

Optimal prefiltering in Iterative Feedback Tuning *

R. Hildebrand[†], A. Lecchini[‡], G. Solari^{*} and M. Gevers^{*}

November 2, 2004

[†]Laboratoire de Modélisation et Calcul (LMC)

Université Joseph Fourier (UJF) Grenoble, France

`roland.hildebrand@imag.fr`

[‡]Department of Engineering, University of Cambridge, UK

`al394@cam.ac.uk`

^{*}Center for Systems Engineering and Applied Mechanics (CESAME)

Université Catholique de Louvain, B-1348 Louvain-la-Neuve, Belgium

`{solari, gevers}@csam.ucl.ac.be`

Abstract

Iterative Feedback Tuning (IFT) is a data-based method for the iterative tuning of restricted complexity controllers. A “special experiment” in which a batch of previously collected output data is fed back at the reference input allows one to compute an unbiased estimate of the gradient of the control performance criterion. We show that, by performing an optimal filtering of the data that are fed back, one can minimize the asymptotic variability of the control performance cost, and hence minimize the average performance degradation that results from the randomness of the data. The expression of the optimal filter is derived, and a simulation illustrates the benefits that result from using this optimal filter as compared to the use of the classical constant filter.

*Paper supported by the the Belgian Programme on Interuniversity Attraction Poles, initiated by the Belgian Federal Science Policy Office and the European Research Network on System Identification (ERNSI) funded by the European Union. The scientific responsibility rests with its authors.

1 Introduction

Iterative Feedback Tuning (IFT) is a data-based iterative scheme to optimize the parameter of a feedback controller according to an LQG cost function. At each iteration an estimate of the gradient of the cost function is constructed from a finite set of data obtained partly from the normal operating condition of the closed loop system and partly from a special experiment in which the output of the plant is fed back in the reference signal of the closed loop. The IFT algorithm converges to an optimal controller as the number of iterations tends to infinity [4]. In a separate contribution we derived an analytical expression for the asymptotic covariance of the controller parameter vector, as well as that of the gradient estimate required for the IFT iterations [3]. The result was derived for the Iterative Feedback Tuning of a disturbance rejection controller.

In this paper we use the expressions derived in [3] to maximize the average asymptotic accuracy of the achieved IFT cost by choice of a prefilter for the reference signal in the special experiment. The prefilter affects the signal to noise ratio in this experiment; it was left as a degree of freedom in the original formulation of IFT. An analytical expression for the optimal prefilter is given. It depends on certain characteristics of the unknown process. However, in the spirit of IFT, these characteristics can be estimated from data collected under normal operating conditions, which are available in large quantities. Thus the computation of the optimal prefilter does not necessitate any special experiment on the process and hence does not impose any additional cost. We also present a Monte Carlo simulation example which confirms the asymptotic covariance expression derived in [3] and shows the benefits obtained by the use of the prefilter.

The paper is organized as follows. In the next section we briefly recall the IFT algorithm for disturbance rejection and results from [3] that we rely upon. This enables us in Section 3 to establish a design criterion for the optimization of the prefilter and, accordingly, derive the optimal prefilter. In Section 4 we present the simulation example. Conclusions are given in the last section.




Figure 1: The control system under normal operating conditions.

2 IFT for disturbance rejection

We assume that the plant to be controlled is a SISO linear time-invariant system; its transfer function is denoted by $G(q)$. The output of the plant is affected by an additive stochastic disturbance $v(t)$ assumed to be quasistationary with zero mean and spectral density $\Phi_v(\omega)$. The transfer function $G(q)$ and the disturbance spectrum $\Phi_v(\omega)$ are unknown. We consider the closed loop system depicted in Figure 1, where $C(q, \rho)$ belongs to a parametrized set of controllers with parameter vector $\rho \in \mathbf{R}^n$. The transfer function from $v(t)$ to $y(t, \rho)$ is denoted by $S(q, \rho)$. In this paper, we consider the situation where the control objective is to perform disturbance rejection only. Thus, normal operating conditions are defined as those for which the reference signal $r(t)$ is set at zero. The goal is to find a minimizer for the cost function

$$J(\rho) = \frac{1}{2} \mathbf{E} [y(t, \rho)^2 + \lambda u(t, \rho)^2] , \quad (1)$$

where $\lambda \geq 0$ is a penalty on the control effort chosen by the user.

The IFT method is an iterative procedure that gives a solution to this problem. It is based on the construction, at each iteration step, of an unbiased estimate of the gradient of $J(\rho)$ from data collected on the plant with some present controller $C(\rho)$ in operation. The cost function $J(\rho)$ is minimized with an iterative stochastic gradient descent scheme of Robbins-Monro type [1]. After each iteration, the controller parameter vector is updated ($\rho_n \rightarrow \rho_{n+1}$), and new data are collected on the closed loop system with this updated controller in order to estimate the next gradient. Under some suitable assumptions [2][4], the sequence of controllers converges to a local minimum of $J(\rho)$.

The IFT parameter update rule is given by

$$\rho_{n+1} = \rho_n - \gamma_n R_n^{-1} \text{est}_N \left[\frac{\partial J}{\partial \rho}(\rho_n) \right] \quad (2)$$

where γ_n is a positive step size and R_n is a positive definite matrix. The reader is referred to [3] [4] for details of the algorithm to construct the gradient estimate. Here it suffices to recall that, in order to construct $\text{est}_N \left[\frac{\partial J}{\partial \rho}(\rho_n) \right]$, first a batch $\{u^1(t, \rho_n), y^1(t, \rho_n)\}_{t=1, \dots, N}$ of N data is collected with the controller $C(q, \rho_n)$ in the loop under normal operating conditions, i.e. with $r(t) = 0$. Then, this batch of data is used to construct the signal sequence $\{r(t) = -K_n(q)y^1(t, \rho_n)\}_{t=1, \dots, N}$ which is applied to the reference input of the system (see Figure 1); this is the “special experiment”, which deviates from normal operating conditions. It produces a second batch of input and output data $\{u^2(t, \rho_n), y^2(t, \rho_n)\}_{t=1, \dots, N}$, which are used to construct the following estimates of the gradients of $u^1(t, \rho_n)$ and $y^1(t, \rho_n)$:

$$\begin{aligned} \text{est} \left[\frac{\partial u^1}{\partial \rho}(t, \rho_n) \right] &= \frac{1}{K_n(q)} \frac{\partial C}{\partial \rho}(q, \rho_n) u^2(t, \rho_n) , \\ \text{est} \left[\frac{\partial y^1}{\partial \rho}(t, \rho_n) \right] &= \frac{1}{K_n(q)} \frac{\partial C}{\partial \rho}(q, \rho_n) y^2(t, \rho_n) . \end{aligned}$$

An estimate of the gradient of $J(\rho)$ at ρ_n is then obtained as

$$\text{est}_N \left[\frac{\partial J}{\partial \rho}(\rho_n) \right] = \frac{1}{N} \sum_{t=1}^N \left[y^1(t, \rho_n) \text{est} \left[\frac{\partial y^1}{\partial \rho}(t, \rho_n) \right] + \lambda u^1(t, \rho_n) \text{est} \left[\frac{\partial u^1}{\partial \rho}(t, \rho_n) \right] \right] . \quad (3)$$

The choice of the prefilter $K_n(q)$ is a degree of freedom of the algorithm and affects the signal to noise ratio in the special experiment.

The sequence of parameters ρ_n converges to a local minimum of $J(\rho)$. In [3], under some suitable technical assumptions, the following characterization of the asymptotic accuracy of the method was established. Let $\bar{\rho}$ denote the local minimizer of $J(\rho)$ to which the sequence ρ_n converges and $H(\bar{\rho})$ denote the Hessian of $J(\rho)$ around $\bar{\rho}$. Let the sequence of step sizes be given by $\gamma_n = \frac{a}{n}$ where a is a positive constant, let the sequence R_n converge to a constant positive definite R , and let the covariance matrix $\mathbf{Cov} \left[\text{est}_N \left[\frac{\partial J}{\partial \rho}(\bar{\rho}) \right] \right]$ be positive definite.

Then, the sequence of random variables $\sqrt{n}(\rho_n - \bar{\rho})$ converges in distribution to a normally

distributed zero mean random variable with covariance matrix

$$\Sigma = a^2 \int_0^\infty e^{At} R^{-1} \mathbf{Cov} \left[est_N \left[\frac{\partial J}{\partial \rho}(\bar{\rho}) \right] \right] R^{-1} e^{A^T t} dt, \quad (4)$$

i.e. $\sqrt{n}(\rho_n - \bar{\rho}) \xrightarrow{D} \mathcal{N}(0, \Sigma)$, where $A = \frac{1}{2}I - aR^{-1}H(\bar{\rho})$.

The asymptotic accuracy of IFT thus depends on the covariance of the gradient estimate. This quantity in turn can be influenced by the prefilter $K_n(q)$. In [3] an asymptotic, with respect to the number of data N , expression for this covariance was derived. It is given by

$$\lim_{N \rightarrow \infty} N \mathbf{Cov} \left[est_N \left[\frac{\partial J}{\partial \rho}(\rho_n) \right] \right] = \Xi[K_n] + \Pi \quad (5)$$

where

$$\Xi[K_n] = \frac{1}{2\pi} \int_{-\pi}^{\pi} \frac{|S(e^{j\omega}, \rho_n)|^4 \Phi_v^2(\omega)}{|K_n(e^{j\omega})|^2} [1 + \lambda |C(e^{j\omega}, \rho_n)|^2]^2 \frac{\partial C}{\partial \rho}(e^{j\omega}, \rho_n) \frac{\partial C^*}{\partial \rho}(e^{j\omega}, \rho_n) d\omega. \quad (6)$$

and Π is a constant matrix which does not depend on the prefilter $K_n(q)$. In [3] a frequency domain expression for Π was also derived. However, for the objective of designing the prefilter the expression of Π is not specifically of interest.

3 Design of the optimal prefilter

In this section we propose a criterion for the design of the prefilter $K_n(q)$ and we deliver the expressions of the corresponding optimal prefilter. We assume that the current controller is near the optimal one. Therefore, in order to construct a design criterion for the prefilter, one can employ the asymptotic covariance results given in Section 3.

3.1 The design criterion

Let us take $\mathbf{E}[J(\rho_n)] - J(\bar{\rho})$ as a measure of quality of the controller $C(\rho_n)$. Note that this measure is positive by definition of $\bar{\rho}$. Denote $\Delta\bar{\rho}_n = \rho_n - \bar{\rho}$, then, expanding $J(\rho)$ into a Taylor series around $\bar{\rho}$ and retaining only terms up to second order, we obtain $\mathbf{E}[J(\rho_n)] - J(\bar{\rho}) \approx 0.5 \cdot \mathbf{E}[\Delta\bar{\rho}_n^T H(\bar{\rho}) \Delta\bar{\rho}_n]$ (notice that $J(\rho)$ is analytical for $\rho = \bar{\rho}$ since it is the integral of

a function that analitically depends on ρ). In the previous section we saw that $\sqrt{n}\Delta\bar{\rho}_n$ is asymptotically normally distributed with zero mean and covariance Σ given by (4). Let us now adopt the choice $R = H(\bar{\rho})$, i.e. we consider a Gauss-Newton scheme. Then the asymptotic covariance of $\sqrt{n}\Delta\bar{\rho}_n$ becomes

$$\Sigma = \frac{a^2}{2a-1} R^{-1} \mathbf{Cov} \left[est_N \left[\frac{\partial J}{\partial \rho}(\bar{\rho}) \right] \right] [R^{-1}]^T. \quad (7)$$

This yields

$$\lim_{n \rightarrow \infty} n \cdot \mathbf{E} [\Delta\bar{\rho}_n^T H(\bar{\rho}) \Delta\bar{\rho}_n] = \frac{a^2}{2a-1} \text{Trace} \left[\mathbf{Cov} \left[est_N \left[\frac{\partial J}{\partial \rho}(\bar{\rho}) \right] \right] [R^{-1}]^T \right]. \quad (8)$$

We shall take this expression as the criterion to be minimized for the design of the optimal prefilter.

There are different methods to satisfy the condition $R = H(\bar{\rho})$. A classical method to obtain a sequence of estimated matrices R_n which converges to the Hessian is to fit a regression model using the gradient estimates obtained in the previous iterations; see [7, 8]. In practice, in order to use (8) as a criterion for the design of the optimal filter, at the iteration n one has to replace the optimal parameter $\bar{\rho}$ on the right-hand side by the current parameter ρ_n , since the prefilter is estimated from data obtained under the current operating conditions (as will be shown in the next subsection). In the same way, R must be replaced by the current estimate of the Hessian R_n . These approximations are reasonable, since the procedure described in the present note is for the case where the current controller is near the optimal controller. It aims at achieving optimal asymptotic accuracy of the expected control performance cost.

3.2 The optimal prefilter

The optimal prefilter at iteration n is the filter K_n that minimizes the criterion (8) with $\bar{\rho}$ replaced by ρ_n : it minimizes the weighted trace of the covariance of the gradient estimate. In order to obtain a bounded solution we need to restrict the gain of the prefilter. A straightforward constraint is a bound on the energy of the reference signal $r_n^2(t)$, i.e. on the input of the second experiment. This bound represents the level of acceptable perturbation to the normal operating

conditions during the second experiment . We thus arrive at the following optimization problem:

$$K_n^{opt} = \arg \min_K \text{Trace} \left\{ R_n^{-1} \mathbf{Cov} \left[est_N \left[\frac{\partial J}{\partial \rho}(\rho_n) \right] \right] \right\} \quad \text{subject to} \quad \mathbf{Var} [r_n^2(t)] \leq \alpha,$$

where α is selected by the user. By using (5), and recalling that Π does not depend on the prefilter, we can rewrite the optimal prefilter design problem as follows:

$$K_n^{opt} = \arg \min_K \text{Trace} \{ R_n^{-1} \Xi(K_n) \} \quad \text{subject to} \quad \mathbf{Var} [r_n^2(t)] \leq \alpha. \quad (9)$$

The explicit solution of this problem is characterized by the following proposition.

Proposition 3.1 *The optimal prefilter solving (9) is given by:*

$$\begin{aligned} |K_n^{opt}(e^{j\omega})|^4 &= \text{const} \cdot |S(e^{j\omega}, \rho_n)|^2 \Phi_v(\omega) [1 + \lambda |C(e^{j\omega}, \rho_n)|^2]^2 \\ &\times \text{Trace} \left\{ R_n^{-1} \frac{\partial C}{\partial \rho}(e^{j\omega}, \rho_n) \frac{\partial C^*}{\partial \rho}(e^{j\omega}, \rho_n) \right\}, \end{aligned} \quad (10)$$

where the constant is determined by the design restriction.

Proof See Appendix. □

In order to compute the optimal prefilter in practice, one needs an estimate of the unknown spectral density $|S(e^{j\omega}, \rho_n)|^2 \Phi_v(\omega)$ of the signal $y(t, \rho_n)$ which is the output of the plant under normal operating conditions, i.e. with zero reference signal. The estimate can be obtained with standard techniques in the time or in the frequency domain [5, 6]. Note that since the data needed to estimate this quantity do not stem from a special experiment they are available in large amounts. In fact periods of normal operating conditions can be interlaced with the IFT special experiments. By assuming these periods to be much longer than the length of the special experiment from which the gradient is estimated, the contribution of the variability in the estimate of $|S(e^{j\omega}, \rho_n)|^2 \Phi_v(\omega)$ to the variability of the gradient estimate can be considered as being negligible. Having an estimate of $|S(e^{j\omega}, \rho_n)|^2 \Phi_v(\omega)$ one can construct the magnitude of the optimal prefilter by calculating the 4-th root of the right-hand side of (10). Then, there exist standard tools to approximate a given magnitude function by a stable minimum phase filter.

4 Simulation example

Consider the system described by

$$G(q) = \frac{q^{-1} - 0.5q^{-2}}{1 - 0.3q^{-1} - 0.28q^{-2}}, \quad \text{with noise spectrum } \Phi_v(\omega) = \left| \frac{1}{1 + 0.9e^{-j\omega}} \right|^2.$$

Let the class of controllers be $C(q, \rho) = \rho^1 + \rho^2 q^{-1}$ and set $\lambda = 0.6$ in (1). The (local) minimizer $\bar{\rho} = [-0.69058 \ 0.33105]$ has been found numerically. The Hessian of $J(\rho)$ at $\bar{\rho}$ is given by

$$H(\bar{\rho}) = \begin{bmatrix} 1.87521 & 0.22642 \\ 0.22642 & 2.44438 \end{bmatrix}.$$

Let us assume that the constraint on the reference signal $r_n^2(t)$ during the IFT procedure is that this signal has to have one half the energy of the output of the first experiment. We quantify the performance improvement obtained with the optimal filter (10) with respect to the constant prefilter when the criterion $\mathbf{E} [\Delta \bar{\rho}_n^T H(\bar{\rho}) \Delta \bar{\rho}_n]$ is used, both filters satisfying the same energy constraint. We consider the IFT procedure with experiment length $N = 512$, step sizes $\gamma_n = 1/n$ and $R_n = H(\bar{\rho})$.

We first calculate the theoretical values of Σ for the two cases by using the numerical value of $H(\bar{\rho})$ and the complete asymptotic expression of Σ derived in [3]. For the constant prefilter we obtain

$$\Sigma = 10^{-2} \cdot \begin{bmatrix} 2.74 & -2.02 \\ -2.02 & 1.71 \end{bmatrix}$$

and correspondingly $\lim_{n \rightarrow \infty} n \cdot \mathbf{E} [\Delta \bar{\rho}_n^T H(\bar{\rho}) \Delta \bar{\rho}_n] = 8.39 \cdot 10^{-2}$. As for the optimal prefilter, we obtain

$$\Sigma = 10^{-2} \cdot \begin{bmatrix} 1.87 & -1.21 \\ -1.21 & 1.16 \end{bmatrix}$$

and $\lim_{n \rightarrow \infty} n \cdot \mathbf{E} [\Delta \bar{\rho}_n^T H(\bar{\rho}) \Delta \bar{\rho}_n] = 5.79 \cdot 10^{-2}$. The improvement in the asymptotic value of $\mathbf{E} [\Delta \bar{\rho}_n^T H(\bar{\rho}) \Delta \bar{\rho}_n]$ between the two cases is 31%.

In order to confirm the validity of these theoretical values, we have performed a Monte-Carlo simulation. The parameter vector ρ_8 has been extracted 1024 times. The extractions have been




Figure 2: the parameter $\bar{\rho}$ (●), the 1024 parameters ρ_8 obtained using the constant filter (×) and the contour lines of $J(\rho)$.




Figure 3: the parameter $\bar{\rho}$ (●), the 1024 parameters ρ_8 obtained using the optimal filter (×) and the contour lines of $J(\rho)$.

performed by starting the IFT procedure at $\rho_0 = \bar{\rho}$, in order to eliminate the transient effect of the initial condition, and then running 8 iterations of the IFT parameter vector update up to the parameter vector ρ_8 . The 1024 parameter vectors obtained this way are shown in Figure 2 for the case of the constant prefilter. The corresponding sampled estimate of Σ is

$$\hat{\Sigma} = 10^{-2} \cdot \begin{bmatrix} 2.64 & -1.94 \\ -1.94 & 1.65 \end{bmatrix}$$

and the corresponding sampled estimate of $8 \cdot \mathbf{E} [\Delta \bar{\rho}_8^T H(\bar{\rho}) \Delta \bar{\rho}_8]$ is $8.38 \cdot 10^{-2}$. The parameter vectors obtained for the case of the optimal prefilter are shown in Figure 3. In this case, the sampled estimate of Σ is

$$\hat{\Sigma} = 10^{-2} \cdot \begin{bmatrix} 1.93 & -1.21 \\ -1.21 & 1.18 \end{bmatrix}$$

and for $8 \cdot \mathbf{E} [\Delta \bar{\rho}_8^T H(\bar{\rho}) \Delta \bar{\rho}_8]$ we obtain $5.98 \cdot 10^{-2}$. The estimated improvement between the two cases hence equals 31.5%, as predicted by the theoretical calculations.

5 Conclusions

In this contribution we have shown how the performance of the IFT algorithm for disturbance rejection can be improved by prefiltering the reference input for the special experiment. The

effect of the prefilter amounts to decreasing the contribution of the process noise in the special experiment to the error in the gradient estimate. Hence the prefilter will be the more effective the bigger the contribution of the process noise in the special experiment is. This is the case if a low energy level of the reference signal for the special experiment is required.

References

- [1] Julius R. Blum. Multidimensional stochastic approximation methods. *Annals of Mathematical Statistics*, 25:737–744, 1954.
- [2] R. Hildebrand, A. Lecchini, G. Solari, and M. Gevers. A convergence theorem for Iterative Feedback Tuning. Technical Report 2003/17, CESAME - UCL, Louvain-la-Neuve, Belgium, 2003.
- [3] R. Hildebrand, A. Lecchini, G. Solari, and M. Gevers. Asymptotic accuracy of Iterative Feedback Tuning. *Submitted to IEEE Trans. on Automatic Control*, 2004.
- [4] H. Hjalmarsson, M. Gevers, S. Gunnarson, and O. Lequin. Iterative Feedback Tuning: theory and applications. *IEEE Control Systems Magazine*, 18(4):26–41, August 1998.
- [5] L. Ljung. *System Identification: Theory for the user*. Prentice Hall, 1999.
- [6] R. Pintelon and J. Schoukens. *System Identification: a frequency domain approach*. IEEE Press, New York, 2001.
- [7] C.Z. Wei. Asymptotic properties of least-squares estimates in stochastic regression models. *The Annals of Statistics*, 13(4):1498–1508, 1985.
- [8] G. Yin. A stopped stochastic approximation algorithm. *Systems and Control Letters*, 11:107–115, 1988.

A Appendix

Proof of Proposition 3.1

We address the following optimization problem. By choice of K_n minimize

$$I_1 = \frac{1}{2\pi} \int_{-\pi}^{\pi} \frac{|S(\rho_n)|^4 \Phi_v^2}{|K_n|^2} [1 + \lambda |C(\rho_n)|^2]^2 \text{Trace} \left\{ R_n^{-1} \frac{\partial C}{\partial \rho}(\rho_n) \frac{\partial C^*}{\partial \rho}(\rho_n) \right\} d\omega$$

subject to

$$I_2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} |K_n|^2 |S(\rho_n)|^2 \Phi_v d\omega \leq \alpha.$$

By the Cauchy-Schwarz inequality we have

$$I_1 I_2 \geq \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} |S(\rho_n)|^3 \Phi_v^{\frac{3}{2}} (1 + \lambda |C(\rho_n)|^2) \sqrt{\text{Trace} \left\{ R_n^{-1} \frac{\partial C}{\partial \rho}(\rho_n) \frac{\partial C^*}{\partial \rho}(\rho_n) \right\}} d\omega \right)^2$$

and equality holds if and only if the integrands in the expressions for I_1 and I_2 are proportional.

It follows that K_n realizes the minimum of I_1 if and only if $I_2 = \alpha$ and there exists a positive constant β such that

$$|K_n|^2 |S(\rho_n)|^2 \Phi_v \equiv \beta \frac{|S(\rho_n)|^4 \Phi_v^2}{|K_n|^2} [1 + \lambda |C(\rho_n)|^2]^2 \text{Trace} \left\{ R_n^{-1} \frac{\partial C}{\partial \rho}(\rho_n) \frac{\partial C^*}{\partial \rho}(\rho_n) \right\}.$$

Proposition 3.1 now follows immediately. □