Appl. Math. and Comp. Sci., 1995, vol.5, No.3, 425–438

ITERATIVE LEARNING CONTROL — AN OVERVIEW OF RECENT ALGORITHMS

DAVID H. OWENS*, NOTKER AMANN*, ERIC ROGERS**

Recent results in an important area of Iterative Learning Control are reviewed. Iterative Learning Control is a new control technique for repetitive systems where the controller learns from previous experience. The paper includes a short survey of work in this area and shows new results that extend previous work. All control algorithms use feedback control instead of the feedforward control often used in Iterative Learning Control. Several important aspects as robustness of learning, convergence conditions, convergence types and limitations due to plant properties are addressed.

1. Introduction

Iterative Learning Control is a new technique to control systems that constantly repeat the same task. One economically important and often studied example is that of robot manipulators (Arimoto, 1990). These execute one task, e.g. following a geometrical trajectory, repetitively. Using normal control, they exhibit the same performance (performance designates the tracking precision of the desired output signal) at each repetition/trial. Motivated by human learning, the basic idea of Iterative Learning Control is to use information from previous executions of a trial in order to *improve* performance from trial to trial (Arimoto *et al.*, 1984). Mathematically expressed, this means that the control input $u_{k+1}(t)$ at the $(k+1)$-th trial to the plant is given as a function of previous inputs and errors (Arimoto, 1990)

$$u_{k+1} = f(u_k, e_k) \tag{1}$$

where e_k denotes the error at trial k.

The intuitive notion of "improving performance progressively" can be refined to a convergence condition on the error, i.e. (in some norm topology)

$$\lim_{k\to\infty} \|e_k(\cdot)\| = 0 \tag{2}$$

This is a stability problem on a 2D-product space, typically of the form $\mathbb{N} \times L_2[0,T]$. This places the analysis of Iterative Learning Control systems firmly outside of the traditional realm of control theory although it is illustrated below how "classical" control theory is a valuable tool in the area. The rigourous analysis of the behaviour of the combined learning and dynamical system is only now emerging (Rogers and Owens, 1992).

* Centre for Systems and Control Engineering, University of Exeter, Exeter EX4 4QF, U.K.
** Advanced Systems Research Group, University of Southampton, Southampton SO9 5NH, U.K.

Previous work can be divided into two classes. In the larger class, non-linear systems and particularly the specific non-linear class of mechanical systems, as found in robotics, are studied. The methods considered there normally make use of special characteristics of mechanical systems. This paper falls into the other class, where linear systems are considered in their generality. This enables the use of well-known analysis and design methods from "classical" control theory. For example, frequency domain methods (Owens and Neuffer, 1992; Padieu and Su, 1990) allow the derivation of convergence conditions as norm bounds on the operator that relates the error from trial k to trial $k+1$ (the so-called "error transition operator").

In contrast to other papers in this area, this paper studies Iterative Learning Control schemes based on a learning law that takes the error of the current trial into account, as in the functional algorithm

$$u_{k+1} = u_k + f(e_{k+1}) \tag{3}$$

as opposed to the normally used scheme (1) which includes the previous trial error only. A discussion of algorithms of the form (1) can be found in Moore (1993). The use of the previous error in the control law corresponds to "feedforward" (trial to trial) control while the algorithm (3) uses a "feedback" (current trial) control. The use of feedback control enables convergence for a wider class of systems (Owens and Neuffer, 1992) and has besides that the usual advantages of feedback control, e.g. stabilisation of the closed loop and the potential for increased robustness.

The outline of the next sections is as follows: in Section 2 high gain controllers are studied, connecting learning control to adaptive stabilisation (Ilchmann, 1993; Owens *et al.*, 1987). Section 3 takes a look at Iterative Learning Control as a form of 2D-systems theory. In Section 4, frequency domain convergence conditions are reviewed and H_∞-optimal Iterative Learning Control is introduced. In the last section, an Iterative Learning Control algorithm that is optimal with regard to the L_2-norm is shown. The comparison of the algorithms addresses the issues of weak versus strong convergence, limitations of Iterative Learning Control due to plant properties and the form of the limit error.

2. High Gain Feedback

In this section, Iterative Learning Control for plants subject to severe uncertainty is studied. This approach is a generalisation of the concepts of universal adaptive stabilisation (Byrnes and Willems, 1984; Ilchmann, 1992) to Iterative Learning Control and makes use of constant, high gain controllers. It is assumed that the plant, given in state-space as

$$\begin{aligned} \dot{x}_k(t) &= A\,x_k(t) + B\,u_k(t) \\ e_k(t) &= r - C\,x_k(t) \end{aligned} \tag{4}$$

where $x_k(t) \in \mathbb{R}^n, u_k(t) \in \mathbb{R}^m, e_k(t) \in \mathbb{R}^m, r \in \mathbb{R}^m$ and $0 \le t \le T < \infty$, is unknown, but satisfies the following properties:

- it is minimum phase

- CB is square, non-singular and has spectrum in the open right half complex plane.

The learning law considered has the form:

$$u_{k+1}(t) = u_k(t) + Ke_{k+1}(t) \tag{5}$$

Under the above conditions and if the scalar, positive gain K is sufficiently large, it can be shown that the error sequence has the following properties:

1. The error satisfies a *slow change condition*:

$$\sum_{k=0}^{\infty} ||e_{k+1} - e_k||^2_{L_2^m[0,T]} < \infty \tag{6}$$

2. It is a *monotonically* non-increasing sequence:

$$||e_{k+1}||^2_{L_2^m[0,T]} \leq ||e_k||^2_{L_2^m[0,T]} \tag{7}$$

3. $e_{k+1} - e_k$ converges in norm to zero as $k \to \infty$ in $L_2(0,\infty) \cap L_\infty(0,\infty)$.

4. The error sequence $\{e_k\}$ is uniformly bounded and it converges to zero in the *weak* topology on $L_2^m[0,T]$, i.e.

$$\lim_{k\to\infty} \langle f, e_k \rangle = 0 \quad \forall f \in L_2^m[0,T] \tag{8}$$

The proofs are given by Owens (1992).

The difficulty of this scheme for practical applications is to find a large enough gain (which depends upon the system in a complex manner). This problem was eased through a recent extension where the control gains K are adaptively changed (Owens, 1993). At each trial, a gain K_{k+1} is employed. This gain is adapted according to

$$K_{k+1} = K_k + ||e_{k+1} - e_k||^2_{L_2^m[0,T]} \tag{9}$$

Under the same conditions as above, it holds additionally to the above properties of $\{e_k\}$ that the gain sequence $\{K_{k+1}\}$ converges to a limit gain K_∞ if the error sequence $\{e_k\}$ is bounded. Sufficient conditions for boundedness of the error sequence is either that the system is SISO and positive real. Or, boundedness can be guaranteed if the initial choice K_1 is sufficiently large. In these cases, it is easy to show that the learning operator has norm of less than one, connecting this analysis to the one of Section 4. Figure 1 shows a simulation of an Iterative Learning Control system using the adaptive gain update. The (positive real) plant in the Laplace-domain is $g(s) = (3s+1)/(s+1)^2$, while the reference signal is $r(t) = 1 - (1+2t)e^{-2t}$. The initial guess for K is $K_1 = 0$ and T is 6 time units.

The adaptation of K_{k+1} parallels the basic control $\dot{k} = y^2$ introduced by Byrnes and Willems (1984) in adaptive stabilisation theory. One advantage of this method is that the plant may be unknown as long as it belongs to the above defined class.

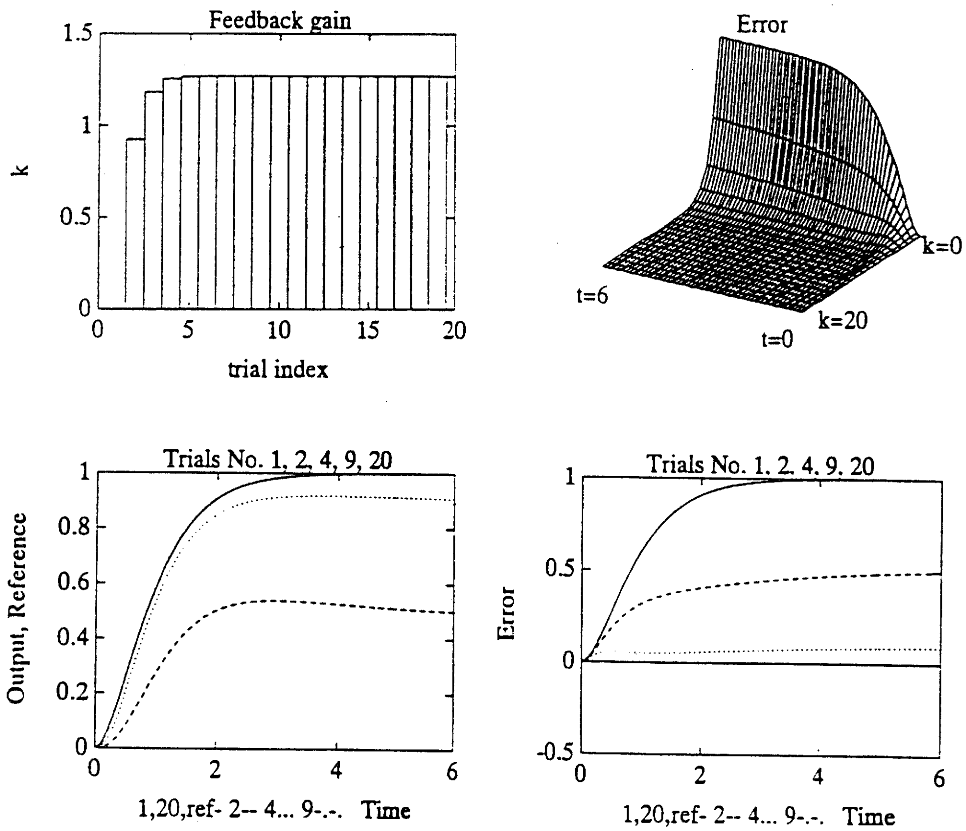

Fig. 1. ILC system using adaptive high gain feedback.

It is however disappointing that only weak convergence is achieved. This means that all Fourier-coefficients of the limit error go to zero, but it could still be the case that the error does not go pointwise to zero as $k \to \infty$.

3. 2D Systems Theory for Iterative Learning Control

Iterative Learning Control under the aspect of 2D-systems theory stresses the repetitive action of the system. The problem of convergence in the k-direction is in this light a stability problem. In (Edwards and Owens, 1982; Rogers and Owens, 1992), a stability theory for repetitive systems is derived. In the following, a general Iterative Learning Control system is transformed into a "linear repetitive system" so that the stability theory of Edwards and Owens (1982) and Rogers and Owens (1992) is applicable.

The learning law of the Iterative Learning Control system considered here uses information from several previous trials, thus making up an n-th order Iterative

Learning Control system. It also incorporates *forgetting factors* α_i so that the inputs of previous trials are weighted. The learning law is thus:

$$u_{k+1}(t) = \sum_{i=1}^{n} \alpha_i u_{k+1-i}(t) + \sum_{i=1}^{n} K_i[e_{k+1-i}](t) + K_0[e_{k+1}](t) \tag{10}$$

The inclusion of several previous inputs and errors brings as advantage that the algorithm is more robust against disturbances that occur only in single trials. To convert this into the general form of a linear repetitive process, the error is considered as the "pass profile" and the reference as disturbance. The system equation for $e_{k+1}(t)$ follows from (10) and the plant equation $e_{k+1}(t) = r(t) - G[u_{k+1}](t)$ as

$$e_{k+1} = (I + GK_0)^{-1} \left\{ \sum_{i=1}^{n} (\alpha_i I - GK_i) e_{k+1-i} + (1 - \sum_{i=1}^{n} \alpha_i) r \right\} \tag{11}$$

Defining a supervector $\hat{e}_{k+1}(t)$ consisting of the errors appearing in (11) is comparable to introducing the (global) state of the 2D-system. The error supervector is given by

$$\hat{e}_k(t) = \begin{bmatrix} e_{k+1-n}(t) \\ \vdots \\ e_{k-1}(t) \\ e_k(t) \end{bmatrix} \tag{12}$$

The notation is further streamlined by introducing the compact notation for the operators $E_i[\cdot](t)$ as $E_0 = (I + GK_0)^{-1}$ and $E_i = \alpha_i I - GK_i$ for $i = 1, \ldots, n$. This yields the final form of the system:

$$\hat{e}_{k+1} = L_T\, \hat{e}_k + b_k \tag{13}$$

with

$$L_T = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ E_0 E_n & E_0 E_{n-1} & E_0 E_{n-2} & \cdots & E_0 E_1 \end{bmatrix} \tag{14}$$

and

$$b_k(t) = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 - \sum_{i=1}^{n} \alpha_i \end{bmatrix} r(t) \tag{15}$$

This is a linear repetitive process with operator L_T acting on the state (the error supervector) $\hat{e}_{k+1}$ and a disturbance $b(t)$ that is the same for all trials k. The stability theory from (Edwards and Owens, 1982; Rogers and Owens, 1992) shows that this system is asymptotically stable (in a precisely defined sense) if and only if the spectral radius of L_T is $\rho(L_T) < 1$. The astonishing result for a finite trial length $T < \infty$ is that the spectral radius $\rho(L_T)$ is given by the maximal eigenvalue of D_{L_T}

$$\rho(L_T) = \max_i \lambda_i(D_{L_T}) \tag{16}$$

where D_{L_T} is the direct feedthrough matrix of the state space realization of the system L_T. If the plant is strictly proper and the controllers are proper, the matrix D_{L_T} is independent of the plant and controllers and is given by

$$D_{L_T} = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ \alpha_n & \alpha_{n-1} & \alpha_{n-2} & \cdots & \alpha_1 \end{bmatrix} \tag{17}$$

and λ_i are the roots of its characteristic polynomial

$$\lambda^n - \alpha_1 \lambda^{n-1} - \ldots - \alpha_{n-1}\lambda - \alpha_n = 0 \tag{18}$$

For example, for a first order system ($n = 1$), this requires that the solution of

$$\lambda - \alpha = 0 \tag{19}$$

has modulus less than unity, i.e. $|\alpha| < 1$ is a necessary and sufficient condition for (asymptotic) stability along the k-axis.

This result supports the specialised result of Arimoto (1984) for his original learning law. He used there an improper controller K and showed that the resulting learning system converges if the controller gain is chosen such that a specific term has norm less than unity. As it turns out, his condition is exactly equal to the condition on the spectral radius as given above.

The counter-intuitive result is that stability is largely independent of the plant and the controllers. This is a direct result of the assumption of a finite trial end-time T because on a finite time interval, a linear system can only produce a bounded output, even if it is unstable. In the stability notion considered here these unstable outputs are still "acceptable". The problem with this result is thus that even if $\{e_k(t)\}$ is guaranteed to converge to a limit $e_\infty(t)$, this terminal error profile might be unstable and/or it might be worse than $e_1(t)$, i.e. that the learning did not result in an improvement of the error. For an acceptable $e_\infty(t)$, additional conditions on E_i, taking the plant and the learning structure into account, have to be satisfied. These additional objectives lead to similar conditions on the norm of the learning operator as the ones reviewed in Section 4.

Further analysis (omitted here for brevity) shows however that there is a trade-off between the magnitude of the limit error $e_\infty(t)$ and the rate of convergence, i.e. the ratio $||e_{k+1} - e_\infty||/||e_k - e_\infty||$. Depending on the relaxation factors α_i, there is either a fast convergence rate to a large terminal error or slow convergence to a small limit error. Specifically, if $\max(\lambda_i)$ is close to unity, a small $e_\infty(t)$ can be expected and if $\max(\lambda_i)$ is close to zero, fast convergence is enforced with a large terminal error. Only if $b_k(t) \equiv 0$ and thus only if $\sum_{i=1}^{n} \alpha_i = 1$ can a terminal error of zero be achieved. It is easy to show that in this case $\lambda = 1$ is a root of (18). Therefore, the spectral radius of L_T can at best be $\rho(L_T) = 1$. The situation is reminiscent of classical control where the inclusion of an integrator into the controller, which is on the border of stability, results in a zero error for constant disturbances.

Addressing algorithm performance is another step that still largely remains to be done. One possible methodology is suggested in the next section for a simpler learning algorithm.

4. H_∞-Optimal ILC

The desire to improve algorithm performance and particularly the convergence rate was addressed through the use of H_∞-optimal control. As is known from (Padieu and Su, 1990; Owens and Neuffer, 1992), the convergence condition in the frequency domain for a learning system on $[0, \infty)$ is that the error transition operator $L(s)$ which relates the error in one trial to the error of the next trials and which is defined by

$$e_{k+1}(s) = L(s)e_k(s) \tag{20}$$

has H_∞-norm less than or equal to one. It is a sufficient condition on finite time intervals. Furthermore, if it has H_∞-norm γ with $\gamma < 1$, the learning rate is exponential and the following estimate for the error norm holds point-wise:

$$|e_{k+1}(j\omega)| = \gamma^{k+1}|e_0(j\omega)| \tag{21}$$

In order to achieve a fast reduction of the error norm, it is thus sufficient to have a learning operator with a small H_∞-norm γ. The strategy for the design of iterative learning controllers is then to construct learning controllers $K(s)$ such that $||L(s)||_\infty = \gamma$ is minimised. For a learning control law of the form

$$u_{k+1}(s) = u_k(s) + K(s)e_{k+1}(s) \tag{22}$$

the error transition operator $L(s)$ is equal to the sensitivity matrix $S(s) = (I + G(s)K(s))^{-1}$. This simple form allows the use of standard algorithms for the design of the H_∞-optimal controller (Doyle *et al.*, 1989). To satisfy regularity conditions of the design algorithm, a mixed sensitivity problem of the form $||S, KS||_\infty$ is sensible to use because the inclusion of the term KS can be justified by the goal to achieve stability robustness of the closed loop with regard to the time axis and additive perturbations to the plant. A frequency dependent weight was included to shape the sensitivity function. In (Owens and Rogers, 1992), this idea is introduced for repetitive systems and the relationship of H_∞-minimised controllers to high gain feedback is discussed.

A practically important question is the necessary shape for the weight function in the H_∞-optimisation problem. Taking the norm of (20) suggests that the maximal singular value $\bar{\sigma}(L(j\omega))$ should be as small as possible for a frequency range as large as possible to achieve a fast reduction of the error. This objective can be refined as explained below.

In order to reduce the error, the weight should be chosen such that the learning operator has small maximal singular values especially at the frequencies where the first errors have large magnitude. The first error $e_0(t)$ is for the particular choice of $u_0(t) \equiv 0$ equal to the reference $r(t)$. In order to satisfy invertibility conditions, the reference signal is usually chosen as the output of a strictly proper system. It has then large magnitude for small frequencies and is above a certain frequency decreasing with at least 20 dB/decade.

The H_∞-optimisation approach, as becomes evident even in the SISO case, can only be successful for a restricted class of systems. It is well-known that for non-minimum phase plants and plants of relative degree greater than one the sensitivity function has always peaks of $\bar{\sigma}(S(j\omega)) > 1$ (Maciejowski, 1989). The studies undertaken showed however that even in these cases the use of iterative learning control can be of practical advantage. This is because even if the error e_k in the limit $k \to \infty$ may become infinite, it can decrease substantially during the first few trials.

The error transition operator can normally through the inclusion of weighting functions in the design criterion be forced to have an amplification of much less than one for small frequencies. If the region where $\bar{\sigma}(S(j\omega)) > 1$ can be pushed to high enough frequencies (this depends on the plant structure, e.g. the location of non-minimum phase zeros), then the application of Iterative Learning Control will reduce $|e_0(j\omega)| = |r(j\omega)|$ and the next few errors as well. Only after a number of trials $|e_k(j\omega)|$ will be noticeably large at the frequencies where it is amplified and small everywhere else. This is a kind of "practical convergence" where during a limited number of trials $|e_k(j\omega)|$ decreases and only then does the expected divergence appear. The algorithm can be terminated at that point with small tracking error.

The conclusions of these considerations is that Iterative Learning Control in the described scenario can be used until the error norm is below a desired tolerance ε or until it increases. As a guideline, it is suggested that the learning controller is designed such that it has a large bandwidth, a small gain especially at low frequencies and no or only a small peak of $\bar{\sigma}(S) > 1$. These are essentially the same "rules of thumb" for the shape of $\bar{\sigma}(S)$ as in classical controller design.

Figure 2 shows a simulation of an Iterative Learning Control system illustrating "practical convergence". Up to the 15-th trial, Iterative Learning Control results in an error reduction and only afterwards increases the error again. If learning is stopped after the 15-th trial, its use has resulted in a decrease of $\max|e_k(t)|$ from $k = 1$ to $k = 15$ of approx. 1/10.

Another possibility of optimised controller design for non-minimum phase plants is to proceed as suggested above, i.e. to design a controller K such that $||S||_\infty$ is minimised. This controller is then used, but the learning law is changed to

$$u_{k+1} = \alpha u_k + K e_{k+1} \tag{23}$$

The factor α is chosen such that $||S||_\infty = 1/\alpha$. Because the error transition operator is proportional to α (see Section 3), this achieves that the H_∞-condition for $L(s)$ is just satisfied. The disadvantage of this method is that the limit error is no longer zero, as discussed in Section 3.

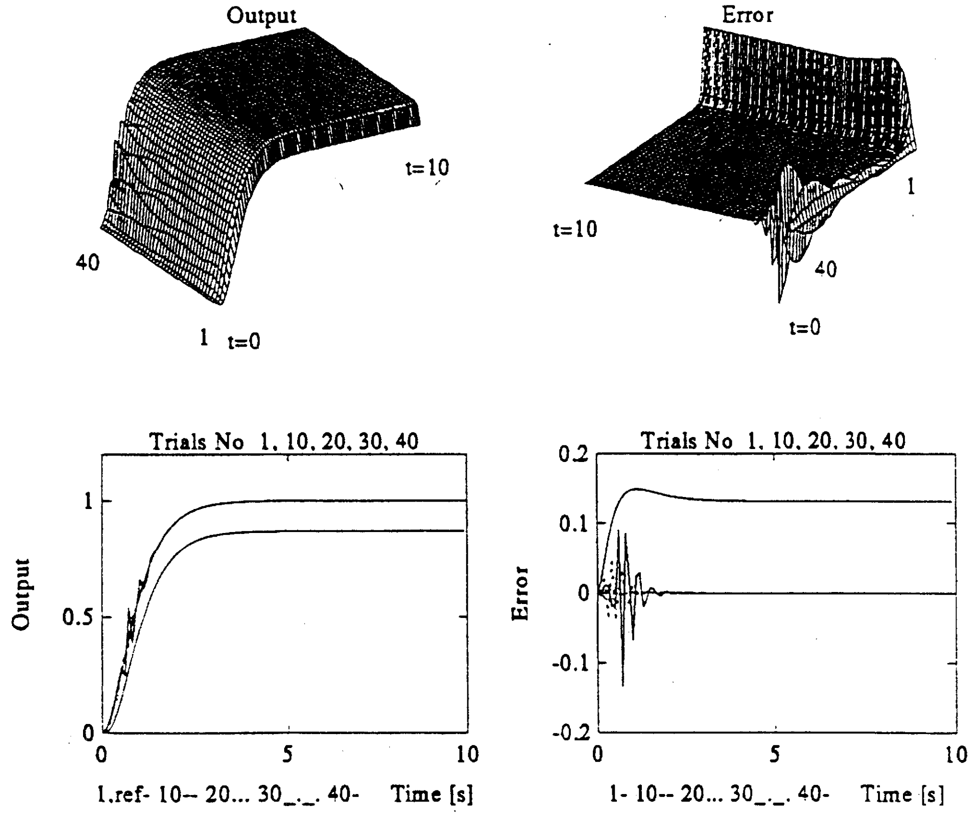

Fig. 2. ILC system showing "practical convergence" (up to the 15th iteration).

5. Norm-Optimal ILC

The newest algorithm for Iterative Learning Control developed by the authors is Norm-Optimal Iterative Learning Control. This algorithm is given in an abstract Hilbert space setting. It uses a "classical" linear quadratic optimality criterion in order to achieve rapid convergence. The introduction of a performance criterion simplifies the two-dimensional stability problem of Iterative Learning Control because the integral criterion 'averages out' the time-dimension and reduces the 2D-convergence problem to the simpler question of convergence of a one-dimensional, scalar sequence. A strictly non-increasing error sequence is achieved through a well-conceived choice of the performance criterion. The algorithm is as follows: at each iteration, the change in input $u_{k+1} - u_k$ is computed as to minimise the cost criterion J_{k+1} with

$$J_{k+1} = ||e_{k+1}||^2 + \lambda ||u_{k+1} - u_k||^2 \tag{24}$$

where the norm $||\cdot||^2$ might be for continuous systems the L_2-norm, i.e. the integral of the squared signal in the interval $[0,T]$, and λ is a positive weight. This approach to generating the control increment aims to reduce the error at each trial whilst bounding the change in the control input. From the formulation of (24) follows immediately that

$$||e_{k+1}||^2 \leq J_{k+1} \leq ||e_k||^2 \tag{25}$$

The monotonic ordering shows that this algorithm really is a descent algorithm. The right inequality is the consequence of optimality: the optimal value of J_{k+1} is less than all non-optimal values and thus less than the value for the input $u_{k+1} = u_k$. For this choice of input, the error e_{k+1} equals e_k. The left inequality follows because norms are always non-negative.

The algorithm is conceptually related to the Levenberg-Marquardt algorithm (Marquardt, 1963) for the minimization of non-linear functionals. Its properties are by now well-known and promise good results for Iterative Learning Control. This algorithm can be seen as a weighted combination of Newton's method and a steepest descent method. The factor λ trades off between these two: as $\lambda \to 0$, it approaches a pure Newton's method and for $\lambda \to \infty$, it goes into a pure steepest descent method.

Using methods from functional analysis, it can be shown that

1) The algorithm converges in the sense that
$$\lim_{k\to\infty} J_k = J_\infty = \lim_{k\to\infty} ||e_k||^2 \tag{26}$$
2) The change in input converges in norm to zero:
$$\lim_{k\to\infty} ||u_{k+1} - u_k|| = 0 \tag{27}$$
3) The error converges weakly to zero in the range of G:
$$\{e_k\} \xrightarrow{w} 0 \tag{28}$$
4) If either $r \in \text{range}(G)$ or range(G) is dense then the error converges in norm to zero:
$$\{e_k\} \longrightarrow 0 \tag{29}$$

The proofs are omitted and may be reviewed in (Amann *et al.*, 1995).

The minimization problem (24) requires only standard methods for its solution (Kwakernaak and Sivan, 1972; Anderson and Moore, 1989). The optimal solution is in abstract form given as

$$u_{k+1} = u_k + G^* e_{k+1} \tag{30}$$

where G^* is the adjoint operator of the plant. This formulation however requires non-causal operations because the adjoint operator is feeding back *future* errors. It can however be recast into a causal algorithm under the assumption of full state knowledge as a consequence of the causality structure of Iterative Learning Control. In Iterative Learning Control, the reference and the previous input are known over the whole time interval before each trial. For the causal formulation, $G^* e_{k+1}$ is given as state-feedback with the well-known, time-dependent Riccati gain matrix $K(t)$ and an additional predictive or feedforward term $\xi_{k+1}(t)$ which is computed as solution of a time-reversed differential equation before each trial. The causal algorithm is then for a continuous plant (A, B, C) as follows

$$u_{k+1}(t) = u_k(t) - \frac{1}{\lambda}B^T\left[K(t)\,x_{k+1}(t) - \xi_{k+1}(t)\right]$$

$$\dot{K}(t) = -A^T K(t) - K(t)A + K(t)\frac{BB^T}{\lambda}K(t) - C^T C, \quad 0 \le t \le T, \qquad K(T) = 0$$

$$\dot{\xi}_{k+1}(t) = -\left(A - \frac{BB^T}{\lambda}K(t)\right)^T \xi_{k+1}(t) - C^T r(t) + K(t)B\,u_k(t), \qquad \xi_{k+1}(T) = 0$$

There are various other aspects of this algorithm which are currently being studied. A rather important one is that the rate of reduction of the error can be arbitrarily changed by the factor λ in (24). The more the change in input is weighted in comparison to the error, the slower the rate of reduction of the error. This can also be observed in simulations. Also, the extension to MIMO systems is straightforward and requires merely the inclusion of weight-matrices into the norms in (24). Other aspects are how to improve robustness with respect to uncertainty of the plant, e.g. by including a learning observer to make all states available. Using different norms can positively affect the numerical conditioning and convergence behaviour, e.g. weighting the error at $t = T$ was found to be advantageous. The behaviour of non-minimum phase plants in this algorithm was studied in (Amann and Owens, 1994). Because the algorithm is given in abstract form, it can as well be applied to discrete-time systems. As last point, the question of whether and how geometrical convergence can be achieved is currently being investigated.

Figure 3 shows a simulation of Iterative Learning Control for the plant $g(s) = (s+0.5)/(s+1)^2$ and the same reference $r(t)$ as in Fig. 1. The error was ten times higher weighted as the input in (24). Full plant and state knowledge was assumed. The good convergence behaviour of this algorithm is evident.

6. Conclusions

The paper presented the current state of research and gave an overview of the advantages and drawbacks of current and proposed learning algorithms. The different methods differed, besides others, in the available knowledge about the plant (the robustness regarding uncertainty about the plant) and the kind of convergence achieved. The analysis for a finite trial duration showed the importance of the studied stability/convergence notion and resulting from that, of the inclusion of the terminal error in the analysis. It showed also that there is an unresolvable trade-off between a zero-terminal error and the speed of how quick (any) terminal error is reached. This points to other (well-known) problems of Iterative Learning Control for "difficult" plants. The largest difficulties are posed by non-minimum phase plants and plants with a high relative degree. One possibility for these systems, suggested e.g. in (Padieu and Su, 1990), is to use a prefilter for the input that enables convergence

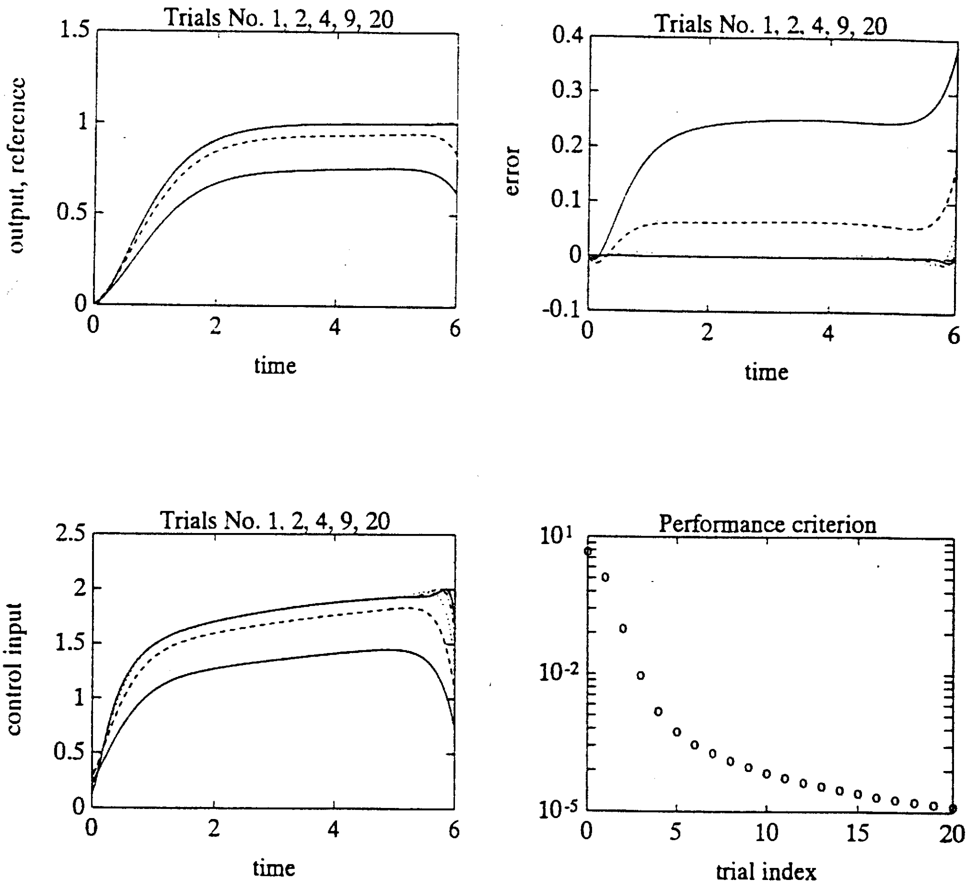

Fig. 3. Simulation of Norm-Optimal Iterative Learning Control for the plant $g(s) = (s+0.5)/(s+1)^2$ and reference $r(t) = 1-(1+2t)e^{-2t}$. Linetypes: trial 1,20: dashed; 2, reference: full; 4: dotted and 9: dot-dashed.

in norm but leads to a non-zero terminal error. Another method was proposed here. It was shown that Iterative Learning Control can be of use for a limited number of trials even for systems that violate convergence conditions (applying per definition to $k \to \infty$).

In the last section, Norm-Optimal Iterative Learning Control was introduced. This algorithm makes good use of the causality structure of Iterative Learning Control and achieves in the limit an error that is minimal in a least-squares sense. It allows good control over the rate of convergence with design parameters but needs the assumption of state-feedback.

Acknowledgements

Current work is supported by the UK Science and Engineering Research Council under contract No. GR/H/48286. Some of the simulations were done by Anja Wahl, an ERASMUS student from the Institut für Systemdynamik und Regelungstechnik, Universität Stuttgart.

References

Amann N., Owens D.H. and Rogers E. (1995): *Iterative learning control using optimal feedback and feedforward actions.* — Submitted to Int. J. Control.

Amann N. and Owens D.H. (1994): *Non-minimum phase plants in iterative learning control.* — Proc. 2nd Int. Conf. *Intelligent Systems Engineering*, Hamburg-Harburg.

Anderson B.D.O. and Moore J.B. (1989): *Optimal Control — Linear Optimal Control.* – Englewood Cliffs: Prentice-Hall.

Arimoto S. (1990): *Learning control theory for robotic motion.* — Int. J. Adaptive Control and Signal Processing, v.4, No.6, pp.543–564.

Arimoto S., Kawamura S. and Miyazaki F. (1984): *Bettering operations of robots by learning.* — J. Robotic Systems, v.1, No.2, pp.123–140.

Byrnes C.I. and Willems J.C. (1984): *Adaptive stabilization of multivariable linear systems.* — Proc. 23rd IEEE Conf. *Decision and Control*, Las Vegas, Nevada, pp.1574–1577.

Doyle J.C., Glover K., Khargonekar P.P. and Francis B.A. (1989): *State-space solutions to standard H_2 and H_∞ control problems.* — IEEE Trans. Automat. Control., v.34, No.8, pp.831–847.

Edwards J.B. and Owens D.H. (1982): *Analysis and Control of Multipass Processes.* — Chicester: Research Studies Press (John Wiley).

Ilchmann A. (1993): *Non-Identifier-Based High-Gain Adaptive Control.* — v.189 of Lecture Notes in Control and Information Sciences, London: Springer-Verlag.

Kwakernaak H. and Sivan R. (1972): *Linear Optimal Control Systems.* — New York: Wiley-Interscience.

Maciejowski J.M. (1989): *Multivariable Feedback Design.* — Wokingham: Addison-Wesley.

Marquardt D.W. (1963): *An algorithm for least-squares estimation of nonlinear parameters.* — J. Soc. Indust. Appl. Math., v.11, No.2 pp.431–441.

Moore K.L. (1993): *Iterative Learning Control for Deterministic Systems.* — Advances in Industrial Control Series, London: Springer-Verlag.

Owens D.H. (1992): *Iterative learning control – convergence using high gain feedback.* — Proc. 31st IEEE Conf. *Decision and Control*, Tucson, AZ, pp.2545–2546.

Owens D.H. (1993): *Iterative learning control using adaptive high gain feedback.* — Proc. 2nd *European Control Conf.*, Groningen, the Netherlands, pp.195–198.

Owens D.H. and Neuffer D. (1992): *Theoretical and computational studies in iterative learning control.* — Internal report series of the Centre for Systems and Control Engineering, University of Exeter, Report No.92/02.

Owens D.H., Prätzel-Wolters D. and Ilchmann A. (1987): *Positive real structure and high gain adaptive stabilization.* — IMA J. Math. Control and Inf., v.4, pp.167–181.

Owens D.H. and Rogers E. (1992): *H-infinity-norm minimisation and the stabilisation of systems with repetitive dynamics.* — Trans. Inst. Measurement and Control, v.14, No.3, pp.126–129.

Padieu F. and Su R. (1990): *An H_∞ approach to learning control systems.* — Int. J. Adaptive Control and Signal Processing, v.4, No.6, pp.465–474.

Rogers E. and Owens D.H. (1992): *Stability Analysis for Linear Repetitive Processes.* — v.175 of Lecture Notes in Control and Information Sciences, Berlin: Springer-Verlag.

Received: October 12, 1994