15th Triennial World Congress, Barcelona, Spain

PREDICTION IN ITERATIVE LEARNING CONTROL SCHEMES VERSUS LEARNING ALONG THE TRIALS

M. Rzewuski * **E. Rogers** * **D. H. Owens** **

* *Dept. of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ UK*
** *Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield S1 3JD, UK*

Abstract: Recent years has seen much progress in the theory and application of iterative learning control schemes for both linear and (classes of) nonlinear dynamics. In the case of the former, many algorithms based on minimizing a suitable cost function have been reported. Here the interest is in the so-called norm optimal approach where the basic philosophy is to compute the control input on the current trial such that the tracking error is reduced in an optimal way without too much deviation from the control input used on the previous trial. This paper compares the performance of a range of controllers arising from use of the norm optimal approach - both stand alone and against alternatives.

Keywords: iterative learning control, optimal algorithms.

1. INTRODUCTION

Iterative learning control (ILC) is a technique to control systems operating in a repetitive mode with the additional requirement that a specified output trajectory $r(t)$ defined over a finite interval $[0, T]$ is followed to high precision. There are numerous examples of such systems including robot manipulators that are required to repeat a given task to high precision, chemical batch processes or, more generally, the class of tracking systems.

Motivated by human learning, the basic of idea of ILC is to use information from previous executions of the task in order to improve performance from trial to trial in the sense that the tracking error is sequentially reduced ((Arimoto *et al.*, 1984; Moore, 1993)). Typical ILC algorithms construct the input to the plant on a given trial from the input used on the previous trial plus an additive increment that is typically a function of past values of the observed output error, i.e. the difference between achieved and desired plant output. The objective of constructing a sequence of input functions $\{u_k(t)\}_k,\ t \in [0, T]$, such that the performance as the task is repeated is gradually improving, can be refined to a convergence condition on the input and error:

$$\lim_{k\to\infty} ||e_k|| = 0, \ \lim_{k\to\infty} ||u_k - u_\infty|| = 0 \quad (1)$$

where $e_k(t)$ is the error on trial k, i.e. difference between $r(t)$ and the system output, $y_k(t)$, and $u_k(t)$ is the input to the system on this trial.

The above definition of convergent learning is a stability problem on an infinite-dimensional two-dimensional (2D)-product space. As such it places the analysis of ILC schemes firmly outside the scope of traditional control theory. In particular, ILC must be studied in the context of fixedpoint problems or, more precisely, linear repetitive processes ((Rogers and Owens, 1992)).

Since the ILC concept was proposed, a very large number of approaches have been proposed but here we consider the very important case when the ILC scheme for linear plant dynamics is designed based on optimal control techniques. One option here is to use a quadratic cost function where the only control input term in this cost function is that to be used on the current trial. Recent years have also seen the emergence of so-called norm optimal ILC which can be realized in terms of current trial mechanisms combined with feedforward of previous trial data. Here the control input term in the cost function is based on the difference between the input signals used on successive trials.

The algorithms which emerge from the norm optimal approach are based on splitting the two-dimensional dynamics into two separate one-dimensional problems. This is done by introducing a performance criterion as the basis of specifying the control input to be used on each trial. In this context, the problem can be interpreted as the determination of the control input on trial $k+1$ that reduces the tracking error in an optimal way but which does not deviate too much from the control input used on trial k. Use of this approach leads to a number of key properties, of which the most important are as follows.

1. Automatic choice of step size (in the iterative procedures needed to solve the optimization problem).

2. Potential for improved robustness through the use of causal feedback of current trial data and feedforward trial data from previous trials.

It is also possible to enhance norm optimal control to include predictive action (Amann *et al.*, 1998), i.e. future predicted error signals are explicitly included in the cost function. In this paper, we use a computer aided analysis tool constructed for optimal control linear model based ILC schemes to undertake a detailed comparative study of norm optimal (including predictive action) ILC controllers for linear plants against those designed using alternative optimal control based approaches. The scope of this study includes the experimental setup reported in (Frueh and Pham, 2000) (thereby giving the option for future experimental verification). We being in the next section by summarizing the necessary background results (for complete details see (Amann *et al.*, 1996*b*; Amann *et al.*, 1998) and the relevant references).

2. BACKGROUND

In this paper we restrict attention to the norm optimal approach in a Hilbert space setting. The 'translation of this setting to the special case of a linear time invariant system is straightforward and will form a starting point for the work reported here. The alternative 'conventional' algorithm for these systems (i.e. a quadratic cost function where the contribution from the control input term is only (explicitly) based on the current trial values) is not detailed here and can be found in, for example, (Frueh and Pham, 2000) and the relevant cited references.

The following is the formal definition of a successful ILC algorithm.

Definition 1. Consider a dynamic system with input u and output y. Let $\mathcal{Y}$ and $\mathcal{U}$ be the output and input function spaces respectively and let $r \in \mathcal{Y}$ be a desired reference trajectory from the system. Then an ILC algorithm is successful if, and only if, it constructs a sequence of control inputs $\{u_k(t)\}_{k\geq 0}$ which, when applied to the system or plant (under identical experimental conditions), produces an output sequence $\{y_k(t)\}_{k\geq 0}$ with the following properties of convergent learning

$$\lim_{k\to\infty} y_k = r, \quad \lim_{k\to\infty} u_k = u_\infty \tag{2}$$

Here convergence is interpreted in terms of the topologies assumed in $\mathcal{Y}$ and $\mathcal{U}$ respectively.

Note: This general description includes linear and nonlinear dynamics, continuous or discrete plants, and time-invariant or time-varying systems.

Now let the space of output signals $\mathcal{Y}$ be a real Hilbert space and $\mathcal{U}$ also be a real (and possibly distinct) Hilbert space of input functions. The respective inner products (denoted by $\langle\cdot,\cdot\rangle$) and norms $||\cdot||^2 = \langle\cdot,\cdot\rangle$ are indexed in a way that reflects the space if it is appropriate to the discussion.

The dynamics of the plant considered here are approximated by a linear model which in operator form can be written as

$$y = Gu + z_0 \tag{3}$$

where no loss of generality arises from setting $z_0 = 0$. Also it is clear that the ILC procedure, if convergent, solves the problem $r = Gu_\infty$ for u_∞ and, if G is invertible, the formal solution is just $u_\infty = G^{-1}r$. A basic premise of the ILC approach is that the direct inversion of G is regarded as an impractical solution because it requires exact knowledge of G and involves derivatives of r. This high-frequency gain characteristic would make such an approach sensitive to noise and other disturbances. Also it can be argued that inversion of the whole plant G is unnecessary as the solution only requires finding the pre-image of r under G.

The above problem is easily seen to be equivalent to finding the minimizing input u_∞ for the optimization problem

$$min_u\{||e||^2 \;:\; e = r - y, y = Gu\} \quad (4)$$

The optimal error $||r - Gu_\infty||^2$ is a measure of how well the ILC algorithm has solved the inversion problem. It also represents the best that the system can do in tracking the signal r. The case of interest here is when the optimal error is zero, i.e. u_∞ is a solution of $r = Gu_\infty$. Also (4) is clearly a singular optimal control problem which by its very nature requires an iterative solution.

There are an infinity of potential iterative procedures for solving (4) and of these the gradient approach has the simplest form and has been extensively investigated in the ILC literature. A gradient based ILC algorithm has the form

$$u_{k+1} = u_k + \epsilon_{k+1} G^* e_k \quad (5)$$

where $G^* : \mathcal{Y} \to \mathcal{U}$ is the adjoint operator to G, and ϵ_{k+1} is a step length to be chosen at each iteration. This general approach suffers from the need to choose this step length on each trial and the feedforward structure of the iteration takes no account of current trial effects - including disturbances and plant modeling errors.

Norm optimal ILC has the following two crucial properties relative to alternative gradient based algorithms
1. Automatic choice of step size.
2. Potential for improved robustness through the use of causal feedback of current trial data and feedforward of data from previous trials.
In particular, the ILC algorithms considered here compute the input on trial $k+1$ as the solution of the minimum norm optimization problem

$$u_{k+1} = \arg\, min_{u_{k+1}}\{J_{k+1}(u_{k+1})\} \quad (6)$$

subject to

$$e_{k+1} = r - y_{k+1},\; y_{k+1} = Gu_{k+1} \quad (7)$$

where the performance index (or optimality criterion) used is

$$J_{k+1}(u_{k+1}) = ||e_{k+1}||^2_{\mathcal{Y}} + ||u_{k+1} - u_k||^2_{\mathcal{U}} \quad (8)$$

The initial control $u_0 \in \mathcal{U}$ can be arbitrary but, in practice, will be a good first guess at the solution of the problem. The relative weighting of reducing the current trial error against minimizing the deviation in the control input signals used on successive passes can be absorbed into the definitions of the norms in $\mathcal{Y}$ and $\mathcal{U}$.

The benefits of this approach are immediate from the simple interlacing result

$$||e_{k+1}||^2 \leq J_{k+1}(u_{k+1}) \leq ||e_k||^2,\; \forall\, k \geq 0 \quad (9)$$

which follows from optimality and the fact that the (non-optimal) choice of $u_{k+1} = u_k$ would lead to the relation $J_{k+1}(u_k) = ||e_k||^2$. This result states that the algorithm is a descent algorithm as the norm of the error is monotonically decreasing in k. Also equality holds if, and only if, $u_{k+1} = u_k$, i.e. when the algorithm has converged and no more input-updating takes place.

The controller on trial $k+1$ is given by

$$u_{k+1} = u_k + G^* e_{k+1},\; \forall k \geq 0 \quad (10)$$

This relationship, together with the error update relation

$$e_{k+1} = (I + GG^*)^{-1} e_k,\; \forall\, k \geq 0 \quad (11)$$

and the recursive input update relation

$$u_{k+1} = (I + G^*G)^{-1}(u_k + G^* r),\; \forall\, k \geq 0 \quad (12)$$

can be used to undertake a detailed analysis of the properties of this class of ILC laws (Amann *et al.*, 1996*b*).

Predictive optimal ILC (Amann *et al.*, 1998) extends the cost function to the form

$$J_{k+1,N}(u_{k+1}) = \sum_{i=1}^{N} \lambda^{i-1}(||e_{k+i}||^2_{\mathcal{Y}} + ||u_{k+i} - u_{k+i-1}||^2_{\mathcal{U}}) \quad (13)$$

This criterion includes the error of the next N trials as well as the corresponding changes in the control input signals The weighting parameter $\lambda > 0$ determines the importance of more distant (future) errors and incremental inputs compared with the current ones. By including more future signals into the performance criterion, the algorithm becomes less ‘short sighted’. The theory given above extends in a natural manner to this case and an obvious question to ask is: when does the extra (computational) cost become worthwhile?

3. ANALYSIS

The implementation of norm optimal and predictive norm optimal ILC schemes is a relatively straightforward exercise and is hence not detailed here. Instead, we proceed to use MATLAB based implementations to discuss and compare the performance of these approaches both stand alone and relative to alternative optimal control based algorithms. The plants used have the following transfer functions

$$G_1(s) = \frac{s+1}{s^2+5s+6} \quad (14)$$

$$G_1(s) = \frac{s-1}{s^2+5s+6} \quad (15)$$

Also we use a minimal state space realization of these processes with state space triple $\{A, B, C\}$, and, in the general case, assume that there are l inputs and m outputs.

In this paper, we will only consider the use of linear quadratic optimal control cost functions in the ILC setting. In particular, we will compare the optimal control based ILC of (Frueh and Pham, 2000) (and others) against norm optimal and predictive norm optimal control. The choice of input and output spaces is as follows

$$\begin{aligned} u \in \mathcal{U} &= L_2^l[0,T] \\ (r, r(T)) \in \mathcal{Y} &= L_2^m[0,T] \times \mathbb{R}^m \end{aligned} \quad (16)$$

Also the inner products in $\mathcal{Y}$ and $\mathcal{U}$ are defined as

$$\begin{aligned} \langle (y_1, z_1), (y_2, z_2) \rangle_{\mathcal{Y}} &= \frac{1}{2} \int_{t=0}^{T} y_1^T(t) Q y_2(t)\, dt \\ &\quad + \frac{1}{2} z_1^T F z_2 \\ \langle u_1, u_2 \rangle_{\mathcal{U}} &= \frac{1}{2} \int_{t=0}^{T} u_1^T(t) R u_2(t)\, dt \end{aligned} \quad (17)$$

where Q and R are symmetric positive definite matrices and F is a symmetric positive semi-definite matrix.

In all cases, the initial conditions are taken to be homogeneous without loss of generality since the plant response to non-zero initial conditions can be absorbed into $r(t)$ in a natural manner. The cost function in each case becomes, with the specified norms in $\mathcal{Y}$ and $\mathcal{U}$, a familiar linear quadratic performance criterion. For example, in the case of algorithms of the form considered by (Frueh and Pham, 2000) and others we have

$$J_{k+1} = \frac{1}{2} \int_{t=0}^{T} \{ e_{k+1}^T Q e_{k+1}(t) + u_{k+1}^T R u_{k+1}(t) \}\, dt \quad (18)$$

In effect, the problem is now a combination of the well known optimal tracking problem (tracking of $r(t)$) and disturbance accommodation problem (regarding $u_k(t)$ as a known disturbance on trial $k+1$) from standard linear systems theory. The cost functions considered in (Frueh and Pham, 2000) and elsewhere result in the following solution algorithm

$$\begin{aligned} u_{k+1} &= u_k + \delta u_{k+1} \\ \delta u_{k+1} &= (R + P^T Q P)^{-1} (P^T Q e_k - R u_k) \end{aligned} \quad (19)$$

(where the matrix P is defined by the plant state space matrices and is not detailed here). In the case of norm optimal control, standard optimal control theory now gives the solution as

$$\begin{aligned} \dot{\psi}_{k+1}(t) &= -A^T \psi_{k+1}(t) - C^T Q e_{k+1}(t) \\ u_{k+1}(t) &= u_k(t) + R^{-1} B^T \psi_{k+1}(t) \\ \psi_{k+1}(T) &= C^T F e_{k+1}(T),\ t \in [0,T] \end{aligned} \quad (20)$$

This representation is non-causal (in the standard sense) but it can be transformed into a causal implementation as detailed next for the case of a relaxation factor α.

Transform the costate vector $\psi_{k+1}(t)$ using

$$\begin{aligned} \psi_{k+1}(t) &= -K(t) \left[x_{k+1}(t) - \alpha x_k(t) \right] \\ &\quad + \zeta_{k+1}(t) \end{aligned} \quad (21)$$

where the feedback gain matrix $K(t)$ satisfies the well known Riccati (matrix) differential equation

$$\begin{aligned} \dot{K}(t) &= -A^T K(t) - K(t) A \\ &\quad + K(t) B^T R^{-1} B^T K(t) - C^T Q C \\ K(T) &= C^T F C \end{aligned} \quad (22)$$

Note that this last equation of independent of the inputs, states and outputs of the system and hence only needs to be computed once before the sequence of trials begin.

The predictive or 'feedforward' term $\zeta_{k+1}(t)$ needs to be computed on each trial using

$$\begin{aligned} \dot{\zeta}_{k+1}(t) &= -(A - BR^{-1}B^T K)^T \zeta_{k+1}(t) \\ &\quad - \alpha C^T Q e_k(t) \\ &\quad + (1-\alpha) K B u_k(t) \\ &\quad - (1-\alpha) C^T Q r(t) \end{aligned} \quad (23)$$

with terminal boundary condition

$$\zeta_{k+1}(T) = C^T F \left[\alpha e_k(T) + (1-\alpha) r(T) \right] \quad (24)$$

The algorithm is now causal since (22) and (23) can be solved off-line by reverse time simulation using available previous trial data.

The following is the final form (with no relaxation factor) of the implementation algorithm for predictive optimal ILC (for complete details see (Amann *et al.*, 1998))

$$\begin{bmatrix} u_{k+1} \\ u_{k+2} \\ \vdots \\ u_{k+N} \end{bmatrix} = \begin{bmatrix} u_k \\ u_k \\ \vdots \\ u_k \end{bmatrix} - R_N^{-1} B_N^T (K(t)$$

$$\times \left(\begin{bmatrix} x_{k+1} \\ x_{k+2} \\ \vdots \\ x_{k+N} \end{bmatrix} - \begin{bmatrix} x_k \\ x_k \\ \vdots \\ x_k \end{bmatrix} \right) - \xi_{k+1,N}(t)) \quad (25)$$

$$\dot{K} = -A_N^T K - K A_N + K B_N R_N^{-1} B_N^T K - C_N^T Q C_N, \quad K(T) = C_N^T(T) F C_N(T) \quad (26)$$

$$\dot{\xi}_{k+1,N}(t) = -\left(A_N - B_N R_N^{-1} B_N^T K \right)^T \xi_{k+1,N}(t) - C_N^T Q_N \begin{bmatrix} e_k(t) \\ e_k(t) \\ \vdots \\ e_k(t) \end{bmatrix} \quad (27)$$

with terminal condition

$$\xi_{k+1,N}(T) = C_N^T F_N \times \left[e_k^T(T) \dots e_k^T(T) \right]^T \quad (28)$$

where A_N is a block diagonal matrix with each diagonal entry equal to the state matrix A, and the matrices B_N and C_N are constructed in an identical manner using the matrices B and C respectively. For the precise form of the weighting matrices Q_N, R_N and F_N, see again (Amann *et al.*, 1998). In the remainder of this paper, we apply these algorithms to the two plants defined by (14) and (15) respectively, where in all simulations $T = 6$, and (as required) $R = 0.1$, $Q = 10$, $F = 0$, $N = 5$. Also $r(t) = \sin t$, $u_0(t) = 0$.

Applying an optimal control algorithm of the form (19) to the minimum phase plant defined by (14) is a straightforward task and it has been reported many times in the literature that this algorithm is capable of very good performance. Our experience shows that in some cases it exhibits slow convergence and involves a matrix inversion which could result in a high computational cost. Also it is easy to demonstrate that this algorithm performs poorly when the plant is non-minimum phase (i.e. here when applied to (15)). Note also that in some cases the control signals required can be very large to the extent that they can be outside the effective operating ranges of many actuators.

The norm optimal algorithm does not involve matrix inversion and, because the cost function penalizes the change in the control input signals between successive trials, the resulting signals demanded are generally much smaller than with an algorithm of the form (19). Overall, we have found that norm optimal control is more effective against non-minimum phase dynamics but still the performance can be rather poor.

Turning to the use of predictive ILC, here we found that this scheme produces much faster convergence but at the expense of higher computational cost. Also the performance of such schemes is significantly dependent on the choice of the extra design parameters, i.e. the prediction horizon N and the weighting factor λ. Space limitations prevent a detailed treatment of this aspect here and the reader is referred to (Rzewuski *et al.*, 2001) which also details all of the simulation studies used in this work. Another advantage of predictive ILC is that it can give much better performance when applied to non-minimum phase processes but note that this could come at the expense of 'very strong' control action. If such control action is unacceptable then it is still essentially an open research question of how to design effective ILC control schemes in such cases.

4. CONCLUSIONS

This paper has undertaken a detailed comparative study of three competing classes of ILC schemes for linear plants based on optimal control based design philosophies. The major conclusion is that those based on the norm optimal approach (including augmentation by prediction) have significant advantages over alternatives. A major reason for this is that in the norm optimal approach it is the difference in the control signals between successive trials which is penalized as opposed to the current trial input alone in the alternatives. The most significant additional factor to be considered before using predictive ILC is the extra computational cost but (in general terms) this should not be excessive unless the plant has a very large number of states, inputs, and outputs.

In all simulation experiments reported here, the continuous-time versions of the algorithms were used. For experimental verification or applications studies, however, this may not be good practice due to high computation costs and many numerical problems that often can be avoided through the use of a discrete time version of a given algorithm. The discrete time version of the norm–optimal ILC algorithm considered here can be found in (Amann *et al.*, 1996*a*).

REFERENCES

Amann, N., D.H. Owens and E. Rogers (1996*a*). Iterative learning control for discrete-time systems with exponential rate of convergence. *Proceedings of The Institution of Electrical Engineers, Part D, Control Theory and Applications* **143(2)**, 217–224.

Amann, N., D.H. Owens and E. Rogers (1996*b*). Iterative learning control using optimal feedback and feedforward actions. *International Journal of Control* **65(2)**, 277–294.

Amann, N., D.H. Owens and E. Rogers (1998). Predictive optimal iterative learning control. *International Journal of Control* **69(2)**, 203–226.

Arimoto, S., S. Kawamura and F. Miyazaki (1984). Bettering operations of robots by learning. *Journal of Robotic Systems* **1**, 123–140.

Frueh, J.A. and M.Q. Pham (2000). Linear quadratic optimal learning control (LQL). *International Journal of Control* **73(10)**, 832–839.

Moore, K.L. (1993). *Iterative Learning Control for Deterministic Systems*. Springer-Verlag. London.

Rogers, E. and D.H. Owens (1992). *Stability Analysis for Linear Repetitive Processes, LNCIS Series, Vol 175*. Springer-Verlag. Berlin.

Rzewuski, M., E. Rogers and D.H. Owens (2001). Comparative studies on linear iterative learning control schemes. *Research Report, Department of Electronics and Computer Science, University of Southampton, UK* pp. 1–56.