



Brief Paper

Direct iterative tuning via spectral analysis[☆]Leonardo C. Kammer^{a,*}, Robert R. Bitmead^b, Peter L. Bartlett^c^a*Pulp and Paper Centre, University of British Columbia, 2385 East Mall, Vancouver, BC V6T 1Z4, Canada*^b*Department of Mechanical & Aerospace Engineering, University of California at San Diego, La Jolla CA 92093-0411, USA*^c*Computer Sciences Laboratory, Research School of Information Sciences & Engineering, Australian National University, Canberra ACT 0200, Australia*

Received 9 March 1998; revised 24 June 1999; received in final form 19 November 1999

Abstract

This paper introduces a nonparametric frequency-domain scheme for tuning closed-loop systems. A linear quadratic criterion is used for guiding the system towards an optimal-regulation state. The computation of tuning directions is performed with unbiased estimates of first and second derivatives of the cost function. In order to provide cautious adjustments of the controller parameters, this scheme introduces a mechanism that guarantees stability of the loop during the tuning procedure. This mechanism uses the generalised stability margin and a metric on transfer functions to impose safe limits on the variation of the control law. © 2000 Elsevier Science Ltd. All rights reserved.

Keywords: Control tuning; Optimisation; Iterative methods; Stability tests; Spectral analysis

1. Overture

Performance improvement of initially stable closed-loop systems is one of the main tasks entrusted to control engineers. Although adaptive controllers could be used to perform this task, they introduce extraneous nonlinearities into the system's behaviour (Mareels & Bitmead, 1986). These nonlinear dynamics are problematic because the control law varies at each sampling time. A safer way to deal with the task of tuning controllers for time-invariant processes is to apply techniques of iterative control. The basic idea is to observe the closed-loop signals, under a fixed controller, for a longer period of time and then to decide on a new control law.

All iterative methods currently available for the accomplishment of the tuning task start with an experiment

performed on the loop. An external excitation is injected into the closed-loop system and two possible ways might be followed: modelling of the plant followed by control design, or direct computation of performance derivatives followed by controller adjustment. This is similar to the distinction between 'indirect' and 'direct' adaptive control. The underlying reasoning for choosing one approach or another includes the degree of confidence the user has in the information obtained from the experiment.

The option of modelling the plant is usually followed by the design of a control law, as in Zang, Bitmead and Gevers (1995), Schrama (1992) and Lee, Anderson, Kosut and Mareels (1993), and this reflects a high degree of confidence in that model. In most of these approaches the structure of the controller is not assumed to be constrained, and the solution might be of high order. When the complexity of the control law is restricted, as in the widely used controllers with proportional, integral and derivative (PID) actions, methods based on derivatives of the performance with respect to the controller parameters become more appropriate, thus avoiding plant or controller model reduction. There are, however, serious difficulties associated with performing parametric plant modelling: structure selection, parameter estimation, model validation, etc.

An alternative to computing a parametric model of the plant is to perform the initial experiment in such a way

[☆]Part of this paper was presented at the IFAC Workshop on Adaptive Systems in Control and Signal Processing, which was held in Glasgow, U.K., during August 1998. This paper was recommended for publication in revised form by Associate Editor G. Dumont under the direction of Editor F.L. Lewis. This work was originally performed when the authors were with the Department of Systems Engineering at The Australian National University. It was supported by the Cooperative Research Centre for Robust and Adaptive Systems.

*Corresponding author. Tel.: +1-604-822-8237; fax: +1-604-822-8563.

E-mail address: kammerl@ppc.ubc.ca (L.C. Kammer).

that the derivatives of the cost function with respect to the controller parameters are obtained directly from (filtered) data. Hjalmarsson, Gevers, Gunnarsson and Lequin (1998) presents a scheme that performs the initial experiment with the injection, at the reference input, of plant-output signals obtained during normal operation of the closed loop. By doing so they are able to compute unbiased estimates of the first derivatives of the cost function. The basic idea is to use those derivatives to compute a direction for adjusting the controller parameters; once the best controller is found in that direction, another experiment has to be performed in order to obtain a new direction of descending cost, reflecting a complete non-reliance on explicit plant modelling.

This reasonably clear distinction among iterative approaches might become blurred when nonparametric models are used within the schemes. This paper presents a method for computing derivatives of linear quadratic (LQ) cost functions with respect to the controller parameters via spectral analysis of closed-loop experimental data. The scheme, referred to as frequency-domain tuning (FDT), does not explicitly model the plant.

In principle, the FDT scheme could be viewed as a frequency-domain variant of the one in Hjalmarsson et al. (1998), but the nonparametric characterisation of the closed-loop system allows for some important improvements over that scheme. In addition to the expressions of first derivatives of the cost function, the FDT scheme computes unbiased estimates of the second derivatives of that function with respect to parameters of the controller. These quantities greatly improve the estimation of the closed-loop behaviour around its operating point, so that tight supervision can be done on the computation of tuning directions. Moreover, first and second derivatives together can be used to detect (local) LQ optimality of the closed-loop system.

A mechanism that guarantees stability of the closed-loop system, during the tuning procedure, is also introduced into the iterative scheme. This mechanism uses the signal spectra to compute a stability margin of the loop, which limits the variation of the controller parameters. The motivation for this implementation comes from examples where gradient-based schemes seem to be improving the system performance but the loop unexpectedly becomes unstable.

2. Scene building

The criterion adopted in this work to guide the tuning procedure is the minimisation of the frequency-weighted linear quadratic cost function:

$$J^{lq} = \lim_{N_t \rightarrow \infty} \left\{ \frac{1}{N_t} \sum_{k=1}^{N_t} [(F_y y_k)^2 + \lambda (F_u u_k)^2] \right\}, \quad (1)$$

where y_t and u_t are, respectively, the output and the input of the plant, λ is a nonnegative scalar that weights the penalty on the control action, and the stable linear filters F_y and F_u introduce frequency-dependent weights on the variance of $\{y_t\}$ and $\{u_t\}$, respectively.

In order to compare performances of different controller tunings, in practice, a finite set of data is collected from the system, resulting in an estimate of (1). Moreover, for the specific purpose of computing a new direction of controller tuning, the cost function will be expressed in the frequency domain, by means of Parseval's Theorem.

$$J^{lq} = \lim_{N_\omega \rightarrow \infty} \left\{ \frac{1}{N_\omega} \sum_{k=1}^{N_\omega} [|F_y|^2 \phi_y(2\pi k/N_\omega) + \lambda |F_u|^2 \phi_u(2\pi k/N_\omega)] \right\} \quad (2)$$

with ϕ_y and ϕ_u being the power spectral densities of the signals $\{y_t\}$ and $\{u_t\}$, respectively.

The derivatives of (2) with respect to the controller parameters can be obtained by spectral analysis of closed-loop signals and full knowledge of the control law, hence there is no explicit need for modelling the plant. Nevertheless, in order to reach this mathematical result the following linear model of the plant is used:

$$y_t = P(q)u_t + H(q)e_t,$$

where $H(q)$ is a stable and stably invertible filter and $\{e_t\}$ is a zero-mean white-noise process of variance σ_e^2 .

The structure of the controller is assumed to have two degrees of freedom, despite only the regulation problem being addressed here. The control action is given by

$$u_t = C_{ff}(q, \rho)r_t - C_{fb}(q, \rho)y_t,$$

where the vector ρ represents the tuning parameters of the controller.

The reference signal, r_t , is assumed to be kept identically zero during normal operation, therefore spectral analyses of the plant output and input signals result in

$$\phi_y(\omega, \rho) = \left| \frac{H(e^{i\omega})}{1 + P(e^{i\omega})C_{fb}(e^{i\omega}, \rho)} \right|^2 \sigma_e^2 \quad (3)$$

and

$$\phi_u(\omega, \rho) = \left| \frac{H(e^{i\omega})C_{fb}(e^{i\omega}, \rho)}{1 + P(e^{i\omega})C_{fb}(e^{i\omega}, \rho)} \right|^2 \sigma_e^2, \quad (4)$$

respectively, provided the closed-loop system is stable.

Finally, the frequency-domain characteristics of the system under stationary excitation through the reference signal are also needed. The cross-spectral density between $\{y_t\}$ and $\{r_t\}$ is given by

$$\phi_{yr}(\omega, \rho) = \frac{P(e^{i\omega})C_{ff}(e^{i\omega}, \rho)}{1 + P(e^{i\omega})C_{fb}(e^{i\omega}, \rho)} \phi_r(\omega), \quad (5)$$

and the cross-spectral density between $\{u_t\}$ and $\{r_t\}$ is

$$\phi_{ur}(\omega, \rho) = \frac{C_{ff}(e^{i\omega}, \rho)}{1 + P(e^{i\omega})C_{fb}(e^{i\omega}, \rho)}\phi_r(\omega). \quad (6)$$

With this set of equations it is possible to proceed with the computation of the derivatives of J^{lq} , given by (2), with respect to the controller parameters.

For the sake of clarity, J^{lq} is expressed in terms of auxiliary functions.

$$J^{lq}(\rho) = \lim_{N_\omega \rightarrow \infty} \left\{ \frac{1}{N_\omega} \sum_{\omega} [g(\omega, \rho)h(\omega, \rho)] \right\},$$

$$g(\omega, \rho) = |F_y(e^{i\omega})|^2 + \lambda |F_u(e^{i\omega})|^2 |C_{fb}(e^{i\omega}, \rho)|^2,$$

$$h(\omega, \rho) = \frac{|H(e^{i\omega})|^2}{|1 + P(e^{i\omega})C_{fb}(e^{i\omega}, \rho)|^2} \sigma_e^2 \quad (= \phi_y(\omega, \rho)). \quad (7)$$

The function arguments were presented above so as to emphasise variable dependencies. In the sequel, the notation is shortened by dropping these arguments.

The derivative of J^{lq} with respect to a parameter ρ_l of the controller is

$$\frac{\partial J^{lq}}{\partial \rho_l} = \lim_{N_\omega \rightarrow \infty} \left\{ \frac{1}{N_\omega} \sum_{\omega} \left[\frac{\partial g}{\partial \rho_l} h + g \frac{\partial h}{\partial \rho_l} \right] \right\},$$

$$\frac{\partial g}{\partial \rho_l} = 2\lambda |F_u|^2 \Re \left(C_{fb}^* \frac{\partial C_{fb}}{\partial \rho_l} \right),$$

$$\frac{\partial h}{\partial \rho_l} = -2\phi_y \Re \left(\frac{P}{1 + PC_{fb}} \frac{\partial C_{fb}}{\partial \rho_l} \right),$$

where C_{fb}^* is the complex conjugate of C_{fb} and $\Re(\cdot)$ returns the real part of its argument. Second derivatives can also be computed from those signals.

$$\frac{\partial^2 J^{lq}}{\partial \rho_l \partial \rho_k} = \lim_{N_\omega \rightarrow \infty} \frac{1}{N_\omega} \sum_{\omega} \left[\frac{\partial^2 g}{\partial \rho_l \partial \rho_k} h + \frac{\partial g}{\partial \rho_l} \frac{\partial h}{\partial \rho_k} + \frac{\partial g}{\partial \rho_k} \frac{\partial h}{\partial \rho_l} + g \frac{\partial^2 h}{\partial \rho_l \partial \rho_k} \right],$$

$$\frac{\partial^2 g}{\partial \rho_l \partial \rho_k} = 2\lambda |F_u|^2 \Re \left(\frac{\partial C_{fb}}{\partial \rho_l} \frac{\partial C_{fb}^*}{\partial \rho_k} + C_{fb}^* \frac{\partial^2 C_{fb}}{\partial \rho_l \partial \rho_k} \right),$$

$$\frac{\partial^2 h}{\partial \rho_l \partial \rho_k} = -2\phi_y \Re \left(\frac{P}{1 + PC_{fb}} \frac{\partial^2 C_{fb}}{\partial \rho_l \partial \rho_k} - 2 \frac{P}{1 + PC_{fb}} \frac{\partial C_{fb}}{\partial \rho_l} \frac{P}{1 + PC_{fb}} \frac{\partial C_{fb}}{\partial \rho_k} - \left| \frac{P}{1 + PC_{fb}} \right|^2 \frac{\partial C_{fb}}{\partial \rho_l} \frac{\partial C_{fb}^*}{\partial \rho_k} \right).$$

By observing the expressions for the first and second derivatives of the cost function, one can verify that almost all terms are known or obtainable from normal operating data. The only term remaining can be obtained via an intrusive experiment performed on the loop by

injecting a stationary non-degenerate signal $\{r_t\}$. Actually, from (5) it is straightforward to compute

$$\frac{P(e^{i\omega})}{1 + P(e^{i\omega})C_{fb}(e^{i\omega}, \rho)} = \frac{\phi_{yr}(\omega, \rho)}{C_{ff}(e^{i\omega}, \rho)\phi_r(\omega)}. \quad (8)$$

3. Production

With the first and second derivatives of J^{lq} we proceed by forming its gradient (*row*) vector, ∇J , and its Hessian matrix, \mathbf{H} . These two quantities, ∇J and \mathbf{H} , are the basis of the methods used for computing the tuning direction.

Unlike adaptive control, where new directions have to be computed in between adjacent samples, the computation of tuning directions in iterative control is performed block-wise and, therefore, is allowed to be very elaborate. This circumstance favours the use of Newton's method, which is accurate but complicated due to a matrix inversion.

Newton's method uses the gradient vector and the Hessian matrix to approximate the local behaviour of the cost function around the current operating point. According to this approximation, the following change in the controller parameters would lead to an operating point with all first derivatives of the cost function identically zero:

$$\Delta \rho = -\mathbf{H}^{-1} \nabla J^T. \quad (9a)$$

That operating point, in the approximation, can be a minimum, a saddle or a maximum, depending whether the Hessian matrix is positive definite, indefinite or negative definite, respectively. In fact, only for the first of these situations is $\Delta \rho$ an appropriate direction of tuning.

A good alternative to Newton's method, when it fails to deliver a direction towards minimum cost, is the combination of the method of Steepest Descent and a *negative curvature* descent direction (Fletcher, 1987, p. 49). Given that the Hessian matrix is available, the following improvement to the original method of Steepest Descent results in a $\Delta \rho$ that corresponds to the minimum approximated cost function in the direction of the gradient vector:

$$\Delta \rho = -\gamma \nabla J^T,$$

$$\gamma = \frac{\nabla J \nabla J^T}{\nabla J \mathbf{H} \nabla J^T}, \quad (9b)$$

as long as γ is positive. This improvement has its basis in the analysis of the original method on a purely quadratic problem. Under some circumstances, the Steepest Descent direction alone makes the tuning scheme extremely slow in terms of performance improvement. In such cases, it is worth alternating those directions with

directions of negative curvature. The authors have experienced very good results by using the direction of the eigenvector associated with the smallest (most negative) eigenvalue of the Hessian matrix.

In the unfortunate situation of γ being negative (indicating lack of local convexity), the original method of Steepest Descent is used instead, that is,

$$\Delta\rho = -\nabla J^T. \quad (9c)$$

In summary, the direction of controller tuning is obtained by one of (9a)–(9c), where the choice depends on local conditions.

It is not likely that the real system conforms entirely to our quadratic approximation, especially when the operating point is far from a local minimum. In practice the parameters of the controller are changed along the line

$$\rho = \rho_0 + \alpha\Delta\rho \quad (10)$$

in the search for a point of minimum cost. The vector ρ_0 contains the initial parameters of the controller, and α is varied along \mathbb{R}^+ .

An initial value for α and its subsequent increments can be chosen on-line, according to several factors:

- the minimum cost function along (10) is expected to be at $\alpha \simeq 1$, except when the original method of Steepest Descent is used;
- the anticipated behaviour of the cost function is given by

$$J^q(\rho) \simeq J^q(\rho_0) + \alpha\nabla J\Delta\rho + \frac{\alpha^2}{2}\Delta\rho^T\mathbf{H}\Delta\rho, \quad (11)$$

- stability of the closed-loop system has to be maintained (see Section 4).

Some basic characteristics of the tuning algorithm provide clues for deciding whether the control law is nearly optimal or not. One example is the anticipated reduction of the cost function given by (11), at $\alpha = 1$. This knowledge is important to stop the tuning process, avoiding unnecessary iterations.

So far the only method available to detect local optimality is the trivial pair of conditions: null gradient with positive definite Hessian matrix (Kammer, 1998). These quantities are an intrinsic part of our iterative scheme, which contrasts with the one in Hjalmarsson et al. (1998). That scheme uses a positive definite approximation of the Hessian matrix and therefore misses the ability to check for local optimality.

4. Rehearsal

The scheme in Hjalmarsson et al. (1998), known as IFT for iterative feedback tuning, has succeeded in improving




Fig. 1. Tuning towards instability.

the performance of several practical applications where low complexity controller structures are tuned according to an LQ criterion. There is, however, a strong issue that has not been addressed so far: sensitivity to overparametrisation, common to all gradient-based tuning methods like IFT, FDT and the approach in Trulsson and Ljung (1985).

Consider, for instance, the process

$$y_t = \frac{0.1(q-0.2)}{(q-0.7)(q-0.5)}u_t + \frac{(q-0.6)(q-0.3)}{(q-0.7)(q-0.5)}e_t,$$

with $\sigma_e^2 = 1$, under the feedback control action

$$u_t = -\frac{0.51q^2 - 0.32q}{q^2 - 1.1q + 0.15}y_t.$$

This stable closed-loop system is tuned according to a linear quadratic criterion with $\lambda = 0.1$ and $F_y \equiv F_u \equiv 1$.

With the true values of ∇J , \mathbf{H} and $J^q(\rho)$, and the tuning direction computed with a positive definite approximation of \mathbf{H} , the line search occurs as depicted in Fig. 1. The cost function is given by the solid line, while the second order approximation of this function is the dotted line. Observe that if α is taken at intervals of 0.02 units, the loop suddenly becomes unstable at $\alpha = 0.18$.

Despite the complexity of the control law being the same as the absolute optimal one, a pole-zero pair tends to cancel near the unit circle, causing the instability. The absolute optimal performance has $J^q = 1.1094$, but under the control law

$$u_t = -\frac{0.4483q}{q - 0.1705}y_t,$$

$J^q(\rho) = 1.1095$, leaving almost no improvements to be obtained by the complementary pole-zero pair.

Such a problem is accentuated as the complexity of the controller increases above the necessary one. The extra pole(s) and zero(es) might tend to cancel each other near the unit circle, causing the closed-loop system to be dangerously close to instability. This effect had been observed previously by Deistler, Dunsmuir and Hannan (1978), but in a different context: estimation of ARMA models.

This example shows that maintenance of closed-loop stability is a strong constraint that should be imposed on

the task of tuning. Although it is always possible to step back to the previous controller parameters, occurrences of instability on the way to the optimal control law tend to reduce confidence in the method being used. It is intuitive to think that the cost function increases as the closed-loop system approaches regions of instability, but that is not always the case, at least not to the degree of tuning discretisation one would expect in practical situations.

The following stability assertions were originally presented in Vinnicombe (1993), but with the plant and the controller interchanged. For a stable closed-loop system, the generalised stability margin is defined as

$$b_{P,C_{fb}} \triangleq \left\| \begin{array}{cc} 1 & C_{fb} \\ \frac{1 + PC_{fb}}{P} & \frac{1 + PC_{fb}}{PC_{fb}} \end{array} \right\|^{-1}_{\infty} \quad (12)$$

and as 0 for unstable systems. Moreover, the *Vinnicombe distance* between two feedback controllers, C_0 ($= C_{fb}(\rho_0)$) and C_1 ($= C_{fb}(\rho_1)$), is defined as

$$\delta_v(C_0, C_1) \triangleq \|(1 + C_1 C_0^*)^{-1/2} (C_1 - C_0) (1 + C_0 C_0^*)^{-1/2}\|_{\infty}, \quad (13)$$

provided the following condition (14), is satisfied, and $\delta_v = 1$ otherwise.

$$\delta_v(C_0, C_1) < 1 \Leftrightarrow \begin{cases} \|1 + C_1^* C_0\| \neq 0 \quad \forall \omega & \text{and} \\ \text{wno}(1 + C_1^* C_0) + \eta(C_0) - \eta(C_1) = 0, \end{cases} \quad (14)$$

where $\text{wno}(\cdot)$ denotes the winding number, or number of counterclockwise encirclements of the origin of the Nyquist plot, and $\eta(\cdot)$ denotes the number of poles outside the unit circle.

Vinnicombe (1993) has shown the following:

Given a plant P and a feedback controller C_0 such that the closed-loop system is stable, then the loop remains stable under any feedback controller, C_1 , satisfying $\delta_v(C_0, C_1) < b_{P,C_0}$.

This result is rather conservative because it gives a sufficient, but not necessary, condition for the stability of the closed-loop system.

The insertion of Vinnicombe’s result into our scheme is straightforward. The generalised stability margin is computed as

$$b_{P,C_{fb}(\rho_0)} = \left\| \begin{array}{cc} \frac{\phi_{ur}(\rho_0)}{C_{ff}(\rho_0)\phi_r} & \frac{\phi_{ur}(\rho_0)C_{fb}(\rho_0)}{C_{ff}(\rho_0)\phi_r} \\ \frac{\phi_{yr}(\rho_0)}{C_{ff}(\rho_0)\phi_r} & \frac{\phi_{yr}(\rho_0)C_{fb}(\rho_0)}{C_{ff}(\rho_0)\phi_r} \end{array} \right\|^{-1}_{\infty}. \quad (15)$$

For each new set of controller parameters, given by (10), Condition (14) is tested, followed by the inequality $\delta_v(C_{fb}(\rho_0), C_{fb}(\rho)) < b_{P,C_{fb}(\rho_0)}$. If any of these tests fails,




Fig. 2. Distance between feedback controllers.

the controller $C_{fb}(\rho)$ is not guaranteed to stabilise the closed-loop system. Then the user has the options of applying the largest α in (10) that satisfies the tests; or performing a new experiment with the current controller, for which a new direction of search and a new stability margin are obtained.

Consider the example above, but now let us augment the tuning scheme with the generalised stability margin. The initial control law results in $b_{P,C_{fb}} = 0.680$. Fig. 2 shows the Vinnicombe distance from the current feedback control law to the initial one, as α increases.

The line search should stop at $\alpha \approx 0.158$ because further controllers are not guaranteed to stabilise the closed-loop system, according to Vinnicombe’s criterion. At that point the tuning proceeds with the computation of a new direction and its corresponding stability margin.

Here we have reached a new parameter that can contribute to the choice of step sizes for α (recall Section 3). Since δ_v is actually a metric, $\delta_v(\alpha)$ is a monotonically increasing function. Therefore, it is straightforward to obtain, via an iterative algorithm, an approximate value of α that corresponds to a given Vinnicombe distance. This means that the line search in (10) could, for instance, be performed on a set of equispaced values of δ_v in $[0, b_{P,C}]$.

The main advantage of this approach is to relate variations in the controller parameters to possible variations in the closed-loop behaviour. This is especially critical when the search direction has negative curvature, implying the absence of predictions about the location of a point of minimum cost.

5. Première

The main ideas presented in this paper are illustrated by the following example. A linear unstable plant, described by

$$y_t = \frac{0.1}{q(q - 1.1)(q^2 - 1.4q + 0.53)} u_t + \frac{q - 0.2}{q - 0.99} e_t,$$

with $\sigma_e^2 = 0.2$, is put under control of a PID controller, in a computer simulation environment. The controller has the three usual parameters to be adjusted: K_p , T_i and T_d for the proportional gain, integral time and derivative time, respectively. Its structure is given by

$$u_t = K_p \frac{q - b_1}{q - 1} r_t - K_p \frac{c_1 q^2 + c_2 q + c_3}{q^2 - (1 + a_1)q + a_1} y_t$$

with

$$a_1 = \frac{T_d}{G_d t_s + T_d}, \quad b_1 = 1 - \frac{t_s}{T_i}$$

$$c_1 = 1 + G_d a_1, \quad c_2 = \frac{t_s}{T_i} - 1 - a_1 - 2G_d a_1,$$

$$c_3 = a_1 \left(1 + G_d - \frac{t_s}{T_i} \right),$$

where t_s , the sampling time, is chosen to be 1 time unit, therefore T_i and T_d are directly given in time units. The variable G_d , known as the derivative filtering parameter, can easily be included in the parameter set, but for this experiment it will be fixed at $G_d = 20$.

The tuning criterion is

$$\hat{J}^{lq} = \frac{1}{N_t} \sum_{k=1}^{N_t} [(F_y y_k)^2 + \lambda (F_u u_k)^2]$$

with $F_y = 1$; $F_u = (q - 1)/q$ due to the integrator in the control law; $\lambda = 1$ and $N_t = 8192$ samples collected during normal operation. The number of samples taken in the computation of \hat{J}^{lq} represents a compromise between tuning speed and accuracy, and it can be made variable, with increasing number of samples as the controller approaches optimal performance.

Each excitation of the loop is made with a sequence of normal operating output data (1024 samples in this example) multiplied by 2, so as to amplify the signal-to-noise ratio. The spectral distribution of normal operating data is usually a good choice for this purpose because the energy of the excitation tends to be concentrated in those frequencies where the noise is higher — an inherent feature of the IFT scheme.

The initial controller parameters correspond to the minimum variance case: $\{K_p: 0.613; T_i: 20.8; T_d: 4.06\}$. Under this control law, the variance of the output signal equals 0.4669, but, since the control action is not penalised, the expected value of $(F_u u_t)^2$ is 0.5227. The optimal ($\lambda = 1$) parameters for this controller structure are $\{K_p: 0.610; T_i: 37.1; T_d: 2.33\}$, which would result in the output variance being 0.5759, the expected value of $(F_u u_t)^2$ being 0.2445 and $J^{lq} = 0.8204$. Due to the finite number of samples taken from the loop, any tuning scheme can only be expected to get near this optimal solution.

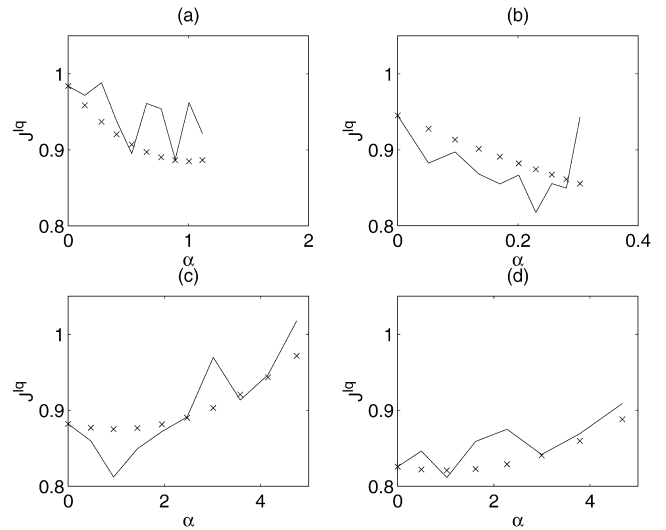


Fig. 3. Cost function along the FDT directions.

The spectral analyses are chosen to be performed on 128 equispaced discrete frequencies. The decisions and tunings performed after each experiment are described below. Measurements of the cost function along the tuning directions are presented in Fig. 3, in solid lines, as well as the anticipated cost function (11), in crosses. The values of α are determined by a set of 10 equispaced values of δ_v in the range $[0, b_{P,C})$.

For the first experiment, the Hessian matrix is found to contain one negative eigenvalue, therefore the direction of Steepest Descent is used. Fig. 3a shows the behaviour of \hat{J}^{lq} as the controller is tuned along that direction. The new parameter set becomes $\{K_p: 0.511; T_i: 20.8; T_d: 4.05\}$.

The following experiments present a positive definite Hessian matrix, implying the use of the Newton method for computing the tuning directions. These line searches are presented in Figs. 3b–d. The tuning finishes with the following parameters for the controller: $\{K_p: 0.619; T_i: 27.5; T_d: 2.58\}$. Further improvements require an increased accuracy in the measurement of \hat{J}^{lq} .

6. Closing the curtains

This paper introduces a model-free tuning scheme based on frequency-domain properties of the closed-loop system's signals. Both first and second derivatives of the cost function, with respect to the controller parameters, are computable from those frequency-domain functions. The actual computation of the cost function is performed directly over time-domain data collected from the loop.

The importance of being able to estimate the Hessian matrix is highlighted throughout the paper; this quantity allows a tight control over the computation of tuning directions, and also provides means for monitoring system optimality. Additionally, the frequency-domain properties are used to compute stability margins and ensure closed-loop stability during the tuning procedure.

References

- Deistler, M., Dunsmuir, W., & Hannan, E. J. (1978). Vector linear time series models: Corrections and extensions. *Advances in Applied Probability*, 10, 360–372.
- Fletcher, R. (1987). *Practical Methods of Optimization* (2nd ed.). Chichester, UK: Wiley.
- Hjalmarsson, H., Gevers, M., Gunnarsson, S., & Lequin, O. (1998). Iterative feedback tuning: Theory and applications. *IEEE Control Systems Magazine*, 18(4), 26–41.
- Kammer, L. C. (1998). *Performance monitoring and assessment of optimality under a linear quadratic criterion*. Ph.D. thesis, Australian National University.
- Lee, W. S., Anderson, B. D. O., Kosut, R. L., & Mareels, I. M. Y. (1993). A new approach to adaptive robust control. *International Journal of Adaptive Control and Signal Processing*, 7, 183–211.
- Mareels, I. M. Y., & Bitmead, R. R. (1986). Nonlinear dynamics in adaptive control: Chaotic and periodic stabilization. *Automatica*, 22(6), 641–655.
- Schrama, R. J. P. (1992). Accurate identification for control: The necessity of an iterative scheme. *IEEE Transactions on Automatic Control*, AC-37(7), 991–994.
- Trulsson, E., & Ljung, L. (1985). Adaptive control based on explicit criterion minimization. *Automatica*, 21(4), 385–399.
- Vinnicombe, G. (1993). Frequency domain uncertainty and the graph topology. *IEEE Transactions on Automatic Control*, AC-38(9), 1371–1382.
- Zang, Z., Bitmead, R. R., & Gevers, M. (1995). Iterative weighted least squares identification and weighted LQG control design. *Automatica*, 31(11), 1577–1594.



Leonardo C. Kammer obtained his Ph.D. in 1998 from the Australian National University. He then worked as a lecturer in the Departamento de Automação e Sistemas at the Universidade Federal de Santa Catarina, Brazil. He has been a Postdoctoral Fellow at the Pulp and Paper Centre since September 1999, where he is conducting research activities in Cross-Directional Control and Performance Assessment and Monitoring.



Robert Bitmead received his Ph.D. from the University of Newcastle in 1980. He has been with the Department of Mechanical and Aerospace Engineering at the University of California, San Diego since July 1999, as part of their new group in Dynamical Systems and Control. Prior to this he was with the Australian National University, Department of Systems Engineering. He is a Fellow of the IEEE and a member of the IFAC Council. He hails originally from Sydney and is General Chair of the IEEE Conference on Decision and Control there from December 12–15, 2000.



Peter Bartlett received the B.Sc. degree in computer science in 1987, the B.E. (Hons I) degree in electrical engineering and the University Medal in 1988, and the Ph.D. in 1992, all from the University of Queensland. He was a lecturer in the Department of Electrical and Computer Engineering at the University of Queensland in 1992. In February 1993 he was appointed as a Research Fellow at the Research School of Information Sciences and Engineering in the Institute for Advanced Studies at the Australian National University, where he is currently Senior Fellow. He is the co-author, with Martin Anthony, of the book “Learning in Neural Networks: Theoretical Foundations”. His research interests include machine learning, computational learning theory, neural networks, and adaptive control.