Performance Monitoring of Control Systems using Likelihood Methods

Matthew L. Tyler [*] Manfred Morari[†]

Chemical Engineering, 210-41, California Institute of Technology, Pasadena, CA 91125

Submitted to *Automatica*

March 8, 1995

Abstract

Evaluating deterioration in performance of control systems using closed loop operating data is addressed. A framework is proposed in which acceptable performance is expressed as constraints on the closed loop transfer function impulse response coefficients. Using likelihood methods, a hypothesis test is outlined to determine if control deterioration has occurred. The method is applied to a simulation example as well as data from an operational distillation column, and the results are compared to those obtained using minimum variance estimation approaches.

1 Introduction

Although enormous research effort has been directed towards design and analysis of controllers, relatively little work has addressed the problem of evaluating the performance of closed loop systems. Nevertheless, automatic monitoring of control loop performance is extremely important for practical control applications, wherein changes in equipment or operating conditions may result in deterioration of a controller which originally functioned well. In a typical chemical manufacturing facility, thousands of control loops are used to track set points and reject disturbances, and manual supervision of each loop is an unwieldy task. As a single control engineer may be responsible for over a thousand controllers, efficient tools are needed to automatically identify controllers which may need to be re-designed.

Recently, attention has been focused on the problem of assessing control performance from closed loop data [3, 4, 16, 7]. The methods developed in the cited references focus on estimating the theoretically limiting variance which can be achieved using feedback and feedforward control, and are only applicable to stable discrete time systems whose only zeros outside the unit disk are at

[*]Supported by Fannie and John Hertz Foundation. e-mail: mlt@aut.ethz.ch.

[†]Corresponding author. Current address: Automatic Control Laboratory, Swiss Federal Institute of Technology (ETH), 8092 Zurich, SWITZERLAND. Phone: 41-1-632-7626, Fax: 41-1-632-1211, e-mail: morari@aut.ethz.ch

infinity, i.e. delays. These methods have found some application to industrial systems [5, 20], and have also been applied to fault detection [6]. Tyler and Morari [19] have extended this approach to unstable systems, and systems with general non-invertible dynamics when the locations of the poles and zeros outside the unit disk are known.

Even though these "minimum variance" methods may provide useful information about achievable performance limits, they have several shortcomings as a tool for evaluating deteriorating control. For example, in many cases, minimum variance does not provide a meaningful measure of control performance. Achieving this theoretical limit may require a controller with high band width, or excessive control action which may result in the violation of robustness conditions. In addition, a shift in the ratio of actual variance to minimum variance may be due to either changes in the controller, changes in the plant, or changes in the disturbance spectrum. Whereas changes in the controller or the plant may merit retuning, changes in the disturbance spectrum may not. Finally, changes in performance may be due to changes in delay or in the location of the non-invertible zeros. Since the minimum variance methods require these parameters to be known, the estimates obtained may be quite poor when incorrect values are used.

In this paper, we take the approach that good performance can be expressed as constraints on the impulse response coefficients. In Section 2, we show how several practical performance criteria can easily be expressed in this paradigm. Once a meaningful performance criterion is established, a hypothesis testing problem is developed to determine if the performance is being achieved. Section 3 shows how this test can be evaluated using a generalized likelihood ratio (GLR) approach, and limitations of direct application of the GLR are enumerated. In Section 4, adjustments to the GLR test are introduced to address these limitations. Extensions to systems subject to command signals are discussed in Section 5. Finally, Section 6 presents examples which compare this method and the minimum variance methods.

2 Performance as impulse response coefficient constraints

In order to evaluate if a given performance criterion is achieved, performance specifications must first be expressed mathematically. For controller synthesis, the performance objectives commonly used are chosen so as to result in a solvable design problem. For example, a linear quadratic objective is often used because the resulting optimal control is easy to compute. Similarly, for robust control of systems with bounded uncertainty, a H_∞-norm performance objective leads to a solvable problem. However, many performance criteria such as settling time and overshoot are not easily translated into these objectives.

Consider systems of the form shown in Figure 1, with e a white noise process. Let ϕ denote the impulse response $[1, \phi_1, \phi_2, \ldots]$ of the closed loop transfer function from e to y. Note that ϕ depends not only on the plant P and the controller C, but also on the disturbance generator W. Also, because the magnitude of the input signal e is unknown, we can always assume $\phi_0 = 1$. Many useful performance criteria can be recast as constraints on the coefficients of the ϕ. We consider some examples here.

1. **Closed loop settling time**. In industry, this criterion is commonly used to assess controller performance, as is evident from the following quotation from Kozub and Garcia [8]:

   Shell control engineers typically aim for a closed loop response with a settling time

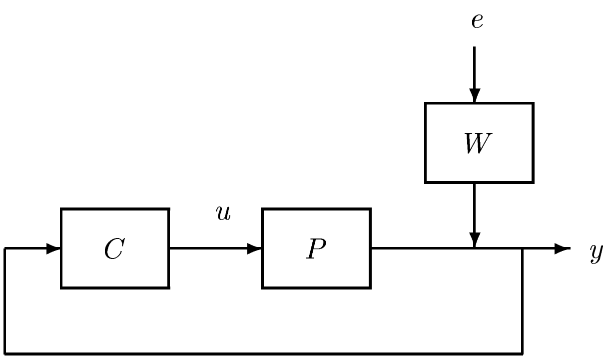

Figure 1: Closed loop system

close to the speed of response of the loop input transfer function when significant upsets occur.

Letting τ denote the response time, this performance objective could be stated as

$$|\phi_t| < \delta\phi_0 \text{ for } t > \tau, \tag{1}$$

where δ is a suitably small constant.

2. **Decay rate**. Alternatively, good performance may be specified by requiring that after t_0 sampling periods, ϕ decay no slower than exponentially with time constant T. This objective translates to the following bounds on ϕ:

$$|\phi_t| < \exp\left(-\frac{t - t_0}{T}\right)\phi_0 \text{ for } t > t_0 \tag{2}$$

3. **Minimum Variance**. For processes whose only non-invertible zeros lie at infinity, the minimum variance closed loop response is a moving average process of order d where d is the number of delays. Minimum variance control could then be expressed by the impulse response constraints:

$$\phi_t = 0 \text{ for } t > d \tag{3}$$

4. **Frequency domain bounds.** Often, it is desired to keep the frequency response of the closed loop ϕ small over a specified frequency range. This is expressed by

$$\left|\sum_{k=0}^{\infty} e^{-jk\omega}\phi_k\right| \leq b(\omega).$$

Although this constraint is non-linear in the coefficients of ϕ, it can be approximated by linear constraints by noting that $|z| < 1$ for z complex is equivalent to $\Re(z)\cos(\theta) + \Im(z)\sin(\theta) < 1$ for $0 \leq \theta < 2\pi$. For ϕ_k real, this transforms the above constraint to the set of linear constraints

$$\sum_{k=0}^{\infty} \phi_k \cos(k\omega + \theta) \leq b(\omega) \text{ for all } \theta \in [0, 2\pi].$$

Although in general this amounts to an infinite number of constraints at each frequency, approximating with a finite discretization of θ is usually acceptable.

5. **Filtered coefficient constraints.** Consider the case where rejection of step-like disturbances is important, but due to controller bandwidth limitations or modeling errors, high frequency oscillations must be tolerated. A meaningful performance criterion should not depend on high frequency phenomena. By constraining the low frequency components of the closed loop impulse response coefficients, this feature may be built into the performance criterion. In general, filtering introduces phase distortion. As a result, the filtered time response may look quite different from the unfiltered time response, even when the time response is band-limited by the filter cutoff frequency. When a filter with linear phase is used, the phase distortion results in a time shift of the time response coefficient for which one can easily compensate [14]. An FIR low pass filter with linear phase can be designed using the method of McClellan and Parks [12], which has been incorporated in MATLAB as the command `remez` [9]. Implementing such a linear filter on the impulse response coefficients will result in a delayed approximation of the response to low frequency disturbances. Thus, meaningful performance constraints will take the form
$$|\phi_k^F| < b_k \phi_D^F \text{ for } k > D,$$
where $\phi_k^F = \sum_{i=0}^{k} F_{k-i}\phi_i$ are the filtered impulse response coefficients, $[F_0, \ldots, F_{2D}]$ is the FIR low pass filter, and D is the delay associated with this filter, that is the slope of the phase.

Each of the above examples involve linear equality or inequality constraints on the closed loop impulse response coefficients of the form:

$$A\phi \leq b. \tag{4}$$

Although ϕ is an infinite dimensional vector, realistic performance criteria will restrict ϕ_t to be arbitrarily small for $t \geq t^*$, where t^* is not too large, implying that ϕ can be accurately approximated by a finite impulse response model of reasonable length when the system is performing satisfactorily. In this paper, we will consider performance criteria of the form (4) with ϕ a t^*-dimensional vector, along with the implicit assumption that $\phi_t = 0$ for $t > t^*$.

Just as a linear quadratic objective leads to straightforward optimal control design, specifying performance criteria as bounds on ϕ results in performance evaluation methods with tractable solutions. In addition, although objectives of the form (4) are simple, they can represent a wide variety of realistic performance measures. Finally, we note that in applying the methods of this paper, effort should be made so that the performance constraints are properly formulated to reflect the desired closed loop properties so that the tests are meaningful.

3 Generalized Likelihood Ratio Test

Performance evaluation of control systems can be viewed as choosing between two hypotheses: 1) The closed loop behavior is satisfactory, and 2) The closed loop behavior is unsatisfactory. For a disturbance and noise free system, the hypothesis test could be quite simple, for example noting the settling time for a set point step change. Unfortunately, most industrial control systems do not satisfy this idealization. Unmeasured disturbances which are stochastic and perhaps time-varying affect the output, and the performance criterion of interest will frequently address the capability of the control system to reject such unknown disturbances.

A common scheme in hypothesis testing methods involves evaluating the likelihood of the observations given each separate hypothesis. The hypothesis which results in the largest likelihood of the observed event is then chosen. This method has found wide use in detection problems [1] and is the basis of many common schemes, such as the cumulative sum, or CUSUM, algorithm.

To address the performance hypothesis test, we propose the following method. Let Θ be a parameterization of closed loop models. The noise signal e generates via Θ the effect of unmeasured disturbances on the output. We will assume that the noise is Gaussian, with zero mean and its covariance σ is part of the model parameterization. This allows the association of a probability density function with the model $\theta \in \Theta$. (In this paper, θ will denote a model, whereas ϕ denotes the impulse response coefficients of the model.) Let $\mathcal{Y}_L(t)$ be the vector containing the L most recent measurements, i.e. $\mathcal{Y}_L(t) = [y(t), y(t-1), \ldots, y(t-L+1)]$. Given a model θ, we denote the probability of $\mathcal{Y}_L(t)$ by $p_\theta(\mathcal{Y}_L(t))$.

Example Suppose $\mathcal{Y}_L(t)$ is generated by a moving average stochastic model,

$$y(t) = e(t) + \sum_{i=1}^{m} b_i e(t-i), \tag{5}$$

where e is Gaussian white noise with covariance σ^2. Then $\theta = [b_1, \ldots, b_m, \sigma^2]$. The probability density function p_θ is given by

$$p_\theta(\mathcal{Y}_L(t)) = \prod_{k=t-L-m}^{t} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\epsilon^2(k,\theta)}{2\sigma^2}\right), \tag{6}$$

with $\epsilon(k,\theta)$ obtained by filtering y as follows

$$\epsilon = \frac{1}{1 + \sum_{i=1}^{m} b_i z^{-i}} y. \tag{7}$$

In order to calculate ϵ, the transfer function $1 + \sum_{i=1}^{m} b_i z^{-i}$ should be invertible. In practice this does not pose a problem. As the signal e is unknown, the closed loop transfer function is essentially fit to the spectrum of the output, and it is well known that for any non-invertible transfer function, there exists an invertible transfer function with the same spectrum.

□

We now consider the problem of evaluating the performance of a closed loop system. Let Θ_0 be a parameterization of all transfer functions corresponding to acceptable performance levels. Because the closed loop system is driven by an unknown signal e, it is possible that any output sequence results from an acceptable model $\theta_0 \in \Theta_0$; however, if we don't restrict the model to correspond to good performance, a more probable e may be found which could have created the observed sequence. By calculating the two probabilities

$$P_0 = P(\mathcal{Y}_L(t) | \text{Good Performance is preserved}),$$

and

$$P_1 = P(\mathcal{Y}_L(t) | \text{Performance not necessarily preserved}),$$

a decision can be made by choosing the most probable hypothesis. To evaluate these probabilities, another parameterization Θ_1 is needed. This parameterization should be large enough that it contains models that correspond to satisfactory performance, as well as models corresponding to unsatisfactory behavior.

Consider now the idealized situation where good performance corresponds to exactly one closed loop model, θ_0, and poor performance corresponds to a different model θ_1. In this simplified case, the likelihood ratio of the two models can easily be calculated using the probability density functions p_{θ_0} and p_{θ_1}:

$$\frac{P_0}{P_1} = \frac{p_{\theta_0}(\mathcal{Y}_L(t))}{p_{\theta_1}(\mathcal{Y}_L(t))}. \tag{8}$$

In a more typical situation, the set of models with good performance Θ_0 contains many models, as does the parameterization Θ_1. In hypothesis testing terminology, this corresponds to a compound hypothesis testing problem. A common way to deal with compound hypothesis testing problems is to replace the single probability density function p_θ by the density function corresponding to the most likely hypothesis in the set, i.e. $\sup_{\theta \in \Theta} p_\theta$. Replacing the densities with with their maximum likelihood estimates results in *Generalized Likelihood Ratio*, or GLR criterion:

$$\frac{P_0}{P_1} = \frac{\sup_{\theta_0 \in \Theta_0} p_{\theta_0}(\mathcal{Y}_L(t))}{\sup_{\theta_1 \in \Theta_1} p_{\theta_1}(\mathcal{Y}_L(t))} \tag{9}$$

In the case of Gaussian white noise, the supremum is achieved by minimizing the sum of squares of the residuals, that is

$$\hat{\theta} = \arg\min_{\theta \in \Theta} V(\mathcal{Y}_L(t), \theta), \ V(\mathcal{Y}_L(t), \theta) = \frac{1}{L} \sum_{k=t-L+1}^{t} \epsilon^2(k, \theta), \tag{10}$$

with ϵ the model residuals producing $\mathcal{Y}_L(t)$. For simplicity of notation, we write $V(\mathcal{Y}, \theta)$ for $V(\mathcal{Y}_L(t), \theta)$. Setting $\sigma^2 = V(\mathcal{Y}, \hat{\theta})$ minimizes the probability. The supremum $p_{\hat{\theta}}(\mathcal{Y}_L(t))$ is then

$$\begin{aligned} p_{\hat{\theta}}(\mathcal{Y}_L(t)) &= \prod_{k=t-L+1}^{t} \frac{1}{\sqrt{2\pi V(\mathcal{Y}, \hat{\theta})}} \exp\left(-\frac{\epsilon^2(k, \hat{\theta})}{2V(\mathcal{Y}, \hat{\theta})}\right), \\ &= \left(\frac{1}{\sqrt{2\pi V(\mathcal{Y}, \hat{\theta})}}\right)^L \exp\left(-\frac{\sum_{k=t-L+1}^{t} \epsilon^2(k, \hat{\theta})}{2V(\mathcal{Y}, \hat{\theta})}\right) \\ &= \left(\frac{1}{\sqrt{2\pi V(\mathcal{Y}, \hat{\theta})}}\right)^L \exp(-L/2). \end{aligned} \tag{11}$$

The GLR then becomes

$$\frac{P_0}{P_1} = \left(\frac{V(\mathcal{Y}, \hat{\theta}_1)}{V(\mathcal{Y}, \hat{\theta}_0)}\right)^{\frac{L}{2}} \tag{12}$$

When $\Theta_0 \subset \Theta_1$, it follows that $V(\mathcal{Y}, \hat{\theta}_1) < V(\mathcal{Y}, \hat{\theta}_0)$, so that the GLR is always less than one. If L is large, then small differences between $V(\mathcal{Y}, \hat{\theta}_1)$ and $V(\mathcal{Y}, \hat{\theta}_0)$ will result in very small values of the GLR. Often, it will be convenient to choose Θ_0 as a subset of Θ_1. Since the GLR will always be less than one in this case, the test can not be applied directly as "Choose the hypothesis which is shown by the GLR to be more likely." In the following section, we discuss several ways to select thresholds for the GLR which also depend on the sample data.

4 Selection of Model Structures Θ_1 and Θ_0

With the performance criterion as in (4), the natural parameterization for Θ_0 is that of moving average models. Although autoregressive models have the advantage that the residuals ϵ are linear functions of the model parameters, thus reducing the estimation of θ_0 to a least squares problem, they have the disadvantage that constraints of the form (4) are quite cumbersome to express in terms of the free model parameters. By contrast, moving average models are more difficult to estimate, but the performance criteria are easily implemented. As efficient numerical methods for estimating moving average model parameters are well known [10, 11, 15], we do not view this difficulty to be overly burdensome. Our experience has shown that including constraints in the estimation of MA coefficients only slightly increases the computational burden. We note that using models with both AR and MA terms, which in most cases gives the most parsimonious fit to a generic time series, also makes implementing constraints of the form (4) difficult. For these reasons, we focus on restricting the structure Θ_0 to be a moving average model.

A more important issue lies in the relation between the two sets Θ_1 and Θ_0. For example, consider the simple case where good performance implies complete settling within t_0 samples. The set of "good" models can be expressed as $\Theta_0 = \{\phi | \phi_t = 0, t > t_0\}$, or in other words, $\theta \in \Theta_0$ is a MA(t_0) model. How should we choose Θ_1? If we choose Θ_1 as the set of MA(t_0) models, then the GLR would be identically 1 for any data set. If we choose Θ_1 as a higher order MA model, than the GLR could be very small even if the data were actually generated by a MA(t_0) model. Although this example employs a simplified performance constraint, an analogous situation will occur for general constraints. For the test to be meaningful, Θ_1 must contain models which are not contained in Θ_0, but if this is true and the model structures are nested, $\Theta_0 \subset \Theta_1$, the GLR will be less than one and frequently quite small.

One solution to this dilemma may be to avoid nested model structures. For many control applications with the controller well designed, the closed loop behavior is accurately modeled by a second order model added to an FIR model whose order equals the delay of the plant. Once Θ_1 is selected, Θ_0 can be fixed as an FIR model whose length is determined by considering the impulse response obtained from $\hat{\theta}_1$ fit to data corresponding to good performance. However, when the underlying unconstrained structure for Θ_0 and the structure of Θ_1 are much different, comparing the two may not be reasonable. Often, considering nested models may be more convenient, and in the following section we discuss three different modifications to the generalized likelihood criterion so that it is meaningful for nested structures.

4.1 Confidence Limit

Consider again the estimation of a MA(t_0) model using a MA(m) structure with $m > t_0$. Applying the methods of [11, 15] to estimating the parameters will result in an unbiased estimate. Therefore, in the limit of infinite data, the $m - t_0$ "extra" parameters in the MA(m) structure will be correctly identified as zero; however, for finite data, the variance of these estimates will be non-zero. Since the variance of the estimated parameters can also be estimated from the data, we can use this information for our criterion. In this section and the next, we discuss two possible approaches of incorporating parameter covariance estimates in the decision criterion.

Let $V(\mathcal{Y}, \hat{\theta}_1)$ be as in (10). The GLR method compares the values of the two quantities $V(\mathcal{Y}, \hat{\theta}_0)$

and $V(\mathcal{Y},\hat{\theta}_1)$. Let Σ_1 be the estimated covariance matrix for the parameters $\hat{\theta}_1$. We know that $V(\mathcal{Y},\hat{\theta}_1) \leq V(\mathcal{Y},\hat{\theta}_0)$ a priori for any nested model structure, so GLR $\gtrless 1$ is not a meaningful test. Instead, consider the following performance evaluation test:

> Is there a model $\theta_1 \in \Theta_1$ such that $P(\theta_1) > r^* P(\hat{\theta}_1)$ and $V(\mathcal{Y},\theta_1) = V(\mathcal{Y},\hat{\theta}_0)$, where P is the probability of the parameter estimates, and r^* is a specified constant on the interval $(0,1]$.

Using standard identification methods, the asymptotic estimate $\hat{\theta}_1$ obtained from minimizing V will be normally distributed with mean $\bar{\theta}_1$ and covariance Σ_1. Using this distribution, the relative likelihood of two models $\hat{\theta}_1$ and $\hat{\theta}_1 + d\theta$ can be calculated:

$$\frac{P(\hat{\theta}_1 + d\theta)}{P(\hat{\theta}_1)} = \frac{\exp(\frac{1}{2}(\hat{\theta}_1 - \bar{\theta}_1)^T \Sigma_1^{-1} (\hat{\theta}_1 - \bar{\theta}_1))}{\exp(\frac{1}{2}(\hat{\theta}_1 + d\theta - \bar{\theta}_1)^T \Sigma_1^{-1} (\hat{\theta}_1 + d\theta - \bar{\theta}_1))}.$$

By assuming $\bar{\theta}_1 = \hat{\theta}_1$, the proposed performance evaluation reduces to the test:

$$\frac{P(\hat{\theta}_1 + d\theta)}{P(\hat{\theta}_1)} = \exp(-\frac{1}{2} d\theta^T \Sigma_1^T d\theta) \gtrless r^*. \tag{13}$$

Choosing a value r^*, and letting $\delta^* = -2\log(r^*)$, this criterion can be expressed as

$$\max_{d\theta^T \Sigma_1^{-1} d\theta \leq \delta^*} V(\mathcal{Y}, \hat{\theta}_1 + d\theta) \gtrless V(\mathcal{Y},\hat{\theta}_0) \tag{14}$$

where the poor performance hypothesis is accepted if $<$ holds, and the good performance if $>$ holds. To apply this test, we could either solve the optimization problem (14), or we could consider an equivalent test,

$$\delta^* \gtrless \delta, \quad \delta = \min \hat{\delta} \text{ s.t. } \max_{d\theta^T \Sigma_1^{-1} d\theta \leq \hat{\delta}} V(\mathcal{Y}, \hat{\theta}_1 + d\theta) = V(\mathcal{Y},\hat{\theta}_0). \tag{15}$$

In general, both of these optimization problems may be very difficult to solve; however, a simple approximation can be made to the latter using that fact that $\hat{\theta}_1$ is an unconstrained optimal point of $V(\mathcal{Y},\theta)$. Consider the second order expansion of $V(\mathcal{Y},\theta)$ about $V(\mathcal{Y},\hat{\theta}_1)$:

$$V(\mathcal{Y},\hat{\theta}_1 + d\theta) = V(\mathcal{Y},\hat{\theta}_1) + V_\theta(\mathcal{Y},\hat{\theta}_1)^T d\theta + \frac{1}{2} d\theta^T V_{\theta\theta}(\mathcal{Y},\hat{\theta}_1) d\theta + O(||d\theta||^3), \tag{16}$$

where $V_\theta(\mathcal{Y},\hat{\theta}_1)$ and $V_{\theta\theta}(\mathcal{Y},\hat{\theta}_1)$ are the gradient and Hessian of $V(\mathcal{Y},\theta)$ respectively evaluated at $\hat{\theta}_1$. Since the matrix $V_{\theta\theta}(\mathcal{Y},\hat{\theta}_1)$ is usually calculated in the optimization routine to find $\hat{\theta}_1$, it is readily available, and also $V_\theta(\mathcal{Y},\hat{\theta}_1) = 0$ since $\hat{\theta}_1$ is a minimum of V. Using this expansion, we see that for $d\theta$ small, the maximization in (14) can be approximated by

$$\begin{aligned}
\max_{d\theta^T \Sigma_1^{-1} d\theta \leq \delta^*} V(\mathcal{Y},\hat{\theta}_1 + d\theta) - V(\mathcal{Y},\hat{\theta}_1) &= \max_{d\theta^T \Sigma_1^{-1} d\theta \leq \delta^*} \frac{1}{2} d\theta^T V_{\theta\theta}(\mathcal{Y},\hat{\theta}_1) d\theta + O(||d\theta||^3), \\
&= \max_{x^T x \leq \delta^*} \frac{1}{2} x^T J V_{\theta\theta}(\mathcal{Y},\hat{\theta}_1) J x + O(||d\theta||^3), \\
&= \delta^* \frac{1}{2} \bar{\lambda}(J V_{\theta\theta}(\mathcal{Y},\hat{\theta}_1) J) + O(||d\theta||^3),
\end{aligned} \tag{17}$$

where $J^2 = \Sigma_1$ and $\bar{\lambda}$ denotes the largest eigenvalue. The matrix J is guaranteed to exist since Σ_1 is symmetric and positive definite. An approximate of the minimal δ^* as in (15) can easily be obtained from this solution:

$$\begin{aligned} \delta &= 2\frac{V(\mathcal{Y},\hat{\theta}_0) - V(\mathcal{Y},\hat{\theta}_1)}{\bar{\lambda}(JV_{\theta\theta}(\mathcal{Y},\hat{\theta}_1)J)} + \frac{O(\|d\theta\|^3)}{\bar{\lambda}(JV_{\theta\theta}(\mathcal{Y},\hat{\theta}_1)J)}, \\ &\approx 2\frac{V(\mathcal{Y},\hat{\theta}_0) - V(\mathcal{Y},\hat{\theta}_1)}{\bar{\lambda}(JV_{\theta\theta}(\mathcal{Y},\hat{\theta}_1)J)}. \end{aligned} \tag{18}$$

If $\delta > \delta^*$, we draw the conclusion that the performance is poor. Rearranging (18) results in a threshold for the GLR test,

$$\left(\frac{V(\mathcal{Y},\hat{\theta}_1)}{V(\mathcal{Y},\hat{\theta}_0)}\right)^{\frac{L}{2}} \gtrless \left(1 - \frac{\delta^*\bar{\lambda}(JV_{\theta\theta}(\mathcal{Y},\hat{\theta}_1)J)}{2V(\mathcal{Y},\hat{\theta}_0)}\right)^{\frac{L}{2}} = T_1(r^*) \tag{19}$$

In summary, this hypothesis test consists of the following steps:

1. Calculate $\hat{\theta}_1 = \sup_{\theta\in\Theta_1} V(\mathcal{Y},\theta)$, as well as the associated covariance matrix Σ_1. Calculate J such that $J^2 = \Sigma_1$.
2. Calculate $\hat{\theta}_0 = \sup_{\theta\in\Theta_0} V(\mathcal{Y},\theta)$.
3. Choose a confidence limit r^* and the associated value $\delta^* = -2\log r^*$.
4. If the GLR exceeds the $T_1(r^*)$,conclude that the performance is satisfactory. Otherwise, conclude that the performance is poor.

4.2 Constraint softening approach

In the previous section, the covariance matrix for the unconstrained parameter estimates $\hat{\theta}_1$ was used to determine a GLR threshold for accepting one hypothesis over another. An alternative approach would be to use the covariance to modify the constraints used for Θ_0. Consider again the example of estimating the parameters for a MA(t_0) time series using a MA(m) model, $m > t_0$. If Θ_1 and Θ_0 are respectively parameterized by the set of all MA(m) and MA(t_0) models, $V(\hat{\theta}_1)$ will be smaller than $V(\hat{\theta}_0)$ due to the extra $m - t_0$ non-zero parameters; however, one would expect that the variance of these parameters might be large compared to the parameter values. In the event that the variance of the parameters exceeds the parameter values, by modifying Θ_0 to be the set of MA(m) models whose tailing $m - t_0$ coefficients are bounded in magnitude by the variance, $\hat{\theta}_1 \in \Theta_0$ implying $\hat{\theta}_0 = \hat{\theta}_1$ and $V(\hat{\theta}_1) = V(\hat{\theta}_0)$.

Let $\sigma_k = \sqrt{\Sigma_1(k,k)}$. Constraints of the form $\phi_k < b_k$ can be softened by using the covariance of the unconstrained estimates, resulting in constraints of the form:

$$\phi_k < b_k + \bar{\delta}^*\sigma_k \tag{20}$$

For more general linear constraints as in (4), we could replace ϕ_k in the constraint equations with $\phi_k + d\phi_k\sigma_k$ with $|d\phi_k| \leq \bar{\delta}^*$. This would transform (4) to the feasibility problem: Does there exist a vector $d\phi = [d\phi_1, \ldots, d\phi_m]$ such that

$$A\phi + Ad\phi \leq b,\ |d\phi_k| \leq \bar{\delta}^*\sigma_k. \tag{21}$$

This feasibility problem can easily be solved via the following linear program:

LP 1

$$\begin{array}{ll} \min_{d\phi} & \bar{\delta}_1 \\ \text{s.t.} & A(\phi + d\phi) \leq b \\ & |d\phi_k| \leq \bar{\delta}_1 \sigma_k \end{array}$$

In summary, this testing scheme consists of the steps:

1. Calculate $\hat{\theta}_1 = \sup_{\theta \in \Theta_1} V(\mathcal{Y}, \theta)$, as well as the associated covariance matrix Σ_1.
2. Choose a threshold $\bar{\delta}^*$.
3. Solve the linear programming problem LP 1 for $\bar{\delta}_1$.
4. If $\bar{\delta} > \bar{\delta}^*$, then conclude the performance is poor. Otherwise, the performance is satisfactory.

We make a brief comparison between this method and the one of the previous section. To implement the former, both an unconstrained and a constrained optimization of $V(\mathcal{Y}, \theta)$ must be solved, as well as an eigenvalue problem, whereas for the latter, only one unconstrained optimization of $V(\mathcal{Y}, \theta)$ is required, but a linear program is also needed. In most cases, solution of LP 1 may be easier than the constrained optimization of V.

Each test measures the distance in the parameter space from the unconstrained optimum of V to another point. The first method uses a weighted 2-norm, where the weight is the full covariance matrix, whereas the second method measures the distance using an ∞-norm, and only the diagonal elements of the covariance are used for weighting. When the covariance matrix is diagonal, these two measurements are equivalent in the following sense. Let $D_2\Theta(\delta)$ and $D_\infty\Theta(\delta)$ be defined in the following way:

$$\begin{aligned} D_2\Theta(\delta) &= \{d\theta \mid d\theta^T \Sigma_1^{-1} d\theta \leq \delta\}, \\ D_\infty\Theta(\delta) &= \{d\theta \mid |d\theta_k| \leq \delta\sigma_k\}, \} \end{aligned}$$

then for any δ_0, there exists δ_1 and δ_2 such that

$$D_2\Theta(\delta_1) \subset D_\infty\Theta(\delta_0) \subset D_2\Theta(\delta_2).$$

However, when Σ_1 is not diagonal, such an equivalence does not exist, and points which are "near" $\hat{\theta}_1$ with the 2-norm measure may not be near with the ∞-norm measure. This will especially be true in the case where Σ_1 is ill-conditioned. In this case, incorporating the directionality of Σ_1 into the linear program may result in a smaller value of $\bar{\delta}$. Since Σ_1 is positive definite and symmetric, it can be factored as $\Sigma_1 = V^T D^2 V$, with V a unitary matrix and D diagonal. Using this factorization, the set $D_2\Theta$ can be expressed as

$$D_2\Theta(\delta) = \{d\theta \| \|D^{-1}V d\theta\|_2^2 \leq \delta\}.$$

From this expression, the infinity norm measure which is equivalent to the norm used in defining $D_2\Theta$ can easily be gleaned, $\|D^{-1}V d\theta\|_\infty$. Minimizing this norm subject to the performance constraints results in the linear program LP 2:

LP 2

$$\begin{array}{ll} \min_{d\phi} & \bar{\delta}_2 \\ \text{s.t.} & A(\phi + d\phi) \leq b \\ & x = D^{-1} V d\phi \\ & |x_k| \leq \bar{\delta}_2 \end{array}$$

Besides the different norms used in the two approaches, another important distinction exists between the two methods. Whereas the latter measures the distance from the unconstrained optimum to the set of good performance models, the former measures the distance to the nearest point which lies on the same contour of $V(\mathcal{Y}, \theta)$ as $V(\mathcal{Y}, \hat{\theta}_0)$. Thus the former method gives a point $\hat{\theta}_1 + d\theta$ which has the same likelihood as $\hat{\theta}_0$, but may not satisfy the constraints, and the latter method gives a point $\hat{\theta}_1 + d\theta$ which satisfies the constraints, but may be less likely than $V(\mathcal{Y}, \hat{\theta}_0)$. This is illustrated in Figure 2. The point A represents the unconstrained minimum of $V(\theta)$, and B the constrained minimum. The set $\{\theta | \theta = \hat{\theta}_1 + d\theta, d\theta^T \Sigma^{-1} d\theta < \delta_1\}$ is contained within the heavy ellipse, and the set $\{\theta | \theta = \hat{\theta}_1 + d\theta, |d\theta_k / \sigma_k| < \delta_2\}$ is the interior of the heavy rectangle. The other curves represent the equicost curves of $V(\theta)$. The latter criterion measures the distance from A to C, the nearest feasible point, and the former criterion measures the distance from A to D, the nearest point with the same cost as the point B.

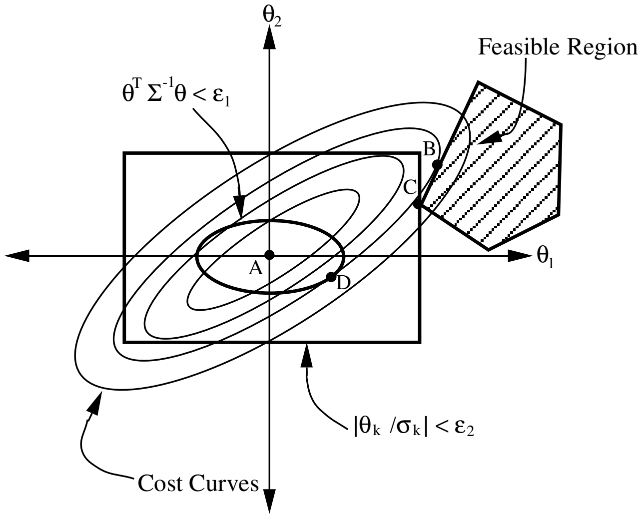

Figure 2: Distance Measures

4.3 Cross-Validation

Recall that the difficulty in using the GLR directly as a decision criterion for nested structures $\Theta_0 \in \Theta_1$ is that optimizing over the larger Θ_1 always gives a smaller value of $V(\mathcal{Y}, \theta)$ than optimizing over Θ_0. If the extra degrees of freedom of the set Θ_1 actually allow a better description of the system, then the model set Θ_0 is inferior to Θ_1. To determine whether added degrees of freedom in Θ_1 are justified by the data, one could use $\hat{\theta}_1$ to calculate residuals from a different data set which was not not used in the search for $\hat{\theta}_1$. This approach is commonly used for selecting model structures for process identification, and is referred to in the literature as *Cross-Validation*. As the proposed performance evaluation essentially consists of choosing between two model structures, Θ_0 and Θ_1, cross-validation tools can be used to test the hypothesis. In this section, we briefly review the general approach, and develop a convenient estimation technique which is valid for model structures containing constraints. We then show how the cross-validation approach can give a threshold for the GLR test.

The idea of using cross-validation for model selection dates to Stone [18]. The basic concept can be described as follows. Given a sequence of data and two separate model structures, first divide the data into two segments. Use one segment to obtain the best possible fit for the chosen structure, and then use the second segment to test the model. The structure which produces the smallest error on the test sequence is chosen.

A generalization of this approach has been presented by Stoica et al.[17]. Let $I = \{1, 2, \ldots L\}$ be an index corresponding to a data sample. Partition the data into k segments I_p such that

$$I_p = \{(p-1)m + 1, \ldots, pm\}. \tag{22}$$

We assume here that $L = km$. When this assumption does not hold, we simply let I_k contain more than m but less than $2m$ points. Also define the following model performance criteria:

$$V(\theta) = \frac{1}{L}\sum_{t\in I} \epsilon^2(t,\theta), \;\; V^p(\theta) = \frac{1}{L}\sum_{t\in I - I_p} \epsilon^2(t,\theta), \tag{23}$$

where for notational simplicity, we have dropped $\mathcal{Y}$ from the argument of V, although it is still implied that V depends on the data. Cross-validation involves first finding $\hat{\theta}_p \in \Theta$ which produces the smallest residual error for $t \in I - I_p$, and then using the residual errors for $t \in I_p$ which were not included in the optimization objective to evaluate the model structure Θ. When this procedure is repeated for each of the data partitions I_p, the the cross-validation criterion can be mathematically expressed by:

$$C_I(\Theta) = \sum_{p=1}^{k} \sum_{t\in I_p} \epsilon^2(t, \hat{\theta}_p), \;\; \hat{\theta}_p = \arg\min_{\theta\in\Theta} V^p(\theta) \tag{24}$$

This measure can be applied to performance evaluations in a straightforward fashion by comparing the values of $C_I(\Theta_0)$ and $C_I(\Theta_1)$. If $C_I(\Theta_0) < C_I(\Theta_1)$, then we can conclude satisfactory performance is maintained, otherwise, deteriorated performance would be concluded.

Although C_I could be calculated exactly, for large k this direct approach would be computationally expensive. Stoica et al. derive an asymptotic approximation for C_I which is much easier to compute for the case where the model structure does not contain constraints on the parameters. Under a relatively mild assumption, a similar approximation can be made for constrained model structures. Let us consider the following assumption:

Assumption: Let Θ be a set whose elements satisfy linear constraints of the form (4), and let $\hat{\theta}_p$ be as in (24). For each $p = 1, \ldots, k$, the active constraints of $\hat{\theta}_p$ are the same as the active constraints of $\hat{\theta}$, where $\hat{\theta}$ minimizes $V(\theta)$.

Then the following theorem holds:

Theorem Let $\hat{\theta} = \arg\min_{\theta\in\Theta} V(\theta)$, and let J index the set of active constraints for $\hat{\theta}$. Then

$$\frac{1}{L} C_I(\Theta) = q_m(\Theta) + O\left(\frac{1}{k^2 m}\right) \tag{25}$$

where

$$\begin{aligned}
q_m(\Theta) &= V(\hat{\theta}) + \frac{4}{L^2}\sum_{p=1}^{k} w_p(\hat{\theta})^T W_J(\hat{\theta}) w_p(\hat{\theta}) \\
w_p(\hat{\theta}) &= \sum_{t\in I_p} \epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta}) \\
W_J(\hat{\theta}) &= \left(V_{\theta\theta}(\hat{\theta})^T \mathcal{N}_{A_J} V_{\theta\theta}(\hat{\theta}) + A_J^T A_J\right)^{-1} V_{\theta\theta}(\hat{\theta})^T \mathcal{N}_{A_J} \\
\mathcal{N}_{A_J} &= (I - A_J^T (A_J A_J^T)^{-1} A_J)
\end{aligned}$$

and A_J is a matrix containing the rows of A indexed by J. When no constraints are active, $W_J = V_{\theta\theta}(\hat{\theta})^{-1}$.

Proof: (follows similar proof in [17]) First, consider the Kuhn-Tucker conditions for an optimum in $V(\theta)$ subject to linear constraints as in (4):

$$V_\theta(\hat{\theta}) + A_J^T \mu_J = 0. \tag{26}$$

Equivalently, this condition could be written as $\mathcal{N}_{A_J} V_\theta(\hat{\theta}) = 0$, or $V_\theta(\hat{\theta})^T \mathcal{N}_{A_J} = 0$.

For k sufficiently large, $\hat{\theta}_p$ is close to $\hat{\theta}$, and we write

$$\epsilon^2(t,\hat{\theta}_p) = \epsilon^2(t,\hat{\theta}) + 2\epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta})^T(\hat{\theta}_p - \hat{\theta}) + O(||\hat{\theta}_p - \hat{\theta}||^2) \tag{27}$$

Using the assumption that the active constraints for $\hat{\theta}_p$ are the same as for $\hat{\theta}$,

$$\begin{aligned}
0 &= \mathcal{N}_{A_J} V_\theta^p(\hat{\theta}_p) \\
&= \mathcal{N}_{A_J}\left(V_\theta^p(\hat{\theta}) + V_{\theta\theta}^p(\hat{\theta})(\hat{\theta}_p - \hat{\theta}) + O(||\hat{\theta}_p - \hat{\theta}||^2)\right) \\
&= \mathcal{N}_{A_J}\left(V_\theta(\hat{\theta}) + -\frac{2}{L}\sum_{t\in I_p}\epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta})\right. \\
&\quad \left. + \left\{V_{\theta\theta}(\hat{\theta}) - \frac{m}{L}\left(\frac{\partial}{\partial\theta}\frac{2}{m}\sum_{t\in I_p}\epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta})\right)\Bigg|_{\theta=\hat{\theta}}\right\}(\hat{\theta}_p - \hat{\theta}) + O(||\hat{\theta}_p - \hat{\theta}||^2)\right) \\
&= \mathcal{N}_{A_J}\left(-\frac{2}{L}\sum_{t\in I_p}\epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta}) + \left[V_{\theta\theta}(\hat{\theta}) + O\left(\frac{1}{k}\right)\right](\hat{\theta}_p - \hat{\theta}) + O(||\hat{\theta}_p - \hat{\theta}||^2)\right)
\end{aligned} \tag{28}$$

To evaluate the order of $\mathcal{N}_{A_J}\left(\frac{1}{m}\sum_{t\in I_p}\epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta})\right)$, note the following:

$$\begin{aligned}
\mathcal{N}_{A_J}\left(\frac{1}{m}\sum_{t\in I_p}\epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta})\right) &= \mathcal{N}_{A_J}\left(E(\epsilon(\hat{\theta})\epsilon_\theta(\hat{\theta})) + O(1/m)\right) \\
&= \mathcal{N}_{A_J}\left(\frac{1}{L}\sum_{t\in I}\epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta}) + O(1/\sqrt{L}) + O(1/\sqrt{m})\right) \\
&= \mathcal{N}_{A_J} V_\theta(\hat{\theta}) + O(1/\sqrt{m}) \\
&= O(1/\sqrt{m})
\end{aligned} \tag{29}$$

It then follows that $||\hat{\theta}_p - \hat{\theta}|| = O(1/(k\sqrt{m}))$. Therefore, we get the relation

$$\mathcal{N}_{A_J} V_{\theta\theta}(\hat{\theta})(\hat{\theta}_p - \hat{\theta}) = \mathcal{N}_{A_J} \left(\frac{2}{L} \sum_{t \in I_p} \epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta}) + O(\frac{1}{k^2\sqrt{m}}) \right). \tag{30}$$

Combined with the constraints $A_J(\hat{\theta}_p - \hat{\theta}) = 0$, we get the following system of equations for $(\hat{\theta}_p - \hat{\theta})$:

$$\begin{bmatrix} \mathcal{N}_{A_J} V_{\theta\theta}(\hat{\theta}) \\ A_J \end{bmatrix} (\hat{\theta}_p - \hat{\theta}) = \begin{bmatrix} \mathcal{N}_{A_J} \left(\frac{2}{L} w_p(\hat{\theta}) + O(\frac{1}{k^2\sqrt{m}}) \right) \\ 0 \end{bmatrix} \tag{31}$$

which has the unique solution

$$\hat{\theta}_p - \hat{\theta} = \frac{2}{L} W_J(\hat{\theta}) w_p + O\left(\frac{1}{k^2\sqrt{m}}\right), \tag{32}$$

Next, noting that $\mathcal{N}_{A_J}(\hat{\theta}_p - \hat{\theta}) = \hat{\theta}_p - \hat{\theta}$, and substituting (32) into (27) gives

$$\begin{aligned}
\frac{1}{L} C_I(\Theta) &= \frac{1}{L} \sum_{p=1}^{k} \sum_{t \in I_p} \left\{ \epsilon^2(t,\hat{\theta}) + 2\epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta})^T \mathcal{N}_{A_J} \left(W_J(\hat{\theta}) \frac{2}{L} w_p(\hat{\theta}) + O\left(\frac{1}{k^2\sqrt{m}}\right)\right)\right\} + O\left(\frac{1}{k^2 m}\right) \\
&= V(\hat{\theta}) + \frac{4}{L^2} \sum_{p=1}^{k} w_p(\hat{\theta})^T W_J(\hat{\theta}) w_p(\hat{\theta}) + \frac{m}{L} \sum_{p=1}^{k} \left(\mathcal{N}_{A_J} \frac{2}{m} \sum_{t \in I_p} \epsilon(t,\hat{\theta})\epsilon_\theta(t,\hat{\theta}) \right)^T \dot{O}\left(\frac{1}{k^2\sqrt{m}}\right) \\
&\quad + O\left(\frac{1}{k^2 m}\right) \\
&= V(\hat{\theta}) + \frac{4}{L^2} \sum_{p=1}^{k} w_p(\hat{\theta})^T W_J(\hat{\theta}) w_p(\hat{\theta}) + O\left(\frac{1}{k^2 m}\right).
\end{aligned} \tag{33}$$

□

By applying the approximate cross-correlation criterion $q_m(\Theta)$ to the model structures Θ_1 and Θ_0, we can determine if the set Θ_1 fits the data $\mathcal{Y}_L(t)$ better than Θ_0. If $q_m(\Theta_1) < q_m(\Theta_0)$, then the performance bounds imposed by Θ_0 are too restrictive. If the inequality is switched, acceptable performance is concluded. This approach can also be interpreted as the selection of a threshold for testing the GLR. The relation $q_m(\Theta_1) < q_m(\Theta_0)$ is equivalent to

$$GLR = \left(\frac{V(\mathcal{Y},\hat{\theta}_1)}{V(\mathcal{Y},\hat{\theta}_0)} \right)^{\frac{L}{2}} < \left(1 - \frac{\sum_{p=1}^{k} \left(w_p(\hat{\theta}_0)^T W_J(\hat{\theta}_0) w_p(\hat{\theta}_0) - w_p(\hat{\theta}_1)^T V_{\theta\theta}(\hat{\theta}_1)^{-1} w_p(\hat{\theta}_1) \right)}{V(\mathcal{Y},\hat{\theta}_0)} \right)^{\frac{L}{2}} = T_2. \tag{34}$$

In summary, a hypothesis test incorporating a cross-validation consists of the following steps:

1. Calculate $\hat{\theta}_1 = \sup_{\theta \in \Theta_1} V(\mathcal{Y}, \theta)$, and $\hat{\theta}_0 = \sup_{\theta \in \Theta_0} V(\mathcal{Y}, \theta)$.
2. Determine the active constraints for $\hat{\theta}_0$, and calculate threshold T_2 as in (34).
3. If the GLR is less than the threshold T_2 conclude poor performance. Otherwise, conclude good performance.

5 Extensions to Systems with Command Signals

The methods developed in the previous sections focus on evaluating the ability of a control system to reject unknown, unmeasured disturbances. We implicitly assumed that no changes in set point occur within the data sequence being tested; however, operating data will often contain set point changes. We could apply the methods of this paper to the signal $y - r$, where r is the reference signal, but this approach has disadvantages. For the GLR method, one essentially seeks to find the most likely unknown sequence ϵ that produces y via some model within the parameterizations Θ_1 and Θ_0. When the set point changes are large compared to the disturbances, a large value of ϵ will be needed to produces $y - r$ when no information on r is included, and this will typically bias the estimation of the underlying closed loop model.

In order to include reference signal data in the performance evaluation scheme, we consider two separate cases. The first corresponds to the case wherein the reference signal excites the system sufficiently for identification purposes, whereas in the second case, it does not. In most circumstances where routine operating data are used, the reference signal will excite the system poorly, and the second approach will be more applicable.

Consider a system as in Figure 1, with the modification that instead of y, the controller C is driven by $y - r$. The closed loop transfer function becomes:

$$y - r = \frac{W}{1+PC}e + \frac{1}{1+PC}r. \tag{35}$$

Previously, we have considered evaluating performance by estimating the transfer function $\frac{W}{1+PC}$ using two model structures, one of which has been constrained to correspond to acceptable performance. When the reference signal r is sufficiently exciting, a similar approach can be taken by specifying a set Θ_0^r which contains a parameterization of transfer functions $\frac{1}{1+PC}$ which corresponds to acceptable performance levels. Similarly, a parameterization Θ_1^r, analogous to Θ_1, is specified, and a generalized likelihood approach test may be developed using the ratio:

$$\frac{P(\mathcal{Y}_L(t)|\mathcal{R}_L(t), \theta \in \Theta_0, \theta^r \in \Theta_0^r)}{P(\mathcal{Y}_L(t)|\mathcal{R}_L(t), \theta \in \Theta_1, \theta^r \in \Theta_1^r)} \tag{36}$$

where $\mathcal{R}_L(t)$ is a vector analogous to $\mathcal{Y}_L(t)$ containing the history of the reference signal.

When the reference signal r is not sufficiently exciting, this approach will work poorly because the unconstrained structure Θ_1^r will tend to over fit the data, resulting in a small GLR. This may result in a poor performance conclusion, when the system is in fact performing satisfactorily. Consider, for example, the special case where the data $\mathcal{R}_L(t)$ contain only one step change in the reference which takes place at time $t_0 \in [t - L + 1, t]$. If Θ_1^r is an FIR(m) model, then the data can be made consistent with $\epsilon(t^*) = 0, t_0 \le t^* \le t_0 + m$, whereas the constrained fit will not be able to make ϵ zero on this range. Since in most practical cases, the signal r will not be sufficiently exciting, another approach is needed.

If the plant and controller were known exactly, one could easily calculate the signal $\frac{1}{1+PC}r$, and then use the previous likelihood approach on the signal z defined by:

$$z = y - r - \frac{1}{1+PC}r = \frac{W}{1+PC}e. \tag{37}$$

Since we are interested in monitoring changes in the performance of the system which could have possibly been caused by changes in P or C, it is unrealistic to assume models are known, in which case z will not be independent of r. Nevertheless, when the performance has not deteriorated, the transfer function $\frac{1}{1+PC}$ will be approximately as designed, and dependence on z of r will be less significant than if the term $\frac{1}{1+PC}r$ were not included. Therefore, when the performance is satisfactory, using z as the input to the monitoring scheme should result in similar diagnosis as in the case where r is constant. On the other hand, when performance has degraded and the transfer function $\frac{1}{1+PC}$ is no longer accurate, both the estimation of $\hat{\theta} \in \Theta_1$ and $\hat{\theta}_0 \in \Theta_0$ are affected by the changes in r, and the constraints of Θ_0 will generally result in larger residuals ϵ than will be needed for the optimal model in Θ_1. Therefore, changes in r will tend to make the GLR smaller, and thus favor a poor performance diagnosis. In the case where the reference signal is not sufficiently exciting, applying the methods of Sections 3 and 4 to the signal z in (37) is recommended.

6 Examples

In this section, we consider examples which serve to demonstrate the advantages of the methods of this paper over minimum-variance estimation methods. Example 1 addresses disturbance rejection for a simple system in which the process parameters as well as the disturbance characteristics may vary. In Example 2, the performance of a distillation column controller is evaluated.

Example 1 Consider a simple model of a stable system, described by the transfer function

$$P = z^{-2}\frac{5}{4}\frac{1-0.6z^{-1}}{1-0.5z^{-1}}.$$

Using the IMC design method outlined in [13], a controller is designed to reject step disturbances. The controller has the form:

$$C = \frac{Q}{1-PQ}, \; Q = \frac{(1-f)}{1-fz^{-1}}\frac{4}{5}\frac{1-0.6z^{-1}}{1-0.5z^{-1}}, \; 0 < f < 1.$$

Increasing the parameter f detunes the controller while increasing robustness to model uncertainty. In this example, $f = 0.3$, which corresponds to tight control which will be nearly minimum variance when there is no plant/model mismatch *and* the disturbances are accurately modeled as integrated white noise.

In order to use the likelihood method to determine system performance, meaningful constraints must be specified. The controller was designed to provide good rejection of step like disturbances. For the nominal plant, the closed loop step response to a unit step disturbance settles to 0.027 within 4 time samples, suggesting the performance specification:

$$|\phi_k| \leq 0.1 \text{ for } k \geq 4. \tag{38}$$

If the disturbance transfer function W were known a priori to be equal to $\frac{1}{1-z^{-1}}$, these constraints would be satisfactory. However, in practice W is not known, and may possibly have a large high frequency component. Assuming we are only interested in the rejection of low frequency disturbances, we should use the filtered coefficient constraint method of Section 2. A twentieth order (length 21) FIR filter was designed so that it had approximately unit amplitude in the range

Table 1: Nominal and Perturbed Plants

Plant	TF	OL gain	M_1	M_2	Diagnosis
P_1	$\frac{5}{4}z^{-2}\frac{1-0.6z^{-1}}{1-0.5z^{-1}}$	1	0.027	0.063	GOOD
P_2	$\frac{5}{7}z^{-2}\frac{1-0.3z^{-1}}{1-0.5z^{-1}}$	1	0.163	0.2197	POOR
P_3	$2z^{-2}\frac{1-0.75z^{-1}}{1-0.5z^{-1}}$	1	0.332	0.355	POOR
P_4	$\frac{7}{4}z^{-2}\frac{1-0.6z^{-1}}{1-0.3z^{-1}}$	1	0.226	0.235	POOR
P_5	$\frac{7}{5}z^{-2}\frac{1-0.75z^{-1}}{1-0.65z^{-1}}$	1	0.039	0.125	GOOD
P_6	$\frac{1}{2}z^{-2}\frac{1-0.6z^{-1}}{1-0.8z^{-1}}$	1	0.317	0.373	POOR
P_7	$\frac{7}{8}z^{-2}\frac{1-0.6z^{-1}}{1-0.5z^{-1}}$	0.7	0.216	0.212	POOR
P_8	$\frac{9}{8}z^{-2}\frac{1-0.6z^{-1}}{1-0.5z^{-1}}$	0.9	0.08	0.08	GOOD

$0 \leq \omega \leq 0.4\pi$ and zero amplitude in the range $0.6\pi \leq \omega \leq \pi$. For systems with bandwidth less than 0.4π, filtering the impulse response coefficients with this filter will result in a sequence approximately equal to the impulse response delayed by 10 samples. Because the impulse response coefficients correspond to the product of W and $\frac{1}{1+PC}$, including the filter implies that the performance criterion depends only on the response of the system to a disturbance of the form FWe rather than We. Combining the settling time criterion (38), a model parameterization for Θ_0 is given by $\theta_0 = [1, \phi_1, \ldots, \phi_{20}]$, and the low pass filter gives the following performance constraints:

$$|\phi_k^F| < 0.1\ \phi_{10}^F \text{ for } k \geq 14,$$

where $\phi_k^F = F_k + \sum_{i=1}^{k} F_{k-i}\theta_i$.

Eight different plants were used in the simulations, as shown in Table 1. The table also shows the open loop steady state gain, the two measures

$$M_1 = \max_{k>=4} |h_k|, \quad M_2 = \max_{k>=4} \frac{|h_{k+10}^F|}{h_{10}^F}, \tag{39}$$

where h and h^F are respectively the unfiltered and filtered closed loop step response coefficients, and the proper diagnosis based on $M_1 < 0.1$. The plant P_1 corresponds to the nominal system. The open loop and closed loop step responses for the plants are shown in Figures 3 through 6.

For each of the plants P_1 through P_8, the closed loop system was simulated using two separate disturbances. The first disturbance was generated by $d_1 = \frac{1}{1-z^{-1}}w$, with w a white noise process. The second disturbance d_2 was obtained by adding to d_1 the sequence $\frac{1-r\cos(\omega)z^{-1}}{1-2r\cos(\omega)z^{-1}+r^2z^{-2}}w$ with $\omega = \frac{2}{3}\pi$, $r = 0.95$, and w the same sequence which generated d_1. The disturbance d_2 has a significant high frequency component. Both disturbances are shown in Figure 7.

For each plant/disturbance pair, the GLR was computed, along with the thresholds T_1 (Eq. (19)) for $r^* = 0.5$ and T_2 (Eq. (34)) for $m = 1$. In addition, $\bar{\delta}_1$ and $\bar{\delta}_2$ were calculated from the linear programming problems LP 1 and LP 2. For the unconstrained estimation of $\hat{\theta}_1$, standard Newton-Rapson methods such as outlined in [15] were used, and the constrained optimization of $\hat{\theta}_0$ was calculated using a quadratic penalty function [2]. Tables 2 and 3 show the decision criteria, along

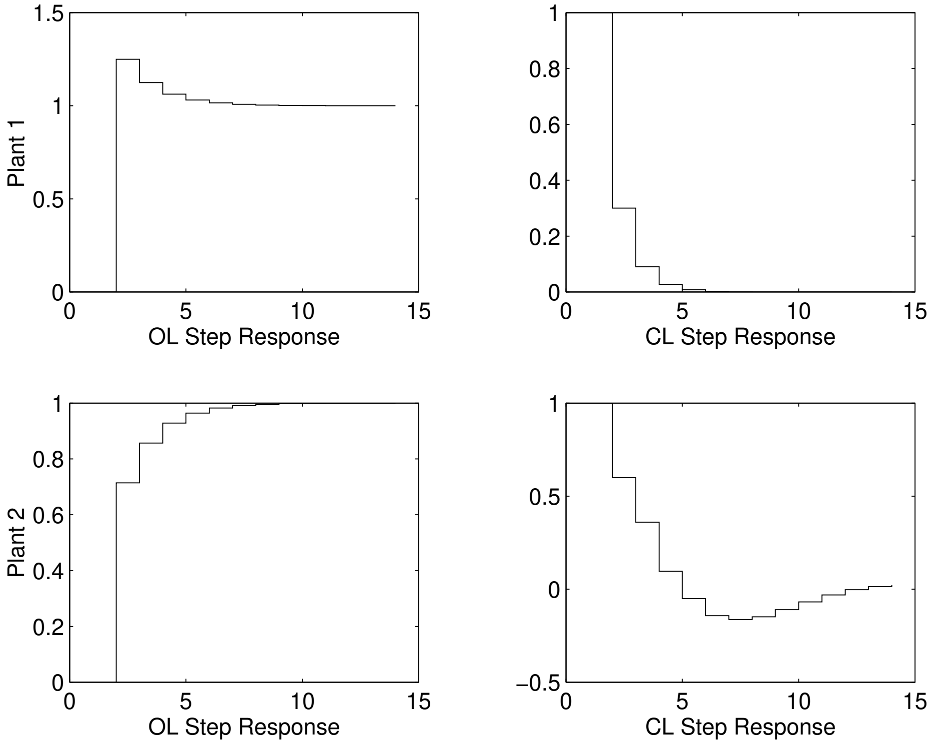

Figure 3: Plants 1 and 2

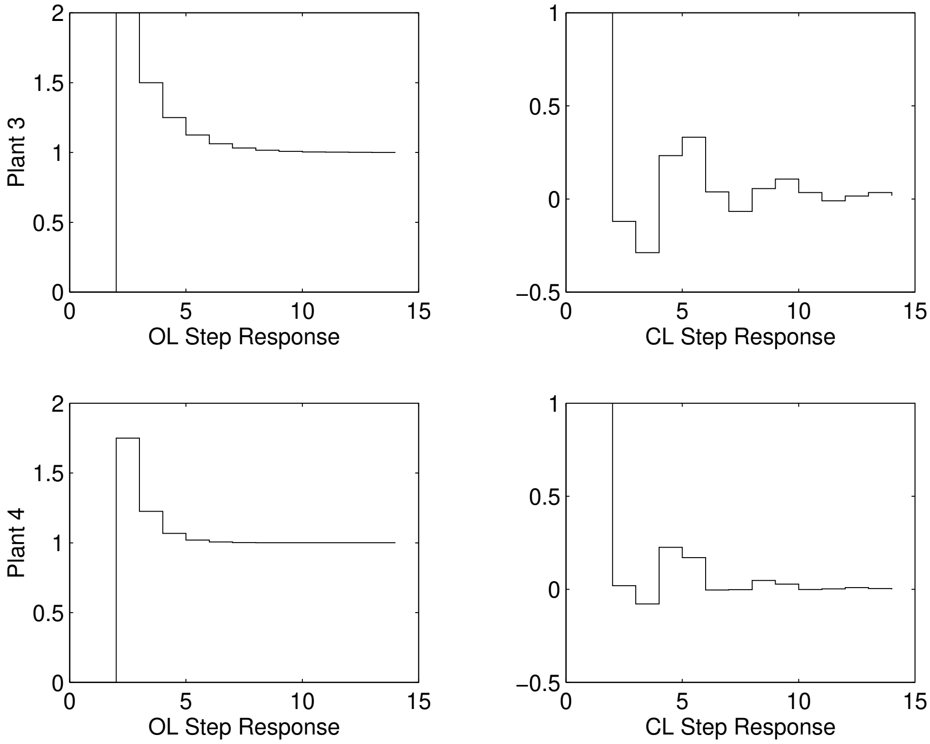

Figure 4: Plants 3 and 4

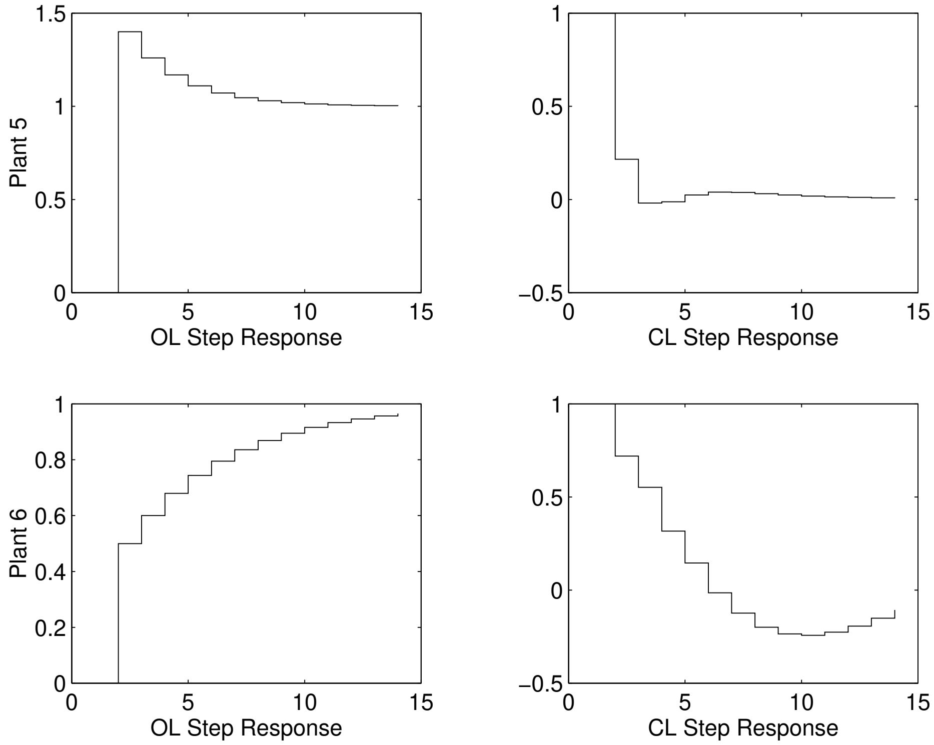

Figure 5: Plants 5 and 6

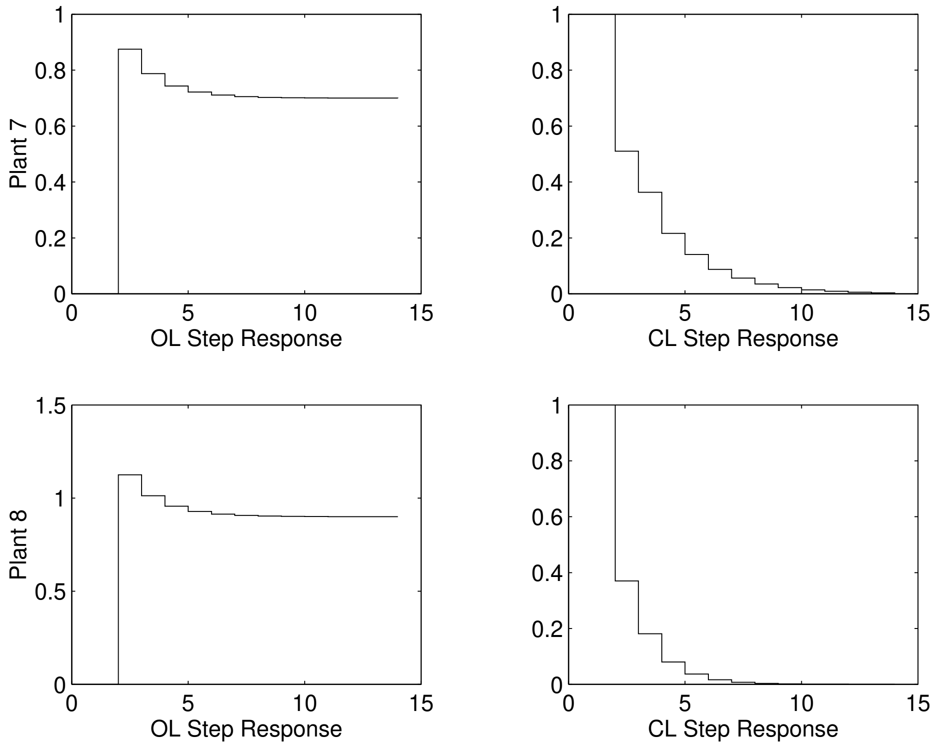

Figure 6: Plants 7 and 8

with the minimum variance (MV) ratio, $||\mathcal{Y}||/\sigma_{mv}$, where σ_{mv} was estimated using the method outlined in [7]. In addition, the Tables show the scaled condition number κ^* for the covariance matrix Σ_1, where κ^* is calculated by

$$\kappa^* = \min_{D_1, D_2} \kappa(D_1 \Sigma_1 D_2),$$

with D_1 and D_2 diagonal matrices, and κ the usual condition number.

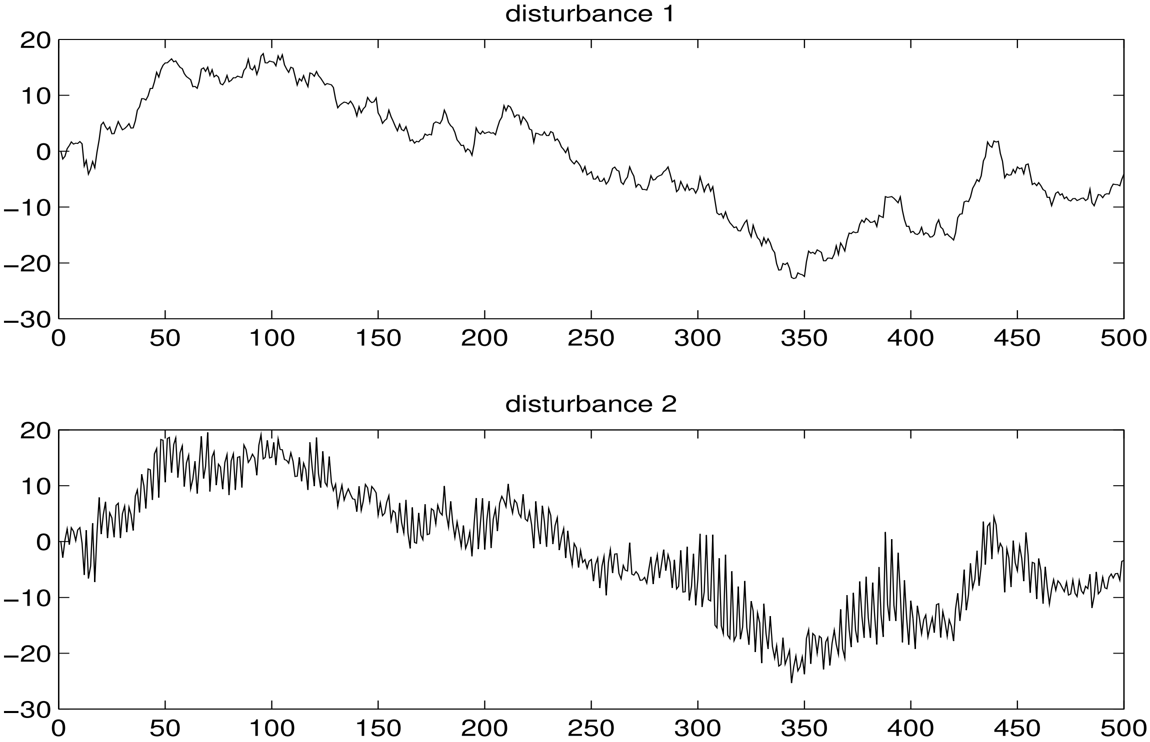

Figure 7: Disturbances

Although Σ_1 was poorly conditioned in each example, the measure $\bar{\delta}_1$ from LP 1 gives similar diagnosis as comparing the GLR to the threshold T_1. The measure $\bar{\delta}_2$ from LP 2 is smaller than $\bar{\delta}_1$, implying that by using the eigenvectors of Σ_1 as coordinates, the perturbation which must be added to $\hat{\theta}_1$ so that the sum lies within Θ_0 is smaller than when the standard coordinate system is used. For the plants P_1, P_3, P_4, P_5, P_6, and P_7, the GLR test gives correct performance diagnosis with either threshold T_1 or T_2, and for both disturbances d_1 and d_2. For the plant P_2, a correct diagnostic conclusion will be drawn when the threshold T_1 is used, but not when threshold the T_2 is used. The performance deterioration for P_2 is the least for all of the poor performing models, so incorrect diagnosis in this case is less severe than it would be in the others. For plant P_8, all cases give correct diagnosis except when the threshold T_1 is used and the disturbance corresponds to d_1; however, the threshold T_1 contains a rather ad hoc tuning parameter r^* or δ^*, and choosing a smaller value of r^* will give a correct diagnosis. That the GLR is less than the threshold T_1 approximately implies that for parameters θ to have generated $\mathcal{Y}$ with the same likelihood as as $\hat{\theta}_0$, θ must be less than half as probable as $\hat{\theta}_1$.

Conclusions are more difficult to draw from the MV ratio. Consider first the case of disturbance d_1. For the plants with good performance, P_5, and P_8, the MV ratio are within 3% of the nominal

Table 2: Diagnostics, disturbance d_1

Plant	MV ratio	GLR	T_1	T_2	$\bar{\delta}_1$	$\bar{\delta}_2$	κ^*
P_1	1.04	0.959	0.736	0.439	0.26	0.15	176
P_2	1.17	0.155	0.734	0.0978	1.5	1.00	153
P_3	1.07	2.46e-05	0.75	0.0965	3.8	2.8	243
P_4	1.03	0.0142	0.746	0.602	2.4	1.6	256
P_5	1.02	0.8	0.735	0.473	0.60	0.38	161
P_6	1.34	5.44e-06	0.745	0.137	3.9	2.8	220
P_7	1.15	0.0293	0.74	0.289	2.6	1.4	198
P_8	1.07	0.707	0.737	0.451	0.7	0.4	182

Table 3: Diagnostics, disturbance d_2

Plant	MV ratio	GLR	T_1	T_2	$\bar{\delta}_1$	$\bar{\delta}_2$	κ^*
P_1	1.9	0.931	0.599	0.426	0.20	0.16	314
P_2	1.71	0.38	0.555	0.156	0.671	0.56	155
P_3	2.16	2.34e-05	0.632	0.375	2.2	2.5	773
P_4	2.22	0.00615	0.625	0.526	1.5	1.4	697
P_5	1.92	0.837	0.609	0.473	0.30	0.30	345
P_6	1.74	0.000542	0.55	0.0881	2.0	1.6	134
P_7	1.83	0.0217	0.574	0.546	1.6	1.1	235
P_8	1.87	0.602	0.59	0.438	0.55	0.41	287

case, P_1; however, two of the poor performance plants, P_3 and P_4, also have MV ratios within 3% of the nominal case, whereas the GLR clearly indicates a deterioration in performance. In addition, although P_2 has the least performance deterioration it terms of the settling criterion, its MV ratio is the largest of any of the plants. In each case, this improper diagnosis is not due to poor estimation of the MV ratio, but rather to the fact that the MV ratio does not provide a good performance measure. In the case of the disturbance d_2, for each plant the MV ratio is significantly larger than when the simulation was carried out with d_1. Additionally, drawing conclusions about which of the models is performing worse, where performance is measured by the settling time criterion, is not possible from the MV ratio.

We would now like to consider four separate scenarios, and use the results in Tables 2 and 3 to evaluate how the GLR methods and the MV method would diagnose changes in process conditions. The scenarios will be:

1. Shift from P_1 to P_4 or P_6, constant disturbance spectrum d_1;
2. Shift from P_1 to P_4 or P_6, constant disturbance spectrum d_2;

3. Shift in disturbance spectrum from d_1 to d_2, constant plant P_1;
4. Shift in disturbance spectrum from d_2 to d_1, shift in plant from P_1 to P_4 or P_6.

Scenario 1 For a shift to P_4, the MV ratio test would indicate a 1% improvement in controller performance, incorrectly suggesting no retuning necessary, whereas for P_6, a 13% deterioration in controller performance would be noted. When the GLR test is applied, a clear deterioration for both cases P_4 and P_6 is correctly diagnosed.

Scenario 2 Here, the opposite situation occurs. Using only the MV ratio, when the plant shifts to P_4, a correct deterioration in control is concluded (increased MV ratio 1.9 to 2.22), but a shift to P_6 gives an incorrect diagnosis (decreased MV ratio 1.9 to 1.74). Again, the GLR test gives a clear poor performance diagnosis in each case.

Scenario 3 When the plant remains unchanged, but the disturbance spectrum changes from d_1 to d_2, the MV ratio indicates a substantial decrease in performance. Although the theoretical MV ratio actual does increase, the increased variance is due to the high frequency component of the disturbance d_2. Retuning the controller would require that the model accurately describe the behavior of the system at high frequency. Because the high frequency modes of the system may be difficult to identify, retuning may not be desirable. On the other hand, the GLR test, through the filtered coefficient constraints, considers only the low frequency phenomena, and indicates that performance remains satisfactory.

Scenario 4 If the disturbance changes from d_2 to d_1, the MV ratio test would indicate substantial improvements in performance, regardless of any changes in the plant. In particular, when the plant shifts from P_1 to either P_4 or P_6, the MV ratio diagnosis would be incorrect. On the other hand, the GLR would correctly diagnose that the performance has deteriorated.

Comparing the threshold T_1 in Tables 2 and 3, we see that lower values of the threshold are obtained in case of d_2. This is expected because although in the case of disturbance d_1, the closed loop system can be very accurately modeled by 20 impulse response coefficients, for disturbance d_2, which contains a slowly decaying, high frequency oscillation, more than 20 coefficients are needed. Since the unconstrained structure Θ_1 is a MA(20) model, the true system with disturbance d_2 is not contained within Θ_1, and the estimated covariance Σ_1 will be larger than in the case of d_1, resulting in the smaller threshold T_1 observed in Table 3.

□

Example 2 This example addresses the control of overhead temperature for a distillation column. The data used have been made publicly available by Shell Research Company, and a full description of the system may be found in [8], wherein it is referred to as *Column 2*. Although the complete data set contained measurements for 80,000 time samples, only the smaller segment of 1000 points shown in Figure 8 are analyzed in this example. From visual inspection, it is clear that a significant change in the overall process occurred after the first 500 samples.

The performance constraints were specified as in the previous example, with the settling time set to 10, which represented the open loop settling time for a step response, as indicated by the step

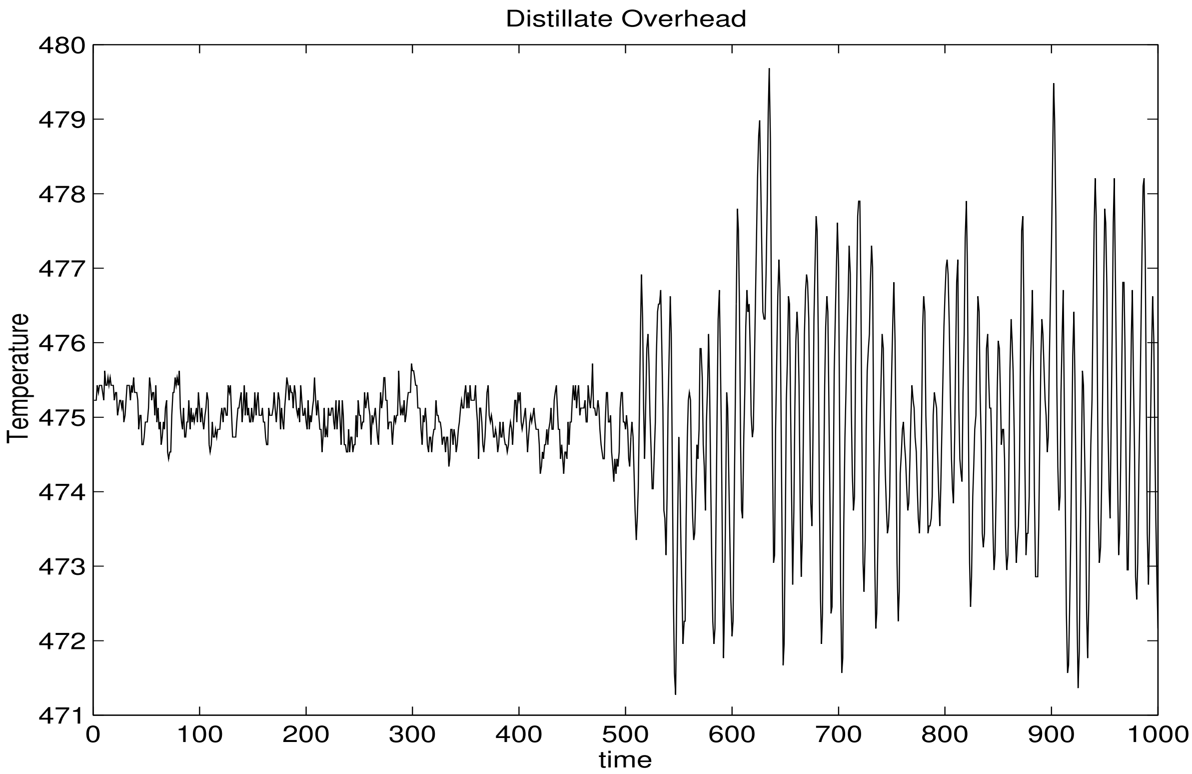

Figure 8: Shell Column 2, Overhead Temperature

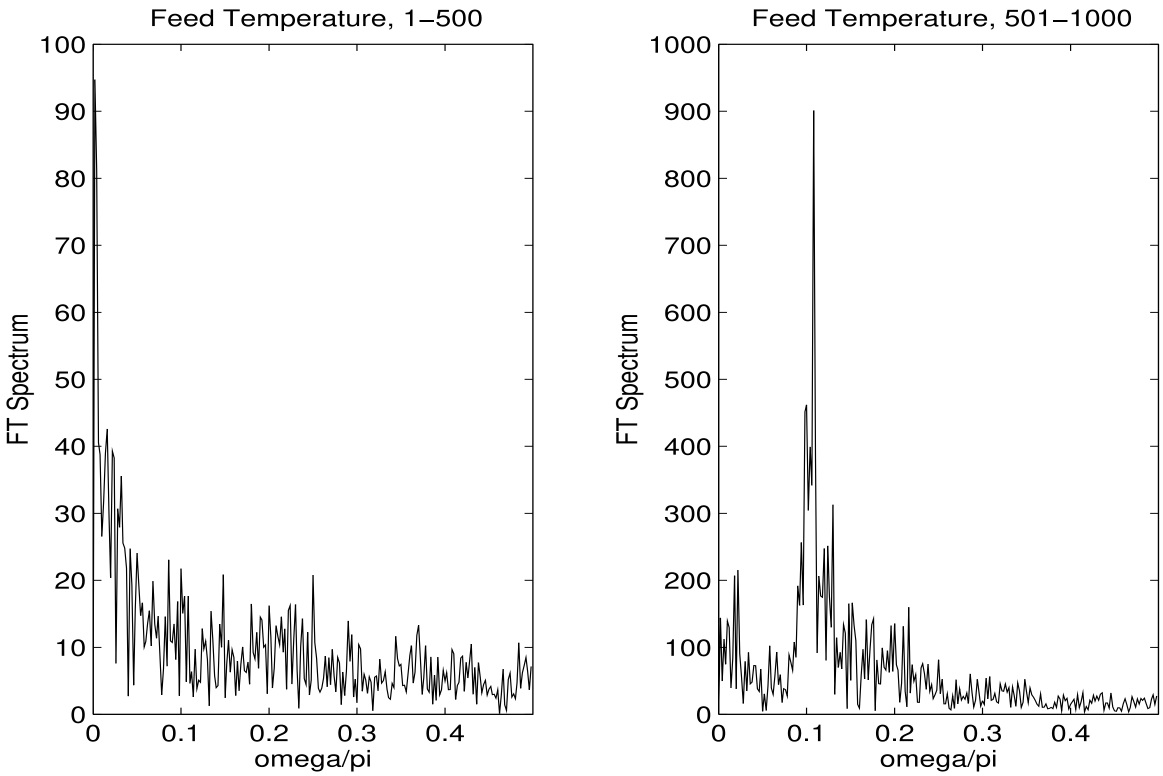

Figure 9: Shell Column 2, Feed Temperature Spectrums

Table 4: Diagnostics, Shell Column 2

Sequence	MV ratio	GLR	$T_1(0.5)/T_1(0.25)$	T_2	$\bar{\delta}_1$	$\bar{\delta}_2$	κ^*
$1 - 500$	1.16	0.46	0.66/0.44	0.44	1.06	0.64	69
$501 - 1000$	1.30	3.6e-28	0.61/0.38	> 1	5.7	5.0	1102

response supplied with the data. Θ_1 and Θ_0 were the set of all MA(25) and the set of all MA(25) models which satisfied the constraints, respectively. Table 4 contains the results of the various diagnostic tests. The threshold T_1 was calculated for both $r^* = 0.5$ and $r^* = 0.25$, and T_2 for $m = 1$. The MV ratio indicates only a slight decrease in the performance of segment 1 compared to segment 2, suggesting that the primary source of the poor quality output is due to an increase in the energy of the disturbances. On the other hand, the GLR tests indicate that for the first segment, the performance is borderline acceptable, whereas for the second test, it is clearly below the acceptable limit.

The data also contain measurements of the feed temperature for the column input. Although this measurement was not used in calculating the control moves, it can provide some insight to the proper diagnosis. The power spectrums for the feed temperature for each of the two data segments are shown in Figure 9. For the segment from $t \in [501, 1000]$, the energy of this disturbance is significantly higher (approximately 60 times). This accounts for much of the increased energy in the overhead temperature. In addition to the increased energy, the second segment also has a large low frequency component at approximately $\omega = 0.1\pi$. Similarly, the unconstrained model $\hat{\theta}_1$ fit to this segment has a low frequency oscillation which does not quickly decay. Since the constrained model $\hat{\theta}_0$ must be small after 10 time samples, $\hat{\theta}_1$ fits the data much better than $\hat{\theta}_0$ and the GLR is small. Thus the GLR test correctly diagnoses that the change in the signal is not only due to increased disturbance energy but also to a change in controller performance due to the disturbance spectrum. The decision whether to retune the controller should depend on several other factors, which may include the time required to retune and knowledge about the transient behavior of the disturbance. For example, if experience or other process knowledge indicated that the disturbance is likely to have the same characteristics for a considerable time, retuning the controller may be profitable.

□

7 Conclusions

In this paper, we have shown how many common and practical controller performance criteria can be expressed as linear constraints on the closed loop impulse response coefficients. Using this type of criterion, performance monitoring can be formulated as a generalized likelihood ratio test. Evaluating the GLR involves solving a constrained as well as an unconstrained model identification problem. In order to evaluate performance, the GLR must be compared to a threshold.

Three approaches to selecting the threshold have been discussed. The first two methods use the covariance of the estimated unconstrained model. The method of Section 4.1 can be interpreted as calculating the maximum relative probability that the system is described by a model with the same likelihood as the optimal constrained model, given the unconstrained estimated is distributed normally with mean $\hat{\theta}_1$ and covariance Σ_1. The method of Section 4.2 tests to see if a set of parameters exists, with each individual parameter within some confidence limit of its optimal estimate, which satisfies the performance bounds. Alternatively, cross-correlations between the parameter estimates can also be incorporated using a modified linear program. Finally, in Section 4.3 a threshold selection with a cross-validation interpretation is derived.

When the operating data used to calculate the performance measure contain set point changes, and the reference signal excites the system sufficiently, the GLR methods can be extended directly by specifying performance bounds for reference tracking, and parameterizing the transfer function from the reference signal to the tracking error. More frequently, changes in the set point will not produce a sufficiently exciting signal, but by subtracting from the tracking error the nominal response to the reference signal, the frequency of false poor performance diagnoses should be decreased.

The examples of Section 6 showed that for a meaningful performance objective, the likelihood ratio methods gave correct diagnoses whereas the minimum variance ratio of Harris [7] did not. In particular, we demonstrated through Example 1 that the MV ratio test can be strongly influenced by changes in the high frequency component of a disturbance which may not merit retuning, and may at the same time be insensitive to model changes which do result in deteriorated controller performance and can be easily rectified by retuning the controller. On the other hand, by properly selecting the performance objective, the GLR test can be made insensitive to irrelevant changes in the disturbance dynamics, while maintaining high sensitivity to model changes. Finally, for the examples considered, the measures $\bar{\delta}_1$ or $\bar{\delta}_2$, while easier to compute, gave diagnoses consistent with those obtained by comparing the GLR to the thresholds T_1 and T_2.

References

[1] M. Basseville and I. Nikiforov. *Detection of Abrupt Changes: Theory and Applications.* Prentice-Hall Information and Systems Series. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1993.

[2] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty. *Nonlinear Programming, Theory and Algorithms.* John Wiley & Sons, second edition, 1993.

[3] L. Desborough and T. Harris. Performance assessment measures for univariate feedback control. *Can. J. of Chem. Eng.*, 70:1186–1197, Dec. 1992.

[4] L. Desborough and T. Harris. Performance assessment measures for univariate feedforward/feedback control. *Can. J. of Chem. Eng.*, 71:605–616, Aug. 1993.

[5] G. A. Dumont. Fifteen years in the life of an adaptive controller. In *IFAC Symposium on Adaptive Systems in Control and Signal Processing*, pages 261–272, Grenoble, France, July 1992. IFAC.

[6] P. S. Fasolo and D. E. Seborg. Monitoring and fault detection for an HVAC control system. In *IFAC International Symposium, ADCHEM 94*, pages 537–542, Kyoto, Japan, 1994.

[7] T. J. Harris. Assessment of control loop performance. *Can. J. of Chem. Eng.*, 67:856–861, 1989.

[8] D. J. Kozub and C. García. Monitoring & diagnosis of automated controllers in the chemical processing industries. In *AIChE Annual Meeting*, St. Louis, MO, 1993.

[9] J. N. Little and L. Shure. *Signal Processing Toolbox.* The Mathworks, Inc., Natick, MA, 1992. Computer Software.

[10] L. Ljung. *MATLAB System Identification Toolbox User's Guide.* The Mathworks, Inc., Sherborn, Massachusetts, 1986.

[11] L. Ljung. *System Identification: Theory for the User.* Prentice-Hall Information and System Sciences Series. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1987.

[12] J. H. McClellan and T. W. Parks. A unified approach to the design of optimum FIR linear phase digital filters. *IEEE Trans. Circ. Theory*, CT-20:697–701, 1973.

[13] M. Morari and E. Zafiriou. *Robust Process Control.* Prentice-Hall, Inc., Englewood Cliffs, N.J., 1989.

[14] A. V. Oppenhiemer and R. W. Schafer. *Discrete-Time Signal Processing.* Prentice-Hall, 1989.

[15] T. Söderström and P. Stoica. *System Identification.* Prentice Hall International Series in Systems and Control Engineering. Prentice Hall, New York, 1989.

[16] N. Stanfelj, T. E. Marlin, and J. F. MacGregor. Monitoring and diagnosing process control performance: The single-loop case. *Ind. Eng. Chem. Res.*, 32:301–314, 1993.

[17] P. Stoica, P. Eykhoff, P. Janssen, and T. Söderström. Model-structure selection by cross validation. *Int. J. of Con.*, 4(6):1841–1878, 1986.

[18] M. Stone. Cross-validatory choice and assessment of statistical predictions. *J. Roy. Statist. Soc. Ser. B*, 36:111–147, 1974.

[19] M. L. Tyler and M. Morari. Performance assessment for unstable and nonminimum-phased systems. In *IFAC workshop on on-line fault detection and supervision in the chemical process industries*, Newcastle-upon-Tyne, England, 1995. To be presented.

[20] B. Weinstein. A sequential approach to the evaluation and optimization of control system performance. In *Proc. American Control Conf.*, pages 2354–2358, Chicago, June 1992.