15th Triennial World Congress, Barcelona, Spain

RANDOMIZED ITERATIVE FEEDBACK TUNING

László Gerencsér * **Zsuzsanna Vágó** **
Håkan Hjalmarsson ***

* *Computer and Automation Institute, Hungarian Academy of Sciences, 13-17 Kende u., Budapest 1111, Hungary*
gerencser@sztaki.hu
** *Computer and Automation Institute, Hungarian Academy of Sciences, 13-17 Kende u., Budapest 1111, Hungary and Pázmány P. Catholic University, Budapest,*
vago@oplab.sztaki.hu
*** *Dept. of Signals, Sensors and Systems The Royal Institute of Technology, S-100 44 Stockholm, Sweden*
hakan.hjalmarsson@s3.kth.se

Abstract. In this contribution we present a controller tuning method which combines some of the advantages of Iterative Feedback Tuning (IFT) with some of the advantages of simultaneous perturbation stochastic approximation (SPSA). In particular this scheme can be shown to converge with geometric rate when a pure gradient search is used and the system is noise free. The number of experiments required to obtain an unbiased estimate of the gradient can be reduced significantly for multi-input multi-output systems. In particular we study the problem where the reference signal is periodic and when the noise is negligible.

Key words. adaptive algorithms, iterative methods

1. INTRODUCTION

Iterative Feedback Tuning (IFT), (Hjalmarsson *et al.*, 1998), is a model free method for tuning, especially, general linear controller structures. The method performs optimization of the controller parameters w.r.t. some design criterion using measurements from the plant collected during a cleverly designed experiment. In standard IFT normal operating conditions are interrupted for ad hoc experiments using specially chosen input signals. An alternative procedure is considered in this paper, assuming that a copy or at least a reasonably good model of the original plant is available. In this case normal operating conditions of the original plant will not be interrupted, and the shadow-plant will be used solely to carry out the gradient experiments. An important feature of IFT is that no modeling of the plant and the disturbances are required. The principle is that the experiments are performed such that an estimate of the gradient of the design criterion is obtained and this estimate is subsequently used to update the parameters of the controller.

During the last years quite some experience has been gained with IFT from, e.g. robust control of a simulation model of a flexible transmission system (Hjalmarsson *et al.*, 1995) and vibration attenuation (Meurers and Veres, 1999). It has also been applied by the chemical multinational S.A. Solvay to tune PID controllers for temperature control in furnaces, in distillation columns, flow control in evaporators etc., see (Hjalmarsson *et al.*, 1998). Common to many of the processes in these applications is that they exhibit some kind of nonlinear behavior and even if IFT was developed for linear time-invariant systems it performs well on these systems as well, see (Hjalmarsson, 1998) and (Sj¨oberg and Bruyne, 1999) for an analysis.

For a single-input/single-output (SISO) system, a maximum of three experiments are required in each iteration regardless of the number of parameters in the controller. For a multi-input/multi-output (MIMO) system the situation is quite different, however. A MIMO system using a controller having n_w inputs and n_u outputs requires $1 + n_u \times n_w$ experiments. Assuming that the experiments are carried out on copies of the original plant, this clearly limits the applicability of the

method for multivariable systems: either we need a large number of copies or we have to use small fragments of time devoted to compute different directional derivatives.

A possible way of reducing the number of experiments is to use a randomization technique that has been first presented in connection with simultaneous perturbation stochastic approximation (SPSA), see (Spall, 1992). The latter method has been developed for function minimization problems where the evaluation of function-values is expensive. In fact, in each iteration only two function values are needed. Recently, it has been shown, (Gerencsér and Vágó, 2001), that the SPSA estimator sequence converges with geometric rate almost surely when the step-size in the algorithm is fix and the function evaluation is noise-free.

Using these ideas we get a randomized IFT where an estimate of gradient is obtained from a physical experiment that gives an approximate value of the directional derivative in a *random direction*. Similar to SPSA, the randomness of this direction implies that an approximately unbiased estimate of the complete gradient can be obtained from this single experiment regardless of the dimension of the system or the controller. In this contribution the focus is on the problem where the reference signal is periodic and when the noise is negligible. It is shown that in this case randomized IFT method inherits the geometric convergence rate of noise-free SPSA.

2. THE CONTROL DESIGN CRITERION

We consider an unknown true system described by the discrete time LTI MIMO model

$$\begin{pmatrix} y_n \\ w_n \end{pmatrix} = G \begin{pmatrix} r_n \\ v_n \\ u_n \end{pmatrix} \tag{1}$$

where t represents the discrete time instances, G is the (generalized) plant consisting of the true plant and possibly some frequency weighting filters and a reference model (see below) and is represented by a transfer function matrix, $r_n \epsilon \mathbf{R}^{n_r}$ represents external signals such as set-points, reference signals etc, $v_n \epsilon \mathbf{R}^{n_v}$ represents unmeasurable signals such as (process) disturbances and noise, $w_n \epsilon \mathbf{R}^{n_w}$ represents the sensed outputs and $u_n \epsilon \mathbf{R}^{n_u}$ represents the control signals. Furthermore $y_n \epsilon \mathbf{R}^{n_y}$ represents the variables that will be included in the control criterion (e.g. measured outputs and control signals). By proper definition of the generalized plant, these signals can be filtered (frequency weighted) versions of measured signals in the real system. By including a reference model in the generalized plant, some of these signals may also be the difference between outputs of the reference model and real plant outputs. In this paper the focus will be on the case when $v_n = 0$, i.e the system is noise-free, and r_n is periodic.

The system is controlled by the following controller

$$u_n = C(\rho) w_n \tag{2}$$

where $C(\rho)$ is a $n_u \times n_w$ transfer function matrix parametrized by some parameter vector $\rho \epsilon \mathbf{R}^{n_\rho}$.

The feedback system is shown in Figure 1. The external signal $z \epsilon \mathbf{R}^{n_u}$ in this figure will be used later.

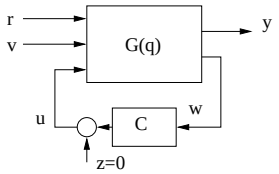

Fig. 1. Feedback system.

Whenever signals are obtained from the closed loop system with the controller $C(\rho)$ operating, we will indicate this by using the ρ-argument; on the other hand, to ease the notation we will from now on, when possible, omit the time argument of the signals. Thus, $y(\rho)$ will denote the signals in the control criterion when the controller (2) is used.

We shall in this paper assume that each channel in the controller is independently parameterized as in Figure 2 where a 2×2 controller is shown and where each block C_{ij} of the controller is parameterized by an independent parameter vector ρ_{ij}.

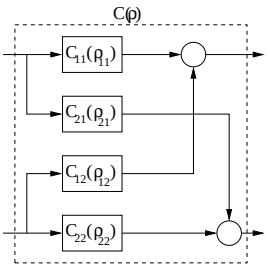

Fig. 2. A controller where each block C_{ij} is independently parameterized.

In the IFT design scheme the following quadratic control criterion is adopted:

$$J(\rho) = \frac{1}{2N} \mathrm{E} \left[\sum_{t=1}^{N} [y_n(\rho)]^T W_n^y y_n(\rho) \right] \tag{3}$$

where $\mathrm{E}[\cdot]$ denotes expectation w.r.t. the disturbance v. The weight matrix W_n^y can be used to minimize the settling time at setpoint changes (Hjalmarsson *et al.*, 1998). This criterion include many commonly used control criteria. In particular it can, with proper definition of the generalized plant handle model reference following.

The optimal controller parameter ρ^* is defined by

$$\rho^* = \arg\min_\rho J(\rho), \tag{4}$$

To carry out the minimization it is necessary to have an expression for the gradient of this criterion with respect to the controller parameters. The novel contribution of the IFT approach (Hjalmarsson *et al.*, 1994) was to show that, in contrast to previous approaches, for SISO systems an approximation of the gradient can be obtained entirely from input-output data collected on the actual closed loop system, by performing one special experiment on the system.

3. ITERATIVE FEEDBACK TUNING

We now address the minimization of $J(\rho)$ given by (3) with respect to the controller parameter vector ρ for a controller of prespecified structure. To facilitate the notation we shall from now on assume that W_n^y equals the identity matrix which we denote by I.

In general the problem of minimizing $J(\rho)$ is not convex and one has to content with a local minimum or even a stationary point. To obtain a stationary point of $J(\rho)$ we would like to find a solution for ρ to the equation

$$0 = \frac{\partial J}{\partial \rho}(\rho) = \frac{1}{N}\mathrm{E}\sum_{t=1}^{N}\left[\frac{\partial y_n}{\partial \rho}(\rho)\right]^T y_n(\rho) \qquad (5)$$

If the gradient $\frac{\partial J}{\partial \rho}$ could be computed, then the solution of (5) would be obtained by the following iterative algorithm:

$$\rho^{k+1} = \rho^k - \gamma_k R_k^{-1}\frac{\partial J}{\partial \rho}(\rho^k), \quad k = 1, 2, \ldots . \qquad (6)$$

Here R_k is some appropriate positive definite matrix, typically a Gauss-Newton approximation of the Hessian of J, while γ_k is a positive real scalar that determines the step-size.

This procedure is not directly implementable since it requires the knowledge of the plant to get the gradient. Equivalently, we need to get an approximate value of

$$\left[\frac{\partial y}{\partial \rho}(\rho^k)\right]^T y(\rho^k).$$

The computation of this quantity has always been the key stumbling block in solving this direct optimal controller parameter tuning problem. The key contribution in the original derivation of IFT (Hjalmarsson *et al.*, 1994) was to show that for single input/single output (SISO) systems these quantities can indeed be obtained by performing experiments on the closed loop system formed by the actual system in feedback with the controller $C(\rho)$.

3.1 *The gradient*

Let us now consider how to obtain the derivative of y w.r.t. the parameter vector ρ_{ij} which parameterizes the ijth block C_{ij} of the controller. We can write the relation between controller block C_{ij} and an arbitrary output y_l as

$$\begin{pmatrix} y_l \\ w_{ij} \end{pmatrix} = \bar{G}_{ijl}\begin{pmatrix} r \\ v \\ u_{ij} \end{pmatrix} \qquad (7)$$

$$u_{ij} = C_{ij} w_{ij} \qquad (8)$$

where $\bar{G}_{ijl}$ does not depend on ρ_{ij} and is obtained from the interconnection between the generalized plant G and all controller blocks except C_{ij}, see Figure 3. The feedback system in Figure 3 has the structure of a SISO system. Differentiating (7)–(8) w.r.t. an arbitrary element of ρ_{ij} gives ($'$ denotes the derivative)

$$\begin{pmatrix} y_l' \\ w_{ij}' \end{pmatrix} = \bar{G}_{ijl}\begin{pmatrix} 0 \\ 0 \\ u_{ij}' \end{pmatrix} \qquad (9)$$

$$u_{ij}' = C_{ij}' w_{ij} + C_{ij} w_{ij}' \qquad (10)$$

These equations represent the same feedback system as in Figure 3 but with the external input $z_{ij} = C_{ij}' w_{ij}$. Hence, the gradient y_l' is obtained by performing a closed loop experiment with all external signals zero except for $z_{ij} = C_{ij}' w_{ij}$ where w_{ij} is first obtained from a closed loop experiment corresponding to (7)–(8), i.e. with r as reference.

Now, using the linearity of the system, the gradient w.r.t. all parameters in C_{ij} can be obtained from these two experiments. The procedure is as follows: 1. Perform a normal experiment with the desired reference signal r, c.f. Figure 1. Measure w and y. 2. For each controller block C_{ij}, c.f. Figure 2, perform an experiment where $r = 0$ and where w_j is injected at the output of controller block C_{ij}, i.e. $z_{ij} = w_j$ in the block diagram in Figure 3.

By filtering the lth output element, y_l^2, of y from this experiment through the gradient of controller block C_{ij} w.r.t. any element of ρ_{ij} (this is a known scalar transfer function which we denote by C_{ij}'), the gradient of y_l w.r.t. any parameter in the vector ρ_{ij} can be obtained. Hence the gradient for each element of y w.r.t. any parameter in the controller block C_{ij} can be computed from this experiment. To compute the gradient w.r.t. to all possible parameters in a full block controller, $n_u \times n_w$ gradient experiments are necessary with this approach giving a total of $1 + n_u \times n_w$ experiments in order to be able to compute all gradients.

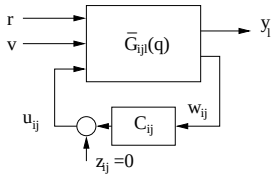

Fig. 3. Feedback system where each channel in the controller is independently parameterized. The generalized plant $\bar{G}_{ijl}$ is independent of the parameter vector ρ_{ij}

4. A MODIFIED SPSA-METHOD

A special procedure for minimizing a cost function the evaluation of which is costly has been developed in (Spall, 1992). In that paper the problem of minimizing a three-times continuously differentiable function $L(\theta)$ was considered, under the condition that the computation of $L(.)$ is expensive. The main idea of the proposed simultaneous perturbation stochastic approximation (SPSA) is to estimate the gradient of $L(.)$ using numerical approximations of directional derivatives in *random directions*. Then, using a simple algebraic operations, we can reconstruct an unbiased estimator of the numerical approximation of the gradient.

SPSA methods have been analyzed under various conditions in (Chen *et al.*, 1999), (Gerencsér, 1999), (Kushner and Yin, 1997) and (Spall, 1992).

Simplified and improved versions have been developed in (Spall, 1997; Spall, 2000). In contrast to the standard setup in SPSA-theory, in IFT the partial or directional derivatives of the cost function *can* be computed experimentally, up to a negligible error, due to "non-stationary initialization". But as we have seen these experiments can be costly for MIMO systems. The novel idea is now to use random directions, as in SPSA, for the gradient estimation to ensure that the number of derivative evaluations is small. In this section we will present this modification of SPSA in a general setup.

Consider the problem of minimizing a three-times continuously differentiable function $L(\theta)$ defined over some p-dimensional bounded open subset D, under the condition that its directional derivatives of are computable exactly via expensive physical experiments. Assume that the gradient L_θ vanishes at single point in D, denoted by θ^*.

We compute directional derivatives in random directions as follows: let θ be a fixed value, and generate random directions by taking a sequence of independent, identically distributed random variables, with time index n, Δ_{ni}, $i = 1, ..., p$. A standard choice is a Bernoulli-sequence, where $\Delta_{ni} = \pm 1$ with equal probability. Introduce the notation $\Delta_n^{-1} = \left[\Delta_{n1}^{-1}, \ldots, \Delta_{np}^{-1}\right]^T$ and define

$$H(n, \theta) = \Delta_n^{-1} \Delta_n^T L_\theta(\theta). \tag{11}$$

It is easy to see that $H(n, \theta)$ is an unbiased estimator of $L_\theta(\theta)$ and. Since $L_\theta(\theta^*) = 0$, we have identically

$$H(n, \theta^*) = 0. \tag{12}$$

The problem of estimating θ^* under the condition $H(n, \theta^*) = 0$ has been considered in a general framework in (Gerencsér and Vágó, 2001). It is easy to see, using the results of (Benveniste *et al.*, 1987) that for a standard stochastic approximation procedure with gain $1/n$ would the asymptotic covariance matrix of the estimator process is 0. Thus the convergence rate is better than the standard rate $n^{-1/2}$. But how much better can it be? A straightforward, but tiresome calculation induces us to consider *fixed gain* stochastic approximation processes of the form

$$\theta_{n+1} = \theta_n - \lambda H(n+1, \theta_n) \qquad \theta_0 = \xi. \tag{13}$$

Fixed gain recursive estimation processes of this general form have been widely used in the engineering literature in connection with the well-known LMS-algorithm of adaptive filtering and the identification of time-varying systems. A *new feature* of the present problem is that the condition $H(n, \theta^*) = 0$ is imposed.

Assuming $\theta^* = 0$ and assuming that H is smooth in θ, an exact linearization around θ^* will lead to the following equivalent form of (13):

$$\theta_{n+1} = \theta_n - \lambda A(n+1, \theta_n)\theta_n. \tag{14}$$

If L is a quadratic function then L_θ is linear and $A(n+1, \theta_n)$ is in fact independent of θ, for which Oseledec's multiplicative ergodic theorem, (Oseledec, 1968; Ragunathan, 1979), is applicable. Thus we get that for any initial condition θ_0 outside of a set of Lebesgue-measure zero

$$\lim_{n \to \infty} \frac{1}{n} \log |\widehat{\theta}_n - \theta^*| = \mu$$

with probability 1. If λ is small enough then (Gerencsér, 1991) implies that $\mu < 0$, thus the convergence rate is geometric.

In the general, non-quadratic case we need to enforce the boundedness of θ_n, say $\theta_n \epsilon D_0$, where $D_0 \subset D$ is some compact domain, by using a resetting mechanism. In (Gerencsér and Vágó, 2001) we presented the surprising result that, in spite of eventual non-linearity in L_θ and resetting, θ_n does converge to θ^* almost surely at a geometric rate. The analysis is based on a discrete-time ODE-method, following (Gerencsér, 1996), in which the θ_n will be compared with the discrete-time deterministic process (z_n) defined by

$$z_{n+1} = z_n + \lambda L_\theta(z_n), \quad z_0 = \xi = \theta \epsilon D_\theta, \tag{15}$$

and thus a random The advantage of the discrete-time ODE-method is that the errors due to conversion from discrete to continuous-time, that show up in standard ODE-methods, do not show up.

The associated continuous-time ODE is defined as

$$\dot{y}_t = \lambda L_\theta(y_t), \qquad y_s = \xi = z \epsilon D_z, \qquad s \geq 0. \tag{16}$$

The solution of (16) will be denoted by $y(t, s, \xi)$.

Condition 4.1. The ordinary differential equation (16) is globally exponentially stable in some compact domain $D_z \subset D$: for some $C_0 > 0$ and $\alpha > 0$ we have for all $0 \leq s \leq t$, $z \epsilon D_z$ $|y(t, s, z)| \leq C_0 e^{-\alpha(t-s)}$. In addition we assume that

$$\|\frac{\partial}{\partial z} y(t, s, z)\| \leq C_0 e^{-\alpha(t-s)}. \tag{17}$$

Resetting: Let D_ξ be a compact domain of possible initial values with interior D_ξ, and let D_θ be a compact truncation domain such that 0ϵ intD_ξ and $S(D_\xi, d') \subset$ intD_θ for some $d' > 0$. Assume that $\xi = \widehat{\theta}_0 \epsilon D_\xi$. At time n we first define a tentative value θ_{n+1-} following (13) as $\theta_{n+1-} = \theta_n - \lambda H(n+1, \theta_n)$, and then we set

$$\begin{aligned} \theta_{n+1} &= \theta_{n+1-} \quad \text{if} \quad \theta_{n+1-} \epsilon D_\theta \\ \theta_{n+1} &= \theta_0 \qquad \text{if} \quad \theta_{n+1-} \notin D_\theta. \end{aligned} \tag{18}$$

Condition 4.2. Let $S(0, r)$ denote the sphere with center in the origin and radius r. It is assumed that for some $0 < r < R$ we have

$$D_\xi \subset S(0, r) \subset S(0, R) \subset D_\theta. \tag{19}$$

Now applying the results of (Gerencsér and Vágó, 2001) we get:

Theorem 4.1. Consider the algorithm

$$\theta_{n+1} = \theta_n - \lambda \Delta_n^{-1} \Delta_n^T L_\theta(\theta) \qquad \theta_0 = \xi. \tag{20}$$

Assume that L is three-times continuously differentiable with bounded derivatives up to second

order in D, and let Conditions 4.1 and 4.2 are satisfied. Let $\xi = \theta_0 \epsilon D_\xi$ and assume that $C_0^3 r/R < 1$. Then, under weak additional conditions on the position of the domain D_z, there exists a γ with $0 < \gamma < 1$ such that for sufficiently small λ we have

$$|\theta_N| \leq C\gamma^{\lambda N}$$

where C is a positive random variable C.

5. RANDOMIZED IFT

We now present a randomized version of the gradient estimation procedure presented in Section 3.1 based on the considerations in the previous section.

Let $v_t = 0$ and (r_t) be a periodic input with period N_0. For any fixed ρ let $x_t^*(\rho)$ and $y_t^*(\rho)$ be the steady-state state and output-processes, respectively. Then the cost function can be written as

$$J(\rho) = \frac{1}{2N_0}\left[\sum_{t=1}^{N_0} [y_n^*(\rho)]^T W_n^y y_n^*(\rho)\right]$$

The cost function and its gradient are computable by physical experiments up to an error due to initial conditions. Let the initial state of the system be ξ and let $\hat{y}_n(\rho, \xi)$ be the corresponding output-process. Let τ be a fixed delay that is introduced to reduce the effect the effect of transients. Define

$$\hat{J}(\rho, \xi) = \frac{1}{2N_0} \sum_{t=\tau+1}^{\tau+N_0} [\hat{y}_n(\rho, \xi)]^T W_n^y \hat{y}_n(\rho, \xi). \quad (21)$$

Then it is easy to see that

$$|\hat{J}(\rho, \xi) - J(\rho)| \leq Ce^{-\alpha\tau}|x^*(\rho) - \xi|$$

with some $\alpha > 0$. A similar estimate holds for the gradient processes.

Let now ρ_n be a sequence of estimators that is used in successive experiments. A key novel assumption is that the system can be *replicated*, and thus one control system is purely used for evaluating the controller performance, while a copy of it is used to evaluate a directional derivative. Let the corresponding state-sequence in the original control system be x_n. Then under mild conditions we have $|x_n - x_n^*(\rho_n)| \leq c|\theta_n - \theta^*|$ with some constant c.

Ultimately we come to consider the perturbed version of (20) where $L_\theta(\theta)$ is replaced by $L_\theta(\theta) + \delta L_{\theta_n}(\theta)$, where the norm of the latter is bounded by $c|\theta_n - \theta^*|$. Note that the perturbation is multiplicative. Taking τ large enough c can be made arbitrarily small, and thus the analysis of (Gerencsér and Vágó, 2001) can be carried over, and the conclusion of Theorem 4.1 remains valid.

6. SIMULATION RESULTS

Most of the simulations have been devoted to test the viability of noise-free SPSA. Testing randomized IFT itself is under progress. The algorithm (20) has been tested for randomly generated functions up to dimension 100. First we considered quadratic functions of the form $L(\theta) = \frac{1}{2}\theta^T A\theta$, with some symmetric positive definite A, which is obtained as follows: first generate a diagonal matrix with diagonal elements chosen according to exponential distribution, and the apply random rotations: a pair of coordinates is chosen randomly, according to uniform distribution, and then a rotation in the selected plane is applied with an angle that is random and has uniform distribution. In each experiment we had $N = 500$ iterations.

We have also considered non-quadratic problems of the form

$$L(\theta) = \frac{1}{2}\theta^T A(\theta)\theta \quad \text{with} \quad A(\theta) = A + u(\theta)u(\theta)^T$$

where $u(\theta) = D\theta$ with some fixed matrix D where the elements of D were chosen uniformly in the range of $[0, c_p]$, where $c_p = 0.1$ was chosen in this experiment. Geometric rate of convergence in this case is ensured by the results of this paper and indeed confirmed by simulation results similar to the quadratic case. We have carried out these experiments with different stepsizes λ, and the top Lyapunov-exponent is approximated by the final value $\log|\theta_N - \theta^*|/N$.

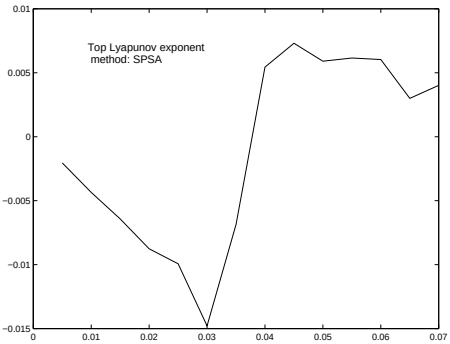

Fig. 4. The optimal top-Lyapunov exponent is $\mu = -0.0148$, SPSA method.

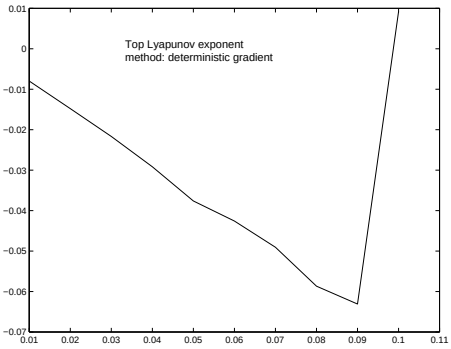

Fig. 5. The optimal top-Lyapunov exponent is $\mu = -0.0631$, gradient method.

In Figure 4. and Figure 5. the approximate top Lyapunov-exponent is plotted against the stepsize λ for an increasingly non-quadratic problem in dimension $p = 20$ for the plain SPSA method and for the gradient method, respectively. A key indicator of the efficiency of the SPSA method is the number of function-value evaluations that is needed to achieve a given accuracy. This is in turn essentially determined by the top-Lyapunov exponent the control of which via the stepsize is therefore a key issue. An adaptive procedure

for choosing the optimal λ has been proposed in (Gerencsér and Vágó, 2001)

Consider now the cases when the SPSA and the gradient method both operate under their respective optimal stepsize for a problem that we described above. Comparing the optimal values of the top Lyapunov exponents it is seen, that the number of experiments to achieve a precision $\varepsilon = 0.01$ is $n = \log(\varepsilon)/\mu = 311$ in the SPSA case and $n = \log(\varepsilon)/\mu \cdot p = 1459$ using gradient method. These results are supported by experimental data. Thus we conclude that SPSA is more economical then similar deterministic methods.

7. CONCLUSION

We have presented a randomized version of Iterative Feedback Tuning (IFT) using ideas of simultaneous perturbation stochastic approximation (SPSA). We have considered noise-free multi-input multi-output systems with periodic reference signal. The directional derivative in a random direction is approximated by a single physical experiment. Under the assumption that the system can be replicated, and the replica is used for the gradient estimation, the proposed method has geometric rate of convergence.

ACKNOWLEDGEMENT

This work was partially supported by the National Research Foundation of Hungary (OTKA) under Grant no. T 032932. The second author gratefully acknowledges the support of the Bolyai János Research Fellowship of the Hungarian Academy of Sciences. The third author gratefully acknowledges the support by MTA SZTAKI's EU Center of Excellence program. The first author expresses his thanks to James C. Spall for cooperating in this research and providing a number of useful comments.

8. REFERENCES

Benveniste, A., M. Métivier and P. Priouret (1987). *Algorithms adaptatifs et approximations stochastiques*. Masson, Paris.

Chen, H.F., T.E. Duncan and B. Pasik-Duncan (1999). A Kiefer-Wolfowitz algorithm with randomized differences. *IEEE Trans. Automat. Contr.* **44**, 442–453.

Gerencsér, L. (1991). Almost sure exponential stability of random linear differential equations. *Stochastics* **36**, 411–416.

Gerencsér, L. (1992). A representation theorem for the error of recursive estimators. In: *Proc. of the 31st IEEE Conference on Decision and Control, Tucson*. pp. 2251–2256.

Gerencsér, L. (1996). On fixed gain recursive estimation processes. *J. of Mathematical Systems, Estimation and Control* **6**, 355–358. Retrieval code for full electronic manuscript: 56854.

Gerencsér, L. (1999). Rate of convergence of moments for a simultaneuous perturbation stochastic approximation method for function minimization. *IEEE Trans. Automat. Contr.* **44**, 894–906.

Gerencsér, L. and Zs. Vágó (2001). A stochastic approximation method for noise free optimization. In *Proceedings of the European Control Conference, ECC'01, Porto*, pages 1496–1500.

Gerencsér, L. and Zs. Vágó (2001). The mathematics of noise-free SPSA. In: *Proceedings of the 40th CDC, 2001*. IEEE.

Hjalmarsson, H. (1998). Control of nonlinear systems using Iterative Feedback Tuning. In: *Proc. 1998 American Control Conference*. Philadelphia. pp. 2083–2087.

Hjalmarsson, H., M. Gevers, S. Gunnarsson and O. Lequin (1998). Iterative Feedback Tuning: theory and applications. *IEEE Control Systems Magazine* **18**(4), 26–41.

Hjalmarsson, H., S. Gunnarsson and M. Gevers (1994). A convergent iterative restricted complexity control design scheme. In: *Proc. 33rd IEEE CDC*. Orlando, Florida. pp. 1735–1740.

Hjalmarsson, H., S. Gunnarsson and M. Gevers (1995). Model-free tuning of a robust regulator for a flexible transmission system. *European Journal of Control* **1**, 148–156.

Kushner, H.J. and G. Yin (1997). *Stochastic Approximation Algorithms and Applications*. Springer Verlag. New York.

Ljung, L. and T. Söderström (1983). *Theory and practice of recursive identification*. The MIT Press.

Meurers, T. and S.M. Veres (1999). Iterative design for vibration attenuation. *International Journal of Acoustics and Vibration* **4**(2), 79–83.

Oseledec, V.I. (1968). A multiplicative ergodic theorem. Lyapunov charasteristic numbers for dynamical systems. *Trans. Moscow Math. Soc.* **19**, 197–231.

Ragunathan, M. S. (1979). A proof of oseledec's multiplicative ergodic theorem. *Israel Journal of Mathematics* **42**, 356–362.

Sjöberg, J. and F. De Bruyne (1999). On a nonlinear controller tuning strategy. In: *14th IFAC World Congress*. Vol. I. Beijing, P.R. China. pp. 343–348.

Spall, J.C. (1992). Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Trans. Automat. Contr.* **37**, 332–341.

Spall, J.C. (1997). A one-measurement form of simultaneous perturbation stochastic approximation. *Automatica* **33**, 109–112.

Spall, J.C. (2000). Adaptive stochastic approximation by the simultaneous perturbation method. *IEEE Trans. Automat. Contr.* **45**, 1839–1853.