Cautious Controller Tuning

Robert R. Bitmead**
Department of Mechanical & Aerospace Engineering
University of California San Diego
9500 Gilman Drive
La Jolla CA 92093-0411, USA
rbitmead@ucsd.edu
Fax: +1 858 822 3107

Brian D.O. Anderson[†]
Department of Systems Engineering
RSISE Building 115
Australian National University
Canberra ACT 0200, Australia
brian.anderson@anu.edu.au

Michel Gevers[‡]
Centre for Systems Engineering & Applied Mechanics
Université Catholique de Louvain
Bâtiment Euler
B-1348 Louvain-La-Neuve, Belgium
gevers@auto.ucl.ac.be

Leonardo C. Kammer
Pulp & Paper Centre
The University of British Columbia
2385 East Mall
Vancouver, BC V6T 1Z4 Canada
kammerl@ppc.ubc.ca

December 5, 2000

Abstract

The problem of controller tuning from one stabilizing value to another with a stability and/or performance guarantee is broached using the ν-metric of Vinnicombe. This is titled *Cautious Controller Tuning* and links the ν-distance between controllers to the stability margin. A number of scenarios are considered which treat cases in which different information is available about the closed-loop system. This includes estimation of the stability margin from closed-loop data and joint model and controller adjustment.

Keywords: Controller Tuning, Robust Control, Adaptive Control

*Research supported by USA National Science Foundation Grant ECS-0070146 and DARPA Grant N00014-00-1-0799

†Research supported by USA Office of Naval Research Grant N00014-97-1-0946

‡This author acknowledges the Belgian Programme on Interuniversity Poles of Attraction, initiated by the Belgian State, Prime Minister's Office for Science, Technology and Culture. The scientific responsibility rests with the authors.

1 Adaptation, Tuning and Iterative Control

Controller Tuning encapsulates a number of areas of control design in which closed-loop data and performance are used as a basis to modify the feedback control law. At the heart of such schemes is the generation of a sequence of controllers and the development of a logical process to attempt to improve performance based on closed-loop data. There are several methods which immediately fall into this category, such as Adaptive Control [ÅW89], Iterative Feedback Tuning [HGGL98], Direct Iterative Tuning [KBB2K, KL2K], Iterative Identification and Control Design [ZBG95]. The differences between schemes rests in their use (or avoidance) of plant models, their computation of adjustment direction from data, and their update rate. We shall treat systems with slow update, in which a large block of data is acquired before a tuning step is made. In this framework, many results from linear theories, both robust control and system identification, carry over.

We study the inclusion of caution into controller adjustment so that, from a stabilizing controller, stability guarantees can be established for the subsequent controllers. Principal among the tools used is the ν-metric of Vinnicombe [Vin93], which describes a stability margin for a multivariable closed loop and then defines allowable variations in the plant or controller for which stability guarantees are made. The ν-metric extends the gap-metric but possesses an important advantage in being simply computable. Vinnicombe's result is that the allowable alteration to a controller with a stability guarantee is related to the current stability margin. This makes practical sense and augments earlier results in the gap metric.

In tandem with the computability of the ν-distance between two controllers, we also study the problem of the estimation of the stability margin from closed-loop operating data, without the need for a plant model. Because we are interested primarily in controller tuning, we formulate the approach in terms of controller variations. However, models of the plant do enter the analysis to facilitate the controller design. We introduce the dual problem of plant model variation to attempt to capture some features of Iterative Identification and Control Design. A number of people have considered cautious controller tuning using the ν-metric [BGP97, ABGK98, AG98, Bom98, KBB2K]. Here we combine and extend those results of the authors.

Controller tuning approaches and methods have arisen is several recent guises. Iterative Feedback Tuning [HGGL98], Direct Iterative Tuning [KBB2K] and Direct Adaptive Control [ÅW89] all are methods which seek to adjust the parameters of a controller using available closed-loop signals to indicate performance, without the use of an intermediary parametric model of the plant. In Iterative Feedback Tuning, an estimate is derived of the gradient of the performance objective function with respect to the controller parameters. This estimate is based on correlations of combinations of signals taken from the closed loop with the current controller. On the basis of this estimate the controller is modified in a negative gradient direction to decrease the objective function until a local minimum is achieved. Direct Iterative Tuning is conceptually akin except for the use of spectrum information, rather than correlations. The main advantage is the ability to develop unbiased estimates of the Hessian matrix to be used in conjunction with the gradient in, say, a Newton scheme. Direct Adaptive Control uses parametrized closed-loop dynamics (poles) to determine modifications for the current parameters. In each of these schemes, a direction of adjustment of the control parameters is determined via gradient or Newton means. The role of the ν-metric is then to provide admissible step-size information so as to constrain the adjustment and guarantee closed-loop stability. Such a metric is expressed in terms of the ν-distance between the new controller and the current one.

Iterative modeling and control design has been espoused by a number of researchers [ZBG95, dCVdH97, BYM92, SB93, LAKM93] as a practically useful approach to the joint development of control-oriented approximate models and model-based controllers. The methods proceed by repeated, interleaved stages of model fitting to closed-loop data with the current controller operating and of model-based control design, sometimes also using frequency weighting based on closed-loop data. There may be several controller re-design steps per identification experiment. Indeed, the *Windsurfer Approach* of [LAKM93] exploits a specific test for the need to improve model quality from design to design. The interleaving of identification and control means that sequences of both models and controllers need to be considered. The role of the ν-metric is to provide stabilization guarantees for both model and controller updates.

The outline of the paper is as follows. Section 2 presents some standard results on the ν-metric and on coprime factor descriptions of linear controllers. A new result on the variation of a stabilizing controller is proven which links the conservative limits of Vinnicombe with the Youla-Kucera parameter. Section 3 deals with the use of closed-loop data to estimate the stability margin of a plant-controller pair. This involves the estimation of an ∞-norm of a transfer function, for which an algorithm is given. Section 4 develops a sequence of problems and results to provide guidelines for the tuning of controllers on-line. An example is worked. Section 5 moves from controller tuning with either no plant model or a fixed model to the situation where both plant model and controller are adjusted using closed-loop data from experiments.

2 Guaranteed stability and Vinnicombe's ν-metric

Cautious controller tuning involves the modification of the current stabilizing controller to a new controller, while providing a guarantee of preservation of stability. One way of interpreting this is to say that the plant simultaneously stabilizes both controllers. This is an inversion of the usual formulation from robust control in which the controller is selected to stabilize simultaneously both the plant and the plant model. The method, of course, works for both formulations.

Let us denote the actual linear plant system by P and denote the sequence of successive controllers by $\{C_i : i = 1, 2, \dots\}$. Suppose that the current controller, C_1, stabilizes P in closed loop. The immediate question is: What conditions must be placed on the next controller, C_2, to guarantee the stability of the feedback pair (P, C_2)?

Glenn Vinnicombe in his 1993 paper [Vin93] provides the following response (adapted for controller tuning). Firstly, we make the following definitions.

Definition 1 (Condition $\mathcal{C}$) *Two controllers, $C_1(s)$ and $C_2(s)$, satisfy Condition $\mathcal{C}$ if*

$$\det(I + C_2^\star C_1)(j\omega) \neq 0,\ \forall\omega \ \textit{and}\ \text{wno}\det(I + C_2^\star C_1) + \eta(C_1) - \eta(C_2) = 0, \tag{1}$$

where wno$(\cdot)$ *indicates the winding number of the Nyquist diagram of the scalar transfer function, evaluated on a contour along the imaginary axis and indented to the right around any pure imaginary poles, and $\eta(C)$ is the number of open right-half-plane poles of* $\det(C)$.[1]

[1] We note that [ZD98] provides an alternative but completely equivalent definition using para-hermitian conjugate $\tilde{C}_2(s) = C_2^T(-s)$ in place of the hermitian conjugate used here, $C_2^\star(s) = \bar{C}_2^T(\bar{s})$, where $(\bar{\cdot})$ represents the complex conjugate.

Definition 2 (ν-metric) *The ν-distance, $\delta_\nu(C_1, C_2)$, between two controllers is given by*

$$\delta_\nu(C_1, C_2) = \begin{cases} \left\| (I + C_2 C_2^\star)^{-\frac{1}{2}} (C_2 - C_1)(I + C_1^\star C_1)^{-\frac{1}{2}} \right\|_\infty, & \text{if Condition } \mathcal{C} \text{ holds} \\ 1, & \text{else} \end{cases} \tag{2}$$

Definition 3 (Generalized Sensitivity Function) *The generalized sensitivity function of the plant-controller feedback pair (P, C) is given by*

$$T(P, C) = \begin{pmatrix} P(I + CP)^{-1}C & P(I + CP)^{-1} \\ (I + CP)^{-1}C & (I + CP)^{-1} \end{pmatrix}. \tag{3}$$

Definition 4 (Generalized Stability Margin) *The generalized stability margin of the plant-controller pair (P, C) is given by,*

$$b_{P,C} = \begin{cases} (\|T\|_\infty)^{-1}, & \text{if } (P, C) \text{ is stable} \\ 0, & \text{else.} \end{cases} \tag{4}$$

Definition 5 (Optimal Stability Margin) *The optimal stability margin of controller C is given by,*

$$b_{opt}(C) = \sup_P b_{P,C}. \tag{5}$$

Theorem 1 (Vinnicombe) *Consider a plant P and two controllers C_1 and C_2, with C_1 stabilizing P. The following results hold.*

[I] *(P, C_2) is stable for all controllers C_2 satisfying $\delta_\nu(C_1, C_2) \le \beta$ if and only if $b_{P,C_1} > \beta$.*

[II] *For a given $\beta < b_{opt}(C_1)$, (P, C_2) is stable for all plants, P, satisfying $\beta < b_{P,C_1} \le b_{opt}(C_1)$ if and only if $\delta_\nu(C_1, C_2) \le \beta$.*

[III] *If $\delta_\nu(C_2, C_1) < 1$ then*

$$\arcsin b_{P,C_2} \ge \arcsin b_{P,C_1} - \arcsin \delta_\nu(C_1, C_2), \tag{6}$$

and

$$\delta_\nu(C_1, C_2) \le \|T(P, C_1) - T(P, C_2)\|_\infty \le \frac{\delta_\nu(C_1, C_2)}{b_{P,C_1} b_{P,C_2}}. \tag{7}$$

The importance of Part [I] of Theorem 1 is that it provides a sufficient condition on a new controller, C_2, that will guarantee its stabilization of the plant P. This condition,

$$\delta_\nu(C_2, C_1) < b_{P,C_1} \tag{8}$$

is characterized by two features:

- the distance from a known stabilizing controller, C_1, which is computable in terms of the frequency responses of the two controllers, subject to the checkable winding number condition,

- the closed-loop performance currently achieved with the controller C_1, which is computed in terms of a norm associated with the frequency responses of sensitivity functions.

The closed-loop performance of the (P,C) pair is in part measured by $T(P,C)$, since the elements of T include both the sensitivity function and complementary sensitivity function, which define the disturbance rejection and tracking performance of the closed-loop, while the other functions describe control and plant gains. Thus, Part [III] of Theorem 1 indicates a bound on the closed-loop performance achieved by the stabilizing modified controller C_2.

In fact, one can see that the theorem effectively provides a sufficient condition for the retention of performance as well as stability. If (8) is strengthened to say

$$\delta_\nu(C_1,C_2) < 0.1\, b_{P,C_1},$$

then the right-hand inequality of (7) yields

$$\begin{aligned} \|T(P,C_1) - T(P,C_2)\|_\infty &\leq 0.1\, \|T(P,C_2)\|_\infty\,, \\ &\leq \frac{1}{9}\, \|T(P,C_1)\|_\infty\,. \end{aligned}$$

One upshot of these results is the intuitive property that well-behaved controllers, as measured by those with smaller $\|T(P,C)\|_\infty$, provide greater scope for variation before striking stability or performance guarantee barriers. This connection between robust stability and robust performance has been explored more fully in [MG90].

Later, we shall be interested in exploring these results in a situation where the transfer function $P(s)$ is not available, but experimental closed-loop data is available; we shall describe then how $b_{P,C}$ can be estimated from closed-loop data, even if the plant P is unknown. The immediate attraction of the ν-metric formulation for the successive design of controllers is that an *a priori* stability guarantee is provided in terms of an estimate of the current closed-loop performance even when the plant is not actually known.

A geometric interpretation of the ν-metric is available for scalar transfer functions as the maximal chordal distance between the projections of $C_1(j\omega)$ and $C_2(j\omega)$ on the Riemannian Sphere, see [Vin93a, ZD98, Bom98], while the generalized stability margin, $b_{P,C}$, is the minimal chordal distance between the projection of $P(j\omega)$ and $\frac{-1}{C(j\omega)}$ on the same sphere. One immediate insight from this interpretation is that the ν-metric is most sensitive to variations between controllers where either $C_1(j\omega)$ or $C_2(j\omega)$ is close to one in magnitude. Bombois [Bom98] uses this property to suggest techniques which keep $\delta_\nu(PC_1,PC_2)$ small in order to maintain frequency responses close at the gain cross-over frequency for the purposes of controller reduction.

It should be stated here that Theorem 1 is inherently conservative, since each part pertains to properties possessed of a class of systems. A direct (but still conservative) variant of this result[ZD98] affords the extension of the inequality (8) to hold on a frequency-by-frequency basis. In this fashion, a frequency function $b_{P,C_1}(j\omega)$ may be defined as the inverse of the maximal singular value of the generalized sensitivity matrix at a particular frequency or, alternatively, as the chordal distance between the projections of $P(j\omega)$ and $\frac{-1}{C(j\omega)}$ on the Riemann sphere. Similarly, one may define $\delta_\nu(C_1,C_2)(j\omega)$ and generalize the stability result as is done in the following Corollary 1. In the sequel for notational convenience we shall use the more conservative result of Theorem 1 while recognizing that this more detailed, relaxed variant is equally applicable. The winding number condition, $\mathcal{C}$, must apply in both cases. The non-conservative, necessary and sufficient conditions for simultaneous stabilization given by the Youla-Kucera parametrization will be discussed shortly. These latter conditions do not lend themselves simply to the introduction of a metric.

Corollary 1 ([ZD98]) *Define the frequency-dependent stability margin of the stable feedback pair* (P, C_1) *by*

$$b_{P,C_1}(\omega) \triangleq \sigma_{\max}^{-1}\left(T(j\omega)\right),$$

where $\sigma_{\max}(\cdot)$ *is the maximal singular value, and define the frequency-dependent distance between controllers at frequency* ω*,*

$$\delta_\nu\left(C_1, C_2, \omega\right) \triangleq \begin{cases} \sigma_{\max}\left((I + C_2(j\omega)C_2^\star(j\omega))^{-\frac{1}{2}}(C_2(j\omega) - C_1(j\omega))(I + C_1^\star(j\omega)C_1(j\omega))^{-\frac{1}{2}}\right) & \text{if Condition } \mathcal{C} \text{ holds} \\ 1 \quad \text{else} \end{cases}.$$

The feedback pair (P, C_2) *is internally stable if* $\delta_\nu(C_1, C_2, \omega) < b_{P,C_1}(\omega)$ *for all* ω*. Further*

$$\arcsin b_{P,C_2}(\omega) \geq \arcsin b_{P,C_1}(\omega) - \arcsin \delta_\nu(C_1, C_2, \omega).$$

We are able to develop the relationship between the difference between successive stability margins and the ∞-norm of the difference between the closed-loop generalized sensitivity functions T.

Corollary 2 *Consider the two stable closed-loop systems* (P, C_1) *and* (P, C_2)*. Then*

$$|b_{P,C_1} - b_{P,C_2}| \leq \delta_\nu(C_1, C_2). \tag{9}$$

Hence

$$|b_{P,C_1} - b_{P,C_2}| \leq \delta_\nu(C_1, C_2) \leq \|T(P, C_1) - T(P, C_2)\|_\infty \leq \frac{\delta_\nu(C_1, C_2)}{b_{P,C_1} b_{P,C_2}}. \tag{10}$$

Proof: The right side of (10) is a restatement of (7) of Theorem 1. Denote $T_1 = T(P, C_1)$, $T_2 = T(P, C_2)$. Since $T_1^{-1} - T_2^{-1} = T_2^{-1}(T_1 - T_2)\, T_1^{-1}$, using the triangle inequality we have

$$\left|\left\|T_1^{-1}\right\| - \left\|T_2^{-1}\right\|\right| \leq \left\|T_1^{-1} - T_2^{-1}\right\| \leq \|T_1 - T_2\| \left\|T_1^{-1}\right\| \left\|T_2^{-1}\right\|.$$

Using the right inequality of (7) yields the result. ∇∇∇

2.1 Coprime-Factorizations and the ν-metric

Our purpose is to consider applications of the ν-metric in the realm of controller tuning either with or without approximate models of the plant being available or adjusted. Because this aim sees us seeking guarantees of remaining within the class of stabilizing controllers, it is useful to introduce coprime factor descriptions of the plants and controllers. Coprime factorizations admit the generation of the entire class of stabilizing controllers for a particular plant and, within the constraints provided by the ν-metric, it makes sense to move within this class where possible. The parameter which completely describes the movement within this class is the Youla-Kucera parameter.

The three results that we recall and/or establish are expressions for $\delta_\nu(C_1, C_2)$ in terms of coprime factors, for $b_{P,C}$ in terms of coprime factors, and for $\delta_\nu(C_1, C_2)$ in terms of the Youla-Kucera parameter. These will be used subsequently for the study of movement between successive controllers. We begin with some standard definitions.

${\mathcal{RH}_\infty}^{p\times m}$ is the space of stable, proper $p \times m$ transfer function matrices. A *right coprime factorization* of a $p \times m$ not-necessarily-stable transfer function matrix P is a fraction, $P = XY^{-1}$, with $X \in {\mathcal{RH}_\infty}^{p\times m}$ and $Y \in {\mathcal{RH}_\infty}^{m\times m}$ with Y invertible at almost every point of the complex plane. A *left coprime factorization* is similarly defined with $P = \tilde{Y}^{-1}\tilde{X}$, $\tilde{X} \in {\mathcal{RH}_\infty}^{p\times m}$ and $\tilde{Y} \in {\mathcal{RH}_\infty}^{p\times p}$ and almost everywhere invertible.

A *right (left) normalized coprime factorization* is a right (left) coprime factorization which satisfies the further condition

$$X^\star X + Y^\star Y = I_m, \quad (\tilde{X}\tilde{X}^\star + \tilde{Y}\tilde{Y}^\star = I_p).$$

A *unit* in ${\mathcal{RH}_\infty}^{m\times m}$ is a transfer function matrix which is stable and proper, and has a stable, proper inverse. That is, if U is a unit in ${\mathcal{RH}_\infty}^{m\times m}$ then U^{-1} also is in ${\mathcal{RH}_\infty}^{m\times m}$. Given a plant P with right coprime factorization $P = XY^{-1}$ and controller C with left coprime factorization $C = \tilde{D}^{-1}\tilde{N}$, we may write the generalized sensitivity function matrix T as

$$T(P,C) = \begin{pmatrix} P(I+CP)^{-1}C & P(I+CP)^{-1} \\ (I+CP)^{-1}C & (I+CP)^{-1} \end{pmatrix} = \begin{pmatrix} X \\ Y \end{pmatrix} \left(\tilde{D}Y + \tilde{N}X\right)^{-1} \begin{pmatrix} \tilde{N} & \tilde{D} \end{pmatrix}. \tag{11}$$

Thus, if and only if $\tilde{D}Y + \tilde{N}X$ is a unit, is the feedback pair (P, C) internally stable, see [GL95].

We first recall the connections between the ν-metric and the normalized coprime factors.

Corollary 3 ([Vin93]) *Suppose there exist normalized coprime factor decompositions of the controllers* $C_1 = N_1 D_1^{-1} = \tilde{D}_1^{-1}\tilde{N}_1$ *and* $C_2 = N_2 D_2^{-1} = \tilde{D}_2^{-1}\tilde{N}_2$. *Then,*

$$\delta_\nu(C_1, C_2) = \begin{cases} \left\|\tilde{D}_2 N_1 - \tilde{N}_2 D_1\right\|_\infty, & \textit{if condition } \mathcal{C}' \textit{ holds} \\ 1, & \textit{else.} \end{cases} \tag{12}$$

Condition $\mathcal{C}'$: $\det(N_2^\star N_1 + D_2^\star D_1)(j\omega) \neq 0 \ \forall\omega$ *and* $\operatorname{wno}\det(N_2^\star N_1 + D_2^\star D_1)(s) = 0$.

Corollary 4 ([Vin93]) *Suppose the feedback pair* (P, C) *is internally stable and suppose there exist normalized coprime factorizations* $P = XY^{-1}$ *and* $C = \tilde{D}^{-1}\tilde{N}$. *Then,*

$$b_{P,C}(\omega) = \frac{1}{\sigma_{\max}\left[(\tilde{D}Y + \tilde{N}X)^{-1}\right]} \tag{13}$$

Proof: We have

$$\begin{aligned}
\lambda_{\max}(TT^\star) &= \lambda_{\max}\left(\begin{pmatrix} X \\ Y \end{pmatrix}\left(\tilde{D}Y + \tilde{N}X\right)^{-1}\begin{pmatrix} \tilde{N} & \tilde{D} \end{pmatrix}\begin{pmatrix} \tilde{N}^\star \\ \tilde{D}^\star \end{pmatrix}\left(\tilde{D}Y + \tilde{N}X\right)^{-\star}\begin{pmatrix} X^\star & Y^\star \end{pmatrix}\right) \\
&= \lambda_{\max}\left(\begin{pmatrix} X \\ Y \end{pmatrix}\left(\tilde{D}Y + \tilde{N}X\right)^{-1}\left(\tilde{D}Y + \tilde{N}X\right)^{-\star}\begin{pmatrix} X^\star & Y^\star \end{pmatrix}\right) \\
&= \lambda_{\max}\left(\left(\tilde{D}Y + \tilde{N}X\right)^{-\star}\begin{pmatrix} X^\star & Y^\star \end{pmatrix}\begin{pmatrix} X \\ Y \end{pmatrix}\left(\tilde{D}Y + \tilde{N}X\right)^{-1}\right) \\
&= \sigma^2_{\max}\left(\left(\tilde{D}Y + \tilde{N}X\right)^{-1}\right).
\end{aligned}$$

$\nabla\nabla\nabla$

We note that these results rely on the coprime factorizations being normalized. The next well-known theorem describes a parameter for moving between stabilizing but unnormalized coprime factorizations. Such a straightforward parametrization of normalized coprime factors is unavailable.

The following theorem is akin to Theorem 1, in that it describes a class of permissible controllers C_2, which, if substituted for a stabilizing controller C_1, ensure that stability is retained. However, Theorem 2 requires full knowledge of the plant, as opposed to Theorem 1 which requires a single number, a stability margin.

Theorem 2 *[GL95] Consider a plant P and controller C_1 with coprime factor descriptions over $\mathcal{RH}_\infty$, $P = XY^{-1} = \tilde{Y}^{-1}\tilde{X}$ and $C_1 = N_1 D_1^{-1} = \tilde{D}_1^{-1}\tilde{N}_1$. Controller C_2 will internally stabilize P if and only if it may be written as the following linear fractional transformation for some $Q \in \mathcal{RH}_\infty{}^{m\times p}$.*

$$\begin{aligned} C_2 &= (N_1 - YQ)(D_1 + XQ)^{-1} \\ &= \left(\tilde{D}_1 + Q\tilde{X}\right)^{-1}\left(\tilde{N}_1 - Q\tilde{Y}\right). \end{aligned} \tag{14}$$

The stable, proper transfer function Q above is known as the *Youla-Kucera Parameter* after the originators of the idea.

These earlier results provide a basis from which to express the ν-metric distance between two stabilizing controllers in terms of their Youla-Kucera parameter and coprime factors.

Corollary 5 *Consider an internally stabilizing feedback pair (P_1, C_1) with coprime factorizations $P_1 = XY^{-1} = \tilde{Y}^{-1}\tilde{X}$ and $C_1 = N_1 D_1^{-1} = \tilde{D}_1^{-1}\tilde{N}_1$, with $\tilde{D}_1 Y + \tilde{N}_1 X = I$ and $\tilde{Y}D_1 + \tilde{X}N_1 = I$ without loss of generality. Suppose that controller C_2 also internally stabilizes P_1 and so from (14) has coprime factorization*

$$C_2 = \left(\tilde{D}_1 + Q\tilde{X}\right)^{-1}\left(\tilde{N}_1 - Q\tilde{Y}\right) = \tilde{D}_2^{-1}\tilde{N}_2, \quad \text{for } Q \in \mathcal{RH}_\infty.$$

Then,

$$\begin{aligned} \delta_\nu(C_1, C_2) &= \begin{cases} \left\| \left(\tilde{D}_2\tilde{D}_2^\star + \tilde{N}_2\tilde{N}_2^\star\right)^{-\frac{1}{2}} Q\left(D_1^\star D_1 + N_1^\star N_1\right)^{-\frac{1}{2}} \right\|_\infty, & \text{if Condition } \mathcal{C} \text{ holds}, \\ 1, & \text{else} \end{cases} \\ &= \begin{cases} \left\| \left((\tilde{D}_1 + Q\tilde{X})(\tilde{D}_1^\star + \tilde{X}^\star Q^\star) + (\tilde{N}_1 - Q\tilde{Y})(\tilde{N}_1^\star - \tilde{Y}^\star Q^\star)\right)^{-\frac{1}{2}} Q\left(D_1^\star D_1 + N_1^\star N_1\right)^{-\frac{1}{2}} \right\|_\infty, \\ \qquad \text{if Condition } \mathcal{C} \text{ holds}, \\ 1, \qquad \text{else} \end{cases} \end{aligned} \tag{15}$$

Define the class of P_1-stabilizing controllers $\{C_\alpha : \alpha \in [0,1]\}$ by

$$C_\alpha = \left(\tilde{D}_1 + \alpha Q\tilde{X}\right)^{-1}\left(\tilde{N} - \alpha Q\tilde{Y}\right).$$

There exists a number $0 < \alpha_0 \leq 1$ such that C_1 and C_α satisfy condition $\mathcal{C}$ for all $\alpha \leq \alpha_0$ and $\delta_\nu(C_1, C_\alpha)$ is given by (15) with Q replaced by αQ.

Proof: From (2) and denoting $C_2 = \tilde{D}_2^{-1}\tilde{N}_2$ we have

$$\left\| (I + C_2 C_2^\star)^{-\frac{1}{2}} (C_2 - C_1)(I + C_1^\star C_1)^{-\frac{1}{2}} \right\|_\infty$$
$$= \left\| \left(\tilde{D}_2^{-1}(\tilde{D}_2\tilde{D}_2^\star + \tilde{N}_2\tilde{N}_2^\star)\tilde{D}_2^{-\star} \right)^{-\frac{1}{2}} \left(\tilde{D}_2^{-1}(\tilde{N}_2 D_1 - \tilde{D}_2 N_1) D_1^{-1} \right) \left(D_1^{-\star}(D_1^\star D_1 + N_1^\star N_1) D_1^{-1} \right)^{-\frac{1}{2}} \right\|_\infty .$$

Considering just the final few terms, we have the following.

$$\begin{aligned} \left[D_1^{-1} \left[D_1^{-\star} \left(D_1^\star D_1 + N_1^\star N_1 \right) D_1^{-1} \right]^{-\frac{1}{2}} \right]^{-1} &= \left[D_1^{-\star} \left(D_1^\star D_1 + N_1^\star N_1 \right) D_1^{-1} \right]^{\frac{1}{2}} D_1 \\ &= \left[\theta^\star \left(D_1^\star D_1 + N_1^\star N_1 \right)^{\frac{1}{2}} D_1^{-1} \right] D_1, \text{for some unitary matrix } \theta, \\ &= \theta^\star \left(D_1^\star D_1 + N_1^\star N_1 \right)^{\frac{1}{2}} . \end{aligned}$$

This yields the right side of (15). Note that, if for hermitian $X > 0$ we define $X^{\frac{1}{2}}$ as the unique positive definite square root of X, then any square Y which satisfies $Y^\star Y = X$ will necessarily be of the form $Y = \theta X^{\frac{1}{2}}$ for some unitary matrix θ.

The set of P_1-stabilizing controllers $\{C_\alpha : \alpha \in [0,1]\}$ contains C_1 at $\alpha = 0$ and C_2 at $\alpha = 1$. For $\alpha = 0$, Condition $\mathcal{C}$ holds and $\delta_\nu(C_1, C_\alpha) = 0$. If Condition $\mathcal{C}$ fails for some $\alpha \in [0,1]$ then there must hold for some smallest α_0 that $\det(I + C_{\alpha_0}^\star C_1)(j\omega) = 0$ and $\delta_\nu(C_1, C_\alpha) = 1$. The continuity of this determinant establishes the result. $\nabla\nabla\nabla$

The parameter α in Corollary 5 defines a homotopy of controllers C_α which stabilize P_1. As $\alpha \to 0$, $\delta_\nu(C_1, C_\alpha) \to 0$ continuously in α. This property will be used in constructing guaranteed stabilizing controllers. Controller C_2 constructed via the Youla-Kucera parameter $C_2 = (\tilde{D}_1 + Q\tilde{X})^{-1}(\tilde{N}_1 - Q\tilde{Y})$, need not lead to $\delta_\nu(C_1, C_2) < 1$. This is illustrated by the dual of the example of [BGB97], in which two plants, $P_1 = \frac{(s-1)}{(s-2)(s-3)}$ and $P_2 = \frac{-1.22}{(s+7.32)}$, are stabilized by the same controller, $C = 5.918$, but do not satisfy Condition $\mathcal{C}$.

Now consider the situation of model-based control design, where we are contemplating changing a known controller, C_1. We have also a known model, P_1, of a partly unknown plant P; we do however know, at least with high accuracy, b_{P,C_1}. We use the model P_1 as a basis for designing an improved controller C_2, but contemplate changing C_1 in the direction of C_2, but not necessarily as far as C_2 itself, in the following way. In case $\delta_\nu(C_1, C_2)$ as given by (15) exceeds b_{P,C_1}, we choose a controller C_α to use instead of C_2, requiring that $\delta_\nu(C_1, C_\alpha) < b_{P,C_1}$. Thus controller tuning may be conducted using nominal model P_1 and actual margin b_{P,C_1}. Corollary 5 allows us to compute allowable limits for C_α in terms of its Youla-Kucera parameter and the initial (P_1, C_1) designed loop. The latter part of the result states that a connected path of controllers may be constructed from C_1 until the condition guaranteeing stability is violated. This holds even if C_2 does not satisfy Condition $\mathcal{C}$ with C_1.

3 Estimating margins from closed-loop data

At the core of the use of the ν-metric for cautious controller or model adjustment is the need to compute or estimate the value of the ∞-norm of the closed-loop achieved sensitivity matrix $T(P, C)$. In this section we consider how this might be done using a finite set of discrete-time experimental data from the closed-loop and without a model for P[2]. The questions that we shall pose deal with two central issues:

- How ought closed-loop signal measurements be used to obtain an approximate overbound for $\|T(P,C)\|_\infty = b_{P,C}^{-1}$ or its frequency-by-frequency equivalent $\sigma_{\max}\big[T(P,C)(e^{j\omega})\big]$?
- What external excitation signals ought to be applied to the closed loop, and at which points in the loop, in order to enhance our ability to estimate this number?

3.1 Estimation of ∞-norms from data

The field of H_∞ System Identification [HJN91, GK92, NG95, MPG95] treats the computation of an entire causal, stable model $\hat{T}(z)$, which is ∞-norm close to the real closed-loop system $T(z)$, and provides a measure of the ∞-norm error, $\|T-\hat{T}\|_\infty \leq \gamma$. This information would suffice to bound $\|T\|_\infty \leq \|\hat{T}\|_\infty + \gamma$. Frequency weighting may be introduced and γ allowed to vary with frequency to yield tighter bounds. The requirement to compute the entire approximant, $\hat{T}$, is much greater than a requirement to estimate the ∞-norm of T. Nevertheless the methods are closely allied.

Evidently, if the error, γ, can be maintained small relative to $\|\hat{T}\|$, then $\|\hat{T}\|$ provides a reasonable estimate of $\|T\|$ without the need for further calculation — especially when other conservative safety factors are built in. From this point we shall assume this to be the case, although the references on H_∞ identification focus strongly on error bounds and we do indicate some connections as we proceed.

Throughout this section, our focus will be on the computation of the stability margin of the true system P with the current controller C_1 without using or generating an open-loop model for P. In this fashion the closed-loop data are allowed to speak for themselves as an indicator of the margin of stability. This data will then be used to constrain admissible variations of the controller, as indicated above.

We record data to allow us to estimate the complete generalized sensitivity function, T. We assume that neither the true plant, P, nor the controller, C, has any internal unstable, pole-zero cancellations. Thus any unstable near pole-zero cancellations between the plant and controller will be visible from the data recorded at the input and output of the plant and the controller, and therefore will appear as instabilities in T.

3.1.1 Estimation of frequency response values

The generalized sensitivity function in all cases of interest to us should be stable. It comprises at least four scalar transfer functions. Let us consider the estimation of its frequency response as a task of estimating that of (a number of) scalar stable transfer functions from input-output data. We consider firstly the problem of estimating a scalar frequency response at a single frequency from input-output data from a stable system. We recast this as a specialized open-loop identification problem before moving on to use these estimates as the basis for estimation of the stability margin.

[2] Without further comment, we note that the continuous-time results so far stated may be used *mutatis mutandis* in the discrete-time domain by calling upon, for example, the bilinear transform to map between the frequency domains. We switch to discrete time here because finite data sets only make sense in sampled data.

Consider the open-loop estimation of a scalar stable transfer function $G(z)$ from input-output data $\{y_t, u_t\}$ generated by

$$y_t = G(z)u_t + v_t,$$

where the disturbance process v_t is generated by filtering zero-mean white noise through a stable system and is independent from u_t. We assume that all the signals are quasi-stationary (so that spectra such as $\Phi_u(\omega)$ and $\Phi_v(\omega)$ are well defined) and we seek to fit a process model, $\hat{G}(z)$, and disturbance model, $\hat{H}(z)$, to minimize the summed squared filtered one-step-ahead prediction error. This is precisely the formalism of Ljung [Ljung99] and we appeal directly to his formulation of this prediction error criterion in the frequency domain to write

$$J = \frac{1}{N}\sum_{t=1}^{N}\left(y_t^f - \hat{y}_{t|t-1}^f\right)^2.$$

Under reasonable conditions such as quasi-stationarity,

$$\lim_{N\to\infty} J_N = \frac{1}{2\pi}\int_{-\pi}^{\pi}\left\{|G(e^{j\omega}) - \hat{G}(e^{j\omega})|^2\Phi_u(\omega) + \Phi_v(\omega)\right\}\frac{|L(e^{j\omega})|^2}{|\hat{H}(e^{j\omega})|^2}\,d\omega, \tag{16}$$

where $L(z)$ is the stable data weighting filter, $y_t^f = L(z)y_t$ is the filtered measurement and $\hat{y}_{t|t-1}^f$ is the one-step-ahead prediction of y_t^f using the model $(\hat{G},\hat{H})$ and the filtered input $u_t^f = L(z)u_t$. This formula describes the distribution in importance of various parts of the frequency range in achieving a fit to G and to the disturbance spectrum Φ_v.

Equation (16) describes the bias of the estimator based on this criterion. In the case where the model structure $(\hat{G},\hat{H})$ allows exact matching of the process, it is also possible to write down an approximate expression for the covariance of the transfer function estimate [Ljung99].

$$\text{Cov}\,\hat{G}_N(e^{j\omega}) \sim \frac{n}{N}\frac{\Phi_v(\omega)|L(e^{j\omega})|^2}{\Phi_u(\omega)|L(e^{j\omega})|^2}, \tag{17}$$

where n is the model order and N is the number of data as above. One might be tempted to cancel $|L(e^{j\omega})|$. We have left it in, since it is important to recognize the effect at frequencies where L is zero.

If we choose data weighting $L(z)$ to be narrow band-pass, then the identification concentrates around the passband frequencies of L. Indeed, if L were infinitesimally narrow-band about frequency ω_0, then fitting with this criterion while using a fixed (parameter-independent) noise model $\hat{H}$ would result in $\hat{G}(e^{j\omega_0}) \approx G(e^{j\omega_0})$, (see (16)) with a variance given by the noise-to-signal ratio at ω_0 divided by the number of data. This latter choice is used as a starting point in H_∞ system identification to produce an approximant with a frequency dependent error bound.

The assumptions needed to produce a close estimate $\hat{G}(e^{j\omega_0})$ of $G(e^{j\omega_0})$ are;

1. that the frequency response of $G(e^{j\omega})$ is smooth at $\omega = \omega_0$,
2. that the input signal has energy in this band, so that $\Phi_u(\omega_0)$ is strictly positive,
3. that the noise spectrum has finite energy in this band, so that $\Phi_v(\omega_0)$ is finite,
4. that a fixed noise model, $\hat{H}(z)$, is used in the identification.
5. that the initial condition effects in $L(z)$ have dissipated,

6. that sufficient data are used to be able to appeal to the statistical features of the Law of Large Numbers and the Central Limit Theorem underpinning the bias formula (16) and the variance formula (17),

7. that the model complexity is sufficient to capture the frequency response in the band, so that the variance formula (17) may be used.

Assumption 1 is a a form of heightened stability requirement on $G(z)$. Assuming the impulse response of G, $\{g_k\}$, to be absolutely summable we have formally

$$\frac{\partial G(e^{j\omega})}{\partial \omega} = \frac{\partial}{\partial \omega}\sum_{k=0}^{\infty} g_k e^{jk\omega} = \sum_{k=0}^{N} jkg_k e^{jk\omega}.$$

Accordingly, if $\sum k|g_k|$ is bounded then smoothness holds. Obviously if $|g_k|$ decays exponentially, there is smoothness with an improved bound for greater degree of stability.

Assumption 5 is frequently overlooked. It becomes increasingly an issue with diminishing bandwidth of filter $L(z)$, since the persistence of initial conditions is related to the bandwidth directly by the time-bandwidth product. It may be combined with Assumption 6 which requires that sufficiently many data are collected to admit the use of these formulæ to describe the bias and variance. In practical terms this means that N should exceed several time-constants of L. Note that, with N sufficiently large, the bias term should dominate the variance term in quantifying the error.

In our case of interest, where we wish to estimate $T(P,C)(e^{j\omega})$, once the individual scalar frequency responses of the elements of $T(P,C)$ are estimated at a particular frequency, ω_i, they may be "stacked up" to provide an estimate of the complete matrix $T(P,C)(e^{j\omega_i})$.

3.1.2 Estimation of margins

The above results concern the estimation of a single point of the frequency response of $G(z)$ or $T(P,C)$ after stacking the element estimates. To move from this point to the estimation of $\sigma_{\max}T(P,C)(e^{j\omega})$ requires the evaluation of the frequency response at a number, M, of points, typically an evenly-spaced grid at frequencies around the unit circle. Then the maximizer over frequency of the maximal singular value of the frequency response $T(P,C)$ is sought. A similar grid of points may also be used as the basis for estimation of a complete transfer function approximant, $\hat{T}(z)$, at all frequencies. Computation of $\|\hat{T}\|_\infty$ provides an alternative method to estimate the margin.

In practice, the filters $L(z)$ are chosen to be the M filters of the M-point Discrete Fourier Transform with centers at $2k\pi/M$. Such a choice of filter family has a number of beneficial features including ease of computation, familiarity, M-sample deadbeat response etc. The bandwidth of these filters is approximately $2\pi/M$ radians per second. Methods such as those described in adaptive frequency-sampling [BA81] can now be applied to conduct the regression estimation as above. Because the data-length, N, is much greater than the filter response time, M, the regression smoothes the filtered data in each bin.

An alternative approach is to choose an integer m so that mM is equal to the data length N. An mM-point DFT is performed on the data yielding mM complex frequency bin values. Of these, the zeroth and then every mth bin value is selected. The bandwidth of each of these filters is $2\pi/mM$ but some straightforward algebra shows that the equivalent filters are compositions of the M-point DFT followed by an m-point moving average. The computation of the frequency response estimate at the sample point proceeds by quotient of the respective bins of filtered input and output.

The kth filter ($k = 0, \ldots, M-1$) of an M-point DFT may be written as,

$$L_{k,M}(z) = \frac{1 - z^{-M+1}}{1 - e^{\frac{j2\pi k}{M}} z^{-1}} = 1 + e^{\frac{j2\pi k}{M}} z^{-1} + e^{\frac{j2\pi k2}{M}} z^{-2} \cdots + e^{\frac{j2\pi k(M-1)}{M}} z^{-M+1}.$$

The kmth filter of an mM-point DFT has the form

$$L_{km,mM}(z) = \sum_{\ell=0}^{mM-1} e^{\frac{j2\pi km\ell}{Mm}} z^{-\ell} = \sum_{i=0}^{m-1} z^{-Mi} \left(\sum_{\ell=0}^{M-1} e^{\frac{j2\pi k\ell}{M}} z^{-\ell} \right) = \sum_{i=0}^{m-1} z^{-Mi} L_{k,M}(z).$$

When combined with the stability margin of G, an estimate may be computed for the accuracy of each sample, which in turn can be used to evaluate an overbound for the total ∞-norm error. This is not pursued further here. As stated above the maximizer of the maximal singular value of the constructed frequency response samples of $T(P,C)$ is used as $b_{P,C}^{-1}$.

The steps to yield the full approximant $\hat{G}(z)$ to $G(z)$ from the frequency response samples are discussed in [HJN91, NG95], in the context of bounded measurement noise only. The frequency response samples are used to generate an impulse response, which is extended to a non-causal ∞-norm-close transfer function. The Nehari Extension is used to approximate this in ∞-norm by a stable, causal system. Given the approximant $\bar{G}(z)$ for each element of $T(P,C)(z)$, we compose an approximant to $T(P,C)$ and then use either $\|T(P,C)\|_\infty$ or $\|T(P,C)\|_\infty + \gamma$ as the estimate of $\|T\|_\infty$, where γ is the error bound carried through the computations. For large numbers of data, N, and large numbers of frequency sampling points, $M \ll N$, we may neglect the γ term provided excitation is adequate.

3.1.3 Signal excitation for margin estimation

There are four transfer functions to be estimated at points around the unit circle to compute $\|T\|_\infty$. As is amply illustrated by the bias and variance formulæ (16-17), broadband, sufficiently rich input signal spectrum content is a critical feature in exposing the transfer function to be estimated. This might be achieved by using, for example, a wideband stochastic input or an almost periodic signal composed of a sufficient set of sinusoids. Absence of input signal power at an estimation frequency can lead to bias and/or produce unmanageable variance. To overcome the possible absence of signal power due to the effects of the controller or of the plant, excitation reference signals should be introduced both before and after the controller[3].

For our purposes, let the input to the controller be $r_t - y_t$ and the input to the plant be $u_t = C(z)(r_t - y_t) + s_t$, where r_t and s_t are external, known excitation signals assumed independent from each other and from the process disturbance v_t. Then

$$y_t = \frac{PC}{1+PC} r_t + \frac{P}{1+PC} s_t + \frac{1}{1+PC} v_t \quad (18)$$

$$u_t = \frac{C}{1+PC} r_t + \frac{1}{1+PC} s_t - \frac{C}{1+PC} v_t. \quad (19)$$

[3]The introduction of excitation signal s_t after the controller is important when the controller has poles on the unit circle, since then the controller's presence prevents identifiability from r_t at these poles of those elements of $T(P,C)$ without C in the numerator. As C is usually known, the need for this additional s_t excitation can be simply evaluated.

Note that each of the transfer functions comprising T is available by regressing y_t or u_t (the closed-loop data) on either r_t or s_t. Further, there is no additional filtering on these signals which might compromise the quality of transfer function estimation. Accordingly, if C is not well known, we propose that the closed-loop experiments be conducted with wideband, independent excitations r_t and s_t. If C is known and does not contain poles on the unit circle, then the excitation signal r_t may be taken as zero and the estimates of $\frac{PC}{1+PC}$ and $\frac{C}{1+PC}$ may be computed from those of $\frac{P}{1+PC}$ and $\frac{1}{1+PC}$.

4 Cautious Controller Adjustment

We have presented in Section 1 results pertaining to the use of the ν-metric in guaranteeing closed-loop stability with a new controller. The condition is expressed in terms of the ν-distance between controllers and the closed-loop stability margin. Section 3 examined the estimation of this stability margin directly from closed-loop data without the need for an intervening plant model. Here we shall explore the combination of these ideas in controller tuning with guarantees of stability.

The approach, following [AG98], is by considering a sequence of indicative problems, which build on earlier results. Recall that the calculation of $b_{P,C}$ from data did not require either the knowledge of the plant or the use of an open-loop plant model. Joint model and controller adjustment with caution will be considered in the next section.

Problem 1 — Known margin

Consider an unknown plant P and known current stabilizing controller C_i with the closed loop yielding a known generalized stability margin b_{P,C_i}. What restriction must be placed on the next controller C_{i+1} in order to guarantee stabilization of P?

The solution to Problem 1 is directly given by Theorem 1: a sufficient condition for the stability of the (P, C_{i+1}) loop is that $\delta_\nu(C_{i+1}, C_i) < b_{P,C_i}$. Data for verifying the inequality are available.

In the event that the controller is locally continuously parametrized by a real vector ρ_i, so that $C_i = C(\rho_i)$, we may seek to modify the controller C_i along the negative gradient direction $-\nabla_\rho J$, where J is the performance objective to be minimized. That, is we adjust the controller via

$$\rho_{i+1} = \rho_i - \beta M \nabla_\rho J(\rho_i), \tag{20}$$

where β is the step-size of the update and M is a non-negative definite matrix, possibly an estimate of the Hessian. This is a version of Problem 1 in which we seek to choose β to ensure that $C_{i+1} = C(\rho_{i+1})$ is sufficiently close to C_i as measured by the ν-metric. In practice this results in the acceptance of a β which achieves

$$0 < \gamma_1 b_{P,C_i} < \delta_\nu(C_{i+1}, C_i) < \gamma_2 b_{P,C_i} < b_{P,C_i},$$

for some predetermined factors $\{\gamma_1, \gamma_2\} \in (0,1)$. This is the approach adopted in [KBB2K] and expanded in [Kam98]. The factor γ_1 seeks to ensure a sufficiently, nontrivial move from C_i, while γ_2 is a conservative safety factor coping with the stability guarantee. In principle we could take $\gamma_2 = 1$.

We note also, from (15) and Corollary 5, that for controller tuning involving adjustment of coprime factors via the Youla-Kucera parameter, one may also generate a class of controllers, $\mathcal{C}_\alpha,\ \alpha \in [0,1]$, where the Youla-Kucera parameter is αQ. For small α, the expression (15) for $\delta_\nu(C_i, C_{i+1})$ has a simple close-to-linear dependence on α, which lends an easy interpretation of the manageable change. Of course, to work with a Youla-Kucera parametrization, one requires coprime factors of the nominal plant model stabilized by C_i.

Problem 2 — Close closed loops

Problem 1 is a direct application of Vinnicombe's result, modified to describe controller tuning rather than robust control, where the one controller, C, is designed to stabilize the model, P_0, and we ask also that it stabilize the true plant, P. Within our controller tuning framework, we now ask about the modification of a controller which stabilizes both a model and the true plant. In particular, we consider the permitted variation of a controller when the true and model-based closed-loop generalized sensitivity functions are close. This is a problem reminiscent of questions arising in closed-loop system identification; some closed-loop identification schemes focus on securing a model P_0 such that $T(P, C_i) - T(P_0, C_i)$ is small in some norm, for example [ZBG95] uses the L_2 norm.

> *Consider an unknown plant P and nominal model P_0 with known controller C_i stabilizing both P and P_0. The achieved generalized sensitivity function $T(P, C_i)$ is unknown, as is the stability margin b_{P,C_i}, but assume that $\|T(P, C_i) - T(P_0, C_i)\|_\infty$ (or a bound on it) is known. What restriction must be placed on the next controller C_{i+1}, which stabilizes P_0, in order to guarantee stabilization of P?*

A sufficient answer is available directly from Theorem 1. We have $\delta_\nu(P, P_0) \le \|T(P, C_i) - T(P_0, C_i)\|$ from an obvious variant of (7). Further, from (6) we have

$$\begin{aligned} \arcsin b_{P,C_i} &\ge \arcsin b_{P_0,C_i} - \arcsin \delta_\nu(P, P_0) \\ \arcsin b_{P,C_{i+1}} &\ge \arcsin b_{P,C_i} - \arcsin \delta_\nu(C_{i+1}, C_i), \end{aligned}$$

whence

$$\begin{aligned} \arcsin b_{P,C_{i+1}} &\ge \arcsin b_{P_0,C_i} - \arcsin \delta_\nu(P, P_0) - \arcsin \delta_\nu(C_{i+1}, C_i) \\ &\ge \arcsin b_{P_0,C_i} - \arcsin \|T(P, C_i) - T(P_0, C_i)\|_\infty - \arcsin \delta_\nu(C_{i+1}, C_i). \end{aligned}$$

To ensure that C_{i+1} stabilizes P we need to ensure that $\arcsin b_{P,C_{i+1}} > 0$, which is achieved provided

$$\arcsin \delta_\nu(C_{i+1}, C_i) \le \arcsin b_{P_0,C_i} - \arcsin \|T(P, C_i) - T(P_0, C_i)\|_\infty.$$

As remarked above, this is a sufficient condition. A more complete answer is available, given in Theorem 3 below. Suppose that we have the following coprime factorizations,

$$C_i = N_i D_i^{-1} = \tilde{D}_i^{-1} \tilde{N}_i, \text{ both normalized,}$$

$$P_0 = X_0 Y_0^{-1} = \tilde{Y}_0^{-1} \tilde{X}_0, \text{ not necessarily normalized.}$$

Since C_i stabilizes P_0 we may also choose the fractional descriptions of P_0 such that without loss of generality,

$$\tilde{D}_i Y_0 + \tilde{N}_i X_0 = \tilde{Y}_0 D_i + \tilde{X}_0 N_i = I.$$

Since P is also stabilized by C_i, we may write it as

$$P = (X_0 - D_i R)(Y_0 + N_i R)^{-1}, \text{ for some (unknown) } R \in \mathcal{RH}_\infty.$$

Now, from (11), we may write

$$\begin{aligned}
T(P_0, C_i) &= \begin{pmatrix} X_0 \\ Y_0 \end{pmatrix} \left(\tilde{D}_i Y_0 + \tilde{N}_i X_0\right)^{-1} \begin{pmatrix} \tilde{N}_i & \tilde{D}_i \end{pmatrix} = \begin{pmatrix} X_0 \\ Y_0 \end{pmatrix} \begin{pmatrix} \tilde{N}_i & \tilde{D}_i \end{pmatrix}. \\
T(P, C_i) &= \begin{pmatrix} X_0 - D_i R \\ Y_0 + N_i R \end{pmatrix} \left(\tilde{D}_i Y_0 + \tilde{N}_i X_0 + (\tilde{N}_i D_i - \tilde{D}_i N_i) R\right)^{-1} \begin{pmatrix} \tilde{N}_i & \tilde{D}_i \end{pmatrix}, \\
&= \begin{pmatrix} X_0 - D_i R \\ Y_0 + N_i R \end{pmatrix} \begin{pmatrix} \tilde{N}_i & \tilde{D}_i \end{pmatrix}.
\end{aligned}$$

It follows that

$$T(P, C_i) - T(P_0, C_i) = \begin{pmatrix} -D_i \\ N_i \end{pmatrix} R \begin{pmatrix} \tilde{N}_i & \tilde{D}_i \end{pmatrix}. \tag{21}$$

From (21) and the normalization property of both factorizations of C_i, there holds

$$\|T(P, C_i) - T(P_0, C_i)\|_\infty = \|R\|_\infty.$$

The Youla-Kucera parameter, R such that $\|R\|_\infty \leq \epsilon$, generates the whole class of systems which satisfy $\|T(P, C_i) - T(P_0, C_i)\|_\infty \leq \epsilon$. Thus to this point, we have proved.

Lemma 1 *Consider an unknown plant P and known model P_0 stabilized by a controller C_i with corresponding generalized sensitivity function matrices $T(P, C_i)$ and $T(P_0, C_i)$. Suppose that $P_0 = X_0 Y_0^{-1} = \tilde{Y}_0^{-1} \tilde{X}_0$ are coprime fractional descriptions of P_0, and that $N_i D_i^{-1}$ is a normalized coprime fractional description of C_i such that $\tilde{Y}_0 D_i + \tilde{X}_0 N_i = I$. Then the set of all P stabilized by C_i and satisfying*

$$\|T(P, C_i) - T(P_0, C_i)\|_\infty \leq \epsilon$$

is given by

$$P = (X_0 - D_i R)(Y_0 + N_i R)^{-1}, \text{ for } R \in \mathcal{RH}_\infty \text{ with } \|R\|_\infty < \epsilon.$$

This leads us directly to the following corollary, which will be used shortly.

Corollary 6 *Consider the two plant models, $P_i = X_i Y_i^{-1} = \tilde{Y}_i^{-1} \tilde{X}_i$ and $P_{i+1} = X_{i+1} Y_{i+1}^{-1} = \tilde{Y}_{i+1}^{-1} \tilde{X}_{i+1}$, which satisfy Condition $\mathcal{C}$ and which both are stabilized by the controller $C_i = N_i D_i^{-1} = \tilde{D}_i^{-1} \tilde{N}_i$, with $\tilde{D}_i Y_i + \tilde{N}_i X_i = I$ and $\tilde{Y}_i D_i + \tilde{X}_i N_i = I$. Then we have*

$$\begin{aligned}
\delta_\nu(P_{i+1}, P_i) &\leq \|T(P_i, C_i) - T(P_{i+1}, C_i)\|_\infty \\
&= \left\| \begin{pmatrix} -D_i \\ N_i \end{pmatrix} R_i \begin{pmatrix} \tilde{N}_i & \tilde{D}_i \end{pmatrix} \right\|_\infty, \\
&= \|R_i\|_\infty, \text{ if the controller coprime factorizations are normalized,}
\end{aligned} \tag{22}$$

Likewise, suppose the two controllers C_i and C_{i+1}, with $C_i = N_i D_i^{-1} = \tilde{D}_i^{-1}\tilde{N}_i$ and $C_{i+1} = N_{i+1}D_{i+1}^{-1} = \tilde{D}_{i+1}^{-1}\tilde{N}_{i+1}$, satisfy Condition $\mathcal{C}$ and both stabilize the plant model $P_{i+1} = X_{i+1}Y_{i+1}^{-1} = \tilde{Y}_{i+1}^{-1}\tilde{X}_{i+1}$ with $\tilde{D}_{i+1}Y_i + \tilde{N}_{i+1}X_i = I$ and $\tilde{Y}_i D_{i+1} + \tilde{X}_i N_{i+1} = I$. Then we have

$$\begin{aligned}
\delta_\nu(C_{i+1}, C_i) &\le \|T(P_{i+1}, C_i) - T(P_{i+1}, C_{i+1})\|_\infty \\
&= \left\| \begin{pmatrix} X_{i+1} \\ Y_{i+1} \end{pmatrix} Q_i \begin{pmatrix} -\tilde{Y}_{i+1} & \tilde{X}_{i+1} \end{pmatrix} \right\|_\infty, \qquad (23) \\
&= \|Q_i\|_\infty, \text{ if the plant coprime factorizations are normalized.} \qquad (24)
\end{aligned}$$

Now, to answer Problem 2 and because no information other than the magnitude of the difference between the generalized sensitivity functions is available, we shall seek the set of controllers which is guaranteed to stabilize the whole class of Ps defined by this difference.

Theorem 3 *Consider an unknown plant P and a known model P_0 both stabilized by a controller C_i. Suppose that we are given a known bound on the differences between the generalized sensitivity functions*

$$\|T(P, C_i) - T(P_0, C_i)\|_\infty \le \epsilon.$$

Let X_0, Y_0, $\tilde{X}_0$, $\tilde{Y}_0$, N_i and D_i be as in Lemma 1. Then the set of controllers, C_{i+1} stabilizing P_0 which are guaranteed to stabilize the class of all possible P satisfying this condition is given by

$$C_{i+1} = (N_i - Y_0 Q)(D_i + X_0 Q)^{-1}, \text{ for } Q \in \mathcal{RH}_\infty \text{ with } \|Q\| \le \epsilon^{-1}.$$

Proof: First we search for one such controller, C_{i+1}. Clearly, since C_{i+1} stabilizes P_0, it has the coprime factorization

$$C_{i+1} = (N_i - Y_0 Q)(D_i + X_0 Q)^{-1} = (\tilde{D}_i + Q\tilde{X}_0)^{-1}(\tilde{N}_i - Q\tilde{Y}_0),$$

for some $Q \in \mathcal{RH}_\infty$. Stability of the (P, C_{i+1}) feedback interconnection is determined by whether the following rational function is a unit:

$$\begin{aligned}
&(\tilde{D}_i + Q\tilde{X}_0)(Y_0 + N_i R) + (\tilde{N}_i - Q\tilde{Y}_0)(X_0 - D_i R) \\
&= (\tilde{D}_i Y_0 + \tilde{N}_i X_0) + Q(\tilde{X}_0 Y_0 - \tilde{Y}_0 X_0) + (\tilde{D}_i N_i - \tilde{N}_i D_i)R + Q(\tilde{X}_0 N_i + \tilde{Y}_0 D_i)R \\
&= I + QR. \qquad (25)
\end{aligned}$$

In order that this be stably invertible for all $R \in \mathcal{RH}_\infty$ with $\|R\|_\infty < \epsilon$ it is evidently necessary and sufficient that $\|Q\|_\infty \le \epsilon^{-1}$. $\nabla\nabla\nabla$

Problem 3 — Moving between stabilizing controllers

The‘ addition of a metric, δ_ν, onto the set of transfer functions means that we can seek to answer questions regarding the properties of this metric space. In particular, if we consider the equivalence class of controllers which stabilize a given plant, we might ask whether it is possible to move continuously between two members of this class without leaving the class, i.e. whether this set is connected. A positive answer would legitimize a strategy of moving between different controllers associated with two sets of experimental conditions, without losing stability en route.

A more refined version of the question asks whether one can move by a series of small steps from an initial controller to a final controller. Why is this an important question? Suppose P is an unknown plant, and we have an initially stabilizing $\underline{C}$. Suppose that if we actually knew P, we would use a controller $\bar{C}$ to achieve a performance specification. Then it is pertinent to ask whether *in principle* a series of identification and controller redesign steps could get us from $\underline{C}$ to $\bar{C}$, with the proviso that at any one step, the change of controller had to be sufficiently small so that stability remained guaranteed. In answering this question, we shall see again that the Youla-Kucera parameter and its connection to the ν-metric via (15) of Corollary 5 provide a central tool. The question is formally stated as follows.

> *Suppose we have a known plant P with two known stabilizing controllers $\underline{C}$ and $\bar{C}$ and known stability margins $b_{P,\underline{C}}$ and $b_{P,\bar{C}}$. Can we move from initial controller $\underline{C}$ to a final controller $\bar{C}$ by a sequence of cautious steps? That is, can we develop a sequence $\{C_0 = \underline{C}, C_1, \ldots, C_n = \bar{C}\}$ of controllers, all stabilizing P, for which $\delta_\nu(C_{i+1}, C_i) < b_{P,C_i}$ for $i = 1, \ldots, n$.*

We first establish that a solution exists to this problem and then show how to construct a solution. Since $(P, \underline{C})$ and $(P, \bar{C})$ are internally stable, we have from Theorem 2 that we may write normalized coprime factorizations $P = XY^{-1} = \tilde{Y}^{-1}\tilde{X}$, and coprime factorizations $\underline{C} = N_0 D_0^{-1} = \tilde{D}_0^{-1}\tilde{N}_0$ and $\bar{C} = (N_0 - YQ_n)(D_0 + XQ_n)^{-1} = (\tilde{D}_0 + Q_n\tilde{X})^{-1}(\tilde{N}_0 - Q_n\tilde{Y})$ with $\tilde{D}_0 Y + \tilde{N}_0 X = I$. Here $Q_n \in \mathcal{RH}_\infty$ is the Youla-Kucera parameter of $\bar{C}$. Further, since

$$T(P, \underline{C}) - T(P, \bar{C}) = \begin{pmatrix} X \\ Y \end{pmatrix} Q_n \begin{pmatrix} -\tilde{Y} & \tilde{X} \end{pmatrix},$$

and, since (X, Y) and $(\tilde{X}, \tilde{Y})$ are normalized, we may write

$$\|T(P, \underline{C}) - T(P, \bar{C})\|_\infty = \|Q_n\|_\infty. \tag{26}$$

From (7) we thus have $\delta_\nu(\underline{C}, \bar{C}) \le \|Q_n\|_\infty$. If we choose as our sequence of controllers $C_i = (N_0 - YQ_i)(D_0 + XQ_i)^{-1}$ where $Q_i = \frac{i}{n}Q_n$ then we have each C_i stabilizing P and

$$\delta_\nu(C_{i+1}, C_i) \le \frac{1}{n}\|Q_n\|_\infty. \tag{27}$$

The existence of a stabilizing sequence is therefore settled, although its computation relies on knowledge of the plant P via its normalized coprime factors.

This sequence of controllers, $\{C_i\}$, has the following property

$$\|T(P, C_i)\|_\infty \le \max\left(\|T(P, \underline{C})\|_\infty, \left\|T(P, \bar{C})\right\|_\infty\right). \tag{28}$$

This follows by writing

$$\begin{aligned} \|T(P, C_i)\|_\infty &= \left\| \begin{pmatrix} X \\ Y \end{pmatrix} \begin{pmatrix} \tilde{N}_0 & \tilde{D}_0 \end{pmatrix} + \frac{i}{n} \begin{pmatrix} X \\ Y \end{pmatrix} Q_n \begin{pmatrix} -\tilde{Y}_0 & \tilde{X}_0 \end{pmatrix} \right\|_\infty \\ &= \left\| \frac{n-i}{n} \begin{pmatrix} X \\ Y \end{pmatrix} \begin{pmatrix} \tilde{N}_0 & \tilde{D}_0 \end{pmatrix} + \frac{i}{n} \begin{pmatrix} X \\ Y \end{pmatrix} \begin{pmatrix} \tilde{N}_0 - Q_n\tilde{Y} & \tilde{D}_0 + Q_n\tilde{X} \end{pmatrix} \right\|_\infty \\ &= \left\| \frac{n-i}{n} T(P, \underline{C}) + \frac{i}{n} T(P, \bar{C}) \right\|_\infty \\ &\le \frac{n-i}{n}\|T(P, \underline{C})\|_\infty + \frac{i}{n}\|T(P, \bar{C})\|_\infty \\ &\le \max\left(\|T(P, \underline{C})\|_\infty, \left\|T(P, \bar{C})\right\|_\infty\right). \end{aligned} \tag{29}$$

Similarly, from (26) we have that

$$\|Q_n\|_\infty = \|T(P,\underline{C}) - T(P,\bar{C})\|_\infty \le \|T(P,\underline{C})\|_\infty + \|T(P,\bar{C})\|_\infty. \tag{30}$$

Suppose, without loss of generality and for the moment, that $b_{P,\underline{C}} \le b_{P,\bar{C}}$. By (29), $b_{P,\underline{C}} \le b_{P,C_i} \le b_{P,\bar{C}}$. If we demand guaranteed safe steps between $\underline{C}$ and $\bar{C}$, then we simply require that $\delta_\nu(C_{i+1}, C_i) \le b_{P,\underline{C}}$ at each step. This may be achieved by demanding $\frac{1}{n}\|Q_n\|_\infty \le b_{P,\underline{C}}$ or $n \ge \|Q_n\|_\infty \|T(P,\underline{C})\|_\infty$ from (27). The argument is trivially different in case $b_{P,\bar{C}} \le b_{P,\underline{C}}$. Appealing to (30), we see that we will achieve safe transition from $\underline{C}$ to $\bar{C}$ provided we select

$$n \ge \max\left(\|T(P,\underline{C})\|_\infty, \left\|T(P,\bar{C})\right\|_\infty\right)\left(\|T(P,\underline{C})\|_\infty + \|T(P,\bar{C})\|_\infty\right).$$

Designs where $\|T(P,C)\|_\infty$ is large are non-robust; to work properly, very accurate models of the plant and implementation of the controller are required. This result shows that the number of steps needed in a safe iterative identification and control algorithm to yield a particular design grows as the square of the measure of design difficulty. The equivalent result holds in describing the number of stages needed in direct controller tuning with caution. This will be further illustrated in a following example.

We note that the above computations have required the use of a known normalized coprime factorization of the plant P to generate the sequence of new controllers. In an adaptive control setting, again a sequence of controllers is generated for a fixed P but this true plant is unknown. These above methods for cautious adjustment of the controller can, however, still be applied using estimates of stability margins and estimated normalized coprime factors of the plant, as described in Subsection 3.1.2.

Kammer and others [Kam98, KBB2K] develop methods to generate non-parametric plant models from closed-loop experimental data using frequency response samples from Fourier analysis of signals. In this way they are able to generate non-parametric descriptions of the plant and even of its coprime factors via sensitivity functions. Thus we may write

$$P = P(I+CP)^{-1} \times \left((I+CP)^{-1}\right)^{-1},$$

so that knowledge of $T(P,C)$ may be interpreted as knowledge of the plant and its coprime factors. In the next section we shall consider the application of the ν-metric in cases where approximate models are provided by successive identification tasks.

4.1 Controller Adjustment Example

Here we develop a computer example to illustrate some features of cautious controller tuning. The objective chosen is to minimize $\|T(P,C)\|_\infty$, or equivalently to maximize the stability margin. The example treats an unknown plant and two sequences of controllers, one of which converges to a local minimum of this objective and the other of which finds the global minimum. We consider controllers parametrized by two parameters, $(\rho_1;\rho_2)$, and presume that the plant is unknown, but that the stability margins can be accurately estimated using the methods of the Section 3.

Consider the following open-loop unstable process:

$$P(z) = \frac{0.1\, z^2\, (z - 0.3)}{(z - 1.08)(z - 0.2 + 0.9i)(z - 0.2 - 0.9i)(z - 0.6 + 0.6i)(z - 0.6 - 0.6i)} \quad (31)$$

$$= \frac{0.1\, z^3 - 0.03\, z^2}{z^5 - 2.68\, z^4 + 3.778\, z^3 - 3.522\, z^2 + 2.025\, z - 0.661}. \quad (32)$$

This process can be stabilized by simple controllers under the structure

$$C(z) = \frac{\rho_1\, z}{z + \rho_2}, \quad (33)$$

where ρ_1 and ρ_2 are the tuning parameters.

The controller starts with a set of stabilizing parameters, $\rho = (\rho_1; \rho_2)$, and ρ is to be safely tuned towards the point of minimum $\|T(P, C_i)\|_\infty$. At each iteration we measure the current stability margin, b_{P,C_i}, and find the complete set of controller parameters satisfying Vinnicombe's criterion

$$\delta_\nu(C_i, C_{i+1}) < b_{P,C_i}. \quad (34)$$

Within this set we select the one which minimizes $\|T(P, C_{i+1})\|_\infty$ using, say, gradient and Hessian information together with the distance constraint of (34).

Figure 1 shows the evolution of the tuning procedure for an initial controller $\rho^{\mathbf{a}} = (1; -0.9)$, which has $b_{P,C_0} = 0.092$. The image displays the following information:

- The axes measure controller parameters ρ_1 and ρ_2.
- For the above plant the boundary of the stabilizing region is depicted by a dotted line.
- The controller parameters which globally minimize $\|T\|_\infty$ are indicated by a circle.
- The initial controller, C_0, parameters are marked with an 'x'.
- The area bounded by a solid line is the complete set of parameters for controllers C_1 satisfying $\delta_\nu(C_0, C_1) \leq b_{P,C_0} = 0.092$.
- The area bounded by the dashed line is the set of controllers satisfying the frequency-by-frequency bound as in Corollary 1.
- Within the class of controller parameters satisfying $\delta_\nu(C_0, C_1) \leq b_{P,C_0}$, a diamond marks that which minimizes $\|T\|_\infty$.

The initial stability margin is $b^a_{P,C} = 0.092$ which limits the safe excursion of the controller parameters to a small region around the initial point, delimited by the solid line. The frequency-by-frequency test is not much less conservative. The point which minimizes the criterion $\|T\|_\infty$ is the initial controller for the next iteration, as presented by Figure 1(b).

The tuning procedure continues until the parameters of the controller reach the point of global maximum of the generalized stability margin, where $\rho = (0.579; -0.680)$ and $b_{P,C} = 0.343$. This is shown in Figure 1(c).

Figure 2 illustrates a similar sequence of controller adjustments commencing from initial condition $\rho = (3.6; 0.8)$. Here we see that because the objective function $\|T\|$ is not convex over the set of controller parameters, tuning

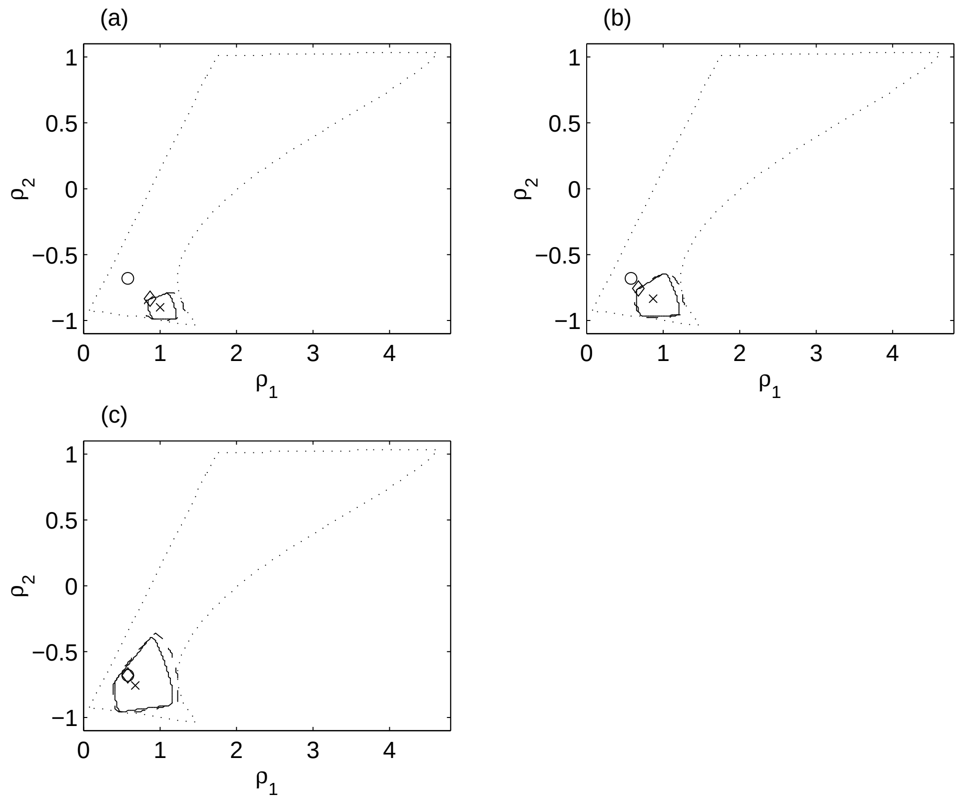

Figure 1: Iterative controller tuning reaching the global maximum of $b_{P,C}$: (a) initial step, (b) second step, (c) third and final step.

can occur to a local minimum instead of the global minimum. A contour plot of the objective function $\|T\| = b_{P,C}^{-1}$ is displayed in Figure 3 and shows the existence of two extremal points for this system.

The conservatism of the ν-metric guarantees is indicated by the small size of the regions enclosed by both the solid and dashed curves within the dotted stability boundaries in these examples. For points near the boundary of the stability region, the guarantees are very limiting. However, as the controller moves to more internal points, the limits relax. This is in accord with the normal expectation that controllers near to a stability boundary need to be adjusted very cautiously.

We note also that the connectivity and convexity properties of controllers suggested by the solution to Problem 3 of this section no longer holds when the controllers are parametrized in the fashion of (33), rather than being parametrized by Youla-Kucera parameter.

This computer example demonstrates that controller caution can be implemented to regulate the evolution of a controller from its initial position towards a newly tuned value. There are no guarantees that the tuning direction is good - that is not our subject here - but the cautious adjustment will maintain closed-loop stability. If the tuning takes the closed loop away from a stability boundary then we expect to see the cautious restrictions relaxed, as happens in the example.

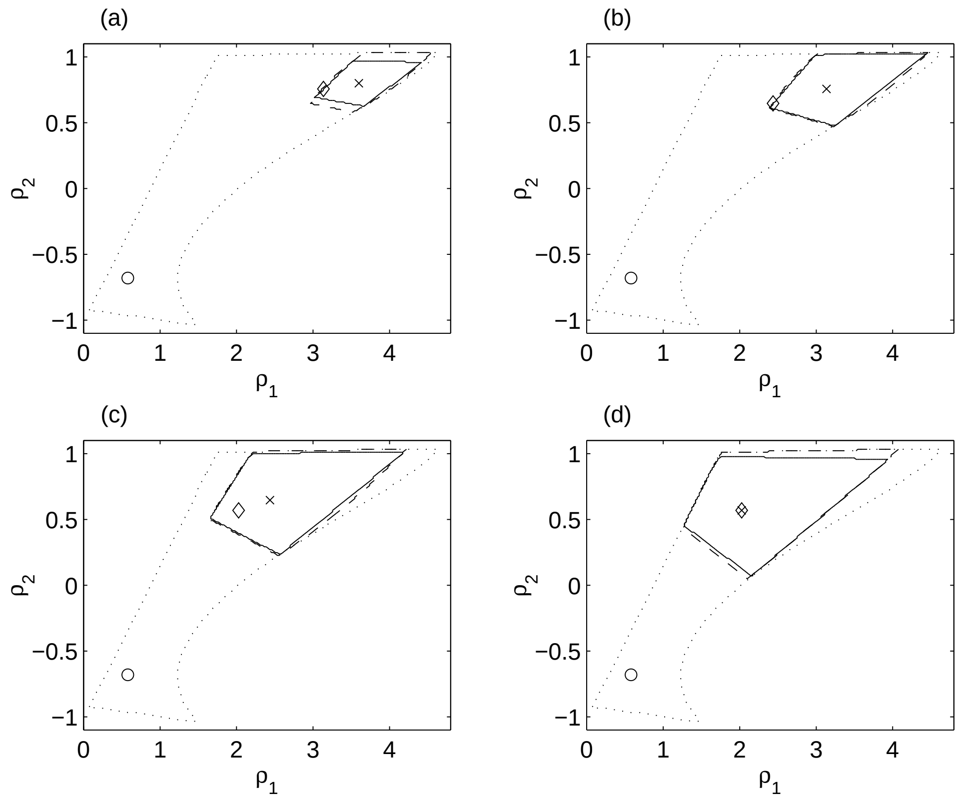

Figure 2: Iterative controller tuning reaching a local maximum of $b_{P,C}$ over four steps (a)–(d).

5 Joint Model and Controller Tuning

Iterative identification and control design is a descendant of adaptive control in which modelling stages are interspersed with control design. Identification stages proceed using blocks of closed-loop data from an experiment operating with the current controller and yield a new plant model. The control design phase uses this model to generate a new controller to be used in the closed loop. Because both models and controllers change during iterative identification and control design, while the true plant remains fixed, the opportunity arises to apply the ν-metric to the cautious adjustment of both phases. We shall actually see that the tunings of model and controller are inter-related, as is foreshadowed in Problem 2 of Section 4.

5.1 Simultaneous Stabilization

Assume that we have an unknown true plant, P, a sequence of feedback controllers, $\{C_i\}$, and a sequence of plant models, $\{P_i\}$, with C_i designed to stabilize P_i. Without loss of generality, we presume that the two sequences are the same length. That is, at each iteration a new model is fitted and a new controller designed. We further assume that the basis of model fitting is that the new model, P_{i+1}, is selected to belong to the class of plants stabilized by the current controller, C_i. Observe that in the case of a controller C_i having unstable poles or non-minimum-phase zeros, this may require that some specific closed loop identification methods be ruled out, as has been shown in [CAG2K]. The sequences of plant models and controllers satisfy the following

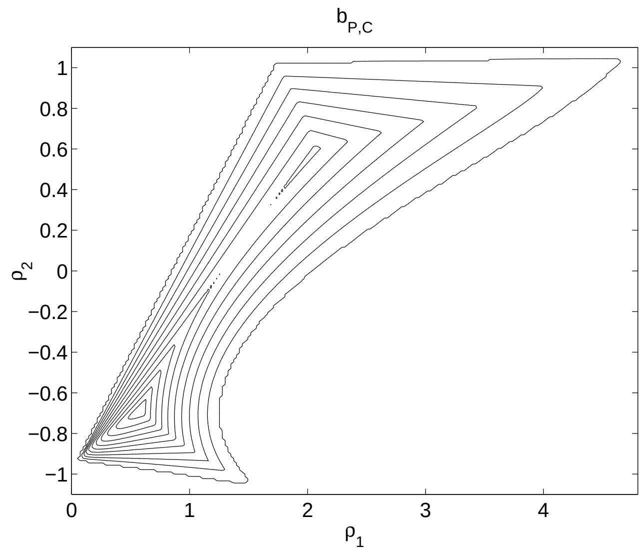

Figure 3: Contour plot of $b_{P,C}$

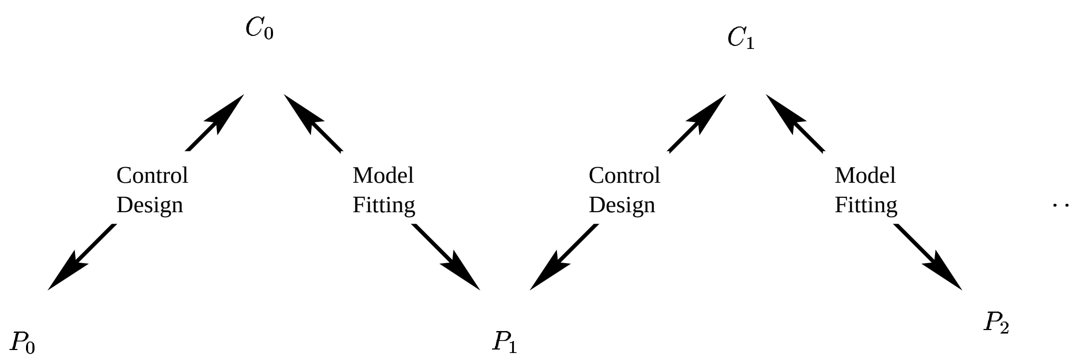

Figure 4: Simultaneous stabilization diagram for Joint Modelling and Control Design. Directed arrows indicate stabilization.

property (depicted in Figure 4).

> [**Simultaneous stabilization property**] *Each plant model, P_{i+1}, is stabilized by controllers C_i and C_{i+1}. Each controller, C_i, stabilizes plant models P_i and P_{i+1}.*

That model P_{i+1} is stabilized by C_{i+1} is a reflection of the property that model-based control design achieves stabilization of the model. The assumption that P_{i+1} is stabilized by C_i is simply checkable, since both are known to the designer. It may also be incorporated directly into the admissible class of models and the chosen identification method, as we shall elaborate in the next subsection. This is particularly relevant in the case of controllers with unstable poles or non-minimum-phase zeros where some of the standard methods are guaranteed to give a model not stabilized by the current controller [CAG2K]. We shall discuss shortly and briefly identification methods which estimate directly the Youla-Kucera parameter of a new model, P_{i+1}, relative to the current model, P_i, and the simultaneously stabilizing controller, C_i.

Following from the above simultaneous stabilization, we have an immediate consequence of Theorem 2, since this theorem describes the class of all stabilizing controllers or all stabilized plant models.

Lemma 2 *Consider plant models $\{P_i\}$ and controllers $\{C_i\}$ satisfying the simultaneous stabilization property. Denote their coprime factorizations*

$$P_i = X_i Y_i^{-1} = \tilde{Y}_i^{-1} \tilde{X}_i, \quad C_i = N_i D_i^{-1} = \tilde{D}_i^{-1} \tilde{N}_i, \tag{35}$$

where $\tilde{D}_i Y_i + \tilde{N}_i X_i = I_m$ and $\tilde{Y}_i D_i + \tilde{X}_i N_i = I_p$. Then we may write the next plant model as

$$P_{i+1} = X_{i+1} Y_{i+1}^{-1} \triangleq (X_i - D_i R_i)(Y_i + N_i R_i)^{-1} = (\tilde{Y}_i + R_i \tilde{N}_i)^{-1} (\tilde{X}_i - R_i \tilde{D}_i) \triangleq \tilde{Y}_{i+1}^{-1} \tilde{X}_{i+1}, \tag{36}$$

for $R_i \in \mathcal{H}_\infty^{p\times m}$. We may write the next controller as

$$C_{i+1} = N_{i+1} D_{i+1}^{-1} \triangleq (N_i - Y_{i+1} Q_i)(D_i + X_{i+1} Q_i)^{-1} = (\tilde{D}_i + Q_i \tilde{X}_{i+1})^{-1} (\tilde{N}_i - Q_i \tilde{Y}_{i+1}) \triangleq \tilde{D}_{i+1}^{-1} \tilde{N}_{i+1}, \tag{37}$$

for $Q_i \in \mathcal{H}_\infty^{m\times p}$.

Lemma 2 describes the class of all possible plant models and controllers satisfying the simultaneous stabilization property. The sequences of plant models $\{P_i\}$ and controllers $\{C_i\}$ satisfy the simultaneous stabilization property if and only if their left coprime factorizations as in (35) satisfy,

$$\begin{pmatrix} X_{i+1} & -Di \\ Y_{i+1} & N_i \end{pmatrix} = \begin{pmatrix} X_i & -D_i \\ Y_i & N_i \end{pmatrix} \begin{pmatrix} I_m & 0_{m\times p} \\ R_i & I_p \end{pmatrix}, \tag{38}$$

$$\begin{aligned} \begin{pmatrix} X_{i+1} & -D_{i+1} \\ Y_{i+1} & N_{i+1} \end{pmatrix} &= \begin{pmatrix} X_{i+1} & -D_i \\ Y_{i+1} & N_i \end{pmatrix} \begin{pmatrix} I_m & -Q_i \\ 0_{p\times m} & I_p \end{pmatrix} & (39) \\ &= \begin{pmatrix} X_i & -D_i \\ Y_i & N_i \end{pmatrix} \begin{pmatrix} I_m & 0_{m\times p} \\ R_i & I_p \end{pmatrix} \begin{pmatrix} I_m & -Q_i \\ 0_{p\times m} & I_p \end{pmatrix} & (40) \\ &= \begin{pmatrix} X_i & -D_i \\ Y_i & N_i \end{pmatrix} \begin{pmatrix} I_m & -Q_i \\ R_i & I_p - R_i Q_i \end{pmatrix}. & (41) \end{aligned}$$

Equivalently, we may write these conditions in terms of right coprime factorizations:

$$\begin{aligned} \begin{pmatrix} \tilde{N}_{i+1} & \tilde{D}_{i+1} \\ -\tilde{Y}_{i+1} & \tilde{X}_{i+1} \end{pmatrix} &= \begin{pmatrix} I_m & Q_i \\ 0_{p\times m} & I_p \end{pmatrix} \begin{pmatrix} I_m & 0_{m\times p} \\ -R_i & I_m \end{pmatrix} \begin{pmatrix} \tilde{N}_i & \tilde{D}_i \\ -\tilde{Y}_i & \tilde{X}_i \end{pmatrix} & (42) \\ &= \begin{pmatrix} I_m - Q_i R_i & Q_i \\ -R_i & I_p \end{pmatrix} \begin{pmatrix} \tilde{N}_i & \tilde{D}_i \\ -\tilde{Y}_i & \tilde{X}_i \end{pmatrix} & (43) \end{aligned}$$

We note that the latter result of (42) may be established from the property that

$$\begin{pmatrix} \tilde{X}_i & -\tilde{Y}_i \\ \tilde{D}_i & \tilde{N}_i \end{pmatrix} \begin{pmatrix} N_i & Y_i \\ -D_i & X_i \end{pmatrix} = \begin{pmatrix} I_p & 0_{p\times m} \\ 0_{m\times p} & I_m \end{pmatrix}.$$

Indeed, stability of the closed loop is implied by these matrices of coprime factors above being invertible in $\mathcal{RH}_\infty$.

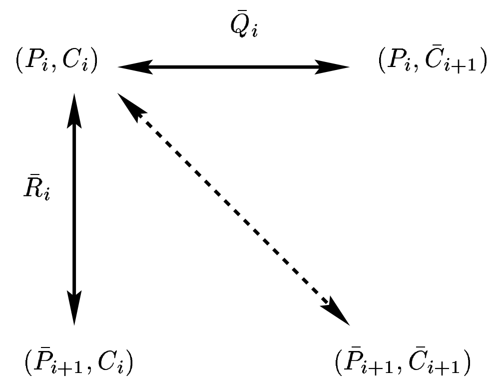

Figure 5: Commuting diagram illustrating the Tay, Moore and Horowitz problem, *cf* Figure 4. The distinction with the Simultaneous Stabilization Property is that C_{i+1} is computed from C_i and P_{i+1} while $\bar{C}_{i+1}$ is computed from C_i and P_i.

The result of Tay, Moore and Horowitz

If, instead of asking that the model sequence $\{P_i\}$ and controller sequence $\{C_i\}$ satisfy the simultaneous stabilization property, we recast the same question with regard to the actual plant, then we arrive at a variant of the result of Tay, Moore and Horowitz [TMH89].

The simultaneous stabilization property deals with plant models stabilizing pairs of controllers and vice versa. Thus pairs (P_i, C_i), (P_{i+1}, C_i), (P_{i+1}, C_{i+1}) are assured stable with interlaced adjustment of P_i and C_i. The problem considered in [TMH89] is,

> *Given that the model P_i has been adjusted using Youla-Kucera parameter $\bar{R}_i$ to $\bar{P}_{i+1}$ so that (P_i, C_i) and $(\bar{P}_{i+1}, C_i)$ are stable, and given that the controller C_i is adjusted using Youla-Kucera parameter $\bar{Q}_i$ to $\bar{C}_{i+1}$ so that (P_i, C_i) and $(P_i, \bar{C}_{i+1})$ are stable, what are the conditions that ensure that $(\bar{P}_{i+1}, \bar{C}_{i+1})$ is stable?*

The elegant solution is that stability of $(\bar{P}_{i+1}, \bar{C}_{i+1})$ is assured if and only if $I + \bar{Q}_i\bar{R}_i$ is a unit, i.e. if $\bar{Q}_i$ stabilizes $\bar{R}_i$ in a feedback loop. This is depicted in Figure 5 by the dashed line representing the stabilization of the altered model, altered controller pair $(\bar{P}_{i+1}, \bar{C}_{i+1})$.

This condition is easily derived by considering the (P_{i+1}, C_{i+1}) closed-loop denominator.

$$\begin{aligned}
\tilde{D}_{i+1}Y_{i+1} + \tilde{N}_{i+1}X_{i+1} &= (\tilde{N}_i - \bar{Q}_i\tilde{Y}_i)(X_i - D_i\bar{R}_i) + (\tilde{D}_i + \bar{Q}_i\tilde{X}_i)(Y_i + N_i\bar{R}_i) \\
&= I + \bar{Q}_i\bar{R}_i.
\end{aligned}$$

Since $\bar{Q}_i$ and $\bar{R}_i$ are stable by construction, $I + \bar{Q}_i\bar{R}_i$ represents the denominator of the $(\bar{Q}_i, \bar{R}_i)$ closed loop.

The simultaneous stabilization property differs from the [TMH89] set-up because of the interleaving of modeling and control design, i.e. the controller C_{i+1} is a Youla-Kucera variation on C_i using P_{i+1}, not P_i. If we denote the linear fractional transformations of [TMH89] above by $\mathcal{L}$, then we have

$$\bar{P}_{i+1} = \mathcal{L}(P_i, C_i, \bar{R}_i), \quad \bar{C}_{i+1} = \mathcal{L}(C_i, P_i, \bar{Q}_i),$$

while our formulation has

$$P_{i+1} = \mathcal{L}(P_i, C_i, R_i), \quad C_{i+1} = \mathcal{L}(C_i, P_{i+1}, Q_i),$$

and the simultaneous stability property is immediate whatever Q_i and R_i are chosen in $\mathcal{RH}_\infty$.

The related question to [TMH89] would be to ask whether C_{i+1} stabilizes P_i. We note without proof that this follows provided $I - R_iQ_i$ is a unit[4] As discussed in the proof of Theorem 3, one immediate way to ensure that $I - R_iQ_i$ is a unit is to ensure that $\|Q_i\|_\infty \|R_i\|_\infty < 1$, which may be interpreted as a joint restriction on the available movement in model and controller.

Rather than considering the stabilization of the sequence of plant *models* by the controllers (which is assured by the Q_i, R_i double Youla-Kucera construction), a more significant question is to ask that the controllers C_{i+1} stabilize the *real* plant P. Then a variation of [TMH89] provides the desired tool.

Theorem 4 *Suppose we have the following coprime factorizations:*

- *the real plant* $P = XY^{-1} = \tilde{Y}^{-1}\tilde{X}$,
- *model* $P_i = X_iY_i^{-1} = \tilde{Y}_i^{-1}\tilde{X}_i$ *and stabilizing controller* $C_i = N_iD_i = \tilde{D}_i^{-1}\tilde{N}_i$ *satisfying* $\tilde{D}_iY_i + \tilde{N}_iX_i = I_m$ *and* $\tilde{Y}_iD_i + \tilde{X}_iN_i = I_p$.

Define the next model P_{i+1} *and controller* C_{i+1} *via the Simultaneous Stabilization Property and Youla-Kucera parameters* R_i *and* Q_i *as in (38—43).*

If controller C_i *stabilizes the real plant* P *in addition to both* P_i *and* P_{i+1}*, then* P_i *and* P_{i+1} *may be written as*

$$\begin{aligned} P_i &= (X - D_i\bar{R}_i)(Y + N_i\bar{R}_i)^{-1} = (\tilde{Y} + \bar{R}_i\tilde{N}_i)^{-1}(\tilde{X} - \bar{R}_i\tilde{D}_i)^{-1}, \\ P_{i+1} &= (X - D_i\bar{R}_{i+1})(Y + N_i\bar{R}_{i+1})^{-1} = (\tilde{Y} + \bar{R}_{i+1}\tilde{N}_{i+1})^{-1}(\tilde{X} - \bar{R}_{i+1}\tilde{D}_i)^{-1}, \end{aligned}$$

where evidently $\bar{R}_{i+1} = \bar{R}_i + R_i$.

Controller C_{i+1} *will also stabilize the real plant* P *if and only if* $I - \bar{R}_{i+1}Q_i$ *is a unit in* $\mathcal{RH}_\infty$.

Proof: The stability of the (P, C_{i+1}) closed loop follows if and only if the matrix $\begin{pmatrix} X & -D_{i+1} \\ Y & N_{i+1} \end{pmatrix}$ is a unit in $\mathcal{RH}_\infty$. We have

$$\begin{aligned} \begin{pmatrix} X_{i+1} & -D_{i+1} \\ Y_{i+1} & N_{i+1} \end{pmatrix} &= \begin{pmatrix} X_i & -D_i \\ Y_i & N_i \end{pmatrix} \begin{pmatrix} I & 0 \\ R_i & I \end{pmatrix} \begin{pmatrix} I & -Q_i \\ 0 & I \end{pmatrix}, \\ &= \begin{pmatrix} X - D_i\bar{R}_i & -D_i \\ Y + N_i\bar{R}_i & N_i \end{pmatrix} \begin{pmatrix} I & 0 \\ R_i & I \end{pmatrix} \begin{pmatrix} I & -Q_i \\ 0 & I \end{pmatrix}, \\ &= \begin{pmatrix} X & -D_i \\ Y & N_i \end{pmatrix} \begin{pmatrix} I & 0 \\ \bar{R}_i + R_i & I \end{pmatrix} \begin{pmatrix} I & -Q_i \\ 0 & I \end{pmatrix}, \\ &= \begin{pmatrix} X & -D_i \\ Y & N_i \end{pmatrix} \begin{pmatrix} I & 0 \\ \bar{R}_{i+1} & I \end{pmatrix} \begin{pmatrix} I & -Q_i \\ 0 & I \end{pmatrix}. \end{aligned}$$

[4] The change of sign here is due to the changing of elements in the commuting diagram Figure 5. (P_{i+1}, C_i) is the top left, (P_{i+1}, C_{i+1}) is the top right controller variation with parameter Q_i. (P_i, C_i) is the bottom left plant variation with parameter $-R_i$. (P_i, C_{i+1}) is the lower right pair whose stabilization is under investigation.

From the simultaneous stabilization property, we have that the left-hand side above is a unit. Similarly, from the stability of (P, C_i) we see that each of the matrices on the right is also a unit.

Now construct the test matrix for stability of (P, C_{i+1}).

$$\begin{pmatrix} X & -D_{i+1} \\ Y & N_{i+1} \end{pmatrix} = \begin{pmatrix} X_{i+1} + D_i \bar{R}_{i+1} & -D_i - X_{i+1} Q_i \\ Y_{i+1} - N_i \bar{R}_{i+1} & N_i - Y_{i+1} Q_i \end{pmatrix}$$
$$= \begin{pmatrix} X_{i+1} & -D_i \\ Y_{i+1} & N_i \end{pmatrix} \begin{pmatrix} I & -Q_i \\ -\bar{R}_{i+1} & I \end{pmatrix}$$
$$= \begin{pmatrix} X_i & -D_i \\ Y_i & N_i \end{pmatrix} \begin{pmatrix} I & 0 \\ -\bar{R}_{i+1} & I \end{pmatrix}, \begin{pmatrix} I & -Q_i \\ 0 & I - \bar{R}_{i+1} Q_i \end{pmatrix}.$$

The the left-hand side of the above equation will be a unit if and only if $I - \bar{R}_{i+1} Q_i$ is a unit. ∇∇∇

Note the requirement that $I - Q_i \bar{R}_{i+1}$ be a unit imposes a further condition on the allowed model and plant variations. A good model of the plant would have $\|\bar{R}_{i+1}\|_\infty$ small and therefore admit larger $\|Q_i\|$ alteration before running into a restriction due to stabilization guarantee if no other information is available about $\bar{R}_{i+1}$. This is entirely analogous to the circumstance described in Theorem 3.

Obtaining models stabilized by the current controller

Since system identification is being introduced, we now consider briefly one approach to ensuring stabilization of the next model by the current controller.

An identification method to move in the class of models stabilized by the current controller is to express the new model P_{i+1} as a perturbation of P_i, which is stabilized by C_i, using the Youla-Kucera parameter. Such methods have been investigated in [HFK89, Schr91, LAKM93] and use the previous model and controller to yield signal filters for direct estimation of the true plant Youla-Kucera parameter, $-\bar{R}_i$. Note, however, that the class of admissible R is $\mathcal{RH}_\infty$, so that R must be stable and proper.

These Youla-Kucera factor identification schemes guarantee that C_i stabilizes P_{i+1}, since these latter models are restricted to the linear fractional transformations which generate the set of all plants stabilized by C_i. The technical issue is then to ensure that the search for Youla-Kucera parameter, R, is limited among stable, proper transfer functions. This can be achieved by judicious selection of model structure, for example, by using finite-impulse-response (moving average) R [TMM97] or by using output error or Box-Jenkins methods. Other direct methods for estimating normalized coprime factors of P from input-output data are considered in [VSDB95, Gu99].

If the controller (or more likely its implementation) is not exactly known to the designer, then we may achieve stabilization of the model by keeping the ν-metric of model variation, $\delta_\nu(P_{i+1}, P_i)$, less than b_{P_i,C_i}, subject to the satisfaction of the winding number condition $\mathcal{C}$ on the two successive models. To achieve this, we may appeal to the duals of the ideas of Problems 1 and 2 in Section 4.

For our purposes here, the major aim is not to generate coprime factor models of the plant but to estimate margins. However, the availability of (not necessarily normalized) coprime factor plant models could well provide a distinct advantage in moving among the class of model-stabilizing controllers using Youla-Kucera parameters.

We note, in summary, that the results of Theorem 3 help to delineate an approach to joint modelling and control design which is capable of guaranteeing the simultaneous stabilization property of the $\{P_i, C_i\}$ and the $\{P, C_i\}$ sequences.

5.2 Cautious Iterative Modeling and Control

As indicated in earlier sections, iterative modeling and control design proceeds by successive steps of model and controller update. Equally, we have seen from the earlier sections that the model fitting phase is necessarily conducted to select a model within the class of P_{i+1} stabilized by the current controller C_i, which stabilizes the true plant P. (We shall return to consider how to ensure this shortly.) In the control design phase, we have portrayed above Youla-Kucera methods to guarantee that C_{i+1} stabilizes P_{i+1}. The fundamental question must be then to ensure that this or a derived C_{i+1} also stabilizes P. For this we shall need to use the techniques described in Problem 2 earlier and Theorem 4 above. The upshot of this analysis is that the cautious phase of these iterative schemes needs to rest with the control update. The requirements of the model step is to capture the closed-loop dynamics well.

An appropriate measure of identification error

We have dealt with the stabilization of the model by the controller. We now look at an appropriate definition of the identification error. The identification objective should be to provide a model, P_{i+1}, first with the properties that;

1. the performances of the feedback loops (P_{i+1}, C_i) and (P, C_i) are close, and
2. there is maximum scope for designing a controller C_{i+1} which stabilizes P_{i+1} and also stabilizes the true plant P.

According to Theorem 3, if we choose P_{i+1} to keep $\|T(P, C_i) - T(P_{i+1}, C_i)\|_\infty = \epsilon$ small, then we correspondingly increase the size of the set of admissible C_{i+1} which will stabilize P. We note that the identification objective to minimize the ∞-norm of the error between the modeled and actual generalized sensitivity functions attempts to capture the closed-loop dynamics of (P, C_i) with the model. This is a closed-loop performance matching model. Note also that Corollary 2 ensures that $|b_{P_{i+1},C_i} - b_{P,C_i}| \leq \delta_\nu(P_{i+1}, P) \leq \epsilon$, so that for small ϵ the model and the plant will be ν-close and exhibit similar stability margins with C_i.

How might this identification objective be achieved? We have already canvassed in Section 3 some existing methods for ∞-norm identification. These might be used here, although it is not clear how to parametrize these solutions to yield efficiently the P_{i+1} component. We note also that ∞-norm model fitting is profligate in its use of data. An alternative approach is to use more familiar least squared prediction error methods to minimize $\|T(P, C_i) - T(P_{i+1}, C_i)\|_2$ and then to use the methods of Section 3 to compute the value of the ensuing ∞-norm, ϵ. Such a methodology for fitting P_{i+1} has the advantage of available algorithms and well-understood behavior.

Notice that we may write (after some algebra),

$$T(P,C_i) - T(P_{i+1},C_i) = \begin{pmatrix} I \\ -C_i \end{pmatrix} (I + PC_i)^{-1}(P - P_{i+1})(I + C_iP_{i+1})^{-1} \begin{pmatrix} C_i & I \end{pmatrix}. \tag{44}$$

That is, the difference between generalized sensitivity functions may be interpreted as a version of the direct closed-loop identification criterion of [ZBG95] weighted by the outer matrices involving C_i. This, in turn, suggests that a variant of standard system identification methods using closed-loop data be used.

Ensuring the controller stabilizes P

With a model minimizing $\|T(P,C_i)-T(P_{i+1},C_i)\|_\infty$, Theorem 3 constructs the class of guaranteed P-stabilizing controllers C_{i+1} using the experimentally computed stability margin b_{P,C_i}. A new candidate controller, $\hat{C}_{i+1}$ is computed using P_{i+1} and expressed as a variation of C_i with Youla-Kucera parameter Q_{i+1}. Controller C_{i+1} is selected with parameter αQ_{i+1} for $\alpha \in (0,1]$, subject to the stability guarantee.

Since the error between the generalized sensitivity functions is minimized by the identified model, the controller class is maximal as measured by the ∞-norm of the control Youla-Kucera parameter. The incorporation of this limit on the magnitude of the controller update (Youla-Kucera) parameter is Cautious Iterative Control. The flexibility in adjusting the controller is limited by the quality of the model fit to the true plant as measured by the ∞-norm of the difference between generalized sensitivity functions.

6 Conclusion

The ν-metric has proven to be a versatile tool in cautious controller tuning. Its interplay with the generalized stability margin and the Youla-Kucera parametrization leads to powerful, and yet intuitive, mechanisms for maintaining closed-loop stability as an initially stabilizing controller is tuned. These associations contrast to a great extent, as the generalized stability margin can be estimated from closed-loop signals without the need for a plant model, while the Youla-Kucera formulation demands complete information about the plant.

We have drawn together several threads;

- the Youla-Kucera parametrization to describe the class of all simultaneously stabilizing controllers for a known plant,
- Vinnicombe's ν-metric formulation to describe the computable distance between two systems and its relation to simultaneous stabilization via the stability margin,
- computation of the stability margin from experimental data, and
- interaction between modeling and control adjustment,

to yield approaches to the understanding and application of cautious controller adjustment. The objective of this paper has been to provide methods for the use of closed-loop data for tuning a stabilizing controller with an *a priori* stability guarantee. We have been able to do this for model-free controller tuning, controllers based on a fixed nominal model, and iterative modeling and control design. Part of the appeal of these methods is their concordance with practical notions of permissible controller variation and model selection when current closed-loop performance is taken into consideration.

References

[8au86] B.D.O. Anderson, R.R. Bitmead, C.R. Johnson Jr, P.V. Kokotovic, R.L. Kosut, I.M.Y. Mareels, L. Praly and B.D. Riedle, *Stability of Adaptive Systems: Passivity and Averaging Analysis*, M.I.T. Press, Cambridge MA, 1986.

[ABGK98] B.D.O. Anderson, X. Bombois, M. Gevers and C. Kulcsár, "Caution in iterative modeling and control design," IFAC Workshop on Adaptive Control and Signal Processing, Glasgow UK, 1998.

[AG98] B.D.O. Anderson and M. Gevers, "Fundamental problems in adaptive control," in *Perspectives in Control: Theory and Applications*, D. Normand-Cyrot (ed), Springer-Verlag, Berlin, 1998.

[ÅW89] K.J. Åström and B. Wittenmark, *Adaptive Control,* Addison-Wesley, New York, 1995.

[BA81] R.R. Bitmead and B.D.O. Anderson, "Adaptive frequency sampling filters", *I.E.E.E. Transactions on Circuits and Systems*, vol. 28, no. 6, pp. 524-534, 1981.

[BGB97] V. Blondel, M. Gevers and R.R. Bitmead, "When is a model good for control design?" *Proc. 36th IEEE Conference on Decision and Control,* San Diego CA, pp. 1283-1288, 1997.

[BGP97] R.R. Bitmead, M. Gevers and A.G. Partanen, "Introducing caution in iterative controller design," IFAC Symposium on System Identification, Fukuoka Japan, 1997.

[Bom98] X. Bombois, "Insights on Vinnicombe's Metric," private notes, May 1998.

[BYM92] D.S. Bayard, Y. Yam, E. Mettler, "A criterion for joint optimization of identification and robust control", *IEEE Trans Automatic Control*, vol. AC-37, pp. 986-991, 1992.

[CAG2K] B. Codrons, B.D.O. Anderson and M. Gevers, "Closed-loop identification with an unstable or nonminimum phase controller," *Proc. IFAC Symp. Systems Identification (SYSID 2000) CDROM,* Santa Barbara CA, Paper ThPM 1-3, 2000.

[dCVdH97] R.A. De Callafon and P.M.J. Van den Hof, "Suboptimal feedback control by a scheme of iterative identification and control design", *Mathematical Modelling of Systems,* vol. 3, pp.77-101, 1997.

[GK92] G. Gu and P.P. Khargonekar, "A class of algorithms for identification in H_∞," *Automatica*, vol. 28, pp. 299-312, 1992.

[GL95] M. Green and D.J.N. Limebeer, *Linear Robust Control*, Prentice-Hall, Englewood Cliffs, 1995.

[Gu99] G. Gu, "Modeling of normalized coprime factors with ν-metric uncertainty," *IEEE Transactions on Automatic Control,* vol. 44, pp. 1498-1511, 1999.

[HFK89] F.R. Hansen, G.F. Franklin and R.L. Kosut, "Closed-loop identification via the fractional representation: experiment design," *Proc. American Control Conference,* Pittsburgh, 1989.

[HGGL98] H. Hjalmarsson, M. Gevers, S. Gunnarsson and O. Lequin, "Iterative feedback tuning : theory and applications", IEEE Control Systems Magazine, vol. 18, No 4, pp. 26-41, August 1998.

[HJN91] A.J. Helmicki, C.A. Jacobson and C.N.Nett, "Control oriented system identification: a worst-case/deterministic approach in H_∞," *IEEE Transactions on Automatic Control,* vol. 36, pp. 1163-1176.

[Kam98] Leonardo C. Kammer, *Performance monitoring and assessment under an LQ criterion*, PhD Thesis, Australian National University, 1998.

[KBB2K] L.C Kammer, R.R Bitmead, and P.L Bartlett, "Direct Iterative Tuning via Spectral Analysis" *Automatica*, vol. 36, pp. 1301-1307, 2000.

[KL2K] A. Karimi and I.D. Landau, "Robust adaptive control of a flexible transmission system using multiple models'," *IEEE Trans. Control Systems Technology,* vol. 8, pp. 321-331, 2000.

[LAKM93] W.S. Lee, B.D.O. Anderson, R.L. Kosut and I.M.Y. Mareels, "A new approach to adaptive robust control," *International Journal of Adaptive Control and Signal Processing,* vol. 7, pp. 183-211, 1993.

[Ljung99] L. Ljung, *System Identification: theory for the user (2nd Edition),* Prentice-Hall, Englewood Cliffs, 1999.

[MG90] D. McFarlane and K. Glover, *Robust controller design using normalized coprime factor plant descriptions,* Lecture Notes in Control & Information Sciences, Springer-Verlag, New York, 1990.

[MPG95] P.M. Mäkilä, J.R. Partington and T.K. Gustafsson, "Worst-case control-relevant identification," *Automatica*, vol. 31, pp. 1799-1819, 1995.

[NG95] B. Ninness and G.C. Goodwin, "Estimation of model quality," *Automatica*, vol. 31, pp. 1771-1797, 1995.

[SB93] R.J.P. Schrama and O.H. Bosgra, "Adaptive performance enhancement by iterative identification and control design", *International Journal of Adaptive Control and Signal Processing,* vol. 7, pp. 475-487, 1993.

[Schr91] R.J.P. Schrama, "Control-oriented approximate closed-loop identification via fractional representation," *Proc. American Control Conference,* Boston, 1991.

[TMH89] T.T. Tay, J.B. Moore and R. Horowitz, "Indirect adaptive techniques for fixed controller performance enhancement," *International Journal of Control*, vol. 50, pp. 1941-1959, 1989.

[TMM97] T.T. Tay, I. M. Y. Mareels and J. B. Moore, *High Performance Control,* Birkhäuser, Boston, 1997.

[Vin93] G. Vinnicombe, "Frequency domain uncertainty and the Graph Topology," *IEEE Transactions on Automatic Control,* vol. 38, pp. 1371-1383, 1993.

[Vin93a] G. Vinnicombe, *Measuring the robustness of feedback systems,* PhD thesis, Cambridge University, 1993.

[VSDB95] P.M.J. Van den Hof, R.J.P. Schrama, R.A. de Callafon and O.H. Bosgra, "Identification of normalized coprime plant factors from closed-loop data," *European Journal of Control,* vol. 1, pp. 62-74, 1995.

[ZBG95] Z. Zang, R.R. Bitmead, and M. Gevers, "Iterative weighted Least Squares Identification and weighted LQG Control Design", *Automatica*, vol 31, No 11, pp. 1577-1594, November 1995.

[ZD98] K Zhou and J.C. Doyle, *Essentials of Robust Control,* Prentice-Hall, Englewood Cliffs, 1998.