7 Windsurfing Approach to Iterative Control Design

Brian D.O. Anderson

Abstract. Iterative identification and control relies on a sequence of controllers, deemed satisfactory for a sequence of models, being connected to a real plant, and thereby at the very least not introducing an unstable closed loop. Given that the models used for control design are not identical to the real plant, a technique is needed to assure stability of the model-derived controllers with the real plant. This is provided by the windsurfer approach to iterative control design, which seeks to model and control over successively larger bandwidths, with any one bandwidth enlargement step being comparatively modest.

7.1 Introduction

What on earth can the "windsurfing approach" to iterative control design be, and how can windsurfing ever have anything at all to do with automatic control? While the reader will have to wait to read some of the details of the answer, let us record here that iterative control is a technology that enables a controller to learn something about the plant being controlled. This is just what a human must do when he or she learns windsurfing, and if one thinks about how the human acquires the necessary knowledge, this can suggest an approach to iterative control.

There were two streams of ideas that provided the genesis of this work. The first arose when the author, working at the time with R. Kosut on adaptive control problems being addressed by "conventional" adaptive control theory, pondered how a human learnt windsurfing, realising that conventional control theory shed no light on the phenomenon. The natural follow-on question is then whether the human learning process can inspire some new style of adaptive control theory. The second stream arose when the author, often working in collaboration with others including especially M. Gevers and R. Kosut, concluded that much adaptive control theory, being advanced, failed to address some very significant practical issues that could even invalidate or make dangerous application of the theory. The follow-on question of how to address these issues also stimulated many ideas.

Adaptive control is technically not the same as iterative control, or iterative control and identification. Nevertheless, they are first cousins at least. Both are concerned with controlling in the presence of plant uncertainty, and seek to learn something about that uncertainty in order to design an appropriate controller. This is in contrast to, say, the approach of robust control, where one attempts to design round the uncertainty, *i.e.*, to provide a fixed control design that will deliver adequate performance no matter what the values of the uncertain parameters are. The difference between adaptive control on the one hand and iterative control and identification on the other essentially rests in one fact. In adaptive control, processes of identification and controller design are normally fully interleaved, *i.e.*, each new measurement is used to improve the parameter estimates, and at each new time instant adjustment of the controller parameters is permitted. In contrast, in iterative control and identification, one identifies over an interval in which the controller stays fixed, then one adjusts the controller, then one identifies again, and so on. The conceptual difference is not great.

Hence in this chapter, although we are technically describing an iterative control and identification method, we see it within a broader framework, and in particular see it as a valid replacement for some adaptive control methodology, which under some circumstances may be flawed.

Let us sketch an outline of this chapter. In Section 7.2, we set out three conceptual difficulties with conventional adaptive control theory. No theory is perfect, least of all that of this chapter. We do not mean to suggest that all adaptive control theory is flawed, or nearly always flawed. But we do wish to make some very cautionary remarks, and to suggest that the approach of this chapter can sometimes circumvent these difficulties.

In Section 7.3, we try to capture some aspects of what the human is doing when learning to windsurf, and then we set up our underlying assumptions for the iterative control and identification algorithm we are going to present. Of course, any set of assumptions can be criticised; there will always be some situations in which they are plainly not fulfilled. The real issue is whether or not there is a substantial set of situations in which they are fulfilled. The assumptions also differ from those in many adaptive control algorithms, especially in relation to fluidity over orders of plants and models of those plants. However, we restrict attention to linear systems and, as it turns out, largely restrict attention to stable plants, or possibly plants with no more instability than a pole at the origin. This is a tighter restriction than applies in many adaptive control contexts, though it is not an absolute one for us.

Section 7.3 also discusses the identification-control design cycle, and the details are fleshed out in Section 7.4. Some of the key questions that arise are:

- how can one design to get a series of closed-loop designs in which variation of one parameter yields a variation of the closed-loop bandwidth?
- how can one identify a possibly unstable plant in a stable closed loop, in the presence of noise that is fed back around the loop?
- how can one build a low-order model that is a good approximation of a high-order model, in the sense of giving close closed-loop transfer function matching with a prescribed controller?

Section 7.5 contains examples, and Section 7.6 discusses the promised extension to open-loop unstable plants. We explain in Section 7.7 just how the conceptual problems of many adaptive control algorithms are avoided by the approach. Section 7.8 gives a flavour of what the next steps might be.

7.2 Fundamental Problems in Adaptive Control

In this section, we will recall certain aspects of conventional adaptive control algorithms, and then isolate three differing problems with these algorithms, that raise fundamental questions about their practical utility.

7.2.1 Typical Adaptive Control Algorithms

For the best part of two decades, many of the ideas regarding adaptive control design have been linked by a common thread, perhaps initiated in book form first by reference [65]. Broadly speaking, at any instant of time, one has a controller attached to a plant. One also has a design objective in mind for the closed-loop, which might be a reference trajectory following, pole positioning, or minimising a linear quadratic index. Measurements are collected on the closed-loop comprising the plant and the controller, and an identification algorithm is run on-line. This identification algorithm may seek to estimate the plant directly or, perhaps using some not particularly obvious parameterisation, seek to estimate the parameters directly of a controller that will achieve a particular design objective. Either way, there is an implicit or explicit plant identification procedure going on.

For the sake of simplicity, let us assume that the identification algorithm focuses directly on estimating plant parameters. Using the current estimate of the plant parameters, the control design algorithm computes the controller that would be appropriate or even optimum for the situation, in the event that the true plant was exactly modelled by the plant parameters. With this controller connected, more measurements are then collected. As more measurements are collected, the plant estimate is updated, the controller is updated using the plant estimate, and so on.

A typical theorem in adaptive control starts with a set of premises concerning the plant. For example, it is very often assumed that the plant is linear and time-invariant; perhaps it has a transfer function of a certain degree or relative degree, perhaps it is minimum-phase, perhaps one knows the sign of the DC gain, and so on. The theorem then postulates that a certain algorithm or combination of algorithms is used to estimate the plant parameters, and to compute the controller from these estimates. The theorem then usually goes on to say that as time tends to infinity, the plant parameter estimates converge to limiting values, the controller parameters converge to limiting values, in the steady state correct behaviour is achieved, and in the process all the signals remain bounded. Of course, not every adaptive control theorem is exactly like this, but many are.

In the later subsections of this section, we indicate some problems that arise with this paradigm. The ideas set out in the later sections in the main avoid these problems.

7.2.2 The Problem of Changing Experimental Conditions, Given Accurate but Inexact Models

Most, but not all, people know that one can find two plants whose Nyquist diagrams or impulse responses are almost indistinguishable, but whose responses become extremely different when the plants are placed in a closed-loop control loop, and an identical controller and input used in both loops.

Indeed, in [141] one reads: "Modelling principle 1: arbitrarily small modelling errors can lead to arbitrarily bad closed-loop performance".

The issue here is associated with change of experimental conditions. Models can only have their quality (as an approximation of whatever it is they are modelling) evaluated for a particular set of experimental conditions, and changing from open-loop operation to closed-loop operation with a specified controller is a particular change in experimental conditions. So, in fact, any change of controller is a change of experimental conditions.

It follows that, with a particular controller in the loop, one could have a good model of a plant secured through identification, where good in this sense means that the closed-loop performances match. But when one changes the controller, perhaps as part of the adaptive control process that calls for a controller change in the light of the current identified model, the performance of a new controller with the identified model may be very different to the performance of the new controller with the true plant. *The model that one has at any instant of time, which may be good with the current controller, is not necessarily guaranteed to be good with a changed controller.* To the extent that in adaptive control or even its first cousin, iterative control and identification, where one changes the controller from time to time, one must reckon with the possibility that with the change of controller, the model will become ineffective. One could even contemplate instability.

7.2.3 Impractical Control Objectives

Suppose one has a state variable model of a plant that is completely controllable and completely observable. There is a well-developed body of theory, the LQG design theory, [8] for example, which sets out how one can design an LQG optimal controller. Such a controller can always be constructed to be stabilising, given the minimality of the state variable model of the plant. However, a celebrated paper [54] establishes that in general for such problems, there are no guaranteed gain or phase margins. The reference, in fact, shows with an example that one can formulate an LQG problem that results in a phase margin of ε degrees, where ε is an arbitrarily small positive number.

Now, consider what would happen if one happened to formulate an adaptive LQG problem where the true but unknown plant coincided with the plant that gave the unacceptably low phase margin of the example. In running an adaptive control problem, one would expect very large transient signal values, one would have great difficulty learning the controller and, in fact, in practical terms one would never expect the algorithm to converge. Signals would just be impossibly large.

Underlying this behaviour is the fact that the control objective posed is an impractical one. Hard enough when the plant is known, the objective becomes effectively impossible to meet in an adaptive situation. This is serious enough.

But what makes this problem even more serious for "conventional" adaptive control paradigms is that it is intrinsic in most adaptive control problems that the plant at the outset is unknown. The unknownness of course is only partial, ***but it may logically prevent the assessment of whether or not the control objective is feasible.*** So, before the adaptive control algorithm has even run any distance, one may not only be posing an impossible control problem, ***but one may lack the knowledge to understand that it is impossible.***

The style of adaptive control algorithms that are the subject of the various theorems mentioned in Section 7.2.1 is such that no hints at all are given for dealing with this sort of difficulty.

7.2.4 Transient Instability Problem

The prototypical adaptive control theorems mentioned in Section 7.2.1 normally include the assertion that all signals in the closed loop remain bounded for all time. Left unmentioned is the issue of how big this bound may be. They allow values of control input 10^6 or 10^9 times the value of the steady-state input appearing in response to a unit reference step because instantaneously unstable closed loops are permitted, *i.e.*, closed loops which, if parameters were frozen, would be unstable. The theoretical assurance is only about $t \to \infty$. ***Without a quantitative indication of what the signal bound might be, the theorems provide insufficient assurance*** to a prospective user of an adaptive controller wanting to rely on the algorithm that is the subject of the theorem.

7.3 High-level Overview of the Windsurfer Approach to Adaptive Control

In this section, we shall seek to capture some aspects of the way a human learns windsurfing, and translate these into a high-level statement of what an adaptive control algorithm may look like.

7.3.1 Learning Windsurfing

How does a human learn windsurfing? It is quite clear that one thing they do not do is identify the coefficients in a transfer function linear model of a windsurfer and set the parameters of the controller (and the human is the controller) to provide good closed-loop control. Some other mechanisms are operating. Any human learns windsurfing by initially coping with very gentle wave conditions, light winds, and non-variable winds. Crudely speaking, the human's first successes on a windsurfer involve low gain, low bandwidth, and low disturbance signals. Increasing experience of the windsurfer progressively expands the bandwidth over which the human can control (thus he or she learns to cope with fast shifts in wind direction, choppy seas), and the human also learns to provide a higher gain loop so that he or she can move faster than a novice, under the same wind and sea conditions. The path from low-gain, low-bandwidth control to high-gain, high-bandwidth control is a progressive one, *i.e.*, as the human learns, there are incremental changes in bandwidth and gain that he or she learns to apply. In effect, the human mentally builds some kind of a model that is initially a very primitive one, and which becomes more and more accurate over a wider and wider bandwidth as experience is acquired. Increasing accuracy of the model allows the higher gain controller.

7.3.2 The Starting Point for Windsurfer Adaptive Control

Let us first remark that the words adaptive control in the term "windsurfer adaptive control" are somewhat of a misnomer. "Iterative identification and control" is more accurate. In adaptive control, one typically contemplates identification and control redesign essentially occurring simultaneously. In discrete time, as each new measurement is collected, there is the opportunity to update the plant estimate, and to update the controller parameters. In iterative identification and control design, on the other hand, one contemplates a period in which the controller stays fixed and identification occurs; then the controller is updated, and a new identification process starts with the new controller. The controller remains fixed during this new identification process. This is how we contemplate the windsurfer approach to adaptive controllers working.

As with any adaptive control problem, one makes some *a priori* assumptions concerning the plant. Different adaptive control problems make different *a priori* assumptions. The ones we shall make here are as follows:

Assumption 7.1. The plant is linear, time invariant, and with a transfer function that has no poles in the right half plane nor on the imaginary axis with the possible exception of the origin.

Assumption 7.2. An initial controller is known, almost certainly of low bandwidth and low authority, which stabilises the plant.

If the plant is open-loop stable, any constant controller with small enough gain will not destabilise it. If the plant has a simple pole at the origin, knowledge of the sign of the residue is enough to enable definition of the controller. If the plant has a multiple pole at the origin, somewhat more information is needed.

What we do not need to make any assumption about is the order of the plant, the relative degree of the plant, and whether it is minimum-phase.

One more assumption is needed:

Assumption 7.3. A rational transfer function model of the plant is available, with no poles in the right half plane or on the imaginary axis, except possibly at the origin. The controller that stabilises the plant also stabilises the model.

Note that the model of the plant could be highly inaccurate. There is no requirement, for example, that the model has the same order of the plant, let alone that the model has a collection of parameters which, if adjusted, would give exact modelling of the plant.

7.3.3 The Identification/Control Design Cycle

Beginning at the initial time, and thereafter at all stages in the algorithm, one has available the true plant (as a physical entity, not in terms of a mathematical description), a model of the plant (which is a mathematical description) and a controller (which is available also as a mathematical description). One postulates also that one has the controller connected to the plant, and that the closed loop is excited, with the signals in it observable. At the same time, a copy of the controller is assumed to be available to be connected to the model, with this latter closed loop excited by the same external signals, apart from possible disturbance signals, as those that excite the true plant–controller combination. The outputs of the two closed loops will, in general, be different as a result of modelling errors, and as a result of disturbance signals associated with the true plant. Both outputs are assumed to be available for measurement.

The various steps of the process are as follows:

Step 1. Using the initial model and the initial controller, and with identical input excitation of the two loops, compare the outputs and assess whether or not the current model is a good model of the plant in the presence of the particular controller.

Step 2. If it is a good model, proceed to Step 3. If not, reidentify the model, using some form of closed-loop identification algorithm (discussion of a procedure appears later).

Step 3. Using the model, design the controller to achieve a somewhat wider closed-loop bandwidth (subsequently, a particular controller design algorithm for doing this will be described).

Step 4. Examine the outputs of the two closed-loop arrangements resulting from the new controller with the true plant and the existing model. If, with this controller, the true model is a good model of the plant, return to the previous Step 3. If not, reidentify.
Step 5. Redesign the controller for the new model, and return to Step 3.

In the above set of steps, we have not set out how the process ends. In fact, it will end when

- either the signal-to-noise ratio is such that effective identification is simply not possible; or
- the controller objectives are being satisfactorily met; or
- with the latest model, the control objectives are manifestly in practice unobtainable.

7.4 More Detailed Description of the Algorithm

The two crucial steps in the algorithm are apparently the controller design step, and the closed-loop identification step. We shall discuss these in Sections 7.4.1 and 7.4.2 below. However, as we will review further below in Section 7.4.3, there is a possible requirement for model reduction. Also, it is desirable to understand the conditions that cause the algorithm to stop in some more detail and this is done in Section 7.4.4.

7.4.1 Internal Model Control Design

The internal model control (IMC) method, set out in, for example, [116] is a control design method that is particularly suited to the task of obtaining a prescribed closed-loop bandwidth when the open-loop plant is stable. In particular, there is a way of parameterising the controller effectively directly in terms of this closed-loop bandwidth.

Let us describe the algorithm. Suppose at some step in the process, one has a model of the plant given by a transfer function G_i. Factor this transfer function as $[G_i]_a [G_i]_m$, where $[G_i]_a$ is the all-pass factor associated with G_i, and $[G_i]_m$ is the minimum-phase factor associated with G_i. This factor may have a zero at infinity; we assume it has no other imaginary axis zero. Then, the controller to achieve a notional bandwidth of λ_i is given by the formula

$$K_i = \frac{Q_i}{1 - G_i Q_i} \tag{7.1}$$

Here, Q_i is defined as

$$Q_i = [G_i]_m^{-1} F_i \tag{7.2}$$

where F_i involves the bandwidth λ_i and an integer n:

$$F_i = \left(\frac{\lambda_i}{s+\lambda_i}\right)^n \tag{7.3}$$

The integer n simply needs to be chosen so that Q_i is proper.

The resulting designed closed-loop transfer function $\bar{T}_i$ is given by

$$\bar{T}_i = \frac{G_i K_i}{1 + G_i K_i} \tag{7.4}$$

If one introduces expressions (7.1) to (7.3) into Equation (7.4) for the closed-loop transfer function, one in fact finds that

$$\bar{T}_i = F_i \left[G_i\right]_a \tag{7.5}$$

What this equation reveals is that *the magnitude response of the designed closed loop is given directly by the magnitude response of F_i*. As one would expect, the unstable zeros of the plant appear in the closed-loop transfer function, but in such a way as to not disturb the magnitude response. The attenuation of the closed-loop response is $20n$ dB/decade above the designed closed-loop bandwidth λ_i.

7.4.2 The Identification Step

Consider Figure 7.1, which depicts the plant with unknown transfer function G, a stabilising controller with known transfer function K, and output measurement noise or disturbance obtained by passing white noise through a filter of transfer function H. Suppose one wants to identify G. The signals r, u and y are available for measurement, and the signals r and d are assumed to be independent. While it is true that

$$y = Gu + d \tag{7.6}$$

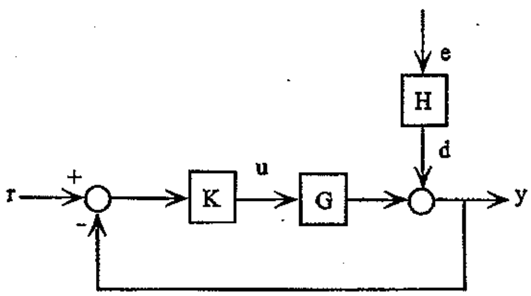

it is not true that u and d are independent. This is because d is fed round to u by the feedback loop and K. Accordingly, many conventional methods of (open-loop) identification cannot be used for identifying G based on measurements of u and y. Additionally, should G be open-loop unstable, some of those methods could also give problems, in the sense that they may presuppose the open-loop stability of G. These few remarks indicate that closed-loop identification is a more difficult task than open-loop identification.

A major step forward in handling closed-loop identification was made by [72]. This work showed how to convert a closed-loop identification problem into a conventional open-loop identification problem, and we shall outline the technique now. Let us suppose that we have available a model G_i which is stabilised by K. Let us also represent G_i as a ratio of stable proper transfer functions, thus

$$G_i = \frac{N_i}{D_i} \tag{7.7}$$

It is also possible to represent K as a ratio of stable proper transfer functions, thus

$$K = \frac{X}{Y} \tag{7.8}$$

Because K is stabilising, it is also possible to require that X, Y satisfy a Bezout identity:

$$N_i X + D_i Y = 1 \tag{7.9}$$

A result originally due to Youla and Kucera, (set out, for example, in [152]) establishes that the set of all plants stabilised by the controller K is precisely the set of transfer functions

$$\mathcal{G} = \left\{ \frac{N_i + RY}{D_i - RX} : R \text{ proper and stable} \right\} \tag{7.10}$$

Evidently, if K stabilises the plant with unknown transfer function G as well as the model with transfer function G_i, the unknown transfer function G must be an element of the set $\mathcal{G}$ of (7.10). Put another way, there must be a proper stable transfer function R such that the following equation holds:

$$G = \frac{N_i + RY}{D_i - RX} \tag{7.11}$$

However, the contribution of [72] was to recognise not only that the task of identifying G was equivalent to the task of identifying R, but that the task of identifying R was one that could be simply cast as an open-loop identification problem.

Consider the arrangement of Figure 7.2, which is in effect a redrawing of Figure 7.1, using the form of G given by (7.11) and introducing the noise d at a slightly different point.

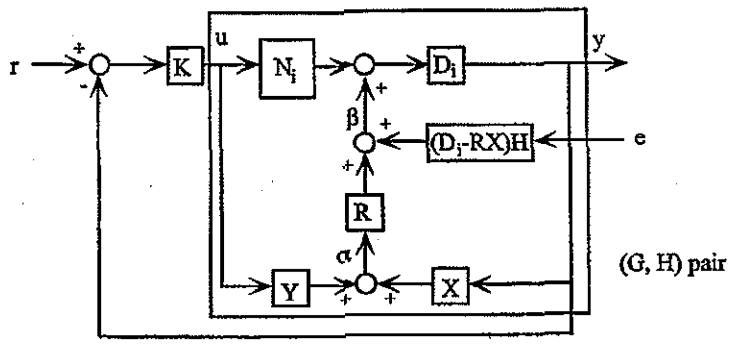

Fig. 7.2. Redrawing of Figure 7.1 using Youla-Kucera parametrisation of plant G

It is not hard to check that the signals α and β can be expressed using known quantities derived from the measurable signals r, u and y:

$$\begin{aligned}\alpha &= Xr\\ \beta &= D_i y - N_i u\end{aligned} \tag{7.12}$$

But more than this is true. As suggested by Figure 7.2, there holds:

$$\beta = R\alpha + (D_i - RX)He \tag{7.13}$$

Notice that because e and r are independent, so must e and α be independent. Also, R must be stable. The task of identifying R is therefore a standard open-loop identification problem, given the availability (from measurements) of α and β, the presumed stability of R and the independence of e and α. Thus, at least in principle, we can see how to solve the closed-loop identification problem.

In open-loop identification, it is typical to contemplate issues of signal-to-noise ratio, use of filters and the like. In this chapter, we will make simple observations on these matters. The first observation can be made in terms of a filtered output error. The unfiltered output error is $\beta - R\alpha$, but we shall look at the filtered output error

$$\xi = Y(\beta - R\alpha) \tag{7.14}$$

[Recall that Y was introduced in (7.8).] Some manipulations will show that:

$$\xi = \left(\frac{GK}{1+GK} - \frac{G_iK}{1+G_iK}\right) r + \frac{1}{1+GK} H e \tag{7.15}$$

This equation reveals several things:

- modulo filtering by Y, the approximation error between the behaviours of G and G_i is not the absolute error, but rather the error associated with the two closed-loop transfer functions. Thus, G_i will be a good model of G with the power norm of ξ (or some other measure of the size of ξ) as the approximation criterion exactly when the two closed-loop transfer functions are close;

- suppose the closed-loop system formed by the unknown plant and the controller has a similar transfer function to the closed-loop system formed by the model and the controller. Then, the signal-to-noise ratio within (7.13) will be small. Accordingly, it will be difficult to identify R. Turning this round, it only makes sense to run an identification step when one knows that there is a genuine closed-loop mismatch;
- as one extends out the bandwidth over which one controls, it is likely to be the case that there is good closed-loop matching over lower bandwidths, and at some stage poorer closed-loop matching near the edge of the new pass band. This suggests that a parameterisation for R should be one that captures well the behaviour near the edge of the current pass band. Thus, R can be of a comparatively low order, even if G_i is of high order, so long as it is able to adequately represent a transfer function that is essentially band-pass in character.

7.4.3 Model Reduction

Recall that in the course of operating the iterative algorithm, one would expect to identify the closed-loop plant (with the aid of an open-loop Youla-Kucera parameter) at least several times. If precautions are not taken, this can result in a dimensionality explosion. To be more precise, suppose that prior to a re-identification of a plant, one is using a model G_i of degree n. Suppose the associated controller also has degree n. It is, in general, possible to find the factorisation representations in Equations (7.7) through (7.9) so that N_i, D_i, X and Y all have degree about n. Suppose that R only has degree 2 (and this may well be an underestimate). Then, as Equation (7.11) shows, the new estimate of the plant is likely to have degree approximately $2n + 2$. Thus, even if a very uncomplicated R is chosen, in an attempt to capture the true plant's contribution to a small part of the pass-band of a closed-loop system, one is faced with a large increase in the degree of the new plant model. It makes sense then to allow for the possibility of a model reduction step at this point. There are many ways to carry out model reduction; suffice to say here that model reduction taking into account the presence of a controller producing a stable closed loop is almost certainly more effective than open-loop model reduction. For a recent treatment of many model reduction ideas, see [119].

7.4.4 Stopping Conditions for the Algorithm

One reason to stop the algorithm would be that the desired closed-loop design has been achieved. However, there are at least two circumstances that might cause premature termination. First, it might be that one is unable to identify accurately, the signal-to-noise ratio being too low. Second, it might be that the most recently identified model and the most recently obtained controller together constitute a closed-loop system of such high sensitivity that further adjustment of the controller in an attempt to produce even wider band performance would be ruled out on the grounds of unacceptably low phase and gain margins, unacceptably high sensitivity and so on.

The paper in [104] analyses those situations that result in difficult or, in practical terms, impossible identification problems. The signal-to-noise ratio associated with the closed-loop output error can be poor because of unstable zeros of the plant that lie inside the pass band, and poor gain or phase margins of the closed-loop. The first phenomenon can only arise when the designed closed-loop bandwidth has been sufficiently pushed out. The second, of course, may well indicate difficulty or inability to expand the bandwidth of the closed-loop system, even given the availability of a good model, perhaps because the sought bandwidth starts to substantially exceed the plant's open-loop bandwidth.

7.5 Examples

Example 7.1. Our first example is drawn from [104]. The plant is a flexible-link robot arm with transfer function

$$G(s) = 0.5196\frac{\prod_{i=1}^{5}(s - z_i)}{\prod_{i=1}^{6}(s - p_i)}$$

The poles are

$$\begin{aligned} p_1, p_2 &= -0.0996 \pm j\,3.0017 \\ p_3, p_4 &= -0.3339 \pm j\,12.131 \\ p_5, p_6 &= -1.845 \pm j\,31.481 \end{aligned}$$

and the zeros are

$$\begin{aligned} z_1 &= -13.162 \\ z_2, z_3 &= -10.646 \pm j\,12.27 \\ z_4, z_5 &= 7.169 \pm j\,11.54 \end{aligned}$$

Evidently, the structure is resonant, with three pole pairs fairly well separated and all lightly damped. The three left half plane zeros are likely to present little problem. The two right half plane zeros will serve to limit the closed-loop bandwidth that can be attained.

The initial model is one that, not unreasonably, captures to a degree the lowest resonance, the roll-off behaviour at high frequencies (20 dB/decade)

– through the insertion of a zero well into the left half plane, and the low frequency gain. Thus

$$G_0(s) = 0.5188 \frac{s + 13.31}{(s + 0.0903 + j\,3.0027)(s + 0.0903 - j\,3.0027)}$$

The frequency response of G and G_0 are shown in Figure 7.3.

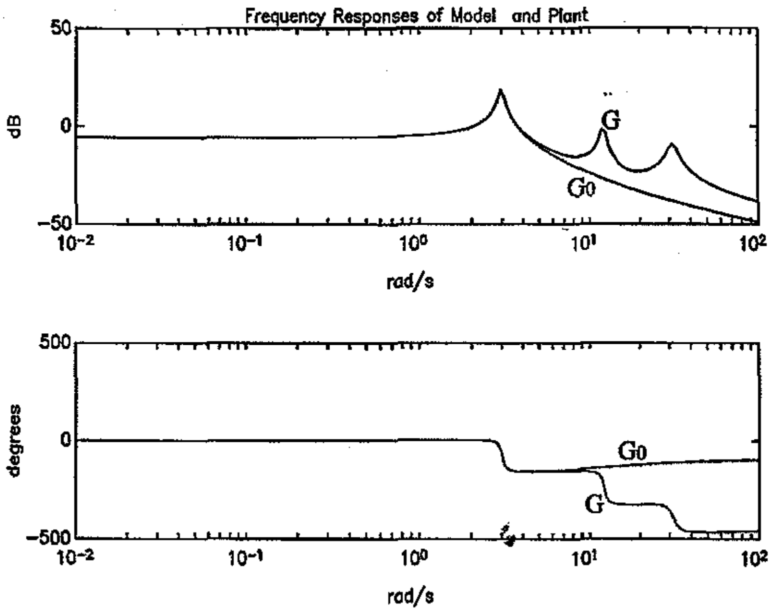

Fig. 7.3. Frequency response of model G_0

A conventional adaptive control algorithm, unless it had built in a sufficiently high degree for estimating the plant, could have great difficulty in coping with unmodelled resonances. One of the features of the windsurfer approach is that *it allows adjustment of the degree of the model transfer function.*

Following the IMC design method, the controller is chosen to secure a closed-loop transfer function $\lambda(s+\lambda)^{-1}$ where λ is variable corresponding to the design closed-loop bandwidth. A first order closed-loop transfer function is acceptable because $G_0(s)$ has no non-minimum-phase zeros, and relative degree one. The controller is second order, with an integrator. As it turns out here, the constant λ defining the bandwidth simply scales the controller transfer function, which is otherwise independent of λ.

For the purposes of identification, the (reference) input consists of just four periods of a zero-mean square wave of amplitude 1 and period π/λ. The plant output was corrupted by zero-mean white noise with standard deviation 0.05. For a validation step (discussed below), the external input can be turned off as appropriate.

Since we have G_0 to begin with, we need to test the closed-loop designs with G_0 and G to validate G_0.

One procedure can be based on looking at power spectra. Recall from (7.13) that we identify R using measurements α, β with

$$\beta = R\alpha + (D_i - RX)He$$

and here e is the perturbing noise. Recall also (7.14) and (7.15), which relate the error in (7.13) to the way the two closed-loop systems (G, K) and (G_0, K) process r:

$$\xi = Y(\beta - R\alpha) = \left(\frac{GK}{1+GK} - \frac{G_0K}{1+G_0K}\right) r + \frac{1}{1+GK}He\,.$$

Evidently, the quantity $\left[GK(1+GK)^{-1} - G_0K(1+G_0K)^{-1}\right] r$ carries the useful information about the existing modelling error, *i.e.*, about R, and the term $(1+GK)^{-1}He$ is the noise component obstructing the determination of R. Call these terms v_0 and ω_0 respectively. Suppose r is a stationary random signal. We can obtain the power spectrum $\Phi_{\omega_0}(\omega)$ by turning r off. We can also determine the power spectrum of ξ:

$$\Phi_{\xi_0}(\omega) = \Phi_{\vartheta_0}(\omega) + \Phi_{\omega_0}(\omega)\,.$$

If this spectrum is significantly larger than $\Phi_{\omega_0}(\omega)$ in some frequency band (especially around the cut-off frequency), then the model G_0 is invalidated in this band.

A second procedure for validation uses correlation techniques, see [108].

Roughly speaking, the cross-correlation method involves plotting (time-lagged) values of the cross-correlation of an input signal and an output error signal. If identification is perfect, the true cross-correlation will be zero. In practice, the sample cross-correlation will not equal the true cross-correlation, and the identification is not perfect. So, non-zero values will be obtained. Should these be within appropriate confidence limits, then that part of the system excited by the known input can be assumed to be well-modelled. No information is gleaned about the noise model; on the other hand, this means that one does not need to try to identify the noise model. Another attractive feature is that the confidence limits are easy to calculate, and are independent of the lag. For details, see [108, pp 511–516].

In terms of our previous notation, this means that we can examine the correlation between α and $\beta - R\alpha$ (when checking the validity of a model of R), or filtered versions of these quantities. Equivalently, one can look at the correlation between r and ξ (or filtered versions of these quantities), and then we are checking more directly the adequacy of G_0 as a model of G, in the presence of a controller K.

When the designed closed-loop bandwidth reaches 1.5 rad/s, the power spectrum validation method reveals no significant difference between $\Phi_{\xi_0}(\omega)$ and $\Phi_{\omega_0}(\omega)$. On the other hand, the method of correlations suggests that G_0 is ceasing to be a good model of G. (Note that at this frequency, one is well short of including the unmodelled resonances in the closed-loop bandwidth. However, the primary resonance, which is in the pass-band, certainly has some modelling error, and this may be responsible.)

The model G_0, however, is retained, and the closed-loop bandwidth pushed out slowly. By the time the bandwidth reaches 3 rad/s, both validation methods indicate poor modelling. The power spectrum method gives a huge discrepancy between $\Phi_{\xi_0}(\omega)$ and $\Phi_{\omega_0}(\omega)$ at 12 rad/s. This is obviously a consequence of G_0 failing to model the resonances in G corresponding to the poles $-0.3339 \pm j\,12.131$, coupled with the fact that this resonance is close enough to the closed-loop cut-off frequency that it shows up.

The fact that $\Phi_{\xi_0}(\omega)$ and $\Phi_{\omega_0}(\omega)$ are so different means that there is adequate signal-to-noise ratio to identify the Youla-Kucera parameter R. Having assumed R is second order and then having identified it, one can validate it by the method of correlations before calculating the new G_1. It passes the test, and so G_1 can be constructed. The G_1 that is found is

$$G_1 = 0.5189 \frac{\prod_{i=1}^{5}(s - z_{i1})}{\prod_{i=1}^{6}(s - p_{i1})}$$

with poles

$$\begin{aligned} p_{11}, p_{21} &= -0.0895 \pm j\,3.0026 \\ p_{31}, p_{41} &= -0.4834 \pm j\,12.03 \\ p_{51}, p_{61} &= -2.475 \pm j\,31.502 \end{aligned}$$

and zeros

$$\begin{aligned} z_{11} &= -12.967 \\ z_{21}, z_{31} &= -7.336 \pm j\,11.05 \\ z_{41}, z_{51} &= 9.098 \pm j\,12.07 \end{aligned}$$

As expected, G_1 has higher order than G_0. Frequency responses of G and G_1 are shown in Figure 7.4. This new model G_1 captures the imaginary parts of the zeros and poles remarkably well.

The closed-loop bandwidth can now be progressively pushed out. Notice that the defined closed-loop transfer function now is the product of an all-pass transfer function and $\lambda(s+\lambda)^{-1}$, because of the non-minimum-phase zeros in $G_1(s)$. When the bandwidth reaches 12 rad/s, both the validation methods suggest that G_1 is inadequate. When re-identification is performed and the controller readjusted for the new model, *performance cannot be improved.* Why? Notice that G and G_1 have unstable zeros with a bandwidth of about 12 rad/s. This will set a "fundamental limit" (see Section 4.4 and [140]) on the

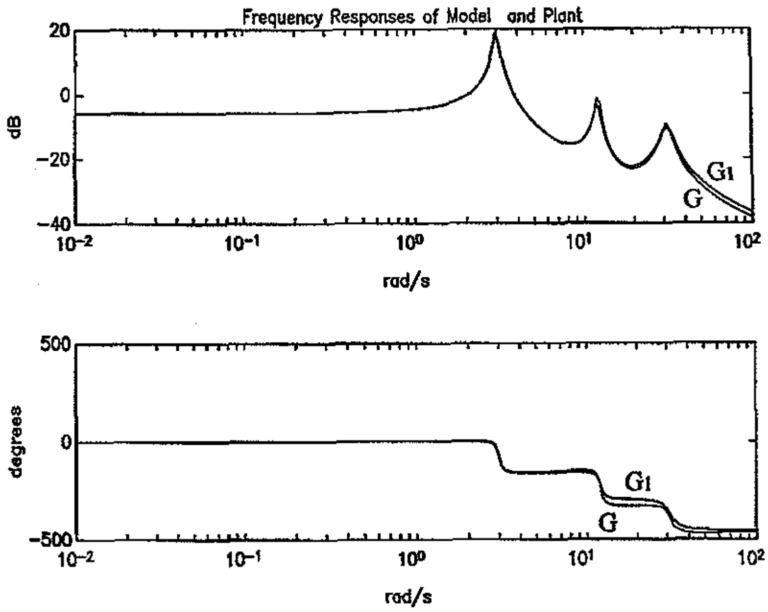

Fig. 7.4. Frequency response of model G_1

practically-achievable closed-loop bandwidth, and even very small modelling errors will not really be able to be properly compensated in a sound design.

The "fundamental limit" concept manifests itself in different ways; one is that, when one is designing at the fundamental limits, the true plant has to be known extremely accurately in order to get the designed-for performance. In our situation, the accuracy requirement for design cannot be matched by the accuracy offered by the identification algorithm.

Example 7.2. Our second example is drawn from [102]. The true plant is

$$G(s) = \frac{9}{(s+1)(s^2+0.06s+9)}$$

with $y = Gu + e$, and disturbance e being zero mean noise with 100 Hz bandwidth and flat spectrum from 0 to 100 Hz of height 0.0025.

The initial model $G_0(s)$ mismodels the DC gain and fails to capture the resonance:

$$G_0(s) = \frac{0.8}{s+1.2}$$

Any adaptive control algorithm that forces the plant model to remain first

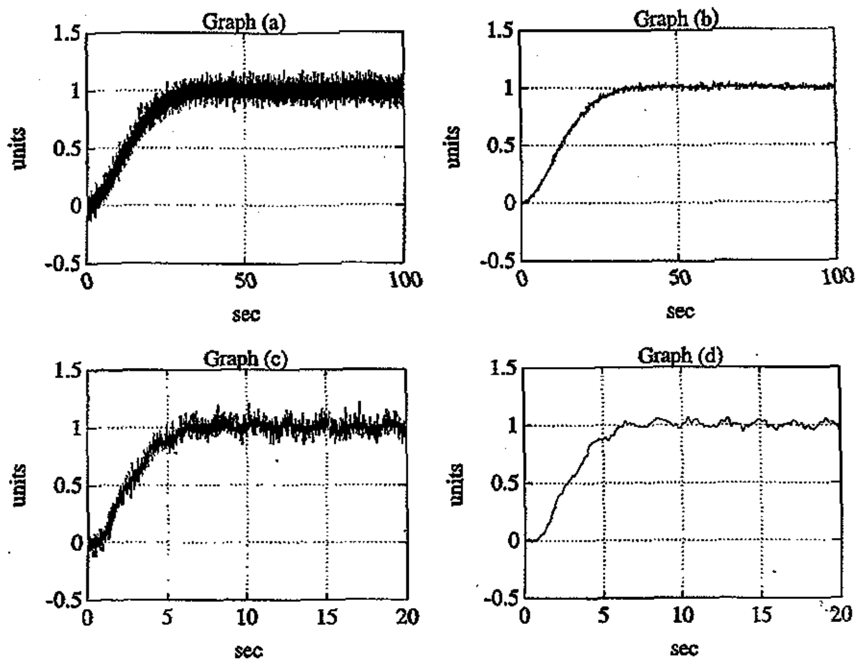

Fig. 7.5. Simulation results 1

order will be very challenged by this example. The design closed-loop transfer function is $\lambda^2(s+\lambda)^{-2}$ with λ variable. Figure 7.5 shows in (a) and (b) the actual closed-loop system unit step responses with $\lambda = 0.1$ rad/s (in (b), low-pass filtering has been applied). In (c) and (d), the bandwidth λ is 0.5 rad/s. Bearing in mind that the ideal response is that from $\lambda^2(s+\lambda)^{-2}$ excited by a unit step, we can conclude that by $\lambda = 0.5$ rad/s, the model is inadequate.

In order to improve the identification, low-amplitude sinusoidal signals near 0.5 rad/s were superimposed on the unit step excitation. A second order R was also assumed. This resulted in a new model, of higher degree than $G_0(s)$:

$$G_1(s) = \frac{0.0625s^2 - 0.34s + 10.28}{s^3 + 1.28s^2 + 9.12s + 10.32}$$

Redesign of the controller using $G_1(s)$ and with $\lambda = 0.5$ rad/s yields closed-loop performance with the true $G(s)$ shown in Figure 7.6. We can increase λ to 2 rad/s without significant degradation.

It turns out that G_1 is a good model of G over this bandwidth, as seen in Figure 7.7.

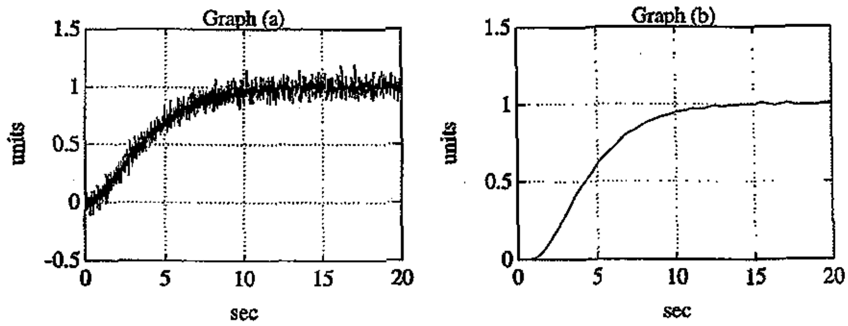

Fig. 7.6. Simulation results 2

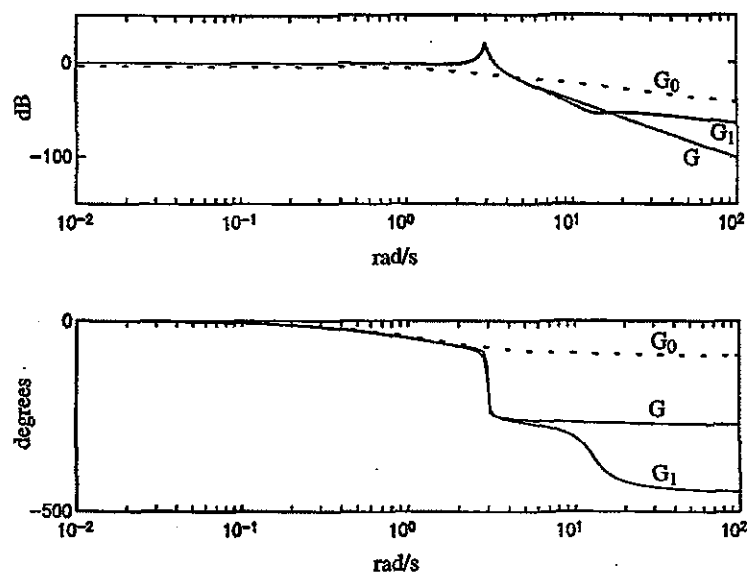

Fig. 7.7. Frequency responses of models and plant

7.6 Unstable Plants

The methods we have outlined up to now depend on the open-loop plant being stable. This is not because of any constraints in the identification process but rather because the IMC design method is far more efficacious for stable plants than for unstable plants [36, 116]. These techniques can often result in the controller having a non-minimum phase-zero, which may lie in the closed-loop pass band. Also, the techniques may involve more than a single parameter, and any parameters that they do involve are not as simply related to the closed-loop bandwidth as is the single parameter in the IMC design procedure for a stable plant. For these reasons, and as explained in more detail in the paper [105], a different structure is used, as depicted in Figure 7.8.

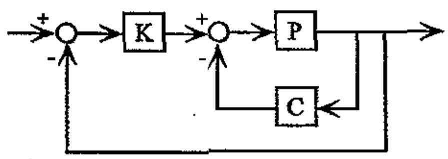

Fig. 7.8. Modification to IMC method for unstable plant P, with C stabilising the inner loop

The inner-loop controller C is chosen in order to stabilise the plant P. The result is an internal closed loop with transfer function $G = P/(1 + PC)$. This inner closed-loop transfer function G is stable, and then IMC methods can be applied to the design of K.

The paper [105] sets out a number of guidelines for the choice of C. For example, should P have an unstable closed loop pole at $+a$, it is desirable for the closed-loop to have a pole at $-a$. Of course, control of an unstable plant is always harder than controlling a stable plant, and there are fundamental limitations expressible with Bode integral theorems, waterbed theorems and the like [51, 140]. Another example of such a difficulty arises when the plant P cannot itself be stabilised by an open-loop-stable C (it is easy to find examples of unstable plants for which there is no controller that is both stable and stabilising). In this case, the transfer function G necessarily acquires a right half plane zero, at any point where C has a right half plane pole. The controller K must then be found to stabilise a stable plant with an unstable zero, and the usual sorts of restrictions start to come into play when the closed-loop bandwidth starts to include that zero.

As these brief remarks should make clear, plant instability can be a major complication for design. Add to this the fact that simple single parameter adjustment of the closed-loop bandwidth is no longer possible, and one can see the situation is a good deal more complex.

7.7 Coping with the "Fundamental Problems"

In Section 7.2, we recorded some fundamental or conceptual problems of much adaptive control theory and thinking. The particular problems flagged were:

- the problem of changing experimental conditions, given accurate but inexact models; the issue is that once you change the controller, the model employed for controller redesign may become invalid;

- the problem of impractical control objectives; the issue is that the possibility or impossibility in practical terms of the control objective may not be assessable at the outset, given the (possibly partial) lack of knowledge of the plant;
- the transient instability problem: saying that all signals are bounded in a theorem does not preclude 1 MA current in a 1 kW motor.

Let us consider how these problems are addressed. The method of addressing the first is somewhat *ad hoc*. When we attempt a controller redesign, we make an adjustment of a bandwidth. If this adjustment is very small, we know that unsuitability of the model will not be an issue. In fact, the quality of the model (as displayed in its ability to match closed-loop performance with the same controller connected to model and to plant) will only degrade smoothly as the bandwidth of the closed-loop system is pushed out, and the controller is changed to do this. Thus, the adaptive algorithm must ensure that the trial expansions of bandwidth are not too large. Nor, of course, should they be so small that the whole process takes too long. An expansion of 20 per cent is a good target to begin with; but note that this could be too big if it suddenly swept into the bandwidth an extremely oscillatory mode. Finally, note that we cannot attempt a re-identification until there is a genuine difference between the two closed-loop transfer functions – otherwise the signal-to-noise ratio will be too small to allow an accurate identification.

How is the problem of impractical control objectives addressed? Here, the approach of this chapter comes in to its own. Effectively, a flag is raised that indicates the impracticality. One has to stop the algorithm early, on account of an inability to identify due to adverse signal-to-noise ratio, or an inability to find a controller giving a satisfactory (including acceptable gain and phase margin) closed-loop design.

Likewise, the problem of transient instability is dealt with well. At each step of the algorithm, it is guaranteed that an unstable (frozen or instantaneous) closed loop cannot be encountered. Smooth expansion of the bandwidth as part of the algorithm produces in the first instance degradation of the quality of the model, and that degradation sets in well before the model–controller closed loop is stable while the real plant–controller one is unstable.

7.8 Some Further Quantitative Directions

In this section, we will discuss briefly a direction in which the ideas of this chapter are being taken further. This concerns the formalising of approximation measures.

In Section 7.2, we alluded to a problem inherent in many adaptive control methodologies, which we termed "the problem of changing experimental conditions, given accurate but inexact models". The idea is that with one controller in place, a model may be a good model of a plant, while with another controller in place, it may not be. The point we want to make here is that this conceptual, rather qualitatively stated, difficulty is capable of quantitative refinement, using the concept of the ν-gap metric [153]. The first application to adaptive control appears in [7].

The question we are going to discuss in quantitative terms is the following. Suppose that G_0 is a good model of G in the presence of a controller K_0. Suppose K_1 is another controller for G_0. To what extent will G_0 be a good model for G in the presence of K_1?

Quantitative treatment of this question requires some preliminary work, to establish certain distance measures.

We shall first explain a distance measure for two transfer functions. At each frequency, we can define a certain distance, termed the *chordal distance*, which is a generalisation of the distance between points associated with the same frequency ω on two Nyquist diagrams. Then we define a single distance, by looking at all frequencies. First, for transfer functions $K_0(j\omega)$ and $K_1(j\omega)$, we have the chordal distance

$$\kappa[K_0, K_1](j\omega) = \frac{|K_1(j\omega) - K_0(j\omega)|}{\sqrt{1 + |K_1(j\omega)|^2}\sqrt{1 + |K_0(j\omega)|^2}} \tag{7.16}$$

If $K_1(j\omega_0) = \infty$, we compute the chordal distance by letting $\omega \to \omega_0$. We also define the ν-gap distance by

$$\delta_\nu(K_0, K_1) = \sup_\omega \kappa[K_0, K_1](j\omega) \tag{7.17}$$

provided

$$1 + K_1(-s)K_0(s) \neq 0\,, \quad \text{for all} \quad s = j\omega \tag{7.18}$$

$$wno[1 + K_1(-s)K_0(s)] + \eta(K_0) - \bar{\eta}(K_1) = 0 \tag{7.19}$$

where $\eta(K_0)$ and $\bar{\eta}(K_1)$ denote the number of open right half plane poles of K_0 and the number of closed right half plane poles of K_1 respectively, while $wnoX(s)$ denotes the number of clockwise encirclements of the origin by $X(s)$ as s moves around the standard Nyquist contour (up the imaginary axis), indented into the right half plane around any $j\omega$-axis pole of $X(s)$. If (7.18) or (7.19) fail, then

$$\delta_\nu(K_0, K_1) = 1 \tag{7.20}$$

Next, we shall look at closed-loop behaviour. In relation to a stable closed-loop plant G and controller K, defined the *generalised sensitivity matrix*,

$$T(G, K) = \begin{bmatrix} G \\ 1 \end{bmatrix} (1 + KG)^{-1} \begin{bmatrix} K & 1 \end{bmatrix} \tag{7.21}$$

Notice that the four entries of T include the conventional sensitivity and complementary sensitivity functions. It is not hard to check that the maximum singular value of T at any frequency is

$$\bar{\sigma}[T(G,K)](j\omega) = \frac{\sqrt{1+|G(j\omega)|^2}\sqrt{1+|K(j\omega)|^2}}{|1+G(j\omega)K(j\omega)|^2} \tag{7.22}$$

A large value of $\bar{\sigma}[T]$ at some frequency generally signals a poor design. In terms of the complementary sensitivity $M, \bar{\sigma}_2[T]$ is given by

$$\bar{\sigma}^2(T) = \left[1 + \frac{1}{|G|^2}\right]\left[|M|^2 + |G|^2\,|1-M|^2\right] \tag{7.23}$$

This reveals, for example, that achieving $|M| \simeq 1$ in frequency bands where $|G|$ is small yields a large $\bar{\sigma}(T)$; and achieving very small $|M|$ when $|G|$ is very large, *i.e.*, suppressing a pole close to the $j\omega$-axis so that it is outside the closed-loop bandwidth, also yields a large $\bar{\sigma}(T)$.

We can say that G_0 is a good model of G when K_0 is connected when for some small ε, there holds,

$$\|T(G_0,K_0) - T(G,K_0)\|_\infty \le \varepsilon \tag{7.24}$$

or possibly

$$\|T(G_0,K_0) - T(G,K_0)\|_\infty \le \varepsilon\,\|T(G_0,K_0)\|_\infty \tag{7.25}$$

or

$$\bar{\sigma}[T(G_0,K_0) - T(G,K_0)](j\omega) \le \varepsilon\bar{\sigma}\,[T(G_0,K_0)]\,(j\omega) \quad \text{for all } \omega \tag{7.26}$$

The question we want to consider now is as follows: suppose one of (7.24) through (7.26) holds, and K_1 is a second controller so that (G_0, K_1) is an acceptable closed loop. What are (sufficient) conditions for (G, K_1) to be, firstly, stable and secondly, a loop with similar behaviour to (G_0, K_1)?

Stability is easiest. According to [153], a sufficient condition is:

$$\delta(K_0, K_1) < 1 \tag{7.27}$$

and

$$\bar{\sigma}[T(G,K_0)](j\omega)\kappa(K_0,K)(j\omega) < 1 \quad \text{for all } \omega \tag{7.28}$$

One has knowledge of $T(G_0, K_0)$; and so given one of (7.24), (7.25) and (7.26), one can check (7.28). Though we will not use the fact, we comment that

$$\bar{\sigma}\,[T(G,K_1)(j\omega)] \le \frac{\bar{\sigma}[T(G,K_0)](j\omega)}{1-\bar{\sigma}[T(G,K_0)](j\omega)\kappa(K_0,K_1)(j\omega)} \tag{7.29}$$

Next, the following inequality is known [7, 153],

$$\bar{\sigma}[T(G,K_1) - T(G,K_0)] \le \frac{\bar{\sigma}[T(G,K_0)]^2 \kappa(K_0,K_1)}{1-\bar{\sigma}[T(G,K_0)]\kappa(K_0,K_1)} \tag{7.30}$$

where we have suppressed dependence on ω.

Accordingly,

$$\begin{aligned} &\bar{\sigma}[T(G,K_1) - T(G_0,K_1)] \\ &\qquad \le \bar{\sigma}[T(G,K_1) - T(G,K_0) + T(G,K_0) \\ &\qquad\quad - T(G_0,K_0) + T(G_0,K_0) - T(G_0,K_1)] \\ &\qquad \le \bar{\sigma}[T(G,K_1) - T(G,K_0)] + \bar{\sigma}[T(G,K_0) - T(G_0,K_0)] \\ &\qquad\quad + \bar{\sigma}[T(G_0,K_0) - T(G_0,K_1)] \end{aligned} \tag{7.31}$$

For simplicity, neglect $\bar{\sigma}[T(G,K_0) - T(G_0,K_0)]$ on the grounds that G_0 is a good model of G with controller K_0 and take $\bar{\sigma}[T(G,K_0)] = \bar{\sigma}[T(G_0,K_0)]$. Then

$$\begin{aligned} &\bar{\sigma}[T(G,K_1) - T(G_0,K_1)] \\ &\qquad \le \frac{\bar{\sigma}[T(G_0,K_0)]^2 \kappa(K_0,K_1)}{1-\bar{\sigma}[T(G_0,K_0)]\kappa(K_0,K_1)} + \bar{\sigma}[T(G_0,K_0) - T(G_0,K_1)] \end{aligned} \tag{7.32}$$

A greater overbound again is possible, since it is known that [7, 153]

$$\kappa(K_0,K_1) \le \bar{\sigma}[T(G_0,K_0) - T(G_0,K_1)]\,. \tag{7.33}$$

There results

$$\begin{aligned} &\bar{\sigma}[T(G,K_1) - T(G_0,K_1)] \\ &\qquad \le \left[\frac{\bar{\sigma}[T(G_0,K_0)]^2}{1-\bar{\sigma}[T(G_0,K_0)]\kappa(K_0,K_1)} + 1\right] \bar{\sigma}[T(G_0,K_0) - T(G_0,K_1)] \end{aligned} \tag{7.34}$$

This formula suggests that changes in $T(G_0,K)$, when K varies from K_0 to K, are reflected proportionally in $T(G,K)$, but with a multiplier. The multiplier may be very large when the initial design $T(G_0,K_0) \simeq T(G,K_0)$ has large $\bar{\sigma}(T)$, *i.e.*, is a poor, or non-robust design. The multiplier may also be large if the changes in K brings the closed-loop system near the stability boundary.

Acknowledgements

Support of the Office of Naval Research, Washington, is gratefully acknowledged. This work also benefited from discussions with R. Bitmead and M. Gevers.