2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

Reduced-order ILC: The Internal Model Principle Reconsidered

Goele Pipeleers and Kevin L. Moore

***Abstract*— When iterative learning control (ILC) is applied to improve a system's tracking performance, the trial-invariant reference input is typically known or contained in a prescribed set of signals. Current ILC algorithms, however, neglect this information and only exploit the trial-invariance of the input signal. In this paper we propose a novel ILC design that explicitly incorporates the additional knowledge on the trial-invariant input. The proposed design approach results in a reduced-order ILC, in the sense that the order of its trial-domain description equals the number of given trial-invariant input signals that are to be tracked. In contrast, current ILC algorithms yield a trial-domain controller of order N, the ILC trial length in discrete time. We discuss the advantages and disadvantages of reduced-order ILC when it is designed to minimize a 2-norm based objective.**

I. Introduction

Iterative learning control (ILC) is an open-loop control strategy that improves the performance of a system executing the same task over and over again by learning from previous iterations/trials [1], [2], [3]. Consider a discrete-time, single-input single-output (SISO), linear time-invariant (LTI) plant $G(q)$ with input $u_l(k)$ and output $y_l(k)$, where k is an independent variable representing time and q is the one-sample-advance operator. The system is commanded to track a given reference command $r(k)$ over and over again, where the trials are labeled by the index l. We assume that each trial has the same length N and that prior to each trial the plant is returned to its zero initial condition [3]. An ILC relies on the repeatability of the input signal to reduce/eliminate the tracking error $e_l(k) = r(k) - y_l(k)$ as $l \to \infty$. To this end, the input $u_{l+1}(k)$ is updated using the input $u_l(k)$ and the error $e_l(k)$ from the previous trial, where the ILC update algorithm is most commonly of the following form:

$$u_{l+1}(k) = Q(q)\left[u_l(k) + L(q)e_l(k)\right] , \qquad (1)$$

with $Q(q)$ and $L(q)$ LTI filters. To achieve superior tracking for $l \to \infty$, ILC relies upon the Internal Model Principle (IMP), which states that if a disturbance/reference signal can be regarded as the output of an autonomous system, including this system in a stable feedback loop guarantees perfect asymptotic rejection/tracking [4], [5]. Although the role of the IMP was recognized early in the development of ILC, even leading to the development of ILC algorithms for rejecting/tracking iteration-varying disturbances/references [6], [7], the full power of this principle was never exploited. As we show below, one consequence of this oversight in the existing ILC literature is that to date all ILC algorithms produce trial-domain dynamics whose order is greater than necessary when the goal is simply to track a prescribed (set of) reference input(s). Our primary contribution here is to show how lower-order trial-domain dynamics result from carefully exploiting knowledge on the reference input.

Goele Pipeleers is with the Department of Mechanical Engineering, Katholieke Universiteit Leuven, Celestijnenlaan 300B, B-3001 Heverlee, Belgium, goele.pipeleers@mech.kuleuven.be.

Kevin L. Moore is with the Division of Engineering, Colorado School of Mines, 1500 Illinois Street, Golden, CO 80401, USA, kmoore@mines.edu.

To explain our contribution in more detail, define the supervectors

$$\mathbf{u}_l = \begin{bmatrix} u_l(0) & u_l(1) & \cdots & u_l(N-1) \end{bmatrix}^T , \qquad (2a)$$
$$\mathbf{y}_l = \begin{bmatrix} y_l(\tau) & y_l(\tau+1) & \cdots & y_l(\tau+N-1) \end{bmatrix}^T , \qquad (2b)$$
$$\mathbf{r} = \begin{bmatrix} r(\tau) & r(\tau+1) & \cdots & r(\tau+N-1) \end{bmatrix}^T , \qquad (2c)$$
$$\mathbf{e}_l = \begin{bmatrix} e_l(\tau) & e_l(\tau+1) & \cdots & e_l(\tau+N-1) \end{bmatrix}^T , \qquad (2d)$$

where τ denotes the relative degree of $G(q)$. In this "lifted notation" [3], [8], the plant $G(q)$ translates into

$$\mathbf{y}_l = \mathbf{G}\mathbf{u}_l , \qquad (3)$$

while the trial-domain description of the ILC algorithm (1) amounts to:

$$\mathbf{u}_{l+1} = \mathbf{Q}(\mathbf{u}_l + \mathbf{L}\mathbf{e}_l) . \qquad (4)$$

The matrices $\mathbf{G}, \mathbf{Q}$ and $\mathbf{L}$ are (as described in more detail below) the Toeplitz matrices formed from the impulse responses of the plant $G(q)$ and the filters $Q(q)$ and $L(q)$, respectively.

It is well-known in the ILC literature that (4) achieves perfect asymptotic rejection/tracking of *any* trial-invariant input if and only if $\mathbf{Q} = I_N$. In this case, the controller (4) can be described in the (trial-domain) state space as

$$\begin{cases} \mathbf{x}_{l+1} = I_N\mathbf{x}_l + \mathbf{L}\mathbf{e}_l \\ \mathbf{u}_l = \mathbf{x}_l \end{cases} . \qquad (5)$$

Since the matrix $\mathbf{L}$ is generally nonsingular, this state-space model is minimal and emphasizes that the ILC is of order N. Consequently, the closed-loop system resulting from the combination of the ILC (4) with the static trial-domain plant (3) is of order N.

In this paper, we show that there also exist ILCs of order less than N that still yield perfect asymptotic tracking for $l \to \infty$. Such an ILC will have the following structure:

$$\begin{cases} \mathbf{x}_{l+1} = I_n\mathbf{x}_l + \mathbf{B_K}\mathbf{e}_l \\ \mathbf{u}_l = \mathbf{C_K}\mathbf{x}_l \end{cases} . \qquad (6)$$

where $n < N$ is the controller order and $\mathbf{C_K}$ is constrained to a specific value (see Section II-C). As we show below, it is possible to track up to n linearly independent prescribed reference signals using a controller of order n.

To summarize, current ILC algorithms achieve tracking of any trial-invariant reference input and are of order N, whereas it is possible to track n particular reference inputs using an ILC algorithm of reduced order $n < N$. In this paper we develop this idea in detail and discuss its implications. The former ILCs are here referred to as full-order, while the latter are called reduced-order ILCs. To allow a comparison between both ILCs, the common full-order norm-optimal ILC design strategy [3], [9] is extended to reduced-order ILCs. Comparison of full-order norm-optimal ILCs and reduced-order norm-optimal ILCs, designed according to the same objective, shows that (i) reduced-order norm-optimal ILCs generally result in simpler learning dynamics and transient behavior; (ii) with a reduced-order norm-optimal ILC the closed-loop stability is more robust to plant model errors; and (iii) only for the reduced-order norm-optimal ILCs a model/plant mismatch degrades the perfect asymptotic tracking performance.

The remaining content of this paper is laid out as follows: Section II introduces some details on the IMP, formulates the ILC problem and details the reduced-order ILC design. Its advantages and disadvantages over full-order ILC are discussed in Section III and illustrated in Section IV by a numerical example. Section V concludes the paper.

To distinguish between time-domain and trial-domain dynamics, plain characters are used for the time domain, while bold characters relate to the trial domain. As such, the symbol q indicates the one-sample-advance operator in the time domain, while the one-sample-advance operator in the trial domain is denoted by $\mathbf{q}$. That is: $qx_l(k) = x_l(k+1)$, while $\mathbf{q}x_l(k) = x_{l+1}(k)$ [1].

II. Methodology

After a brief discussion on reference signal generation (Section II-A), this section formulates the ILC design problem (Section II-B), presents the general structure of a reduced-order ILC (Section II-C), and details its 2-norm optimal design methodology used to illustrate the ideas (Section II-D).

A. Trial-invariant Signal Generation

To track a signal generated by an autonomous system, the IMP tells us to embed that autonomous system in a stable closed-loop system. Let us analyse how a trial-invariant signal can be produced by an autonomous system. First, consider the signal generator $\mathbf{\Sigma}_I(\mathbf{q})$ shown in Figure 1(a), where $\mathbf{q}$ denotes the one-trial-advance operator. Determined by its initial condition $\boldsymbol{\xi}_0$, this system can generate any trial-invariant signal in $\mathbb{R}^N$, as it yields $\mathbf{w}_l = \boldsymbol{\xi}_0$ for all $l = 0, 1, \ldots$ Next, consider the signal generator $\mathbf{\Sigma}_{\mathbf{W}}(\mathbf{q})$ shown in Figure 1(b). The trial-invariant signals generated by $\mathbf{\Sigma}_{\mathbf{W}}(\mathbf{q})$ are restricted to the range of $\mathbf{W} \in \mathbb{R}^{N\times n}$, where $n \leq N$. That is, they equal $\mathbf{w}_l = \mathbf{W}\boldsymbol{\xi}_0$ for $l = 0, 1, \ldots,$ and some arbitrary initial condition $\boldsymbol{\xi}_0 \in \mathbb{R}^n$. Thus, by the IMP, if we embed the system shown in Figure 1(a) inside a stable closed loop, the resulting system will be able to track any trial-invariant reference input, whereas if we embed the system of Figure 1(b), the resulting system will only be able to track reference inputs in the range of $\mathbf{W} \in \mathbb{R}^{N\times n}$.

[1] Notice that the boldfaced $\mathbf{q}$ notation is equivalent to the w-operator introduced in [10] and developed in [8], [11].

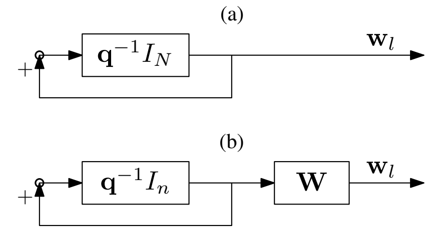

Fig. 1. (a) Generator $\mathbf{\Sigma}_I(\mathbf{q})$ of arbitrary trial-invariant signals in $\mathbb{R}^N$; and (b) generator $\Sigma_{\mathbf{W}}(\mathbf{q})$ of arbitrary trial-invariant signals in the range of $\mathbf{W} \in \mathbb{R}^{N\times n}$. Symbol $\mathbf{q}$ denotes the one-trial-advance operator.

It is readily verified that the full-order ILC (5) contains the signal generator $\mathbf{\Sigma}_I(\mathbf{q})$, while we will show in Section II-C that with a proper design of $\mathbf{C_K}$, the reduced-order ILC (6) embeds $\mathbf{\Sigma}_{\mathbf{W}}(\mathbf{q})$ into the closed-loop system.

B. Problem Formulation

The ILC design is considered in discrete time, where the discrete time instants are labeled by $k = 0, 1, \ldots$ Each trial comprises N time samples and prior to each trial the plant is returned to the same initial conditions, which are here assumed zero without loss of generality [3]. We distinguish between the plant and its model. The discrete-time plant is denoted by $G(q)$, it has relative degree τ and its impulse response is indicated by $g(k)$. The plant model is denoted by $\hat{G}(q)$ and is assumed to have the same relative degree as the plant. The model's impulse response is indicated by $\hat{g}(k)$. The ILC design is formulated in trial domain according to Figure 2, where the supervector signals are defined in (2). Reformulating the plant's convolution relation

$$y_l(k) = \sum_{i=\tau}^{k} g(i)\, u_l(k-i)\ ,$$

in terms of the supervectors $\mathbf{u}_l$ and $\mathbf{y}_l$ yields the following trial-domain plant $\mathbf{G}$:

$$\mathbf{y}_l = \underbrace{\begin{bmatrix} g(\tau) & 0 & \cdots & 0 \\ g(\tau+1) & g(\tau) & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ g(\tau+N-1) & \cdots & g(\tau+1) & g(\tau) \end{bmatrix}}_{\mathbf{G}} \mathbf{u}_l\ .$$

In a similar way, the trial-domain plant model $\hat{\mathbf{G}}$ is derived from $\hat{g}(k)$.

Next, consider the exogenous input signal $\mathbf{w}_l$, which combines the reference input and output disturbances. It is trial-invariant: $\mathbf{w}_l = \mathbf{w}$, for $l = 0, 1, \ldots,$ and confined to the range of $\mathbf{W} \in \mathbb{R}^{N\times n}$. That is, $\mathbf{w}_l$ corresponds to the autonomous output of $\mathbf{\Sigma}_{\mathbf{W}}(\mathbf{q})$, shown in Figure 1(b), from an arbitrary initial condition $\boldsymbol{\xi}_0 \in \mathbb{R}^n$. The matrix $\mathbf{W}$ is assumed to have full column rank and hence, $n \leq N$. An ILC

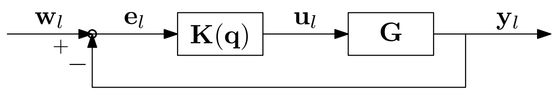

Fig. 2. Trial-domain formulation of the ILC problem, where $\mathbf{w}_l$, $\mathbf{e}_l$, $\mathbf{u}_l$ and $\mathbf{y}_l$ correspond to the supervectors of the exogenous input, tracking error, control signal and plant output, respectively, and $\mathbf{G}$ denotes the lifted system matrix. An ILC corresponds to a trial-domain feedback controller $\mathbf{K}(\mathbf{q})$ that yields perfect asymptotic tracking for trial-invariant inputs $\mathbf{w}_l = \mathbf{w}$.

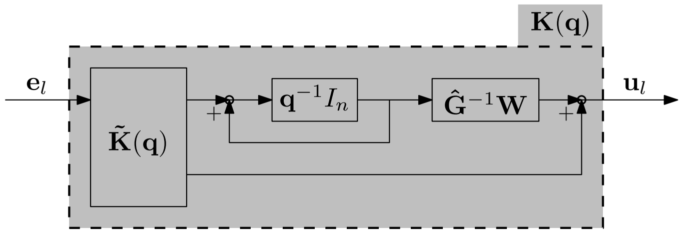

Fig. 3. Structure of an ILC $\mathbf{K}(\mathbf{q})$ that achieves perfect asymptotic tracking of trial-invariant inputs in the range of $\mathbf{W}$.

corresponds to a trial-domain feedback controller $\mathbf{K}(\mathbf{q})$ that yields an internally stable closed-loop system and guarantees perfect asymptotic tracking of the considered trial-invariant inputs $\mathbf{w}$, i.e. $\lim_{l\to\infty} \mathbf{e}_l = 0$.

C. Internal Model Principle

The IMP [4], [5] states that $\mathbf{K}(\mathbf{q})$ achieves perfect asymptotic tracking of all inputs $\mathbf{w}_l$ that can be generated by $\mathbf{\Sigma_W}(\mathbf{q})$ if and only if it admits a structure as shown in Figure 3. The design of the controller part $\tilde{\mathbf{K}}(\mathbf{q})$ is free as long as it guarantees internal closed-loop stability. The controller structure of Figure 3 can also be understood from the interpolation constraints [12]. Perfect asymptotic tracking of trial-invariant inputs in the range of $\mathbf{W}$ requires the closed-loop sensitivity to have n zeros at $\mathbf{q} = 1$ with input zero directions spanning the range of $\mathbf{W}$. To this end, the loop transfer matrix must have n poles at $\mathbf{q} = 1$ with output pole directions spanning the same subspace of $\mathbb{R}^N$. The multiple poles at $\mathbf{q} = 1$ are created by enclosing $\mathbf{q}^{-1}I_n$ in a positive feedback loop, while the corresponding output pole directions are determined by the blocks on the right-hand side of this loop. Consequently, the output pole directions are determined by the series connection of $\hat{\mathbf{G}}^{-1}\mathbf{W}$ from the controller, and the plant $\mathbf{G}$. Hence, in the case of a perfect model, $\mathbf{G}\hat{\mathbf{G}}^{-1}\mathbf{W} = \mathbf{W}$ and perfect asymptotic tracking of all $\mathbf{w}_l$ generated by $\mathbf{\Sigma_W}(\mathbf{q})$ is achieved. Section III below discusses the effect of a model/plant mismatch, i.e. $\hat{\mathbf{G}} \neq \mathbf{G}$.

D. Norm-optimal Design of $\tilde{\mathbf{K}}(\mathbf{q})$

As noted above, $\tilde{\mathbf{K}}(\mathbf{q})$ in Figure 3 is free as long as it guarantees internal closed-loop stability. In this paper we design $\tilde{\mathbf{K}}(\mathbf{q})$ in accordance with the full-order norm-optimal ILC design [3], [9]. To accomplish this, $\tilde{\mathbf{K}}(\mathbf{q})$ is set equal to a trial-invariant filter with no direct feed-through term, i.e. no current-iteration contribution:

$$\tilde{\mathbf{K}}(\mathbf{q}) = \begin{bmatrix} \mathbf{L} \\ 0 \end{bmatrix} ,$$

whereby the overall ILC $\mathbf{K}(\mathbf{q})$ amounts to

$$\mathbf{K}(\mathbf{q}) \, : \, \begin{cases} \mathbf{x}_{l+1} \, = \, \mathbf{x}_l \, + \, \mathbf{L}\mathbf{e}_l \\ \mathbf{u}_l \, = \, \hat{\mathbf{G}}^{-1}\mathbf{W}\mathbf{x}_l \end{cases} , \tag{7}$$

and the overall closed-loop sensitivity $\mathbf{S}(\mathbf{q})$ is given by

$$\mathbf{S}(\mathbf{q}) = I_N - \mathbf{G}\hat{\mathbf{G}}^{-1}\mathbf{W}\left(\mathbf{q}I_n - I_n + \mathbf{L}\mathbf{G}\hat{\mathbf{G}}^{-1}\mathbf{W}\right)^{-1}\mathbf{L} \, . \tag{8}$$

By substituting $\mathbf{W} = \hat{\mathbf{G}}$, the controller $\mathbf{K}(\mathbf{q})$ given by (7) reverts to the full-order ILC (5).

The matrix $\mathbf{L}$ is computed such that $\mathbf{x}_{l+1}$ minimizes the objective J_{l+1} for given $\mathbf{x}_l$ and $\mathbf{e}_l$:

$$J_{l+1}(\mathbf{x}_{l+1}) = \mathbf{e}_{l+1}^T\Gamma\mathbf{e}_{l+1} + (\mathbf{u}_{l+1} - \mathbf{u}_l)^T\Lambda(\mathbf{u}_{l+1} - \mathbf{u}_l) \, , \tag{9}$$

where the relations $\mathbf{u}_l = \hat{\mathbf{G}}^{-1}\mathbf{W}\mathbf{x}_l$, $\mathbf{u}_{l+1} = \hat{\mathbf{G}}^{-1}\mathbf{W}\mathbf{x}_{l+1}$ and $\mathbf{e}_{l+1} = \mathbf{w} - \mathbf{W}\mathbf{x}_{l+1}$ should be substituted. As in full-order ILC, a quadratic term in $\mathbf{u}_{l+1}$ can be added to J_{l+1}, but it is chosen here not to do so, since this would no longer yield perfect asymptotic tracking [3]. The $\mathbf{x}_{l+1}$ that minimizes (9) equals $\mathbf{x}_{l+1} = \mathbf{x}_l + \mathbf{L}\mathbf{e}_l$ with

$$\mathbf{L} = (\mathbf{W}^T\Gamma\mathbf{W} + \mathbf{W}^T\hat{\mathbf{G}}^{-T}\Lambda\hat{\mathbf{G}}^{-1}\mathbf{W})^{-1}\mathbf{W}^T\Gamma \, . \tag{10}$$

Again, substituting $\mathbf{W} = \hat{\mathbf{G}}$ yields the more commonly known full-order ILC design [3], [9]. As shown in [9], by selecting Γ as a scaled identity matrix, the optimal solution (10) guarantees $\|\mathbf{e}_{l+1}\| \leq \|\mathbf{e}_l\|$ for all $l = 0, 1, \ldots$

III. Reduced-order versus Full-order ILC

This section discusses the advantages and disadvantages of reduced-order ILC, $n < N$, compared to full-order ILC $n = N$. This discussion applies to norm-optimal ILCs, designed according to section II-D with the same objective (9).

A. Zeros and Poles

As reflected in the terminology, the key difference between reduced-order ILCs and full-order ILCs is their order and, consequently, the order of the closed-loop system. With full-order ILC, $\mathbf{S}(\mathbf{q})$ is of order N with N zeros at $\mathbf{q} = 1$. The design (10) guarantees that the N closed-loop poles are stable provided that $\mathbf{G} = \hat{\mathbf{G}}$. That is, (10) guarantees that the eigenvalues of $(I - \mathbf{L}\mathbf{W})$ are contained in the open unit disc. However, the closed-loop poles are generally scattered throughout this disc, which translates into complex and non-intuitive closed-loop dynamics. Reduced-order ILC results in a n'th order closed loop with $\mathbf{S}(\mathbf{q})$ featuring n zeros at $\mathbf{q} = 1$ with input zero directions spanning the range of $\mathbf{W}$. On account of the reduced system order, more intuitive closed-loop dynamics generally result compared to full-order ILC.

In addition to the lower closed-loop order, Eq. (8) reveals that with a reduced-order ILC the dynamic part of $\mathbf{S}(\mathbf{q})$ only manifests for inputs in the range of $\mathbf{L}^T$, leaving inputs in the orthogonal subspace of $\mathbb{R}^N$ unaffected. By selecting Γ as a scaled identity matrix, the range of $\mathbf{L}^T$ corresponds to the range of $\mathbf{W}$, as is clear from Eq. (10).

B. Performance Under Model/Plant Mismatch

In the case of a model/plant mismatch, i.e. $\hat{\mathbf{G}} \neq \mathbf{G}$, the input directions of the sensitivity's n zeros at $\mathbf{q} = 1$ are given by $\mathbf{G}\hat{\mathbf{G}}^{-1}\mathbf{W}$. For reduced-order ILC, these directions generally don't span the range of $\mathbf{W}$ and as a result, the perfect asymptotic tracking of inputs $\mathbf{w}_l$ generated by $\mathbf{\Sigma_W}(\mathbf{q})$ is compromised. For full-order ILC, on the other hand, the subspaces spanned by the columns of $\mathbf{G}\hat{\mathbf{G}}^{-1}\mathbf{W}$ and $\mathbf{W}$ do coincide as they both equal $\mathbb{R}^N$. Hence, even in the presence of a model/plant mismatch, the full-order ILC still yields perfect asymptotic tracking of all trial-invariant inputs; a property sometimes referred to as robustly stable output regulation [5].

C. Stability Under Model/Plant Mismatch

A model/plant mismatch endangers closed-loop stability more in the case of a full-order ILC compared to a reduced-order ILC. The closed-loop poles correspond to the eigenvalues of

$$I_n - \mathbf{L}\mathbf{G}\hat{\mathbf{G}}^{-1}\mathbf{W} = I_n - \mathbf{L}\mathbf{W} + \mathbf{L}\underbrace{(\hat{\mathbf{G}} - \mathbf{G})\hat{\mathbf{G}}^{-1}}_{\boldsymbol{\delta}\hat{\mathbf{G}}}\mathbf{W}\ ,$$

and hence, with full-order ILC, they are affected by all the singular values of the relative plant difference $\boldsymbol{\delta}\hat{\mathbf{G}}$. In the reduced-order case, they are only affected by the singular values of $\boldsymbol{\delta}\hat{\mathbf{G}}$ in the input range $\mathbf{W}$ and output range $\mathbf{L}$. Robust closed-loop stability requires only these singular values to be small, which is less stringent a condition than requiring $\boldsymbol{\delta}\hat{\mathbf{G}}$ to be small. The less stringent stability condition can also be understood from the controller's state-space model (7). A reduced-order ILC only responds to tracking errors $\mathbf{e}_l$ in the range of $\mathbf{L}^T$ and can only generate control signals $\mathbf{u}_l$ in the range of $\hat{\mathbf{G}}^{-1}\mathbf{W}$. This explains respectively the output and input range in which an accurate plant model is required.

D. Time-domain Implementation

This section elaborates on the time-domain formulation of the ILC (7), by reformulating the state-space model (7) as a trial-domain difference equation, similar to (4). The output equation of (7) allows reconstructing $\mathbf{x}_l$ from $\mathbf{u}_l$:

$$\mathbf{x}_l = \mathbf{W}^\dagger \hat{\mathbf{G}}\mathbf{u}_l\ , \tag{11}$$

where $\mathbf{W}^\dagger \in \mathbb{R}^{n\times N}$ is a pseudo-inverse of $\mathbf{W}$, i.e. an arbitrary matrix that satisfies $\mathbf{W}^\dagger\mathbf{W} = I_n$. With the help of (11), the state-space model (7) is reformulated as

$$\mathbf{u}_{l+1} = \hat{\mathbf{G}}^{-1}\mathbf{W}\mathbf{W}^\dagger\hat{\mathbf{G}}\mathbf{u}_l + \hat{\mathbf{G}}^{-1}\mathbf{W}\mathbf{L}\mathbf{e}_l\ . \tag{12}$$

Since the matrices $\hat{\mathbf{G}}^{-1}\mathbf{W}\mathbf{W}^\dagger\hat{\mathbf{G}}$ and $\hat{\mathbf{G}}^{-1}\mathbf{W}\mathbf{L}$ are not Toeplitz and not lower-triangular, the time-domain description of (12) involves noncausal, linear time-varying filters. An additional difference with the full-order ILC (4), is that the matrices $\hat{\mathbf{G}}^{-1}\mathbf{W}\mathbf{W}^\dagger\hat{\mathbf{G}}$ and $\hat{\mathbf{G}}^{-1}\mathbf{W}\mathbf{L} \in \mathbb{R}^{N\times N}$ are of rank n instead of N. This rank-deficiency allows reducing the computational complexity of (12).

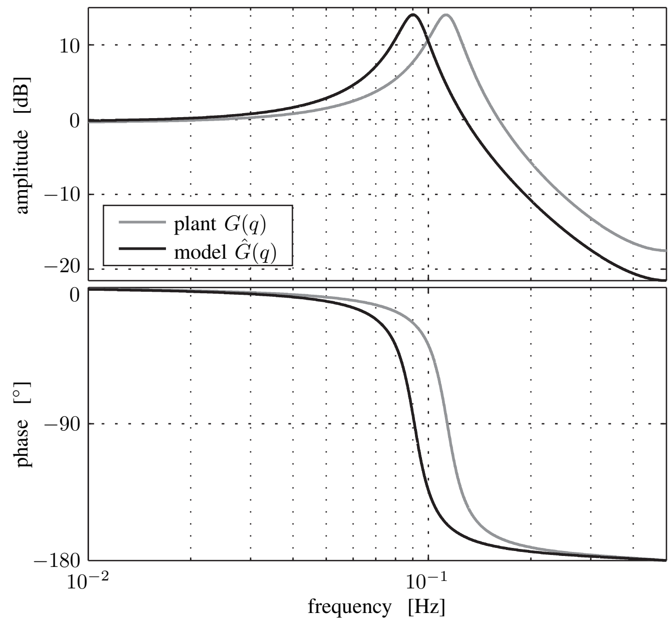

Fig. 4. FRFs of the time-domain plant $G(q)$ and its model $\hat{G}(q)$.

IV. Simulation Results

This section illustrates the differences between reduced-order and full-order ILC by comparing their norm-optimal solutions for the numerical example presented in Section IV-A. This comparison is first performed under the assumption that the actual plant $\mathbf{G}$ equals the plant model $\hat{\mathbf{G}}$ (Section IV-B), while this assumption is dropped in Section IV-C to reveal the different robustness properties of the controllers.

A. Numerical Example

ILC is applied to improve the tracking of a given trial-invariant reference $\mathbf{r}$, which comprises $N = 40$ time samples and corresponds to the black line shown in Figure 6 below. The time-domain plant $G(q)$ and its model $\hat{G}(q)$ are given by:

$$G(q) = \frac{0.436q}{q^2 - 1.412q + 0.867}\ , \tag{13a}$$

$$\hat{G}(q) = \frac{0.292q}{q^2 - 1.592q + 0.892}\ , \tag{13b}$$

and Figure 4 shows their frequency response functions (FRFs). Below, two norm-optimal ILCs are compared, where $\Gamma = I_N$ and $\Lambda = 1.5I_N$ are used in (9). The first ILC, indicated by $\mathbf{K}_{\text{fo}}(\mathbf{q})$, is the full-order solution for $n = N$ and $\mathbf{W}_{\text{fo}} = \hat{\mathbf{G}}$. The corresponding matrix $\mathbf{L} = \mathbf{L}_{\text{fo}}$ is computed according to (10):

$$\mathbf{L}_{\text{fo}} = (\hat{\mathbf{G}}^T\hat{\mathbf{G}} + 1.5I_N)^{-1}\hat{\mathbf{G}}^T\ .$$

The second ILC, indicated by $\mathbf{K}_{\text{ro}}(\mathbf{q})$, is the reduced-order solution for $n = 1$ and $\mathbf{W}_{\text{ro}} = \mathbf{r}/\|\mathbf{r}\|$. According to (10), the corresponding matrix $\mathbf{L} = \mathbf{L}_{\text{ro}}$ equals:

$$\mathbf{L}_{\text{ro}} = \left(1 + 1.5\mathbf{W}_{\text{ro}}^T\hat{\mathbf{G}}^{-T}\hat{\mathbf{G}}^{-1}\mathbf{W}_{\text{ro}}\right)^{-1}\mathbf{W}_{\text{ro}}^T\ =\ \rho\mathbf{W}_{\text{ro}}^T\ .$$

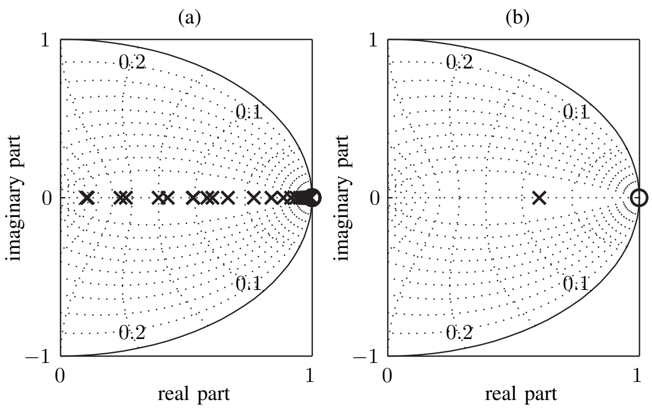

Fig. 5. Poles and zeros of the closed-loop sensitivity for the plant model $\hat{\mathbf{G}}$ and (a) the full-order ILC $\mathbf{K}_{\text{fo}}(\mathbf{q})$; and (b) the reduced-order ILC $\mathbf{K}_{\text{ro}}(\mathbf{q})$.

B. Evaluation for $\hat{\mathbf{G}}$

This section compares $\mathbf{K}_{\text{fo}}(\mathbf{q})$ and $\mathbf{K}_{\text{ro}}(\mathbf{q})$ for the plant model $\hat{\mathbf{G}}$, or equivalently, temporary assumes that the actual plant $\mathbf{G}$ equals the model $\hat{\mathbf{G}}$. The closed-loop sensitivities corresponding to $\mathbf{K}_{\text{fo}}(\mathbf{q})$ and $\mathbf{K}_{\text{ro}}(\mathbf{q})$ are respectively indicated by $\mathbf{S}_{\text{fo}}(\mathbf{q})$ and $\mathbf{S}_{\text{ro}}(\mathbf{q})$, and for $\mathbf{G} = \hat{\mathbf{G}}$ they equal

$$\begin{aligned} \mathbf{S}_{\text{fo}}(\mathbf{q}) &= I_N - \hat{\mathbf{G}}(\mathbf{q}I_N - I_N + \mathbf{L}_{\text{fo}}\hat{\mathbf{G}})^{-1}\mathbf{L}_{\text{fo}} , \\ \mathbf{S}_{\text{ro}}(\mathbf{q}) &= I_N - \mathbf{W}_{\text{ro}}(\mathbf{q} - 1 + \rho)^{-1}\rho\mathbf{W}_{\text{ro}}^T . \end{aligned}$$

Figure 5 shows the corresponding poles and zeros. Sensitivity $\mathbf{S}_{\text{fo}}(\mathbf{q})$ has $N = 40$ zeros at $\mathbf{q} = 1$, and N poles corresponding to the eigenvalues of $I_N - \mathbf{L}_{\text{fo}}\hat{\mathbf{G}}$, which are scattered throughout the open unit disc. The reduced-order result $\mathbf{S}_{\text{ro}}(\mathbf{q})$, on the other hand, has only $n = 1$ zero at $\mathbf{q} = 1$ and $n = 1$ pole at $\mathbf{q} = 1 - \rho$. Moreover, the dynamic part of $\mathbf{S}_{\text{ro}}(\mathbf{q})$ only manifests for inputs in the range of $\mathbf{W}_{\text{ro}}$, producing an output signal in the same subspace of $\mathbb{R}^N$. The input directions orthogonal to $\mathbf{W}_{\text{ro}}$ are not affected by the reduced-order ILC. That is: $\mathbf{S}_{\text{ro}}(\mathbf{q})\mathbf{W}_{\text{ro}}^{\perp} = \mathbf{W}_{\text{ro}}^{\perp}$, with the columns of $\mathbf{W}_{\text{ro}}^{\perp} \in \mathbb{R}^{N\times(N-n)}$ spanning the null-space of $\mathbf{W}_{\text{ro}}$.

Figure 6 shows for both ILCs the evolution of the plant output $y_l(k)$ for $l = 0, 1, \ldots, 9$, while the black curve corresponds to $r(k)$. Since both ILCs achieve perfect asymptotic tracking, $\lim_{l\to\infty} y_l(k) = r(k)$. The black curves in Figure 8 show the evolution of the norm of the corresponding tracking error, i.e. $\|\mathbf{e}_l\|$, as a function of l. For the full-order ILC $\mathbf{K}_{\text{fo}}(\mathbf{q})$ the transient tracking behavior is affected by all the closed-loop poles, and since some of these poles lie closely to the unit circle, $\mathbf{S}_{\text{fo}}(\mathbf{q})$ features very slow convergence of some characteristics of $r(k)$. As revealed by Figure 6(a), the overall behavior of $r(k)$ is accurately tracked within a few iterations, while slow convergence is observed on the last time samples where $r(k) = 1$. Figure 8(a) shows that this results in a fast initial decrease of $\|\mathbf{e}_l\|$, which levels off as l increases.

For the reduced-order ILC $\mathbf{K}_{\text{ro}}(\mathbf{q})$ the transient tracking behavior is determined by the $n = 1$ closed-loop poles and zeros. Hence, $\|\mathbf{e}_l\|$ decays according to $(1-\rho)^l$, as is confirmed in Figure 8. Since in addition, the input and output directions of the dynamic part of $\mathbf{S}_{\text{ro}}(\mathbf{q})$ coincide, both the tracking error $\mathbf{e}_l$ and the plant output $\mathbf{y}_l$ are proportional to $\mathbf{r}$. This is clearly observed in Figure 6(b).

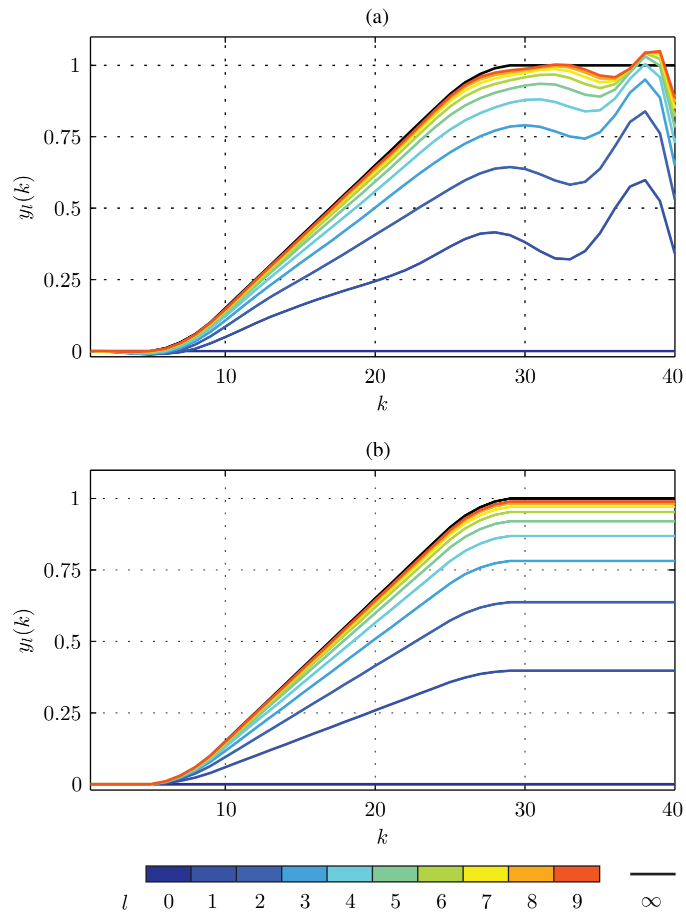

Fig. 6. Evolution of the plant output $y_l(k)$ as a function of l for the plant model $\hat{\mathbf{G}}$ and (a) the full-order ILC $\mathbf{K}_{\text{fo}}(\mathbf{q})$; and (b) the reduced-order ILC $\mathbf{K}_{\text{ro}}(\mathbf{q})$. The black line corresponds to $r(k)$, and for both ILCs $\lim_{l\to\infty} y_l(k) = r(k)$.

C. Evaluation for $\mathbf{G}$

In this section, the assumption that $\mathbf{G} = \hat{\mathbf{G}}$ is dropped and the ILCs are evaluated for $\mathbf{G}$ instead of $\hat{\mathbf{G}}$.

As argued in Section III, a model/plant mismatch endangers closed-loop stability more in the case of a full-order ILC compared to a reduced-order ILC. This section confirms this statement, since evaluated for $\mathbf{G}$ the closed-loop system with $\mathbf{K}_{\text{fo}}(\mathbf{q})$ is unstable, while it is stable for $\mathbf{K}_{\text{ro}}(\mathbf{q})$. Figure 7 shows the corresponding pole-zero maps of the closed-loop sensitivity. Comparison with Figure 5 reveals that for $\mathbf{K}_{\text{fo}}(\mathbf{q})$ some closed-loop poles are significantly affected by the model/plant mismatch, while for $\mathbf{K}_{\text{ro}}(\mathbf{q})$ this effect is minor.

For $\mathbf{K}_{\text{ro}}(\mathbf{q})$, the closed-loop sensitivity still has a zero at $\mathbf{q} = 1$, but due to $\mathbf{G} \neq \hat{\mathbf{G}}$, the corresponding input direction is no longer aligned to $\mathbf{r}$. Consequently, $\mathbf{r}$ is no longer perfectly tracked for $l \to \infty$. This is confirmed by the grey curve in Figure 8(b), which shows the corresponding $\|\mathbf{e}_l\|$

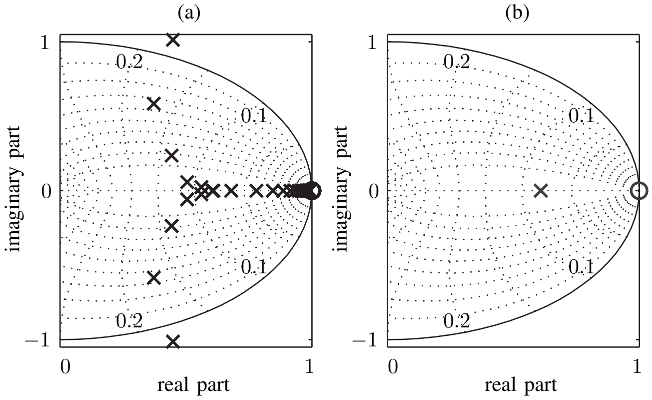

Fig. 7. Poles and zeros of the closed-loop sensitivity for the plant $\mathbf{G}$ and (a) the full-order ILC $\mathbf{K}_{\mathrm{fo}}(\mathbf{q})$; and (b) the reduced-order ILC $\mathbf{K}_{\mathrm{ro}}(\mathbf{q})$.

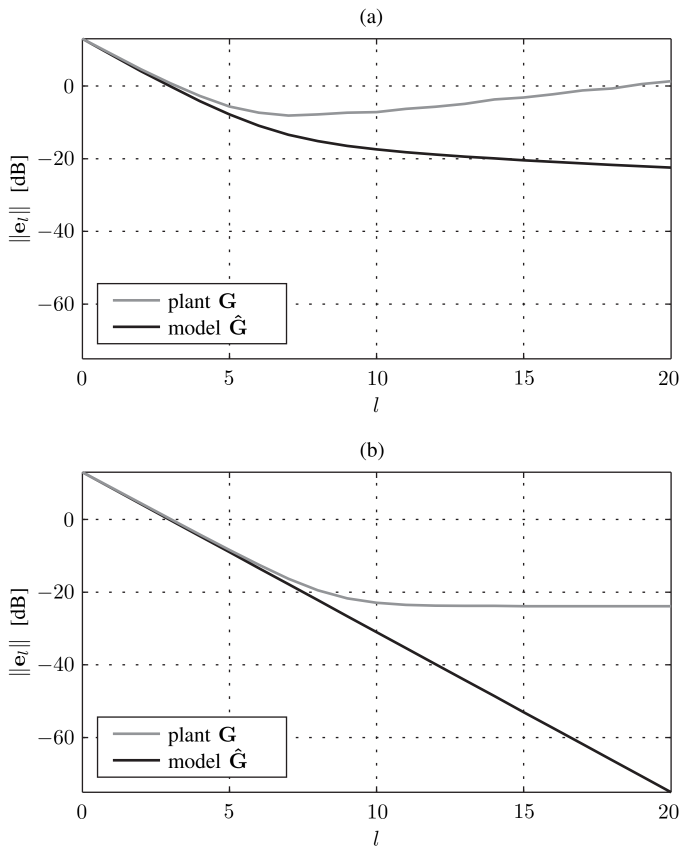

Fig. 8. Evolution of $\|\mathbf{e}_l\|$ as a function of l for (a) the full-order ILC $\mathbf{K}_{\mathrm{fo}}(\mathbf{q})$; and (b) the reduced-order ILC $\mathbf{K}_{\mathrm{ro}}(\mathbf{q})$. The black lines relate to the closed-loop system with the plant model $\hat{\mathbf{G}}$, while the grey lines relate to the actual plant $\mathbf{G}$.

as a function of l. The grey curve in Figure 8(a) confirms the closed-loop instability when $\mathbf{K}_{\mathrm{fo}}(\mathbf{q})$ is evaluated for $\mathbf{G}$.

V. Conclusions

This paper presents a novel ILC design that allows exploiting the direction of the input signals in addition to their trial-invariance. To this end, a reduced-order trial-invariant signal generator is included in the ILC, whereby the controller order is less than the number of samples per trial. The ILCs are therefore called reduced-order ILCs, while the current ILCs are referred to as full-order. The reduced-order ILCs are here designed in accordance with the common full-order norm-optimal ILC design.

It is illustrated that reduced-order norm-optimal ILCs result in simpler (more intuitive) learning dynamics and a more desirable transient learning behavior compared to full-order norm-optimal ILCs. In addition, a model/plant mismatch affects both types of ILC in a different way: with a reduced-order norm-optimal ILC the closed-loop stability is more robust to plant model errors then with a full-order norm-optimal ILC. On the other hand, if robust stability is achieved, the robust performance is slightly better for full-order norm-optimal ILCs. In future work, the robust reduced-order ILC design for plant uncertainties will be considered.

VI. Acknowledgement

Goele Pipeleers is a Postdoctoral Fellow of the Research Foundation Flanders (FWO-Vlaanderen). This work benefits from K.U.Leuven-BOF PFV/10/002 Center-of-Excellence Optimization in Engineering (OPTEC), the Belgian Programme on Interuniversity Attraction Poles, initiated by the Belgian Federal Science Policy Office (DYSCO), research project G.0422.08 of the Research Foundation Flanders (FWO-Vlaanderen), and project IWT-SBO 80032 (LeCoPro) of the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen).

References

[1] S. Arimoto, S. Kawamura, and F. Miyazaki, "Bettering operation of robots by learning," *Journal of Robotic Systems*, vol. 1, no. 2, pp. 123–140, 1984.

[2] K. Moore, *Iterative Learning Control for Deterministic Systems*. London, Great Brittain: Springer-Verlag, 1993.

[3] D. Bristow, M. Tharayil, and A. Alleyne, "A survey of iterative learning control," *IEEE Control Systems Magazine*, vol. 26, no. 3, pp. 96–114, 2006.

[4] B. A. Francis and W. M. Wonham, "The internal model principle of control theory," *Automatica*, vol. 12, no. 5, pp. 457–465, 1976.

[5] A. Saberi, A. A. Stoorvogel, and P. Sannuti, *Control of Linear Systems with Regulation and Input Contraints*. London, Great Britain: Springer-Verlag, 2000.

[6] K. Moore and Y. Chen, "An investigation on the monotonic convergence of an high-order iterative learning update law," in *Proc. of the 15th IFAC World Congress*, Barcelona, Spain, July 2002.

[7] ——, "A separative high-order framework for monotonic convergent iterative learning controller design," in *Proc. of the 2003 American Control Conference*, Denver, Co., June 2003, pp. 3644–3649.

[8] K. Moore, "A matrix-fraction approach to higher-order iterative learning control: 2-D dynamics through repetition-domain filtering," in *Proc. of the 2nd International Workshop on Multidimensional (nD) Systems*, Lower Selesia, Poland, June 2000, pp. 99–104.

[9] D. Owens and J. Hätönen, "Iterative learning control – an optimization paradigm," *Annual Reviews in Control*, vol. 29, no. 1, pp. 57–70, 2005.

[10] K. Moore, "Multi-loop control approach to designing iterative learning controllers," in *Proc. of the 37th IEEE Conference on Decision and Control*, Tampa, Fl., December 1998, pp. 666–671.

[11] D. Owens, E. Rogers, and K. Moore, "Analysis of linear iterative learning scheme using repetitive process theory," *Asian Journal of Control*, vol. 4, no. 1, pp. 90–98, 2002.

[12] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control - Analysis and Design*. Chichester, West Sussex, England: John Wiley & Sons, 2005.