

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

Alexey V. Morozov*, Jesse B. Hoagg† and Dennis S. Bernstein‡

Department of Aerospace Engineering, The University of Michigan, Ann Arbor, MI 48109-2140

We present 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. We explore computationally the algorithm's performance and robustness in the presence of errors in the required modeling information.

I. Introduction

Plants with nonminimum-phase zeros present a major challenge to direct adaptive control. In fact, many direct adaptive control methodologies rely on the assumption that the plant is minimum phase.¹⁻⁵ 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.^{6,7} In the present paper, we review a discrete-time adaptive controller that addresses these challenges. Next, 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.^{2,4,8-10} In addition, there are adaptive control algorithms that are unique to discrete-time.^{4,11-13} In Refs. 4, 11, discrete-time adaptive control laws are presented for stabilization and command following of minimum-phase systems. An extension is given in Ref. 12, which addresses the combined stabilization, command following, and disturbance rejection problem. Note that the results of Refs. 4, 11, 12 are restricted to minimum-phase systems.

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.¹⁴⁻¹⁶ Furthermore, 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. 12, while extensions to the nonminimum-phase case are described in Refs. 15, 16. Retrospective cost adaptive control uses a retrospective performance measure, in which the performance measure 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 Refs. 14, 15 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. 16 are derived by minimizing a cumulative retrospective cost function, which is a function of the retrospective

*Graduate student.

†Postdoctoral Research Fellow, AIAA Member.

‡Professor, AIAA Member.

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. 17, the NASA generic transport model in Ref. 18, and flow control problems in Ref. 19.

The goal of the present paper 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. 16. In particular, we computationally explore the algorithm's behavior in the presence of inaccuracies in the required model information. In Section IV, we explore computationally the controller's sensitivity to the accuracy of the required model information, specifically, the nonminimum-phase zeros from the control to the performance. In Section V, we explore the effect of sample frequency on the controller's performance for sampled-data systems.

II. Problem Formulation

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

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

$$y(k) = Cx(k) + Du(k) + D_2w(k), \quad (2)$$

$$z(k) = E_1x(k) + E_2u(k) + E_0w(k), \quad (3)$$

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 performance z are available for feedback; however, we do not assume that a direct measurement of the exogenous signal w is available; however, a measurement of w can be used if it is 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_1x follow the command signal $-E_0w$. 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_1x . 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.

We represent (1) and (3) as the time-series model from u and w to z given by

$$z(k) = \sum_{i=1}^n -\alpha_i z(k-i) + \sum_{i=d}^n \beta_i u(k-i) + \sum_{i=0}^n \gamma_i w(k-i), \quad (4)$$

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 $\beta_d = H_d$.

III. Review of the Cumulative Retrospective Cost Adaptive Controller

In this section, we review the cumulative retrospective cost adaptive controller presented in Ref. 16. Consider a strictly proper time-series controller of order n_c , such that the control $u(k)$ is given by

$$u(k) = \sum_{i=1}^{n_c} M_i(k)u(k-i) + \sum_{i=1}^{n_c} N_i(k)y(k-i), \quad (5)$$

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 (5) can be expressed as

$$u(k) = \theta(k)\phi(k), \quad (6)$$

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

$$\phi(k) \triangleq \begin{bmatrix} y(k-1) \\ \vdots \\ y(k-n_c) \\ u(k-1) \\ \vdots \\ u(k-n_c) \end{bmatrix} \in \mathbb{R}^{n_c(l_u+l_y)}. \quad (7)$$

Next, we define the retrospective performance

$$\hat{z}(\hat{\theta}, k) \triangleq z(k) + \sum_{i=d}^{\nu} \bar{\beta}_i [\hat{\theta} - \theta(k-i)] \phi(k-i), \quad (8)$$

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 .¹⁶ In this paper, we let $\beta_d, \dots, \beta_\nu$ 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. 16. 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 $\bar{\beta}_d = H_d$.

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

$$\hat{z}(\hat{\Theta}, k) = z(k) + \sum_{i=d}^{\nu} \Phi_i^T(k) [\hat{\Theta} - \Theta(k-i)] = z(k) - \sum_{i=d}^{\nu} \Phi_i^T(k) \Theta(k-i) + \Psi^T(k) \hat{\Theta}, \quad (9)$$

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

$$J(\hat{\Theta}, k) \triangleq \sum_{i=0}^k \lambda^{k-i} \hat{z}^T(\hat{\Theta}, i) R \hat{z}(\hat{\Theta}, i) + \lambda^k (\hat{\Theta} - \Theta(0))^T Q (\hat{\Theta} - \Theta(0)), \quad (10)$$

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 (10) is minimized by a recursive least-squares (RLS) algorithm with a forgetting factor.^{2,4,5} Therefore, $J(\hat{\Theta}, k)$ is minimized by the adaptive law

$$\Theta(k+1) = \Theta(k) - P(k)\Psi(k) [\lambda R^{-1} + \Psi^T(k)P(k)\Psi(k)]^{-1} z_R(k), \quad (11)$$

$$P(k+1) = \frac{1}{\lambda} P(k) - \frac{1}{\lambda} P(k)\Psi(k) [\lambda R^{-1} + \Psi^T(k)P(k)\Psi(k)]^{-1} \Psi^T(k)P(k), \quad (12)$$

where $P(0) = Q^{-1}$, $\Theta(0) \in \mathbb{R}^{n_c l_u(l_y+l_u)}$, and the retrospective performance measure $z_R(k) \triangleq \hat{z}(\Theta(k), k)$. Note that $z_R(k)$ is computable from (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 (11), (12), and

$$u(k) = \theta(k)\phi(k) = \text{vec}^{-1}(\Theta(k))\phi(k). \quad (13)$$

The key feature of the retrospective cost adaptive control (RCAC) algorithm is the use of the retrospective performance (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 (11)-(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.

IV. Sensitivity to the Accuracy of the Nonminimum-Phase Zero Estimates

In this section, we investigate the sensitivity of the cumulative retrospective cost adaptive controller (11)-(13) to the accuracy of the estimates of the nonminimum-phase zeros (i.e., the parameters $\bar{\beta}_d, \dots, \bar{\beta}_\nu$). We systematically test RCAC's performance with a collection of stable nonminimum-phase plants.

For all examples in this section, we consider the single-input, single-output command following problem, where the control objective is to have the plant output $Cx(k) = E_1x(k)$ follow a sinusoid of frequency $\pi/6$ radians per sample with amplitude 10. Note that the frequency of the sinusoidal command is not assumed to be known. RCAC 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.

IV.A. Sensitivity with one nonminimum-phase zero

First, we investigate RCAC's performance when used to control a stable, third-order system with one nonminimum-phase (NMP) zero whose location is uncertain. More specifically, consider the single-input, single-output plant with poles at 0.5, 0.6, and 0.7; relative degree $d = 2$; first nonzero Markov parameter $H_d = 1$; and one NMP zero. We investigate the performance for a range of NMP zero locations, specifically, the NMP zero is varied from 1.1 to 8.1. The controller (11)-(13) requires estimate of d , H_d , and the NMP zero. We assume that d and H_d are known exactly, while the estimate of the NMP zero is varied from 1.1 to 8.1 (regardless of the actual value of the NMP zero).

Figure 1 shows the logarithm of the transient performance (i.e., $\max |z(k)|$ over the entire simulation) and the logarithm of the steady-state performance (i.e., $\max |z(k)|$ over the last 100 steps of the simulation) for various locations of the NMP zero and various estimates the NMP zero used by RCAC. Note that the diagonal of Figure 1(a) and Figure 1(b) corresponds to the nominal cases, that is, the cases where the estimate of the NMP zero is the actual NMP zero location, whereas the off-diagonals correspond to the off-nominal cases, that is, the cases where the estimate of the NMP zero is not equal to the actual NMP zero location. Figure 1 suggests that RCAC is more robust to overestimating the NMP zero location than it is to underestimating the NMP zero location. Furthermore, Figure 1 suggests that RCAC is more robust to errors in the NMP zero estimate as magnitude of the NMP zero increases. Note that the white portions of Figure 1 correspond to the cases where the closed-loop system is unstable; these cases occur when the estimate of the NMP zero is much larger or much smaller than the actual NMP zero location.

Next, we explore RCAC's performance on a collection of stable second-order systems, where there is uncertainty in the estimate of the NMP zero. We consider single-input, single-output plants with a single NMP zero located at 2, $d = 1$, $H_d = 1$, and two poles located at the same point along the real axis. In particular, we consider 61 plants with this structure where the double real poles are evenly spaced on the interval $[-1, 1]$.

For each of the 61 plants, we let the estimate of the NMP zero (which is located at 2) vary from 1 to 12 and simulate the closed-loop system. Figure 2 shows the logarithm of the transient performance and the logarithm of the steady-state performance for various locations of the double real poles and various estimates of the NMP location. This figure suggests that the plants with positive double poles are, in general, more robust to errors in the NMP zero estimate than the plants with negative double poles. Note that the plant

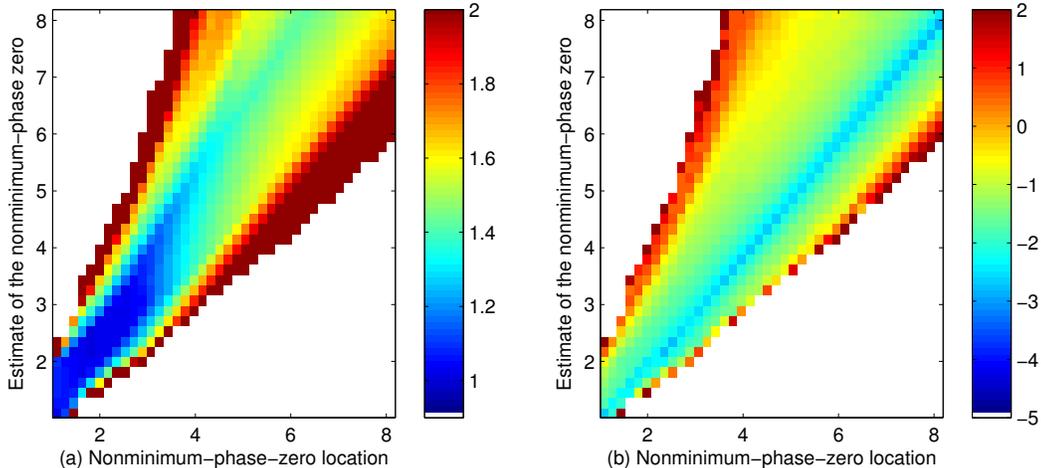


Figure 1. A single-input, single-output plant with poles at 0.5, 0.6, and 0.7; relative degree $d = 2$; first nonzero Markov parameter $H_d = 1$; and one NMP zero is connected in feedback with the adaptive controller (11)-(13). The NMP zero is varied from 1.1 to 8.1. We assume that d and H_d are known exactly, while the estimate of the NMP zero is varied from 1.1 to 8.1. The x -axis corresponds to the actual location of the NMP zero, while the y -axis corresponds to the estimate of the NMP zero used by RCAC. The color in (a) is the logarithm of the transient performance, whereas the color in (b) is the logarithm of the steady state performance. Figure 1 suggests that RCAC is more robust to overestimating the NMP zero location than it is to underestimating the NMP zero location.

with a double pole at -1 appears to be most sensitive to the accuracy of the NMP zero estimate; however, such a plant is unlikely to occur in a sampled-data system because a plant with a double pole at -1 occurs only if a continuous-time system with poles at $\pm j\omega$ is sampled with a sample frequency $\omega/2$.

The previous example explored second-order plants with double poles on the real axis; in this example, we extend the results to the case where a plant has two complex conjugate poles. More specifically, we consider single-input, single-output plants with a single NMP zero located at 2, $d = 1$, $H_d = 1$, and a pair of stable complex conjugate poles. In particular, we consider 2601 plants with this structure where the plant poles are located at $re^{\pm j\theta}$, where r is varied from 0 to 1 with 51 increments and θ is varied from 0 to π with 51 increments.

For each of the 2601 plants, we let the estimate of the NMP zero (which is located at 2) vary from 1.6 to 3.6 and simulate the closed-loop system. Figures 3 and 4 show the maximum and minimum values of the estimate of the NMP zero, respectively, for which the closed-loop adaptive system is stable. The x -axes and y -axes of Figures 3 and 4 correspond to the location of the plant poles in the complex plane. Furthermore, the colors used in the plots correspond to the maximum and minimum values of the NMP zero estimate for which the closed-loop adaptive system is stable. These figures extend the results shown in Figure 2 and suggest that the plants with poles that have negative real parts are more sensitive to error in the NMP zero estimate. Note that the majority of the complex plane is colored in white, which indicates that these closed-loop systems are stable over the entire range of NMP zero estimates that were tested in this example.

IV.B. Sensitivity with one nonminimum-phase zero for randomly generated systems

In this section, we investigate RCAC's performance when used to control randomly generated plants with varying system order, relative degree, pole locations, zero locations, and first nonzero Markov parameter. Specifically, a set of 50 stable, single-input, single-output plants with one NMP 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 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 NMP 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, RCAC requires estimates of d , H_d , and the location of the NMP zero. Next, for each of the 50 plants, we let the estimate of the NMP zero (which is located at 2) vary from 1.4 to 3.8 and simulate the closed-loop system. Figure 5 is a histogram showing the percent

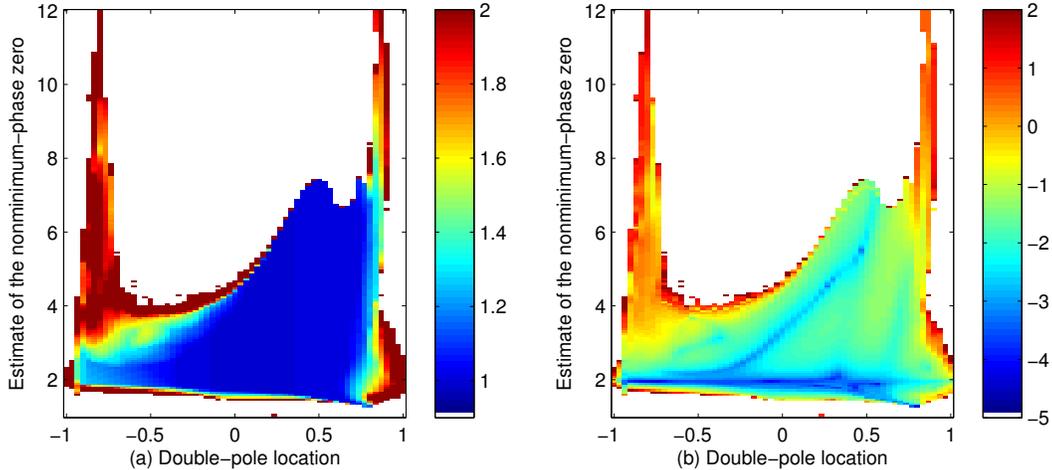


Figure 2. For 61 second-order plants with double poles along the real axis, we let the estimate of the NMP zero (which is located at 2) vary from 1 to 12 and simulate the closed-loop system. The x -axis corresponds to the location of the double poles, while the y -axis corresponds to the estimate of the NMP zero used by RCAC. The colors in (a) is the logarithm of the transient performance, whereas the colors in (b) is the logarithm of the steady-state performance. This figure suggests that the plants with positive double poles are, in general, more robust to errors in the NMP zero estimate than the plants with negative double poles.

of the 50 plants that have unstable closed-loop responses (i.e., the closed-loop performance is unbounded) as the estimate of the NMP zero varies from 1.4 to 3.8. Note that all 50 plants result in stable closed-loop responses when the estimate of the NMP zero is set at the actual value 2. Furthermore, Figure 5 indicates that the stability of the closed-loop system is less sensitive to overestimating the location of the NMP zero than it is to underestimating the location of the NMP zero.

Next, we explore the impact of system order and relative degree on RCAC's robustness to the accuracy of the NMP zero estimate. We consider 10 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 NMP zero at 2. Next, for plant, we let the estimate of the NMP zero (which is located at 2) vary from 1.4 to 2.6 and simulate the closed-loop system.

Figure 6 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 NMP zero varies from 1.4 to 2.6. Note that all plants result in stable closed-loop responses when the estimate of the NMP zero is set at the actual value 2. Figure 6 suggests that the stability of the closed-loop is less sensitive to the NMP zero estimate for plants with either low order or high relative degree. Accordingly, the cases that are least sensitive to the NMP zero estimate are situated on the diagonal of Figure 6.

Next, we explore, in more detail, the 50 randomly generated second-order systems shown in the bottom right of Figure 6. For these 50 randomly generated plants, we let the estimate of the NMP zero (which is located at 2) vary from 1.4 to 10 and simulate the closed-loop system. Figure 7 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 NMP zero varies from 1.4 to 10. Figure 7 indicates that the none of the 50 plants have an upward margin in excess of 10, which corresponds to 5 times the true value of the NMP zero. Note that finite margins can be observed for other values of n and d shown in Figure 6.

IV.C. Sensitivity with two nonminimum-phase zeros

In this section, we explore the sensitivity of RCAC to the location of NMP zero estimates for plants with two NMP zeros. First, 50 stable plants with NMP 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 RCAC, while the NMP zero estimates are varied from 1 through 4. Figure 8 shows the percent of the 50 plants that have unstable closed-loop responses as the estimates of the NMP zeros vary. Figure 8 shows that the slope on the downward margin is steeper than the slope on the upward margin, suggesting that the performance of RCAC with two NMP zeros is similar to the single NMP zero case. Figure 9 demonstrates

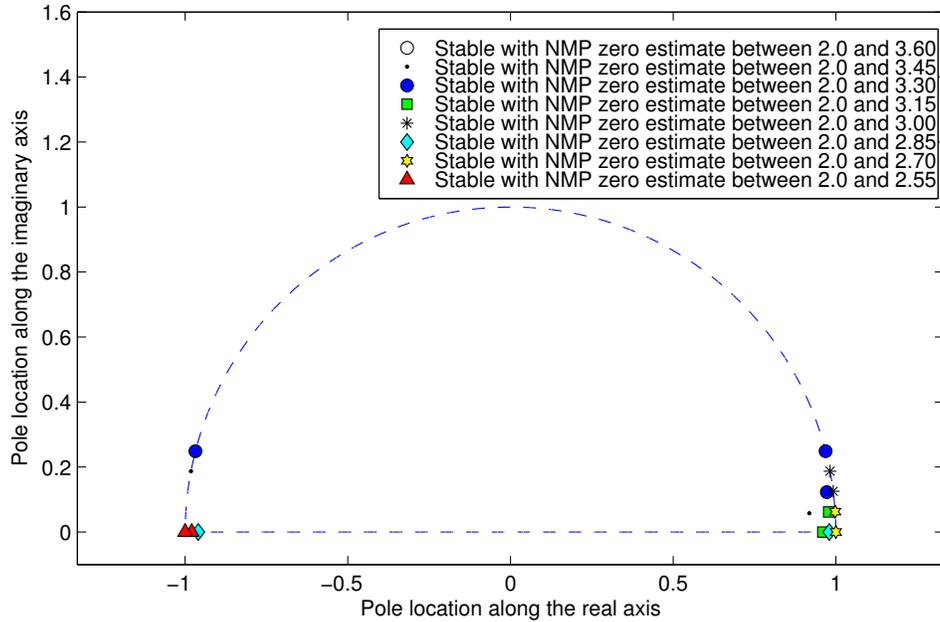


Figure 3. For 2601 second-order plants with a pair of stable complex conjugate poles, we let the estimate of the NMP zero (which is located at 2) vary from 2.0 to 3.6 and simulate the closed-loop system. The x -axis and y -axis correspond to the location of the plants' complex conjugate poles in the complex plane. Note that the plot is symmetric about the real axis; however, only the top half of the complex plane is shown. The colored regions represent the maximum value of the NMP zero estimate for which the closed-loop adaptive system is stable. The plants with poles that have real parts close to ± 1 are most sensitive to overestimating the NMP zero location. Note that the majority of the complex plane is colored in white, which indicates that these closed-loop systems are stable for NMP zero estimates up to (and potentially above) 3.6.

that RCAC is more robust to overestimating the location of the NMP zeros than it is to underestimating the location of the NMP zeros.

Next, we examine the sensitivity of RCAC to the location of complex NMP zero estimates. We randomly generate 100 stable plants with complex NMP zeros at $2 \pm j$, orders 3 through 10, and relative degrees 2 through $n - 2$. These plants are simulated in feedback with RCAC, while the NMP zero estimates are varied, in the complex plane, from 1 through $4 \pm 3j$. Figure 10 shows the percent of the 100 plants that have unstable closed-loop responses as the estimates of the NMP zeros vary. The color in Figure 10 indicates the percent of plants that unstable closed-loop responses.

Next, the 100 plants are simulated in feedback with RCAC, while the real part, imaginary part, phase, and magnitude of the NMP zero estimates are varied independently. For example, we simulate the closed-loop systems, where the estimate of the imaginary parts of the NMP zeros are ± 1 (i.e., the actual values), while the estimate of the real parts of the NMP zeros are varied from 1 to 4. Figure 11(a) shows the percent of the 100 plants that have unstable closed-loop responses. Figure 11(b) shows the percent of the 100 plants that have unstable closed-loop responses when the real parts of the NMP zeros are 2 (i.e., the actual value), while the estimate of the imaginary parts of the NMP zeros are varied from 0 to 3. Figures 11(c) and (d) show similar plots, where the phases of the NMP zeros estimates are set to the actual value while the magnitudes are varied, and where the magnitudes of the NMP zero estimates are set to the actual value while the phases are varied. Figures 11 (a)-(d) suggest that RCAC is less sensitive to underestimating the phase than to overestimating the phase, but RCAC is more sensitive to underestimating magnitude than it is to overestimating magnitude.

V. Sampled-Data Systems

In this section, we investigate the sensitivity of RCAC to sampling frequency for continuous-time systems that are discretized using a zero-order hold. More specifically, we randomly generate 100 stable, single-input, single-output, continuous-time plants with one NMP zero located at 2. For each of the 100 continuous-time

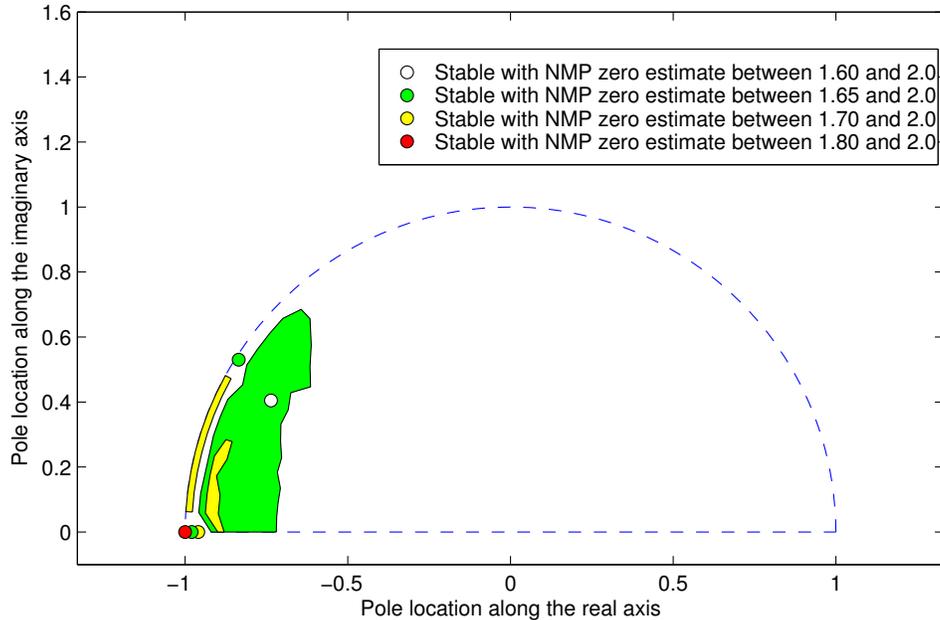


Figure 4. For 2601 second-order plants with a pair of stable complex conjugate poles, we let the estimate of the NMP zero (which is located at 2) vary from 2.0 to 1.6 and simulate the closed-loop system. The x -axis and y -axis correspond to the location of the plants' complex conjugate poles in the complex plane. Note that the plot is symmetric about the real axis; however, only the top half of the complex plane is shown. The colored regions represent the minimum value of the NMP zero estimate for which the closed-loop adaptive system is stable. The plants with poles that have negative real parts (more specifically, real parts closed to -1) are most sensitive to underestimating the NMP zero location. Note that the majority of the complex plane is colored in white, which indicates that these closed-loop systems are stable for NMP zero estimates down to (and potentially below) 1.6.

plants, the system order is generated from the uniform discrete integer distribution on the interval $[2, 6]$, the relative degree is generated from the uniform discrete integer distribution on the interval $[1, n]$, the first nonzero Markov parameter is 1, the plant poles are stable with magnitude less than 1, and the plant zeros (other than the NMP zero located at 2) are stable with magnitude less than 1. Note that discretizing these continuous-time plants with a zero-order hold may result in additional NMP zeros due to sampled-data effects.²⁰ In this case, estimates of these NMP sampling zeros are used by RCAC.

Next, for each plant, let F_{\min} be two times the maximum pole magnitude. For each continuous-time plant, we discretize the system using a zero order hold, where the sample frequency is varied from F_{\min} to $10F_{\min}$. Finally, for each plant and each sample frequency, we simulate the closed-loop system using RCAC, where the estimates of d , H_d , and the NMP zeros are set to the actual values. Figure 12(a) shows the logarithm of the transient performance (i.e., $\max |z(k)|$ over the entire simulation) for various sample frequencies, and Figure 12(b) shows the logarithm of the steady-state performance (i.e., $\max |z(k)|$ over the last 100 steps of the simulation) for various sample frequencies. Figure 12 suggests that RCAC is robust to the choice of sample frequency.

VI. Conclusion

In this paper, we reviewed the cumulative retrospective cost adaptive controller, which was presented in Ref. 16, for stabilization, command following, and disturbance rejection. The algorithm requires limited information about the open-loop system. Specifically, it requires knowledge of the first nonzero Markov parameter and the nonminimum-phase zeros from the control to the performance measurement. We computationally explored RCAC's behavior in the presence of inaccuracies in the estimates of the nonminimum-phase zeros from the control to the performance. In particular, the numerical results suggest that RCAC is, in general, more robust to overestimating the location of the nonminimum-phase zeros than it is to underestimating the location of the nonminimum-phase zeros. For complex nonminimum-phase zeros, the numerical results suggest that RCAC is less sensitive to underestimating the phase than to overestimating the phase, but

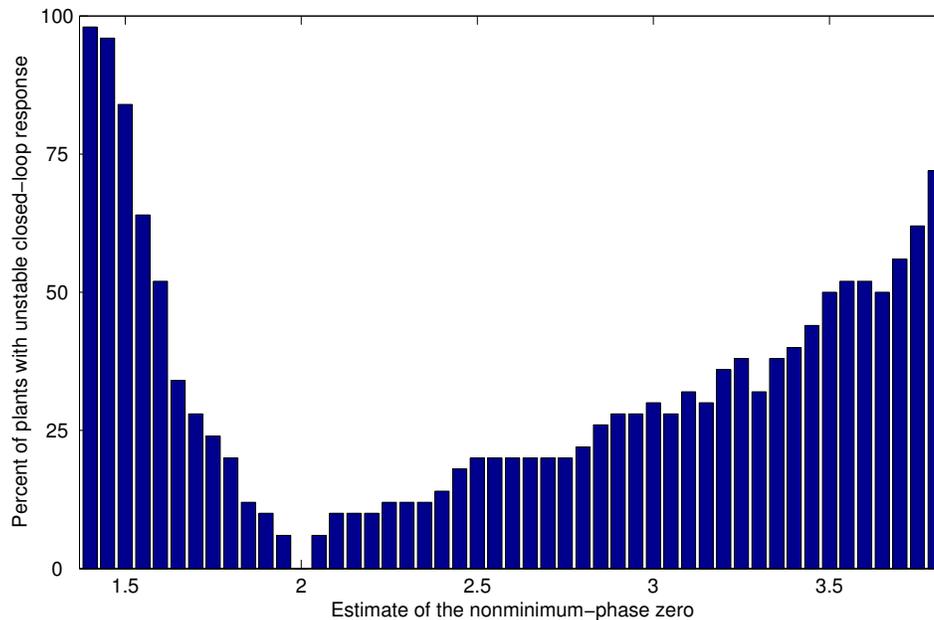


Figure 5. For 50 randomly generated plants, all of which have a NMP zero at 2, the adaptive controller (11)-(13) uses an estimate of the NMP zero that varies from 1.4 to 3.8. This histogram shows the percent of the 50 plants that have unstable closed-loop responses for different estimates of the NMP zero. For these 50 random plants, the stability of the closed-loop system is less sensitive to overestimating the location of the NMP zero than it is to underestimating the location of the NMP zero.

RCAC is more sensitive to underestimating magnitude than it is to overestimating magnitude. Furthermore, the numerical results demonstrate that RCAC is, in general, more robust to errors in the nonminimum-phase zero estimates for systems with lower order or higher relative degree. Finally, we investigated the performance of RCAC on sampled-data systems for a range of sample frequencies. These results demonstrate that RCAC maintains acceptable performance even at slow rates close to the Nyquist frequency of the fastest mode. Unlike continuous-time control techniques, which require post-design controller discretization, RCAC does not require fast sampling in order to ensure stability and performance.

References

- ¹Narendra, K. S. and Annaswamy, A. M., *Stable Adaptive Systems*, Prentice Hall, 1989.
- ²Åström, K. J. and Wittenmark, B., *Adaptive Control*, Addison-Wesley, 2nd ed., 1995.
- ³Ioannou, P. and Sun, J., *Robust Adaptive Control*, Prentice Hall, 1996.
- ⁴Goodwin, G. C. and Sin, K. S., *Adaptive Filtering, Prediction, and Control*, Prentice Hall, 1984.
- ⁵Tao, G., *Adaptive Control Design and Analysis*, Wiley, 2003.
- ⁶Fuentes, R. J. and Balas, M. J., "Direct adaptive disturbance accommodation," *Proc. IEEE Conf. Dec. Contr.*, Sydney, Australia, December 2000, pp. 4921-4925.
- ⁷Hoagg, 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.
- ⁸Johansson, R., "Global Lyapunov stability and exponential convergence of direct adaptive control," *Int. J. Contr.*, Vol. 50, 1989, pp. 859-869.
- ⁹Akhtar, 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.
- ¹⁰Akhtar, 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.
- ¹¹Goodwin, G. C., Ramadge, P. J., and Caines, P. E., "Discrete-time multivariable adaptive control," *IEEE Trans. Autom. Contr.*, Vol. 25, 1980, pp. 449-456.
- ¹²Hoagg, 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.
- ¹³Bayard, D. S., "Extended horizon liftings for stable inversion of nonminimum-phase systems," *IEEE Trans. Autom. Contr.*, Vol. 39, 1994, pp. 1333-1338.

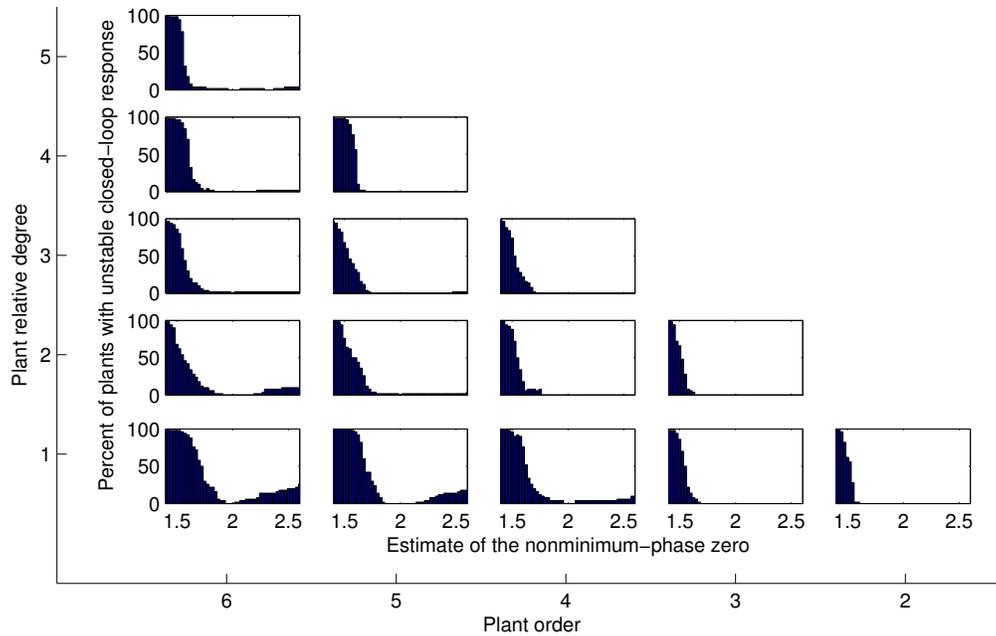


Figure 6. For 50 randomly generated plants for each order and relative degree specified per subplot, all of which have a NMP zero at 2, the adaptive controller (11)-(13) uses an estimate of the NMP zero that varies from 1.4 to 2.6. These histograms show the percent of the 50 plants that have unstable closed-loop responses for different estimates of the NMP zero. The robustness to the estimate of the NMP zero increases in rows from left to right (with decreasing order), and in columns from bottom to top (with increasing relative degree).

¹⁴Venugopal, R. and Bernstein, D. S., "Adaptive disturbance rejection using ARMARKOV/Toeplitz Models," *IEEE Trans. Contr. Sys. Tech.*, Vol. 8, 2000, pp. 257-269.

¹⁵Santillo, M. A. and Bernstein, D. S., "Adaptive Control Based on Retrospective Cost Optimization," *AIAA J. Guid. Contr. Dyn.*, Vol. 33, 2010, pp. 289-304.

¹⁶Hoagg, J. B. and Bernstein, D. S., "Cumulative retrospective cost adaptive control with RLS-based optimization," *Proc. Amer. Contr. Conf.*, June 2010.

¹⁷Hoagg, 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.

¹⁸Holzel, 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.

¹⁹Fledderjohn, 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.

²⁰Åström, K. J., Hagander, P., and Sternby, J., "Zeros of Sampled Systems," *Automatica*, Vol. 20, 1984, pp. 31-38.

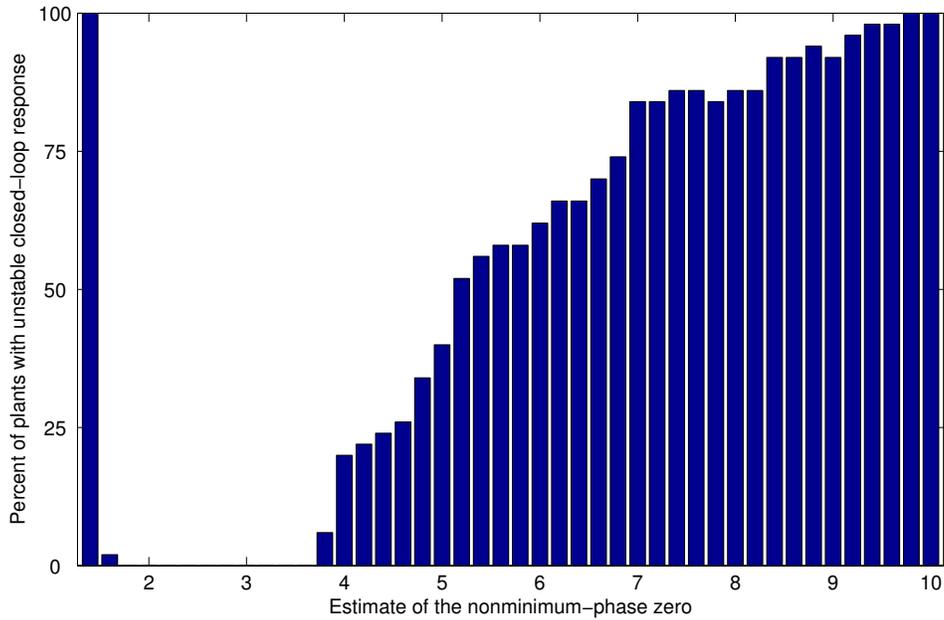


Figure 7. For 50 randomly generated plants with $n = 2$ and $d = 1$, all of which have a NMP zero at 2, the adaptive controller (11)-(13) uses an estimate of the NMP zero that varies from 1.4 to 10. This histogram shows the percent of the 50 plants that have unstable closed-loop responses for different estimates of the NMP zero. For these 50 random plants, the stability of the closed-loop system is less sensitive to overestimating the location of the NMP zero than it is to underestimating the location of the NMP zero. However, none of the 50 plants have an upward margin in excess of 10, which corresponds to 5 times the true value of the NMP zero.

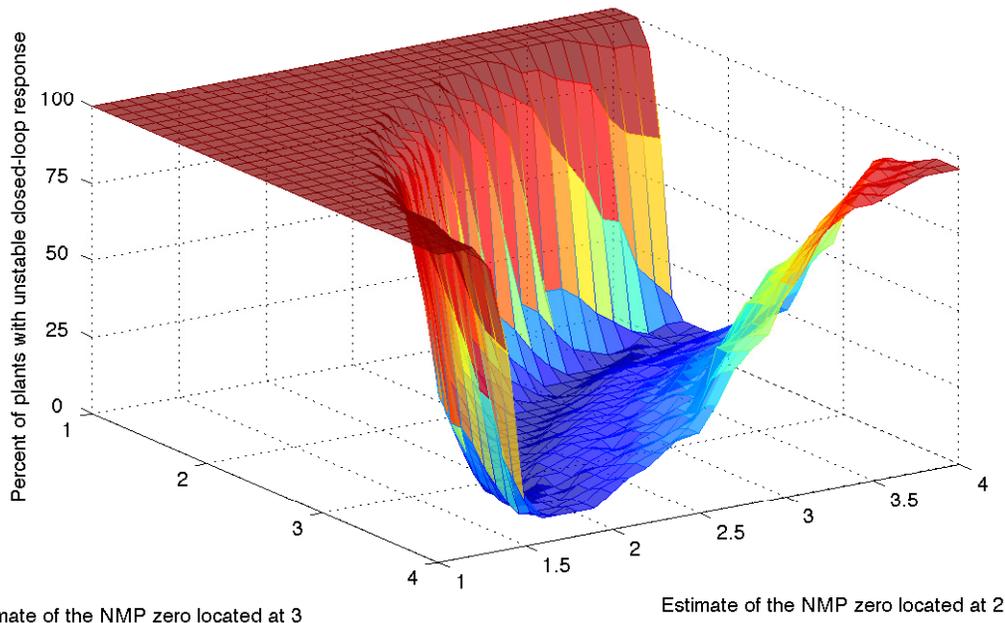


Figure 8. For 50 randomly generated plants, all of which have a NMP zeros at 2 and 3, the adaptive controller (11)-(13) uses estimates of the NMP zeros that vary from 1 to 4 for each NMP zero. The steep slope on the downward margin and shallower slope on the upward margin suggests that the performance of RCAC in the two-NMP-zeros case is similar to the single-NMP-zero case. That is, RCAC is more robust to overestimating the location of the NMP zeros, than it is to underestimating the location of the NMP zeros.

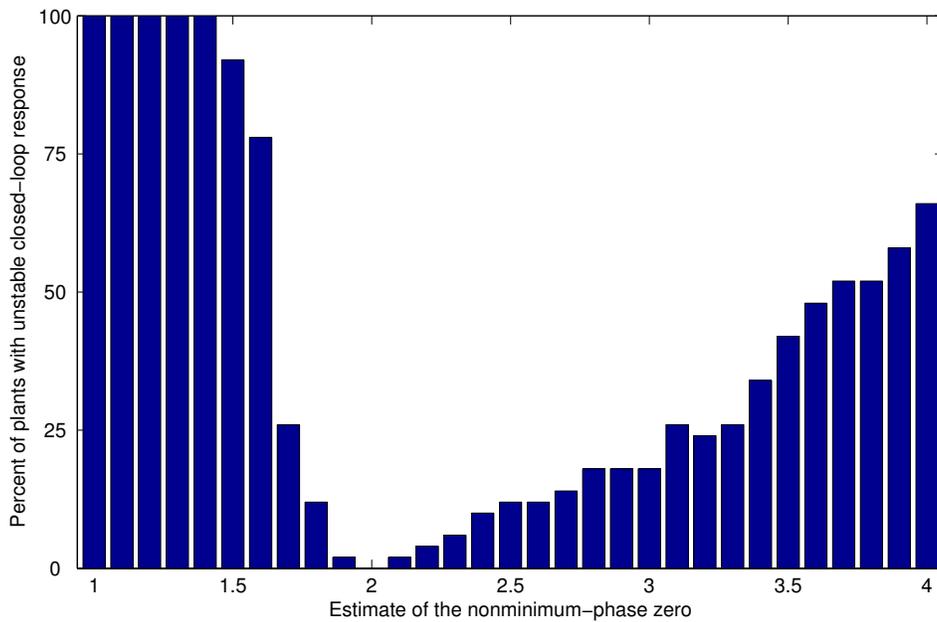


Figure 9. For 50 randomly generated plants, all of which have a NMP zeros at 2 and 3, the adaptive controller (11)-(13) uses the actual value of the NMP zero located at 3 and an estimate of the NMP zero located at 2 that varies from 1 to 4. Similar to the trend observed in the single-NMP-zero case, the stability of the closed-loop system is less sensitive to overestimating the location of the NMP zero than it is to underestimating the location of the NMP zero.

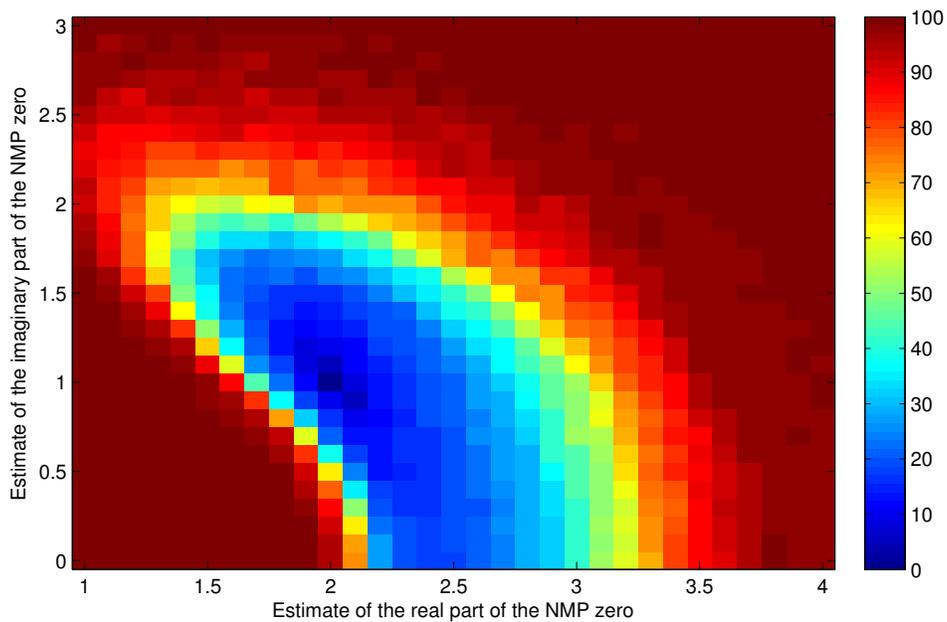


Figure 10. For 100 randomly generated plants, all of which have a NMP zeros at $2 \pm j$, the adaptive controller (11)-(13) uses estimates of the NMP zeros that vary from 1 to $4 \pm 3j$. The color indicates the percent of plants that unstable closed-loop responses.

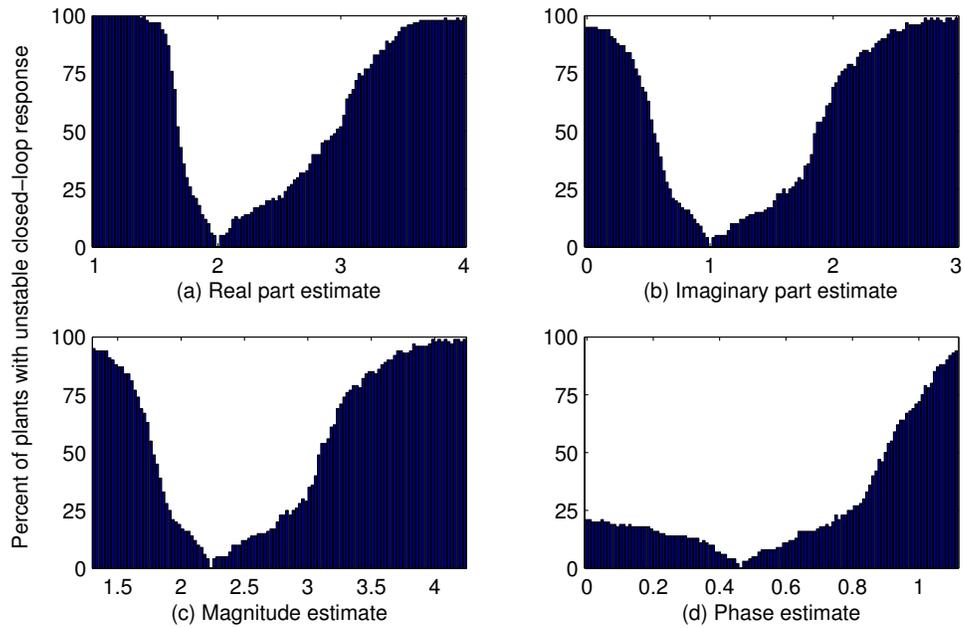


Figure 11. For 100 randomly generated plants, all of which have a NMP zeros at $2 \pm j$, the adaptive controller (11)-(13) uses estimates of the NMP zeros that (a) have the actual values for the imaginary parts while the real parts are varied, (b) have the actual values for the real parts while the imaginary parts are varied, (c) have the actual values for the phases while the magnitudes are varied, and (d) have the actual values for the magnitudes while the phases are varied. These results suggest that RCAC has similar sensitivity to errors in the estimates of the real and imaginary parts of the NMP zeros. However, the results suggest that RCAC is less sensitive to underestimating the phase than to overestimating the phase, but RCAC is more sensitive to underestimating magnitude than it is to overestimating magnitude.

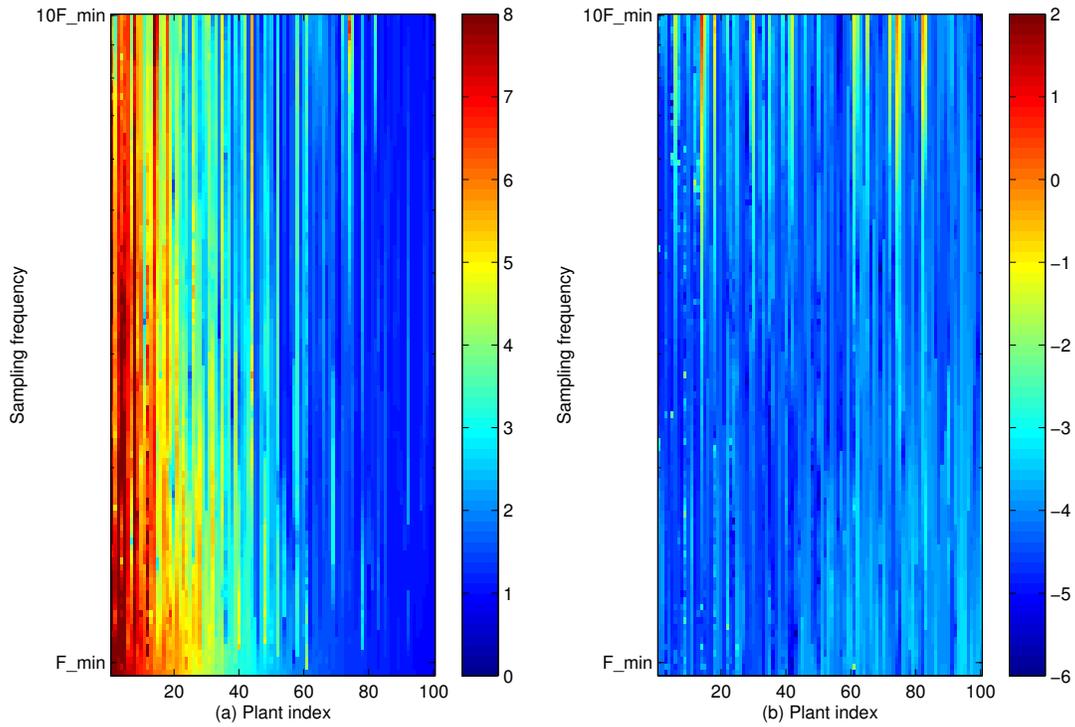


Figure 12. One hundred randomly generated, stable, continuous-time plants with one NMP zero located at 1 are discretized using sampling frequencies that vary from F_{\min} to $10F_{\min}$. The adaptive controller (11)-(13) is simulated in feedback using the actual value of the relative degree, first nonzero Markov parameter, and NMP zero estimate. The color scale in (a) and (b) corresponds to the logarithm of the transient performances and steady-state performances, respectively. These results suggest that RCAC is robust to the choice of sample frequency.