

Adaptive Control and Data Assimilation for Nonlinear Systems

by

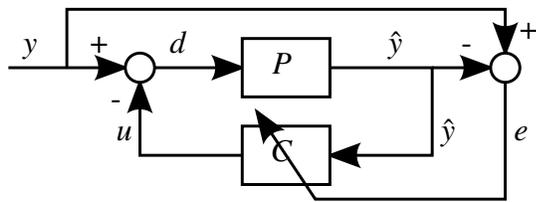
Alexey V. Morozov

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Aerospace Engineering)
in the University of Michigan
2013

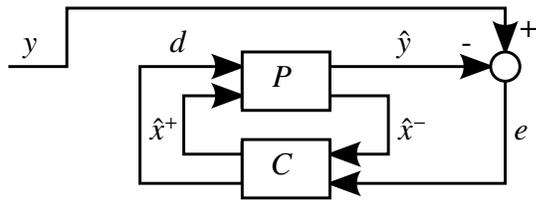
Doctoral Committee:

Professor Dennis S. Bernstein, Chair
Professor Pierre T. Kabamba
Professor Ilya V. Kolmanovsky
Associate Professor Aaron J. Ridley

(A) adaptive control for comand following:



(B) data assimilation with driver estimation:



This dissertation in a nutshell: This dissertation draws a connection between two mathematical problems: adaptive control and data assimilation. The problem of adaptive control seeks to find controller C would make the output \hat{y} of plant P follow the reference input y . Similarly, the problem of data assimilation seeks to find states \hat{x} and drivers d to drive error $e = y - \hat{y}$ to zero. Chapters 2-4 describe the former problem, whereas chapters 5 and 6 describe the latter.

©Alexey V. Morozov

2013

Dedication

To Mom, Dad, and Lera, who gave me hope every time I was about to give up and believed in me even when I did not.

Acknowledgments

I would like to thank Professor Bernstein for guiding my research and tirelessly correcting my papers, Professor Ridley for his guidance on coding, research, and academia in general, Professors Girard, Kabamba, and Kolmanovsky for providing invaluable comments on this dissertation, Nancy Collins, Timothy Hoar, and Jeffrey Anderson for their tremendous help in writing the DART-GITM interface, Denise Phelps for guiding me through the entirety of my PhD process, Astrik Tenney, Sharon Shoham, Dana O’Neill, Emily Froimson, Jack Kent Cooke, Michelle McLeod, Camie Wood, and many more Cookies at the Jack Kent Cooke Foundation for giving me a chance to pursue my dreams, Jesse Hoagg and Tony D’Amato for helpful conversations about adaptive and internal model control, James Coughlin, Asad Ali, Gerardo Cruz, Kenny Mitchell, Amy Wu, Matthew Robertson, Aisyah Wan, Laura Haskins, Rawan Sobuh, Ruben Vandenheede, Yusuf Rida, Zach Tubbs, Derrick Cheung, Timothy Nash, Vince Marshall, Uroš Kalabić, Derek Dalle, and many other teachers, friends and family for teaching me so much and being awesome human beings!

TABLE OF CONTENTS

Dedication	ii
Acknowledgments	iii
List of Figures	vi
Abstract	xiii
 Chapter	
1 Introduction	1
2 Retrospective Cost Adaptive Control for Nonminimum-Phase Systems with Uncertain Nonminimum-Phase Zeros Using Convex Optimization	10
2.1 Introduction	10
2.2 Problem Formulation	13
2.3 Review of RCAC	15
2.4 Constrained Convex Optimization	18
2.5 Numerical Examples	22
2.5.1 Third-Order Plant with a Known NMP Zero	23
2.5.2 Third-Order Plant with an Uncertain NMP Zero	23
2.5.3 Effect of Varying γ	26
2.5.4 Fourth-Order Plants with Uncertain NMP Zeros	27
2.6 Conclusion	28
3 A Computational Study of the Performance and Robustness Properties of Retrospective Cost Adaptive Control	32
3.1 Introduction	32
3.2 Problem Formulation	34
3.3 Review of the Cumulative Retrospective Cost Adaptive Controller	36
3.4 Sensitivity to the Accuracy of the Required Model Information	39
3.4.1 Sensitivity to the Accuracy of the Nonminimum-Phase Zero Es- timate	39
3.4.2 Sensitivity to the Accuracy of the Estimate of the First Nonzero Markov Parameter	48
3.5 Sampled Data Systems	49
3.6 Conclusion	51

4 Retrospective Adaptive Control of a Planar Multilink Arm with Nonminimum-Phase Zeros	54
4.1 Introduction	54
4.2 Nonlinear Equations of Motion	56
4.3 Linearized Equations of Motion	61
4.4 Nonminimum-Phase Zero for the N -Link System	64
4.5 Numerical Examples	66
4.6 Conclusion	68
5 Assimilating Neutral Density Measurements into the Global Ionosphere-Thermosphere Model using the Ensemble Adjustment Kalman Filter	71
5.1 Introduction	71
5.2 Methodology	74
5.2.1 Global Ionosphere-Thermosphere Model	74
5.2.2 EAKF: Data Assimilation without Driver Estimation	76
5.2.3 EAKF: Data Assimilation with Driver Estimation	79
5.2.4 EAKF: Data Assimilation with Driver Estimation, Ensemble Inflation and Model Uncertainty	82
5.2.5 EAKF: Localization	83
5.3 Results	85
5.3.1 Simulated data above the subsolar point	87
5.3.2 Simulated CHAMP data	88
5.3.3 Real CHAMP data	89
5.4 Conclusions	90
6 Assimilating Total Electron Content Measurements into GITM using EAKF	95
6.1 Introduction	96
6.2 Methodology	97
6.2.1 Total Electron Content Measurements	97
6.2.2 DART-GITM Interface Experimental Setup	99
6.3 Results	100
6.3.1 Assimilating simulated data above North America and Europe without estimating $F_{10.7}$	100
6.3.2 Assimilating simulated data above North America and Europe with estimating $F_{10.7}$	102
6.3.3 Assimilating real GPS data without estimating $F_{10.7}$	103
6.3.4 Assimilating real GPS data with estimating $F_{10.7}$	103
6.4 Conclusions	104
7 Conclusions	108
Bibliography	111

LIST OF FIGURES

1.1	Feedback architecture for the command following problem. The goal is to find controller C , which would make plant output \hat{y} track the desired trajectory y . Describing the transfer function from y to \hat{y} in the frequency domain demonstrates that this goal is achieved when $C = 1 - P^{-1}$	1
2.1	Spectral radius $\rho(K)$ is plotted versus maximum singular value $\sigma_{\max}(K)$ for 10^5 10^{th} order random monic polynomials. The plot shows that the constraint (2.21) is conservative, since a polynomial of high enough order can be found with companion maximum singular value arbitrarily far to the right of the $y = x$ line. This fact is further demonstrated in Figure 2.3.	20
2.2	Roots of the selected second order polynomials are shown in (a) and the singular values of the corresponding (correspondence is demonstrated by color) companion matrix (K) are shown versus the polynomial roots in (b). Note that the axes in (b) are swapped as compared to Figure 2.1. (b) demonstrates that in the worst case scenario, the singular value of a companion matrix K is 2.414 times greater than the absolute value of the root of the polynomial (1.0).	21
2.3	Maximum singular value for n^{th} order monic polynomial with n roots at 1. The conservatism of the maximum singular value increases as the order of the matrix increases.	22
2.4	RCAC performance, control, and controller coefficients are shown as a function of time for the transfer function (2.24).	24
2.5	Evolution of the RCAC controller poles is shown as a function of time in terms of color. Later pole locations are colored with warmer colors and the final locations are marked with a red square. After the controller is turned on, the poles settle in about 400 steps.	25
2.6	CC-RCAC performance, control, and controller coefficients are shown as a function of time for the transfer function (2.24).	26
2.7	Evolution of the CC-RCAC controller poles is shown as a function of time in terms of color. After the controller is turned on, the poles settle in about 200 steps.	27
2.8	For the third-order plant (2.24) with a nonminimum-phase zero whose position is given by the horizontal axis, RCAC uses the estimate of the nonminimum-phase zero whose location is given by the vertical axis. The color in (a) shows the logarithm of the transient performance, whereas (b) shows the logarithm of the steady-state performance.	28

2.9	For the third-order plant with a nonminimum-phase zero whose position is given by the horizontal axis, CC-RCAC with $\gamma = 1.1$ uses the estimate of the nonminimum-phase zero whose location is given by the vertical axis.	29
2.10	For the $G_3(z)$ with a nonminimum-phase zero whose position is given by the horizontal axis, CC-RCAC uses the exact location of the nonminimum-phase zero and a value of γ that varies from 1.1 to 4.1. For 20 values of γ , the best performance is achieved with $\gamma = 1.1$	30
2.11	For 20 random fourth-order plants with a nonminimum-phase zero at 2, RCAC uses an estimate of the nonminimum-phase zero whose location is given by the horizontal axis.	30
2.12	For 20 random fourth-order plants, CC-RCAC with $\gamma = 1.1$ uses an estimate of the nonminimum-phase zero whose location is given by the horizontal axis.	31
3.1	For 40 third order plants with a nonminimum-phase zero whose position is given by the x-axis, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1.1 to 8.1. The color in (a) shows logarithm of the transient performance, whereas in (b) the logarithm of the steady state performance is shown. For these 40 plants, the stability of the closed-loop system is less sensitive to overestimating the location of the nonminimum-phase zero than it is to underestimating it. In addition, the plants with NMP zeros further out on the real axis have greater stability margins than those with NMP zeros closer to 1.	41
3.2	For 50 randomly generated plants, all of which have a nonminimum-phase zero at 2, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1.4 to 3.8. This histogram shows the percent of the 50 plants that have unstable closed-loop responses for different estimates of the nonminimum-phase zero. For these 50 random plants, the stability of the closed-loop system is less sensitive to overestimating the location of the nonminimum-phase zero than it is to underestimating the location of the nonminimum-phase zero.	42
3.3	For 50 randomly generated plants with fixed order and relative degree (per subplot), all of which have a nonminimum-phase zero at 2, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1.4 to 2.6. These histograms show the percent of the 50 plants that have unstable closed-loop responses for different estimates of the nonminimum-phase zero. The robustness to the estimate of the nonminimum-phase zero decreases in rows from right to left (with increasing order), and in columns from top to bottom (with decreasing relative degree).	43
3.4	For 50 randomly generated plants, all of which have a nonminimum-phase zero at 2, the adaptive controller uses an estimate of the nonminimum-phase zero that varies from 1.4 to 10. For all 50 plants, there are finite bounds on the required accuracy of the nonminimum-phase zero.	44

3.5	For 61 second order plants with double poles at locations indicated by the x-axis, all of which have a nonminimum-phase zero at 2, the adaptive controller uses an estimate of the nonminimum-phase zero that varies from 1 to 12 (y-axis). The color map represents the logarithm of transient performance and logarithm of steady state performance in subfigures (a) and (b), respectively. The adaptive controller (3.11)-(3.13) is more robust for plants with positive double-poles.	45
3.6	Downward stability margin. For 2601 second order plants with a nonminimum-phase zero at 2, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1.6 to 3.6. The colored regions represent plant-pole areas where the closed loop response was unstable at the estimate of the nonminimum-phase zero shown in the legend. Due to the nature of the complex variables this plot is symmetric, but only the top half is shown. The plants with poles whose real part is close to -1 are less robust to the underestimation of the nonminimum phase zero than the rest of the plants with poles inside the unit circle.	46
3.7	Upward stability margin. For 2601 second order plants with a nonminimum-phase zero at 2, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1.6 to 3.6. The colored regions represent plant-pole areas where the closed loop response was unstable at the estimate of the nonminimum-phase zero shown in the legend. Due to the nature of the complex variables this plot is symmetric, but only the top half is shown. The plants which are least robust to the overestimation of the nonminimum phase zero are the ones with poles close to ± 1	47
3.8	For 50 randomly generated plants, all of which have a nonminimum-phase zeros at 2 and 3, the adaptive controller (3.11)-(3.13) uses estimates of the nonminimum-phase zeros which vary from 1 to 4 each. The steep slope on the downward margin and shallower slope on the upward margin suggest that the performance of the algorithm in 2-nonminimum-phase-zeros case is similar to the single nonminimum-phase zero case, but this conjecture is further investigated in the following plots.	48
3.9	For 50 randomly generated plants, all of which have a nonminimum-phase zeros at 2 and 3, the adaptive controller (3.11)-(3.13) uses an estimate of the second nonminimum-phase zero that varies from 1 to 4. The estimate of the first nonminimum-phase zero is fixed at 3. Just like in the single nonminimum-phase zero case, the algorithm is more robust to overestimating the location of the nonminimum-phase zero, than to underestimating it.	49
3.10	For 100 randomly generated plants, all of which have a nonminimum-phase zeros at $2 \pm 1j$, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1 to $4 \pm 3j$. The color corresponds to the percent of plants that have unstable closed loop performance.	50

3.11	For 100 randomly generated plants, all of which have a nonminimum-phase zeros at $2 \pm 1j$, the adaptive controller (3.11)-(3.13) uses an estimate of the real part of the nonminimum-phase zero in (a), imaginary part in (b), magnitude in (c) and phase in (d). The algorithm seems to be equally robustness to changes in estimates of the real and imaginary parts as well as magnitude. Additionally, in this particular setup, the algorithm is more robust to underestimating phase than overestimating it.	51
3.12	For 100 randomly generated stable plants, the adaptive controller (3.11)-(3.13) uses the scaled value of the first nonzero Markov parameter (given by the x-axis) and exact values for the nonminimum-phase zeros. The color scale in (a) corresponds to logarithm of the transient performance and in (b) it corresponds to logarithm of steady state performance. The lesser slope of the histogram on the right than on the left suggest that the algorithm is more sensitive to underestimating the first nonzero Markov parameter than to overestimating it.	52
3.13	For 100 randomly generated plants, the adaptive controller (3.11)-(3.13) uses the exact value of the first nonzero Markov parameter and nonminimum-phase zeros. The color scale in (a) corresponds to logarithm of the transient performance and in (b) it corresponds to logarithm of steady state performance, as established in Section 3.4. The algorithm performs best at the sampling period corresponding to the minimum sampling frequency, as defined in section 3.5.	53
4.1	N-link planar arm system. All motion is in the horizontal plane.	57
4.2	<i>Command following for the linearized two-link system:</i> The adaptive control (3.11)-(3.13) uses knowledge of the nonminimum-phase zero to track a sinusoid with unknown frequency and amplitude. The adaptive control is turned on after 7.5 seconds and drives the performance to zero.	67
4.3	<i>Command following for the nonlinear two-link system:</i> The adaptive control (3.11)-(3.13) uses knowledge of the linearized nonminimum-phase zero to track a sinusoid with unknown frequency and amplitude. The adaptive control is turned on after 7.5 seconds and drives the performance to zero. The performance with the nonlinear system is comparable to the linear case shown in Figure 4.2.	68
4.4	<i>Disturbance rejection for the linearized two-link system:</i> The adaptive control (3.11)-(3.13) uses knowledge of the nonminimum-phase zero to reject an unknown sinusoidal disturbance acting on both joints of the two-link mechanism. The adaptive control is turned on after 7.5 seconds and drives the performance to zero.	69
4.5	<i>Disturbance rejection for the nonlinear two-link system:</i> The adaptive control (3.11)-(3.13) uses knowledge of the nonminimum-phase zero to reject an unknown sinusoidal disturbance acting on both joints of the two-link mechanism. The adaptive control is turned on after 7.5 seconds and drives the performance to zero. The performance with the nonlinear system is comparable to the linear case shown in Figure 4.4.	70

5.1	(a) shows the neutral density at about 400 km altitude at 02:32UT 12/1/2002 as well as CHAMP and GRACE orbital trajectories. CHAMP (red) and GRACE (yellow) orbits for the chosen time period do not pass near the subsolar point. The GITM gridpoint closest to Ann Arbor, MI (green) is referred to in subsequent figures. (b) demonstrates the model-data mismatch between no-DA GITM output (blue) and CHAMP measurement (red). GITM without data assimilation underestimates neutral density compared to CHAMP measurements.	75
5.2	Time evolution of the EAKF state estimate for the system (5.1), (5.2) without driver estimation. (a) shows the true state (black dashed), the state measurement (red), the state estimate initial distribution (green), and the state estimate mean (dark blue) and spread (light blue, standard deviation). (b) shows the true driver, driver estimate distribution, mean, and variance in the same colors as (a). The posterior ensemble mean converges to the true state, whereas the prior mean deviates during the model propagation step due to the incorrect driver estimate.	79
5.3	Time evolution of EAKF state and driver estimates for the system (5.1), (5.2) using driver estimation, but without ensemble inflation. Due to ensemble model assuming no dynamics for the driver, ensemble spread (light blue) approaches zero. Accordingly, the mean estimates converge to fixed values and do not track the true state and driver.	81
5.4	Time evolution of EAKF state and driver estimates for the system (5.1), (5.2) when driver estimation and ensemble inflation are used. Using variance inflation of $\lambda = 2.0$ results in an average decrease of 17% in the state tracking error.	81
5.5	Time evolution of EAKF state and driver estimates for the system (5.1), (5.2) using the ensemble model (5.16) where driver estimation and ensemble inflation are used. With ensemble inflated using $\lambda = 2.0$, the driver estimate compensates for model uncertainty and allows minor degradation in state estimate (1%) as compared to the perfect model case.	83
5.6	The localization region for the measurement located above Ann Arbor, MI. Markers are placed at GITM cell center locations. The dark red markers correspond to measurements having a direct effect on the state variable, whereas transparent markers correspond to measurements that have no effect on the state variable. The state variables affected by the measurement above Ann Arbor lie in a projection of a small circle region centered in Ann Arbor.	84
5.7	Variance of the sample variance as a function of the number of ensemble members when $\sigma = 1$. As the number of ensembles members increases, the accuracy of variance estimate increases.	86
5.8	Time evolution of the simulated (red) and estimated (blue) mass density above the subsolar point (measurement location) in (a) and at the 400 km gridpoint closest to Ann Arbor, MI in (b). Estimated mass density above the subsolar point converged to the true value within 3 hours, whereas density at the Ann Arbor gridpoint needed 12 hours to converge.	88

5.9	(a) shows the time evolution of mass density at CHAMP locations demonstrates that ρ converged to the true value within 9 hours when measurements were taken at the subsolar point. (b) shows mass density estimate at GRACE locations and demonstrates the same rate of convergence. (c) and (d) reinterpret (a) and (b) by plotting the orbit averages of the errors between simulated and estimated data. This figure demonstrates that estimated mass density converged to true mass density at locations other than the measurement location (subsolar point).	89
5.10	Time evolution of $F_{10.7}$ for the truth simulation (dashed black), nominal no-DA simulation (solid black), and the ensemble mean (dark blue). The ensemble mean oscillated around the true value and converged to the true value by the end of the simulation.	90
5.11	Orbit averaged mass density absolute percentage errors at (a) CHAMP and (b) GRACE locations. Simulated CHAMP data were used as measurements in the EAKF, whereas the simulated GRACE values were only used for validation. Density estimate at CHAMP locations dropped immediately to about 3%, whereas the estimate at GRACE locations decreased only after about 9 hours due to altitude difference between these satellites. The bias in densities at both locations was decreased and the final error level was lower at the measurement location (CHAMP) than at the validation location (GRACE).	91
5.12	Time evolution of $F_{10.7}$ for the true, nominal and ensemble simulations. The ensemble mean estimate initially overestimates the true value, but then converges within 15 hours.	92
5.13	Orbit averaged mass density measurements at CHAMP and GRACE are shown in (a) and (b) respectively, and absolute percentage errors are shown in (c) and (d) respectively. Real CHAMP data were used as measurements in the EAKF, whereas GRACE data were used only for validation. The bias in density at CHAMP locations was decreased over the second day from 57% to 7%, whereas the bias at GRACE location was only decreased by a smaller amount (from 70% to 52%, see text for relevant discussion).	93
5.14	Time evolution of $F_{10.7}$ for the true, nominal and ensemble simulations. $F_{10.7}$ estimate took on higher values than what was measured by NOAA, in order to drive GITM closer to the CHAMP data.	94
6.1	VTEC snapshot at 00:30 UT on Dec. 1 2002. VTEC data cover the whole globe, but are more sparse above oceans than above land.	98
6.2	VTEC cropped simulated data at 00:30 UT on Dec. 1 2002. The cropping region locations were chosen to form rectangular patches over Europe and North America.	98
6.3	Initial distribution and time evolution of $F_{10.7}$ estimates. The $F_{10.7}$ estimates were initialized as $\mathcal{N}(130 \text{ SFU}, 5 \text{ SFU})$ and were kept constant for the entire simulation.	100

6.4	Time evolution of electron density at the (40.0°N, 80.0°W, 411 km) gridpoint. Discontinuities in the blue line at each step are a result of plotting prior and posterior density estimate at each step. The slope of blue line between assimilation steps is indicative of GITM influence on the estimates during the model-propagation periods.. Ensemble mean and spread increase after 12UT (07LT) and reach their maximum values around 19UT (14UT).	101
6.5	VTEC error snapshot at 22:30 UT on Dec. 1 2002. The error is smallest where data is available (note the darker area above the United States).	101
6.6	Time evolution of VTEC error average over the whole globe. Average error does not surpass 2 TECU.	102
6.7	Assimilating simulated data with driver estimation. (a) shows the time evolution of $F_{10.7}$ estimate, (b) shows the time evolution of electron density at the (40.0°N, 80.0°W, 411 km) gridpoint, (c) shows the VTEC error snapshot at 22:30 UT on Dec. 1 2002, and (d) shows time evolution of average VTEC error. The $F_{10.7}$ estimate did not converge to the true value, but overall error levels were still comparable to the results of Section 6.3.2.	105
6.8	Assimilating real data without driver estimation. For plot-type descriptions see Figure 6.7. (c) confirms the rich/poor data region trend by showing low error values above the US and Europe and high values above South Pacific. Without driver estimation, error leveled out at about 8 TECU.	106
6.9	Assimilating real data with driver estimation. For plot-type descriptions see Figure 6.7. Adding driver estimation had an overall stabilizing effect to the data assimilation process, as can be seen in fewer upward spikes in (b), smaller errors over South Pacific Ocean in (c), and lower and flatter average error response in (d) when compared to not using driver estimation.	107

ABSTRACT

This dissertation draws a connection between the problem of adaptive command following and data assimilation with driver estimation. For these two different problems, two different approaches are adapted - retrospective cost adaptive control (RCAC) and ensemble adjustment Kalman filter (EAKF). Both methods are first applied to linear minimum phase examples to demonstrate the modeling requirements and operating principles. Next, RCAC robustness to the accuracy of nonminimum phase (NMP) zero is investigated for plants of varying order and relative degree and is then improved through a convex constraint on the controller poles. It is demonstrated that the convex constraint significantly improves the transient and steady state command following and disturbance rejection performance. RCAC is then shown to successfully solve the command following problem for a nonlinear system (planar multilink arm) in a particular range of command amplitudes.

This nonlinear example serves as a nice introduction to the data assimilation problem for a particular space weather model, which can be thought of as a generalization of the command following problem. The command in this case comes from real satellite data and the goal is to make a space weather model track the conditions observed by the satellites. The particular model used in this project is the global ionosphere-thermosphere model (GITM). GITM's states (densities, temperatures and velocities for multiple species over a three-dimensional grid) and inputs (solar flux index F10.7) are estimated using the

EAKF. The assimilation technique is first demonstrated on the simulated data and is then applied to real satellite data. Successful assimilation of neutral density measurements from the CHAMP satellite as well as total electron content measurements from the GPS satellites is demonstrated.

CHAPTER 1

Introduction

We all seek to achieve something. Pilots seek to maintain the heading and altitude of their planes, weathermen seek to improve the predictive capabilities of their models, and car drivers aim to keep their cars in the proper lane. Operators of these systems seek to find a control command, which when applied to the plant (the system to be controlled) would achieve the desired trajectory. Knowledge of the plant is beneficial for trajectory planning because it allows the operator to anticipate the reaction of the plant to any given command. For example, a car driver might anticipate certain amount of side-slip through an icy turn and pre-compensate for it by cutting to the inside of the curve. In a way, these anticipatory precautions are a form of model inversion, that is designing the input signal in such a way as to cancel the effects of the plant. This concept is better represented graphically as the feedback architecture shown in Figure 1.1.

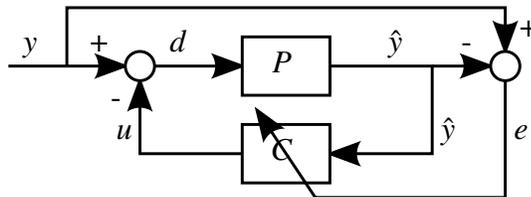


Figure 1.1: Feedback architecture for the command following problem. The goal is to find controller C , which would make plant output \hat{y} track the desired trajectory y . Describing the transfer function from y to \hat{y} in the frequency domain demonstrates that this goal is achieved when $C = 1 - P^{-1}$.

The ideal controller for this architecture can be derived by recalling the goal of this experiment (drive \hat{y} to y), that is $\hat{y} = y$. Then it can be noted that \hat{y} is the result of applying d to the plant P , which in frequency domain can be denoted by $\hat{y} = Pd$, where P is a transfer function from d to \hat{y} either in discrete or continuous time. For simplicity of this introduction, we consider only single input single output (SISO) plants and controllers. Accordingly, $d = y - u$, and $u = C\hat{y}$, which together with the plant equation yields $\hat{y} = \frac{P}{1+PC}y$. To satisfy the original goal of $\hat{y} = y$ it is necessary that $\frac{P}{1+PC} = 1$, that is $C = 1 - P^{-1}$. In this case, the “ideal controller” transfer function must equal to one minus the inverse of the plant transfer function.

However, this kind of direct inversion is often impractical, as it requires perfect knowledge of the plant, might result in non-causal controllers (controllers that require knowledge of future commands or use exact derivatives), or unstable controllers (controllers producing unbounded output u when input y is bounded).

Early work on model inversion was performed by Brockett and Mesarovic [1], who described necessary and sufficient conditions for the existence of the inverse system for SISO linear and nonlinear time-invariant systems and introduced the concept of functional reproducibility (ability to produce particular outputs, similar to target path controllability [2]). This work inspired Silverman [3], Dorato [4], and Sain and Massey [5] to simplify these necessary and sufficient conditions and design more efficient algorithms for finding the inverse system. In particular, Silverman extended the results of Brockett and Mesarovic to time-varying linear systems with multiple inputs and multiple outputs (MIMO). On the other hand, Sain and Massey extended the findings of [1] to discrete time systems and introduced the concept of k -integral inverse, where k refers to the minimum number of integrations in the plant that no controller can remove without using exact derivatives. This work was continued by Willsky [6], who produced stronger versions of the necessary and sufficient conditions; Orner, Emre and Huseyin [7, 8], who demonstrated how transfer function coefficients can be used to check existence of a left k -integral inverse without

converting the model to state space; Hirschorn [9], who generalized the findings to the nonlinear systems; and Isidori and Byrnes [10], who showed how nonlinear systems can be controlled by using input-output feedback linearization.

The approaches mentioned above have the restricting assumption of a minimum phase plant, that is plant has no zeros in the open right half plane for linear systems and stable zero dynamics for nonlinear systems (for the definition of zero dynamics, see [11]). This limitation is crucial for many control applications and several solutions to this problem have been proposed. Devasia, Chen, and Paden [12] demonstrated a stable but non-causal inversion-based approach. This approach outperformed the Byrnes-Isidori regulator in terms of transient response when preview was available, but resulted in similar performance otherwise. Huang [13] solved the tracking problem for a class of nonlinear systems by decomposing the zero dynamics into stable and unstable parts. Similarly, Doyle, Allgower, and Morari [14] described a method that decomposes the nonlinear plant into minimum and nonminimum phase (NMP) parts, the latter one of which is approximated by a minimum phase system. This minimum phase approximation is then used in the controller design, which is then applied to the original nonlinear system. The main difference between this approach and that of [10] is the use of control derivatives (\dot{u} , \ddot{u} , etc.) to transform the state variables into a canonical form. The new expression for the states is then modified in a fashion similar to the reflection of unstable zeros for linear plants [15] to approximate the NMP part of the plant by a minimum-phase system. The provided examples demonstrate that the algorithm achieves near ideal tracking regardless of the command amplitude given that the required derivatives can be obtained.

However, one aspect that [12] and [14] did not take into consideration was the effect of disturbances on the inversion process. This issue is addressed by George, Verhaegen, and Scherpen in [16], where the Kalman filter is used to estimate the unknown input needed to track the desired output. This work is an extension of earlier results by Hou and Patton [17], who combined the unknown input decoupling approach of Hou and Muller [18]

with the innovative approach to deriving optimal filters by Kailath [19], and Anderson and Moore [20]. Another limiting factor of the previous approaches is the need to know the plant exactly. To address this issue, adaptive control and adaptive disturbance rejection (or noise cancellation) methods have been developed. The references for initial developments in active noise control are given by Burgess in [21], which also demonstrates a novel finite impulse response (FIR) adaptive noise controller. In particular, Burgess took the least-mean-squares (LMS) based adaptive echo/noise canceler used in signal processing (Sondhi and Presti [22] and Widrow *et al.* [23]) and applied it in a control context. This approach was further improved by relaxing the FIR-restriction on the controller, thereby allowing for infinite impulse response (IIR) controllers, as shown by Eriksson [24]. The IIR approach is demonstrated to improve disturbance rejection performance since it is capable of removing the poles introduced by feedback. As a continuation of the overall shift to the IIR approaches, an ARMARKOV/Toeplitz adaptive approach is proposed by Venugopal and Bernstein for disturbance rejection in [25] and for command following in [26]. The major differences of this approach from the approaches mentioned above include needing only secondary path numerator model of the plant, not requiring knowledge of the disturbance spectrum, and updating the controller based on the past error instead of the predicted future error. In particular, [27] and [28] show that the ARMARKOV controller attenuates disturbance almost twice as well as the filtered-u IIR LMS (FURLMS, [29]) algorithm (attenuations of 89.5 dB and 56.3 dB, respectively).

The ARMARKOV approach was further developed by Hoagg, Santillo, and Bernstein in [30,31], where plant Markov parameters were used instead of the numerator coefficients in controller design and introducing a retrospective correction filter (RCF). RCF pre-filters the measurements based on the difference between the actual past control inputs and re-computed past control inputs assuming the current controller was used in the past. This concept is implicit in the original formulation of [25], but its effects are more extensively investigated in [31]. This study demonstrates that the modified retrospective approach is

effective for disturbance rejection and command following problems for unstable, MIMO, and/or NMP plants. One improvement of this algorithm is described in Chapter 2, where a convex constraint is imposed on the poles of the controller to prevent them from attempting to cancel the NMP zeros. This constraint in addition to meeting the goal of restraining the poles inside the unit disk has the beneficial side effects of decreasing the transients experienced during adaptation and improving robustness to the knowledge of the NMP zero locations.

The next step in the evolution of retrospective cost adaptive controllers (RCAC) was the cumulative modification proposed by Hoagg and Bernstein in [32]. One of the innovations in this version is the recursive least squares update law, which improved the transient performance of the algorithm. The minimal information required by this algorithm is the first nonzero Markov parameter and the NMP zeros of the control-performance transfer function. Chapter 3 of this dissertation describes the degree of preciseness of the NMP zero locations required for the stability of RCAC. The performance of this version of RCAC for nonlinear command following is investigated in Chapter 4, where a model of a planar multilink arm is derived. The linearized transfer functions between noncollocated torque inputs and link angular positions are shown to be nonminimum phase. This example successfully demonstrated that RCAC can be applied to nonlinear command following and disturbance rejection problems when command and disturbance amplitudes are within certain bounds.

This application inspired another nonlinear command following (or, more broadly, data assimilation with driver estimation) application. Data assimilation with driver estimation can be thought of as a generalization of command following, since in addition to finding a control input required to drive the output of the system to the desired trajectory it computes such system states, which would bring system output even closer to the desired trajectory than control input alone ever could. Accordingly, data assimilation and driver estimation are complementary in achieving output tracking. To provide an analogy for this symbiotic relationship we revisit the car driver example, where the goal is to keep the car in the middle

of the lane. In this example, the car operator has a noisy (perhaps due to degraded vision) measurement of the car's position relative to the center of the lane and can influence this position via the steering wheel. On one hand, driver estimation (or command following) would consist of determining how to turn the steering wheel to get the car closer to the center of the lane in the future. On the other hand, data assimilation can be thought of operator having the power to reset car's position at any instance (for example, every minute), but keep the steering wheel in the same position throughout the whole experiment. Therefore, the only way the operator can influence the car's position during pure data assimilation (without driver estimation) is through, figuratively speaking, picking it up and placing it a desired distance away from the middle of the lane. At this point a somewhat unrelated question of 'Would the car operator put the car right in the middle of the lane the first time he had an opportunity?' comes up. The answer to this question depends on the car. If the car is old or has incorrect wheel alignment, it might drift to the left or to the right even with zero steering input. With such a car, operator is better off placing the car slightly off center to pre-compensate for this drifting effect (thereby inverting the system). This mental experiment is meant to provide an intuitive feel for the relationship between the data assimilation and driver estimation, namely that they complement each other by giving the operator both powers - to reset car's position as well as the steering input at every step.

In the aforementioned project inspired by the nonlinear linkage experiment, the system to be driven was a model of the upper atmosphere - the global ionosphere-thermosphere model (GITM, [33]). The goal of this project was to find inputs and states of GITM that would result in simulated GITM outputs tracking real satellite data. The motivation for this project came from the solar storms of October 2003 (Halloween Storms). A solar storm is an event of extreme solar activity resulting in elevated amount of energetic particles and electromagnetic waves affecting Earth. While the October 2003 storms did not leave 55 million people without power as was the case during the Northeast Blackout two months earlier, they demonstrated how little is known about the solar wind effects on the upper

atmosphere. In particular, the Halloween Storms damaged some instruments on NASA's Advanced Composition Explorer (ACE) satellite, temporarily disabled NASA-ESA Solar and Heliospheric Observatory (SOHO), and caused a power outage in Sweden [34–36]. As a result, the NASA Solar Dynamics Observatory (SDO) was launched in 2010 and is capable of providing solar-flare warnings four days before solar wind particles reach Earth. Nonetheless, even with such an advance warning, the capability to predict how a particular flare will affect Earth's thermosphere, ionosphere, and magnetosphere is not mature enough [37]. This fact is particularly apparent in the realm of orbit determination and conjunction analysis, where thermospheric drag can play a vital role [38]. Uncertainty in drag comes from inaccuracies in thermospheric models, which can be improved through data assimilation, that is, comparing model output to real data and adjusting model states and inputs to compensate for observed discrepancies.

One approach to data assimilation is ensemble Kalman filtering, which is a technique that uses multiple instances of the nonlinear model with different initial conditions to estimate the state of the system. The motivation for creating ensemble filters arose from several practical issues with the extended Kalman filter (EKF, [39]) and the unscented Kalman filter (UKF, [40]). One of the issues with the EKF is that it requires model linearization at every step, which might not be possible if the model does not have a simple analytic description (if it is a complex computational fluid dynamics simulation, for example). UKF, on the other hand, does not require linearization, but does require propagation of $2n$ sigma-points, where n is length of the state vector of the nonlinear system. Atmospheric models can have tens of millions of state variables, so UKF would have to propagate tens of millions of model copies, which is impractical with current computers.

To address these issues, the ensemble Kalman filter (EnKF) was proposed as described in [41, 42]. This filter saw many applications ranging from cardiac computer tomography [43] to fishery stock collapse problem [44] to assimilating data from the atmosphere on Mars [45]. In this dissertation, a filter similar to EnKF (ensemble adjustment Kalman

filter, EAKF, [46]) is used. EAKF (as well as many other filters) is implemented in the data assimilation research testbed (DART, [47]) written and maintained by the staff of the national center for atmospheric research (NCAR). A large part of the effort described in this dissertation went into interfacing DART and GITM and some software aspects of this interface are dispersed throughout this text. The minor difference between the EnKF and the EAKF is the fact that EAKF does not perturb the measurements to prevent the ensemble spread from going to zero (filter divergence). Filter divergence is said to occur when all ensemble members converge to the same trajectory and ignore the incoming measurements. EAKF solves this problem by artificially increasing the ensemble spread, thereby introducing a “healthy” level of disagreement between the ensemble members, which can be related to covariance resetting.

A novel approach to this problem is described in Chapter 5. The new approach consists of choosing different inflation values for different estimates. For example, the external driver estimates require more inflation than internal state variables, since the drivers are not updated by the model and therefore are more uncertain. In particular, we show that setting the driver variance to a constant value was beneficial as it made sure that the driver variance would not increase without bound, which can be the case when the covariance matrix is scaled up at every step. The differential inflation technique improves the filter convergence speed as well as the steady state performance, when compared to the equilateral inflation. Chapter 5 describes the assimilation of the Challenging Minisatellite Payload (CHAMP) neutral density measurements into GITM to estimate the solar activity index $F_{10.7}$, whereas Chapter 6 demonstrates assimilation of a more global measurement - total electron content. Both of these chapters demonstrate an improvement in model-data agreement through driver estimation.

In summary, the main contributions of this dissertation can be presented in the following list.

- **Description** of the Retrospective Cost Adaptive Control (RCAC) stability margins

for plants with uncertain nonminimum-phase zeros.

- **Introduction** of a convex constraint on the controller pole locations to improve RCAC transient and steady-state performance.
- **Modeling** of a nonlinear system and description of the achievable amplitude and frequency control ranges.
- **Modification** of the Data Assimilation Research Testbed (DART) to interface it with the Global Ionosphere-Thermosphere Model (GITM).
- **Development** of a novel inflation technique for the Ensemble Adjustment Kalman Filter (EAKF) as applied to GITM for purposes of data assimilation and driver estimation.
- **Introduction** of an ability to assimilate Total Electron Content (TEC) measurements into the DART-GITM interface.

CHAPTER 2

Retrospective Cost Adaptive Control for Nonminimum-Phase Systems with Uncertain Nonminimum-Phase Zeros Using Convex Optimization

The goal of this Chapter is to describe how the robustness of RCAC to uncertainty in the locations of the nonminimum-phase zeros can be increased. Specifically, a convex constraint is imposed on the poles of the controller in order to prevent the adaptive controller from attempting to cancel the nonminimum-phase zeros. Numerical results show that, when constrained convex optimization is used at each step, the transient response is improved and the adaptive controller has increased robustness to uncertainty in the locations of the nonminimum-phase zeros. The results of this Chapter are published in collaboration with D'Amato, Hoagg, and Bernstein in [48] and are further investigated in Chapter 3.

2.1 Introduction

Nonminimum-phase zeros are a major impediment to achievable performance in feedback control. While all fields of science and technology that work with dynamical systems are familiar with poles, the field of control is unique in recognizing the role of zeros in sys-

tems with inputs and outputs. The multiple ways in which zeros impact the operation of control systems are discussed in [49]. In the classical case of root locus analysis, the attraction of poles to zeros limits the magnitude of the feedback gain. The same phenomenon occurs in LQG control for SISO and MIMO systems, where, in the high-authority limit, the controller obtained from the Riccati equations drives some of the closed-loop poles to the open-left-half-plane reflections of the open-loop nonminimum-phase (NMP) zeros. Nonminimum-phase zeros thus limit the achievable bandwidth and control authority. In addition, real nonminimum-phase zeros are responsible for initial undershoot and direction reversals in the step response. These issues are well understood for continuous-time systems; for sampled-data systems, analogous phenomena occur.

Nonminimum-phase zeros are especially challenging in adaptive control since an adaptive controller may attempt to cancel a nonminimum-phase zero. Such pole-zero cancellation is impossible, and thus, in fixed-gain control, the use of a nonminimum-phase zero to cancel an unstable pole is well known to be ineffective. In adaptive control, the reverse situation occurs, namely, an adaptive controller may attempt to cancel a nonminimum-phase zero that limits its performance; however, the nonminimum-phase zero cannot be removed. The detrimental affect of this attempt to cancel a nonminimum-phase zero is the fact that the controller may seek to apply unbounded control effort, thereby destabilizing the closed-loop system.

In the present chapter we focus on retrospective cost adaptive control (RCAC) [50, 51]. RCAC uses a retrospective cost functional defined in terms of a surrogate performance variable, which is based on measured data over a past window of operation. In effect, the retrospective cost functional “looks backward” over the window of data in order to determine a controller modification that would have improved the past performance. This retrospective cost functional is optimized at each step in order to update the controller. The algorithm thus seeks the controller that achieves the best performance in terms of a prior window of operation as determined by the retrospective cost. The approach of “looking

backward” rather than forward (as in the case, for example, of model predictive control) allows RCAC to control the system under minimal modeling information.

As shown in [50, 52], RCAC has the ability to adaptively control nonminimum-phase systems if the locations of the nonminimum-phase zeros are known. Modeling information that captures the locations of the nonminimum-phase zeros (either SISO or MIMO) is included in the matrix \bar{B}_{zu} , as described in the next section. The matrix \bar{B}_{zu} can be defined in terms of the Markov parameters of the transfer function from the control input to the performance variable. The Markov parameters are coefficients of the Laurent expansion of the transfer function expressed in terms of powers of $1/z$. The Laurent expansion provides a convergent series for the transfer function outside of the spectral radius of the plant; consequently, this series automatically captures all nonminimum-phase zeros outside of the spectral radius. Alternatively, if the nonminimum-phase zeros are known, then their values can be used directly in \bar{B}_{zu} in place of a finite number of Markov parameters. Consequently, identification of nonminimum-phase zeros is of interest in practice [53].

The above discussion leads to the following question: How does RCAC exploit knowledge of the nonminimum-phase zeros in order to avoid any attempt to cancel them? This question is addressed in [52], where it is shown that the nonminimum-phase zeros appear in the numerator of a filter that processes the data in the regressor used in the controller update. This filter removes spectral content corresponding to the locations of the nonminimum-phase zeros, thus avoiding the possibility of having the adaptive controller attempt to cancel a nonminimum-phase zero.

Since RCAC requires knowledge of the nonminimum-phase zeros of the plant, it is of interest to determine the required accuracy of this information. Numerical examples in [50, 54] suggest a negative result, namely, that there may exist plants for which the robustness to uncertainty in the locations of the nonminimum-phase zeros may be arbitrarily small. This lack of robustness is manifested by the increasingly larger transients that arise as the nonminimum-phase-zero locations become increasingly uncertain. This negative re-

sult is consistent with [55, 56], namely, that adaptive control must confront plants that are inherently difficult to control. RCAC shows that this difficulty is inherent in the modeling information relating specifically to the nonminimum-phase zeros (if any are present).

In the present chapter our goal is to develop a technique that increases the robustness of RCAC to uncertainty in the locations of the nonminimum-phase zeros. To do this, we consider an extension of RCAC, where the minimization of the retrospective cost is performed subject to a constraint on the allowable locations of the controller poles. A convex constraint on eigenvalue locations is given in [57] and is used in [58] for model identification with guaranteed stability. However, this approach cannot be used with RCAC since RCAC updates the coefficients of the denominator of the controller transfer function, rather than the entries of an unstructured dynamics matrix. We thus use this polynomial to construct a companion matrix. Since a bound on the spectral radius of the companion matrix does not provide a convex constraint on the coefficients of the polynomial, we bound the spectral radius with a matrix norm, which defines a convex constraint. Although bounding the spectral radius with a matrix norm introduces conservatism, this conservatism has minimal effect since the magnitude of the bound on the matrix norm can be adjusted as a design parameter.

2.2 Problem Formulation

Consider the multi-input, multi-output discrete-time system

$$x(k + 1) = Ax(k) + Bu(k) + D_1w(k), \quad (2.1)$$

$$y(k) = Cx(k) + Du(k) + D_2w(k), \quad (2.2)$$

$$z(k) = E_1x(k) + E_2u(k) + E_0w(k), \quad (2.3)$$

where $x(k) \in \mathbb{R}^n$, $y(k) \in \mathbb{R}^{l_y}$, $z(k) \in \mathbb{R}^{l_z}$, $u(k) \in \mathbb{R}^{l_u}$, $w(k) \in \mathbb{R}^{l_w}$, and $k \geq 0$. Our goal is to develop an adaptive controller that generates a control signal u that minimizes the performance variable z in the presence of the exogenous signal w . We assume that measurements of the output y and the performance variable z are available for feedback; however, we assume that a direct measurement of the exogenous signal w is not available.

Note that w can represent either a command signal to be followed, an external disturbance to be rejected, or both. For example, if $D_1 = 0$, $E_2 = 0$, and $E_0 \neq 0$, then the objective is to have the output $E_1 x$ follow the command signal $-E_0 w$. On the other hand, if $D_1 \neq 0$, $E_2 = 0$, and $E_0 = 0$, then the objective is to reject the disturbance w from the performance measurement $E_1 x$. The combined command following and disturbance rejection problem is addressed when D_1 and E_0 are block matrices. Lastly, if D_1 and E_0 are empty matrices, then the objective is output stabilization, that is, convergence of z to zero.

The performance variable z can include the feedthrough term $E_2 u$. This term allows us to design an adaptive controller where the performance variable z to be minimized can include a weighting on control authority. For example, if $E_1 = [\hat{E}_1^T \ 0]^T$, $E_2 = [0 \ \hat{E}_2^T]^T$, and $E_0 = [\hat{E}_0^T \ 0]^T$, then the performance z consists of the components $z_1 \triangleq \hat{E}_1 x + \hat{E}_0 w$ and $z_2 \triangleq \hat{E}_2 u$. In this case, the goal is to minimize a weighted combination of z_1 and z_2 , where z_1 is the weighted state performance and z_2 is the weighted control authority.

We represent (2.1) and (2.3) as the time-series model from u and w to z given by

$$z(k) = \sum_{i=1}^n -\alpha_i z(k-i) + \sum_{i=d}^n \beta_i u(k-i) + \sum_{i=0}^n \gamma_i w(k-i), \quad (2.4)$$

where $\alpha_1, \dots, \alpha_n \in \mathbb{R}$, $\beta_d, \dots, \beta_n \in \mathbb{R}^{l_z \times l_u}$, $\gamma_0, \dots, \gamma_n \in \mathbb{R}^{l_z \times l_w}$, and the relative degree d is the smallest non-negative integer i such that the i th Markov parameter, either $H_0 \triangleq E_2$ if $i = 0$ or $H_i \triangleq E_1 A^{i-1} B$ if $i > 0$, is nonzero. Note that $\beta_d = H_d$.

2.3 Review of RCAC

In this section we give a brief overview of the RCAC. Full details are given in [31]. RCAC depends on several parameters that are selected *a priori*. Specifically, n_c is the controller order, p is the data window size, and μ is the number of Markov parameters. The adaptive update law is based on a quadratic cost function, which involves a time-varying weighting parameter $\zeta(k) > 0$, referred to as the *learning rate* since it affects the convergence speed of the adaptive control algorithm.

We use a strictly proper time-series controller of order n_c such that the control $u(k)$ is given by

$$u(k) = \sum_{i=1}^{n_c} M_i(k)u(k-i) + \sum_{i=1}^{n_c} N_i(k)y(k-i), \quad (2.5)$$

where $M_i \in \mathbb{R}^{l_u \times l_u}$, $i = 1, \dots, n_c$, and $N_i \in \mathbb{R}^{l_u \times l_y}$, $i = 1, \dots, n_c$, are given by an adaptive update law. The control can be expressed as

$$u(k) = \theta(k)\phi(k), \quad (2.6)$$

$$\theta(k) \triangleq \left[N_1(k) \cdots N_{n_c}(k) \ M_1(k) \cdots M_{n_c}(k) \right] \quad (2.7)$$

is the *controller parameter block matrix*, and the *regressor vector* $\phi(k)$ is given by

$$\phi(k) \triangleq \begin{bmatrix} y(k-1) \\ \vdots \\ y(k-n_c) \\ u(k-1) \\ \vdots \\ u(k-n_c) \end{bmatrix} \in \mathbb{R}^{n_c(l_u+l_y)}. \quad (2.8)$$

For positive integers p and μ , we define the *extended performance vector* $Z(k)$, and the *extended control vector* $U(k)$ by

$$Z(k) \triangleq \begin{bmatrix} z(k-1) \\ \vdots \\ z(k-p_c) \end{bmatrix}, \quad U(k) \triangleq \begin{bmatrix} u(k-1) \\ \vdots \\ u(k-p_c) \end{bmatrix}, \quad (2.9)$$

where $p_c \triangleq n + \mu + p - 1$.

From (2.6), it follows that the extended control vector $U(k)$ can be written as

$$U(k) \triangleq \sum_{i=1}^{p_c} L_i \theta(k-i) \phi(k-i), \quad (2.10)$$

where

$$L_i \triangleq \begin{bmatrix} \mathbf{0}_{(i-1)l_u \times l_u} \\ I_{l_u} \\ \mathbf{0}_{(p_c-i)l_u \times l_u} \end{bmatrix} \in \mathbb{R}^{p_c l_u \times l_u}. \quad (2.11)$$

We define the *surrogate performance vector* $\hat{Z}(\hat{\theta}(k), k)$ by

$$\hat{Z}(\hat{\theta}(k), k) \triangleq Z(k) - \bar{B}_{zu} (U(k) - \hat{U}(k)), \quad (2.12)$$

where

$$\hat{U}(k) \triangleq \sum_{i=1}^{p_c} L_i \hat{\theta}(k) \phi(k-i), \quad (2.13)$$

and $\hat{\theta}(k) \in \mathbb{R}^{l_u \times [n_c(l_u + l_y)]}$ is the *surrogate controller parameter block matrix*. The block-

Toeplitz *surrogate control matrix* \bar{B}_{zu} is given by

$$\bar{B}_{zu} \triangleq \begin{bmatrix} 0_{l_z \times l_u} & \cdots & 0_{l_z \times l_u} & H_d & \cdots \\ \vdots & \ddots & \ddots & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots \\ 0_{l_z \times l_u} & \cdots & 0_{l_z \times l_u} & 0_{l_z \times l_u} & \cdots \\ \cdots & H_\mu & 0_{l_z \times l_u} & \cdots & 0_{l_z \times l_u} \\ \ddots & \ddots & \ddots & \ddots & \vdots \\ \ddots & \ddots & \ddots & \ddots & \vdots \\ \cdots & 0_{l_z \times l_u} & H_d & \cdots & H_\mu \end{bmatrix}, \quad (2.14)$$

where the *relative degree* d is the smallest positive integer i such that the i th Markov parameter $H_i \triangleq C_0 A_0^{i-1} B_0$ is nonzero. The leading zeros in the first row of \bar{B}_{zu} account for the nonzero relative degree d . The algorithm places no constraints on either the value of d or the rank of H_d or \bar{B}_{zu} .

We now consider the cost function

$$J(\hat{\theta}, k) \triangleq \hat{Z}^T(\hat{\theta}, k) \hat{Z}(\hat{\theta}, k) + \zeta(k) \text{tr} \left[(\hat{\theta} - \theta)^T (\hat{\theta} - \theta) \right], \quad (2.15)$$

where the positive scalar $\zeta(k)$ is the learning rate. Substituting (2.12) into (2.15), the cost function can be written as the quadratic form

$$J(\hat{\theta}, k) = \left(\text{vec } \hat{\theta} \right)^T A(k) \text{vec } \hat{\theta} + b^T \text{vec } \hat{\theta} + c(k), \quad (2.16)$$

where

$$\begin{aligned}
D(k) &\triangleq \sum_{i=1}^{p_c} \phi^T(k-i) \otimes (\bar{B}_{zu} L_i), \\
f(k) &\triangleq Z(k) - \bar{B}_{zu} U(k), \\
A(k) &\triangleq D^T(k) D(k) + \zeta(k) I_{n_c l_u(l_u+l_y)}, \\
b(k) &\triangleq 2D^T(k) f(k) - 2\zeta(k) \text{vec } \theta(k), \\
c(k) &\triangleq f(k)^T f(k) + \zeta(k) \text{tr} [\theta^T(k) \theta(k)].
\end{aligned} \tag{2.17}$$

Since $A(k)$ is positive definite, $J(\hat{\theta}, k)$ has the strict global minimizer

$$\hat{\theta}(k) = -\frac{1}{2} \text{vec}^{-1}(A(k)^{-1} b(k)). \tag{2.18}$$

The controller gain update law is $\theta(k+1) = \hat{\theta}(k)$.

2.4 Constrained Convex Optimization

In this section we extend RCAC by using constrained convex optimization instead of (2.18) to update $\theta(k) \in \mathbb{R}^{1 \times 2n_c}$. For simplicity, we consider only SISO systems. The denominator coefficients of the controller can be constructed from the last n_c entries of $\theta(k)$ as

$$\text{den}(\theta(k)) \triangleq [1 \quad -M_1 \quad -M_2 \quad \cdots \quad -M_{n_c}]. \tag{2.19}$$

The roots of the monic polynomial whose coefficients are given by $\text{den}(\theta(k))$ are the pole locations of the adaptive controller at step k .

In order to prevent the poles of the adaptive controller from approaching the nonminimum-phase-zero locations, we constrain the poles to lie inside a disk centered at the origin of the complex plane. Accordingly, we modify the problem of minimizing (2.16) by imposing an

additional constraint on the companion-form matrix

$$K \triangleq \begin{bmatrix} M_1 & M_2 & \dots & M_{n_c-1} & M_{n_c} \\ 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 0 \end{bmatrix}. \quad (2.20)$$

We constrain the poles to a disk by bounding the spectral radius of K by a matrix norm, which is a convex function and thus defines a convex region in terms of the denominator coefficients $\text{den}(\theta(k))$. Various matrix norms can be used to bound the spectral radius. For example, every equi-induced norm provides an upper bound, see Corollary 9.4.5 of [59]. One such norm is the maximum singular value of K . Accordingly, the constraint we use is given by

$$\sigma_{\max}(K) \leq \gamma, \gamma > 1. \quad (2.21)$$

To investigate the conservatism of this bound when applied to matrices of form (2.20), we generate 10^5 10^{th} -order monic polynomials whose last ten coefficients are taken from a standard normal distribution. For each polynomial we compute the spectral radius $\rho(K)$ and the maximum singular value $\sigma_{\max}(K)$. Each dot in Figure 2.1 corresponds to a polynomial. Figure 2.1 suggests that the constraint (2.21) is conservative. For the rest of the chapter we consider only the maximum singular value to bound the spectral radius. The use of alternative bounds, which include maximum absolute value norm, row-norm, and Frobenius-norm, is left for future work.

Further investigation shows that if the focus of the conservatism study is restricted to second order polynomials with roots within unit circle, then the conservatism of the bound is more clearly shown in Figure 2.2. Figure 2.2(b) shows the ratio of the maximum

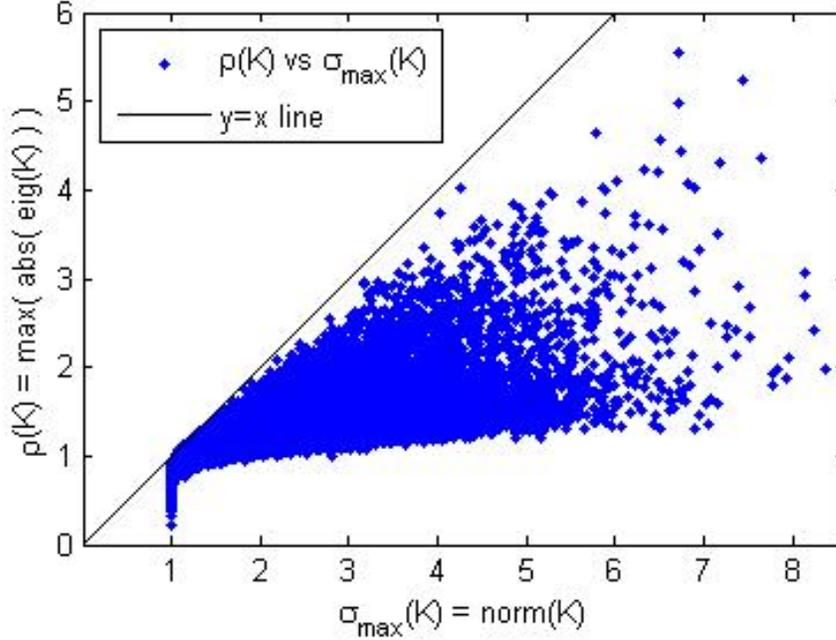


Figure 2.1: Spectral radius $\rho(K)$ is plotted versus maximum singular value $\sigma_{\max}(K)$ for 10^5 10^{th} order random monic polynomials. The plot shows that the constraint (2.21) is conservative, since a polynomial of high enough order can be found with companion maximum singular value arbitrarily far to the right of the $y = x$ line. This fact is further demonstrated in Figure 2.3.

singular value of a companion matrix to the corresponding polynomial root magnitude and demonstrates that in the worst case scenario this ratio might reach 2.414. Furthermore, Fact 5.11.30 of [59] allows for exact computation of the maximum singular value of a companion matrix directly from the polynomial coefficients as given by

$$\sigma_{\max} = \sqrt{\frac{1}{2} \left[1 + \sum_{i=1}^{n_c} M_i^2 + \sqrt{\left(1 + \sum_{i=1}^{n_c} M_i^2 \right)^2 - 4M_{n_c}^2} \right]}. \quad (2.22)$$

By computing (2.22) for all polynomials with roots shown in Figure 2.2(a) it was found that the maximum singular value among all these polynomials is achieved with roots of $[-1,-1]$ or $[1,1]$, both resulting in a value of $\sqrt{\frac{1}{2}(6 + \sqrt{32})} \approx 2.414$. We conjecture that the n^{th} order polynomials with all roots at the plus or minus 1 locations will have the largest

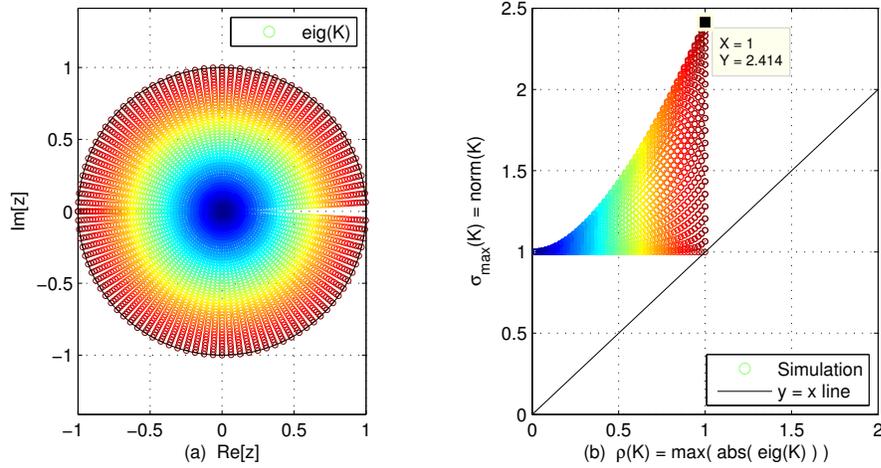


Figure 2.2: Roots of the selected second order polynomials are shown in (a) and the singular values of the corresponding (correspondence is demonstrated by color) companion matrix (K) are shown versus the polynomial roots in (b). Note that the axes in (b) are swapped as compared to Figure 2.1. (b) demonstrates that in the worst case scenario, the singular value of a companion matrix K is 2.414 times greater than the absolute value of the root of the polynomial (1.0).

maximum singular value out of all the polynomials of order n . Once the roots of n^{th} order polynomial are constrained to be at +1, (2.22) can be simplified by noting that in this case $b \triangleq 1 + \sum_{i=1}^{n_c} M_i^2 = \frac{(2n)!}{(n!)^2}$ and realizing that $(\frac{b}{2} + \frac{\sqrt{b^2-4}}{2})$ is a complete square, as given by

$$\sigma_{max}^{wc} = \sqrt{\frac{(2n)! - 2(n!)^2}{4(n!)^2}} + \sqrt{\frac{(2n)! + 2(n!)^2}{4(n!)^2}}, \quad (2.23)$$

where wc superscript denotes our conjecture that the polynomials with unity roots result in the worst case scenario, that is the largest singular values of the corresponding companion matrix. (2.23) is further demonstrated in Figure 2.3 as a function of polynomial order and is shown to be bounded from above by 2^n . Figure 2.3 demonstrates that the norm constraint is more conservative for higher order polynomials than for the lower order polynomials.

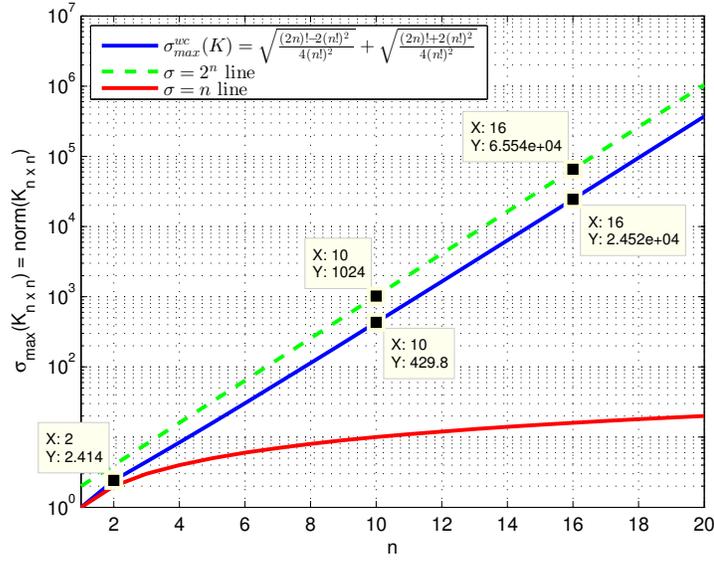


Figure 2.3: Maximum singular value for n^{th} order monic polynomial with n roots at 1. The conservatism of the maximum singular value increases as the order of the matrix increases.

2.5 Numerical Examples

We now demonstrate the performance of convex-constrained-RCAC (CC-RCAC) with constrained convex optimization on command following problems. The package CVX [60, 61] is used to minimize (2.16) subject to the constraint (2.21). For all examples in this section, the control objective is to have the plant output $y(k)$ follow a sinusoid with amplitude 1 and frequency $\frac{\pi}{10} \frac{\text{rad}}{\text{samp}} = 18 \frac{\text{deg}}{\text{samp}}$, except where noted otherwise. The adaptive controller (2.6) is implemented in feedback with $n_c = 10$, $\mu = 1$, $p = 1$, $\zeta \equiv 1$, and $\theta(0) = 0$. Also, we assume that the relative degree d and the first nonzero Markov parameter H_d are known.

The closed-loop is simulated for 1000 steps. Initial conditions are generated at the beginning of each simulation from a Gaussian distribution with mean 0 and variance 0.3. The transient and the steady-state performances are the two performance metrics used to compare RCAC with CC-RCAC. By transient performance we mean the maximum of the absolute value of the performance variable $z(k)$, and by steady-state performance we mean the maximum of the absolute value of the performance variable over the last 100 steps of

the simulation.

2.5.1 Third-Order Plant with a Known NMP Zero

We apply RCAC to a third-order plant with transfer function given by

$$G_3(z) \triangleq \frac{z - 1.4}{z^3 - 1.7z^2 + 1.2z - 0.35}. \quad (2.24)$$

RCAC requires knowledge of relative degree, the first nonzero Markov parameter, and the NMP zero. For this example the exact values of these parameters are assumed to be known, and therefore $d = 2$, $H_d = 1$, and $\bar{B}_{zu} = [0 \ 0 \ 1 \ -1.4]$. For RCAC Figure 2.4 shows the performance $z(k)$, control input $u(k)$, and controller coefficients $\theta(k)$. After the controller is turned on at $k = 100$, the performance variable $z(k)$ approaches zero in about 200 steps, and the controller coefficients converge in about 400 steps. Figure 2.5 shows the evolution of the controller poles, which settle in about 400 steps.

Under the same assumptions we apply CC-RCAC with $\gamma = 1.1$. For CC-RCAC Figure 2.6 shows the performance $z(k)$, control input $u(k)$, and controller coefficients $\theta(k)$. After the controller is turned on, the performance variable $z(k)$ approaches zero and the controller coefficients converge in about 200 steps. Figure 2.7 shows the evolution of the controller poles, which settle in about 200 steps. For this example Figures 2.4-2.7 show that, compared to RCAC, CC-RCAC provides improved performance in terms of decreasing the transient magnitude by about 70 percent and decreasing the settling time by about 50 percent.

2.5.2 Third-Order Plant with an Uncertain NMP Zero

We now compare the performance of RCAC with CC-RCAC on the plant (2.24), but with a NMP zero whose location is varied from 1.1 to 4.1. In addition, we also test the robustness of both algorithms to uncertainty in the estimate of the NMP zero. Therefore, the estimate

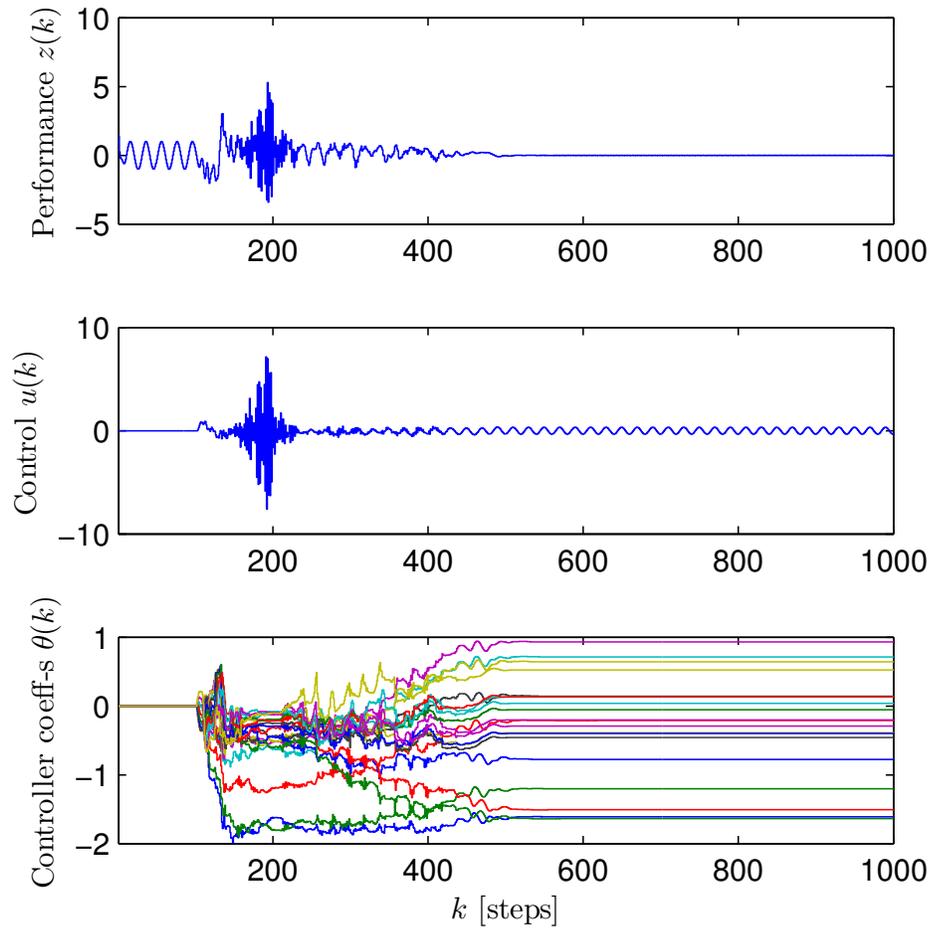


Figure 2.4: RCAC performance, control, and controller coefficients are shown as a function of time for the transfer function (2.24).

of the NMP zero is varied from 1.1 to 4.1 regardless of the location of the actual NMP zero. The true values for d and H_d are provided to the algorithm.

Figure 2.8 shows the transient and steady-state performances for various locations of the nonminimum-phase zero and its estimates using RCAC. Note that in some cases (particularly small values of the estimate of the NMP zero locations and large actual NMP zero locations) the closed-loop becomes unstable as signified by white in these plots. Additionally, the color map for these plots is saturated so that dark red corresponds to transient performance greater than 100. Also note that, the diagonal in both plots corresponds to

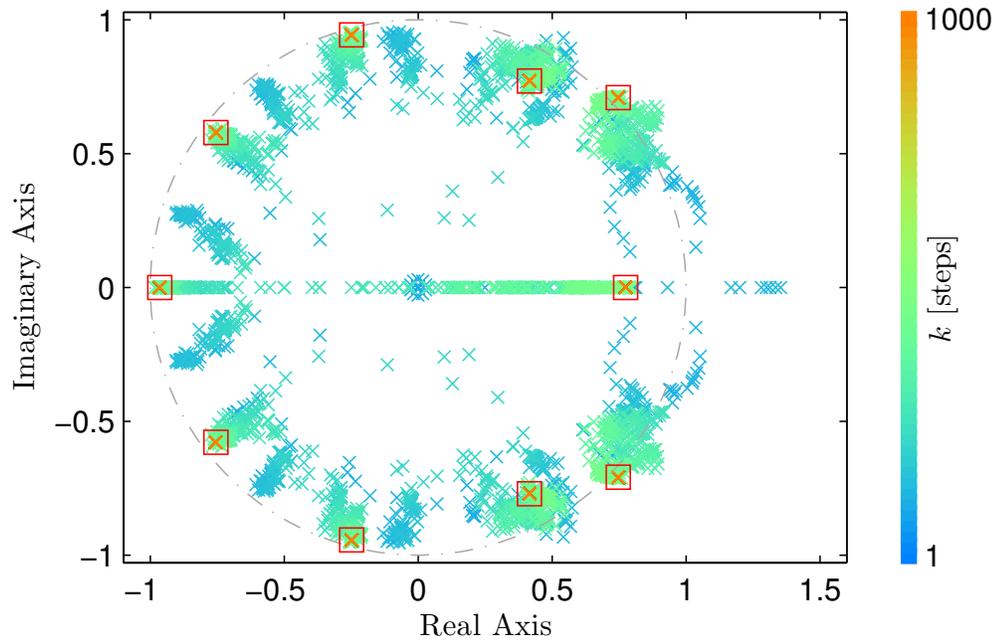


Figure 2.5: Evolution of the RCAC controller poles is shown as a function of time in terms of color. Later pole locations are colored with warmer colors and the final locations are marked with a red square. After the controller is turned on, the poles settle in about 400 steps.

the nominal case, i.e., the zero estimate location is at the actual zero location. Figure 2.8 indicates that RCAC is more robust to overestimating the NMP zero location than underestimating it and that when the NMP zero is further out on the real axis, RCAC has greater stability margins than when the NMP zero is closer to the unit circle.

Next, Figure 2.9 shows the CC-RCAC transient and steady-state performance and indicates that, like RCAC, CC-RCAC is more robust to overestimating the NMP zero location than underestimating it. Also, CC-RCAC has wider stability margins when compared to RCAC since the white area in Figure 2.9 is reduced by about 70 percent as compared to Figure 2.8.

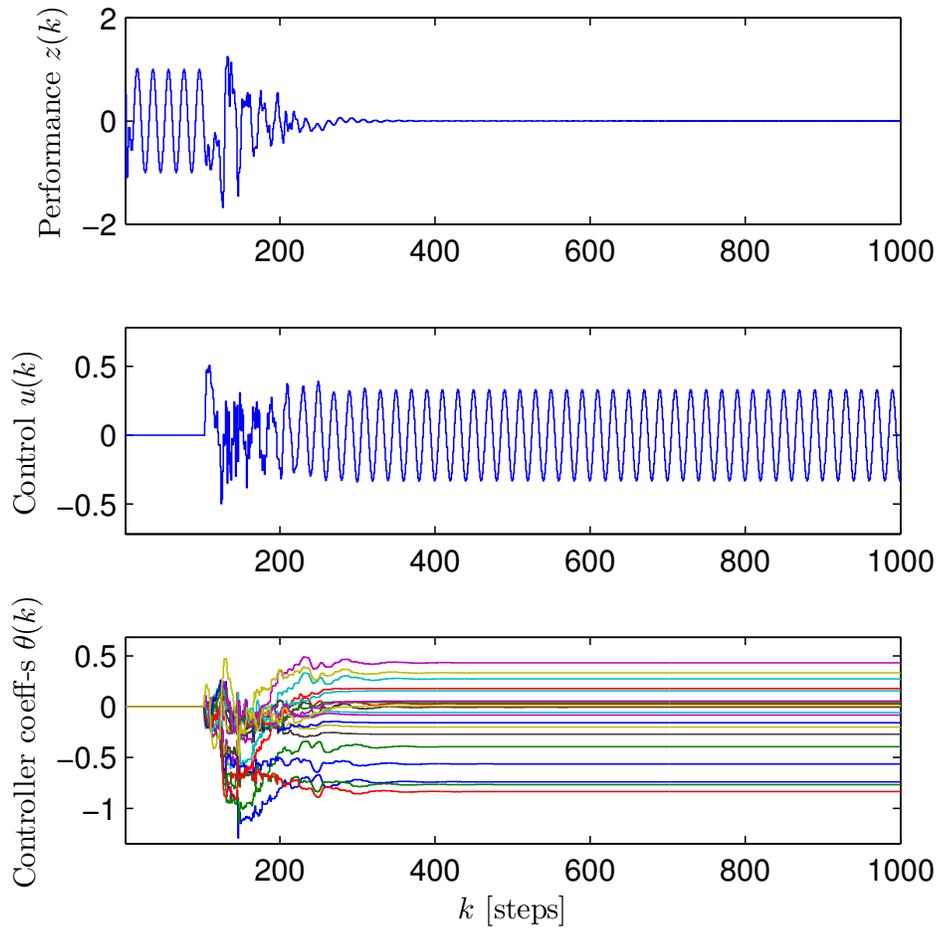


Figure 2.6: CC-RCAC performance, control, and controller coefficients are shown as a function of time for the transfer function (2.24).

2.5.3 Effect of Varying γ

We now examine the performance of CC-RCAC for various values of γ , more precisely, from 1.1 to 2.1. The third-order plant (2.24) with the NMP zero location varied from 1.1 to 4.1, is used. The exact location of the NMP zero is provided to the algorithm. Figure 2.10 shows the transient and steady-state performances for various values of γ and known NMP zero locations. The strip above each plot shows the RCAC performance for various NMP zero locations. Figure 2.10 suggests that the best performance for each location of the NMP zero is achieved at the lowest boundary of these plots, namely, $\gamma = 1.1$. $\gamma = 1.1$ was

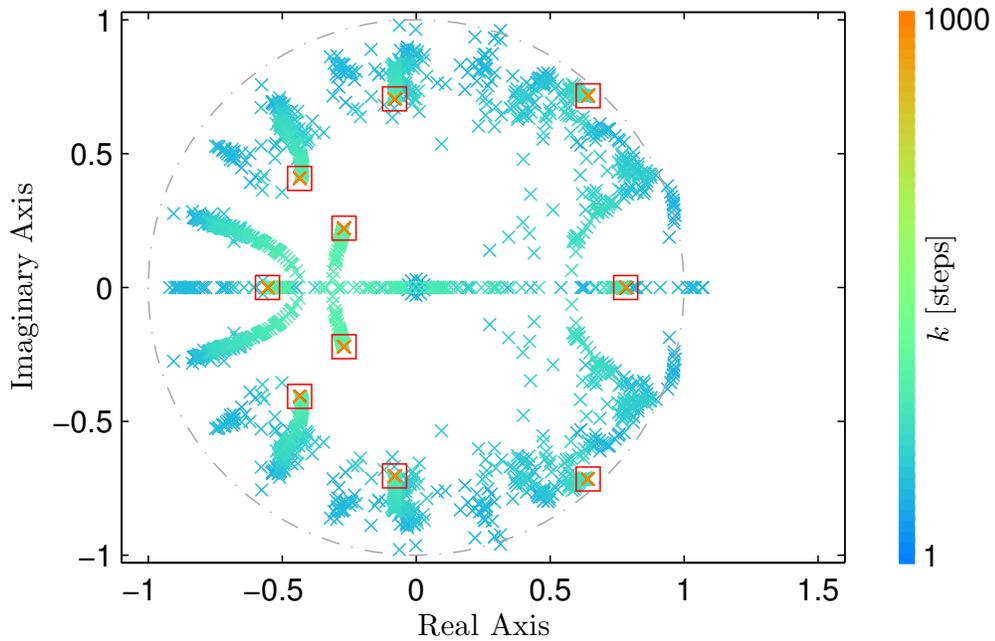


Figure 2.7: Evolution of the CC-RCAC controller poles is shown as a function of time in terms of color. After the controller is turned on, the poles settle in about 200 steps.

found to work satisfactorily in all cases tested. Additionally, plants with NMP zeros closer to the unit circle have smaller transients and smaller steady-state errors. Lastly, it can be seen that CC-RCAC produces improved transient and steady-state performances compared to RCAC.

2.5.4 Fourth-Order Plants with Uncertain NMP Zeros

In this section we compare RCAC and CC-RCAC performance on 20 fourth-order plants with uncertain NMP zeros. The plants used for this example have poles generated from uniform random distribution, first nonzero Markov parameter equaling 1, and the nonminimum-phase zero location of 2. The command following problem for these plants is simulated in feedback with RCAC, while the estimate of the NMP zero is varied from 1.4 to 2.6. Figure 2.11 shows the resulting transient and steady-state performance. Next, the same set of plants is simulated in feedback with CC-RCAC with $\gamma = 1.1$. Figure 2.12 shows the result-

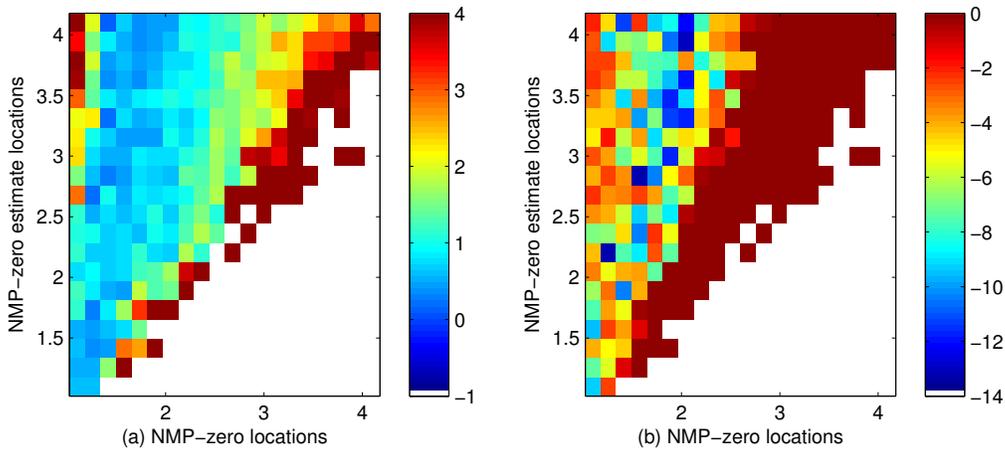


Figure 2.8: For the third-order plant (2.24) with a nonminimum-phase zero whose position is given by the horizontal axis, RCAC uses the estimate of the nonminimum-phase zero whose location is given by the vertical axis. The color in (a) shows the logarithm of the transient performance, whereas (b) shows the logarithm of the steady-state performance.

ing transient and steady-state performance. Figures 2.11 and 2.12 suggest that both RCAC and CC-RCAC have problems adapting to plants 7, 9, 11 and 19, since the steady-state performance for both versions of the algorithm is relatively high, as compared to the rest of the plants. Additionally, both versions of the algorithm are more robust to overestimating the location of the nonminimum-phase zero than to underestimating it. However, these Figures also suggest that, on average, CC-RCAC produces smaller transient responses, as compared to RCAC. Lastly, CC-RCAC has greater stability margin, as shown by the fact that there are fewer white regions in Figure 2.12 than in Figure 2.11.

2.6 Conclusion

Retrospective cost adaptive control (RCAC) is applicable to command following and disturbance rejection problems under minimal modeling assumptions, namely, knowledge of the relative degree, first nonzero Markov parameter, and nonminimum-phase zeros. In RCAC, the controller is updated by optimizing a surrogate performance variable that is used to define a retrospective cost. The retrospective cost uses knowledge of the nonminimum-phase

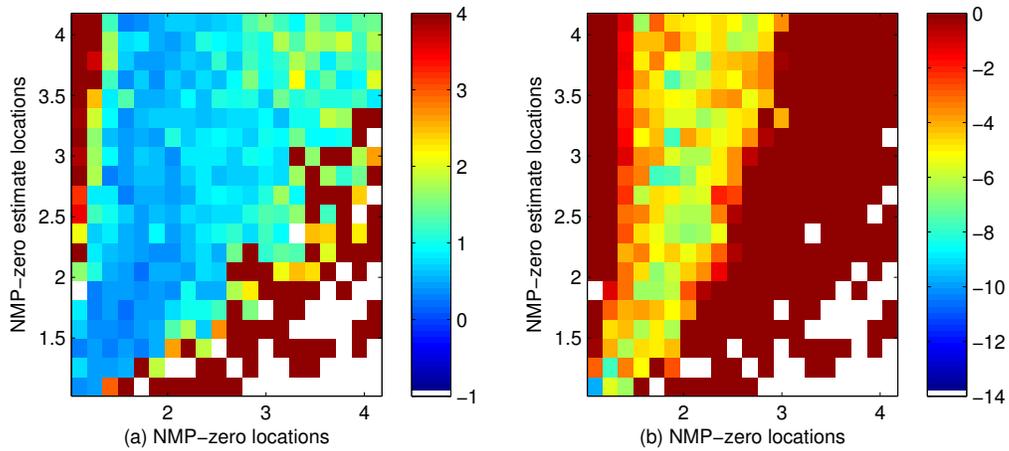


Figure 2.9: For the third-order plant with a nonminimum-phase zero whose position is given by the horizontal axis, CC-RCAC with $\gamma = 1.1$ uses the estimate of the nonminimum-phase zero whose location is given by the vertical axis.

zeros to prevent unstable pole-zero cancellation. The goal of this chapter is to increase the robustness of RCAC to uncertainty in the nonminimum-phase zero locations. To do this, we extend RCAC to include a convex constraint on the locations of the controller poles. Convex optimization is then used to optimize the retrospective cost subject to this constraint. The resulting convex-constrained retrospective cost adaptive controller (CC-RCAC) was found to have improved transient and steady-state performance as well as improved robustness to uncertainty in the locations of the nonminimum-phase zeros.

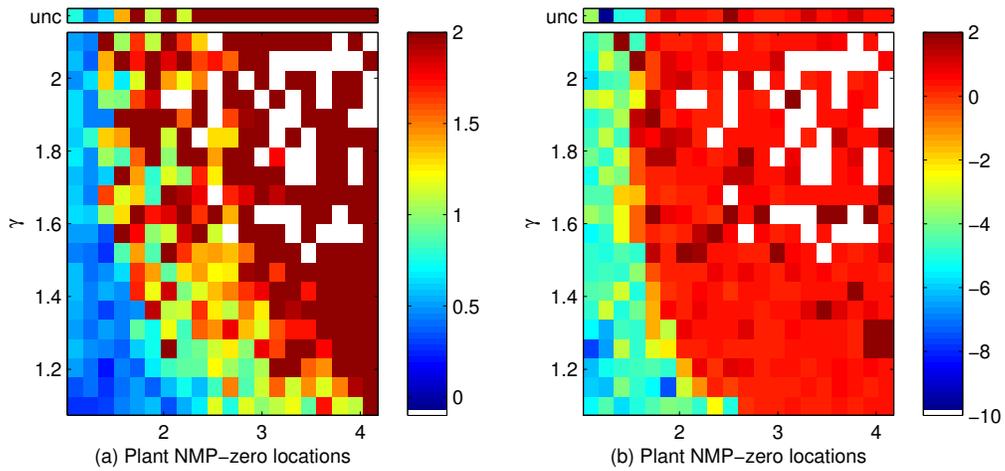


Figure 2.10: For the $G_3(z)$ with a nonminimum-phase zero whose position is given by the horizontal axis, CC-RCAC uses the exact location of the nonminimum-phase zero and a value of γ that varies from 1.1 to 4.1. For 20 values of γ , the best performance is achieved with $\gamma = 1.1$.

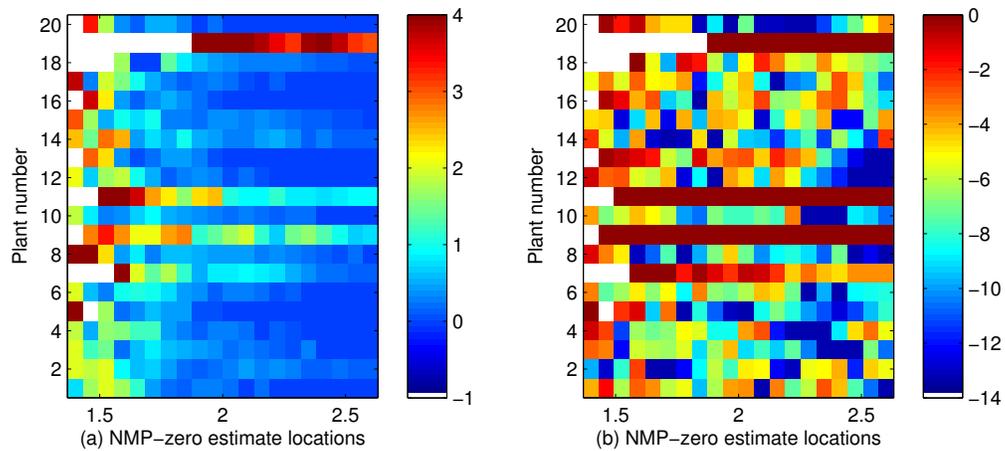


Figure 2.11: For 20 random fourth-order plants with a nonminimum-phase zero at 2, RCAC uses an estimate of the nonminimum-phase zero whose location is given by the horizontal axis.

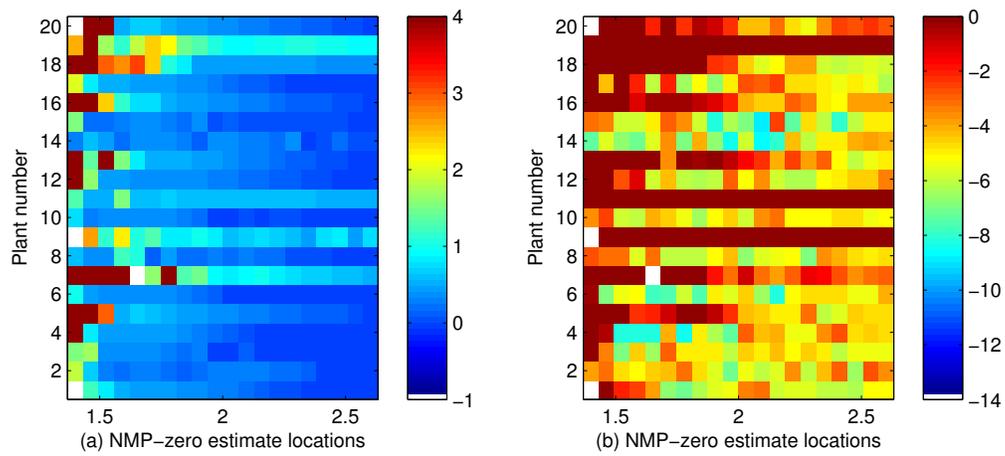


Figure 2.12: For 20 random fourth-order plants, CC-RCAC with $\gamma = 1.1$ uses an estimate of the nonminimum-phase zero whose location is given by the horizontal axis.

CHAPTER 3

A Computational Study of the Performance and Robustness Properties of Retrospective Cost Adaptive Control

This Chapter presents a computational study of a discrete-time adaptive control algorithm that is effective for multi-input, multi-output systems that are either minimum phase or nonminimum phase. The adaptive control algorithm requires limited model information, specifically, the first nonzero Markov parameter and the nonminimum-phase transmission zeros of the transfer function from the control signal to the performance measurement. Furthermore, the adaptive control algorithm is effective for stabilization, command following, and disturbance rejection. For command following and disturbance rejection, the algorithm does not require knowledge of the command or disturbance spectrum. The algorithm's performance and robustness in the presence of errors in the required modeling information is explored. The results of this Chapter are published in collaboration with Hoagg and Bernstein in [54].

3.1 Introduction

The existence of nonminimum-phase zeros is a major challenge in direct adaptive control. In fact, many direct adaptive control methodologies rely on the assumption that the plant

is minimum phase [62–66]. Furthermore, for command following and disturbance rejection, many adaptive control methodologies rely on the assumptions that the command and disturbance spectra are known or that the disturbances are measured [67, 68]. Sensitivity to process noise and sensor noise is an additional challenge in adaptive control, which may be dealt with using robust adaptive control laws [64, 69]. In the present chapter, we review a discrete-time adaptive controller that addresses these challenges. Furthermore, we computationally explore the algorithm’s performance and robustness in the presence of errors in the required model information.

Discrete-time versions of many continuous-time algorithms are available in the literature [63, 65, 70–72]. In addition, there are adaptive control algorithms that are unique to discrete-time [65, 73–75]. In Ref. [65, 73], discrete-time adaptive control laws are presented for stabilization and command following of minimum-phase systems based on the assumption that the command signals are known *a priori* and that an ideal tracking controller exists. An extension is given in Ref. [74], which addresses the combined stabilization, command following, and disturbance rejection problem. Note that the results of Ref. [65, 73, 74] are restricted to minimum-phase systems. For nonminimum-phase systems, Ref. [75] shows that periodic control may be used; however, this adaptive control scheme requires periods of open-loop operation.

Discrete-time adaptive controllers using a retrospective cost are known to be effective for stabilization, command following, and disturbance rejection for systems that are either minimum phase or nonminimum phase provided that knowledge of the nonminimum-phase zeros is available [32, 76, 77]. Furthermore, these retrospective cost adaptive controllers are effective for command following and disturbance rejection where the spectrum of the commands and disturbances is unknown and the disturbance is unmeasured. Proof of stability and convergence in the minimum-phase case is given in Ref. [74], while extensions to the nonminimum-phase case are described in Ref. [32, 77]. Retrospective cost adaptive control uses a retrospective performance measurement, in which the performance measurement is

modified based on the difference between the actual past control inputs and the recomputed past control inputs, assuming that the current controller had been used in the past.

The adaptive laws of Ref. [76, 77] are derived by minimizing an instantaneous retrospective cost, which is a function of the retrospective performance at the current time step, whereas the adaptive laws of Ref. [32] are derived by minimizing a cumulative retrospective cost function, which is a function of the retrospective performance at the current time step and all previous time steps. Retrospective cost adaptive controllers have been demonstrated on various experiments and applications, including the Air Force's deployable optical telescope testbed in Ref. [78], the NASA generic transport model in Ref. [79], and flow control problems in Ref. [80].

The goal of the present chapter is to examine the robustness (e.g., stability margins) and performance (e.g., transient behavior and steady state behavior) of the cumulative retrospective cost adaptive controller presented in Ref. [32]. In particular, we computationally explore the algorithm's behavior in the presence of inaccuracies in the required model information. In Section 3.4, we explore computationally the controller's sensitivity to the accuracy of the required model information, specifically, the first nonzero Markov parameter and nonminimum-phase zeros from the control to the performance. Sampled data systems are examined in detail in Section 3.5.

3.2 Problem Formulation

Consider the multi-input, multi-output discrete-time system

$$x(k+1) = Ax(k) + Bu(k) + D_1w(k), \quad (3.1)$$

$$y(k) = Cx(k) + Du(k) + D_2w(k), \quad (3.2)$$

$$z(k) = E_1x(k) + E_2u(k) + E_0w(k), \quad (3.3)$$

where $x(k) \in \mathbb{R}^n$, $y(k) \in \mathbb{R}^{l_y}$, $z(k) \in \mathbb{R}^{l_z}$, $u(k) \in \mathbb{R}^{l_u}$, $w(k) \in \mathbb{R}^{l_w}$, and $k \geq 0$. Our goal is to develop an adaptive controller that generates a control signal u that minimizes the performance z in the presence of the exogenous signal w . We assume that measurements of the output y and the performance z are available for feedback; however, we assume that a direct measurement of the exogenous signal w is not available.

Note that w can represent either a command signal to be followed, an external disturbance to be rejected, or both. For example, if $D_1 = 0$, $E_2 = 0$, and $E_0 \neq 0$, then the objective is to have the output $E_1 x$ follow the command signal $-E_0 w$. On the other hand, if $D_1 \neq 0$, $E_2 = 0$, and $E_0 = 0$, then the objective is to reject the disturbance w from the performance measurement $E_1 x$. The combined command following and disturbance rejection problem is addressed when D_1 and E_0 are block matrices. Lastly, if D_1 and E_0 are empty matrices, then the objective is output stabilization, that is, convergence of z to zero.

The performance variable z can include the feedthrough term $E_2 u$. This term allows us to design an adaptive controller where the performance z to be minimized can include a weighting on control authority. For example, if $E_1 = [\hat{E}_1^T \ 0]^T$, $E_2 = [0 \ \hat{E}_2^T]^T$, and $E_0 = [\hat{E}_0^T \ 0]^T$, then the performance z consists of the components $z_1 \triangleq \hat{E}_1 x + \hat{E}_0 w$ and $z_2 \triangleq \hat{E}_2 u$. In this case, the goal is to minimize a weighted combination of z_1 and z_2 , where z_1 is the weighted state performance and z_2 is the weighted control authority.

We represent (3.1) and (3.3) as the time-series model from u and w to z given by

$$z(k) = \sum_{i=1}^n -\alpha_i z(k-i) + \sum_{i=d}^n \beta_i u(k-i) + \sum_{i=0}^n \gamma_i w(k-i), \quad (3.4)$$

where $\alpha_1, \dots, \alpha_n \in \mathbb{R}$, $\beta_d, \dots, \beta_n \in \mathbb{R}^{l_z \times l_u}$, $\gamma_0, \dots, \gamma_n \in \mathbb{R}^{l_z \times l_w}$, and the relative degree d is the smallest non-negative integer i such that the i th Markov parameter, either $H_0 \triangleq E_2$ if $i = 0$ or $H_i \triangleq E_1 A^{i-1} B$ if $i > 0$, is nonzero. Note that $\beta_d = H_d$.

3.3 Review of the Cumulative Retrospective Cost Adaptive Controller

In this section, we present an adaptive control algorithm for the general control problem represented by (3.1)-(3.3). We use a strictly proper time-series controller of order n_c , such that the control $u(k)$ is given by

$$u(k) = \sum_{i=1}^{n_c} M_i(k)u(k-i) + \sum_{i=1}^{n_c} N_i(k)y(k-i), \quad (3.5)$$

where, for all $i = 1, \dots, n_c$, $M_i : \mathbb{N} \rightarrow \mathbb{R}^{l_u \times l_u}$ and $N_i : \mathbb{N} \rightarrow \mathbb{R}^{l_u \times l_y}$ are determined by the adaptive control law presented below. The control (3.5) can be expressed as

$$u(k) = \theta(k)\phi(k), \quad (3.6)$$

where

$$\theta(k) \triangleq \begin{bmatrix} N_1(k) & \cdots & N_{n_c}(k) & M_1(k) & \cdots & M_{n_c}(k) \end{bmatrix},$$

and

$$\phi(k) \triangleq \begin{bmatrix} y(k-1) \\ \vdots \\ y(k-n_c) \\ u(k-1) \\ \vdots \\ u(k-n_c) \end{bmatrix} \in \mathbb{R}^{n_c(l_u+l_y)}. \quad (3.7)$$

Next, we define the retrospective performance

$$\hat{z}(\hat{\theta}, k) \triangleq z(k) + \sum_{i=d}^{\nu} \bar{\beta}_i [\hat{\theta} - \theta(k-i)] \phi(k-i), \quad (3.8)$$

where $\nu \geq d$, $\hat{\theta} \in \mathbb{R}^{l_u \times (n_c(l_y + l_u))}$ is an optimization variable used to derive the adaptive law, and $\bar{\beta}_d, \dots, \bar{\beta}_\nu \in \mathbb{R}^{l_z \times l_u}$. The parameters ν and $\bar{\beta}_d, \dots, \bar{\beta}_\nu$ must capture the information included in the first nonzero Markov parameter and the nonminimum-phase zeros from u to z [32]. In this chapter, we let $\bar{\beta}_d, \dots, \bar{\beta}_\nu$ be the coefficients of the portion of the numerator polynomial matrix $\beta(\mathbf{z}) \triangleq \mathbf{z}^{n-d} \beta_d + \mathbf{z}^{n-d-1} \beta_{d+1} + \dots + \mathbf{z} \beta_{n-1} + \beta_n$ that includes the nonminimum-phase transmission zeros. More specifically, let $\beta(\mathbf{z})$ have the polynomial matrix factorization $\beta(\mathbf{z}) = \beta_U(\mathbf{z}) \beta_S(\mathbf{z})$, where $\beta_U(\mathbf{z})$ is an $l_z \times l_u$ polynomial matrix of degree $n_U \geq 0$ whose leading matrix coefficient is β_d , $\beta_S(\mathbf{z})$ is a monic $l_u \times l_u$ polynomial matrix of degree $n - n_U - d$, and each Smith zero of $\beta(\mathbf{z})$ counting multiplicity that lies on or outside the unit circle is a Smith zero of $\beta_U(\mathbf{z})$. Next, we can write $\beta_U(\mathbf{z}) = \beta_{U,0} \mathbf{z}^{n_U} + \beta_{U,1} \mathbf{z}^{n_U-1} + \dots + \beta_{U,n_U-1} \mathbf{z} + \beta_{U,n_U}$, where $\beta_{U,0} \triangleq \beta_d$. In this case, we let $\nu = n_U + d$ and for $i = d, \dots, n_U + d$, $\bar{\beta}_i = \beta_{U,i-d}$. For other choices of the parameters ν and $\bar{\beta}_d, \dots, \bar{\beta}_\nu$, see Ref. [32].

Note that if the transfer function from u to z is minimum phase, that is, the invariant zeros of (A, B, E_1, E_0) are contained inside of the unit circle, then $\nu = d$ and $\bar{\beta}_d = H_d$. The minimum-phase case with $z = y$ is considered in Ref. [74] using a gradient-based adaptive law rather than the RLS-based adaptive law considered in this chapter. Under the minimum-phase assumption, Ref. [74] proves asymptotic convergence of z to zero

Next, defining $\hat{\Theta} \triangleq \text{vec } \hat{\theta} \in \mathbb{R}^{n_c l_u (l_y + l_u)}$ and $\Theta(k) \triangleq \text{vec } \theta(k) \in \mathbb{R}^{n_c l_u (l_y + l_u)}$, it follows that

$$\hat{z}(\hat{\Theta}, k) = z(k) + \sum_{i=d}^{\nu} \Phi_i^T(k) [\hat{\Theta} - \Theta(k-i)] = z(k) - \sum_{i=d}^{\nu} \Phi_i^T(k) \Theta(k-i) + \Psi^T(k) \hat{\Theta}, \quad (3.9)$$

where, for $i = d, \dots, \nu$,

$$\Phi_i(k) \triangleq \phi(k-i) \otimes \bar{\beta}_i^T \in \mathbb{R}^{(n_c l_u (l_y + l_u)) \times l_z},$$

where \otimes represents the Kronecker product, and

$$\Psi(k) \triangleq \sum_{i=d}^{\nu} \Phi_i(k).$$

Now, define the cumulative retrospective cost function

$$J(\hat{\Theta}, k) \triangleq \sum_{i=0}^k \lambda^{k-i} \hat{z}^T(\hat{\Theta}, i) R \hat{z}(\hat{\Theta}, i) + [\hat{\Theta} - \Theta(0)]^T Q [\hat{\Theta} - \Theta(0)], \quad (3.10)$$

where $\lambda \in (0, 1]$, and $R \in \mathbb{R}^{l_z \times l_z}$ and $Q \in \mathbb{R}^{(n_c l_u (l_y + l_u)) \times (n_c l_u (l_y + l_u))}$ are positive definite. Note that λ serves as a forgetting factor, which allows more recent data to be weighted more heavily than earlier data.

The cumulative retrospective cost function (3.10) is minimized by a recursive least-squares (RLS) algorithm with a forgetting factor [63, 65, 66]. Therefore, $J(\hat{\Theta}, k)$ is minimized by the adaptive law

$$\Theta(k+1) = \Theta(k) - P(k) \Psi(k) \left[\lambda R^{-1} + \Psi^T(k) P(k) \Psi(k) \right]^{-1} z_R(k), \quad (3.11)$$

$$P(k+1) = \frac{1}{\lambda} P(k) - \frac{1}{\lambda} P(k) \Psi(k) \left[\lambda R^{-1} + \Psi^T(k) P(k) \Psi(k) \right]^{-1} \Psi^T(k) P(k), \quad (3.12)$$

where $P(0) = Q^{-1}$, $\Theta(0) \in \mathbb{R}^{n_c l_u (l_y + l_u)}$, and the retrospective performance measurement $z_R(k) \triangleq \hat{z}(\Theta(k), k)$. Note that $z_R(k)$ is computable from (3.9) using measured signals z , y , u , θ , and the matrix coefficients $\bar{\beta}_d, \dots, \bar{\beta}_\nu$. The cumulative retrospective cost adaptive control law is thus given by (3.11), (3.12), and

$$u(k) = \theta(k) \phi(k) = \text{vec}^{-1}(\Theta(k)) \phi(k). \quad (3.13)$$

The key feature of the adaptive control algorithm is the use of the retrospective performance (3.9), which modifies the performance variable $z(k)$ based on the difference between the actual past control inputs $u(k-d), \dots, u(k-\nu)$ and the recomputed past control inputs

$\hat{u}(\hat{\Theta}, k-d) \triangleq \text{vec}^{-1}(\hat{\Theta})\phi(k-d), \dots, \hat{u}(\hat{\Theta}, k-\nu) \triangleq \text{vec}^{-1}(\hat{\Theta})\phi(k-\nu)$, assuming that the current controller $\hat{\Theta}$ had been used in the past. Note that the cumulative retrospective cost adaptive controller (3.11)-(3.13) requires the matrix coefficients $\bar{\beta}_d, \dots, \bar{\beta}_\nu$, which are estimates of the first nonzero Markov parameters and the nonminimum-phase zeros.

3.4 Sensitivity to the Accuracy of the Required Model Information

In this section, we investigate the sensitivity of the cumulative retrospective cost adaptive controller (3.11)-(3.13) to the accuracy of the estimates of the first nonzero Markov parameter and the nonminimum-phase zeros (i.e., the parameters $\bar{\beta}_d, \dots, \bar{\beta}_\nu$). We systematically test the adaptive controller's performance with a range of stable minimum-phase as well as nonminimum-phase plants.

3.4.1 Sensitivity to the Accuracy of the Nonminimum-Phase Zero Estimate

For all examples in this section, we consider the command following problem, where the control objective is to have the plant output follow a sinusoid of frequency $\pi/6$ radians per sample with amplitude 10. The adaptive controller (3.11)-(3.13) is implemented in feedback with $\lambda = 1$, $R = 1$, $n_c = 2n - d + 4$, $P(0) = I_{2n_c}$, and $\theta(0) = 0$. Also, we assume that the relative degree d and the first nonzero Markov parameter are known exactly, that is, we let $\nu = d + 1$ and $\bar{\beta}_d = H_d$. The plants have the initial condition of 1 for every state.

3.4.1.1 Sensitivity with one nonminimum-phase zero

First, numerical tests are performed on 40 third order stable nonminimum-phase plants to explore the algorithm's robustness to the location of the nonminimum-phase zero estimate.

The single-input, single-output plants used for this experiment have poles at 0.5, 0.6 and 0.7, relative degree 2, first nonzero Markov parameter equaling 1, and the nonminimum-phase zero location is varied from 1.1 to 8.1. For each particular plant the adaptive algorithm requires estimates of the nonminimum-phase (NMP) zero, relative degree and first nonzero Markov parameter. Accordingly, the exact values of the relative degree and first nonzero Markov parameter are provided to the algorithm, while the estimate of the NMP zero is varied from 1.1 to 8.1. Next, the adaptive algorithm is connected in feedback and the closed loop is simulated. Figure 3.1 shows the logarithm of the transient and steady state performances for various locations of the nonminimum-phase zero and its estimates. More precisely, by transient performance we mean the maximum of the absolute value of the performance variable $z(k)$, and by steady state performance we mean the maximum of the absolute value of the performance variable in the last hundred steps of the simulation. Note that the diagonal in both plots corresponds to the nominal case, i.e. zero estimate location is at the actual zero location. Also, when the estimate is much greater than the actual location (upper left corner) or much lower (lower right corner), the closed loop becomes unstable and this is signified by white color in these plots. Additionally, the color map for these plots is saturated in that dark red color not only corresponds, for example, to transient performance of 10^2 , but also all transients greater than 10^2 . Lastly, Figure 3.1 suggests that the adaptive controller is more robust to overestimating the NMP zero location than underestimating it and that the plants with NMP zeros that are further out on the real axis have greater stability margins than those that have NMP zero closer to the unit circle.

Next, numerical tests are conducted on a more heterogeneous set of plants with varying orders, relative degrees, pole locations and first nonzero Markov parameters. Specifically, a set of 50 stable, single-input, single-output plants with one nonminimum-phase zero located at 2 are randomly generated. For each of the 50 plants, the system order n is generated from the uniform discrete integer distribution on the interval $[2, 10]$, the relative degree d is generated from the uniform discrete integer distribution on the interval $[1, n - 1]$, the first

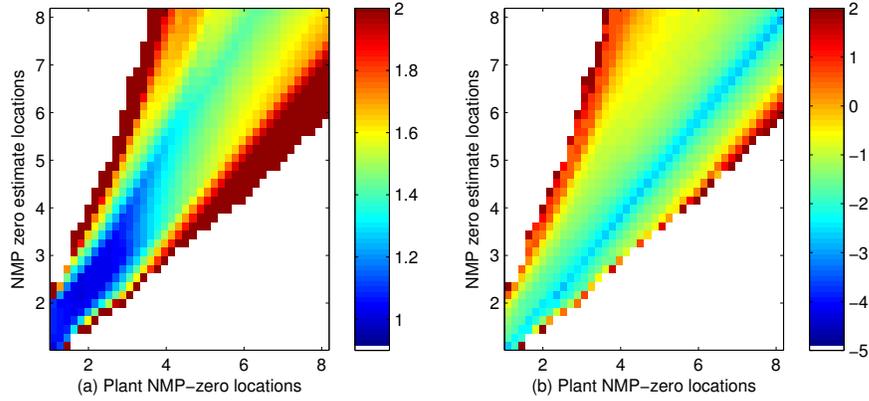


Figure 3.1: For 40 third order plants with a nonminimum-phase zero whose position is given by the x-axis, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1.1 to 8.1. The color in (a) shows logarithm of the transient performance, whereas in (b) the logarithm of the steady state performance is shown. For these 40 plants, the stability of the closed-loop system is less sensitive to overestimating the location of the nonminimum-phase zero than it is to underestimating it. In addition, the plants with NMP zeros further out on the real axis have greater stability margins than those with NMP zeros closer to 1.

nonzero Markov parameter H_d is generated from the uniform distribution on the interval $[-5, 5]$, the plant poles are generated in a pseudo-random manner and constrained to lie within the unit circle, and the plant zeros (other than the nonminimum-phase zero located at 2) are generated in a pseudo-random manner and constrained to lie within the unit circle.

For each of the 50 randomly generated plants, the adaptive controller requires estimates of the relative degree, the first nonzero Markov parameter, and the location of the nonminimum-phase zero. Next, for each of the 50 plants, we let the estimate of the nonminimum-phase zero (which is located at 2) vary from 1.4 to 3.8 and simulate the closed-loop system. Figure 3.2 is a histogram showing the percent of the 50 plants that have unstable closed-loop responses (i.e., the closed-loop performance is unbounded) as the estimate of the nonminimum-phase zero varies from 1.4 to 3.8. Note that all 50 plants result in stable closed-loop responses when the estimate of the nonminimum-phase zero is set at the actual value 2. Furthermore, Figure 3.2 indicates that the stability of the closed-loop system is less sensitive to overestimating the location of the nonminimum-phase zero

than it is to underestimating the location of the nonminimum-phase zero.

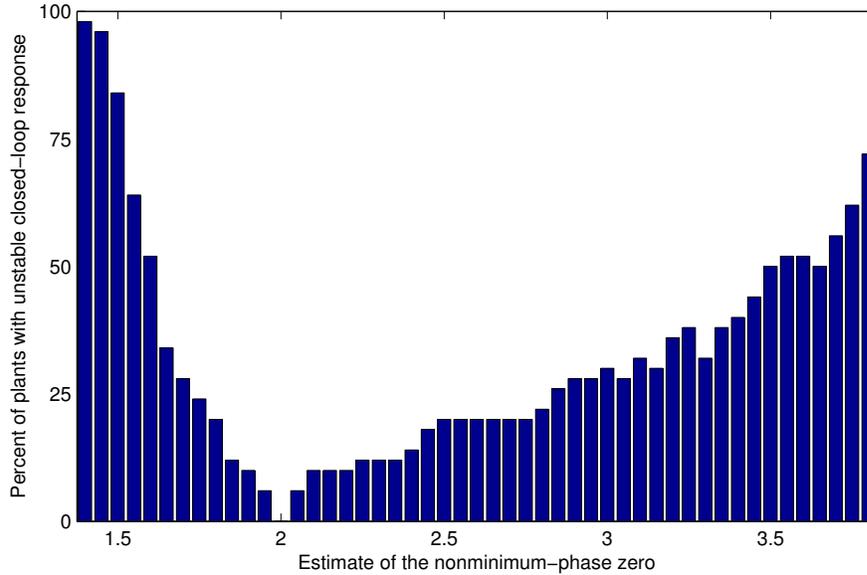


Figure 3.2: For 50 randomly generated plants, all of which have a nonminimum-phase zero at 2, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1.4 to 3.8. This histogram shows the percent of the 50 plants that have unstable closed-loop responses for different estimates of the nonminimum-phase zero. For these 50 random plants, the stability of the closed-loop system is less sensitive to overestimating the location of the nonminimum-phase zero than it is to underestimating the location of the nonminimum-phase zero.

Next, we explore the impact of system order and relative degree on the algorithm’s robustness to the accuracy of the nonminimum-phase zero estimate. We consider ten cases. Specifically we consider the cases where system order $n = 2, \dots, 6$ and relative degree $d = 1, \dots, n - 1$. For each case, we generate 50 random stable plants with fixed order and relative degree, and having a nonminimum-phase zero at 2. The plants are simulated with the adaptive controller in feedback, while letting the estimate of the nonminimum-phase zero vary from 1.4 to 2.6. The resulting stability histograms are shown in Figure 3.3. The algorithm tends to have better robustness for plants with lower order and higher relative degree. Accordingly, the most robust cases are situated on the diagonal of Figure 3.3.

Next, we explore, in more detail, the stability margins for second order systems. The

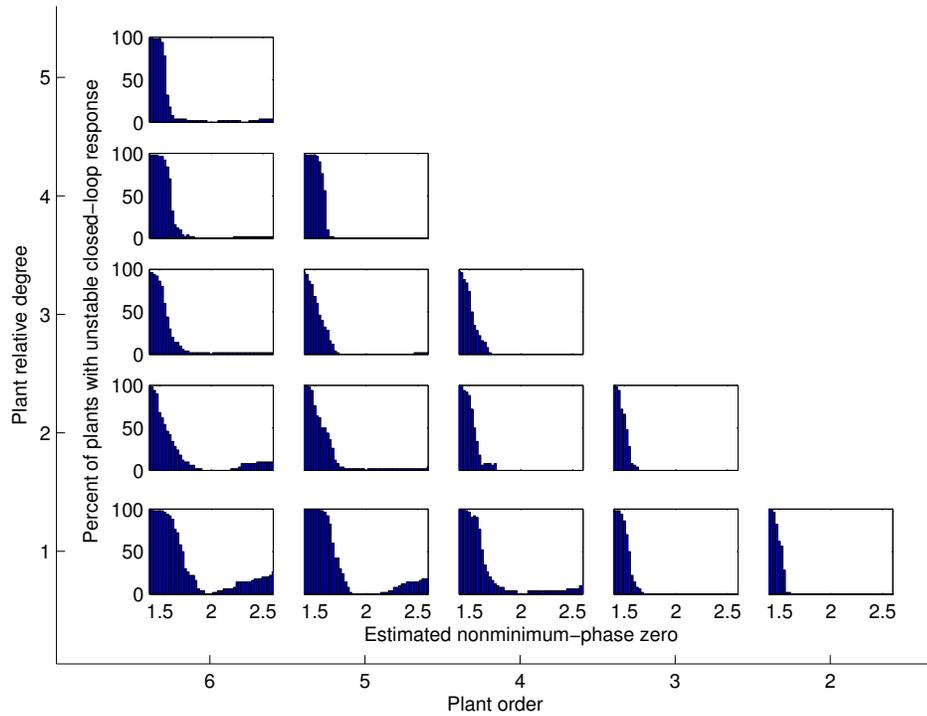


Figure 3.3: For 50 randomly generated plants with fixed order and relative degree (per subplot), all of which have a nonminimum-phase zero at 2, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1.4 to 2.6. These histograms show the percent of the 50 plants that have unstable closed-loop responses for different estimates of the nonminimum-phase zero. The robustness to the estimate of the nonminimum-phase zero decreases in rows from right to left (with increasing order), and in columns from top to bottom (with decreasing relative degree).

same 50 randomly generated $n = 2$, $d = 1$ plants from the previous experiment are simulated while the estimate of the nonminimum-phase zero is varied from 1.4 to 10. The stability histogram is shown in Figure 3.4, demonstrating that none of the 50 plants have upward stability margin in excess of 10 (which corresponds to 5 times the true value of the nonminimum-phase zero location). Note that similar finite stability margins can be observed for other values of n and d .

Now, we explore the effect of the plant pole locations on the algorithm's robustness to the accuracy of the nonminimum-phase zero estimate. A set of 61 second order plants with nonminimum-phase zero at 2, first Markov parameter equaling 1, and double real poles at

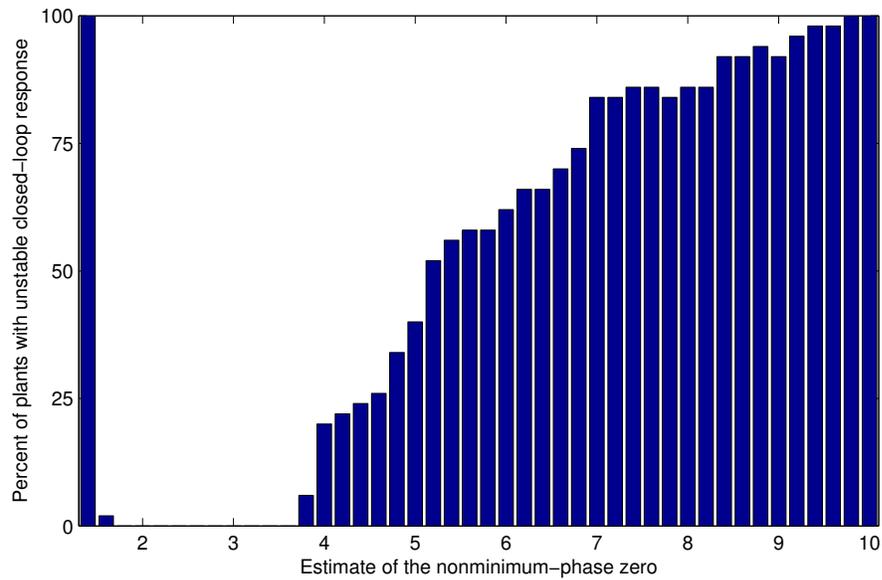


Figure 3.4: For 50 randomly generated plants, all of which have a nonminimum-phase zero at 2, the adaptive controller uses an estimate of the nonminimum-phase zero that varies from 1.4 to 10. For all 50 plants, there are finite bounds on the required accuracy of the nonminimum-phase zero.

equal increments ($1/30$) between -1 and 1 were generated. These plants are simulated with the adaptive controller in feedback, while the estimate of the nonminimum-phase zero is varied from 1 to 12 . The logarithm of the transient performance and logarithm of the steady state performance for each plant is plotted against the double-pole location, as shown in Figure 3.5. This figure indicates that plants with positive double poles are generally more robust than the ones with negative poles. However, there are plants with poles at -0.8 that have high upward stability margin, but the transient and residual error is much larger as compared to the nominal case. The plant with a double pole at -1 has the worst robustness to the accuracy of the nonminimum-phase zero estimate.

While the previous example looked at second order plants with double poles on the real axis, the following example extends these results to the case where a plant has two complex conjugate poles. A set of 2601 plants is generated, where each plant has two poles at $re^{\pm j\theta}$, where r is varied from 0 to 1 with 51 increments and θ is varied from 0

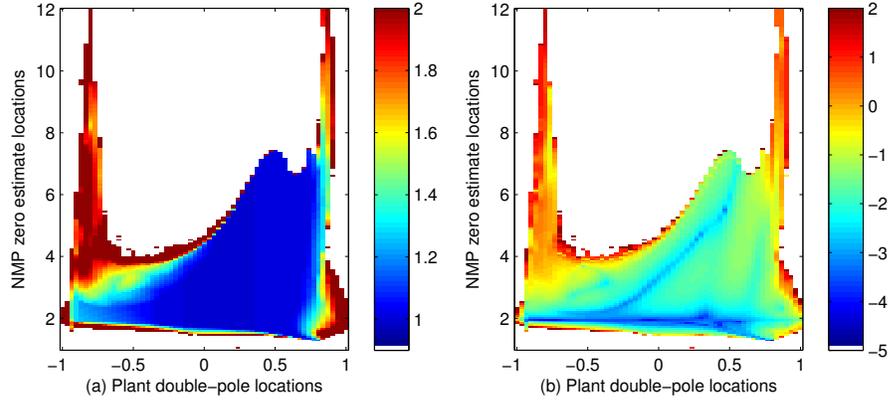


Figure 3.5: For 61 second order plants with double poles at locations indicated by the x-axis, all of which have a nonminimum-phase zero at 2, the adaptive controller uses an estimate of the nonminimum-phase zero that varies from 1 to 12 (y-axis). The color map represents the logarithm of transient performance and logarithm of steady state performance in subfigures (a) and (b), respectively. The adaptive controller (3.11)-(3.13) is more robust for plants with positive double-poles.

to π with 51 increments. These plants have a nonminimum zero at 2 and first Markov parameter equaling 1. Figures 3.7 and 3.6 show the upward and downward stability margins respectively. The color-coding corresponds to the estimate of the nonminimum-phase zero, where the closed-loop system becomes unstable. The plants with poles that have negative real parts are less robust than plants with poles that have positive real parts.

3.4.1.2 Sensitivity with two nonminimum-phase zeros

In this section, we explore the sensitivity of the adaptive controller (3.11)-(3.13) to the location of nonminimum-phase zero estimates for plants with two nonminimum-phase zeros. First, 50 stable plants with nonminimum-phase zeros at 2 and 3 are randomly generated with orders 3 through 10 and relative degrees 2 through $n - 2$. Furthermore, these plants are simulated in feedback with adaptive controller (3.11)-(3.13), while the nonminimum-phase zero estimates are varied from 1 through 4. Figure 3.8 shows the percent of the 50 plants that have unstable closed-loop responses as the estimates of the nonminimum-phase zeros vary. Figure 3.8 shows that the slope on the downward margin is steeper than the

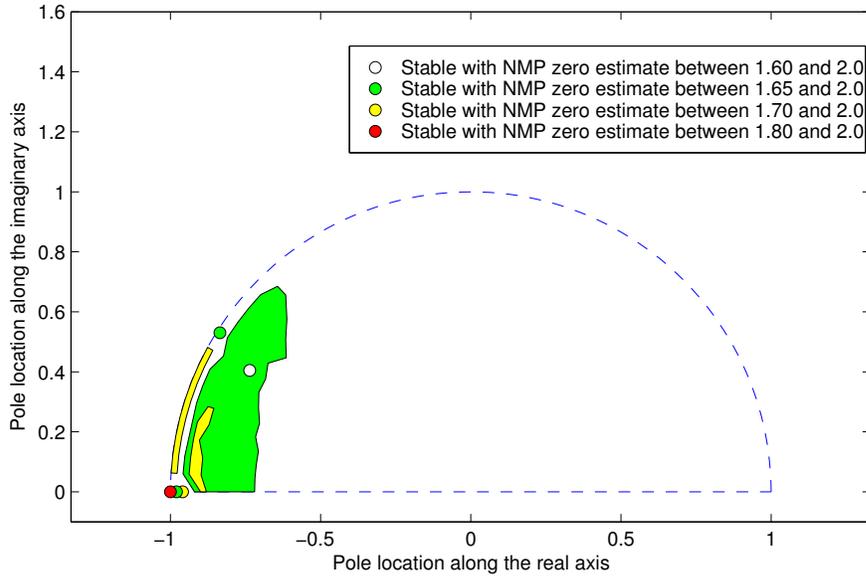


Figure 3.6: Downward stability margin. For 2601 second order plants with a nonminimum-phase zero at 2, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1.6 to 3.6. The colored regions represent plant-pole areas where the closed loop response was unstable at the estimate of the nonminimum-phase zero shown in the legend. Due to the nature of the complex variables this plot is symmetric, but only the top half is shown. The plants with poles whose real part is close to -1 are less robust to the underestimation of the nonminimum phase zero than the rest of the plants with poles inside the unit circle.

slope on the upward margin, suggesting that the performance of the algorithm with two nonminimum-phase zeros is similar to the single nonminimum-phase zero case. Figure 3.9 demonstrates that the algorithm is more robust to overestimating the location of the nonminimum-phase zero, than to underestimating it even in the two nonminimum-phase zeros case.

In the next example we look at the sensitivity of the algorithm to locations of complex conjugate nonminimum-phase zero estimates. For each of the 100 stable plants with nonminimum-phase zeros at $2 \pm 1j$, orders 3 through 10, relative degrees 2 through $n - 2$, the adaptive controller (3.11)-(3.13) is connected in feedback. These plants are simulated while the nonminimum-phase zero estimates are varied from 1 through $4 \pm 3j$. Resulting stability plot is shown in Figure 3.10. The color in this image corresponds to the percent of

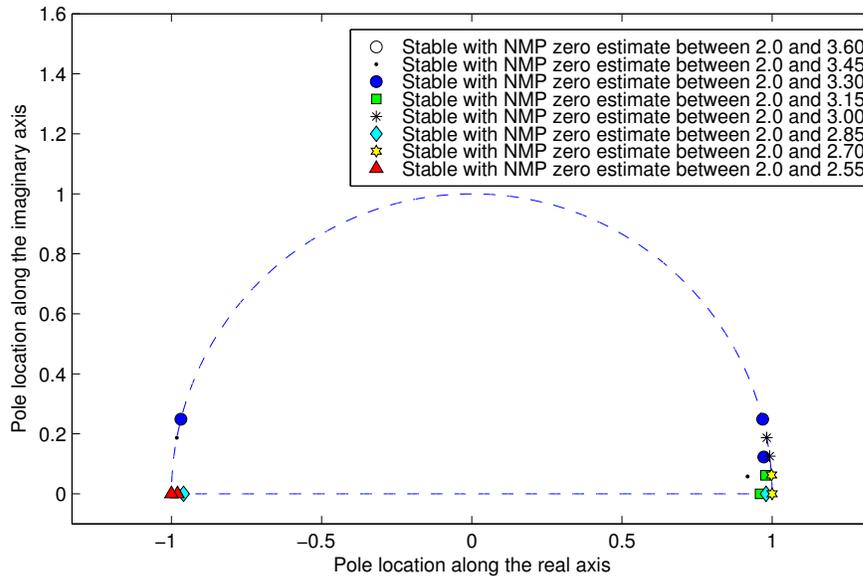


Figure 3.7: Upward stability margin. For 2601 second order plants with a nonminimum-phase zero at 2, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1.6 to 3.6. The colored regions represent plant-pole areas where the closed loop response was unstable at the estimate of the nonminimum-phase zero shown in the legend. Due to the nature of the complex variables this plot is symmetric, but only the top half is shown. The plants which are least robust to the overestimation of the nonminimum phase zero are the ones with poles close to ± 1 .

plants that have unstable closed loop performance. Next, the same 100 plants are simulated while the estimates of the nonminimum-phase zeros are varied in phase (but the magnitude is fixed) as shown in Figure 3.11 (d). Similar plots were generated for magnitude, real, and imaginary parts of the nonminimum-phase zero estimates and are shown in Figure 3.11 (a)-(c). From these plots, we conclude that the algorithm is more robust to underestimating phase than overestimating it, but is more robust to overestimating magnitude than underestimating it.

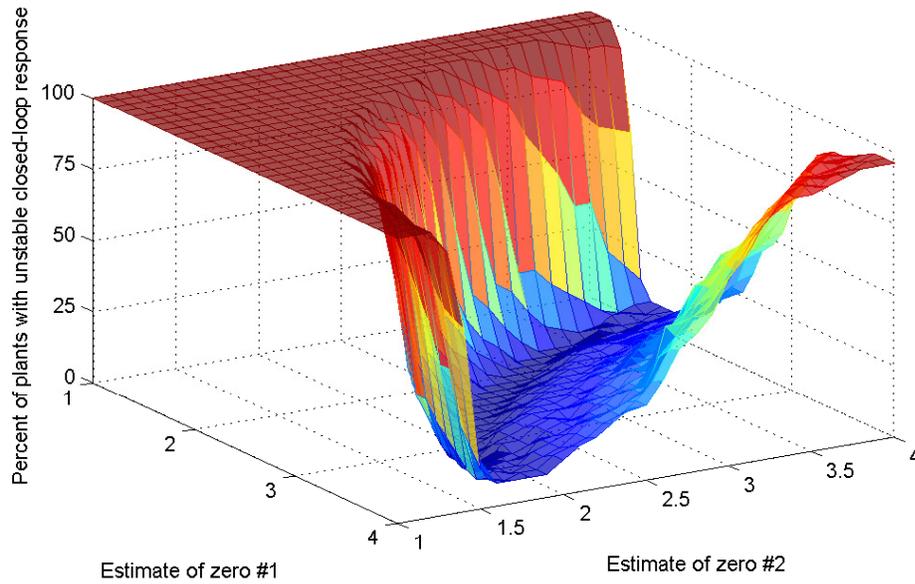


Figure 3.8: For 50 randomly generated plants, all of which have a nonminimum-phase zeros at 2 and 3, the adaptive controller (3.11)-(3.13) uses estimates of the nonminimum-phase zeros which vary from 1 to 4 each. The steep slope on the downward margin and shallower slope on the upward margin suggest that the performance of the algorithm in 2-nonminimum-phase-zeros case is similar to the single nonminimum-phase zero case, but this conjecture is further investigated in the following plots.

3.4.2 Sensitivity to the Accuracy of the Estimate of the First Nonzero Markov Parameter

In this experiment, we explore the sensitivity of the adaptive controller (3.11)-(3.13) to the estimate of the first nonzero Markov parameter. Accordingly, 100 stable plants of orders 2 through 6, relative degrees 1 through $n - 1$, with zeros uniformly distributed on interval [02] and with first nonzero Markov parameter equalling 1, were generated. Furthermore, these plants are simulated in feedback with adaptive controller (3.11)-(3.13), while the first nonzero Markov parameter estimates are varied from 0.1 to 10 and the exact locations of the nonminimum-phase zeros are supplied to the algorithm. Figure 3.12 shows the percent of plants with unstable closed loop response as the estimates vary. This histogram might look symmetric, but since it is plotted against a logarithmic scale it is indeed not symmetric when

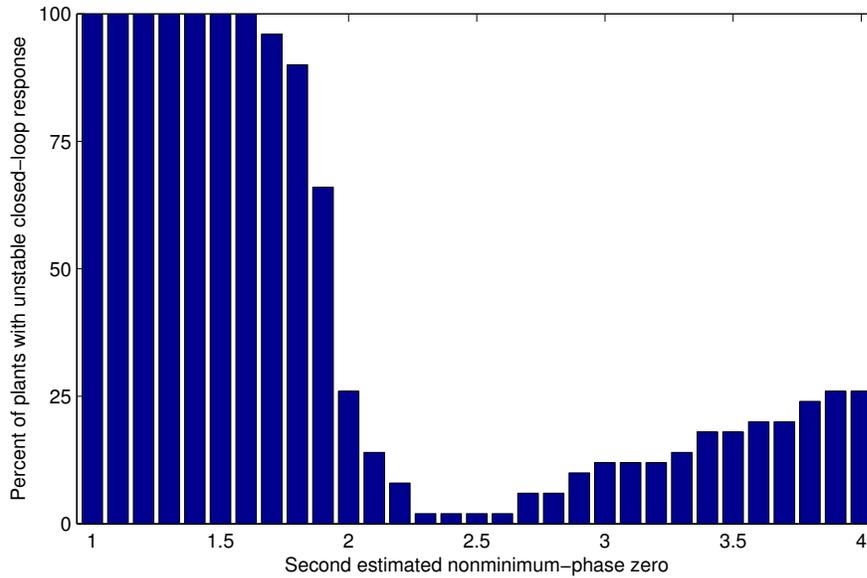


Figure 3.9: For 50 randomly generated plants, all of which have a nonminimum-phase zeros at 2 and 3, the adaptive controller (3.11)-(3.13) uses an estimate of the second nonminimum-phase zero that varies from 1 to 4. The estimate of the first nonminimum-phase zero is fixed at 3. Just like in the single nonminimum-phase zero case, the algorithm is more robust to overestimating the location of the nonminimum-phase zero, than to underestimating it.

viewed on the linear scale. Actually, the slope on the downward margin is steeper than the slope of the upward margin, suggesting that the algorithm is more robust to overestimating the first nonzero Markov parameter, than to underestimating it for both the minimum-phase and nonminimum-phase plants.

3.5 Sampled Data Systems

In this section, we investigate the sensitivity of the cumulative retrospective cost adaptive controller (3.11)-(3.13) to the increase in sampling rate beyond what will be defined next as minimum sampling frequency. We define minimum sampling frequency F_{min} in rad/sec as twice the magnitude of the pole furthest out from the origin. We test the adaptive controller's performance on a set of continuous nonminimum-phase stable plants discretized

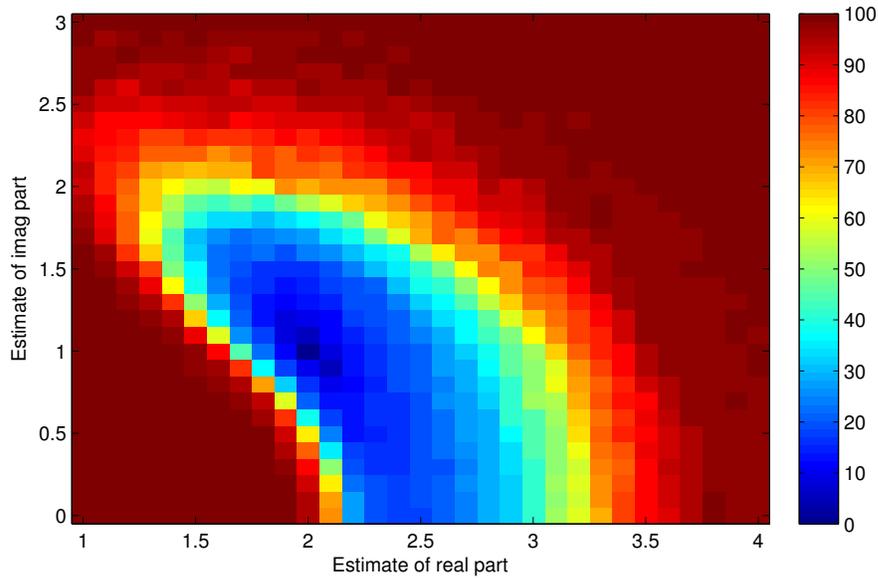


Figure 3.10: For 100 randomly generated plants, all of which have a nonminimum-phase zeros at $2 \pm 1j$, the adaptive controller (3.11)-(3.13) uses an estimate of the nonminimum-phase zero that varies from 1 to $4 \pm 3j$. The color corresponds to the percent of plants that have unstable closed loop performance.

at rates in interval $[F_{min} \ 10F_{min}]$. We vary the sampling rate to see if sampling at higher rates implies better performance. First, 100 random stable continuous plants with poles and minimum-phase zeros having magnitude less than 1, nonminimum-phase zero at 1 and first nonzero Markov Parameter equalling 2 are generated. Next, these plants are discretized using zero order hold and are simulated in feedback with the adaptive controller. Figure 3.13 shows logarithm of transient and steady state performances as a function of the plant and sampling frequency. The plants are sorted in the descending order of the transient performance in the nominal case. Figure 3.13 suggests that for certain plants as the sampling frequency is increased, the transient performance improves, whereas the steady state performance degrades.

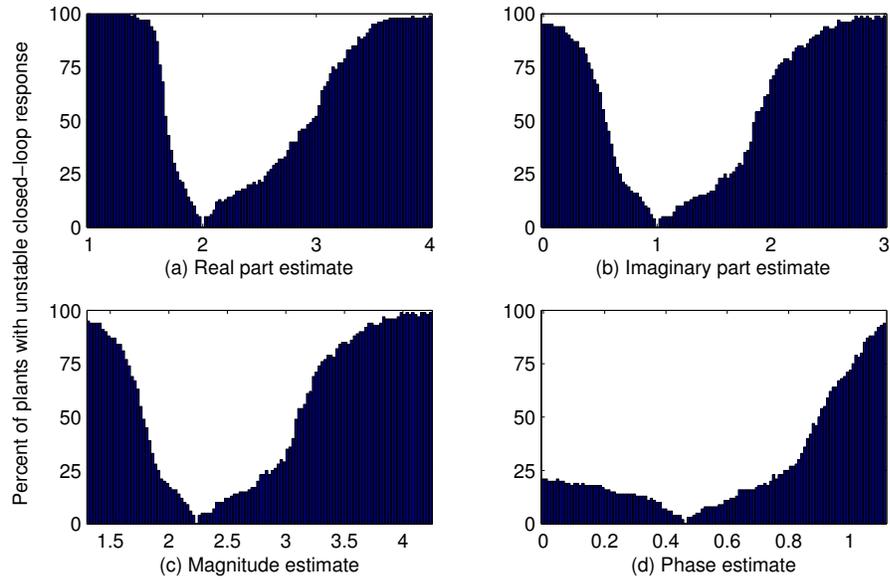


Figure 3.11: For 100 randomly generated plants, all of which have a nonminimum-phase zeros at $2 \pm 1j$, the adaptive controller (3.11)-(3.13) uses an estimate of the real part of the nonminimum-phase zero in (a), imaginary part in (b), magnitude in (c) and phase in (d). The algorithm seems to be equally robustness to changes in estimates of the real and imaginary parts as well as magnitude. Additionally, in this particular setup, the algorithm is more robust to underestimating phase than overestimating it.

3.6 Conclusion

In this chapter, we have presented a direct adaptive control algorithm for stabilization, command following, and disturbance rejection. The algorithm requires limited information about the open-loop system. More specifically, it requires knowledge of the first nonzero Markov parameter and the nonminimum-phase zeros from the control to the performance measurement. The adaptive law minimizes a cumulative retrospective cost function using a recursive least-squares algorithm. The algorithm has been shown to be effective through simulation.

We found out that the algorithm is more robust to overestimating the location of the nonminimum-phase zero in the case of plants with single NMP zero. Also, the order and relative degree of the plant do have an effect on the closed loop stability - the algorithm has

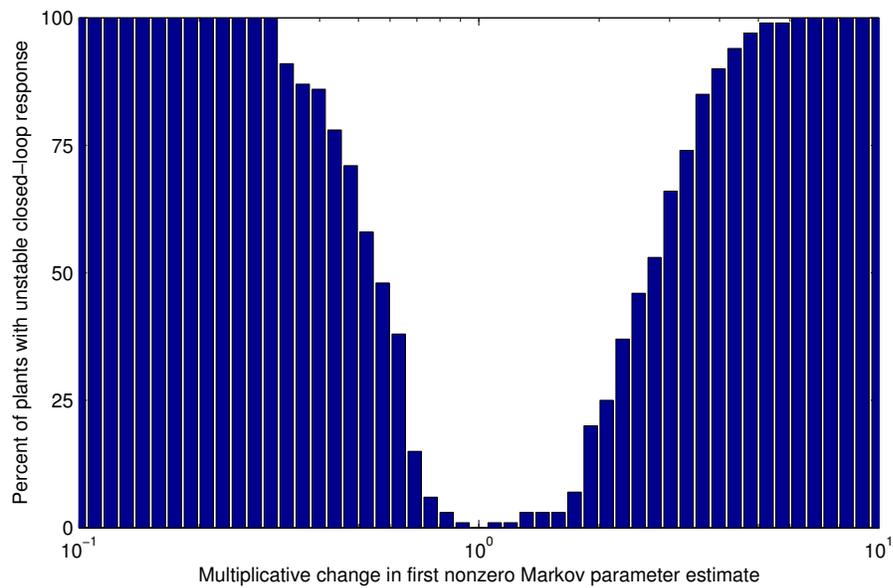


Figure 3.12: For 100 randomly generated stable plants, the adaptive controller (3.11)-(3.13) uses the scaled value of the first nonzero Markov parameter (given by the x-axis) and exact values for the nonminimum-phase zeros. The color scale in (a) corresponds to logarithm of the transient performance and in (b) it corresponds to logarithm of steady state performance. The lesser slope of the histogram on the right than on the left suggest that the algorithm is more sensitive to underestimating the first nonzero Markov parameter than to overestimating it.

greater stability margins for plants with low order and high relative degree. We also found that for plants with two real nonminimum-phase zeros, the algorithm is more robust to overestimating the location of either zero independently, but not overestimating both zeros at the same time. Additionally, for plants with two complex nonminimum-phase zeros, the algorithm was most forgiving for underestimating phase of the zero locations as opposed to overestimating it.

Next, the algorithm has been shown to have a certain amount of robustness to uncertainty in the first nonzero Markov parameter. It is more robust to overestimating the first nonzero Markov parameter, than underestimating it.

Lastly, we investigated the effect of increasing the sampling frequency for sampled data systems on the algorithm's performance. It was found that the algorithm performs well if

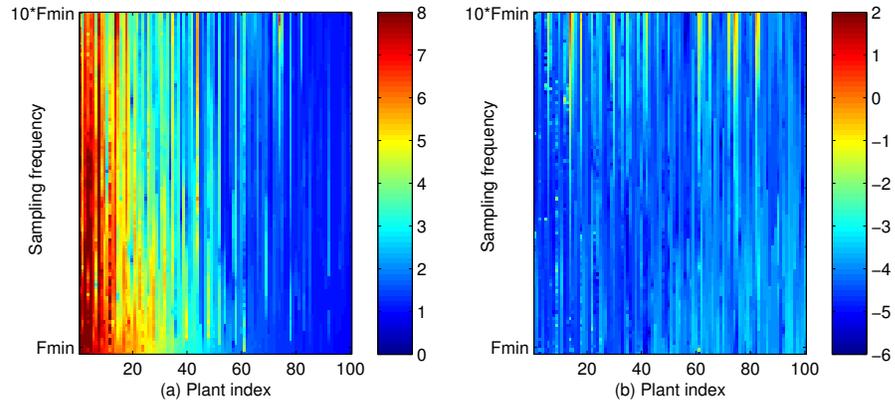


Figure 3.13: For 100 randomly generated plants, the adaptive controller (3.11)-(3.13) uses the exact value of the first nonzero Markov parameter and nonminimum-phase zeros. The color scale in (a) corresponds to logarithm of the transient performance and in (b) it corresponds to logarithm of steady state performance, as established in Section 3.4. The algorithm performs best at the sampling period corresponding to the minimum sampling frequency, as defined in section 3.5.

the used sampling frequency isn't greater than five times the minimum sampling frequency, as defined in section 3.5.

CHAPTER 4

Retrospective Adaptive Control of a Planar Multilink Arm with Nonminimum-Phase Zeros

This chapter addresses the problem of adaptive command following and disturbance rejection for a nonlinear planar multilink mechanism interconnected by torsional springs and dashpots. It considers a nonlinear multilink mechanism where a control torque is applied to the hub of the multilink mechanism, and the objective is to control the angular position of the tip, which is separated from the hub by N links. Accordingly, the nonlinear equations of motion for the N link mechanism are derived. Then, these equations of motion are linearized and it is demonstrated to have nonminimum-phase zeros when the control torque and angular position sensor are not colocated. A retrospective cost adaptive controller is used to control this mechanism, which is effective for nonminimum-phase systems if an estimate of the nonminimum-phase zeros is available. Both command following and disturbance rejection problems are considered, where the spectrum of the commands and disturbance are unknown. The results of this chapter are published in collaboration with Hoagg and Bernstein in [81] and are extended in [82].

4.1 Introduction

Nonminimum-phase zeros present a fundamental impediment to the achievable performance of a closed-loop system, limiting the bandwidth and, in the case of positive ze-

ros, causing initial undershoot or direction reversals under step inputs [83, p. 289], [84, 85]. Nonminimum-phase zeros are also challenging for adaptive control methods, which typically assume that the plant is minimum phase [86]. For discrete-time systems with nonminimum-phase zeros, the adaptive control method in [32,87] requires that the nonminimum-phase zero locations be known.

In view of these challenges, it is of interest to determine physical properties that give rise to nonminimum-phase zeros. It is known that the transfer function of a flexible structure with co-located force actuation and velocity sensing is positive real and thus minimum phase [88]. This property suggests that noncollocation is the underlying cause of nonminimum-phase zeros. It was shown in [89], however, that, for a string of translating masses interconnected by springs and dashpots, the noncollocated transfer functions between every pair of masses are minimum phase. Therefore, noncollocation per se is not the source of nonminimum phase zeros.

A vehicle with rear-wheel steering, or, equivalently, a car driving in reverse, exhibits initial undershoot in the sense that the driver initially moves in the direction that is opposite to the ultimate direction of motion. This example, as well as the examples in [90, 91], suggest that nonminimum phase zeros may arise from a combination of noncollocation and rotational motion.

In place of the translating masses considered in [89], we thus consider a planar multilink mechanism with rotating masses interconnected by torsional stiffnesses and dashpots. This mechanism can be viewed as a lumped approximation of a flexible rotating arm, whose dynamics and control are widely studied for applications such as space structures and hard drives [92, 93].

The multilink mechanism is nonlinear, and thus the derivation of its equations of motion is more complicated than the case of translating masses considered in [89], whose dynamics are linear. Analysis of the zeros of the rotating masses must therefore be based on a linearized model. A related analysis is given in [90].

For the linearized model of the rotating masses we show that the damping and stiffness matrices have the same form as in the case of translating masses. However, the key difference between the translational and rotational cases is the inertia matrix, which is diagonal for the translating masses but nondiagonal for the rotating masses. With this distinction in mind, the first objective of this chapter is to revisit the analysis of [89] and show how the off-diagonal entries of the inertia matrix for the rotating masses give rise to nonminimum-phase zeros.

Next, we consider adaptive control of the planar multilink mechanism using the approach of [32]. Since this method requires knowledge of the nonminimum-phase zeros, we assume that this information is available, either by analytical modeling or system identification [53]. We then apply the retrospective adaptive control algorithm of [32] on both the linearized and nonlinear system and assess the resulting performance for problems of command following and disturbance rejection.

4.2 Nonlinear Equations of Motion

In this section, we derive the nonlinear equations of motion for an N -link planar arm system by using Lagrange's equations. First, we define the parameters of the system. Let p_1 be the point where the first link is connected to the horizontal plane, and, for $n = 2, \dots, N$, let p_n be the point where the n^{th} link is connected to the $(n - 1)^{\text{th}}$ link. Next, for $n = 1, \dots, N$, let q_n be the center of mass of the n^{th} link. Furthermore, for $n = 1, \dots, N$, let m_n be the mass of the n^{th} link, let l_n be the length of the n^{th} link, let c_n be the damping at the joint p_n , let k_n be the stiffness of the joint p_n , and let $I_n \triangleq \frac{1}{12}m_n l_n^2$ be the moment of inertia of the n^{th} link about q_n .

Next, we define the inertial frame F_A with orthogonal unit vectors $(\hat{i}_A, \hat{j}_A, \hat{k}_A)$, where \hat{i}_A and \hat{j}_A lie in the plane of motion of the N -link planar arm, and \hat{k}_A is orthogonal to the plane of motion. For simplicity, we assume that the origin of F_A is located at p_1 . In addition, for

$n = 1, \dots, N$, let F_{B_n} be a frame attached to the n^{th} link. More specifically, F_{B_n} is a body-fixed frame which rotates as the n^{th} link rotates. For $n = 1, \dots, N$, let F_{B_n} have orthogonal unit vectors $(\hat{i}_{B_n}, \hat{j}_{B_n}, \hat{k}_{B_n})$, where \hat{i}_{B_n} is in the direction from p_1 to q_1 , \hat{j}_{B_n} is orthogonal to \hat{i}_{B_n} and in the plane of motion, and \hat{k}_{B_n} is orthogonal to the plane of motion. Note that, for all $n = 1, \dots, N$, $\hat{k}_{B_n} = \hat{k}_A$. Finally, for $n = 1, \dots, N$, let θ_n be the angle from \hat{i}_A to \hat{i}_{B_n} . The N -link planar arm is shown in Figure 4.1. To construct the Lagrangian for the N -

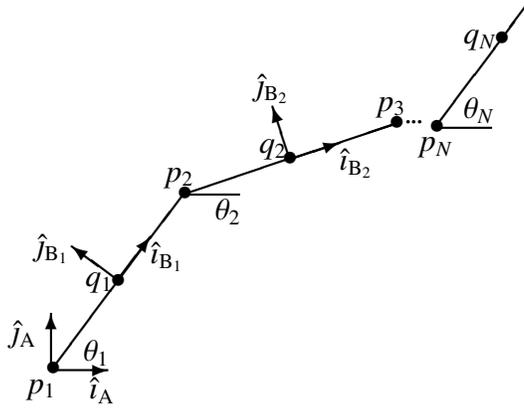


Figure 4.1: N -link planar arm system. All motion is in the horizontal plane.

link system, we derive expressions for kinetic and potential energies, and thus we require the translational and rotational velocities of each linkage. For $n = 1, \dots, N$, the rotational velocity of F_{B_n} with respect to F_A resolved in F_A is given by $\omega_n \triangleq \vec{\omega}_{B_n/A} \Big|_A = \begin{bmatrix} 0 & 0 & \dot{\theta}_n \end{bmatrix}^T$. Furthermore, for $n = 1, \dots, N$, the orientation matrix of F_{B_n} with respect to F_A is given by

$$\mathcal{O}_{B_n/A} = \begin{bmatrix} \cos(\theta_n) & \sin(\theta_n) & 0 \\ -\sin(\theta_n) & \cos(\theta_n) & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (4.1)$$

Next, for $n = 1, \dots, N$, let \vec{r}_{q_n/p_1} be the position vector from p_1 to q_n . For $n = 1, \dots, N$,

the velocity of q_n relative to p_1 with respect to F_A is given by

$$\vec{V}_{q_n/p_1/A} \stackrel{\Delta}{=} \overset{A\bullet}{\vec{r}}_{q_n/p_1} = \overset{A\bullet}{\vec{r}}_{q_n/p_n} + \sum_{i=1}^{n-1} \overset{A\bullet}{\vec{r}}_{p_{i+1}/p_i}, \quad (4.2)$$

where $\overset{A\bullet}{\vec{r}}$ denotes the derivative of \vec{r} taken in the frame F_A . Next, we apply the transport theorem to each term in (4.2), which yields

$$\begin{aligned} \vec{V}_{q_n/p_1/A} &= \left(\overset{B_n\bullet}{\vec{r}}_{q_n/p_n} + \vec{\omega}_{B_n/A} \times \vec{r}_{q_n/p_n} \right) \\ &\quad + \sum_{i=1}^{n-1} \left(\overset{B_i\bullet}{\vec{r}}_{p_{i+1}/p_i} + \vec{\omega}_{B_i/A} \times \vec{r}_{p_{i+1}/p_i} \right). \end{aligned}$$

Note that, for $n = 1, \dots, N$, \vec{r}_{q_n/p_n} is fixed relative to F_{B_n} , and thus $\overset{B_n\bullet}{\vec{r}}_{q_n/p_n} = 0$. Furthermore, note that for $i = 1, \dots, N-1$, \vec{r}_{p_{i+1}/p_i} is fixed relative to B_i , and thus $\overset{B_i\bullet}{\vec{r}}_{p_{i+1}/p_i} = 0$. Therefore,

$$\vec{V}_{q_n/p_1/A} = \vec{\omega}_{B_n/A} \times \vec{r}_{q_n/p_n} + \sum_{i=1}^{n-1} \vec{\omega}_{B_i/A} \times \vec{r}_{p_{i+1}/p_i}.$$

For $n = 1, \dots, N$, resolving $\vec{V}_{q_n/p_1/A}$ in frame F_A yields

$$\begin{aligned} \vec{V}_{q_n/p_1/A} \Big|_A &= \vec{\omega}_{B_n/A} \Big|_A \times \vec{r}_{q_n/p_n} \Big|_A \\ &\quad + \sum_{i=1}^{n-1} \vec{\omega}_{B_i/A} \Big|_A \times \vec{r}_{p_{i+1}/p_i} \Big|_A, \end{aligned} \quad (4.3)$$

where, for $n = 1, \dots, N$,

$$\vec{\omega}_{B_n/A} \Big|_A^\times = \begin{bmatrix} 0 & -\dot{\theta}_n & 0 \\ \dot{\theta}_n & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad (4.4)$$

$$\vec{r}_{q_n/p_n} \Big|_A = \mathcal{O}_{B_n/A} \begin{bmatrix} \frac{l_n}{2} & 0 & 0 \end{bmatrix}^T, \quad (4.5)$$

and, for $n = 1, \dots, N-1$,

$$\vec{r}_{p_{n+1}/p_n} \Big|_A = \mathcal{O}_{B_n/A} \begin{bmatrix} l_n & 0 & 0 \end{bmatrix}^T. \quad (4.6)$$

Furthermore, for $n = 1, \dots, N$, define $V_n \triangleq \left\| \vec{V}_{q_n/p_1/A} \Big|_A \right\|$. For demonstration, it follows from (4.3)-(4.6) that

$$\begin{aligned} \vec{V}_{q_1/p_1/A} \Big|_A &= \begin{bmatrix} 0 & -\dot{\theta}_1 & 0 \\ \dot{\theta}_1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \mathcal{O}_{B_1/A} \begin{bmatrix} \frac{l_1}{2} \\ 0 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} -\frac{1}{2}l_1 \sin(\theta_1)\dot{\theta}_1 \\ \frac{1}{2}l_1 \cos(\theta_1)\dot{\theta}_1 \\ 0 \end{bmatrix}, \end{aligned}$$

and thus $V_1 = \frac{1}{2}l_1\dot{\theta}_1$. Following this same procedure for $n \geq 2$, yields, for $n = 1, \dots, N$,

$$\begin{aligned} V_n &= \left[\frac{1}{4}l_n^2\dot{\theta}_n^2 + \sum_{i=1}^{n-1} (l_i^2\dot{\theta}_i^2 + l_n l_i \dot{\theta}_n \dot{\theta}_i \cos(\theta_i - \theta_n)) \right. \\ &\quad \left. + 2 \sum_{i \neq j}^{n-1} l_i l_j \dot{\theta}_i \dot{\theta}_j \cos(\theta_i - \theta_j) \right]^{1/2}. \end{aligned} \quad (4.7)$$

For $n = 1, \dots, N$, the kinetic energy of the n^{th} link is

$$\begin{aligned}
T_n &\triangleq \frac{1}{2}m_n V_n^2 + \frac{1}{2}I_n \|\omega_n\|^2 \\
&= \frac{m_n}{2} \left(\frac{1}{3}l_n^2 \dot{\theta}_n^2 + \sum_{i=1}^{n-1} (l_i^2 \dot{\theta}_i^2 + l_n l_i \dot{\theta}_n \dot{\theta}_i \cos(\theta_i - \theta_n)) \right. \\
&\quad \left. + 2 \sum_{i \neq j}^{n-1} l_i l_j \dot{\theta}_i \dot{\theta}_j \cos(\theta_i - \theta_j) \right), \tag{4.8}
\end{aligned}$$

and the total kinetic energy is defined by $T \triangleq \sum_{n=1}^N T_n$. Next, for $n = 1, \dots, N$, the potential energy of the n^{th} link is

$$U_n \triangleq \begin{cases} \frac{1}{2}k_1 \theta_1^2, & n = 1, \\ \frac{1}{2}k_n (\theta_{n-1} - \theta_n)^2, & n > 1, \end{cases} \tag{4.9}$$

and the total potential energy is defined by $U \triangleq \sum_{n=1}^N U_n$.

Thus, the Lagrangian for the N -link system is $L \triangleq T - U$. Next, for $n = 1, \dots, N$, let F_{c_n} be the dissipative torque resulting from the damping at joint p_n , that is,

$$F_{c_n} \triangleq \begin{cases} \frac{1}{2}c_1 \dot{\theta}_1^2, & n = 1, \\ \frac{1}{2}c_n (\dot{\theta}_{n-1} - \dot{\theta}_n)^2, & n > 1. \end{cases} \tag{4.10}$$

Furthermore, for $n = 1, \dots, N$, let u_n be an external torque applied at p_n . Therefore, for $n = 1, \dots, N$ the nonlinear equations of motion are given by

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{\theta}_n} - \frac{\partial L}{\partial \theta_n} + \frac{\partial F_{c_n}}{\partial \dot{\theta}_n} = u_n. \tag{4.11}$$

Now, we specialize to the case where $N = 2$. In this case, the Lagrangian is

$$\begin{aligned}
L &= \frac{1}{2}m_1 \left(\frac{1}{3}l_1^2 \dot{\theta}_1^2 \right) - \frac{1}{2}k_1 \theta_1^2 - \frac{1}{2}k_2 (\theta_1 - \theta_2)^2 \\
&\quad + \frac{1}{2}m_2 \left(\frac{1}{3}l_2^2 \dot{\theta}_2^2 + l_1^2 \dot{\theta}_1^2 + l_1 l_2 \dot{\theta}_1 \dot{\theta}_2 \cos(\theta_1 - \theta_2) \right), \tag{4.12}
\end{aligned}$$

and it follows from (4.11) and (4.12) that the equations of motion are given by

$$\begin{aligned}
u_1 = & \left(\frac{1}{3}m_1l_1^2 + m_2l_1^2\right)\ddot{\theta}_1 + \frac{1}{2}m_2l_1l_2 \sin(\theta_1 - \theta_2)\dot{\theta}_2^2 \\
& + \frac{1}{2}m_2l_1l_2 \cos(\theta_1 - \theta_2)\ddot{\theta}_2 + (k_1 + k_2)\theta_1 - k_2\theta_2 \\
& + (c_1 + c_2)\dot{\theta}_1 - c_2\dot{\theta}_2,
\end{aligned} \tag{4.13}$$

$$\begin{aligned}
u_2 = & \left(\frac{1}{3}m_2l_2^2\right)\ddot{\theta}_2 - \frac{1}{2}m_2l_1l_2 \sin(\theta_1 - \theta_2)\dot{\theta}_1^2 \\
& + \frac{1}{2}m_2l_1l_2 \cos(\theta_1 - \theta_2)\ddot{\theta}_1 - k_2\theta_1 + k_2\theta_2 \\
& - c_2\dot{\theta}_1 + c_2\dot{\theta}_2.
\end{aligned} \tag{4.14}$$

4.3 Linearized Equations of Motion

In this section, we derive linearized equations of motion for N -link system. First, we linearize the equations of motion for the two-link case. Then, we linearize the equations of motion for the three-link case. Finally, we generalize the linear equations of motion to the N -link case.

First, define

$$\Theta \triangleq \begin{bmatrix} \theta_1 & \dots & \theta_N \end{bmatrix}^T, \quad \Upsilon \triangleq \begin{bmatrix} u_1 & \dots & u_N \end{bmatrix}^T.$$

We linearize about the $(\Theta, \dot{\Theta}) \equiv 0$ equilibrium. Note that if, for all $n = 1, \dots, N$, $k_n > 0$, then $(\Theta, \dot{\Theta}) \equiv 0$ is the only unforced equilibrium of the N -link system. Let $\delta\Theta$ be the linear approximation of Θ around the equilibrium $(\Theta, \dot{\Theta}) \equiv 0$. To obtain the linearization, we use the small angle approximations $\sin(\theta_1 - \theta_2) \approx \delta\theta_1 - \delta\theta_2$, $\cos(\theta_1 - \theta_2) \approx 1$.

Linearizing the two-link system, with nonlinear equations of motion (4.13) and (4.14), about $(\Theta, \dot{\Theta}) \equiv 0$ yields

$$M\delta\ddot{\Theta} + C_d\delta\dot{\Theta} + K\delta\Theta = \Upsilon, \tag{4.15}$$

where

$$M \triangleq \begin{bmatrix} (\frac{m_1}{3} + m_2)l_1^2 & \frac{m_2}{2}l_2l_1 \\ \frac{m_2}{2}l_1l_2 & \frac{m_2}{3}l_2^2 \end{bmatrix},$$

$$C_d \triangleq \begin{bmatrix} c_1 + c_2 & -c_2 \\ -c_2 & c_2 \end{bmatrix}, K \triangleq \begin{bmatrix} k_1 + k_2 & -k_2 \\ -k_2 & k_2 \end{bmatrix}.$$

Similarly, linearizing the three-link system about $(\Theta, \dot{\Theta}) \equiv 0$ yields (4.15), where

$$M \triangleq \begin{bmatrix} (\frac{m_1}{3} + m_2 + m_3)l_1^2 & (\frac{m_2}{2} + m_3)l_2l_1 & \frac{m_3}{2}l_1l_3 \\ (\frac{m_2}{2} + m_3)l_1l_2 & (\frac{m_2}{3} + m_3)l_2^2 & \frac{m_3}{2}l_2l_3 \\ \frac{m_3}{2}l_1l_3 & \frac{m_3}{2}l_2l_3 & \frac{m_3}{3}l_3^2 \end{bmatrix},$$

$$C_d \triangleq \begin{bmatrix} c_1 + c_2 & -c_2 & 0 \\ -c_2 & c_2 + c_3 & -c_3 \\ 0 & -c_3 & c_3 \end{bmatrix},$$

$$K \triangleq \begin{bmatrix} k_1 + k_2 & -k_2 & 0 \\ -k_2 & k_2 + k_3 & -k_3 \\ 0 & -k_3 & k_3 \end{bmatrix}.$$

Finally, extending this technique, we obtain the linearization for the N -link system,

which is given by (4.15), where

$$\begin{aligned}
 M &\triangleq \begin{bmatrix} \gamma_{1,1} & \cdots & \gamma_{1,N} \\ \vdots & \ddots & \vdots \\ \gamma_{N,1} & \cdots & \gamma_{N,N} \end{bmatrix}, \\
 C_d &\triangleq \begin{bmatrix} c_1 + c_2 & -c_2 & 0 & \cdots & 0 \\ -c_2 & c_2 + c_3 & -c_3 & \cdots & 0 \\ 0 & -c_3 & c_3 + c_4 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & c_N \end{bmatrix}, \\
 K &\triangleq \begin{bmatrix} k_1 + k_2 & -k_2 & 0 & \cdots & 0 \\ -k_2 & k_2 + c_3 & -k_3 & \cdots & 0 \\ 0 & -k_3 & k_3 + k_4 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & k_N \end{bmatrix},
 \end{aligned}$$

where, for $g = 1, \dots, N$,

$$\gamma_{g,g} \triangleq \left(\frac{m_g}{3} + \sum_{i=g+1}^N m_i \right) l_g^2, \quad (4.16)$$

for $g = 1, \dots, N$ and $h = g + 1, \dots, N$,

$$\gamma_{g,h} \triangleq \left(\frac{m_h}{2} + \sum_{i=h+1}^N m_i \right) l_g l_h, \quad (4.17)$$

and, for $g, h = 1, \dots, N$, $\gamma_{h,g} = \gamma_{g,h}$.

4.4 Nonminimum-Phase Zero for the N -Link System

In this section, we prove that, for the two-link system, the linear transfer function from u_1 to $\delta\theta_2$ has one nonminimum phase zero. In fact, this transfer function has one positive zero. For the N -link system, we numerically demonstrate that the linear transfer function from u_1 to $\delta\theta_N$ (i.e., from the hub to the tip of the multilink mechanism) has $N - 1$ nonminimum-phase zeros.

For the N -link system, the linearized equations of motion (4.15) can be written as

$$\begin{bmatrix} \delta\dot{\Theta} \\ \delta\ddot{\Theta} \end{bmatrix} = A \begin{bmatrix} \delta\Theta \\ \delta\dot{\Theta} \end{bmatrix} + B\Upsilon, \quad (4.18)$$

where

$$A \triangleq \begin{bmatrix} 0_{N \times N} & I_N \\ -M^{-1}K & -M^{-1}C_d \end{bmatrix}, B \triangleq \begin{bmatrix} 0_{N \times N} \\ M^{-1} \end{bmatrix}.$$

Next, for $n = 2, \dots, N$, the transfer function from u_1 to $\delta\theta_n$ is given by

$$G_n(s) \triangleq \frac{\delta\theta_n(s)}{u_1(s)} = C_n(sI_N - A)^{-1}B_1, \quad (4.19)$$

where

$$C_n \triangleq \begin{bmatrix} 0_{1 \times n-1} & 1 & 0_{1 \times 2N-n} \end{bmatrix}, B_1 \triangleq B \begin{bmatrix} 1 \\ 0_{1 \times N-1} \end{bmatrix}.$$

For the two-link case (i.e., $N = 2$), the transfer function from u_1 to $\delta\theta_2$ can be expressed as

$$G_2(s) = \frac{\delta\theta_2(s)}{u_1(s)} = \frac{a_2s^2 + a_1s + a_0}{b_4s^4 + b_3s^3 + b_2s^2 + b_1s + b_0},$$

where the coefficients $a_0, \dots, a_2, b_0, \dots, b_4$ depend on the physical parameters of the system. More specifically, the numerator coefficients of $G_2(s)$ are given by $a_2 = -18l_1l_2m_2, a_1 = 36c_2, a_0 = 36k_2$. Since the zeros of $G_2(s)$ are the roots of the quadratic polynomial $a_2s^2 + a_1s + a_0$, we can solve for these roots expressed in the physical parameters of the system. More specifically, the quadratic polynomial $a_2s^2 + a_1s + a_0$ has the roots $z_{c,1} = \frac{c_2 + \sqrt{c_2^2 + 2k_2l_1l_2m_2}}{l_1l_2m_2}$ and $z_{c,2} = \frac{c_2 - \sqrt{c_2^2 + 2k_2l_1l_2m_2}}{l_1l_2m_2}$. Since the physical parameters l_1, l_2, m_2, c_2 , and k_2 are positive, it follows that $z_{c,1}$ is positive and $z_{c,2}$ is negative. Thus, we conclude that $G_2(s)$ has one nonminimum-phase zero.

For the N -link case, where $N > 2$, we conduct a numerical study to investigate the properties of the zeros of the transfer function from u_1 to $\delta\theta_N$. In particular, we let $N = 3, \dots, 10$, and for each value of N , we randomly generate 10,000 multilink systems. For each of the multilink systems, the masses m_1, \dots, m_N , the stiffnesses k_1, \dots, k_N , the damping coefficients c_1, \dots, c_N , and the lengths l_1, \dots, l_N are sampled from a uniformly generated random variable on the interval $(0, 100]$. Next, we compute the linearized transfer function $G_N(s)$ from u_1 to $\delta\theta_N$. For $N = 3, \dots, 10$, all 10,000 randomly generated multilink systems have $N - 1$ nonminimum-phase zeros in the transfer function $G_N(s)$. In fact, all of the randomly generated multilink systems have $N - 1$ positive zeros in the transfer function $G_N(s)$. Future work will include a proof of the conjecture that, for an N -link system, the linearized transfer function $G_N(s)$ from the control torque at the hub to the angular position of the N^{th} link has $N - 1$ positive zeros.

Next, we discretize $G_2(s)$ using a zero-order hold on the inputs. For this example, we consider the system parameters given by $m_1 = 2$ kg, $m_2 = 1$ kg, $l_1 = 3$ m, $l_2 = 2$ m, $k_1 = 7 \frac{N-m}{rad}$, $k_2 = 5 \frac{N-m}{rad}$, $c_1 = 10 \frac{kg-m^2}{rad}$, and $c_2 = 1 \frac{kg-m^2}{rad}$.

Discretizing $G_2(s)$ using a zero-order hold on the inputs results in a discrete-time transfer function, which also has one nonminimum-phase zero. The location of this nonminimum-phase zero depends on the sampling time used for the discretization. The discrete-time nonminimum-phase zero of $G_2(z)$ with the system parameters above and sampled at a rate

of 20Hz is located at approximately 1.08. Furthermore, note that the discrete-time system has one zero, which results from discretization effects. In this case, the sampled-data zero is located at -0.94 .

4.5 Numerical Examples

In this section, we use the retrospective cost adaptive controller (3.11)-(3.13) to control the linearized and nonlinear two-link system. In particular, we consider both the command following and disturbance rejection problems for the linearized and nonlinear two-link system. We assume that u_1 is the only available control input. We consider the two-link system with parameters shown in Section 4.4. The adaptive controller (3.11)-(3.13) is implemented in feedback at 20Hz with $\lambda = 0.99$, $R = 1$, $n_c = 8$, $P(0) = 10^{16}I_{16}$, and $\theta_c(0) = 0$. These values were found experimentally to yield the best performance for all cases with command amplitude of no more than 0.4 rad. Additionally, for each example, the system is allowed to run open-loop for 7.5 seconds and then the adaptive controller is turned on.

First, numerical simulations are performed using the linearized and nonlinear two-link system to assess the adaptive control's performance on a command following problem. The control objective is for θ_2 to track a 0.8Hz sinusoid with a magnitude of 0.3rad. We assume that the relative degree d and the first nonzero Markov parameter are known, that is, we let $\nu = d + 1$ and $\bar{\beta}_d = H_d$. In this example, $d = 1$ and $H_d = -0.00032$. In addition, we assume that the location of the nonminimum-phase zero is known, but no other information about the system is assumed to be known. Figure 4.2 shows that the adaptive controller drives performance variable z to zero.

Next, we implement the adaptive controller in feedback with the nonlinear plant, using the estimate of the nonminimum-phase zero obtained from the linearized two-link system. Figure 4.3 shows that the adaptive controller drives the performance variable z toward zero, and the performance is comparable to the linear case shown in Figure 4.2. We simulated the

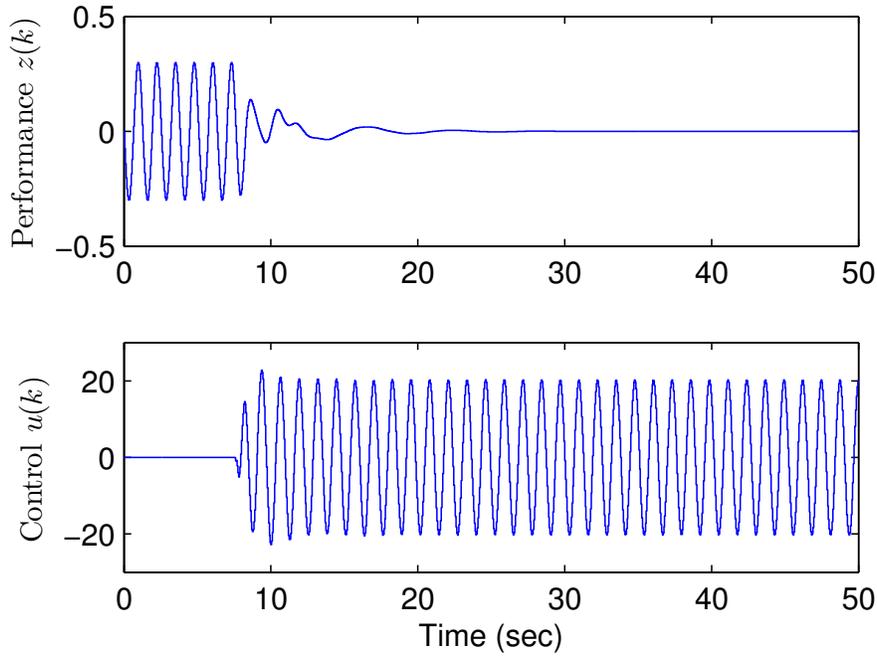


Figure 4.2: *Command following for the linearized two-link system:* The adaptive control (3.11)-(3.13) uses knowledge of the nonminimum-phase zero to track a sinusoid with unknown frequency and amplitude. The adaptive control is turned on after 7.5 seconds and drives the performance to zero.

nonlinear two-link system with physical parameters given in Section 4.4 and the adaptive controller in feedback for various command amplitudes, and we found that the adaptive controller is able to drive z toward zero for all command amplitudes less than 0.4 rad (or 23 degrees). Command amplitudes greater than 0.4 rad resulted in unacceptable controller performance.

Next, we consider the disturbance rejection problem, where the control objective is to drive θ_2 to zero, while a 1.6Hz sinusoidal disturbance is applied at both p_1 and p_2 . The magnitudes of the disturbances at p_1 and p_2 are 0.2 rad and 0.4 rad, respectively. We assume that the relative degree d , the first nonzero Markov parameter, and the location of the nonminimum-phase zero are known, but no other information about the system is assumed to be known. Figure 4.4 shows that the adaptive controller is able to reject the disturbance from θ_2 , and thus drives z to zero.

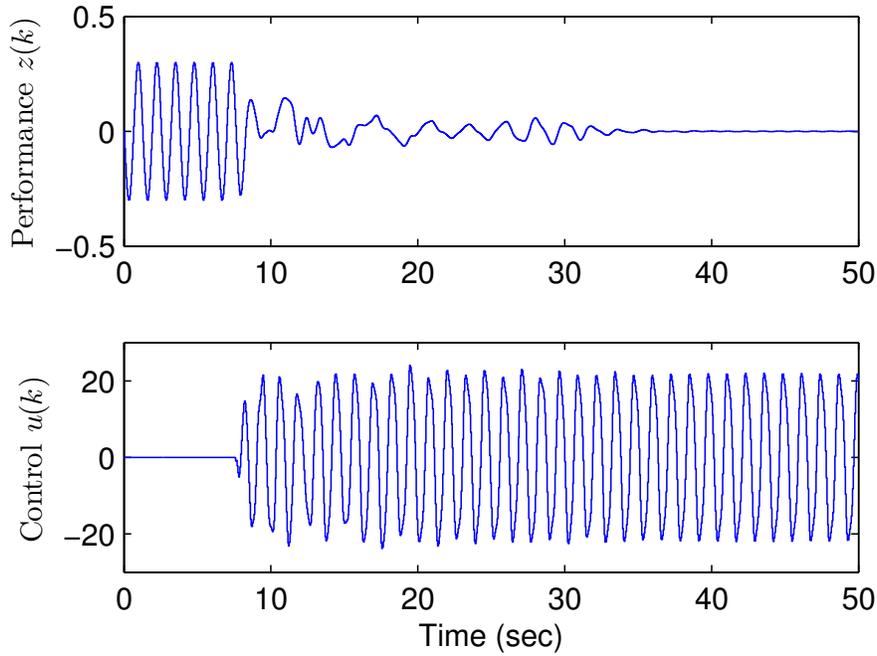


Figure 4.3: *Command following for the nonlinear two-link system:* The adaptive control (3.11)-(3.13) uses knowledge of the linearized nonminimum-phase zero to track a sinusoid with unknown frequency and amplitude. The adaptive control is turned on after 7.5 seconds and drives the performance to zero. The performance with the nonlinear system is comparable to the linear case shown in Figure 4.2.

Next, we implement the adaptive controller in feedback with nonlinear plant, using the estimate of the nonminimum-phase zero obtained from the linearized two-link system. Figure 4.5 shows that the adaptive controller drives z toward zero, and the performance is comparable to the linear case shown in Figure 4.4.

4.6 Conclusion

In this chapter, we investigated a nonlinear planar multilink mechanism that is interconnected by torsional springs and dashpots. More specifically, we considered a nonlinear multilink mechanism where a control torque is applied to the hub of the multilink mechanism, and the objective is to control the angular position of the tip, which is separated from the hub by N links. We derived the nonlinear equations of motion, linearized these

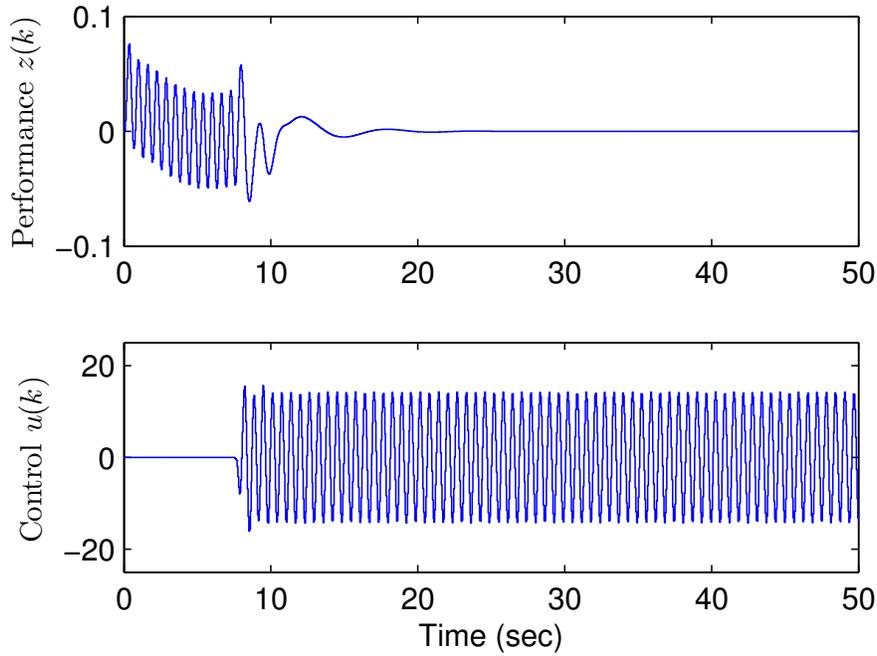


Figure 4.4: *Disturbance rejection for the linearized two-link system:* The adaptive control (3.11)-(3.13) uses knowledge of the nonminimum-phase zero to reject an unknown sinusoidal disturbance acting on both joints of the two-link mechanism. The adaptive control is turned on after 7.5 seconds and drives the performance to zero.

equations of motion, and conjectured that the linear transfer function from the hub to the tip of the multilink mechanism has $N - 1$ nonminimum-phase zeros. Finally, we implemented a retrospective cost adaptive controller [32] to control the multilink mechanism. We demonstrated both command following and disturbance rejection where commands and disturbances had unknown spectra. One possible application of the findings of this chapter is disturbance rejection as applied to the flexible ABB robotic arm [94, 95].

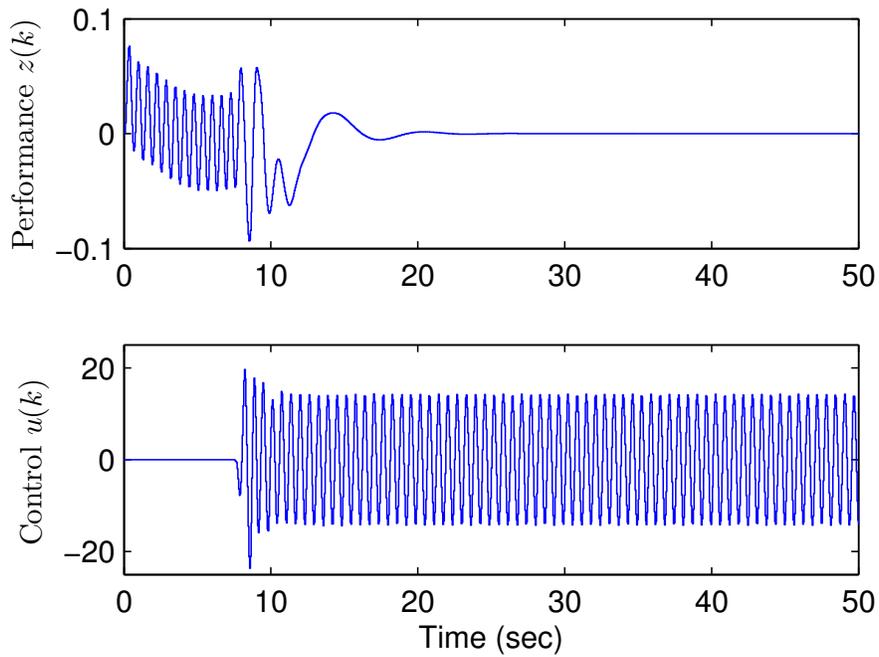


Figure 4.5: *Disturbance rejection for the nonlinear two-link system:* The adaptive control (3.11)-(3.13) uses knowledge of the nonminimum-phase zero to reject an unknown sinusoidal disturbance acting on both joints of the two-link mechanism. The adaptive control is turned on after 7.5 seconds and drives the performance to zero. The performance with the nonlinear system is comparable to the linear case shown in Figure 4.4.

CHAPTER 5

Assimilating Neutral Density Measurements into the Global Ionosphere-Thermosphere Model using the Ensemble Adjustment Kalman Filter

This chapter considers another nonlinear system (a space weather model) in the context of command following (or, more broadly, state estimation or data assimilation with driver estimation problem). In particular, *in situ* density measurements from a particular satellite (CHAMP) are used as a desired trajectory to be tracked by GITM's thermospheric density output. To accomplish this goal, both GITM states and inputs are estimated thereby giving the method more flexibility than estimating just the states or just the inputs. The results of this chapter are in preparation to be published in collaboration with Ridley, Bernstein, Collins, Hoar, and Anderson in [96].

5.1 Introduction

Solar radiation is the largest source of energy in the thermosphere, which can significantly affect the neutral mass density, ρ , which in turn influences the drag experienced by objects in low-Earth orbit. Uncertainty in drag translates to uncertainty in position that may even result in a loss of spacecraft [97]. One way of decreasing drag uncertainty is by obtaining more precise estimates of the neutral density from thermospheric models. An immediate

and cost effective way of improving accuracy of these models is data assimilation.

For example, assimilation of azimuth, elevation, range and range-rate measurements from the calibration satellites into the high-accuracy satellite drag model (HASDM) resulted in about 25% bias reduction as compared to the case with no data assimilation (*no-DA*) [97].

Using an alternative data source, [37] showed that assimilating Challenging Minisatellite Payload (CHAMP, [98]) and Gravity Recovery and Climate Experiment (GRACE, [99]) mass density measurements resulted in about 50% bias reduction for the Coupled Thermosphere Ionosphere Plasmasphere Electrodynamics (CTIPe, [100]) model, as compared to the no-DA case. A modified Kalman filter was used to perform data assimilation with a prior error covariance matrix constructed using empirical orthogonal functions with a maximum-likelihood update method.

CTIP [101] was also used by [102], which interpreted the results of radio tomographic imaging performed on the U.S. Navy Ionospheric Measuring System and compared these results with GPS total electron content (TEC) measurements. The authors concluded that, while the small-scale details within the data were beyond the capabilities of the model, the general features captured by the model aid the interpretation of the tomography results.

Similarly, Utah State University's Global Assimilation of Ionospheric Measurements (USU GAIM, [103, 104]) framework also demonstrated how initial model bias can be removed using an approximate Kalman filter. The Kalman filter was used to assimilate the slant total electron content (TEC) measurements and ionosonde density profiles into the ionosphere plasmasphere model (IPM) to estimate one of the model drivers, namely, the equatorial vertical plasma drift. A reduced state Kalman filter was implemented through numerical linearization of reduced-state IPM at each time step. The authors concluded that, while this linearization might require thousands of model runs, each model run is parallel and computationally efficient.

Likewise, the performance of a low-resolution version of the Jet Propulsion Labora-

tory/University of Southern California Global Assimilation Ionospheric Model (JPL/USC GAIM) is described by [105]. The GPS slant TEC measurements from 200 ground GPS sites were used to estimate the single-ion 3D density state using a band-limited Kalman filter approximation. The results of the assimilation were validated against the withheld ground measurements, and the vertical TEC measurements from the JASON satellite [106]. The study concluded that, when apparent bias in the JASON measurements was removed, the GAIM performance was improved from 7 total electron content units (TECU) to about 5.3 TECU.

An application of data assimilation for the Global Ionosphere-Thermosphere Model (GITM) is described by [107,108]. In particular, [107] uses the localized unscented Kalman filter (LUKF) to assimilate electron number density, ion temperature, and ion velocity into the one-dimensional (1D) GITM, described in detail in [33]. That study concluded that using a small localization region (about 11 altitude cells out of 50) in LUKF is comparable to the nominal unscented Kalman filter (UKF) performance without the added computational cost. On the other hand, [108] demonstrates how the GITM data assimilation setup can be extended to include three-dimensional (3D) GITM using the ensemble Kalman filter (EnKF, [41]). More precisely, 7 ensemble members were used in the EnKF to assimilate the simulated electron number density, ion temperature, neutral density, and (separately) TEC at six fixed locations.

The present work extends the preliminary findings of [108] by increasing the ensemble size and adding driver estimation. Driver estimation is sought to improve bias removal in the GITM neutral mass density estimates as compared to the CHAMP measurements. Following the example of [37], the GRACE measurements are used to validate the results, though in this case the GRACE neutral density measurements are not scaled and are used at the altitude of measurements (about 498 km). The CHAMP measurements are assimilated into GITM via the ensemble adjustment Kalman filter (EAKF, [46]). EAKF was chosen over the EnKF as it does not perturb measurements but otherwise provides similar per-

formance. The EAKF, implemented in the Data Assimilation Research Testbed (DART), which is written and maintained by National Center for Atmospheric Research [47], is used in this study. DART includes the EnKF and EAKF, as well as several other filters and has been used in numerical weather prediction [109, 110], aerosol modeling systems [111], and upper atmosphere modeling [112].

This chapter is organized as follows. Section 5.2.1 introduces GITM and demonstrates the need for data assimilation by highlighting the mismatch between ρ produced by GITM and ρ measured by CHAMP when bias is added to GITM. EAKF is introduced independently of GITM in sections 5.2.2-5.2.4 to separate general data assimilation concepts from model-dependent concepts. Section 5.2.5 addresses localization. The results of performing data assimilation with GITM are shown in section 5.3.

5.2 Methodology

5.2.1 Global Ionosphere-Thermosphere Model

GITM is a global model of the upper atmosphere, whose state variables include neutral and ion densities, temperatures, and velocities, as well as electron density and temperature. One notable feature of GITM is that it produces nonhydrostatic solutions [113] by solving a vertical momentum equation for each neutral species. The species momenta are then coupled using friction terms such that in the lower thermosphere the major species (for example, N_2) tend to force the minor species (for example, NO, O) to be nonhydrostatic. The relaxation of the hydrostatic assumption is useful for modeling the auroral region due to the presence of strong heating effects.

One of the main differences between the upper atmosphere and troposphere is the effect of drivers. The time evolution of the troposphere is determined mainly by its initial conditions. In the upper atmosphere, however, the highly time-dependent drivers play a more important role, thereby making the upper atmosphere a contractive system [114, 115]. For

example, the relationship between neutral temperature and solar radio flux at the wavelength of 10.7 cm ($F_{10.7}$) can be shown to have a positive correlation. This implies that, as $F_{10.7}$ increases, the neutral temperature will increase in the atmosphere region closest to the sun. Various studies [116] have shown that, when the solar flux increases dramatically (as in a solar flare), the density on the day side responds rapidly; in contrast, the rest of the atmosphere responds with a time delay as the perturbation wave propagates to the night side. It can also be shown that, as temperature increases, the volume of the atmosphere increases and pushes the upper atmosphere higher in altitude. Accordingly, a satellite orbiting at a fixed altitude would observe an increase in ρ on the day side. The direct relationship between ρ and $F_{10.7}$ is exploited in section 5.3.1, where the density at about 400 km above the subsolar point is assimilated to estimate $F_{10.7}$.

While models can provide insight into the dynamics of the upper atmosphere, it is often necessary to validate and calibrate the model dynamics using measurements. The CHAMP

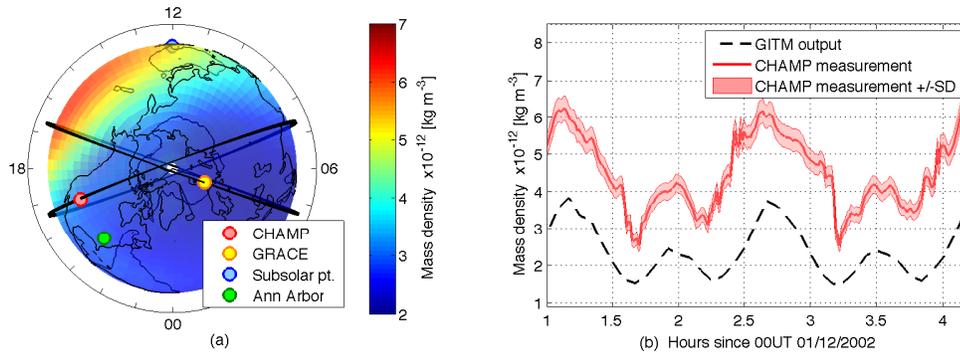


Figure 5.1: (a) shows the neutral density at about 400 km altitude at 02:32UT 12/1/2002 as well as CHAMP and GRACE orbital trajectories. CHAMP (red) and GRACE (yellow) orbits for the chosen time period do not pass near the subsolar point. The GITM gridpoint closest to Ann Arbor, MI (green) is referred to in subsequent figures. (b) demonstrates the model-data mismatch between no-DA GITM output (blue) and CHAMP measurement (red). GITM without data assimilation underestimates neutral density compared to CHAMP measurements.

satellite, which orbited from 2000 to 2010, provided measurements of the upper atmospheric state. One measurement was of the acceleration experienced by the satellite, which can be used to infer drag. Since drag is proportional to the mass density of the thermo-

sphere, mass density can be estimated from CHAMP measurements [117, 118]. Neutral density readings from CHAMP are separated in time by approximately 47 sec on average. The CHAMP trajectory (87.3° inclination with an average altitude of about 400 km during the simulated days) is shown in Figure 5.1(a). Figure 5.1(b) shows that no-DA GITM ($F_{10.7} = 142$ SFU) underestimates ρ by about 2×10^{-12} kg m⁻³ compared to CHAMP measurements between 01 UT and 04 UT. One reason for this mismatch is that the high-latitude drivers (electric potential and aurora) are intentionally held steady during this time period to cause GITM temperatures to drop, thereby forcing a bias in GITM neutral density. It is shown in [119] that variations in high-latitude drivers can cause heating. Therefore, by purposefully holding the drivers steady, the atmosphere cools off, and the neutral density at CHAMP altitudes decreases. The following sections establish how model bias can be removed by estimating drivers, whereas section 5.3 applies this technique and demonstrates how GITM estimate of neutral density at CHAMP locations can be brought closer to CHAMP measurements by estimating $F_{10.7}$.

5.2.2 EAKF: Data Assimilation without Driver Estimation

This section demonstrates data assimilation on a linear system without the knowledge of the time-varying driver. In the next subsection the same example is revisited with driver estimation. Consider the linear system

$$s_k = 0.5s_{k-1} + u_{k-1}, \quad u_k = 1.0 + \sin(0.4k), \quad (5.1)$$

$$y_k = s_k + v_k, \quad v_k \sim N(0, 0.3), \quad (5.2)$$

where k is the time step, s_k is the true state of the system, u_k is the unknown time-varying driver, y_k is the state measurement corrupted by noise v_k sampled from a zero-mean normal distribution with known variance $R_k = 0.3$. The goal of this experiment is to estimate the state s_k based on the noisy measurements y_k without knowledge of the driver u_k . For the

i^{th} ensemble member, $\hat{s}_{k,i}^+$ denotes the posterior state estimate just after the measurement y_k is assimilated, whereas $\hat{s}_{k,i}^-$ denotes the prior state estimate just before assimilating y_k . The prior state estimate is calculated from the posterior estimate at the previous step by propagating the dynamics, as given by

$$\hat{s}_{k,i}^- = 0.5\hat{s}_{k-1,i}^+ + \hat{u}_{k-1,i}, \quad (5.3)$$

where, for each ensemble member, $\hat{u}_{k-1,i}$ is a constant value selected randomly from a normal distribution (driver estimation is explored in section 5.2.3). EAKF assumes that $\hat{s}_{k,i}^-$ and y_k are normal random variables whose statistics can be described by their means and variances. The prior state variance can thus be approximated by the sample variance

$$\sigma^2[\hat{s}_k^-] = \sum_{i=1}^N (\hat{s}_{k,i}^- - \mu[\hat{s}_k^-])^2 / (N - 1), \quad (5.4)$$

where $\mu[\hat{s}_k^-]$ is the prior ensemble sample mean. EAKF generates the posterior ensemble estimate by applying Bayes theorem. Accordingly, the posterior probability density function (PDF) is equal to the normalized product of the prior and the measurement likelihood [46]. Since the product of two normal random variables is also a normal random variable, the posterior state estimate PDF is also normal, and its mean and variance are given by

$$\mu[\hat{s}_k^+] = \sigma^2[\hat{s}_k^+] \left[\frac{\mu[\hat{s}_k^-]}{\sigma^2[\hat{s}_k^-]} + \frac{y_k}{R_k} \right] \quad (5.5)$$

$$= \frac{R_k}{\sigma^2[\hat{s}_k^-] + R_k} \mu[\hat{s}_k^-] + \frac{\sigma^2[\hat{s}_k^-]}{\sigma^2[\hat{s}_k^-] + R_k} y_k, \quad (5.6)$$

$$\sigma^2[\hat{s}_k^+] = \left[\frac{1}{\sigma^2[\hat{s}_k^-]} + \frac{1}{R_k} \right]^{-1}. \quad (5.7)$$

The final step of EAKF scales and shifts the prior ensemble to match the new mean and variance, as given by

$$\hat{s}_{k,i}^+ = \sqrt{\frac{\sigma^2[\hat{s}_k^+]}{\sigma^2[\hat{s}_k^-]}}(\hat{s}_{k,i}^- - \mu[\hat{s}_k^-]) + \mu[\hat{s}_k^+]. \quad (5.8)$$

To demonstrate the overall performance of (5.1)-(5.8), 20 ensemble members are initialized from a zero-mean normal initial distribution with variance of 0.4 for both the state and driver estimates. The initial variance is chosen to be greater than the noise variance (R_k) so as to initially weight the measurements y_k more heavily. More precisely, $\sigma^2 > R_k$ implies that the measurement weight $\sigma^2/(\sigma^2 + R_k)$ is greater than the prior mean weight $R_k/(\sigma^2 + R_k)$ in (5.6). Equations (5.1)-(5.8) are propagated forward in time for 50 steps, while the measurements y_k are assimilated at every step.

Figure 5.2 shows the resulting state and driver estimates as a function of time. Along with plotting the true state (s_k), Figure 5.2(a) shows the measurement (y_k), the initial ensemble distribution for the posterior state estimate (\hat{s}^+), and the mean and spread of the ensemble ($\mu[\hat{s}]$ and $\sigma[\hat{s}]$). The line depicting the ensemble mean (blue) consists of prior and posterior estimate, thereby resulting in a discernible discontinuity at every step. Figure 5.2(a) demonstrates that, although the posterior mean approaches the true state s_k , the prior mean deviates away from s_k due to an incorrect driver value used by the ensemble members.

The performance can be better quantified in terms of estimation errors. The prior state root mean square percentage error (RMSPE) is defined as

$$\text{RMSPE}^- = 100 \times \frac{\sqrt{\mu[(\mu[\hat{s}_k^-] - s_k)^2]}}{\sqrt{\mu[s_k^2]}}, \quad (5.9)$$

and the posterior RMSPE is defined by replacing the minus superscripts by plus superscripts. Calculating RMSPE for the estimates shown in Figure 5.2(a) results in prior and

posterior state RMSPE values of 50.92% and 26.41%, respectively. To decrease the influence of a particular noise realization, this experiment was repeated using different noise realizations 1,000 times, and the averaged state RMSPE was found to be 62.40%. The next section shows how this value can be reduced by estimating the driver.

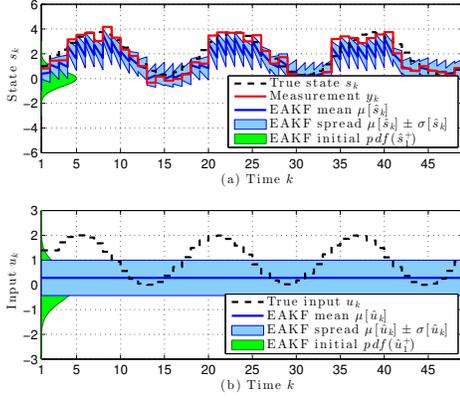


Figure 5.2: Time evolution of the EAKF state estimate for the system (5.1), (5.2) without driver estimation. (a) shows the true state (black dashed), the state measurement (red), the state estimate initial distribution (green), and the state estimate mean (dark blue) and spread (light blue, standard deviation). (b) shows the true driver, driver estimate distribution, mean, and variance in the same colors as (a). The posterior ensemble mean converges to the true state, whereas the prior mean deviates during the model propagation step due to the incorrect driver estimate.

5.2.3 EAKF: Data Assimilation with Driver Estimation

This section demonstrates the effect of introducing driver estimation into the data assimilation framework. To perform driver estimation, the augmented state vector is defined as $x_k = [s_k; u_k]$. Since state augmentation makes the problem of estimating x_k multivariate, the EAKF equations take on a matrix form, as derived in [46].

The matrix formulation of EAKF can be described in joint state-measurement space [46, 120], where the joint state vector is defined as $z_k = [x_k; y_k] = [s_k; u_k; y_k]$. Then,

(5.3)-(5.8) for the i^{th} ensemble member become

$$\hat{z}_{k,i}^- = [0.5\hat{s}_{k-1,i}^+ + \hat{u}_{k-1,i}^+; \quad \hat{u}_{k-1,i}^+; \quad \hat{s}_{k,i}^+], \quad (5.10)$$

$$P_k^- = \sum_{i=1}^N (\hat{z}_{k,i}^- - \mu[\hat{z}_k^-])(\hat{z}_{k,i}^- - \mu[\hat{z}_k^-])^T / (N - 1), \quad (5.11)$$

$$\mathcal{A}_k = (\mathcal{F}_k^T)^{-1} \mathcal{G}_k^T (U_k^T)^{-1} B_k^T (\mathcal{G}_k^T)^{-1} \mathcal{F}_k^T, \quad (5.12)$$

$$P_k^+ = [(P_k^-)^{-1} + \mathcal{H}^T R_k^{-1} \mathcal{H}]^{-1} = \mathcal{A}_k P_k^- \mathcal{A}_k^T, \quad (5.13)$$

$$\mu[\hat{z}_k^+] = P_k^+ [(P_k^-)^{-1} \mu[\hat{z}_k^-] + \mathcal{H}^T R_k^{-1} y_k], \quad (5.14)$$

$$\hat{z}_{k,i}^+ = \mathcal{A}_k^T (\hat{z}_{k,i}^- - \mu[\hat{z}_k^-]) + \mu[\hat{z}_k^+], \quad (5.15)$$

where $\hat{y}_k \in \mathbb{R}^m$, $\hat{z}_k \in \mathbb{R}^{n+m}$, \mathcal{F}_k is obtained from the singular value decomposition (SVD) $P_k^- = \mathcal{F}_k D_k \mathcal{F}_k^T$, \mathcal{G}_k is a square root of D_k , U_k is obtained from the SVD $\mathcal{G}_k^T \mathcal{F}_k^T \mathcal{H}^T R_k^{-1} \mathcal{H} \mathcal{F}_k \mathcal{G}_k = U_k J_k U_k^T$, B_k is a square root of $(I + J_k)^{-1}$, and $\mathcal{H} = [0_{m \times n} \quad I_{m \times m}]$.

Note that (5.10) assumes no model uncertainty, that is, it assumes no mismatch between the model used to generate the truth data (truth model) and the model used to propagate ensemble members (ensemble model). Additionally, the ensemble model assumes no dynamics for the driver, that is, the prior driver at the current step $\hat{u}_{k,i}^-$ is equal to the posterior driver at the previous step $\hat{u}_{k-1,i}^+$.

To demonstrate driver estimation, (5.10)-(5.15) are applied to the linear system (5.1) and (5.2). As in the previous section, the ensemble model is initialized from 20 different initial conditions for both $\hat{s}_{1,i}^+$ and $\hat{u}_{1,i}^+$, which are taken from a zero-mean normal distribution with variance 0.4. Figure 5.3 shows that both the state and the driver estimates converge to constant values and the ensemble spread converges to zero. The ensemble spread converges to zero due to the fact that the ensemble model assumes no dynamics for the driver. The resulting average RMSPE value of 55.45% is an improvement from the values achieved in the previous section by about 7%. Ensemble spread converging to zero and ensemble mean converging to a fixed value are referred to as filter divergence.

To remedy filter divergence, the prior ensemble variance can be inflated as given by

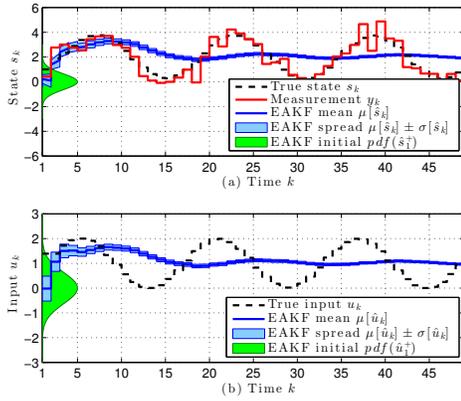


Figure 5.3: Time evolution of EAKF state and driver estimates for the system (5.1), (5.2) using driver estimation, but without ensemble inflation. Due to ensemble model assuming no dynamics for the driver, ensemble spread (light blue) approaches zero. Accordingly, the mean estimates converge to fixed values and do not track the true state and driver.

$\hat{z}_{k,i}^- = \sqrt{\lambda}(\hat{z}_{k,i}^- - \mu[\hat{z}_k^-]) + \mu[\hat{z}_k^-]$, where $\lambda > 1$ [121]. This directly increases both the ensemble spread at every step and makes state and driver estimates more uncertain, thereby allowing them to be updated by a greater increment. Figure 5.4 demonstrates that $\lambda = 2.0$ decreases the state RMSPE by about 17% to 38.63%. The driver estimate has approximately the same period and amplitude as the actual driver, but is phase shifted.

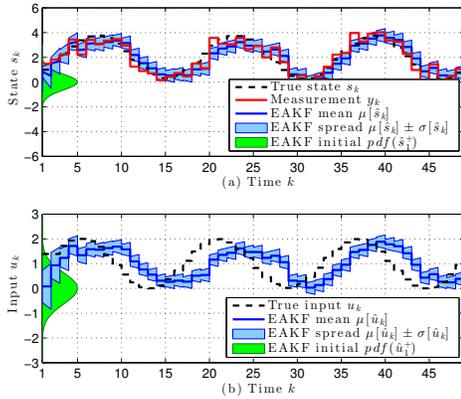


Figure 5.4: Time evolution of EAKF state and driver estimates for the system (5.1), (5.2) when driver estimation and ensemble inflation are used. Using variance inflation of $\lambda = 2.0$ results in an average decrease of 17% in the state tracking error.

5.2.4 EAKF: Data Assimilation with Driver Estimation, Ensemble Inflation and Model Uncertainty

We now consider the case where the ensemble model (5.10) is different from the truth model (5.1), (5.2). For example, consider the ensemble model

$$\hat{z}_{k,i}^- = [0.1\hat{s}_{k-1,i}^+ + \hat{u}_{k-1,i}^+ \quad \hat{u}_{k-1,i}^+ \quad \hat{s}_{k,i}^+]^T, \quad (5.16)$$

where the only difference between this ensemble model and the truth model (5.1) is in the state dynamics coefficient, which is 0.1 for the ensemble model and 0.5 for the truth model. Since the truth model and the ensemble model are different, the driver estimate has to correct not only for the uncertainty in the measurement y_k , but also the model uncertainty [122]. As a result $\hat{u}_{k-1,i}^+$ might not converge to the true driver value, but instead to a value that can drive the state of the ensemble model to the true state. Ensemble inflation has more influence in the uncertain model case as compared to the accurate model case since the driver has to compensate for model uncertainty in addition to the measurement uncertainty. Accordingly, using $\lambda = 2.0$ for the ensemble model (5.16) to assimilate measurements from the truth model (5.1), (5.2), results in the performance shown in Figure 5.5. The value of λ is chosen as a compromise between ensemble spread and state tracking. Increasing λ increases the ensemble spread, which in turn allows for large instantaneous change in the estimates. On one hand, a large spread is undesirable since it implies large uncertainty in the estimates, whereas, on the other hand, a small spread slows down the updating of the estimates, and may result in filter divergence.

In this case, the state RMSPE is degraded by 1% (with a new value of 39.59%) when compared to the accurate model case. The driver estimate had to deviate farther from true value than in the previous section, but was able to compensate for model uncertainty as shown by relatively small degradation in state RMSPE when compared to values in the previous section.

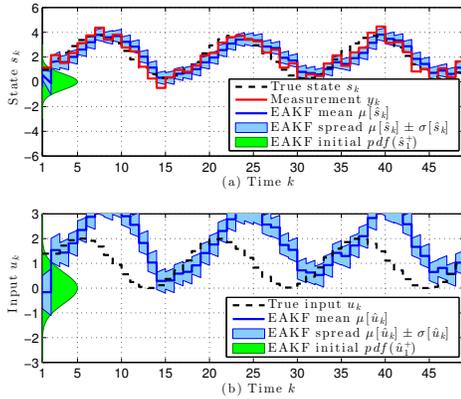


Figure 5.5: Time evolution of EAKF state and driver estimates for the system (5.1), (5.2) using the ensemble model (5.16) where driver estimation and ensemble inflation are used. With ensemble inflated using $\lambda = 2.0$, the driver estimate compensates for model uncertainty and allows minor degradation in state estimate (1%) as compared to the perfect model case.

5.2.5 EAKF: Localization

The previous section considered only a single state and a single driver. Within GITM there are $72 \times 36 \times 50$ grid points with 35 state variables at each grid point (densities, temperatures, and velocities for multiple species), all of which are estimated. Using the equations derived in the previous section for problems of this size is not only prohibitively expensive computationally, but also can produce physically meaningless solutions arising from correlations between distant states and measurements. If these correlations are retained, the state estimates could be driven in the wrong direction. Accordingly, the influence of the measurements on the state variables at the current step is spatially confined. This region is defined as a smooth function of the distance between the measurement and the state variable. This localization function takes values between 0 and 1, and multiplies the correlation between the measurement and the state variable [123]. A value of 1 corresponds to direct connection between the measurement and the state, whereas a value of 0 corresponds to no connection. The values between 0 and 1 vary smoothly as a function of distance. The piecewise continuous function (4.10) in [124] is used with a half width of 0.6 radians in the

horizontal direction and 100 km in altitude for all experiments in this chapter. The resulting localization region is shown in Figure 5.6 for a measurement above Ann Arbor, MI, where the intensity of red indicates the value of the localization function.

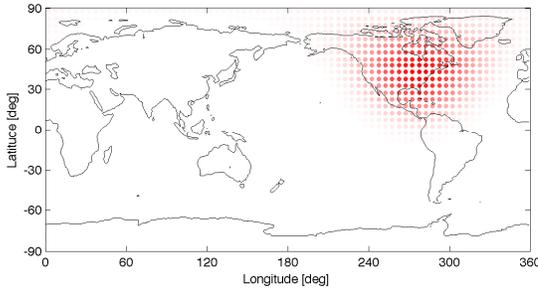


Figure 5.6: The localization region for the measurement located above Ann Arbor, MI. Markers are placed at GITM cell center locations. The dark red markers correspond to measurements having a direct effect on the state variable, whereas transparent markers correspond to measurements that have no effect on the state variable. The state variables affected by the measurement above Ann Arbor lie in a projection of a small circle region centered in Ann Arbor.

In this chapter, localization is used for all GITM variables except the driver estimate $\hat{F}_{10.7}$, which is affected by all measurements. This implies that, for example, the measurements on the night side of the Earth have full impact on the driver.

Although data assimilation can be performed at every step in simple models, this approach becomes impractical in the DART-GITM interface due to the cost of stopping and restarting GITM. Accordingly, even though GITM uses 2 s time step and CHAMP data is available about every 47 s, the assimilation window (time between assimilation steps, t_{aw}) was chosen to be 30 min, which is a compromise between runtime and performance.

One current feature of DART-GITM interface is that it does not interpolate measurements in time. In other words, all the measurements from current time minus half the assimilation window to current time plus half the assimilation window are used without modification as if they occurred at the current time.

5.3 Results

All of the data assimilation experiments performed on GITM were run for a period of two days (Dec 1-2 2002). For the simulated measurement experiments the measurement error variance R_k was set to a constant value of $2.6 \times 10^{-13} \text{ kg m}^{-3}$, which was the average CHAMP measurement error for the simulated dates. For the real CHAMP measurement experiment, R_k was set to the measurement error values provided in the CHAMP data files. In this section, the driver estimate was inflated to have a constant variance; the state estimate covariance was not inflated. More precisely, the entry corresponding to the variance of $\hat{F}_{10.7}$ in the prior covariance matrix (5.11) was set to the fixed value $\sigma_i^2 = 49 \text{ SFU}^2$ just before computing (5.11). Keeping the driver variance constant was found to outperform state inflation since the driver estimate variance in the state inflation case would take multiple days to grow to the desired level. On the other hand, driver-only inflation reaches these levels immediately and keeps the variance of the driver estimate constant.

While the EAKF assumes that the state estimates are normally distributed, temperatures, densities and $F_{10.7}$ cannot take on negative values and hence cannot be normal. Accordingly, when estimates of these state variables become negative, their values are set at half the value at the previous step. Therefore, the initial $\hat{F}_{10.7}$ ensemble distribution was chosen to be normal with mean of 130 SFU and variance of 25 SFU^2 . This distribution created spread in the initial conditions for other variables since all ensemble members were pre-spun for two days (11/29-12/01) prior to the start of the data assimilation. The mean was chosen to be below the true value to avoid inadvertently compensating for GITM-CHAMP ρ mismatch. More precisely, since the true $F_{10.7}$ was about 142 SFU during Dec 1-2 2002 and GITM underestimated ρ at CHAMP locations, choosing an initial mean $\hat{F}_{10.7}$ to be about 220 SFU would fix the bias. Instead, the mean of 130 SFU was chosen to provide for a more challenging setup. The variance of the initial $\hat{F}_{10.7}$ distribution was chosen as to create the initial variance in ρ to be greater than R_k in the beginning of the assimilation to give the measurements more weight (see section 5.2.2 for related discussion).

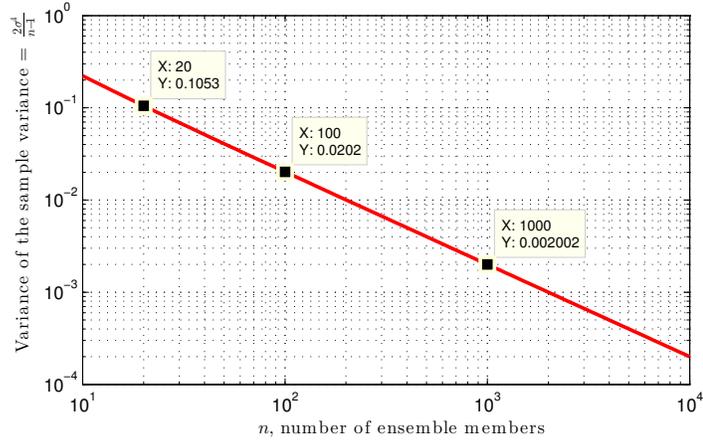


Figure 5.7: Variance of the sample variance as a function of the number of ensemble members when $\sigma = 1$. As the number of ensembles members increases, the accuracy of variance estimate increases.

Finally, the number of ensemble members was fixed at 20 for this study due to the initial limit on computational resources. More precisely, 5° GITM runs on 32 processors, which for 20 instances entails 640 processors. To facilitate the simulations, NASA’s Pleiades supercomputer was used. One way to gauge the accuracy of the estimate of the ensemble variance is by considering the distribution of sample variance, which can be shown to be scaled chi-squared distribution with mean σ^2 and variance $\frac{2\sigma^4}{n-1}$ if n samples come from a normal distribution with variance σ^2 [125, 126]. Figure 5.7 demonstrates variance of the sample variance as a function of the number of the ensemble members and shows that, with 20 ensemble members the variance of the ensemble variance reaches 0.1053 when the true variance is 1. It also demonstrates that increasing the number of ensemble members improves the accuracy of the variance estimate in hyperbolic fashion since a tenfold increase in number of ensemble members reduces variance of ensemble variance by a factor of ten.

5.3.1 Simulated data above the subsolar point

As a first example, a simulated measurement of the density ρ_{sp} at about 400 km above the subsolar point is obtained from a GITM truth simulation. Since it is impractical to measure ρ_{sp} in practice, this example illustrates EAKF in the case where the measurement is directly linked to the quantity to be estimated. We thus record ρ_{sp} at 1-min intervals from the truth model for use in the EAKF as measurements. Truth GITM is simulated for 2 days (Dec 1-2 2002) with the true $F_{10.7}$ (mean value of 142 SFU). With data assimilation performed every 30 min, Figure 5.8 shows results from the truth model and ensemble estimates at two locations, namely, the subsolar point (measurement location) and the GITM grid point closest to Ann Arbor, MI (diagnostics location). Figure 5.8(b) demonstrates that state estimates (blue) at the subsolar point converge to the true state (red). The measurement uncertainty used in this experiment is the average value of CHAMP measurement error data (standard deviation) for the Dec 1-3 2002 and is shown in light red. The solid black line shows the estimated density at the subsolar point without data assimilation with $F_{10.7} = 130$ SFU (the mean of the initial $\hat{F}_{10.7}$ distribution). The assimilated mean deviates from this passive trajectory and converges to approximately the true state within 3 hours.

Figure 5.8(b) shows that even at a location different from the measurement location (the GITM grid point closest to Ann Arbor, MI), the mean state estimate (blue, $\hat{\rho}_{aa}$) converged to the true state (dashed black, ρ_{aa}). At this location the convergence is slower since Ann Arbor only entered the subsolar point localization region between 12 and 24UT each day.

Figures 5.9(a) and 5.9(b) show mass density estimates at CHAMP and GRACE locations ($\hat{\rho}_{ch}$ and $\hat{\rho}_{gr}$, respectively). The inset in (a) demonstrates that by about 09UT 12/1/2002 CHAMP ensemble mean $\mu[\hat{\rho}_{ch}]$ converged to the true value ρ_{ch} . Similarly, (a) leads us to conclude that convergence took place along the GRACE trajectory as well. The

absolute percentage error (APE) values are computed for the (a) and (b), as

$$\text{APE}_k^- = 100 \times \frac{|\mu[\hat{\rho}_{ch,k,i}^-] - \rho_{ch,k}|}{|\rho_{ch,k}|}, \quad (5.17)$$

and are shown in (c) and (d). These plots demonstrate that the subsolar point measurement not only removed the bias at the subsolar point, but also at CHAMP and GRACE locations. More precisely, the RMSPE^- at CHAMP locations over the second day decreased from 36% to 3% when compared to the bias in the nominal case (the truth model with $F_{10.7} = 142$ SFU and no-DA GITM using $F_{10.7} = 130$ SFU). Similarly, the RMSPE^- at the GRACE location decreased from 43% to 4%. Finally, Figure 5.10 shows that the $F_{10.7}$ estimate converged to the true value by the end of the simulated period.

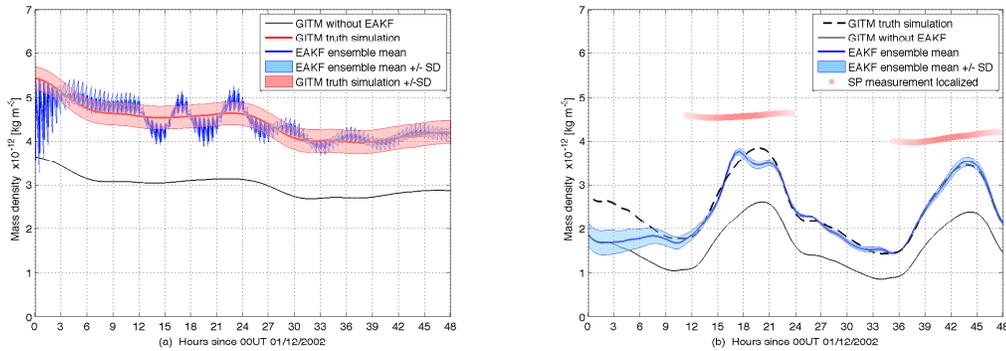


Figure 5.8: Time evolution of the simulated (red) and estimated (blue) mass density above the subsolar point (measurement location) in (a) and at the 400 km gridpoint closest to Ann Arbor, MI in (b). Estimated mass density above the subsolar point converged to the true value within 3 hours, whereas density at the Ann Arbor gridpoint needed 12 hours to converge.

5.3.2 Simulated CHAMP data

As a more physically realistic case, we now use measurements at the CHAMP locations. Figure 5.1 shows that the CHAMP orbit does not pass through the subsolar point, and hence this case is more challenging than the previous example. Figure 5.11 demonstrates

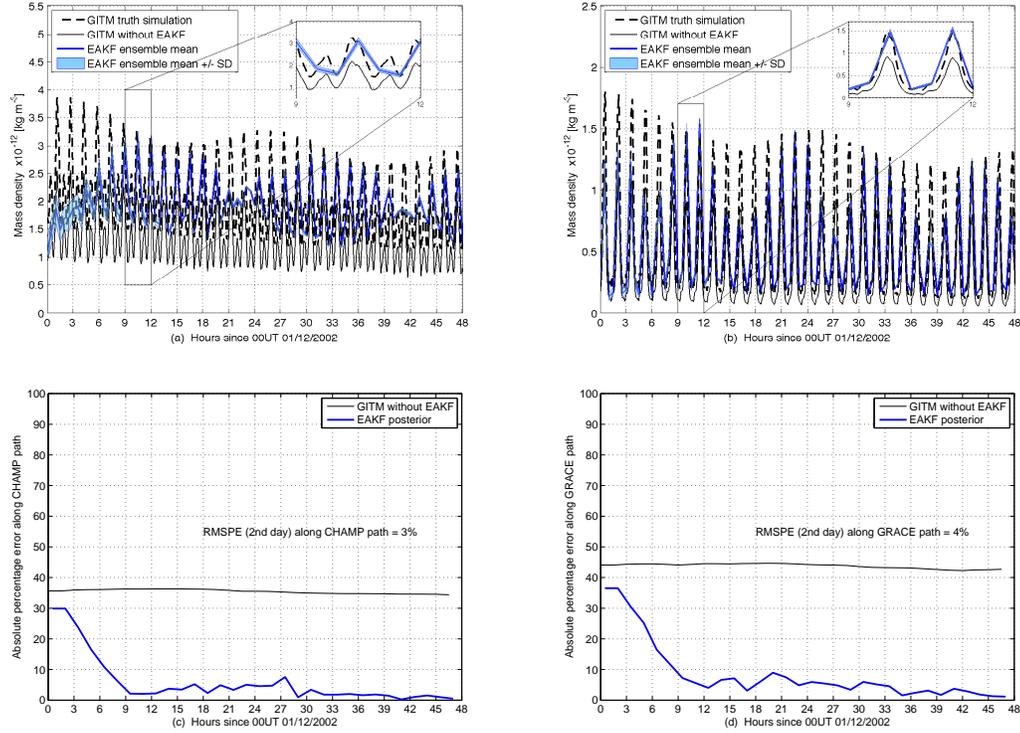


Figure 5.9: (a) shows the time evolution of mass density at CHAMP locations demonstrates that ρ converged to the true value within 9 hours when measurements were taken at the subsolar point. (b) shows mass density estimate at GRACE locations and demonstrates the same rate of convergence. (c) and (d) reinterpret (a) and (b) by plotting the orbit averages of the errors between simulated and estimated data. This figure demonstrates that estimated mass density converged to true mass density at locations other than the measurement location (subsolar point).

that even in this seemingly harder case, EAKF was able to decrease the prior RMSPE from 36% to 2% along CHAMP trajectory and from 43% to 4% along GRACE trajectory compared to the no-DA case. The values of $\hat{F}_{10.7}$ shown in Figure 5.12 demonstrate that the filter initially overestimated the true value, but then converged close to the true value within 15 hours. Indeed, the convergence was significantly faster than in the previous example.

5.3.3 Real CHAMP data

We now consider mass density data and uncertainty measured by CHAMP. Because the true state was unknown, data from GRACE was used as the validation metric. Since GITM with

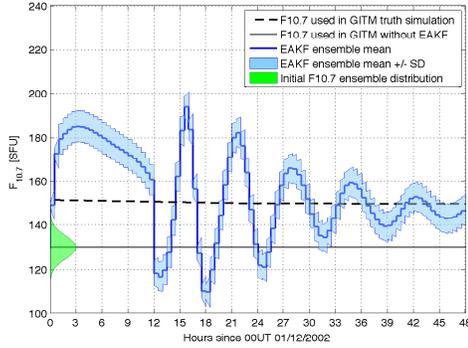


Figure 5.10: Time evolution of $F_{10.7}$ for the truth simulation (dashed black), nominal no-DA simulation (solid black), and the ensemble mean (dark blue). The ensemble mean oscillated around the true value and converged to the true value by the end of the simulation.

proper high-latitude driving has a relatively low bias compared to CHAMP data, GITM simulations were conducted with constant high-latitude driving, to intentionally introduce potential bias. Figure 5.13 shows that the bias at the CHAMP locations with data assimilation and driver estimation is reduced from 73% to 7%. At the GRACE satellite location, however, the bias was only marginally reduced. Even though the GRACE RMSPE is reduced from 76% to 52%, the large remaining bias suggests that either more than one driver needs to be estimated in order to remove bias at multiple locations or that more data needs to be assimilated to remove the bias over the whole thermosphere. Figure 5.14 shows the $F_{10.7}$ values used to achieve the aforementioned performances and demonstrates that the driver had to take on higher values (about 220 SFU instead of nominal 142 SFU) to compensate for the model bias.

5.4 Conclusions

This study described the preliminary implementation of the DART-GITM interface. In particular, it demonstrated how both simulated and real CHAMP ρ measurements can be assimilated into GITM. In addition to estimating GITM states such as densities, temperatures, and velocities, this interface estimated one of GITM drivers, namely, $F_{10.7}$. This

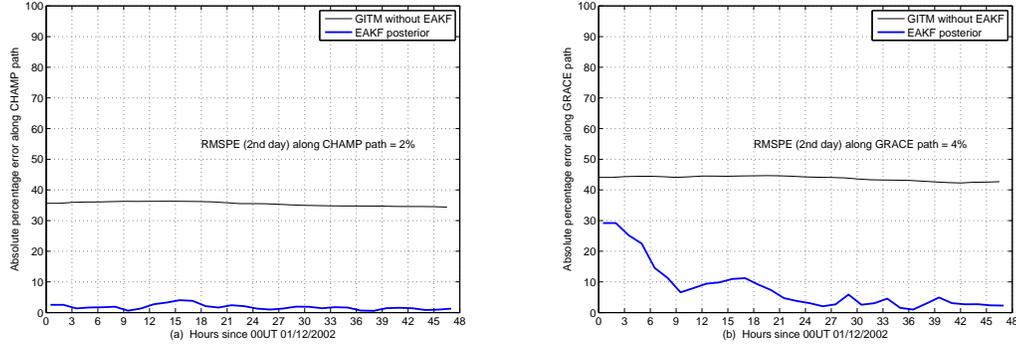


Figure 5.11: Orbit averaged mass density absolute percentage errors at (a) CHAMP and (b) GRACE locations. Simulated CHAMP data were used as measurements in the EAKF, whereas the simulated GRACE values were only used for validation. Density estimate at CHAMP locations dropped immediately to about 3%, whereas the estimate at GRACE locations decreased only after about 9 hours due to altitude difference between these satellites. The bias in densities at both locations was decreased and the final error level was lower at the measurement location (CHAMP) than at the validation location (GRACE).

driver was not estimated for the purposes of knowing $F_{10.7}$ more precisely, but instead for the purposes of driving GITM’s estimate of ρ closer to CHAMP measurements. Accordingly, it was found that in the simulated CHAMP measurements case, the good estimate of $F_{10.7}$ was able to drive GITM to the point of decreasing the bias in the simulated mass densities at the CHAMP locations from 36% to 2% and from 43% to 4% at the validation locations (GRACE orbit). This technique produced greater decrease when real CHAMP data were used as measurements in EAKF and GRACE data were used for validation. More precisely, assimilating real CHAMP data and estimating $F_{10.7}$ decreased the mass density bias from 58% to 7% for CHAMP locations and from 77% to 52% for GRACE locations when compared to not performing data assimilation. The relatively large final GRACE bias (52% compared to CHAMP’s 7%) can be explained by the way these data were derived. In particular, [37] mentions that CHAMP and GRACE data were found to be inconsistent when compared to the CTIpe outputs and attributes this disparity to the “uncertainty of the drag coefficient assumptions employed in the retrieval as well as the different orbital altitudes.” The authors then scaled the GRACE densities by an arbitrary number to remove the

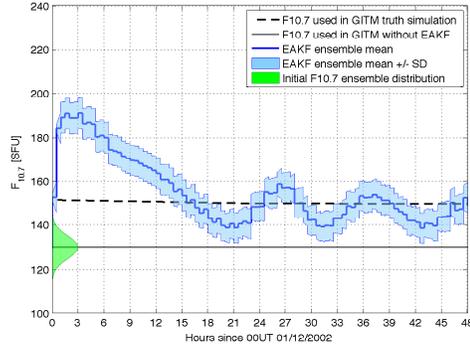


Figure 5.12: Time evolution of $F_{10.7}$ for the true, nominal and ensemble simulations. The ensemble mean estimate initially overestimates the true value, but then converges within 15 hours.

apparent inconsistency. This solution might have reduced the observed GRACE RMSPE in this study, but is left for future work.

These preliminary results can be further improved by tuning parameters and relaxing of some assumptions. In particular, the interface parameters that require tuning include the number of ensemble members, the data assimilation temporal window, and observational error variance R . The possible interface improvements mentioned above include using log-normal distributions instead of normal distributions for densities and temperatures, defining a localization region for the driver, and implementing temporal localization of the measurements. Also, investigating full state inflation in conjunction with constant driver inflation is subject for future work.

Future work will include estimating more drivers, such as the auroral power and cross polar cup potential, and assimilating more measurements, such as the total electron content measurements made by GPS, which will help to extend this study from a thermosphere into the ionosphere.

Lastly, since this study considered only the geomagnetically calm period, it would be interesting to see whether the current setup is capable of producing equally good estimates during a geomagnetic storm.

The main goal in this demonstration was to explore whether it is possible to remove

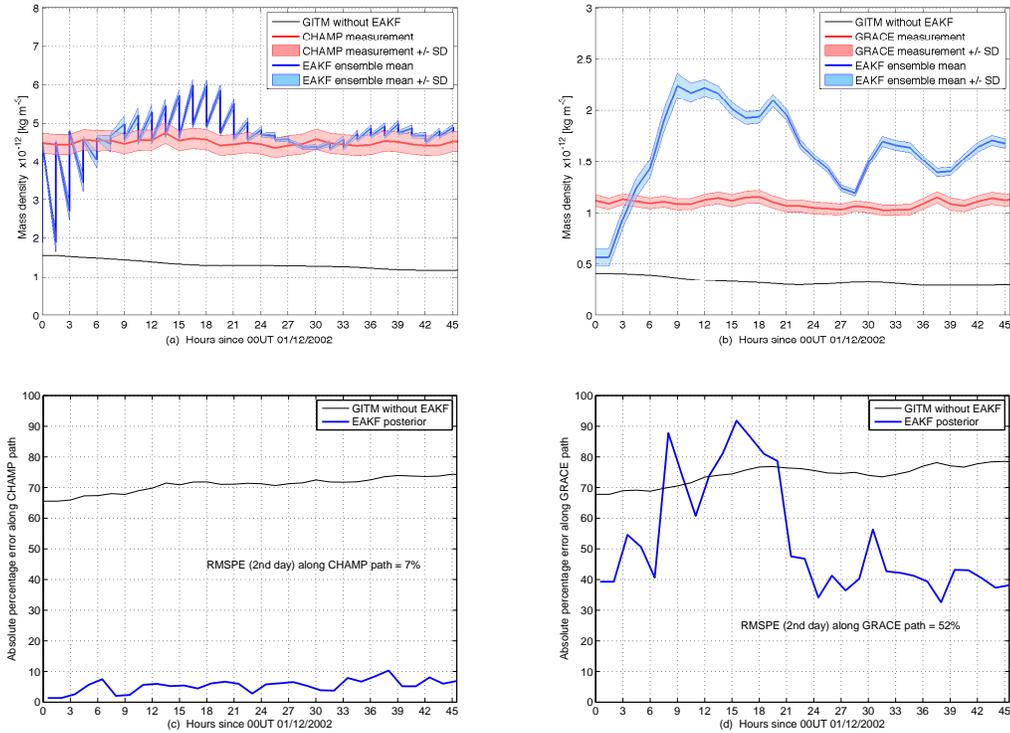


Figure 5.13: Orbit averaged mass density measurements at CHAMP and GRACE are shown in (a) and (b) respectively, and absolute percentage errors are shown in (c) and (d) respectively. Real CHAMP data were used as measurements in the EAKF, whereas GRACE data were used only for validation. The bias in density at CHAMP locations was decreased over the second day from 57% to 7%, whereas the bias at GRACE location was only decreased by a smaller amount (from 70% to 52%, see text for relevant discussion).

the bias in a model using data assimilation along with driver estimation. The last example demonstrates that at the location where data is ingested, the bias is almost completely removed (due mostly to the data assimilation), while, at other locations, the bias is reduced. This shows that the technique works but might be dependent on choosing the correct driver to estimate.

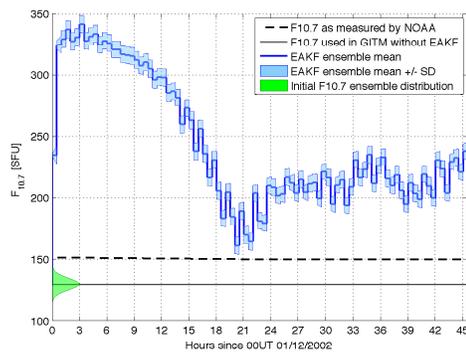


Figure 5.14: Time evolution of $F_{10.7}$ for the true, nominal and ensemble simulations. $F_{10.7}$ estimate took on higher values than what was measured by NOAA, in order to drive GITM closer to the CHAMP data.

CHAPTER 6

Assimilating Total Electron Content Measurements into GITM using EAKF

This chapter extends the findings of Chapter 5 by considering measurements coming from multiple satellites. In particular, the total electron content (TEC) measured by the global positioning system (GPS) satellites is assimilated into GITM. Two major differences between this chapter and Chapter 5 are: GPS measurements provide a much more global coverage than the CHAMP measurements and GPS measurements are related to the state of the ionosphere, whereas the CHAMP measurements are related to the state of the thermosphere. The first distinction dictates the technical details of choosing the EAKF localization parameters, whereas the second one has more general implications of choosing whether or not to estimate GITM drivers and if so, what drivers to estimate. The second distinction is very similar to the distinction made in the previous chapter between thermosphere and troposphere - the former one is strongly driven and is weakly dependant on the initial conditions, whereas the second one is weakly driven and is strongly dependent on the local structure of the initial conditions (akin to turbulence). Accordingly, this chapter investigates whether it is necessary to adjust GITM drivers in addition to GITM ionospheric states or whether it is enough to update GITM ionospheric states only.

6.1 Introduction

Knowing the structure of the ionosphere is crucial to multiple applications. Perhaps the most familiar one to most people would be GPS turn-by-turn navigation used in modern cars. This technology relies heavily on measuring the time it took for the signal to travel from the satellite to the receiver. This time is then used to calculate the distance between the satellite and the receiver. If only a simplistic model of 'time times speed' is used for computing distance, the resulting position uncertainty will be unacceptable for navigation purposes since the signal gets delayed by the ionosphere [127]. A variety of ionospheric models can be used to account for this delay. Empirical models can provide satisfactory estimates of the ionosphere during geomagnetically calm times, but provide unacceptable results during solar storms [128]. Therefore, this chapter uses a physics-based ionosphere model (GITM, [33]) and demonstrates how GITM can be used to estimate GPS ionospheric delays. More precisely, this chapter demonstrates how data assimilation can be used to improve GITM's electron density specification as to improve GITM's forecasting ability in terms of predicting ionospheric conditions. To accomplish this goal, the vertical total electron content (VTEC) measurements are assimilated by the ensemble adjustment Kalman filter (EAKF), which is a part of the Data Assimilation Research Testbed (DART, [47]).

This chapter first describes some traits of VTEC measurements in section 6.2.1 and describes the filter-model setup in Section 6.2.2. Then, the results of assimilating simulated VTEC data are described in Section 6.3.1. Section 6.3.2 also describes assimilating simulated data, but adds an extra degree of freedom by estimating a GITM driver (solar index $F_{10.7}$). Similarly, sections 6.3.3 and 6.3.4 describe the state-only and state-and-driver estimation experiments by using real VTEC data.

6.2 Methodology

6.2.1 Total Electron Content Measurements

The total electron content measurements are calculated from the frequency signals broadcasted by GPS (see, for example, [129, 130] and references given within). The *slant* total electron content (STEC) measurements represent an integral of electron density along the path from a ground receiver to a visible GPS satellite. Since the units of electron density are number of electrons per meter cubed ($\#e\ m^{-3}$), the units of STEC are $\#e\ m^{-2}$, which is scaled down by defining $1\ TECU = 10^{16}\ \#e\ m^{-2}$. The *vertical* total electron content (VTEC) measurements are calculated by projecting STEC measurements to the local-vertical axis via a simple cosine rule [129]. The daily VTEC data sets are available from the Madrigal open source database project [131] and have 5 minute latency. Most VTEC values lie in the [0 20] TECU range, although there are values above that range with extreme outliers of 300 TECU. The number of active ground stations varies on a per-day basis, but during the dates of this study (December 1 2002), about 200 stations were operational. The number of visible satellites changes much more rapidly and is different for each station, but is about 6 on average. Combining these figures and assuming each satellite provides data about every minute results in 6000 VTEC data points every 5 minutes. This amount of data is about 1000 times larger than what was available in the previous chapter (about 6.3 points every 5 minutes). To highlight the amount of available data, Figure 6.1 shows a snapshot of the geographical distribution of these measurements and highlights the fact that more measurements are available above land than above ocean.

With this distribution in mind, we first construct a simple test case using simulated data to prove the assimilation technique. Accordingly, sections 6.3.1 and 6.3.2 use simulated VTEC data recorded from a GITM “truth” simulation with $F_{10.7}=148.0$ SFU. The locations for retrieving data from the truth simulation were selected to be somewhat representative of the real data distribution and have simple shapes to make data-rich areas distinct from

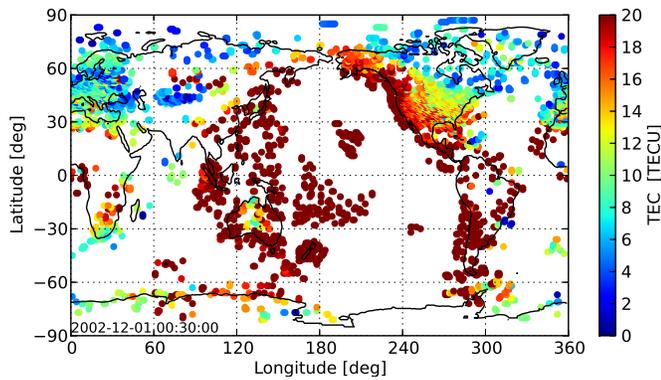


Figure 6.1: VTEC snapshot at 00:30 UT on Dec. 1 2002. VTEC data cover the whole globe, but are more sparse above oceans than above land.

areas with no data. Accordingly, longitude-latitude rectangles over Europe ($40^{\circ} - 60^{\circ}\text{N}$, $0^{\circ} - 40^{\circ}\text{E}$) and North America ($30^{\circ} - 50^{\circ}\text{N}$, $80^{\circ} - 120^{\circ}\text{W}$) with 2° increments in both dimensions were selected and are shown in Figure 6.2. Since the underlying resolution of the truth simulation was only 20° longitude by 10° latitude as is discussed in the next section, the truth data were interpolated via cubic splines to the 2° resolution.

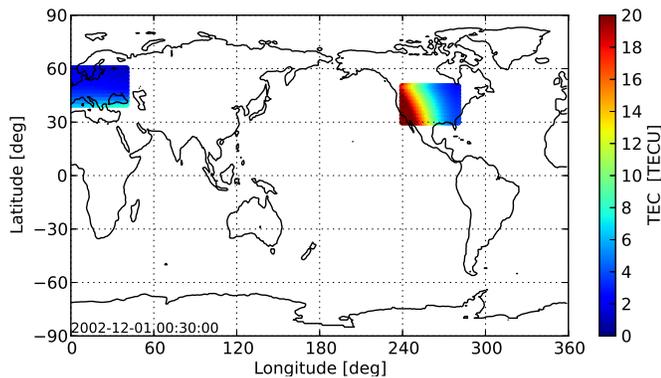


Figure 6.2: VTEC cropped simulated data at 00:30 UT on Dec. 1 2002. The cropping region locations were chosen to form rectangular patches over Europe and North America.

6.2.2 DART-GITM Interface Experimental Setup

The DART-GITM interface is extensively described in Chapter 5. This section describes some minor tuning and source-code changes made for this particular problem. GITM source code was not modified, whereas DART had to be augmented with a forward operator module, which describes how to compute VTEC from electron density. In a nutshell, the forward operator interpolates the three-dimensional global electron density field in longitude and latitude resulting in a vertical column of e-density values. This column is then integrated with respect to altitude using the trapezoidal rule and the result is multiplied by 10^{-16} to convert it to TEC units.

Coarse resolution was used for both truth and ensemble simulations (20° in longitude and 10° in latitude) as proving the technique was more important to this chapter than quantifying the behavior of small-scale features of the ionosphere. Using coarse resolution allowed for decreased simulation time, as simulating 20 ensemble members with the assimilation window of 30 min required only 9 wall-clock hours on 20 processors to simulate 24 virtual-hours, as opposed to the same amount of time on 640 processors to produce the high-resolution results of the previous chapter. The EAKF horizontal localization halfwidth was set at 0.6 radians as to ensure that state variables in regions with sparse data-coverage (oceans) would be affected by the assimilation. The vertical localization halfwidth was set to 100 km, but does not have any effect on assimilation since VTEC has no vertical coordinate. The whole-state (as opposed to a single state variable or single driver) inflation value λ introduced in Section 5.2.3 was set to 1.01 for all examples in this chapter as to prevent filter divergence. In addition to that, the $F_{10.7}$ inflation value $\sigma_i^2 = 1 \text{ SFU}^2$ was used in sections 6.3.2 and 6.3.4. The date chosen for this study (December 1 2002) was selected since the level of solar activity (about 150 SFU) was favorable for driver estimation, by which it is meant that it was not close to the GITM minimum allowable value (about 70 SFU). Lastly, the simulated VTEC measurements were assigned uncertainty (standard deviation) equalling 5% of the simulated measurement in sections 6.3.1 and 6.3.2, whereas the real

VTEC uncertainty was used with real measurements in sections 6.3.3 and 6.3.4.

6.3 Results

6.3.1 Assimilating simulated data above North America and Europe without estimating $F_{10.7}$

As a first example, simulated data above North America and Europe is assimilated to estimate GITM states without estimating GITM drivers. Figure 6.3 highlights the fact that the $F_{10.7}$ was not estimated and the initial $F_{10.7}$ distribution was chosen to be normal with mean of 130 SFU (which is about 18SFU below the value used in the truth simulation) and standard deviation of 5 SFU. Before proceeding to demonstrate the VTEC results, we turn

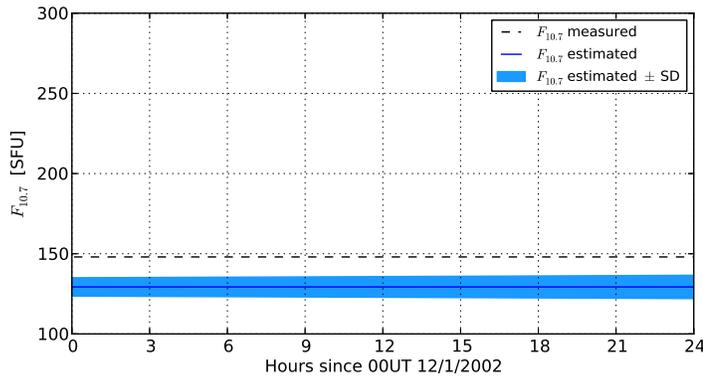


Figure 6.3: Initial distribution and time evolution of $F_{10.7}$ estimates. The $F_{10.7}$ estimates were initialized as $\mathcal{N}(130 \text{ SFU}, 5 \text{ SFU})$ and were kept constant for the entire simulation.

our attention to the variable from which VTEC is calculated, namely electron density. Electron density is three-dimensional and varies over time, so to simplify the plotting, values at one gridpoint are selected. In particular, electron density at the (40.0°N, 80.0°W, 411 km) GITM gridpoint is shown in Figure 6.4. This figure also demonstrates that the ensemble mean and spread increase starting 12UT (07LT) and reach their maximum around 19UT (14LT). Figure 6.4 also shows that the ensemble mean (solid blue) captures the overall

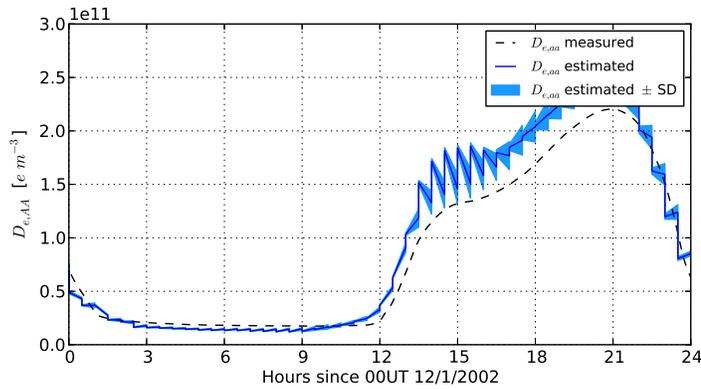


Figure 6.4: Time evolution of electron density at the (40.0°N, 80.0°W, 411 km) gridpoint. Discontinuities in the blue line at each step are a result of plotting prior and posterior density estimate at each step. The slope of blue line between assimilation steps is indicative of GITM influence on the estimates during the model-propagation periods.. Ensemble mean and spread increase after 12UT (07LT) and reach their maximum values around 19UT (14UT).

shape of the true electron density time profile (dashed black), but possibly not the magnitude.

Next, VTEC for both true and ensemble simulations is computed. Then, the absolute error ($|VTEC_{GITM} - VTEC_{EAKF}|$) is computed. The interpolated version of the VTEC error is shown in Figure 6.5. This figure demonstrates that the error is smallest in the data-rich

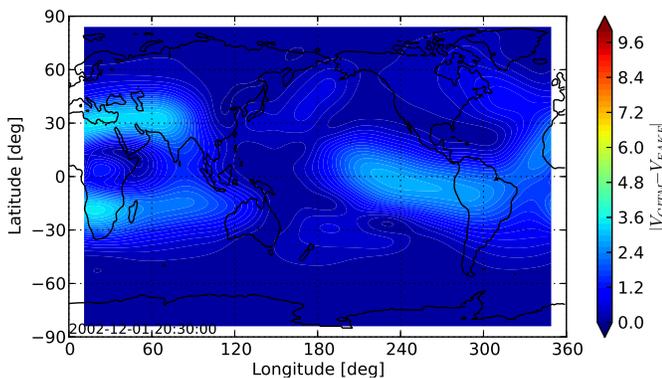


Figure 6.5: VTEC error snapshot at 22:30 UT on Dec. 1 2002. The error is smallest where data is available (note the darker area above the United States).

areas, as represented by the darker area above the United States, and is largest in areas with

poor data coverage, particularly the Pacific Ocean and Central Africa. However, Figure 6.5 is only a snapshot and does not provide a complete picture of the error evolution over time. To compensate for this issue, Figure 6.6 shows the error averaged over the whole globe as a function of time and demonstrates that in this simple case average error does not increase above 2 TECU.

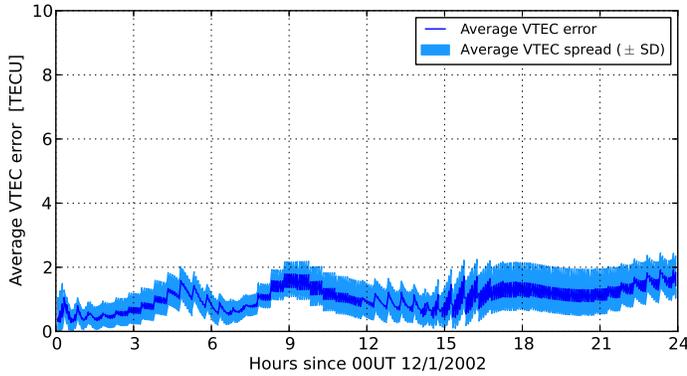


Figure 6.6: Time evolution of VTEC error average over the whole globe. Average error does not surpass 2 TECU.

6.3.2 Assimilating simulated data above North America and Europe with estimating $F_{10.7}$

This section introduces driver estimation in the simulated data case. Figure 6.7(a) shows the estimated $F_{10.7}$ and demonstrates that the driver estimate did not converge to the true value in the time allotted. However, when driver estimation is combined with data assimilation, the overall error levels do not surpass 4 TECU, as shown in (d). Figure 6.7(c) shows a VTEC snapshot similar to that of Figure 6.5, but the distinction between data-rich and data-poor areas is less distinct in this case, which is partially caused by introducing driver estimation.

6.3.3 Assimilating real GPS data without estimating $F_{10.7}$

This section introduces assimilation of real GPS data, but does not employ driver estimation. Figure 6.8 shows the $F_{10.7}$ estimate in (a), electron density gridpoint estimate in (b), VTEC error snapshot in (c), and VTEC error time evolution in (d). In particular, the error shown in (c) differs from the definitions in the previous sections and is defined as the absolute difference between the GPS data linearly interpolated to GITM gridpoints and the EAKF estimate ($|VTEC_{GPS} - VTEC_{EAKF}|$). This error is shown in Figure 6.8(c), which follows the trend set by the previous sections of having lowest VTEC error in the data-rich areas (North America and Europe) and highest in data-poor areas (South Pacific Ocean). Multiple upward spikes in Figure 6.8(b) imply that EAKF attempted to magnify GITM's electron density values, but during the model-propagation step these values slid back down due to, possibly, insufficient forcing. Overall, the VTEC error stayed at the level of about 8 TECU, which is 4 times as high as the simulated data case without driver estimation.

6.3.4 Assimilating real GPS data with estimating $F_{10.7}$

This section completes this study by adding driver estimation to the real-data-estimation case. Figure 6.9(a) demonstrates that in this case, driver estimate varied substantially more than in the simulated data case of Section 6.3.2 to compensate for larger spread of the data. Accordingly, the data assimilation part of EAKF did not have to adjust GITM state as much, which can be seen by a reduced number of upward spikes in Figure 6.9(b) as compared to Figure 6.8(b). This cooperation of data assimilation and driver estimation resulted in improved error distribution shown in (c), as shown by lower error values over South Pacific Ocean. Additionally, introducing driver estimation provided for a lower and flatter average error dynamics, as shown in Figure 6.8(d).

6.4 Conclusions

This chapter extended the findings of Chapter 5 into the ionosphere by assimilating GPS VTEC measurements. In output tracking (or command following) terms, this chapter described how GITM output can be made to track a desired trajectory (real satellite data) or, equivalently, how GITM can be inverted to back out driver values from the measurements. This task was shown to be achievable either through pure data assimilation, where model drivers are kept constant, or through data assimilation augmented with driver estimation.

While the simulated driver estimation case had worse average error than the case without driver estimation, this might change with different initial ensemble distributions. More precisely, while completing this study, the authors found that initializing $F_{10.7}$ at 130 SFU with 5 SFU standard deviation did not result in sufficient ensemble spread. Therefore, two possible solutions could be approached in future work: choosing a larger standard deviation for the initial $F_{10.7}$ distribution or choosing a different driver that has stronger effect on the electron density. In the real data case, however, driver estimation improved the overall assimilation performance by decreasing average VTEC error. A particular benefit of introducing driver estimation comes from decreased estimation error in sparse data regions.

Future directions for this project include investigating the effect of increasing the initial driver distribution standard deviation, selecting different (and possibly multiple) drivers, varying various filter tuning parameters, increasing model resolution, assimilating slant total electron data, and repeating these experiments during geomagnetic storms.

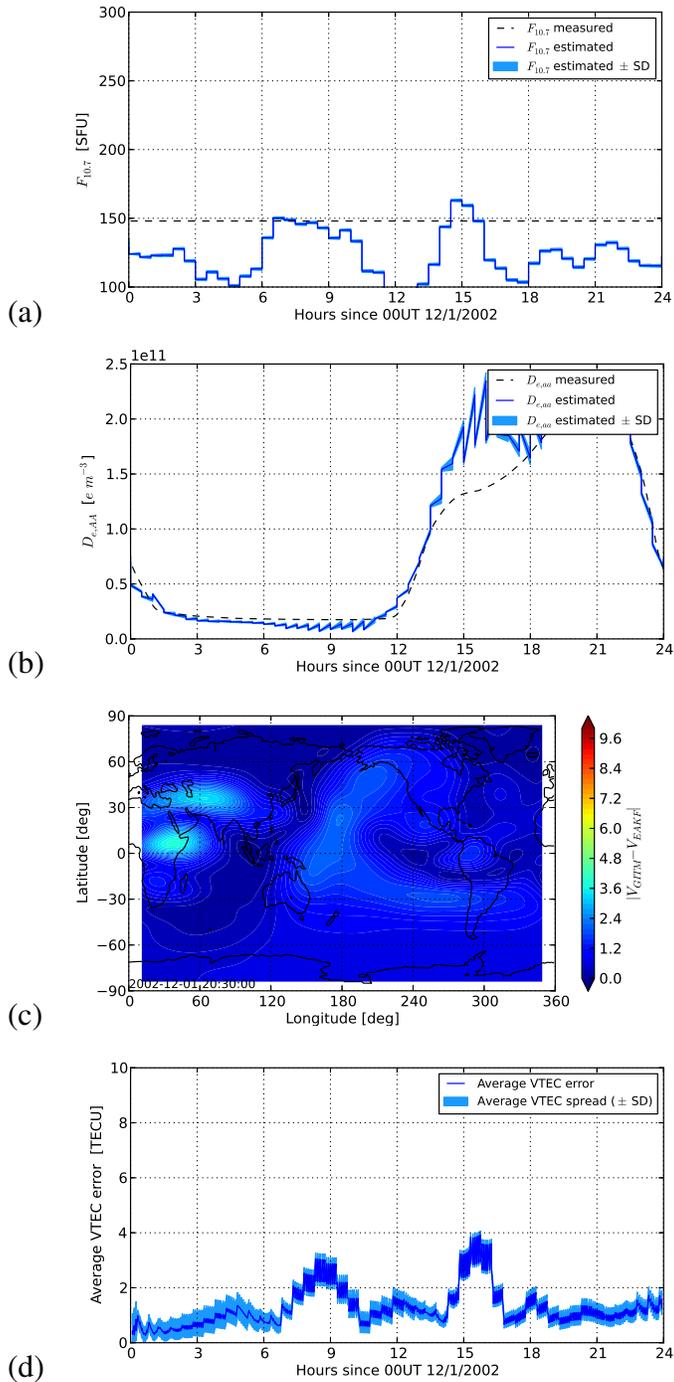


Figure 6.7: Assimilating simulated data with driver estimation. (a) shows the time evolution of $F_{10.7}$ estimate, (b) shows the time evolution of electron density at the (40.0°N, 80.0°W, 411 km) gridpoint, (c) shows the VTEC error snapshot at 22:30 UT on Dec. 1 2002, and (d) shows time evolution of average VTEC error. The $F_{10.7}$ estimate did not converge to the true value, but overall error levels were still comparable to the results of Section 6.3.2.

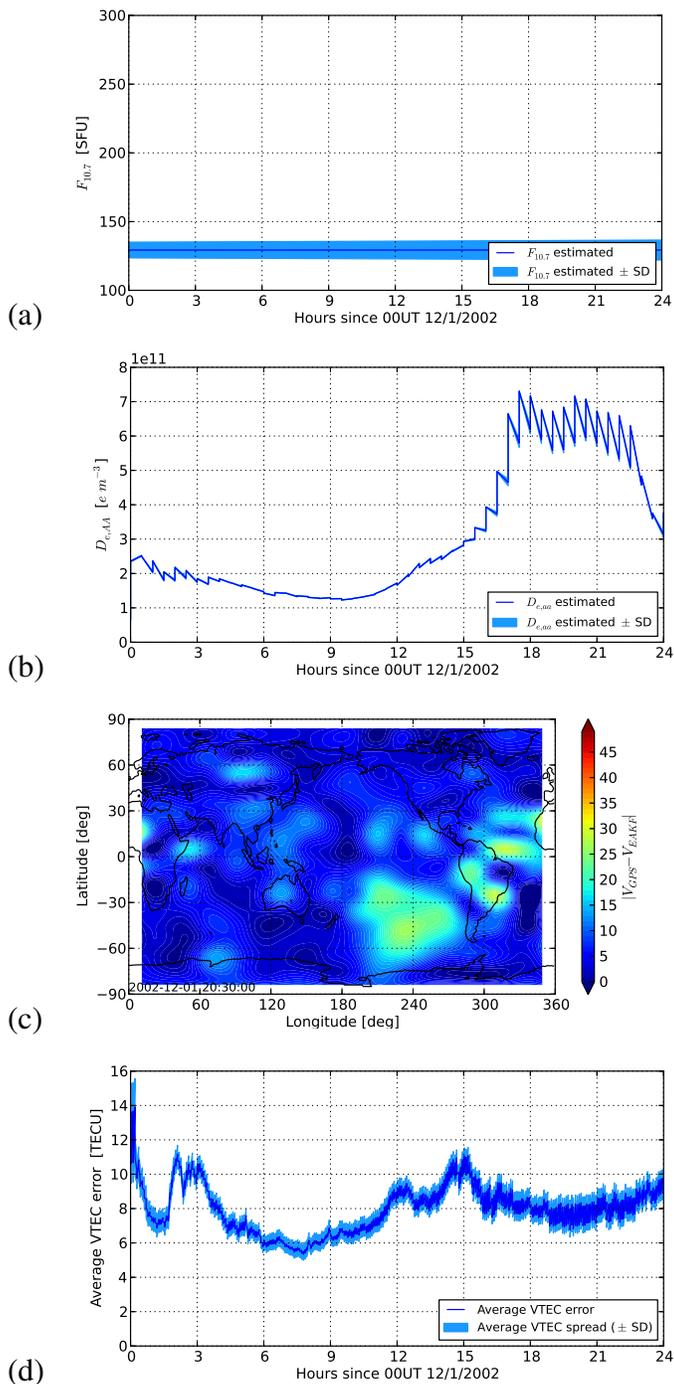


Figure 6.8: Assimilating real data without driver estimation. For plot-type descriptions see Figure 6.7. (c) confirms the rich/poor data region trend by showing low error values above the US and Europe and high values above South Pacific. Without driver estimation, error leveled out at about 8 TECU.

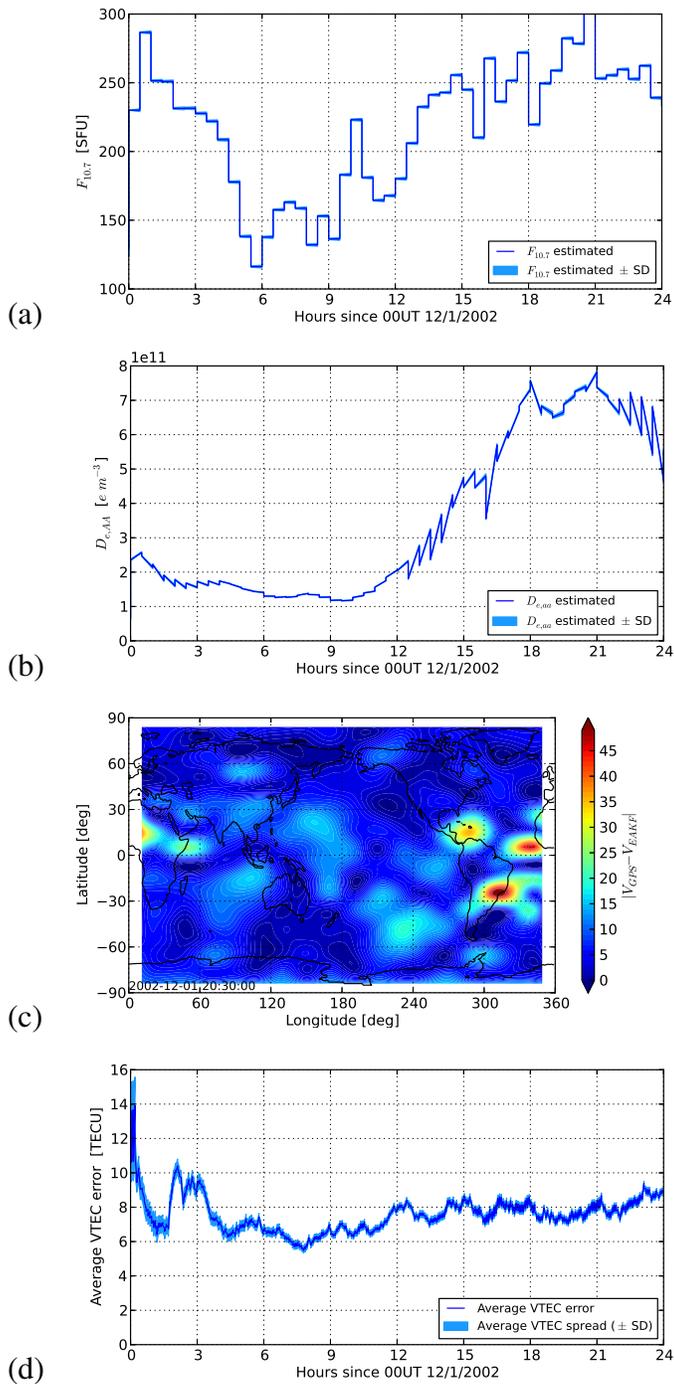


Figure 6.9: Assimilating real data with driver estimation. For plot-type descriptions see Figure 6.7. Adding driver estimation had an overall stabilizing effect to the data assimilation process, as can be seen in fewer upward spikes in (b), smaller errors over South Pacific Ocean in (c), and lower and flatter average error response in (d) when compared to not using driver estimation.

CHAPTER 7

Conclusions

This chapter describes the main contributions of this dissertation, but does so in a slightly nontraditional, retrospective way. In particular, instead of starting at the beginning of the dissertation and going forward, this chapter starts at the end and goes backward. The benefit of this direction reversal is twofold: it allows the earliest results to be reinterpreted from the standpoint of the latest results and it does the opposite of the introduction, thereby allowing to go back from specific to general topics (which is perhaps reminiscent of Crab Canon from [132]).

Accordingly, we start with Chapter 6, which demonstrated the assimilation process for a very special type of measurements (GPS total electron content) to improve ionospheric estimation accuracy of a space weather model (GITM). It showed that real data is best assimilated with driver estimation enabled, as drivers affect even the regions of the model where data is sparse. This insight can be carried over to the simulated data case, which at first seemed to suggest that $F_{10.7}$ estimation might not be useful, but now just seems like it requires greater initial ensemble spread. More general future research directions include assimilating slant total electron content data and total electron content occultation measurements between satellites.

Chapter 5 considered the case where a much smaller data set coming from a single satellite (CHAMP neutral density) was assimilated into GITM. A large part of this chapter was devoted to describing and validating the DART-GITM interface created for the pur-

poses of this study. In particular, it was found that during the dates in this study, CHAMP was almost in the terminator orbit, making it more difficult to estimate $F_{10.7}$ as compared to dates when CHAMP would fly directly above the subsolar point. Even in the presence of this difficulty, the assimilation was shown to perform well at removing the discrepancy between the data and the model. One major changes inspired by developments of chapter 6 is defining neutral density as a measurement inside the DART-GITM interface instead of making it a state variable. While this distinction might seem somewhat technical, it affects the source code implementation side of this project by reducing how much of DART and GITM source code needs to be modified to create another state variable. Arguably this reduction is offset by the need to write a forward operator relating the density of various GITM species to the neutral density, but the benefits of such an operator seem to outweigh the costs.

Another possible change to DART-GITM interface that became apparent at the later stages of the project described in Chapter 6 pertains mostly to GITM and deals with how GITM writes its output. Currently, unformatted binary files are used for GITM output and restart files. The alternative solution is to use netCDF files, which are designed with large atmospheric applications in mind and are much quicker to access.

Progressing backwards towards the beginning of the dissertation we encounter Chapter 4, which described a nonlinear application of RCAC. More precisely, it began with proving the technique on a linear approximation of the plant (much like chapters 5 and 6 started with simulated data) and then used the NMP zero of the linearization to control the nonlinear system. Perhaps the most obvious extension of this project inspired by space weather projects would be figuring out how nonlinear the linkage system is. In a way it is obvious that it is strongly nonlinear as shown by the equations of motion, but on the other hand it is not clear “how far” the system is from being linear. Essentially, it would be of interest to come up with a nonlinearity index (or use an existing one [133]), where a linear system would get an index of zero, the same linear system with a small nonlinearity added would

get a small positive value, and so on. This index could provide insight into how “easy” it is to control certain nonlinear system or class of nonlinear systems.

The next (or previous) chapter considered was Chapter 3, which described RCAC robustness to the NMP zero locations. In particular it was found that the algorithm is more robust to overestimating NMP zero locations than to underestimating them and that there exist inherent difficulty in controlling plants for which the NMP zero needs to be known. One possible extension of this project inspired by the space weather projects is to use real systems instead of randomized plants. This idea is inspired by the model and data libraries that the space weather community (and numeric weather prediction community in general) maintains. Similarly, it could be beneficial to have a publicly accessible database of real systems. This virtual library could have several divisions, either by application area or by plant characteristics. The more specimens such library would have, the more credibility the simulation results based on it would have.

Lastly, Chapter 2 introduced a method to improve robustness of the RCAC to the location of NMP zeros. In particular, this modification was found to improve transient performance and increase robustness to NMP zero locations for the command following and disturbance rejection problems. One possible extension of this work is to design a new version of RCAC around this convex constraint.

BIBLIOGRAPHY

- [1] Brockett, R. and Mesarovic, M., “The reproducibility of multivariable systems,” *Journal of Mathematical Analysis and Applications*, Vol. 11, 1965, pp. 548–563. 2
- [2] Wohltmann, H.-W., “Target path controllability of linear time-varying dynamical systems,” *IEEE Trans. Autom. Contr.*, Vol. 30, No. 1, Jan, pp. 84–87. 2
- [3] Silverman, L., “Inversion of multivariable linear systems,” *IEEE Trans. Autom. Contr.*, Vol. 14, No. 3, 1969, pp. 270–276. 2
- [4] Dorato, P., “On the inverse of linear dynamical systems,” *Systems Science and Cybernetics, IEEE Transactions on*, Vol. 5, No. 1, 1969, pp. 43–48. 2
- [5] Sain, M. and Massey, J., “Invertibility of linear time-invariant dynamical systems,” *IEEE Trans. Autom. Contr.*, Vol. 14, No. 2, 1969, pp. 141–149. 2
- [6] Willsky, A., “On the invertibility of linear systems,” *IEEE Trans. Autom. Contr.*, Vol. 19, No. 3, Jun, pp. 272–274. 2
- [7] Orner, P., “Construction of inverse systems,” *IEEE Trans. Autom. Contr.*, Vol. 17, No. 1, Feb, pp. 151–153. 2
- [8] Emre, E. and Huseyin, O., “Invertibility criteria for linear multivariable systems,” *IEEE Trans. Autom. Contr.*, Vol. 19, No. 5, Oct, pp. 609–610. 2
- [9] Hirschorn, R. M., “Invertibility of nonlinear control systems,” *SIAM Journal on Control and Optimization*, Vol. 17, No. 2, 1979, pp. 289–297. 3
- [10] Isidori, A. and Byrnes, C. I., “Output regulation of nonlinear systems,” *IEEE Trans. Autom. Contr.*, Vol. 35, No. 2, 1990, pp. 131–140. 3
- [11] Byrnes, C. I. and Isidori, A., “A frequency domain philosophy for nonlinear systems, with applications to stabilization and to adaptive control,” *Decision and Control, 1984. The 23rd IEEE Conference on*, Vol. 23, IEEE, 1984, pp. 1569–1573. 3
- [12] Devasia, S., Chen, D., and Paden, B., “Nonlinear inversion-based output tracking,” *IEEE Trans. Autom. Contr.*, Vol. 41, No. 7, 1996, pp. 930–942. 3
- [13] Huang, J., “Output regulation of nonlinear systems with nonhyperbolic zero dynamics,” *IEEE Trans. Autom. Contr.*, Vol. 40, No. 8, 1995, pp. 1497–1500. 3

- [14] Doyle III, F. J., Allgower, F., and Morari, M., "A normal form approach to approximate input-output linearization for maximum phase nonlinear SISO systems," *IEEE Trans. Autom. Contr.*, Vol. 41, No. 2, 1996, pp. 305–309. 3
- [15] Garcia, C. E. and Morari, M., "Internal model control. A unifying review and some new results," *Industrial & Engineering Chemistry Process Design and Development*, Vol. 21, No. 2, 1982, pp. 308–323. 3
- [16] George, K., Verhaegen, M., and Scherpen, J. M., "Stable inversion of MIMO linear discrete time nonminimum phase systems," *Proc. 7th Mediterranean Conference on Control and Automation*, 1999, pp. 267–281. 3
- [17] Hou, M. and Patton, R., "Optimal filtering for systems with unknown inputs," *IEEE Trans. Autom. Contr.*, Vol. 43, No. 3, Mar, pp. 445–449. 3
- [18] Hou, M. and Muller, P., "Disturbance decoupled observer design: A unified viewpoint," *IEEE Trans. Autom. Contr.*, Vol. 39, No. 6, 1994, pp. 1338–1341. 3
- [19] Kailath, T., "An innovations approach to least-squares estimation—Part I: Linear filtering in additive white noise," *IEEE Trans. Autom. Contr.*, Vol. 13, No. 6, 1968, pp. 646–655. 4
- [20] Anderson, B. D. and Moore, J. B., *Optimal filtering*, Vol. 11, Prentice-hall Englewood Cliffs, NJ, 1979. 4
- [21] Burgess, J. C., "Active adaptive sound control in a duct: A computer simulation," *The Journal of the Acoustical Society of America*, Vol. 70, 1981, pp. 715. 4
- [22] Sondhi, M. and Presti, A., "A self-adaptive echo canceller," *Bell Syst. Tech. J.*, Vol. 45, 1966, pp. 1851–1854. 4
- [23] Widrow, B., Glover Jr, J. R., McCool, J. M., Kaunitz, J., Williams, C. S., Hearn, R. H., Zeidler, J. R., Dong Jr, E., and Goodlin, R. C., "Adaptive noise cancelling: Principles and applications," *Proceedings of the IEEE*, Vol. 63, No. 12, 1975, pp. 1692–1716. 4
- [24] Eriksson, L., Allie, M., and Greiner, R., "The selection and application of an IIR adaptive filter for use in active sound attenuation," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, Vol. 35, No. 4, 1987, pp. 433–437. 4
- [25] Venugopal, R. and Bernstein, D. S., "Adaptive Disturbance Rejection Using AR-MARKOV System Representations," *IEEE Trans. Contr. Sys. Tech.*, Vol. 8, 2000, pp. 257–269. 4
- [26] Venugopal, R. and Bernstein, D. S., "Adaptive tracking using ARMARKOV/Toeplitz models," *American Control Conference, 1998. Proceedings of the 1998*, Vol. 1, IEEE, 1998, pp. 144–148. 4

- [27] Van Pelt, T., Venugopal, R., and Bernstein, D., “Experimental comparison of adaptive cancellation algorithms for active noise control,” *Control Applications, 1997., Proceedings of the 1997 IEEE International Conference on*, IEEE, 1997, pp. 559–564. 4
- [28] Ben-Amara, F., Kabamba, P. T., Ulsoy, A. G., Bernstein, D. S., et al., “A comparison of two adaptive algorithms for disturbance cancellation,” *Annual Reviews in Control*, Vol. 21, 1997, pp. 67–78. 4
- [29] Kuo, S. M. and Morgan, D. R., *Active Noise Control Systems*, Wiley, New York, 1996. 4
- [30] Hoagg, J. B. and Bernstein, D. S., “Discrete-time adaptive feedback disturbance rejection using a retrospective performance measure,” *Proc. ACTIVE 04*, Williamsburg, VA, 2004. 4
- [31] Santillo, M. A. and Bernstein, D. S., “A Retrospective Correction Filter for Discrete-time Adaptive Control of Nonminimum Phase Systems,” *Proc. Conf. Dec. Contr.*, Cancun, Mexico, December 2008, pp. 690–695. 4, 15
- [32] Hoagg, J. B. and Bernstein, D. S., “Cumulative retrospective cost adaptive control with RLS-based optimization,” *American Control Conference (ACC), 2010*, IEEE, 2010, pp. 4016–4021. 5, 33, 34, 37, 55, 56, 69
- [33] Ridley, A., Deng, Y., and Tth, G., “The global ionosphere–thermosphere model,” *Journal of Atmospheric and Solar-Terrestrial Physics*, Vol. 68, No. 8, 2006, pp. 839–864. 6, 73, 96
- [34] Toth, G., De Zeeuw, D. L., Gombosi, T. I., Manchester, W. B., Ridley, A. J., Sokolov, I. V., and Roussev, I. I., “Sun-to-thermosphere simulation of the 28–30 October 2003 storm with the Space Weather Modeling Framework,” *Space Weather*, Vol. 5, No. 6, 2007, pp. S06003. 7
- [35] Doherty, P., Coster, A. J., and Murtagh, W., “Space weather effects of October–November 2003,” *GPS Solutions*, Vol. 8, No. 4, 2004, pp. 267–271. 7
- [36] Wik, M., Pirjola, R., Lundstedt, H., Viljanen, A., Wintoft, P., and Pulkkinen, A., “Space weather events in July 1982 and October 2003 and the effects of geomagnetically induced currents on Swedish technical systems,” *Annales Geophysicae*, Vol. 27, Springer-Verlag GmbH, 2009, pp. 1775–1787. 7
- [37] Matsuo, T., Fedrizzi, M., Fuller-Rowell, T. J., and Codrescu, M. V., “Data assimilation of thermospheric mass density,” *Space Weather*, Vol. 10, May 2012, pp. S05002. 7, 72, 73, 91
- [38] Newman, L. K., “The NASA robotic conjunction assessment process: Overview and operational experiences,” *Acta Astronautica*, Vol. 66, No. 7, 2010, pp. 1253–1261. 7

- [39] Jazwinski, A. H., *Stochastic processes and filtering theory*, Dover Publications, 2007. 7
- [40] Julier, S. J. and Uhlmann, J. K., “New extension of the Kalman filter to nonlinear systems,” *AeroSense’97*, International Society for Optics and Photonics, 1997, pp. 182–193. 7
- [41] Evensen, G., “Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics,” *Journal of Geophysical Research*, Vol. 99, 1994, pp. 10,143–10,162. 7, 73
- [42] Evensen, G., “The ensemble Kalman filter: Theoretical formulation and practical implementation,” *Ocean dynamics*, Vol. 53, No. 4, 2003, pp. 343–367. 7
- [43] George, A., Butala, M., Frazin, R., Kamalabadi, F., and Bresler, Y., “Time-resolved cardiac CT reconstruction using the ensemble Kalman Filter,” *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on*, IEEE, 2008, pp. 1489–1492. 7
- [44] Poudel, D., Sandal, L., and Kvamsdal, S., “Analyzing Risk of Stock Collapse in a Fishery under Stochastic Profit Maximization,” *NHH Dept. of Finance & Management Science Discussion Paper*, , No. 4, 2012. 7
- [45] Hoffman, M. J., Greybush, S. J., John Wilson, R., Gyarmati, G., Hoffman, R. N., Kalnay, E., Ide, K., Kostelich, E. J., Miyoshi, T., and Szunyogh, I., “An ensemble Kalman filter data assimilation system for the Martian atmosphere: Implementation and simulation experiments,” *Icarus*, Vol. 209, No. 2, 2010, pp. 470–481. 7
- [46] Anderson, J. L., “An Ensemble Adjustment Kalman Filter for Data Assimilation,” *Monthly Weather Review*, Vol. 129, No. 12, Dec 2001, pp. 2884–2903. 8, 73, 77, 79
- [47] Anderson, J., Hoar, T., Raeder, K., Liu, H., Collins, N., Torn, R., and Avellano, A., “The Data Assimilation Research Testbed: A Community Facility,” *Bulletin of the American Meteorological Society*, Vol. 90, No. 9, Sep 2009, pp. 1283–1296. 8, 74, 96
- [48] Morozov, A. V., D’Amato, A. M., Hoagg, J. B., and Bernstein, D. S., “Retrospective cost adaptive control for nonminimum-phase systems with uncertain nonminimum-phase zeros using convex optimization,” *American Control Conference (ACC), 2011*, IEEE, 2011, pp. 1188–1193. 10
- [49] Hoagg, J. B. and Bernstein, D. S., “Nonminimum-Phase Zeros: Much to Do about Nothing,” *IEEE Contr. Sys. Mag.*, Vol. 27, No. June, 2007, pp. 45–57. 11
- [50] Santillo, M. A. and Bernstein, D. S., “Adaptive Control Based on Retrospective Cost Optimization,” *AIAA J. Guid. Contr. Dyn.*, Vol. 33, 2010, pp. 289–304. 11, 12

- [51] Hoagg, J. B. and Bernstein, D. S., “Cumulative Retrospective Cost Adaptive Control with RLS-Based Optimization,” *Proc. Amer. Contr. Conf.*, Baltimore, MD, June 2010, pp. 4016–4021. 11
- [52] Hoagg, J. B. and Bernstein, D. S., “Retrospective Cost Adaptive Control for Nonminimum-Phase Discrete-Time Systems Part 1: The Ideal Controller and Error System; Part 2: The Adaptive Controller and Stability Analysis,” *Proc. Conf. Dec. Contr.*, Atlanta, GA, December 2010, pp. 893–904. 12
- [53] Fledderjohn, M. S., Holzel, M. S., Morozov, A. V., Hoagg, J. B., and Bernstein, D. S., “On the Accuracy of Least Squares Algorithms for Estimating Zeros,” *Proc. Amer. Contr. Conf.*, Baltimore, MD, June 2010, pp. 3729–3734. 12, 56
- [54] Morozov, A. V., Hoagg, J. B., and Bernstein, D. S., “A Computational Study of the Performance and Robustness Properties of Retrospective Cost Adaptive Control,” *AIAA Guid. Nav. Contr. Conf.*, Toronto, August 2010, AIAA-2010-8011-911. 12, 32
- [55] Anderson, B. D. O., “Challenges of Adaptive Control—Past, Permanent, and Future,” *Annual Reviews of Control*, Vol. 32, 2008, pp. 123–135. 13
- [56] Anderson, B. D. O., “Failures of Adaptive Control Theory and Their Resolution,” *Comm. Information Systems*, Vol. 5, 2005, pp. 1–20. 13
- [57] Maciejowski, J. M., “Guaranteed stability with subspace methods,” *Sys. Contr. Lett.*, Vol. 26, 1995, pp. 153–156. 13
- [58] Lacy, S. L. and Bernstein, D. S., “Subspace Identification with Guaranteed Stability Using Constrained Optimization,” *IEEE Trans. Autom. Contr.*, Vol. 48, 2003, pp. 1259–1263. 13
- [59] Bernstein, D. S., *Matrix Mathematics*, Princeton University Press, 2nd ed., 2009. 19, 20
- [60] Grant, M. and Boyd, S., “CVX: Matlab Software for Disciplined Convex Programming, version 1.21,” <http://cvxr.com/cvx>, Aug. 2010. 22
- [61] Grant, M. and Boyd, S., “Graph implementations for nonsmooth convex programs,” *Recent Advances in Learning and Control*, edited by V. Blondel, S. Boyd, and H. Kimura, Lecture Notes in Control and Information Sciences, Springer-Verlag Limited, 2008, pp. 95–110, http://stanford.edu/~boyd/graph_dcp.html. 22
- [62] Narendra, K. S. and Annaswamy, A. M., *Stable Adaptive Systems*, Prentice Hall, 1989. 33
- [63] Åström, K. J. and Wittenmark, B., *Adaptive Control*, Addison-Wesley, 2nd ed., 1995. 33, 38
- [64] Ioannou, P. and Sun, J., *Robust Adaptive Control*, Prentice Hall, 1996. 33

- [65] Goodwin, G. C. and Sin, K. S., *Adaptive Filtering, Prediction, and Control*, Prentice Hall, 1984. 33, 38
- [66] Tao, G., *Adaptive Control Design and Analysis*, Wiley, 2003. 33, 38
- [67] Fuentes, R. J. and Balas, M. J., “Direct adaptive disturbance accommodation,” *Proc. IEEE Conf. Dec. Contr.*, Sydney, Australia, December 2000, pp. 4921–4925. 33
- [68] Hoagg, J. B. and Bernstein, D. S., “Direct Adaptive Command Following and Disturbance Rejection for Minimum Phase Systems with Unknown Relative Degree,” *Int. J. Adaptive Contr. Signal Processing*, Vol. 21, 2007, pp. 49–75. 33
- [69] Rohrs, C., Valavani, L., Athans, M., and Stein, G., “Robustness of continuous-time adaptive control algorithms in the presence of unmodeled dynamics,” *IEEE Trans. Autom. Contr.*, Vol. 30, 1985, pp. 881–889. 33
- [70] Johansson, R., “Global Lyapunov stability and exponential convergence of direct adaptive control,” *Int. J. Contr.*, Vol. 50, 1989, pp. 859–869. 33
- [71] Akhtar, S. and Bernstein, D. S., “Logarithmic Lyapunov functions for direct adaptive stabilization with normalized adaptive laws,” *Int. J. Contr.*, Vol. 77, 2004, pp. 630–638. 33
- [72] Akhtar, S. and Bernstein, D. S., “Lyapunov-stable discrete-time model reference adaptive control,” *Int. J. Adaptive Contr. Signal Proc.*, Vol. 19, 2005, pp. 745–767. 33
- [73] Goodwin, G. C., Ramadge, P. J., and Caines, P. E., “Discrete-time multivariable adaptive control,” *IEEE Trans. Autom. Contr.*, Vol. 25, 1980, pp. 449–456. 33
- [74] Hoagg, J. B., Santillo, M. A., and Bernstein, D. S., “Discrete-Time Adaptive Command Following and Disturbance Rejection with Unknown Exogenous Dynamics,” *IEEE Trans. Autom. Contr.*, Vol. 53, 2008, pp. 912–928. 33, 37
- [75] Bayard, D. S., “Extended horizon liftings for stable inversion of nonminimum-phase systems,” *IEEE Trans. Autom. Contr.*, Vol. 39, 1994, pp. 1333–1338. 33
- [76] Venugopal, R. and Bernstein, D. S., “Adaptive disturbance rejection using AR-MARKOV/Toeplitz Models,” *IEEE Trans. Contr. Sys. Tech.*, Vol. 8, 2000, pp. 257–269. 33, 34
- [77] Santillo, M. A. and Bernstein, D. S., “A retrospective correction filter for discrete-time adaptive control of nonminimum-phase systems,” *Proc. IEEE Conf. Dec. Contr.*, Cancun, Mexico, 2008, pp. 690–695. 33, 34
- [78] Hoagg, J. B., Lacy, S. L., and Bernstein, D. S., “Broadband adaptive disturbance rejection for a deployable optical telescope testbed,” *Proc. Amer. Contr. Conf.*, Portland, OR, 2005, pp. 4953–4958. 34

- [79] Holzel, M. S., Santillo, M. A., Hoagg, J. B., and Bernstein, D. S., “Adaptive Control of the NASA Generic Transport Model Using Retrospective Cost Optimization,” *Proc. AIAA Guid. Nav. Contr. Conf.*, August 2009, AIAA-2009-5616. 34
- [80] Fledderjohn, M. S., Cho, Y.-C., Hoagg, J. B., Santillo, M. A., Shyy, W., and Bernstein, D. S., “Retrospective Cost Adaptive Flow Control Using a Dielectric Barrier Discharge Actuator,” *Proc. AIAA Guid. Nav. Contr. Conf.*, August 2009, AIAA-2009-5857. 34
- [81] Morozov, A. V., Hoagg, J. B., and Bernstein, D. S., “Retrospective Adaptive Control of a Planar Multilink Arm with Nonminimum-Phase Zeros,” *Proc. Conf. Dec. Contr.*, Atlanta, GA, December 2010. 54
- [82] Isaacs, M. W., Hoagg, J. B., Morozov, A. V., and Bernstein, D. S., “A numerical study on controlling a nonlinear multilink arm using a retrospective cost model reference adaptive controller,” *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, IEEE, 2011, pp. 8008–8013. 54
- [83] Kwakernaak, H. and Sivan, R., *Linear Optimal Control Systems*, Wiley, 1972. 55
- [84] Skogestad, S. and Postlethwaite, I., *Multivariable Feedback Control*, Wiley, New York, 2nd ed., 2005. 55
- [85] Hoagg, J. B. and Bernstein, D. S., “Nonminimum-Phase Zeros: Much to Do about Nothing,” *IEEE Contr. Sys. Mag.*, Vol. 27, June 2007, pp. 45–57. 55
- [86] Ioannou, P. and Fidan, B., *Adaptive Control Tutorial*, SIAM, Philadelphia, 2006. 55
- [87] Santillo, M. A. and Bernstein, D. S., “Adaptive Control Based on Retrospective Cost Optimization,” *AIAA J. Guid. Contr. Dyn.*, Vol. 33, 2010, pp. 289–304. 55
- [88] Hong, J. and Bernstein, D. S., “Bode Integral Constraints, Colocation, and Spillover in Active Noise and Vibration Control,” *IEEE Trans. Contr. Sys. Tech.*, Vol. 6, 1998, pp. 111–120. 55
- [89] Hoagg, J. B., Chandrasekar, J., and Bernstein, D. S., “On the Zeros, Initial Undershoot, and Relative Degree of Lumped-Mass Structures,” *ASME J. Dynamical Systems, Measurement, and Contr.*, Vol. 129, 2007, pp. 493–502. 55, 56
- [90] Miu, D. K., *Mechatronics*, Springer-Verlag, New York, 1993. 55
- [91] Maslen, E. H., “Positive real zeros in flexible beams,” *Shock and Vibration*, Vol. 2, 1995, pp. 429–435. 55
- [92] Abramovitch, D. and Franklin, G., “A Brief History of Disk Drive Control,” *IEEE Contr. Sys. Mag.*, Vol. 22, no. 3, 2002, pp. 28–42. 55
- [93] Rigney, B. P., Pao, L. Y., and Lawrence, D. A., “Nonminimum Phase Dynamic Inversion for Settle Time Applications,” *IEEE Trans. Contr. Sys. Tech.*, Vol. 17, 2009, pp. 989–1005. 55

- [94] Somolinos, J., Feliu, V., and Sánchez, L., “Design, dynamic modelling and experimental validation of a new three-degree-of-freedom flexible arm,” *Mechatronics*, Vol. 12, No. 7, 2002, pp. 919–948. 69
- [95] Moberg, S., Öhr, J., and Gunnarsson, S., “A benchmark problem for robust control of a multivariable nonlinear flexible manipulator,” *Proceedings of 17th IFAC World Congress*, 2008, pp. 1206–1211. 69
- [96] Morozov, A. V., Ridley, A. J., Bernstein, D. S., Collins, N., Hoar, T. J., and Anderson, J. L., “Data Assimilation and Driver Estimation for the Global Ionosphere-Thermosphere Model Using the Ensemble Adjustment Kalman Filter,” *Journal of Atmospheric and Solar-Terrestrial Physics*, 2013, in preparation. 71
- [97] Storz, M., Bowman, B., Branson, M., Casali, S., and Tobiska, W., “High accuracy satellite drag model (HASDM),” *Advances in Space Research*, Vol. 36, No. 12, 2005, pp. 2497–2505. 71, 72
- [98] Reigber, C., Lühr, H., and Schwintzer, P., “CHAMP mission status,” *Advances in Space Research*, Vol. 30, No. 2, 2002, pp. 129–134. 72
- [99] Tapley, B., Bettadpur, S., Watkins, M., and Reigber, C., “The gravity recovery and climate experiment: Mission overview and early results,” *Geophys. Res. Lett.*, Vol. 31, No. 9, 2004, pp. L09607. 72
- [100] Millward, G., Müller-Wodarg, I., Aylward, A., Fuller-Rowell, T., Richmond, A., and Moffett, R., “An investigation into the influence of tidal forcing on F region equatorial vertical ion drift using a global ionosphere-thermosphere model with coupled electrodynamics,” *Journal of Geophysical Research. A. Space Physics*, Vol. 106, 2001, pp. 24. 72
- [101] Millward, G., Moffett, R., Quegan, S., and Fuller-Rowell, T., “A coupled thermosphere-ionosphere-plasmasphere model (CTIP),” *Handbook of Ionospheric Models*, 1996, pp. 239–279. 72
- [102] Kersley, L., Pryse, S., Denton, M., Bust, G., Fremouw, E., Secan, J., Jakowski, N., and Bailey, G., “Radio tomographic imaging of the northern high-latitude ionosphere on a wide geographic scale,” *Radio Science*, Vol. 40, No. RS5003, 2005, pp. 1–9. 72
- [103] Schunk, R., Scherliess, L., Sojka, J., Thompson, D., Anderson, D., Codrescu, M., Minter, C., Fuller-Rowell, T., Heelis, R., Hairston, M., et al., “Global assimilation of ionospheric measurements (GAIM),” *Radio Science*, Vol. 39, No. 1, 2004, pp. RS1S02. 72
- [104] Scherliess, L., Schunk, R., Sojka, J., and Thompson, D., “Development of a physics-based reduced state Kalman filter for the ionosphere,” *Radio Science*, Vol. 39, No. 1, 2004, pp. RS1S04. 72

- [105] Mandrake, L., Wilson, B., Wang, C., Hajj, G., Mannucci, A., and Pi, X., “A performance evaluation of the operational Jet Propulsion Laboratory/University of Southern California global assimilation ionospheric model (JPL/USC GAIM),” *Journal of geophysical research*, Vol. 110, No. A12, 2005, pp. A12306. 73
- [106] Perbos, J., Escudier, P., Parisot, F., Zaouche, G., Vincent, P., Menard, Y., Manon, F., Kunstmann, G., Royer, D., and Fu, L., “Jason-1: Assessment of the System Performances Special Issue: Jason-1 Calibration/Validation,” *Marine Geodesy*, Vol. 26, No. 3-4, 2003, pp. 147–157. 73
- [107] Kim, I., Pawlowski, D., Ridley, A., and Bernstein, D., “Localized data assimilation in the ionosphere-thermosphere using a sampled-data unscented Kalman filter,” *American Control Conference, 2008*, IEEE, 2008, pp. 1849–1854. 73
- [108] Kim, I., *Large scale data assimilation with application to the ionosphere-thermosphere*, ProQuest, 2008. 73
- [109] Li, J. and Liu, H., “Improved hurricane track and intensity forecast using single field-of-view advanced IR sounding measurements,” *Geophys. Res. Lett.*, Vol. 36, 2009, pp. L11813. 74
- [110] Stensrud, D. and Gao, J., “Importance of horizontally inhomogeneous environmental initial conditions to ensemble storm-scale radar data assimilation and very short-range forecasts,” *Monthly Weather Review*, Vol. 138, No. 4, 2010, pp. 1250–1272. 74
- [111] Khade, V., Hansen, J., Reid, J., and Westphal, D., “Ensemble filter based estimation of spatially distributed parameters in a mesoscale dust model: experiments with simulated and real data,” *Atmospheric Chemistry and Physics Discussions*, Vol. 12, 2012. 74
- [112] Lee, I. T., Matsuo, T., Richmond, A. D., Liu, J. Y., Wang, W., Lin, C. H., Anderson, J. L., and Chen, M. Q., “Assimilation of FORMOSAT-3/COSMIC electron density profiles into a coupled thermosphere/ionosphere model using ensemble Kalman filtering,” *Journal of Geophysical Research*, Vol. 117, No. A10, 10 2012. 74
- [113] Deng, Y., Richmond, A., Ridley, A., and Liu, H., “Assessment of the non-hydrostatic effect on the upper atmosphere using a general circulation model (GCM),” *Geophysical Research Letters*, Vol. 35, No. 1, 2008, pp. L01104. 74
- [114] Sontag, E., “Contractive systems with inputs,” *Perspectives in Mathematical System Theory, Control, and Signal Processing*, 2010, pp. 217–228. 74
- [115] Russo, G., Di Bernardo, M., and Sontag, E., “Global entrainment of transcriptional systems to periodic inputs,” *PLoS computational biology*, Vol. 6, No. 4, 2010, pp. e1000739. 74
- [116] Pawlowski, D. and Ridley, A., “Modeling the thermospheric response to solar flares,” *Journal of Geophysical Research*, Vol. 113, No. A10, 2008, pp. A10309. 75

- [117] Sutton, E., “Normalized force coefficients for satellites with elongated shapes,” *Journal of Spacecraft and Rockets*, Vol. 46, No. 1, 2009, pp. 112–116. 76
- [118] Sutton, E., Nerem, R., and Forbes, J., “Density and winds in the thermosphere deduced from accelerometer data,” *Journal of Spacecraft and Rockets*, Vol. 44, No. 6, 2007, pp. 1210–1219. 76
- [119] Yiğit, E. and Ridley, A., “Effects of high-latitude thermosphere heating at various scale sizes simulated by a nonhydrostatic global thermosphere–ionosphere model,” *Journal of Atmospheric and Solar-Terrestrial Physics*, Vol. 73, No. 5, 2011, pp. 592–600. 76
- [120] Tarantola, A., “Inverse problem theory and methods for model parameter estimation,” *Siam, Philadelphia*, 2005. 79
- [121] Anderson, J. L. and Anderson, S. L., “A Monte Carlo Implementation of the Nonlinear Filtering Problem to Produce Ensemble Assimilations and Forecasts,” *Monthly Weather Review*, Vol. 127, No. 12, Dec 1999, pp. 2741–2758. 81
- [122] Godinez, H. and Koller, J., “Localized adaptive inflation in ensemble data assimilation for a radiation belt model,” *Space Weather: The International Journal of Research and Applications*, Vol. 10, No. null, 2012, pp. S08001. 82
- [123] Anderson, J. and Collins, N., “Scalable implementations of ensemble filter algorithms for data assimilation,” *Journal of Atmospheric and Oceanic Technology*, Vol. 24, No. 8, 2007, pp. 1452–1463. 83
- [124] Gaspari, G. and Cohn, S. E., “Construction of correlation functions in two and three dimensions,” *Quarterly Journal of the Royal Meteorological Society*, Vol. 125, No. 554, 1999, pp. 723–757. 83
- [125] Casella, G. and Berger, R., *Statistical inference*, Duxbury advanced series in statistics and decision sciences, Thomson Learning, 2002. 86
- [126] Knight, K., *Mathematical Statistics*, Texts in Statistical Science Series, Taylor & Francis, 2010. 86
- [127] Coco, D. S., Coker, C., Dahlke, S. R., and Clynch, J. R., “Variability of GPS satellite differential group delay biases,” *Aerospace and Electronic Systems, IEEE Transactions on*, Vol. 27, No. 6, 1991, pp. 931–938. 96
- [128] Burešová, D. and Laštovička, J., “Pre-storm enhancements of foF2 above Europe,” *Advances in Space Research*, Vol. 39, No. 8, 2007, pp. 1298–1303. 96
- [129] Ciruolo, L. and Spalla, P., “Comparison of ionospheric total electron content from the Navy Navigation Satellite System and the GPS,” *Radio Science*, Vol. 32, 1997, pp. 1071–1080. 97

- [130] Arikan, F., Erol, C., and Arikan, O., "Regularized estimation of vertical total electron content from Global Positioning System data," *Journal of Geophysical Research*, Vol. 108, No. A12, 2003, pp. 1469. 97
- [131] Rideout, W. and Coster, A., "Automated GPS processing for global total electron content data," *GPS Solutions*, Vol. 10, No. 3, 2006, pp. 219–228. 97
- [132] Hofstadter, D. R., *Godel, Escher, Bach*, Basic books, 2000. 108
- [133] Sinclair, A. J., Hurtado, J. E., and Junkins, J. L., "Nonlinearity index of the cayley form," *The Journal of the Astronautical Sciences*, Vol. 54, No. 3, 2006, pp. 619–634. 109