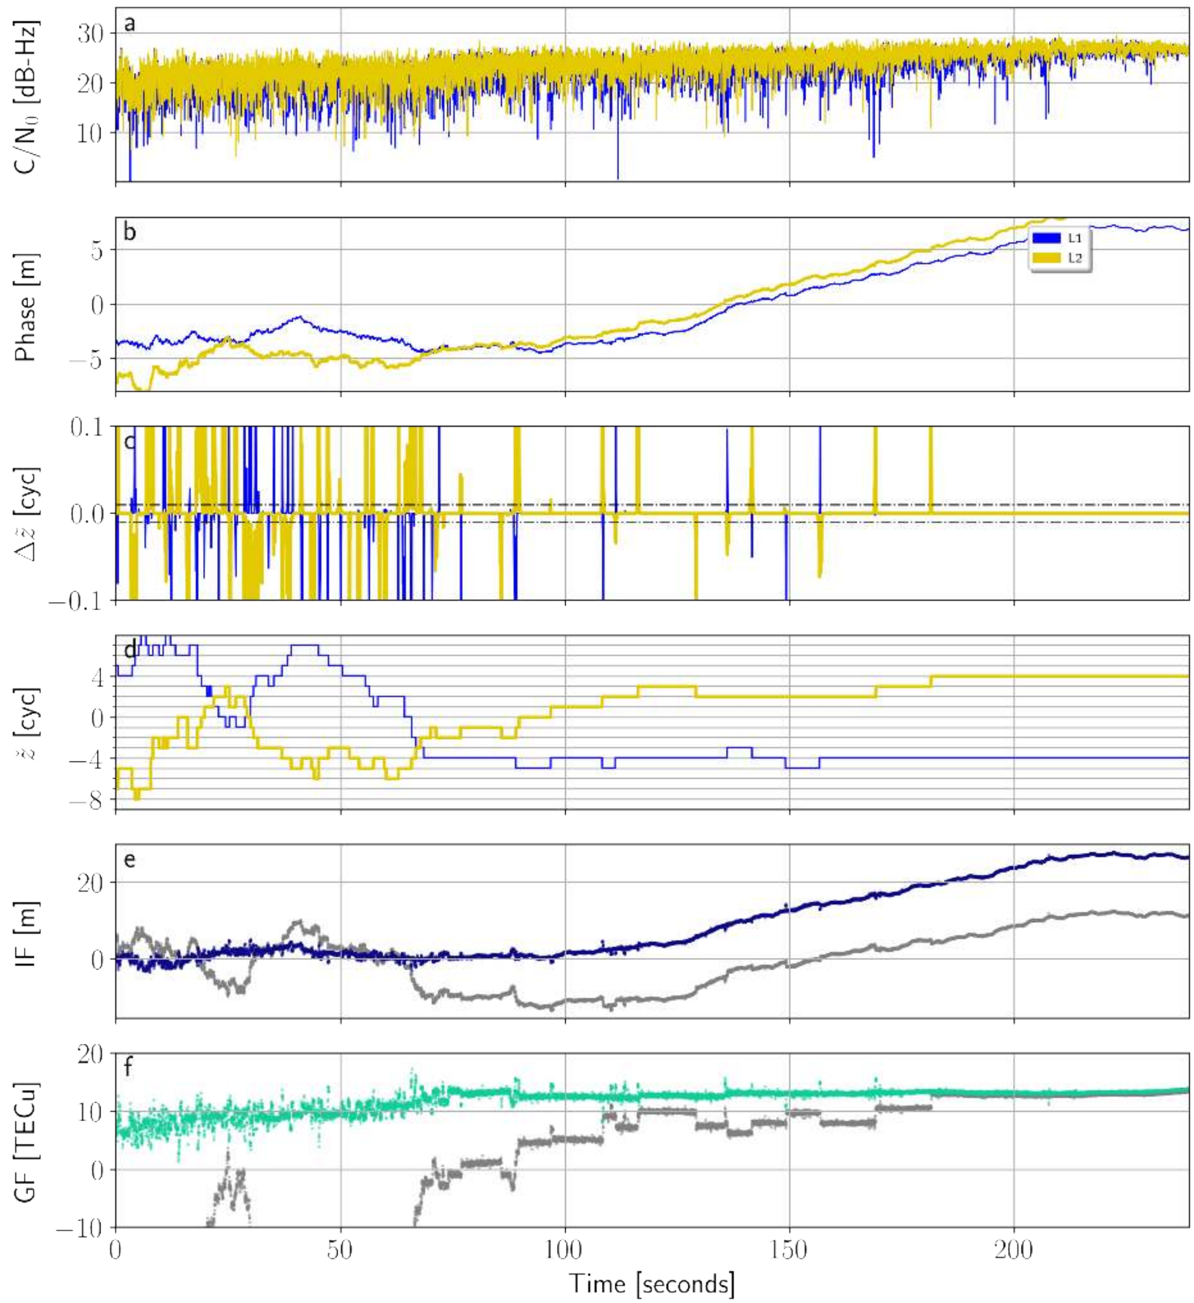

Figure 5.15: Dual-frequency detection and estimation results for the GNSS-R ocean reflection dataset. Similar to Figure 5.14, panels a and b show the C/N_0 and phase, respectively. Panel c shows the float estimates of the slip amplitude sequences $\boldsymbol{\Delta\tilde{z}}$ obtained through the sparse estimation method. Panel d shows the MAP estimate of slip bias sequences. Panels e and f show the IF and GF phase combinations respectively.

instead of integer-cycle slips. To accommodate for this, as briefly mention in Section 4.2, we simply adjust $\mathbf{B}$ to use 1/2 wavelengths. Again, we hand tune the covariance parameters (provided in Table 5.3) to achieve decent results with reasonable runtime. Note that in the single-frequency case, specification of both $\mathcal{G}$ and $\mathcal{I}$ phase components is somewhat redundant since the single-frequency signal measurements cannot take advantage of the dispersive nature of the ionosphere effect. Nevertheless, slips using a single frequency can still be estimated and the model does not require any special modification. Figure 5.16 shows the results of applying the detection and estimation procedures. In panel c, we see very rapid occurrence of detected slips over very short spans of time. Panel d shows the estimated slip bias sequence that noticeably follows the phase trend in the raw measurements, which is presumably due to cycle slips. Despite the very low C/N_0 level, we are able to estimate slips in the single-frequency data because the phase components have little variation. Panel e shows the corrected phase measurements, now showing much smoother behavior. With only one signal frequency, it is difficult to assess the fidelity of the corrected phase. One way to do this would be to compare the corrected phase with its expected variations due to physical parameters, e.g. in this case due to tropospheric water vapor content. For instance, the in [85] the cycle slip correction results for ocean-reflected signals are validated by comparing corrected phase variations with sea-surface height variations obtained via an independent altimetry experiment. Such validation is an interesting topic for future study.

5.8 Summary and Discussion

In this chapter, we presented a comprehensive approach for detecting and estimating cycle slips. We introduced a novel method for batch cycle slip detection over a window of measurements, and we assessed the performance of the approach by applying it to high-rate multi-frequency measurements in both simulated and real data. The results of applying this method to simulated scintillation data sets indicate that it can correctly estimate the most difficult-to-detect slip amplitudes and that the MAP estimate can correctly identify at least 80% of the cycle slips. For the real datasets, we applied the method for dual- and triple-frequency GPS measurements for the cases

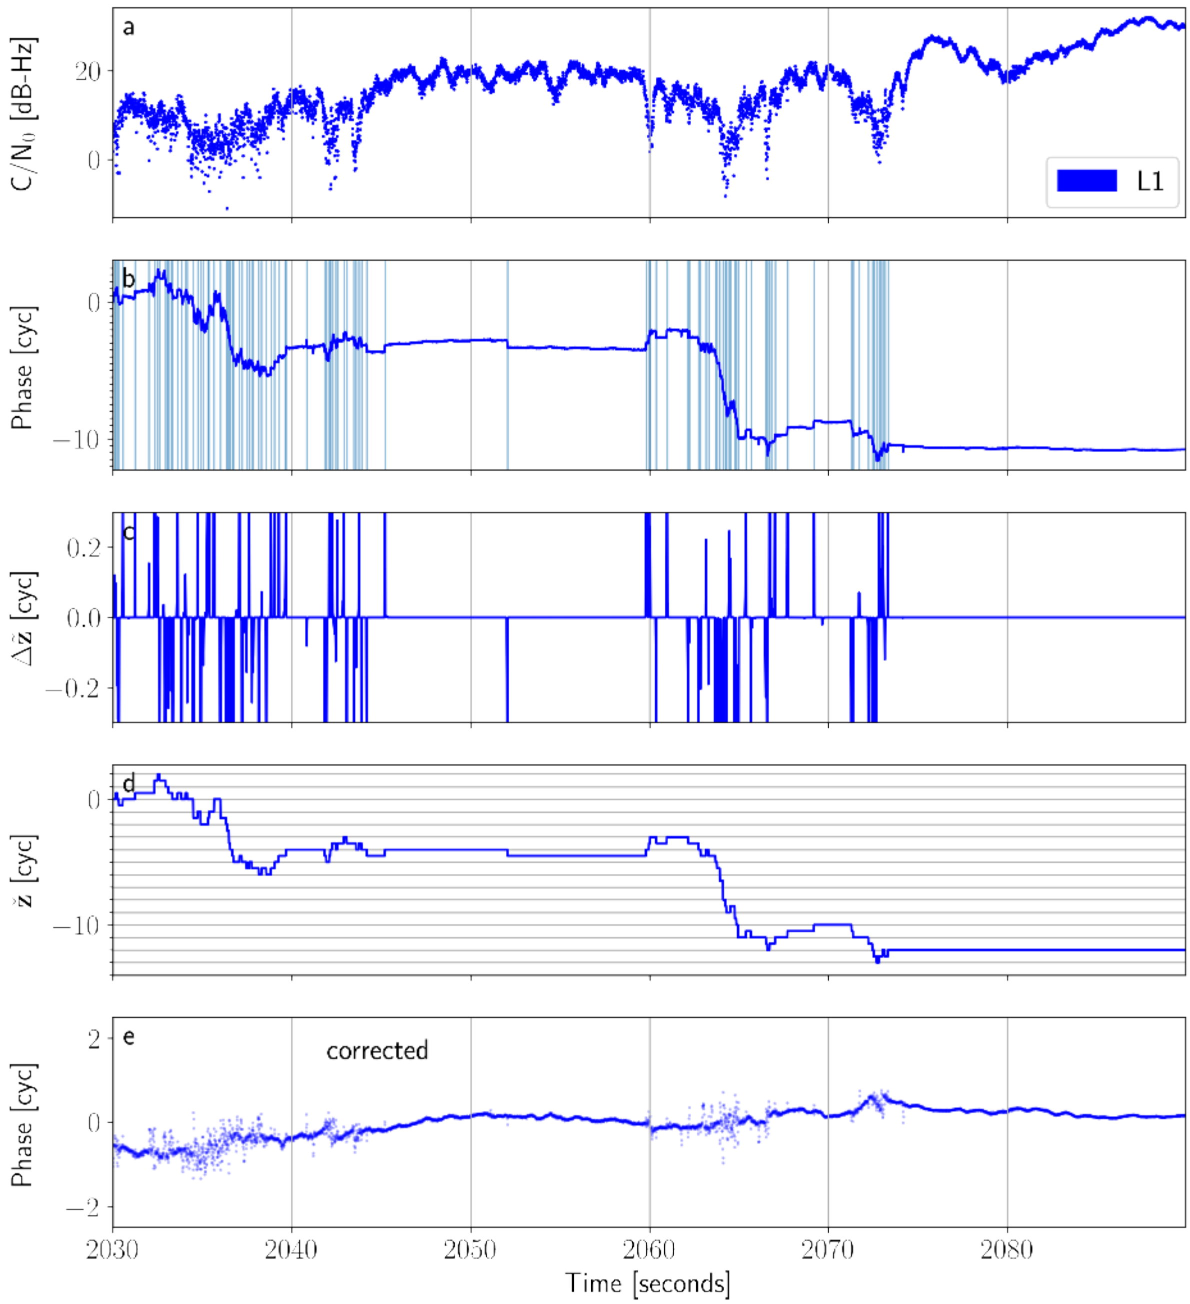

Figure 5.16: Single-frequency detection and estimation results for the Hawaii mountaintop RO dataset. Similar to Figure 5.14, panels a and b show the C/N$_0$ and phase, respectively. Panel c shows the float estimates of the slip amplitude sequences $\Delta\tilde{\mathbf{z}}$ obtained through the sparse estimation method. Panel d shows the MAP estimate of slip bias sequences.

of diffractive ionosphere scintillation and weak ocean reflection. There is no pure truth reference with which to validate results on the real scintillation measurements, however, the smoothing of the corrected phase combinations suggests that the algorithm is effective. When the algorithm does fail, it appears to be associated with estimations of L1 slips over simultaneous slips on L2 and L5.

While we mostly focused on ionosphere scintillation in our results, we intend for this work to be applicable in a wide variety of harsh signal environments. GNSS remote sensing is ripe with opportunity to apply this method, including semi-coherent ocean reflections used in GNSS-R, or tropospheric scintillation effects on low-elevation and radio occultation signals. In addition to remote-sensing applications, this work may also be useful for batch estimation of cycle slips occurring in weak signals on dynamic platforms. A major part of the effort in this work was devoted to identifying appropriate covariance hyperparameters for modeling phase components and measurement noise. The key to generalizing this algorithm to other harsh signal scenarios is finding appropriate hyperparameters, although the values we use in this work should be a good starting point. We also made a number of stated and implicit assumptions about how to characterize the phase components, such as our use of the Matérn kernel with $\nu = 3/2$ and our characterization of noise and unmodeled errors, which in this study was held constant over time. The approach we presented can easily be adapted to incorporate a time-varying model for noise variance, which may help improve mitigation capabilities under certain scenarios. It will be important to reassess these assumptions and characterizations of the signal phase components and errors for these new scenarios.

Chapter 6

Summary and Conclusion

In this dissertation, we explored the origin and nature of cycle slips in GNSS phase measurements, with an emphasis on cycle slip occurrence under harsh signal conditions. This problem is challenging because of 1) the occurrence of numerous consecutive cycle slips, 2) presence of large, uncertain phase component variations, and/or 3) presence of excessive noise. In chapter 2, we focused on characterization of cycle slips in simulations. We showed how the effects of phase transitions and noise can create many consecutive cycle slips, and we quantified the rate of this slip occurrence for the case of diffractive ionosphere scintillation. In chapter 3, we took a closer look at cycle slips in real multi-frequency phase measurements collected during scintillation events and we provided a performance analysis of two algorithms designed to mitigate their occurrences. The first of these algorithms used measurement combinations to detect and estimate cycle slips and yielded a very high rate of false detection when applied to one of the scintillation datasets. The second algorithm was a state-space sequential algorithm that made use of the signal C/N_0 to adaptively filter the phase measurements and detect cycle slips. While this method performed significantly better than the method using phase combinations, it still introduced a large number of false cycle slips into the resulting measurements. The filtering algorithm is of particular interest because it had previously been applied with very good performance to scenarios with low signal C/N_0. We claim that the case of ionosphere diffraction, and any harsh signal conditions that contain phase transitions, will be difficult for sequential algorithms to address. This claim stems from our characterization of phase transitions in Chapter 2, where we saw how an "up-close" view could not necessarily reveal whether or not canonical fading corresponds to a cycle slip. Our conclusion from these results was that effective cycle slip mitigation under harsh signal scenarios would require use of an extended window of high-rate measurements when detecting and estimating slips.

In order to quantify how long a window and how high a sampling rate are necessary, in Chapter 4 we introduce a general probabilistic model for estimating cycle slips in GNSS carrier phase measurements. In particular, we modeled the non-dispersive and refractive ionosphere phase components as stationary Gaussian processes. We discussed noise modeling and how to choose covariance parameters. Then, we used our model to illustrate how window duration and sampling rate impact false detection rates under 4 hypothetical scenarios, showing that we can greatly decrease false detection probabilities for a single slip when using 20 Hz (or faster) multi-frequency measurements over at least 16 seconds. However, our characterization from chapters 2 and 3 demonstrate that it is possible for multiple cycle slips to occur over such timescales. This led us into the culminating problem from this dissertation: how to reliably detect and estimate the occurrence of multiple cycle slips under harsh signal conditions. In Chapter 5, we provided our approach, the steps of which we summarize as 1) compute the sparse float estimate, 2) use the sparse float estimate to detect slip occurrences, 3) compute the reduced system model for the detected slips, 4) search for the best cycle slip amplitudes. The results demonstrate that this algorithm is capable of accurately detecting and estimating cycle slip occurrences in both simulated and real datasets, including both single-, dual-, and triple-frequency measurements.

The approach we introduce has an advantage over other cycle slip mitigation techniques because it uses all the relevant measurements both before and after slip occurrences when estimating their amplitudes. Our development and algorithms also allow for its use with high-rate measurements and any number of signal frequencies. However, this power and flexibility come at a cost. The algorithm is fairly computationally expensive compared to most other cycle slip algorithms. One of the major challenges in this work was in finding ways to make the algorithm computationally feasible for large windows of high-rate measurements without compromising its power and flexibility. Our use of stationary Gaussian processes, the L1-norm regularization optimization approach during detection, and our new approaches to searching for integer solutions in high dimensions (SASOC, ASUM) were key elements to our success in overcoming this challenge. Also, unlike many other mitigation strategies, the algorithm we present is designed for batch post-processing. Exploring how well the model principles from Chapter 4 can translate to sequential or near-real-time processing will be an interesting topic for future study.

While our model and algorithms are powerful and perform well, we still made a number of assumptions that are possibly too rigid or over-simplifying. One of these was our assumption about noise modeling. We mentioned more sophisticated options for noise modeling in the case of ionosphere scintillation, e.g. those in [73]. We expect that using a noise covariance that adapts based on estimates of signal $\mathrm{C/N_0}$, similar to the cycle slip filtering algorithm [85] or many carrier tracking approaches (c.f. [95], [73], [83]), could lead to better performance under certain circumstances. Also, our modeling of the time series $\mathcal{G}(t)$ and $\mathcal{I}(t)$ relied on assumptions about the behavior of these components. In particular, it will be valuable for future studies to more fully characterize when they can adequately be modeled as stationary and what are appropriate covariance models. As we discussed in Chapter 4, for the case of ionosphere scintillation, diffractive fluctuations can be highly correlated. It would be interesting to see to what extent diffraction can be absorbed into the refractive ionosphere phase component and how that affects its covariance structure.

Finally, while throughout this dissertation we heavily focused on the mitigation of the cycle slips themselves, it is important to also consider the actual end-goal that motivated us to address these cycle slips in the first place: obtaining accurate, slip-free phase measurements for our precision remote-sensing or navigation applications. Throughout our results section, we showed the "smoothed" phase combinations obtained after applying a correction using the estimated cycle slips, which shows some indication that we have removed the slip-induced measurement biases. It will also be important for some applications to know when cycle slips were corrected, as well as any potential statistics for the reliability of that correction. We presented limited results on the theoretical performance of the algorithm with an emphasis on ionosphere diffraction conditions since those are the ones we can simulate. It will be valuable to further investigate algorithm performance under other harsh signal scenarios.

Bibliography

[1] Gerd Ascheid and Heinrich Meyr. Cycle slips in phase-locked loops: A tutorial survey. IEEE transactions on communications, 30(10):2228–2241, 1982.

[2] S Banville, RB Langley, S Saito, and T Yoshihara. Handling cycle slips in GPS data during ionospheric plasma bubble events. Radio Science, 45(06):1–14, 2010.

[3] Simon Banville and Richard B Langley. Mitigating the impact of ionospheric cycle slips in GNSS observations. Journal of Geodesy, 87(2):179–193, 2013.

[4] Alvaro Barbero and Suvrit Sra. Fast newton-type methods for total variation regularization. In Proceedings of the 28th International Conference on Machine Learning (ICML-11), pages 313–320. Citeseer, 2011.

[5] Alvaro Barbero and Suvrit Sra. Modular proximal optimization for multidimensional total-variation regularization. The Journal of Machine Learning Research, 19(1):2232–2313, 2018.

[6] L Bastos and H Landau. Fixing cycle slips in dual frequency kinematic GPS-applications using kalman filtering. manuscripta geodaetica. 1988.

[7] L Bastos and H Landau. Fixing cycle slips in dual-frequency kinematic GPS-applications using kalman filtering. 1998.

[8] Geoffrey Blewitt. An automatic editing algorithm for GPS data. Geophysical research letters, 17(3):199–202, 1990.

[9] Brian Breitsch and Yu Morton. Triple-frequency GNSS cycle slip detection performance in the presence of diffractive ionosphere scintillation. In 2020 IEEE/ION Position, Location and Navigation Symposium-PLANS 2020. IEEE, 2020.

[10] Brian Breitsch, Yu T Morton, Charles Rino, and Dongyang Xu. GNSS carrier phase cycle slips due to diffractive ionosphere scintillation: Simulation and characterization. IEEE Transactions on Aerospace and Electronic Systems, 2020.

[11] Brian Breitsch, Yu T Morton, Dongyang Xu, and Rong Yang. Ionosphere scintillation-induced phase transitions in triple-frequency GPS measurements. In Proceedings of the 2020 International Technical Meeting of The Institute of Navigation, pages 912–919, 2020.

[12] Brian Breitsch, Yang Wang, and Y Morton. Performance of cycle slip filtering algorithm during ionosphere scintillation. In Proceedings of the 33rd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2020), pages 3132–3139, 2020.

[13] Brian W Breitsch. Master's Thesis: Linear combinations of GNSS phase observables to improve and assess TEC estimation precision. 2017.

[14] Changsheng Cai, Zhizhao Liu, Pengfei Xia, and Wujiao Dai. Cycle slip detection and repair for undifferenced GPS observations under high ionospheric activity. GPS solutions, 17(2):247–260, 2013.

[15] Sebastien Carcanague. Real-time geometry-based cycle slip resolution technique for single-frequency ppp and rtk. In Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012), pages 1136–1148, 2012.

[16] Charles S Carrano and Keith M Groves. Temporal decorrelation of GPS satellite signals due to multiple scattering from ionospheric irregularities. In Proceedings of the 23rd International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2010), pages 361–374, 2010.

[17] Charles S Carrano, Keith M Groves, William J McNeil, and Patricia H Doherty. Direct measurement of the residual in the ionosphere-free linear combination during scintillation. In Proceedings of the 2013 Institute of Navigation ION NTM meeting, pages 585–596, San Diego, CA, 2013.

[18] Marc Cocard, Stéphanie Bourgon, Omid Kamali, and Paul Collins. A systematic investigation of optimal carrier-phase combinations for modernized triple-frequency GPS. Journal of Geodesy, 82(9):555–564, 2008.

[19] Paul De Jonge, CCJM Tiberius, et al. The lambda method for integer ambiguity estimation: implementation aspects. Publications of the Delft Computing Centre, LGR-Series, 12(12):1–47, 1996.

[20] Maria Clara de Lacy, Mirko Reguzzoni, and Fernando Sansò. Real-time cycle slip detection in triple-frequency GNSS. GPS solutions, 16(3):353–362, 2012.

[21] Maria Clara de Lacy, Mirko Reguzzoni, Fernando Sansò, and Giovanna Venuti. The bayesian detection of discontinuities in a polynomial regression and its application to the cycle-slip problem. Journal of Geodesy, 82(9):527–542, 2008.

[22] MC De Lacy, Fernando Sanso, G Rodriguez-Caderot, and AJ Gil. The bayesian approach applied to GPS ambiguity resolution. a mixture model for the discrete–real ambiguities alternative. Journal of geodesy, 76(2):82–94, 2002.

[23] Alison de Oliveira Moraes, Emanoel Costa, Mangalathayil Ali Abdu, Fabiano S Rodrigues, Eurico Rodrigues de Paula, Kelias Oliveira, and Waldecir João Perrella. The variability of low-latitude ionospheric amplitude and phase scintillation detected by a triple-frequency GPS receiver. Radio Science, 52(4):439–460, 2017.

[24] Alison de Oliveira Moraes, Fabiano da Silveira Rodrigues, Waldecir João Perrella, and Eurico Rodrigues de Paula. Analysis of the characteristics of low-latitude GPS amplitude scintillation measured during solar maximum conditions and implications for receiver performance. Surveys in geophysics, 33(5):1107–1131, 2012.

[25] Delft Geodetic Computing Centre. LAMBDA Software Package. Delft University of Technology, Delft, Netherlands, 2020.

[26] Chenlong Deng, Jianhui Cui, Weiming Tang, Xuan Zou, and Liang Shu. Reliable real-time triple-frequency cycle slip detection and recovery with adaptive detection thresholds. Measurement Science and Technology, 30(5):055007, 2019.

[27] Michael J Dunn and DAF DISL. Global positioning system directorate systems engineering & integration interface specification IS-GPS-200, 2012.

[28] Z Gao, W Wang, S Ji, and Z Wang. Kinematic cycle slip detection and correction for carrier phase based navigation applications in urban environment in case of ultra high rate GNSS observations. Survey Review, 47(344):371–378, 2015.

[29] Javier G Garcia, Pedro A Roncagliolo, and Carlos H Muravchik. A bayesian technique for real and integer parameters estimation in linear models and its application to GNSS high precision positioning. IEEE Transactions on Signal Processing, 64(4):923–933, 2015.

[30] Vadim E Gherm, Nikolay N Zernov, and Hal J Strangeways. Propagation model for transionospheric fluctuating paths of propagation: Simulator of the transionospheric channel. Radio Science, 40(1):1–9, 2005.

[31] Vadim E Gherm, Nikolay N Zernov, and Hal J Strangeways. Effects of diffraction by ionospheric electron density irregularities on the range error in GNSS dual-frequency positioning and phase decorrelation. Radio Science, 46(03):1–10, 2011.

[32] Clyde C Goad. Precise positioning with the GPS. In Applied Geodesy, pages 17–30. Springer, 1987.

[33] Jacob Goldberger and Amir Leshem. A gaussian tree approximation for integer least-squares. In Advances in Neural Information Processing Systems, pages 638–645, 2009.

[34] Georgii Ksenofontovich Golubev and Ekaterina A Krymova. On interpolation of smooth processes and functions. Problems of Information Transmission, 49(2):127–148, 2013.

[35] Lingyong Huang, Zhiping Lu, Guojun Zhai, Yongzhong Ouyang, Motao Huang, Xiuping Lu, Taiqi Wu, and Kaifeng Li. A new triple-frequency cycle slip detecting algorithm validated with BDS data. GPS solutions, 20(4):761–769, 2016.

[36] Todd E Humphreys, Mark L Psiaki, and Paul M Kintner. Modeling the effects of ionospheric scintillation on GPS carrier phase tracking. IEEE Transactions on Aerospace and Electronic Systems, 46(4):1624–1637, 2010.

[37] Todd E Humphreys, Mark L Psiaki, Brent M Ledvina, Alessandro P Cerruti, and Paul M Kintner. Data-driven testbed for evaluating GPS carrier tracking loops in ionospheric scintillation. IEEE Transactions on Aerospace and Electronic Systems, 46(4):1609–1623, 2010.

[38] Norbert Jakowski, Christoph Mayer, Volker Wilken, and Mohammed M Hoque. Ionospheric impact on GNSS signals. Física de la Tierra, 20:11, 2008.

[39] HW Janes, RB Langley, and SP Newby. Analysis of tropospheric delay prediction models: comparisons with ray-tracing and implications for GPS relative positioning. Bulletin géodésique, 65(3):151–161, 1991.

[40] Shahram Jazaeri, Alireza Amiri-Simkooei, and Mohammad Ali Sharifi. On lattice reduction algorithms for solving weighted integer least squares problems: comparative study. GPS solutions, 18(1):105–114, 2014.

[41] Yu Jiao, Yu T Morton, Steven Taylor, and Wouter Pelgrum. Characterization of high-latitude ionospheric scintillation of GPS signals. Radio Science, 48(6):698–708, 2013.

[42] Yu Jiao, Dongyang Xu, Yu Morton, and Charles Rino. Equatorial scintillation amplitude fading characteristics across the GPS frequency bands. Navigation: Journal of The Institute of Navigation, 63(3):267–281, 2016.

[43] Yu Jiao, Dongyang Xu, Charles L Rino, Yu T Morton, and Charles S Carrano. Multi-frequency GPS signal strong equatorial ionospheric scintillation simulator: Algorithm, performance, and characterization. IEEE Transactions on Aerospace and Electronic Systems, 54(4):1947–1965, 2018.

[44] JM Juan, A Aragon-Angel, J Sanz, Guillermo González-Casado, and Adrià Rovira-Garcia. A method for scintillation characterization using geodetic receivers operating at 1 hz. Journal of Geodesy, 91(11):1383–1397, 2017.

[45] José Miguel Juan, Jaume Sanz, Guillermo González-Casado, Adrià Rovira-Garcia, Adriano Camps, Jaume Riba, José Barbosa, Estefania Blanch, David Altadill, and Raul Orus. Feasibility of precise navigation in high and low latitude regions under scintillation conditions. Journal of Space Weather and Space Climate, 8:A05, 2018.

[46] George S Kimeldorf and Grace Wahba. A correspondence between bayesian estimation on stochastic processes and smoothing by splines. The Annals of Mathematical Statistics, 41(2):495–502, 1970.

[47] PM Kintner, BM Ledvina, ER De Paula, and IJ Kantor. Size, shape, orientation, speed, and duration of GPS equatorial anomaly scintillations. Radio Science, 39(2):1–23, 2004.

[48] Jan Kouba. A guide to using international GNSS service (IGS) products, 2009.

[49] Pan Li, Xinyuan Jiang, Xiaohong Zhang, Maorong Ge, and Harald Schuh. Kalman-filter-based undifferenced cycle slip estimation in real-time precise point positioning. GPS Solutions, 23(4):99, 2019.

[50] Zhizhao Liu. A new automated cycle slip detection and repair method for a single dual-frequency GPS receiver. Journal of Geodesy, 85(3):171–183, 2011.

[51] Khurram Mazher and Muhammad Tahir. Small cycle slip detection using singular spectrum analysis. In 2016 24th European Signal Processing Conference (EUSIPCO), pages 1053–1057. IEEE, 2016.

[52] Y Miao, ZW Sun, and SN Wu. Error analysis and cycle-slip detection research on satellite-borne GPS observation. Journal of Aerospace Engineering, 24(1):95–101, 2011.

[53] Oliver Montenbruck, Urs Hugentobler, Rolf Dach, Peter Steigenberger, and André Hauschild. Apparent clock variations of the Block IIF-1 (SVN62) GPS satellite. GPS solutions, 16(3):303–313, 2012.

[54] Viet Khoi Nguyen, Adria Rovira-Garcia, José Miguel Juan, Jaume Sanz, Guillermo González-Casado, Tung Hai Ta, et al. Measuring phase scintillation at different frequencies with conventional GNSS receivers operating at 1 hz. Journal of Geodesy, 93(10):1985–2001, 2019.

[55] Carey E Noll. The crustal dynamics data information system: A resource to support scientific analysis using space geodesy. Advances in Space Research, 45(12):1421–1440, 2010.

[56] Alexander Pavelyev, Jens Wickert, Christoph Reigber, Torsten Schmidt, Yuei-An Liou, Chen-Young Huang, Stanislav Matyugov, and Dmitrii Pavelyev. Amplitude variations in CHAMP radio occultation signal as an indicator of the ionospheric activity. In Earth Observation with CHAMP, pages 431–440. Springer, 2005.

[57] Mark L Psiaki, Gary S Bust, Alessandro P Cerruti, Paul M Kintner Jr, and Steven P Powell. Diffraction tomography of the disturbed ionosphere based on GPS scintillation data. Proc. of the ION GNSS 2008, pages 16–19, 2008.

[58] Mark L Psiaki, T Humphreys, A Cerruti, S Powell, and P Kintner. Tracking L1 C/A and L2C signals through ionospheric scintillations. In ION GNSS, volume 20, pages 25–28. Citeseer, 2007.

[59] Chuang Qian, Hui Liu, Ming Zhang, Bao Shu, Longwei Xu, and Rufei Zhang. A geometry-based cycle slip detection and repair method with time-differenced carrier phase (TDCP) for a single frequency global position system (GPS)+ BeiDou navigation satellite system (BDS) receiver. Sensors, 16(12):2064, 2016.

[60] Todd Richert and Naser El-Sheimy. Optimal linear combinations of triple frequency carrier phase data from future global navigation satellite systems. GPS Solutions, 11(1):11–19, 2007.

[61] Charles Rino, Brian Breitsch, Yu Morton, and Charles Carrano. GNSS signal phase, TEC, and phase scintillation. to be submitted Navigation, J. Institute of Navigation, 2020.

[62] Charles Rino, Brian Breitsch, Yu Morton, Yu Jiao, Dongyang Xu, and Charles Carrano. A compact multi-frequency GNSS scintillation model. Navigation: Journal of The Institute of Navigation, 65(4):563–569, 2018.

[63] Charles Rino, Yu Morton, Brian Breitsch, and Charles Carrano. Stochastic TEC structure characterization. Journal of Geophysical Research: Space Physics, 2019.

[64] Carolyn Roesler, Y Wang, Y Jade Morton, and R Steve Nerem. Coherent GPS reflections over ocean surface. In IGARSS 2020-2020 IEEE International Geoscience and Remote Sensing Symposium, pages 6218–6221. IEEE, 2020.

[65] A Rovira-Garcia, JM Juan, J Sanz, and G González-Casado. The geodetic detrending technique: enabling high-accuracy navigation under scintillation. In 2020 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE), pages 1–4. IEEE, 2020.

[66] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. Physica D: nonlinear phenomena, 60(1-4):259–268, 1992.

[67] Dietrich Ryter and Heinrich Meyr. Theory of phase tracking systems of arbitrary order: Statistics of cycle slips and probability distribution of the state vector. IEEE Transactions on Information Theory, 24(1):1–7, 1978.

[68] Stefan Schaer, Gerhard Beutler, and Markus Rothacher. Mapping and predicting the ionosphere. In Proceedings of the 1998 IGS Analysis Center Workshop, Darmstadt, February 9-11, 1998, 1998.

[69] Ivan W Selesnick, Harry L Graber, Douglas S Pfeil, and Randall L Barbour. Simultaneous low-pass filtering and total variation denoising. IEEE Transactions on Signal Processing, 62(5):1109–1124, 2014.

[70] John Gordon Skellam. A probability distribution derived from the binomial distribution by regarding the probability of success as variable between the sets of trials. Journal of the Royal Statistical Society. Series B (Methodological), 10(2):257–261, 1948.

[71] Andrew M Smith, Cathryn N Mitchell, Robert J Watson, Robert W Meggs, Paul M Kintner, Kirsti Kauristie, and Farideh Honary. GPS scintillation in the high arctic associated with an auroral arc. Space Weather, 6(3), 2008.

[72] Nadezda Sokolova, Aiden Morrison, James T Curran, and Michele Bavaro. Phase scintillation decorrelation impact on multi-frequency users. In 2017 IEEE Aerospace Conference, pages 1–7. IEEE, 2017.

[73] V Sreeja, M Aquino, and ZG Elmas. Impact of ionospheric scintillation on GNSS receiver tracking performance over latin america: Introducing the concept of tracking jitter variance maps. Space Weather, 9(10), 2011.

[74] Noah Stephens-Davidowitz. On the Gaussian Measure Over Lattices. PhD thesis, New York University, 2017.

[75] Gilbert Strang. A proposal for toeplitz matrix calculations. Studies in Applied Mathematics, 74(2):171–176, 1986.

[76] Robert Tausworthe. Cycle slipping in phase-locked loops. IEEE Transactions on Communication Technology, 15(3):417–421, 1967.

[77] Peter Teunissen. The lambda method for the GNSS compass. Artificial Satellites, 41(3):89–103, 2006.

[78] Peter JG Teunissen. Least-squares estimation of the integer GPS ambiguities. In Invited lecture, section IV theory and methodology, IAG general meeting, Beijing, China, 1993.

[79] PJG Teunissen. The least-squares ambiguity decorrelation adjustment: a method for fast GPS integer ambiguity estimation. Journal of geodesy, 70:1–2, 1995.

[80] PJG Teunissen. Statistical GNSS carrier phase ambiguity resolution: a review. In Proceedings of the 11th IEEE Signal Processing Workshop on Statistical Signal Processing (Cat. No. 01TH8563), pages 4–12. IEEE, 2001.

[81] PJG Teunissen and PF De Bakker. Next generation GNSS single receiver cycle slip reliability. In VII Hotine-Marussi Symposium on Mathematical Geodesy, pages 159–164. Springer, 2012.

[82] AJ Van Dierendonck, John Klobuchar, and Quyen Hua. Ionospheric scintillation monitoring using commercial single frequency C/A code receivers. In Proc. of ION GPS, volume 93, pages 1333–1342, 1993.

[83] Jordi Vilà-Valls, Pau Closas, Carles Fernández-Prades, Jose A López-Salcedo, and Gonzalo Seco-Granados. Adaptive GNSS carrier tracking under ionospheric scintillation: Estimation vs. mitigation. IEEE Communications Letters, 19(6):961–964, 2015.

[84] Kan Wang and Markus Rothacher. Ambiguity resolution for triple-frequency geometry-free and ionosphere-free combination tested with real data. Journal of geodesy, 87(6):539–553, 2013.

[85] Yang Wang, Brian Breitsch, and Yu T Jade Morton. A state-based method to simultaneously reduce cycle slips and noise in coherent GNSS-R phase measurements from open-loop tracking. IEEE Transactions on Geoscience and Remote Sensing, 2020.

[86] Yang Wang, Rong Yang, and Y Jade Morton. Kalman filter-based robust closed-loop carrier tracking of airborne GNSS radio-occultation signals. IEEE Transactions on Aerospace and Electronic Systems, 56(5):3384–3393, 2020.

[87] Robert Ware and Frank Lad. Approximation of posterior means and variances of the digitised normal distribution using continuous normal approximation. Technical report, Citeseer, 2003.

[88] Albert D Wheelon. Electromagnetic Scintillation: Volume 2, Weak Scattering. Cambridge University Press, 2003.

[89] Christopher KI Williams and Carl Edward Rasmussen. Gaussian processes for machine learning, volume 2. MIT press Cambridge, MA, 2006.

[90] Guorui Xiao, Michael Mayer, Bernhard Heck, Lifen Sui, Tian Zeng, and Dongming Zhao. Improved time-differenced cycle slip detect and repair for GNSS undifferenced observations. GPS Solutions, 22(1):1–13, 2018.

[91] Dongyang Xu. GPS Equatorial Ionospheric Scintillation Signals Simulation, Characterization, and Estimation. PhD thesis, Colorado State University. Libraries, 2019.

[92] Dongyang Xu and Yu Morton. A semi-open loop GNSS carrier tracking algorithm for monitoring strong equatorial scintillation. IEEE Transactions on Aerospace and Electronic Systems, 54(2):722–738, 2017.

[93] Dongyang Xu, Yu T. Jade Morton, Charles L Rino, Charles S Carrano, and Yu Jiao. A two-parameter multifrequency GPS signal simulator for strong equatorial ionospheric scintillation: modeling and parameter characterization. Navigation, pages 1–15, 2020.

[94] Xiaofei Xu, Zhixi Nie, Zhenjie Wang, and Yuanfan Zhang. A modified TurboEdit cycle-slip detection and correction method for dual-frequency smartphone GNSS observation. Sensors, 20(20):5756, 2020.

[95] Rong Yang, Dongyang Xu, and Yu T Morton. Generalized multi-frequency GPS carrier tracking architecture: design and performance analysis. IEEE Transactions on Aerospace and Electronic Systems, 2019.

[96] Rong Yang, Xingqun Zhan, Yang Wang, Yu Morton, and Jennifer S Haase. Hybrid carrier tracking and position determination using the low elevation satellite signals. In Proceedings of the 2020 International Technical Meeting of The Institute of Navigation, pages 289–299, 2020.

[97] Nikolay N Zernov, Vadim E Gherm, and Hal J Strangeways. Further determinations of strong scintillation effects on GNSS signals using the hybrid scintillation propagation model. Radio Science, 47(04):1–6, 2012.

[98] Guochao Zhang, Qingming Gui, Songhui Han, Jun Zhao, and Wenhua Huang. A bayesian method of GNSS cycle slips detection based on arma model. In 2017 Forum on Cooperative Positioning and Service (CPGPS), pages 219–222. IEEE, 2017.

[99] Xiaohong Zhang and Pan Li. Benefits of the third frequency signal on cycle slip correction. GPS solutions, 20(3):451–460, 2016.

[100] Yunong Zhang and William E Leithead. Approximate implementation of the logarithm of the matrix determinant in gaussian process regression. Journal of Statistical Computation and Simulation, 77(4):329–348, 2007.

[101] Dongsheng Zhao, Craig M Hancock, Gethin Wyn Roberts, and Shuanggen Jin. Cycle slip detection during high ionospheric activities based on combined triple-frequency GNSS signals. Remote sensing, 11(3):250, 2019.

[102] Dongsheng Zhao, Gethin Wyn Roberts, Craig M Hancock, Lawrence Lau, and Ruibin Bai. A triple-frequency cycle slip detection and correction method based on modified HMW combinations applied on GPS and BDS. GPS solutions, 23(1):22, 2019.

[103] Qile Zhao, Binzi Sun, Zhiqiang Dai, Zhigang Hu, Chuang Shi, and Jingnan Liu. Real-time detection and repair of cycle slips in triple-frequency GNSS measurements. GPS Solutions, 19(3):381–391, 2015.

Appendix A

Linear Combinations

Here we derive the analytical expressions for orthogonal IF, GF, and GIF combinations of triple-frequency carrier phase measurements. The derivation stem from work originally done in [13]. Recall that GF and IF combinations satisfy:

$$\sum_k c_k = 0 \quad \text{GF} \tag{A.1}$$

$$\sum_k \beta_k c_k = 0 \quad \text{IF} \tag{A.2}$$

Then the coefficients satisfying GF and IF constraints for triple-frequency can be found by taking the cross product $\mathbf{1} \times \beta$, where $\mathbf{1} = [1, 1, 1]$ and $\beta = [\beta_1, \beta_2, \beta_3]$:

$$\mathbf{c}_{\text{GIF}} = \begin{bmatrix} \beta_2 - \beta_3 \\ \beta_3 - \beta_1 \\ \beta_1 - \beta_2 \end{bmatrix} \tag{A.3}$$

Note, $\mathbf{c}_{\text{GIF}}$ is defined by a linear subspace of $\mathbb{R}^3$ and so its GIF properties are preserved under scalar multiplication. The IF and GF coefficients that are orthogonal to these GIF coefficients can then be found by taking $\mathbf{1} \times \mathbf{c}_{\mathrm{GIF}}$ and $\beta \times \mathbf{c}_{\mathrm{GIF}}$:

$$\mathbf{c}_{\mathrm{IF}} = \begin{bmatrix} \beta_2(\beta_1 - \beta_2) - \beta_3(\beta_3 - \beta_1) \\ \beta_3(\beta_2 - \beta_3) - \beta_1(\beta_1 - \beta_2) \\ \beta_1(\beta_3 - \beta_1) - \beta_2(\beta_2 - \beta_3) \end{bmatrix} \tag{A.4}$$

$$\mathbf{c}_{\mathrm{GF}} = \begin{bmatrix} 2\beta_1 - \beta_2 - \beta_3 \\ 2\beta_2 - \beta_3 - \beta_1 \\ 2\beta_3 - \beta_1 - \beta_2 \end{bmatrix} \tag{A.5}$$

$$\tag{A.6}$$

Again, the properties of these coefficients are invariant under scalar multiplication.

Appendix B

Hyperparameter Estimation

We want to find a set of values for the covariance hyperparameters σ_ϵ^2, $\sigma_\mathcal{I}^2$, and $\rho_\mathcal{I}$ that generalize a large set of simulated scintillation measurements. To do this, we optimize the average log-likelihood of measurements over a set of simulations. In particular, we perform 50 5-minute simulations of scintillation measurements. From each set of simulated measurements we subtract off the truth reference of slip occurrences to obtain measurements without cycle slips. We consider a grid of parameter values for each of σ_ϵ^2, $\sigma_\mathcal{I}^2$, and $\rho_\mathcal{I}$ and evaluate the log of Equation 4.19 for a model $\mathbf{Q_y}$:

$$2 \log p(\mathbf{y}|\Delta\mathbf{z}_{\text{true}}) = -\log|\mathbf{Q_y}| - (\mathbf{y} - \mathbf{B}\Delta\mathbf{z}_{\text{true}})^T \mathbf{Q_y}^{-1} (\mathbf{y} - \mathbf{B}\Delta\mathbf{z}_{\text{true}}) + C \tag{B.1}$$

Here, C is constant with respect to the hyperparameter values. Computing the log-determinant of $\mathbf{Q_y}$ is non-trivial, and we apply the approximation algorithm from [100]. We choose the values that minimize the average log-likelihood over all 50 simulations. Figure B.1 shows a slice of the optimization landscape corresponding to σ_ϵ and $\rho_\mathcal{I}$ for Window 1. The location corresponding to our chosen hyperparameter values are indicated with a white dot. The mostly uni-modal optimization landscape indicates that the chosen parameters should be valid for performing GP regression on this type of data.

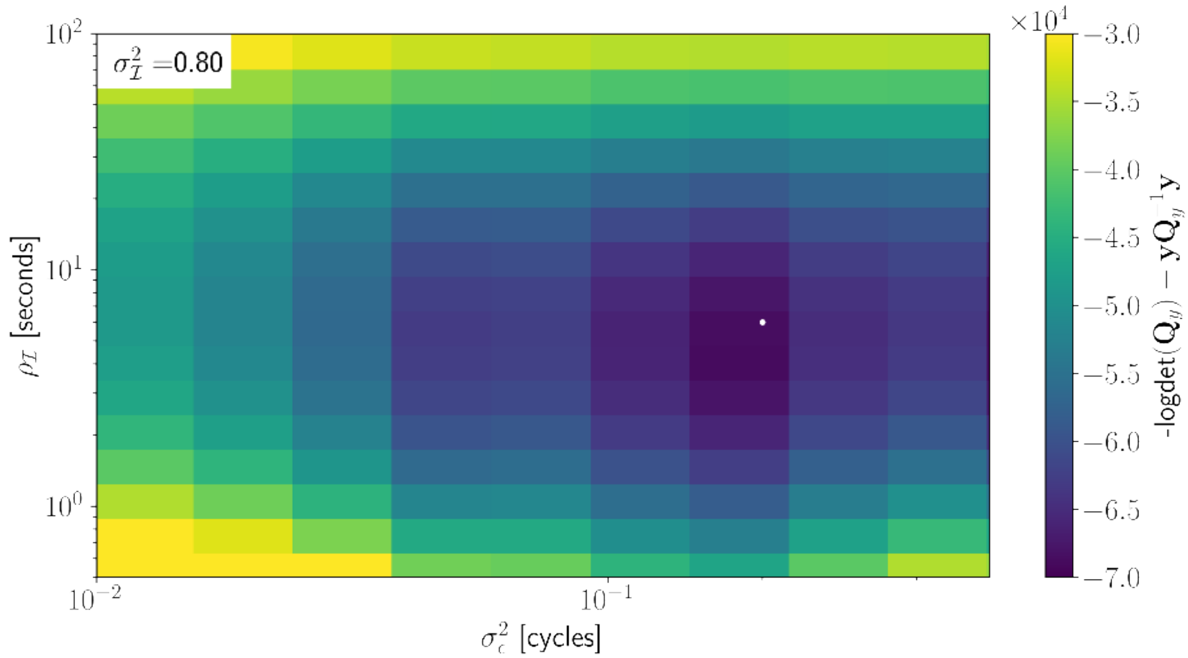

Figure B.1: Grid of average log-likelihood for corrected simulated scintillation measurements corresponding to Window 1 parameters. The white dot indicates the location for parameter values that minimize this average log-likelihood.

Appendix C

Algebra for MAP and ILS Equivalence

In Chapter 4, we show the equivalence between MAP and ILS estimates assuming a uniform (non-informative) prior for $\Delta\mathbf{z}$. To show this, we used the following proportionality with respect to argument $\Delta\mathbf{z}_i$:

$$\mathcal{N}(\mathbf{y}; \mathbf{B}\Delta\mathbf{z}_i, \mathbf{Q_y}) \propto \mathcal{N}(\Delta\mathbf{z}_i; \mu_{\Delta\hat{\mathbf{z}}}, \mathbf{Q}_{\Delta\hat{\mathbf{z}}}) \tag{C.1}$$

For the sake of completeness, here we derive this proportionality in more detail. First we recall the definitions for $\mathbf{Q}_{\Delta\hat{\mathbf{z}}}$ and $\mu_{\Delta\hat{\mathbf{z}}}$, which are provided in Equations 4.31 and 4.32 and which we provide here for convenience:

$$\mathbf{Q}_{\Delta\hat{\mathbf{z}}} = \left(\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{B}\right)^{-1} \tag{C.2}$$

$$\mu_{\Delta\hat{\mathbf{z}}} = \mathbf{Q}_{\Delta\hat{\mathbf{z}}}\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{y} \tag{C.3}$$

Now consider the exponent for the expression $\mathcal{N}(\mathbf{y}; \mathbf{B}\Delta\mathbf{z}_i, \mathbf{Q_y})$:

$$\left(\mathbf{y} - \mathbf{B}\Delta\mathbf{z}_i\right)^T \mathbf{Q_y}^{-1} \left(\mathbf{y} - \mathbf{B}\Delta\mathbf{z}_i\right) \tag{C.4}$$

$$= \mathbf{y}^T\mathbf{Q_y}^{-1}\mathbf{y} - 2\Delta\mathbf{z}_i^T\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{y} + \Delta\mathbf{z}_i^T\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{B}\Delta\mathbf{z}_i \tag{C.5}$$

$$= \left(\Delta\mathbf{z}_i - \left(\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{B}\right)^{-1}\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{y}\right)^T \left(\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{B}\right) \left(\Delta\mathbf{z}_i - \left(\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{B}\right)^{-1}\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{y}\right) \tag{C.6}$$

$$= \left(\Delta\mathbf{z}_i - \mathbf{Q}_{\Delta\hat{\mathbf{z}}}\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{y}\right)^T \mathbf{Q}_{\Delta\hat{\mathbf{z}}}^{-1} \left(\Delta\mathbf{z}_i - \mathbf{Q}_{\Delta\hat{\mathbf{z}}}\mathbf{B}^T\mathbf{Q_y}^{-1}\mathbf{y}\right) \tag{C.7}$$

$$= \left(\Delta\mathbf{z}_i - \mu_{\Delta\hat{\mathbf{z}}}\right)^T \mathbf{Q}_{\Delta\hat{\mathbf{z}}}^{-1} \left(\Delta\mathbf{z}_i - \mu_{\Delta\hat{\mathbf{z}}}\right) \tag{C.8}$$

It follows that the exponents of the two densities are equal when evaluated at appropriate inputs. Therefore the terms in Equation C.1 are proportional.