UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering, Science and Mathematics
School of Electronics and Computer Science

The Robust Stability of Iterative Learning Control

by

Richard Stephen Bradley

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

21st September, 2010

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE & MATHEMATICS
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

THE ROBUST STABILITY OF ITERATIVE LEARNING CONTROL

by Richard Stephen Bradley

This thesis examines the notion of the long term robust stability of iterative learning control (ILC) systems engaged in trajectory tracking, using a robust stability theorem based on a biased version of the nonlinear gap metric. This is achieved through two main results:

The first concerns the establishment of a nonlinear robust stability theorem, where signals are measured relative to a given trajectory. Although primarily motivated by ILC, the theorem provided is applicable to a wider range of problems. This is due to its development being made independently of any particular signal space, provided the space is furnished with a definition of causality. The theorem's formulation therefore permits its implementation on single- or multi-dimensional problems in a variety of different settings. Necessitated by an ILC constraint concerning reference signals, the trajectory that stability is measured relative to must often lie outside the signal space that is chosen. The robust stability theorem is therefore devised to address signals that lie in extended signal spaces throughout. Additionally, the theorem is applicable to nonlinear systems, and it is shown that the biased gap measure collapses to the standard nonlinear gap measure when the bias is set to zero. It is also shown to collapse to the classical linear gap when restricting the analysis to certain linear systems.

The second result applies the robust stability theorem to ILC using a 2D signal space. Initially the subject of ILC is reviewed and some of the problems associated with controllers are described; in particular the issues of long-term stability and the criteria for convergence. ILC algorithms expressed in the 'lifted system' or 'supervector' formulation are discussed and then analysed using the biased robust stability theorem. Results are presented regarding the use of filtering in ILC algorithms to aid robustness, and also the robustness of inverse model-based techniques. The robust stability tool applied to ILC in this thesis is not restricted in its analysis to algorithms which are 'causal' (in an ILC sense); and is based on a general unstructured uncertainty model in contrast to the existing literature, whereby uncertainties are typically constrained to additive, multiplicative or parametric models.

Contents

List of Figures

<u>DECLARATION OF AUTHORSHIP</u>

I, Richard Stephen Bradley, declare that the thesis entitled

The Robust Stability of Iterative Learning Control

and the work presented in the thesis are both my own, and have been generated by me as a result of my own original research. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University;

- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;

- Where I have consulted the published work of others, this is always clearly attributed;

- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;

- I have acknowledged all main sources of help;

- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;

- None of this work has been published before submission.

Signed

Date:

Chapter 1

Introduction

In tracking tasks, controllers are developed such that the output of a plant can closely follow a given trajectory. A controller will be designed to produce the input signals required for the plant to give this correct output signal. The controller may also attempt to minimise other properties such as the magnitude of actuator signals, the time taken for the output to reach its desired value or the effect of disturbances.

Classical control is based on the principle of feedback, where output errors are fed, through a controller, back to the input. This allows more accurate control of the output signal and a greater ability to reduce the effects of unwanted disturbances and plant uncertainty. The cost of this feedback is the possibility of causing instability, where bounded inputs to the system are able to produce unbounded outputs, even for a priori stable systems.

Within classical control, once a controller is implemented on a deterministic, disturbance free plant the closed-loop system is fixed and deterministic (assuming time invariant system dynamics) — the same signal applied to the system in the same initial state will result in the same output. When a process is repeated it always leads to identical errors being present at the output.

Intelligent controllers attempt to improve this error profile by adapting as they are running; the controller adjusting itself in order to improve system performance. *If a classical controller always produces the same error when a process is performed, can learning this error and adapting the control reduce it?*

The work of Wiener on cybernetics and Turing on artificial intelligence in the 1940s developed the idea of machines mimicking man's ability to learn. Following this, Rosenblatt's invention of the perceptron led to the development of adaptive control in the mid-1950s; the original aim being the solution of problems such as the changing dynamics of aircraft in flight (Åström and Wittenmark, 1989). With adaptive control, the control scheme is adjusted to maintain stability as the plant changes.

This led to other control strategies with different goals and plants in mind. Repetitive control (born in the early-1980s) had the aim of controlling systems such as power supplies and hard disks, where disturbances and reference signals are periodic. By adding an extra constraint — that the plant is returned to the same state at the start of each period — iterative learning control is characterised.

1.1 Iterative Learning Control

Humans acquire skills by repeating tasks and learning from mistakes. Iterative learning control (ILC) attempts to employ this approach to improve a system's performance as it repeats a task. The system is provided with an input signal in an attempt to follow a given trajectory, and the resulting error is measured. After a finite period of time the system is reset and an 'improved' input signal is calculated using data from the previous trial. The system then attempts the same task using the new input signal, the error is again measured, and the process continues. This update at each iteration aims to reduce the error at the output.

ILC is distinct from other learning controllers in the way that algorithms exploit repetitiveness. The classical ILC requirements are that the desired reference trajectory is identical for each iteration, that the system state is always returned to the same point for the start of the next iteration, and that the plant dynamics do not change. At first these appear to heavily restrict the possible applications of ILC, but in many of the manufacturing industries robots are required to build the same products again and again. Improving the accuracy of these robots using ILC could reduce the amount of time required for each product build, increasing the cost efficiency of the robots. It is also possible that similar performance could be achieved with less expensive hardware. As an example, ILC could compensate for a lack of stiffness in robot arms.

One of the theoretical demands within ILC stems from the two-dimensional attributes of ILC systems. Within classical feedback control, all the system information is propagating in one direction: time. As time progresses the output develops depending on the input in a manner governed by the feedback loop. For ILC signals develop in both time and iteration directions. The system runs for a finite period of time and resets. A separate rule then defines how the results from that trial will affect the system along the following trial. Two control laws are therefore entwined together, increasing the complexity of any analysis. This also renders traditional definitions of stability insufficient. As an example, within any finite length of time a linear plant's output will remain bounded, regardless of any conventional interpretation of its stability. Another major problem is that many ILC systems are non-linear by their very nature and so analysis can be difficult even for specific cases.

Learning controllers are also prone to long term stability problems: the error decreasing

for many iterations before suddenly increasing, leading to instability. Figure 1.1 shows the error norm per iteration for a simulated ILC system. The error starts by decreasing but then diverges after a number of iterations. (The system in question is an inverse model based ILC system, where there exists a small perturbation between the plant being controlled and the model used to develop the learning controller. Further details of the simulation are given in Chapter 6.)

This type of instability is poorly understood and so is a serious drawback of using learning control. A theory is therefore needed to guarantee stability of this type of controller. This thesis considers the possibility of using the gap metric as a robust stability tool to examine this.
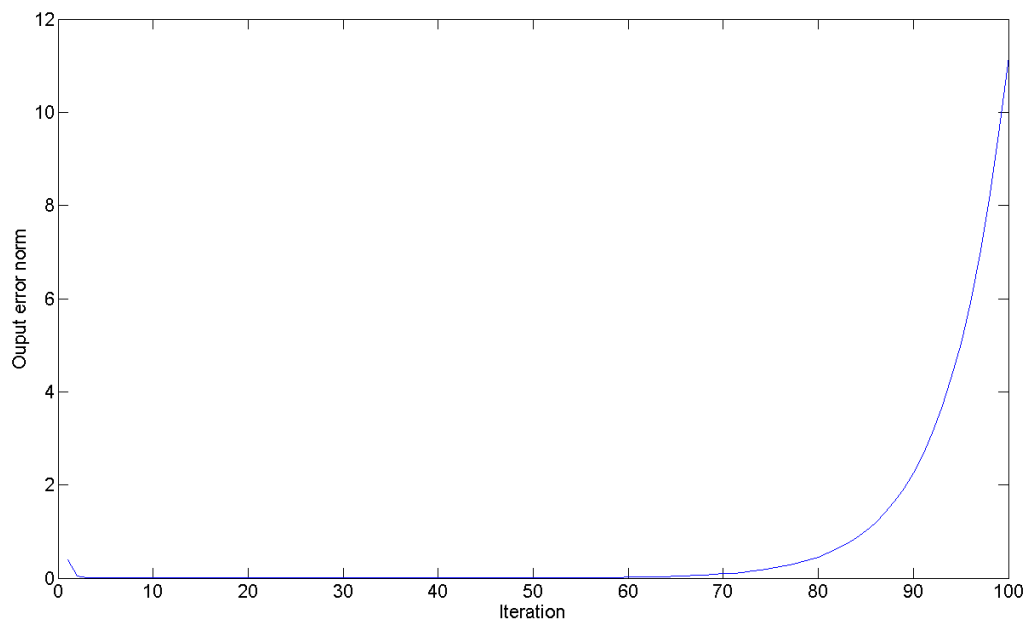
FIGURE 1.1: Output error per iteration for an ILC system suffering from long term instability. (For more detail see Chapter 6.)

1.2 Robust Stability

In order to design a controller for a plant a mathematical model of the plant is usually derived and then a controller built to fit this model. The closed-loop system given by this model and controller will be designed to meet certain desired criteria; in this case stability. However, no model is a perfect reproduction of its parent plant, and so the modelled closed-loop system must be sufficiently robust to accept the perturbations encountered when the controller is implemented on the real plant.

If a system is stable, a stability margin depicts the amount that the plant can differ

and still have closed-loop stability guaranteed. The gap metric introduces a method of measuring the distance between two plants, and so provides a quantitative description of a set of plants able to be stabilised by a robust controller.

Consider the physical plant P_1 and its model P. The Controller C is developed such that the feedback system $[P, C]$ given in Figure 1.2 is stable. The aim is then to prove that P can be replaced by P_1 and the closed-loop $[P_1, C]$ is also stable. This is guaranteed when the gap between P and P_1 is less than the robust stability margin of $[P, C]$.
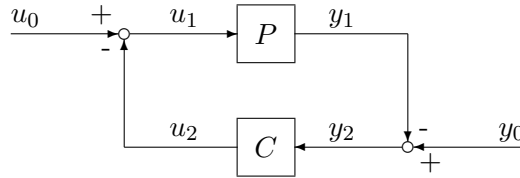
FIGURE 1.2: Feedback configuration $[P, C]$

As an example, Figure 1.3 (from Ratcliffe, 2005) shows the Bode plot for the x-axis of a gantry robot discussed in Example 2.1 later. The blue line shows the gain and phase plots generated by feeding signals of various frequencies into the plant and measuring the corresponding output. A linear model is then fitted to these results to approximate the real plant. The red line shows the plot for a 1^{st} order plant model. This model was then used to develop a controller; however, any controller that performed well on the model must be robust enough to perform well on the physical plant.

This is where the gap metric can be employed. It can be used to prove that the model captures enough information so that any controller that works on the model will also work on the real plant it represents. This case is examined more fully in Example 2.1 and again in Chapter 6.

Here an affiliation of the gap metric and ILC seems apt. A learning controller is often used when it is desired to control a poorly modelled plant; using the controller to compensate for the user's lack of knowledge of the plant. The gap metric could be used, with the stability margin around the modelled system, to determine whether the real system is close enough to still be guaranteed stable.

The gap metric is well established within classical control applications. However, as discussed, ILC has a particular 2D structure to which the standard robust stability theorem does not apply. Therefore the gap metric requires some re-formulation in order to link the two subjects in a meaningful way.

The aim of this thesis is therefore to present work linking the gap metric to iterative learning control such that the robust stability of ILC algorithms engaged in trajectory tracking can be proven in a two-dimensional setting. This is done through the use of
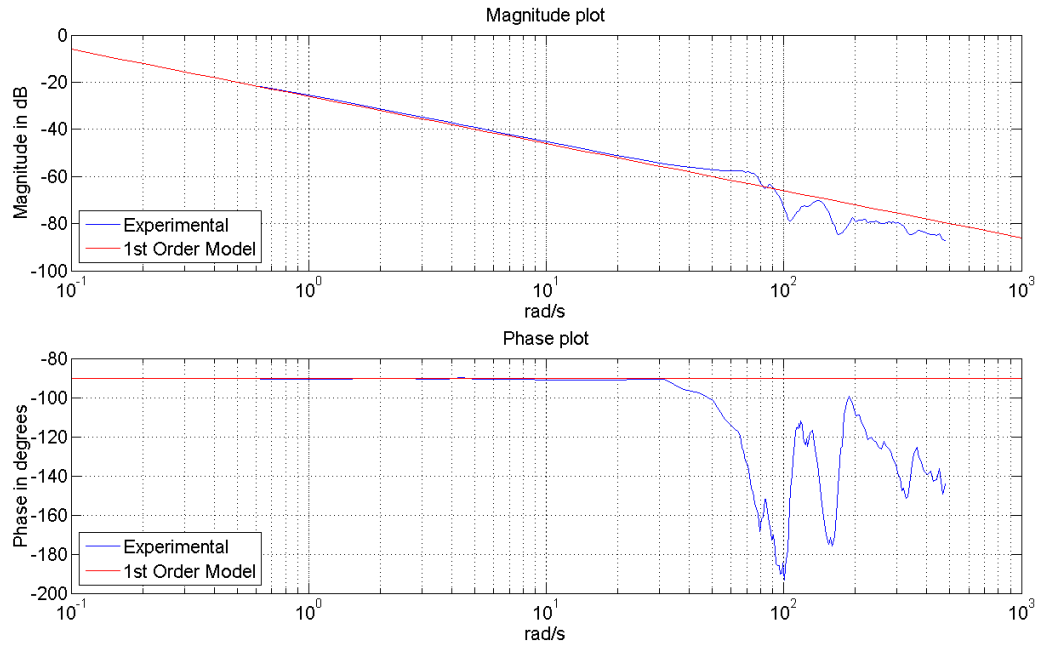
FIGURE 1.3: Bode plot for measured system and 1st order model

a biased gap to measure the distance between two plants, a 2D biased robust stability margin, and a robust stability theorem that brings them both together.

Once this is achieved the results are used along with an example ILC system to demonstrate issues such as the trade-off between robustness and performance.

1.3   Thesis Overview

This thesis begins with an introduction to robust stability. This covers the preliminary notation used within the thesis, including details on signal spaces and norms. The motivation and concept of robust stability is then given and the small gain theorem discussed, leading to robust stability results based on various uncertainty models, namely additive, multiplicative, inverse multiplicative and coprime factor uncertainties.

Following this, Chapter 3 presents the subject of ILC. This begins with a brief history of the area and an explanation of some of the algorithms. Continuous- and discrete-time ILC is examined along with some results on stability and convergence. The lifted system form is also introduced as it has a large bearing on the approach used later in the thesis. The problem of the long term instability of iterative learning control schemes is investigated including some of the possible causes and available solutions. Current papers on robustness for ILC are then reviewed and some examples of physical ILC implementations given.

Chapter 4 presents the gap metric. Some of the theorems pertaining to the gap metric

as a measure for robust stability are portrayed in their linear and non-linear settings, and a theorem given for the stability of systems where a zero input does not give a zero output. Problems that occur when the gap is applied to ILC are expounded, particularly the problem of repeating signals being infinite under certain norms. Although ILC is not specifically addressed in the chapter all of the work is generalised to enable its application to any signal space containing a compatible definition of causality, with a 2D signal space being the main concern. The crux of the chapter is a theorem concerning the robust stability of systems engaged in trajectory tracking, without restrictions on linearity.

In Chapter 5 previous work is detailed relating the gap and ILC. Then the robust stability theorem of Chapter 4 is used in a 2D setting suitable for the examination of ILC, investigating the robustness of a generalised ILC system engaged in trajectory tracking. The robust stability margin of ILC systems in lifted system form is calculated. This is done in a 2D setting and related to a 2D gap measure, providing a complete 2D robust stability theorem pertaining to lifted system ILC. This 2D gap is related back to the standard linear gap enabling ILC robustness to be addressed in a similar way to standard robustness problems.

Certain implications that follow from the 2D robust stability result are given in Chapter 6, including a view on inverse model based ILC and the use of filtering to increase robustness. Examples are given based on simulations of a modelled plant. The chapter attempts to explain in a mathematical framework some of the processes that have been observed on physical ILC systems, including the problem of long term instability.

The thesis is then concluded and some avenues down which the work can be continued are given.

Chapter 2

An Introduction to Robust Stability

When a controller is designed for a plant, a mathematical model of the plant is usually first developed. This enables the designer to ensure that the controller will be capable of stabilising the plant before it is physically implemented. Many plants are modelled as finite-dimensional, second order differential equations. If the models given by these equations are non-linear then they may also be simplified to linear time invariant models to facilitate controller design or simulation. This then becomes the nominal plant and a controller is built to stabilise it, leaving a *nominally stable closed-loop system.*

There are usually other performance requirements, beyond stability, that influence the design of a controller. Often the plant's output is required to follow a reference trajectory or a certain variable is required to be minimised, and the controller is designed such that the system achieves these objectives. This gives rise to the concept of performance: a measurement of how well the system is capable of accomplishing the given task. In order to quantify this performance some method of calculating the size of signals therefore needs to be defined. This takes the form of a norm on the inputs and outputs.

One way of improving the performance of a plant is to introduce feedback. Feedback techniques monitor the output of the plant (compare it with the desired output if required) and then transmit a signal back to the input. Hopefully this signal will then assist in keeping the output close to its objective. The result can increase the robustness and performance of a closed-loop system by removing some of the output's dependance on the dynamic behaviour of the plant. The real plant should therefore be able to differ slightly from it's nominal model without significantly affecting the performance.

However, if the real plant and the model differ excessively then a feedback system may become unstable, even if both the plant are controller are stable by themselves. It is therefore desired to know in what ways they differ, and by how much.

No matter how good a model of a plant is there will always be some mismatch between the modelled and physical systems. These arise in a variety of different ways, for example:

i) a non-linear plant may be assumed to be linear within the region of operation

ii) time delays may be considered negligible and left out of the model

iii) a high order system may be represented with a lower order one to simplify calculations

iv) the plant dynamics may change over time, and this change may be left out of the model

v) the plant may have simply been measured badly, resulting in an inaccurate model

vi) it is physically impossible to actually measure the high frequency behaviour of a plant by identification, and hence there is always a degree of model mismatch at higher frequencies.

These simplifications of the plant may produce equations that are easier to work with but also add uncertainty to the model which, as explained above, could compromise stability under feedback control. It is therefore prudent to work with a plant model that also incorporates a description of the uncertainty between it and the real plant.

Example 2.1. *An example of some of the above is given here using the gantry robot from Ratcliffe (2005). This robot is explained in more detail in Chapter 6 where it is used to examine and demonstrate the gap metric work developed in Chapters 4 and 5.*

When the robot was analysed it was deemed sufficient to create a linear model. In order to do this frequency response tests were undertaken for a range of frequencies from 0.628rad/s to 477rad/s. This is shown in Figure 2.1. A least-mean-square optimisation was then used to find an appropriate transfer function to fit the results (see Ratcliffe, 2005, for details). Although this transfer function was 21^{st} order it was still not exactly equal to the results gathered (also given in Figure 2.1). This was then reduced to a 1^{st} order model consisting of a gain, approximated from the low frequency gain of the measured plant, and an integrator; which was then used for simulation. This model adequately reflects the real system at low frequencies but fails to capture some of the higher frequency effects (see Figure 2.2).[1]

When the apparatus was disassembled and moved to a different location the frequency response tests were re-run and the new Bode plot (Figure 2.3) was found to differ significantly from the old at high frequency. This is due to the interaction between the robot and its mounting — on the new mounting the plant had different resonances than it did

---

[1]The derivations of both 1^{st} and 21^{st} order models are given in more detail in Chapter 6. Equation 6.1 describes the 21^{st} order model and equation 6.3 the 1^{st} order model.
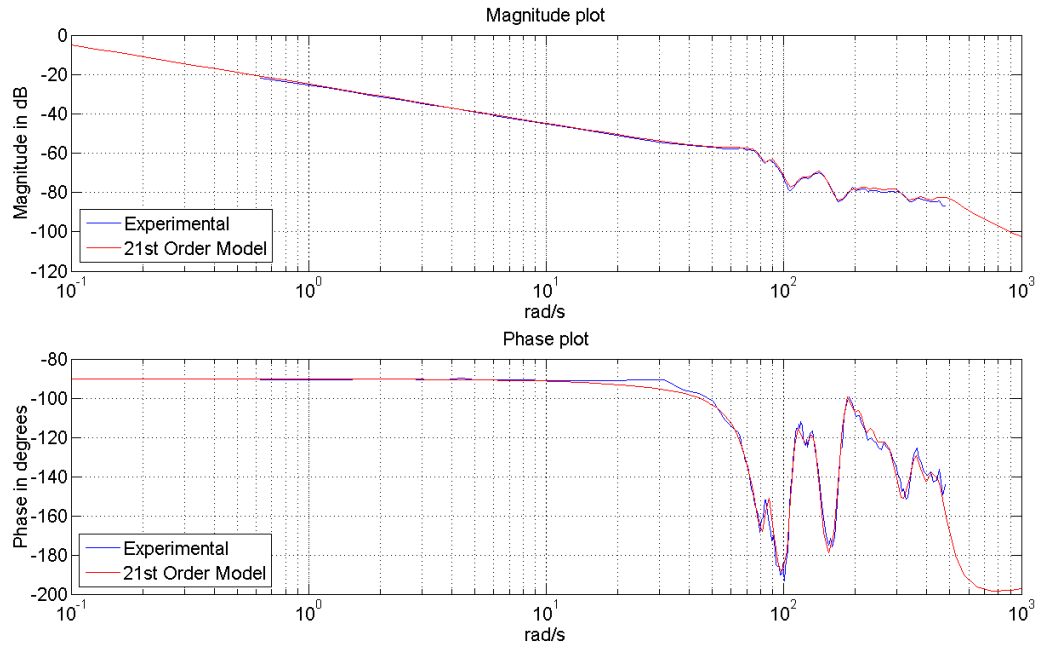
FIGURE 2.1: Gantry robot measurement data and 21st order model
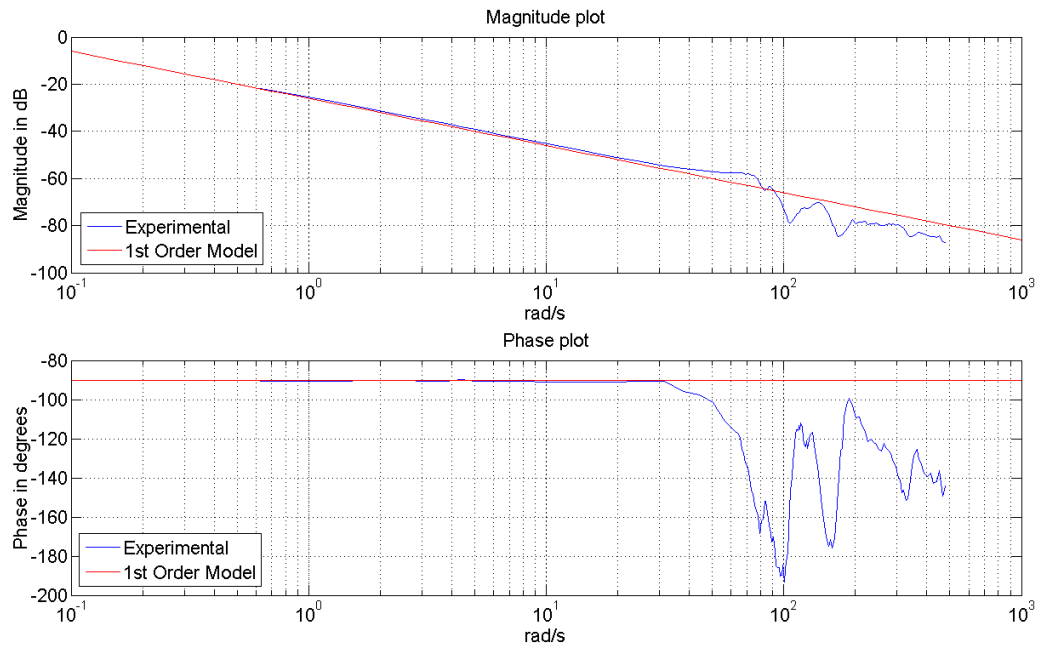
FIGURE 2.2: Gantry robot measurement data and 1st order model

on the old. In both cases the same mathematical model was used to describe the plant as its low frequency behaviour was preserved. The higher frequency behaviour that is failed to be captured by this model in each case could then be regarded as the unmodelled dynamics.

In situations such as this it would be useful to have a notion of how far the model can be simplified and still be an adequate enough representation, and whether or not a single model could describe both plants. Robustness analysis attempts to address this issue.
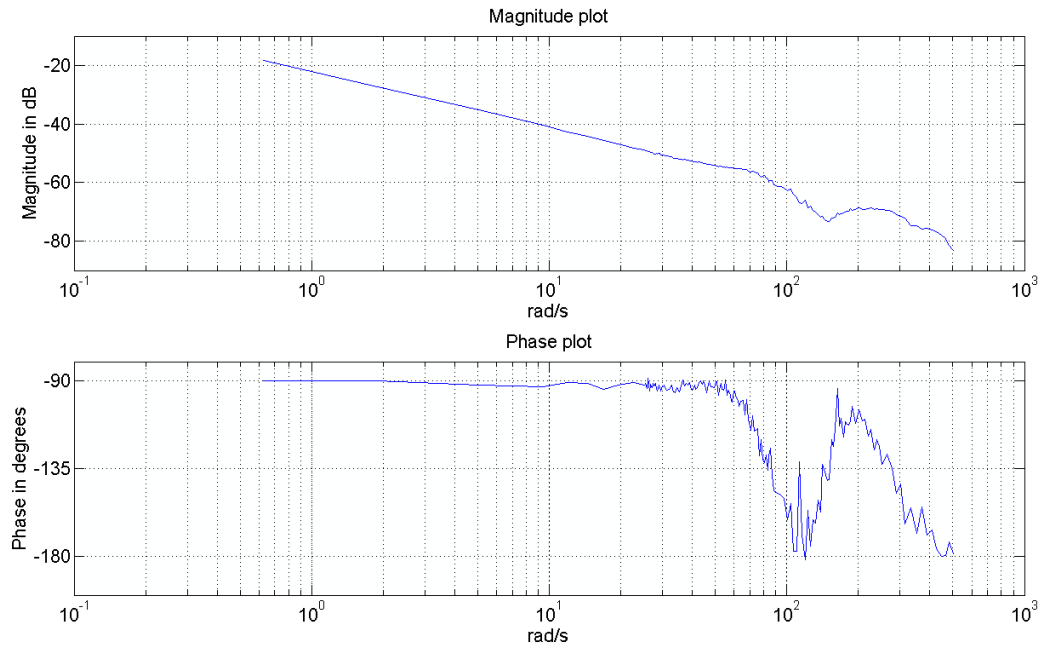
FIGURE 2.3: Gantry robot measurement data after moving

One way to improve on this is to develop a description of the difference between two plants in some abstract sense. If this distance (referred to as a 'perturbation') is small enough then it would be hoped that both plants are able to be stabilised by a single controller. Many different models exist to describe these perturbations, and some will be described in this chapter. Before this can be done the notation used within this thesis will be introduced.

2.1 Notation

2.1.1 An introduction to metric spaces

When dealing with signals and operators it is often important to have some measure of their size and a notion of distance between them. For this reason, this section will briefly introduce the concept of metric and normed vector spaces.

A metric space is a set with a metric defined on it. This metric provides a notion of distance between any two elements within the set. Formally a metric space is the pair (X, δ), where X is the set and δ the defined metric. This notation will usually be dropped as the set and metric are generally obvious. Here $\delta \colon X \times X \to \mathbb{R}$ is a function obeying the following axioms[2]:

For all $x, y, z \in X$

  i) $\delta(x, y) \geq 0$

  ii) $\delta(x, y) = 0 \qquad \Leftrightarrow \qquad x = y$

  iii) $\delta(x, y) = \delta(y, x)$

  iv) $\delta(x, z) \leq \delta(x, y) + \delta(y, z)$.

A vector space V over the field $\mathbb{R}$ is a nonempty set of elements (vectors) combined with two operations: vector addition and scalar multiplication; where addition $V \times V \to V$ is denoted by $x + y$ where $x, y \in V$, and multiplication $\mathbb{R} \times V \to V$ is denoted by ax where $a \in \mathbb{R}$ and $x \in V$. These obey the following axioms.

For all $x, y, z \in V$ and $a, b \in \mathbb{R}$

  i) $x + y = y + x$ (addition is commutative)

  ii) $x + (y + z) = (x + y) + z$ (addition is associative)

  iii) $\exists 0 \in V$ such that $x + 0 = x$ (zero vector)

  iv) $\exists -x \in V$ such that $x + (-x) = 0$ (additive inverse)

  v) $a(x + y) = ax + ay$ (distributive over vector addition)

  vi) $(a + b)x = ax + bx$ (distributive over field addition)

  vii) $a(bx) = (ab)x$ (compatability with scalar multiplication)

  viii) $\exists 1 \in \mathbb{R}$ such that $1x = x$.   (identity)

Normed vector spaces $(V, \| \cdot \|)$ are vector spaces together with the notion of 'size' as measured by a norm, denoted $\| \cdot \|$, which obeys the following axioms:

For all $x, y \in V$ and $\alpha \in \mathbb{R}$

  i) $\|x\| \geq 0$

  ii) $\|x\| = 0 \qquad \Leftrightarrow \qquad x = 0$

---

[2] $X \times X$ denotes a pair of points in X. If $x, y \in X$ then $(x, y) \in X \times X$

iii) $\|\alpha x\| = |\alpha| \|x\|$

iv) $\|x + y\| \le \|x\| + \|y\|$.

The distance between two elements is given by $\delta(x, y) = \|x - y\|$; this forms a metric. Where different normed spaces are used, subscript after the $\|\cdot\|$ or δ will denote the space or distance measure in question.

2.1.2 Lebesgue p-spaces

The spaces mainly used in this thesis are Lebesgue p-spaces of $\mathbb{R}^n$ valued 'functions'[3] (signals) on $\mathbb{R}+$, denoted by $L^p(\mathbb{R}+)$; with the norm given by equation 2.1.

$$\|u(\cdot)\|_{L^p(\mathbb{R}+)} := \begin{cases} \left(\int_0^\infty |u(t)|^p \, \mathrm{d}t \right)^{\frac{1}{p}} & 0 < p < \infty \\ \sup_{0 \le t \le \infty} |u(t)| & p = \infty \end{cases} \tag{2.1}$$

The set of signals in the space $L^p(\mathbb{R}+)$ are all those for which the above norm is finite.

It is useful to define another norm, $L^p[0, T]$ for measuring signals $u \colon [0, T] \to \mathbb{R}^n$ over the finite timespan $[0, T]$:

$$\|u(\cdot)\|_{L^p[0,T]} := \begin{cases} \left(\int_0^T |u(t)|^p \, \mathrm{d}t \right)^{\frac{1}{p}} & 0 < p < \infty \\ \sup_{0 \le t \le T} |u(t)| & p = \infty \end{cases} . \tag{2.2}$$

Analogous to the continuous $L^p(\mathbb{R}+)$ space is the discrete $l^p(\mathbb{N})$ space, with norm

$$\|u(\cdot)\|_{l^p(\mathbb{N})} := \begin{cases} \left(\sum_{k=0}^\infty |u(k)|^p \right)^{\frac{1}{p}} & 0 < p < \infty \\ \sup_{k \in \mathbb{N}} |u(k)| & p = \infty \end{cases} . \tag{2.3}$$

(In this thesis the set of natural numbers $\mathbb{N}$ will include 0.)

In Chapter 5, when dealing with iterative learning control, we will be using 2D signal spaces that take account of the time and iteration nature of the systems we will be examining.

---

[3]Formally the elements of $L^p(\mathbb{R}+)$ are equivalence classes of functions, where f and g lie in the same equivalence class if they differ only on a set of measure zero (see Rudin, 1987).

2.1.3   Extended spaces

It will often be necessary to examine unbounded signals. Since these will not be included in the standard signal space an extension is required (Zames, 1966). The extended Lebesgue p-space, $L_e^p(\mathbb{R}+)$ is defined by

$$L_e^p(\mathbb{R}+) := \{\, f \colon \mathbb{R}+ \to \mathbb{R}^n \,\mid\, T_\tau f \in L^p(\mathbb{R}+), \quad \forall \tau > 0 \,\} \tag{2.4}$$

where T_τ is the truncation operator,

$$T_\tau f(t) := \begin{cases} f(t) & t \le \tau \\ 0 & t > \tau \end{cases}. \tag{2.5}$$

This extended space allows us to consider unstable mappings $L^p \to L_e^p$ from bounded signals onto signals which are potentially unbounded after infinite time (the $\mathbb{R}+$ is often dropped from the notation).

A few points should be noted about extended signal spaces:

i) The extended signal space contains both bounded and unbounded signals, i.e. $L^p \subset L_e^p$.

ii) A signal which lies in the extended signal space but not in the standard signal space will have an infinite norm.

iii) We can expand the definition of extended signal spaces to cover other spaces such as l^p, however signal spaces defined on a finite timespan (such as $L^p[0,T]$ given above) cannot have an extension.

iv) A signal's boundedness is entirely dependant on the norm used to measure it and therefore the signal space in which it is defined. As an example, consider a continuous signal of constant magnitude over t: this signal is bounded under L^∞ and so can lie within the L^∞ space, however its norm under L^2 is infinite unless truncated at any point on t and so lies within L_e^2 but not L^2.

v) Signals that have *finite escape times* are those that become pointwise unbounded before infinite time and therefore do not lie in an extended signal space with L^∞-norms. An example of such a signal is the tan function under L^∞ as it is unbounded at $(n - 0.5)\pi$ for all $n \in \mathbb{Z}$.

Since this thesis makes use of the truncation definition throughout, we now define signal spaces as normed vector spaces provided with a definition of truncation.

Definition 2.1. A **signal space** is a normed vector space with a defined truncation, which is also 'truncation complete', that is for any $x \in \mathcal{W}$ and any $\tau \in \mathrm{dom}(x)$ then $T_\tau x \in \mathcal{W}$.

2.1.4   Induced norms and gains

Let $\mathcal{X}$ and $\mathcal{Y}$ be normed vector spaces.

The induced norm of an operator $G\colon \mathcal{X}_e \to \mathcal{Y}_e$ is defined by equation 2.6. This can be interpreted as the highest gain of the operator.

$$\|G\| := \sup_{\substack{x\in\mathcal{X}_e,\ \tau>0 \\ T_\tau x\neq 0}} \frac{\|T_\tau Gx\|_\mathcal{Y}}{\|T_\tau x\|_\mathcal{X}} \tag{2.6}$$

For operators with superlinear growth this norm will be infinite and it may therefore be more useful to use gain functions, where the gain is measured with respect to the size of the input. $g[G]\colon \mathbb{R}+ \to \mathbb{R}+$ with

$$g[G](\alpha) := \sup\{\|Gx\| : \|x\| \le \alpha\}. \tag{2.7}$$

2.1.5   Notation in the frequency domain

Although the bulk of the material in this thesis is based on time domain concepts, occasionally it will be more appropriate to consider certain themes in the frequency domain. This section will therefore introduce some of the equivalences between time and frequency domain spaces and measures. This material is based on definitions given in Vinnicombe (2001).

First it is required to define a new space of functions $\mathcal{H}^2$ in the frequency domain. This space consists of the Laplace transforms of all functions in $L^2(\mathbb{R}+)$. Therefore a signal $u \in L^2(\mathbb{R}+)$ if and only if its Laplace transform $\hat{u} = \mathscr{L}(u) \in \mathcal{H}^2$. For u and $\hat{u}$ defined in this manner Parseval's theorem states that $\|u\|_{L^2(\mathbb{R}+)} = \|\hat{u}\|_{\mathcal{H}^2}$.

Operators on $\mathcal{H}^2$ that map any signal in $\mathcal{H}^2$ to one in $\mathcal{H}^2$ can be considered stable since any bounded energy input produces a bounded energy output. The space consisting of all of these operators is a Hardy space denoted $\mathcal{H}^\infty$, with the norm of any element $P \in \mathcal{H}^\infty$ given by:

$$\|P\|_\infty = \sup_{u\in\mathcal{H}^2,\ u\neq 0} \frac{\|Pu\|_{\mathcal{H}^2}}{\|u\|_{\mathcal{H}^2}}. \tag{2.8}$$

Note that this $\mathcal{H}^\infty$-norm is the induced $\mathcal{H}^2$-norm[4], and since the norm of functions in $\mathcal{H}^2$ is the same as the L^2-norm of their inverse Laplace transform we can also state that:

$$\|P\|_\infty = \|\mathscr{L}^{-1}P\mathscr{L}\|_{L^2(\mathbb{R}+)} = \sup_{u\in L^2(\mathbb{R}+),\ u\neq 0} \frac{\|\mathscr{L}^{-1}P\mathscr{L}u\|_{L^2(\mathbb{R}+)}}{\|u\|_{L^2(\mathbb{R}+)}}. \tag{2.9}$$

---

[4]The norm of operators in $\mathcal{H}^\infty$ is denoted $\|\cdot\|_\infty$ to distinguish it from the induced norm notation used primarily in this thesis.

Yet another way to define the $\mathcal{H}^\infty$-norm uses the maximum singular value, denoted $\bar{\sigma}$:

$$\|P\|_\infty = \sup_{s:\,\Re(s)>0} \bar{\sigma}(P(s)). \tag{2.10}$$

The set of operators that are finite in this norm and are also rational are denoted $\mathcal{RH}^\infty$. For any $P \in \mathcal{RH}^\infty$ the supremum of equation 2.10 occurs on the boundary and so the equation can be simplified to:

$$\|P\|_\infty = \max_{\omega \in \mathbb{R} \cup \infty} \bar{\sigma}(P(j\omega)). \tag{2.11}$$

2.1.6 Feedback system configuration

Before demonstrating some of the methods of modelling uncertainty it is necessary to explain some preliminary notation regarding feedback systems. The plants considered here will be within feedback interconnections of the form shown in Figure 2.4.
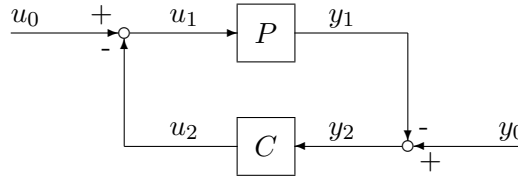
FIGURE 2.4: Feedback configuration $[P, C]$

The closed-loop system will be denoted by $[P, C]$, and defined as the set of equations:

$$
\begin{aligned}
P : \mathcal{U}_e &\to \mathcal{Y}_e & u_0 &= u_1 + u_2 \\
C : \mathcal{Y}_e &\to \mathcal{U}_e & y_0 &= y_1 + y_2 \\
& & y_1 &= Pu_1 \\
& & u_2 &= Cy_2. \tag{2.12}
\end{aligned}
$$

Here $\mathcal{U}$ and $\mathcal{Y}$ are appropriate signal spaces (such as $L^p(\mathbb{R}+)$), $u_0 \in \mathcal{U}$, $y_0 \in \mathcal{Y}$ and $\mathcal{W} = \mathcal{U} \times \mathcal{Y}$. The internal signals lie in extensions of $\mathcal{U}$ and $\mathcal{Y}$ as per equation 2.4 (i.e. $u_1, u_2 \in \mathcal{U}_e$; $y_1, y_2 \in \mathcal{Y}_e$) and $\mathcal{W}_e = \mathcal{U}_e \times \mathcal{Y}_e$. At first, attention will also be restricted to the case where $P0 = 0$ and $C0 = 0$.

Definition 2.2. The system governed by the set of equations $[P, C]$ is said to be **globally well-posed** if for all $(u_0, y_0) \in \mathcal{W}$ there exists unique $((u_1, y_1), (u_2, y_2)) \in \mathcal{W}_e \times \mathcal{W}_e$.

For a globally well-posed closed-loop system $[P, C]$, $H_{P,C}$ is defined as the mapping from

the external to internal signals,

$$H_{P,C} \colon \mathcal{W} \to \mathcal{W}_e \times \mathcal{W}_e : \begin{pmatrix} u_0 \\ y_0 \end{pmatrix} \mapsto \left(\begin{pmatrix} u_1 \\ y_1 \end{pmatrix}, \begin{pmatrix} u_2 \\ y_2 \end{pmatrix} \right). \tag{2.13}$$

Definition 2.3. For a globally well-posed system given by the set of equations $[P, C]$, **stability** is attained when, for all $(u_0, y_0) \in \mathcal{W}$, $H_{P,C}(u_0, y_0) \in \mathcal{W} \times \mathcal{W}$ (i.e. the internal signals are bounded, $((u_1, y_1), (u_2, y_2)) \in \mathcal{W} \times \mathcal{W}$, and so no signal lies in $\mathcal{W}_e \setminus \mathcal{W}$). This can be termed bounded-input bounded-output (BIBO) stability.

Definition 2.4. A system is **gain stable** when the induced norm of $H_{P,C}$ is finite:

$$\|H_{P,C}\| < \infty. \tag{2.14}$$

BIBO stability and gain stability are equivalent for finite-dimensional linear time-invariant systems. For non-linear systems this equivalence may not hold, e.g. zero input disturbances could produce bounded, non-zero internal signals, giving an infinite induced norm but still possessing bounded input bounded output (BIBO) stability. We will return to this point in Section 4.3.

2.1.7 Small gain theorem

The basis for a range of robust stability results is the small gain theorem (Zhou et al., 1996). Introduced in Zames (1966), it provides an open-loop condition for the closed-loop stability of a feedback loop. The paper examines the plant and controller as abstract mathematical mappings from input to output and gives conditions on these maps that guarantee stability. The results are not restricted to stable systems and so extended signal spaces and truncations are used to examine signals that become unbounded after infinite time.

Theorem 2.5. (From Theorem 1 of Zames, 1966.) *For the globally well-posed feedback system given in Figure 2.4, the closed-loop system is gain stable if $\|P\| \cdot \|C\| < 1$.*

For proofs of these results see Zames (1966) or Zhou et al. (1996).

The small gain theorem forms the fundamental basis of latter robustness results (Vinnicombe, 2001; Zhou et al., 1996).

2.1.8 Robustness, stability and performance

Given a nominal plant model and a description of the model's uncertainty, we can develop an uncertainty model set that contains the nominal model and all other possible models

given that uncertainty. By that definition, this uncertainty model set will include the real plant, although its exact description would not be known. If we can then satisfy certain conditions on all elements of this set, we can guarantee these conditions for the real plant. Notions of stability and performance can therefore give rise to four key properties for a given feedback system (Zhou et al., 1996):

For an uncertainty model set $\mathcal{P}_\Delta$ and a set of performance objectives, where $P \in \mathcal{P}_\Delta$ is the nominal plant model and K is a controller,

1) Nominal stability:

   If K stabilises the nominal model P then the system $[P, K]$ is nominally stable.

2) Robust stability:

   Robust stability is attained when the controller K stabilises all plants belonging to the set $\mathcal{P}_\Delta$.

3) Nominal performance:

   Nominal performance is achieved when the objectives are satisfied for the system $[P, K]$.

4) Robust performance:

   Robust performance is achieved when the objectives are satisfied for the system $[P_i, K]$ for every $P_i \in \mathcal{P}_\Delta$.

A plant is therefore modelled and a controller built to fit this model. This modelled plant and controller is nominally stable and so it is desired that the controller simultaneously stabilises a wider set of plants around the model. With a suitable choice of uncertainty description the original plant should lie somewhere within this set and so if robust stability is achieved then the controller will successfully stabilise the original plant.

In practice stabilisation is not the only objective in controller design. Usually some performance objective is desired, such as the accurate tracking of a trajectory. In this case the controller should fulfil this objective for all plants within the set to ensure that the performance condition is guaranteed for the controller operating on the original plant.

This can also take place in the opposite direction: a controller may be constructed for a plant model in order to maximise the set of plants around it that can also be stabilised. It would then be established whether the original plant lies within this set.

A notion of conservativeness arises when considering robustness. Given a real plant $P_1 \in \mathcal{P}_\Delta$ and its model $P \in \mathcal{P}_\Delta$, a robust controller will be designed to stabilise all plants within $\mathcal{P}_\Delta$. It is conceivable that P and P_1 are in fact very close and yet $\mathcal{P}_\Delta$ is a much larger set. In this case designing a controller for all of $\mathcal{P}_\Delta$ may give good robustness results but also lead to poor performance for P_1. It may be wiser to shrink $\mathcal{P}_\Delta$ to a

more appropriate set and potentially achieve better performance results. In order to develop the best controller for the task in mind it is therefore good practice to have an idea of the level and form of uncertainty existing between real and modelled plants. It is natural to assume that, given an accurate description of the difference between the two, the set of plants can be narrowed and conservativeness minimised. The method of modelling the uncertainty is therefore of prime importance.

2.2   Modelling Uncertainty

In this thesis a real physical plant will be considered to consist of its mathematical model with some form of perturbation. This section will introduce some of the methods used to model these plant perturbations in the frequency domain. The different models are useful for different types of model uncertainty and whether it is the sensitivity function or complementary sensitivity function that is being analysed (Doyle et al., 1990; Vinnicombe, 2001).

As in the previous section the uncertainty model set shall be denoted $\mathcal{P}_\Delta$. Fundamentally there are two different types of uncertainty: structured and unstructured. Structured uncertainties have an explicit form defining the plant model set. An example of a structured uncertainty set (from Doyle et al., 1990) is the set of plant models

$$\mathcal{P}_\Delta = \left\{ P_1 \in \mathcal{R} \;\middle|\; P_1 = \frac{1}{s^2 + as + 1}, \; a_{min} \leq a \leq a_{max} \right\}. \tag{2.15}$$

This set has a highly restrictive form. In reality, the physical plant in question would likely be missing from $\mathcal{P}_\Delta$ — high frequency unmodelled dynamics and any extra poles and zeros being absent from the set. It is much more desirable for the uncertainty to be at least partially unstructured. The following are useful tools as they provide disk-like uncertainty around the nominal plant model. This means stability guarantees can be expressed for any plant within a highly general perturbation framework around the nominal plant.

2.2.1   Additive and multiplicative

Unstructured uncertainty representations are more useful as no uncertainty structure is implied, the set of plants is simply confined to a neighbourhood of the nominal model (Zhou et al., 1996). This allows the inclusion of perturbations such as parameter changes or neglected high frequency dynamics.

For a nominal plant $P(s) \in \mathcal{R}$ additive uncertainty is described by the set

$$\mathcal{P}_\Delta = \left\{ P_1 \in \mathcal{R} \;\mid\; P_1 = P + \Delta, \; \|W_1 \Delta W_2\|_\infty < 1, \; \Delta, W_1, W_2 \in \mathcal{H}^\infty \right\}. \tag{2.16}$$

At frequencies where the plant is well known W_1 and W_2 are large, forcing Δ small (Vinnicombe, 2001). Figure 2.5 shows a block diagram of this additive uncertainty representation.
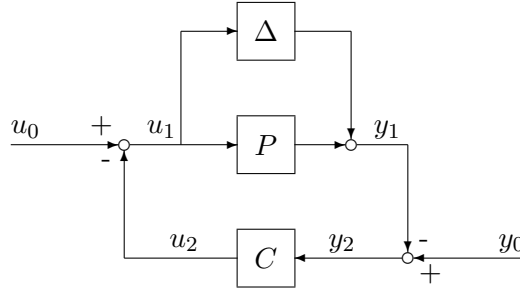
FIGURE 2.5: Additive uncertainty block diagram representation

The representation is known as 'disk-like uncertainty', which comes from viewing the set on a Nyquist diagram for a single-input-single-output (SISO) system. The Nyquist curve of any plant within $\mathcal{P}_\Delta$ lies within the set consisting of the union of disks with centre $P(j\omega)$ and radius $\frac{1}{w(j\omega)}$ over all frequencies, where $w = W_1 W_2$.

The small gain theorem can be used to show that the plant C stabilises all $P_1 \in \mathcal{P}_\Delta$ provided $[P, C]$ is stable and $\|W_2^{-1} C(I - PC)^{-1} W_1^{-1})\|_\infty < 1$. One significant weakness of this type of uncertainty is that it precludes a change in the number of right half plane poles between P and P_1.

Similar to additive perturbations, multiplicative perturbations take the structure of $P_1 = (I + \Delta)P$. This comes from the definition of the perturbation as a stable transfer function of the form

$$\Delta = \frac{P_1}{P} - I \tag{2.17}$$

with the system diagram as Figure 2.6.
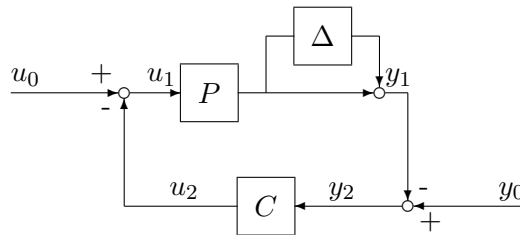
FIGURE 2.6: Multiplicative uncertainty block diagram representation

This gives us very similar results to the additive case. For the set of plants

$$\mathcal{P}_\Delta = \{ P_1 \in \mathcal{R} \mid P_1 = (I + \Delta)P, \ \|W_1 \Delta W_2\|_\infty < 1, \ \Delta, W_1, W_2 \in \mathcal{H}^\infty \} \qquad (2.18)$$

for any $P_1 \in \mathcal{P}_\Delta$, the closed-loop system $[P_1, C]$ is stable provided $[P, C]$ is stable and $\|W_2^{-1} PC(I - PC)^{-1} W_1^{-1})\|_\infty < 1$.

These models are well suited to dealing with higher frequency errors (Vinnicombe, 2001; Zhou et al., 1996), which are generally also harder to model with structured uncertainty methods. Example of the source of such errors from Zhou et al. (1996) are infinite-dimensional electro-mechanical resonance, time-delays and diffusion processes.

Both of these models usually provide fairly coarse approximations leading to conservativeness (Doyle et al., 1990), although they also provide very strong results and are therefore useful tools for studying robustness — in ILC there are a number of recent papers that base robust stability results on multiplicative and additive perturbations, see Section 3.7 for some examples.

Analysis is not always restricted to unknown plant perturbations, it is often wished to model complex non-linear elements in a simple linear fashion and use robustness results such as conic sector non-linearity (see Zames, 1966) to prove stability.

2.2.2 Inverse multiplicative

Inverse multiplicative uncertainty is denoted by $P_1 = (I - \Delta)^{-1}P$, $\Delta \in \mathcal{RH}^\infty$ (Figure 2.7). Under the conditions that $[P, C]$ is stable and $\|W_2^{-1}(I - PC)^{-1} W_1^{-1})\|_\infty < 1$, the set of plants guaranteed stable in feedback with the controller C is therefore given by

$$\left\{ P_1 \in \mathcal{R} \mid P_1 = (I - \Delta)^{-1}P, \ \|W_1 \Delta W_2\|_\infty < 1, \ \Delta, W_1, W_2 \in \mathcal{H}^\infty \right\}. \qquad (2.19)$$

Unlike the previous two models, here the number and location of right half plane poles can change between plants; although they must share the same right half plane zeros (Vinnicombe, 2001). Inverse multiplicative uncertainty is more suited to the analysis of lower frequency perturbations such as parametric uncertainty.

2.2.3 Coprime factor

The key features of all of the above uncertainty models can be described by coprime factor uncertainty (Georgiou and Smith, 1990; Vinnicombe, 2001; Xie, 2004; Doyle et al., 1990). For this, the plant is first written down as a ratio of elements $N, M \in \mathcal{RH}^\infty$. These elements define a *right coprime factorisation* of a plant P when $P = NM^{-1}$ and there exists $X, Y \in \mathcal{H}^\infty$ such that $XM + YN = I$. This coprime factorisation is also
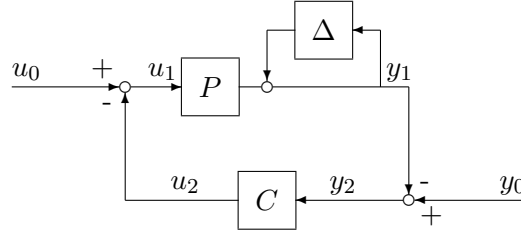
FIGURE 2.7: Inverse multiplicative uncertainty block diagram representation

considered *normalised* if $M^*M + N^*N = I$, where M^* denotes the conjugate transpose of M.

In this case we will drop the weights W_1 and W_2. Let $P = NM^{-1}$ be a normalised right coprime factorisation of P. A set of perturbed plants then takes the form

$$\mathcal{P}_\Delta = \left\{ P_1 \in \mathcal{R} \ \middle| \ P_1 = (N + \Delta_N)(M + \Delta_M)^{-1}, \ \left\| \begin{bmatrix} \Delta_N \\ \Delta_M \end{bmatrix} \right\|_\infty < \frac{1}{\gamma} \right\}, \tag{2.20}$$

shown by Figure 2.8 in block diagram form.
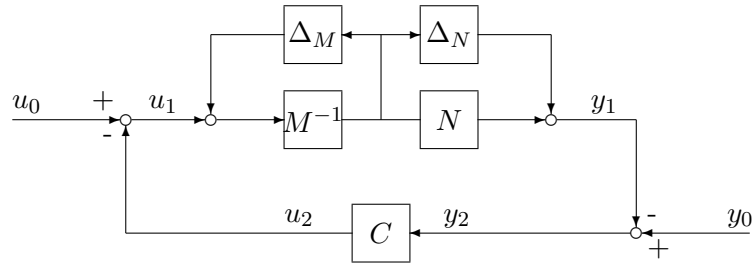
FIGURE 2.8: Coprime factor uncertainty block diagram representation

Similar to previous results, $[P_1, C]$ is stable for all $P_1 \in \mathcal{P}_\Delta$ provided that $[P, C]$ is stable and $\|M^{-1}(I - CP)^{-1}[C \ -I]\|_\infty \leq \gamma$ (Vinnicombe, 2001).

2.3   Summary

This chapter has defined the notation necessary for this thesis and introduced the concept of uncertainty modelling. Additive, multiplicative, inverse multiplicative and coprime factor uncertainty representations have also been described.

Chapter 4 will expand the uncertainty modelling work to the gap metric. It will be shown that the coprime factor uncertainty detailed above is a special case of the linear

gap (which is in turn a special case of the non-linear gap). For systems in $\mathcal{R}$ the directed gap between the two plants in equation 2.20, can be written:

$$\vec{\delta}_0(P, P_1) = \inf \left\{ r \in \mathbb{R} \ \middle| \ r = \left\| \begin{array}{c} \Delta_N \\ \Delta_M \end{array} \right\|, \ P_1 = \frac{N + \Delta_N}{M + \Delta_M}, \right.$$
$$\left. P = NM^{-1}, \ \begin{pmatrix} \Delta_N \\ \Delta_M \end{pmatrix} \in \mathcal{H}_\infty, \ M^*M + N^*N = I \right\}. \quad (2.21)$$

The gap between two plants must be invariant to the order in which they are expressed and so the non-directed gap is given by the maximum of both directed gaps:

$$\delta_0(P, P_1) = \max \left\{ \vec{\delta}_0(P, P_1), \vec{\delta}_0(P_1, P) \right\}. \quad (2.22)$$

Before this, the following chapter will introduce the subject of iterative learning control.

Chapter 3

Iterative Learning Control

In classical control theory, feedback is used to improve the performance of a plant, and the controller's parameters are tuned to improve the way the system reacts to its environment. The most widely used type of control in industry is PID control, where proportional, integral and differential terms are tuned to maximise system performance. Iterative learning control (ILC) has a different feedback configuration.

Given a tracking task of finite duration in time, where it is desired to repeat this task over many iterations, ILC can be implemented in order to improve the tracking performance. A learning algorithm takes data from past trials and uses it to improve the tracking on the current trial. Within traditional ILC systems the plant is run open-loop for each iteration and the plant's input signal is calculated from the previous iteration and the error resulting at the output from the previous iteration. The aim of the whole process is for the output to converge to the set task over progressive trials.

Figure 3.1 shows a typical ILC set-up. It is required that the output signal y_k should repeatedly track a reference signal y_{ref} which is T seconds long, where $k \geq 0$ is the trial number. The plant may not be known exactly and so simply applying the input signal $u_k^* = P^{-1} y_{ref}$ may not be possible. Instead an iterative learning (IL) controller is implemented that attempts to find the signal u_k^* (or a suitable approximation thereof) by applying an input signal trial-after-trial and updating it each time based on the error at the output. At first an input signal u_1 is chosen (this would likely be an initial guess at u_k^*). This signal is applied to the plant and the plant operates open-loop. At time T the system is reset and an ILC algorithm calculates a new input signal from a combination of the input that was given to the plant, the reference signal and the output of the plant. The method of generating this new signal depends on the algorithm used. Once this new signal has been generated the whole process is started again, the new input is applied and the results used to calculate the next. This can all happen without any explicit knowledge of the plant characteristics.

A deterministic plant and controller will always produce the same error under classical
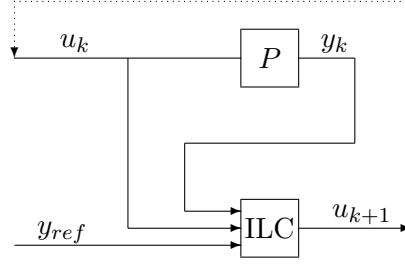
FIGURE 3.1: ILC system block diagram

control methods (neglecting disturbances). However under iterative learning control, by the process of determining the following control input, errors are fed forward in order to develop input signals that provide closer tracking. It is therefore possible for an ILC implementation to achieve more accurate control than is possible with classical techniques, often with less knowledge of the plant.

3.1 History

The first theory of iterative learning control is regarded to have been made in Uchyama (1978) although its publication in a Japanese journal meant that the subject did not really gain acclaim until 1984, when three papers appeared on what is now known as iterative learning control (Arimoto et al., 1984; Casalino and Bartolini, 1984; Craig, 1984). The aim of the papers was the same — improved control of robotic manipulators. This section will give a brief explanation of these three papers, explain their aims, proposed solutions, and results.

3.1.1 Three early ILC papers

Often cited as the birth of ILC, Arimoto et al. (1984) concerned the possibility of mimicking in robots the human ability of learning movement through repeating the same action. A robot's operation is improved using data from the previous operation.

What is now known as the D-type algorithm was developed, the update rule shown below. (Here the plant P maps u_k to y_k, and it is desired for the output y_k to follow the trajectory y_{ref}.)

$$u_{k+1}(t) = u_k(t) + \gamma \dot{e}_k(t)$$
$$e_k(t) = y_{ref}(t) - y_k(t) \tag{3.1}$$

The plant is given an input signal u_k and its output, y_k, is recorded. This is subtracted from the desired trajectory to form an error trajectory which is then differentiated, weighted by γ, and added to u_k to form the input for the next trajectory. If designed

properly, over successive iterations the motion should asymptotically approach the desired trajectory.

The paper proves the uniform convergence of this algorithm in terms of a norm on the error signal for a class of linear systems. In particular, the analysis concentrates on robot manipulators consisting of a linear servomechanism with a single degree of freedom. The 'Arimoto type' algorithms are used extensively in this thesis and elsewhere in the field of ILC and are explained in much more detail in Section 3.2.

Casalino and Bartolini (1984) begins by explaining three methodologies for the control of highly non-linear servomechanisms:

- Assume the system is linear

  This can work well over a small operating area; however, as the demands of the system increase, or it is required to operate in an area where the linear model is inadequate, the approach can fail.

- Use a non-linear model

  If the non-linearity inherent within the plant is conserved in the model very good performance can be attained, although this is counter-balanced by a more costly control system and a need for a more accurate knowledge of plant and load.

- Use adaptive techniques

  This grants the use of a simple plant model and allows an intelligent controller to compensate for the non-linear plant dynamics. Although not addressed in Casalino and Bartolini (1984), there are often some significant drawbacks to intelligent control designs. Most notably there is a poor level of understanding as to the operation of some adaptive control methodologies. ILC naturally falls into this category and, as explained later, can suffer stability issues. This instability is not very well understood; in many cases it only becomes apparent after long periods of apparent convergence.

Casalino and Bartolini (1984) takes greater account of the plant model than Arimoto et al. (1984) in calculating the input for the next iteration. Also — rather than just using the velocity error — the position, velocity and acceleration errors are all used within the update algorithm.

The aim of Casalino and Bartolini (1984) is the control of robot manipulators of the form

$$A(\theta)\ddot{\theta} + B(\theta, \dot{\theta}) + C(\theta) = M(\theta, \dot{\theta}, \ddot{\theta}), \tag{3.2}$$

$\theta \in \mathbb{R}^n$ being a vector of angular joint positions, $M \in \mathbb{R}^n$ a vector of actuator torques, A as the inertia matrix, B representing Coriolis and centrifugal forces and C gravitational forces.

The feedback law used is of the structure

$$M = m + Q - K\ddot{\theta} - L\dot{\theta} - P\theta, \tag{3.3}$$

where Q is a constant vector of external signals required to set the initial position for all the trials; K, L and P are acceleration, velocity and position constants; and m is the driving signal at the heart of the ILC algorithm. The signal m along trial j is given by

$$m_j = m_{j-1} - K\delta\ddot{\theta}_{j-1} - P\delta\dot{\theta}_{j-1} - L\delta\theta_{j-1}, \tag{3.4}$$

where $\delta\ddot{\theta}_{j-1}$, $\delta\dot{\theta}_{j-1}$ and $\delta\theta_{j-1}$ are acceleration, velocity and position errors along trial $j-1$ ($\delta\theta_{j-1} = \theta^* - \theta_{j-1}$ where θ^* is the desired trajectory).

Bounded input bounded output stability is proven possible for this class of plants, advocating a trial and error approach to finding the gains K, P and L that provide a sufficient reduction in error.

A more in-depth look into this algorithm will not be given here, suffice to say that equation 3.4 is very similar in structure to equation 3.1 in the algorithm of Arimoto et al. but with added position and acceleration error terms. These added terms mimic the plant in order to make best use of available knowledge of it, allowing the learning to compensate for the perturbations between the plant and plant model. Casalino and Bartolini (1984) also makes the crucial point that a major downside to using this type of control is that, once initiated, as the algorithm tunes the input signal in order to produce the reference at the output, the whole scheme must be restarted when attempting to follow a different reference trajectory.

Published in the same year, Craig (1984) describes a similar adaptive scheme. It is proposed that the best design of controller will make use of all possible information about the plant. Craig (1984) therefore splits the plant into two; a rigid body model, and the remainder of the plant remaining unmodelled. Effects such as friction, being difficult to model, are then present in the unmodelled dynamics, which an adaptive technique attempts to learn.

The methodology given is based on the manipulator model similar to that of Casalino and Bartolini (1984):

$$\tau(\theta, \dot{\theta}, \ddot{\theta}) = M(\theta)\ddot{\theta} + V(\theta, \dot{\theta}) + F(\theta, \dot{\theta}), \tag{3.5}$$

where $\tau(\theta, \dot{\theta}, \ddot{\theta})$ is the $N \times 1$ vector of torques at the joints, $M(\theta)$ an $N \times N$ mass matrix of the manipulator, $V(\theta, \dot{\theta})$ an $N \times 1$ vector of velocity and gravitational effects, $F(\theta, \dot{\theta})$ an $N \times 1$ vector of frictional effects, and N the number of joints.

The control law is then given by

$$\tau = \hat{M}(\theta) \left[\ddot{\theta}_d + K_v \dot{\epsilon} + K_p \epsilon \right] + \hat{V}(\theta, \dot{\theta}) + \hat{F}_k, \tag{3.6}$$

where K_v and K_p are diagonal gain matrices, $\hat{M}$ and $\hat{V}$ are approximations of M and V, $\hat{F}_k$ is the feedforward torque after trial k, θ_d is the desired trajectory, and ϵ are the servo errors $\epsilon = \theta_d - \theta$.

The torque is updated from the torque of the previous trial and the error between measured and desired signals convolved with a filter P.

$$\hat{F}_{k+1} = \hat{F}_k + \hat{M}P * (\theta_d - \theta_k) \tag{3.7}$$

The paper shows the ability of this algorithm to follow arbitrary paths (within the limits of the actuators). It also suggests the use of the method for verification of a plant model — the accuracy of a plant model would be reflected in the size of the feedforward torque required to correct it.

Unlike the work of Arimoto et al. (1984), both Casalino and Bartolini (1984) and Craig (1984) embed the iterative update into a standard feedback structure. Arimoto et al. (1984) instead implements a feedforward technique that effectively runs the plant open-loop and updates the input after each trial.

The main difference between Casalino and Bartolini (1984) and Craig (1984) is that the latter has no acceleration error term in the feedback equation 3.6 whereas the former does (equation 3.3). This is due to its assumption that M is modelled sufficiently by the rigid body model which is the only term containing acceleration. Craig (1984) also only iteratively updates the position error; Casalino and Bartolini (1984) feeds forward errors for position, velocity and acceleration.

3.1.2 ILC since conception

Since 1984 the field of ILC has expanded significantly. Much research has been done on new algorithms, proving their stability, analysing and improving convergence rates, reducing the overall error, applying ILC to different classes of plants and also on practical implementation. The following section will describe some of the core ILC algorithms in discrete and continuous time, and explain how these algorithms are developed into more complex learning structures, often utilising work done within other branches of control theory. Some studies on robustness have also been undertaken and are discussed later in the chapter.

3.2   D-type and P-type ILC in Continuous Time

Define a plant by the convolution equation

$$y(t) = g(t) + \int_0^t H(t,\tau)u(\tau)\,\mathrm{d}\tau, \tag{3.8}$$

where $u(t), y(t), g(t) \in \mathbb{R}^r$ and $H(t,\tau) \in \mathbb{R}^{r \times r}$.

As explained in Section 3.1.1, the D-type algorithm from Arimoto et al. (1984) is given by the update rule

$$u_{k+1}(t) = u_k(t) + \Gamma(t)\dot{e}_k(t)$$
$$e_k(t) = y_{ref}(t) - y_k(t), \tag{3.9}$$

where the desired trajectory is given by y_{ref}, and $\Gamma(t) \in \mathbb{R}^{r \times r}$ is a function that determines the learning rate.

A norm that operates on r-vector-valued functions on $[0,T]$ is defined by

$$\|e(\cdot)\|_\lambda = \sup_{0 \le t \le T} \{e^{-\lambda t} \max_{1 \le i \le r} |e_i(t)|\} \qquad \lambda > 0. \tag{3.10}$$

The following theorem (from Arimoto et al., 1984) provides a condition for the convergence of the closed-loop system of equations 3.8 and 3.9 in this norm.

Theorem 3.1. (From Theorem 1 of Arimoto et al., 1984.) *For the plant given by equation 3.8 and controller by equation 3.9, if $y_{ref}(0) = y_k(0)$, $\|I - H(t,t)\Gamma(t)\| < 1$, and a given input $u_0(t)$ is continuous on $[0,T]$ then there exists positive constants λ and p_0 such that*

$$\|\dot{e}_{k+1}\|_\lambda \le p_0 \|\dot{e}_k\|_\lambda \qquad and \qquad 0 \le p_0 < 1 \tag{3.11}$$

for $k = 0, 1, 2, \ldots$, where the matrix norm $\| \cdot \|$ of an $r \times r$ matrix G with elements g_{ij} is given by

$$\|G\| = \max_{1 \le i \le r} \left\{ \sum_{j=1}^r |g_{ij}| \right\}. \tag{3.12}$$

Corollary 3.2. *Given that $\|\dot{e}_k\|_\lambda \to 0$ as $k \to \infty$ then $|e_k(t)| \to 0$ for all $t \in [0,T]$ as $k \to \infty$.*

Proof. Using mean value theorem, for all $t \in (0,T]$ and $k \in \mathbb{N}$ there exists $\zeta \in (0,t]$ and

$c \in \mathbb{R}$ such that

$$\left| \frac{e_k(t) - e_k(0)}{t} \right| = |\dot{e}_k(\zeta)|$$

$$\leq \|\dot{e}_k\|_\infty \leq e^{\lambda T} \|\dot{e}_k\|_\lambda$$

$$\leq c\|\dot{e}_k\|_\lambda \tag{3.13}$$

Finally, since $\|\dot{e}_k\|_\lambda \to 0$ and $e_k(0) = 0$, we observe that $|e_k(t)| \to 0$ for all $t \in (0, T]$, as required. $\qquad \square$

The error is therefore not proven to decay monotonically. The constant c above can be very large, allowing a huge error that eventually converges to zero after many iterations. In Amann (1996) it is explained that the iterative process concentrates the learning towards the start of the cycle, gradually pushing the reduction in error further along the trial.

Figure 3.2 shows a simulation of the Arimoto algorithm for the (relative degree 1) plant $(s+1)/(s^2 + s + 1)$. The figure shows the convergence of the algorithm to the reference signal. Figure 3.3 shows the $L^\infty[0, T]$ norm of the error for each iteration, the error monotonically decreasing.

After the development of the D-type update rule came the P-type rule, using the error as opposed to its derivative (Arimoto et al., 1985):

$$u_{k+1}(t) = u_k(t) + \Gamma(t)e_k(t) \tag{3.14}$$

with the error defined in the same way. Arimoto et al. (1985) explains that the class of systems for which the P-type algorithm converges is smaller than that for the D-type.

P-type is inherently easier to implement as no differentiation is required; differentiation amplifies noise as noise often contains high frequency components (Kawamura et al., 1988).

It is not immediately clear of the need to take the derivative of the error in the original algorithm and why the D-type ILC is capable of stabilising a larger class of systems. Sugie and Ono (1991) justifies the choice between P- and D-type algorithms by explaining that the derivative order must match the relative degree of the system. This is due to the convergence condition's dependance on the 'direct transmission term'. Consider a relative degree one continuous time plant given by

$$\dot{x}(t) = Ax(t) + Bu(t)$$

$$y(t) = Cx(t) \tag{3.15}$$

In order for the system to track a reference trajectory y_{ref} we wish to find the input such that $y(t)$ approaches y_{ref}. On implementation of an IL controller the output error
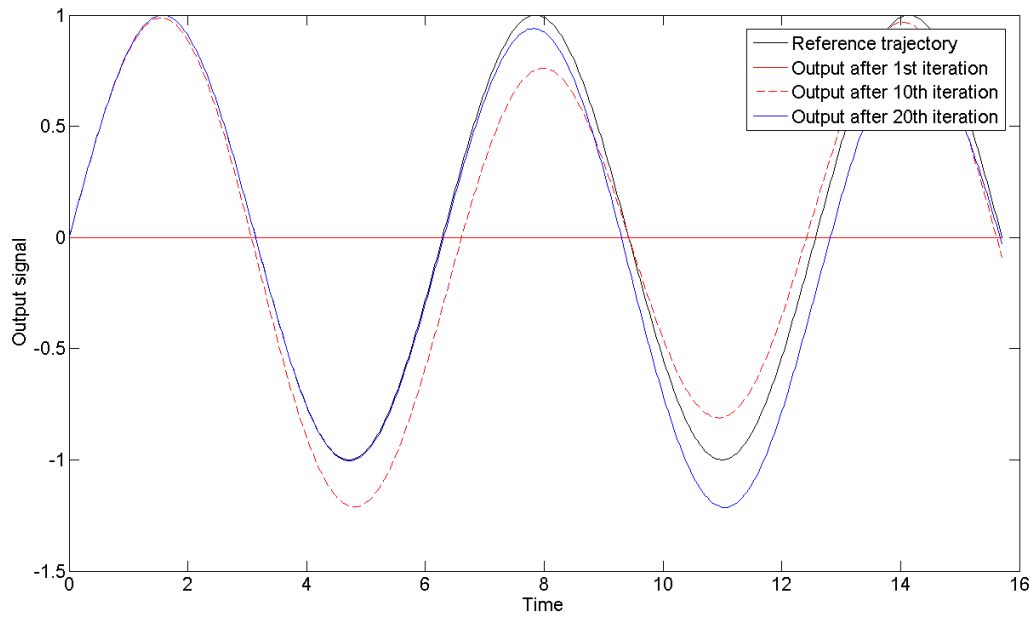
FIGURE 3.2: Simulation of the Arimoto algorithm
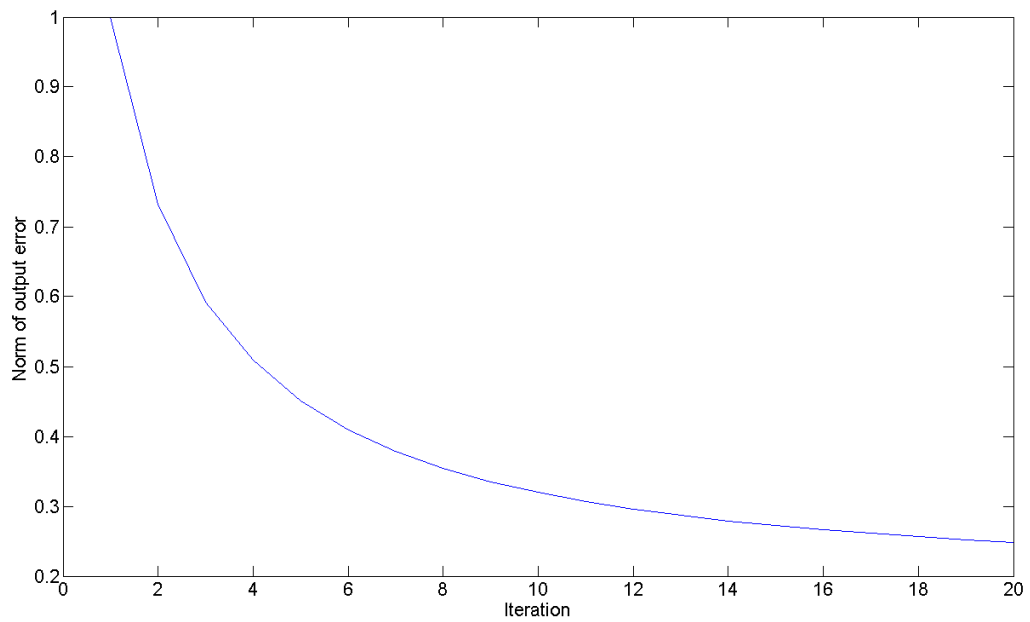
FIGURE 3.3: Simulation of the Arimoto algorithm — $L^{\infty}[0, T]$ norm of the output signal at each iteration

after an iteration is given by $y_{ref} - y_k$ and is fed forward to the next iteration. Since the input signal is integrated once by the plant before reaching the output, this must be compensated for by once differentiating the error before it is fed forward to the input. This holds for higher relative degree systems also — in order to guarantee convergence the derivative of the error must be taken repeatedly in order to propagate the first available Markov parameter through the system.

This demonstrates the logic behind the D-type algorithm's development before the P-type: the control of robotic manipulators was the original aim, these being of relative degree one required one differentiation in order to achieve stability. The reason such systems are generally relative degree one can be illustrated using a simple motor as an example. A voltage will be applied to a motor which affects its velocity, however it is standard to measure displacement which is the integral of velocity and so the resulting system will be relative degree one.

Arimoto et al. (1985) explains convergence criteria for D- and P-type algorithms and also a PD-type algorithm, where u_{k+1} is a weighted sum of P and D terms (this also appears in Kawamura et al., 1985).

3.3   ILC in Discrete Time

ILC is more often implemented in discrete rather than continuous time. This is not simply down to the increase in digitisation that is typical of modern control but comes from ILC's fundamental requirement of a memory component; as it is necessary to store trajectories somehow in order to perform calculations on them. These operations are done by computer and so signals are stored digitally.

3.3.1   Lifted system representation

Several ILC papers use what is known as the lifted system or supervector representation to examine discrete-time plants and controllers within ILC. Within this framework the plant and controller are written as matrices and the problems is transferred from a 2D (time and iteration) problem into a 1D (iteration) problem. (Phan et al., 2000; Hätönen, 2004; Norrlöf and Gunnarsson, 2002; Hätönen et al., 2004; Norrlöf, 2000b; Harte et al., 2005; Hakvoort et al., 2008; Moore, 2001; Owens and Feng, 2003.)

Define a discrete time plant obeying

$$x(t+1) = Ax(t) + Bu(t)$$
$$y(t) = Cx(t) + Du(t), \tag{3.16}$$

with $u(t) \in \mathbb{R}^{r_1}, y(t) \in \mathbb{R}^{r_2}$, $x(t) \in \mathbb{R}^m$ and A, B, C and D of appropriate dimensions. The plant is required to track a reference trajectory $y_{ref}(\cdot)$ with $y_{ref}(t) \in \mathbb{R}^{r_2}$ and $t \in \{0, 1, 2, \ldots, N\}$.

If the initial state vector $x(0)$ is set to be 0 then the plant can be written as

$$
\begin{pmatrix} y(0) \\ y(1) \\ y(2) \\ \vdots \\ y(N) \end{pmatrix} = \begin{pmatrix} D & 0 & 0 & \cdots & 0 \\ CB & D & 0 & \cdots & 0 \\ CAB & CB & D & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ CA^{N-1}B & CA^{N-2}B & CA^{N-3}B & \cdots & D \end{pmatrix} \begin{pmatrix} u(0) \\ u(1) \\ u(2) \\ \vdots \\ u(N) \end{pmatrix}. \quad (3.17)
$$

This matrix is lower-triangular and Toeplitz and will be denoted by P, all the terms within the matrix are the Markov parameters of the system, and terms above the major diagonal are all zero. We are here considering a general plant description where, unlike many ILC papers, the 'D' Markov parameter is retained (for a system with relative degree one or higher D will equal zero). The reason for its retention is mainly to simplify notation and avoid having to explicitly state the relative degree of a system when applying stability and robustness theorems — the relative degree of a system plays an important part in the analysis of this thesis and this is easiest to express using this notation.

The representation can be used in the following way, letting $k \in \mathbb{N}$ denote the trial. The trial dynamics of the plant can now be written (with P as the $(N+1) \times (N+1)$ matrix in equation 3.17),

$$
\begin{pmatrix} y_k(0) \\ y_k(1) \\ y_k(2) \\ \vdots \\ y_k(N) \end{pmatrix} = P \begin{pmatrix} u_k(0) \\ u_k(1) \\ u_k(2) \\ \vdots \\ u_k(N) \end{pmatrix}, \quad (3.18)
$$

with the update rule for trial $k + 1$ written in the same form,

$$
\begin{pmatrix} u_{k+1}(0) \\ u_{k+1}(1) \\ u_{k+1}(2) \\ \vdots \\ u_{k+1}(N) \end{pmatrix} = Q \left(\begin{pmatrix} u_k(0) \\ u_k(1) \\ u_k(2) \\ \vdots \\ u_k(N) \end{pmatrix} + L \begin{pmatrix} e_k(0) \\ e_k(1) \\ e_k(2) \\ \vdots \\ e_k(N) \end{pmatrix} \right), \quad (3.19)
$$

where $e_k(\cdot) = y_{ref}(\cdot) - y_k(\cdot)$ and $y_{ref}(\cdot)$ is the desired trajectory. This principle of this update rule is that L is a matrix of learning gains denoting how the error is fed back to the next input signal, and Q is a filter used to tune the asymptotic properties of the

system. By combining both update rule and plant equations we obtain:

$$u_{k+1}(\cdot) = Q(I - LP)u_k(\cdot) + QLy_{ref}(\cdot). \tag{3.20}$$

This equation can then be used to analyse various ILC strategies.

(For the multiple-input-multiple-output (MIMO) case, where $u(t) \in \mathbb{R}^{r_1}$ and $y(t) \in \mathbb{R}^{r_2}$ in equation 3.16, with $r_1, r_2 > 1$, the matrices L and Q will be of the appropriate sizes.)

3.3.2   D-type and P-type ILC in discrete-time

The discrete-time equivalent of the continuous-time P-type algorithm is simple to construct. Rather than a signal being multiplied by a gain we have (in the lifted system form) a vector being multiplied by a scaled identity matrix. We can therefore define the algorithm using equation 3.19 with Q defined as the $(N+1) \times (N+1)$ identity and L as the same identity multiplied by a scalar, γ. This P-type learning matrix is denoted L_P:

$$L_P = \begin{pmatrix} \gamma & 0 & 0 & \cdots \\ 0 & \gamma & 0 & \cdots \\ 0 & 0 & \gamma & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{3.21}$$

In order to define the D-type ILC algorithm in discrete-time we can choose to approximate the continuous-time differentiation using a forward finite difference method (Saab, 2001). Taking the differential of a signal f as

$$\frac{\mathrm{d}f(t)}{\mathrm{d}t} = \lim_{\Delta \to 0} \frac{f(t + \Delta) - f(t)}{\Delta}, \tag{3.22}$$

by letting Δ instead equal the sample time period T we arrive at

$$\frac{\mathrm{d}f(t)}{\mathrm{d}t} \approx \frac{f(t + T) - f(t)}{T}. \tag{3.23}$$

If we transfer this into the lifted system form, including the gain γ, we have the learning matrix L_D as

$$L_D = \begin{pmatrix} -\frac{\gamma}{T} & \frac{\gamma}{T} & 0 & \cdots \\ 0 & -\frac{\gamma}{T} & \frac{\gamma}{T} & \cdots \\ 0 & 0 & -\frac{\gamma}{T} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \tag{3.24}$$

The above definition of a discrete D-type ILC algorithm is based on approximating continuous-time differentiation. A more intuitive approach would be to construct the al-

gorithm separately based on the discrete-time equivalent of the goals that the continuous-time algorithm was developed for. We shall now look at the simplest continuous-time relative degree 1 plant, an integrator $\frac{1}{s}$, and ignore any robustness issues and noise. Implementing the continuous D-type algorithm with a gain of 1 on such a plant we will converge to the reference trajectory in a single iteration. In discrete-time, the simplest relative degree 1 system consists of a time delay of one time step[1]. The L matrix that allows this plant to converge in a single iteration is an $(N+1) \times (N+1)$ matrix defined as follows (with $T = 1$):

$$
L = \begin{pmatrix} 0 & 1 & 0 & \cdots \\ 0 & 0 & 1 & \cdots \\ 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \tag{3.25}
$$

Setting Q to the identity and substituting L and Q into equation 3.20 we see that $u_{k+1}(\cdot) = P^{-1}y_{ref}(\cdot)$ for all k. It could therefore be argued that this is a more appropriate discrete-time equivalent to the continuous D-type ILC algorithm.

Note that both the continuous D-type algorithm and the discrete D-type algorithm of equation 3.25 have their respective plants' inverses present in their algorithms: in continuous time the plant $\frac{1}{s}$ and controller s; and in discrete time the L matrix multiplied by the plant matrix leaves the identity ($LP = I$).

This version of the discrete D-type algorithm appears in various sources including Owens and Hätönen (2005) and Moore (1993); and in Lin (2006) and Freeman (2004) is referred to as the discrete-time version of the D-type Arimoto algorithm. It is also similar to some other algorithms given in literature associated with anticipatory ILC (Wang, 1999, 2000), and a reference shift algorithm (Cai, 2009).

In Wang (1999) it is pointed out that D-type algorithms are harder to implement (in continuous time) than P-type and also have problems with the noise amplification caused by differentiation. It is also pointed out that the same powerful convergence results that exist for D-type are not shared by P-type algorithms. The proposed solution is an algorithm, based on the P-type learning law, that includes an anticipatory term:

$$
u_{k+1}(t) = u_k(t) + L(\cdot)(e(t + \Delta)). \tag{3.26}
$$

If the algorithm was implemented in discrete-time with Δ equal to one time-step then the algorithm matches the alternative D-type described above. This idea also appears in Barton et al. (2000), where the time delay is taken to be the system's settling time; Freeman et al. (2005c), where the delay is applied to a non-minimum phase system; and Ma et al. (1993), where it is used with current iteration feedback (see Section 3.4.6).

---

[1]Relative degree in discrete-time can be interpreted as the delay (in time-steps) between input and output (Jang et al., 1994).

3.3.3 Convergence in the lifted system format

Before any proofs for convergence can be related it is necessary to point out a property of the lifted system notation used above.

In the examples above we have only examined the top left corner of the matrices. Looking at the right-most column of the plant matrix for a non-relative degree 0 plant reveals a column of zeros (due to the D Markov parameter being zero). It is evident from equation 3.18 that, as we are dealing with a finite trial length, if we have a relative degree n plant then the inputs for the final n time-steps will have no bearing on the plant's output. With this in mind we now re-examine the update equation:

$$u_{k+1}(\cdot) = Q(I - LP)u_k(\cdot) + QLy_{ref}(\cdot). \tag{3.27}$$

We construct a matrix G_n consisting of an $(N + 1 - n) \times (N + 1 - n)$ identity matrix with an extra n columns of zeros on the right:

$$G_n = \underbrace{\begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}}_{(N+1-n)\times(N+1-n)} \underbrace{\begin{matrix} 0 & \cdots & 0 \\ 0 & \cdots & 0 \\ 0 & \cdots & 0 \\ \vdots & \cdots & \vdots \\ 0 & \cdots & 0 \end{matrix}}_{(N+1-n)\times n}. \tag{3.28}$$

Where P denotes the lifted system matrix of a relative degree n plant, with G_n defined above and its transpose as G_n^T, it is sufficient to re-write equation 3.27 in the following form:

$$u_{k+1}(\cdot)|_{t\in[0,T-n]} = G_n Q(I - LP)G_n^T u_k(\cdot)|_{t\in[0,T-n]} + G_n QLG_n^T y_{ref}(\cdot)|_{t\in[0,T-n]}. \tag{3.29}$$

Most of the ILC literature deals with this issue at the outset by a change in notation. Rather than the plant P mapping $[u(0), \cdots, u(N)]$ to $[y(0), \cdots, y(N)]$, for a relative degree 1 plant, it is said to map $[u(0), \cdots, u(N-1)]$ to $[y(1), \cdots, y(N)]$ and so the equivalent lifted system matrix has the CB terms along the major diagonal and all the lifted matrices are one element smaller in both row and column (Phan et al., 2000). Both methods are equivalent since they lead to precisely the same update equation.

Definition 3.3. For an ILC system **asymptotic stability** is defined as the existence of $\overline{u} \in \mathbb{R}$ such that

$$\|u_k(t)\|_{l^\infty} \leq \overline{u} \qquad \forall t \in \{0, 1, 2, \ldots, N\}, k \in \mathbb{N} \tag{3.30}$$

and also the existence of $u_\infty(t) \in \mathbb{R}^r$ for each $t \in [0, N]$ such that

$$u_\infty(t) = \lim_{k \to \infty} u_k(t) \qquad \forall t \in \{0, 1, 2, \ldots, N\}. \tag{3.31}$$

Asymptotic stability on $[0, N - n]$ is defined as the existence of $\overline{u} \in \mathbb{R}$ such that

$$\|u_k(t)\|_{l^\infty} \le \overline{u} \qquad \forall t \in \{0, 1, 2, \ldots, N - n\}, k \in \mathbb{N} \tag{3.32}$$

and also the existence of $u_\infty(t) \in \mathbb{R}^r$ for each $t \in [0, N - n]$ such that

$$u_\infty(t) = \lim_{k \to \infty} u_k(t) \qquad \forall t \in \{0, 1, 2, \ldots, N - n\}. \tag{3.33}$$

Note that this definition relates to stability and not necessarily exact convergence to the reference signal. For that it would also be required that $u_\infty(\cdot) = P^{-1} y_{ref}(\cdot)$. Note also that the second definition only requires convergence of the input for the first $N - n$ time-steps for a relative degree n system. This is due to the redundancy of the last n time-steps in calculating the output of a system with relative degree n.

Various different versions and adaptions of the following theorem can be found in numerous ILC papers and theses including Freeman et al. (2008), Norrlöf and Gunnarsson (2002), Norrlöf and Gunnarsson (2005), van de Wijdeven (2008) and Hätönen (2004). There are also resemblances in earlier functional analysis work from Edwards and Owens (1982).

Theorem 3.4. *An ILC system given by equations 3.16 to 3.20, with the plant P of relative degree n and G_n given by equation 3.28, is asymptotically stable on $[0, N - n]$ if and only if $|\lambda_i(G_n Q(I - LP) G_n^T)| < 1$, where $\{\lambda_0, \ldots, \lambda_{N-n}\}$ is the set of all eigenvalues of $G_n Q(I - LP) G_n^T$.*

Using the eigenvalues in this way simply shows that the system converges and the input signal along the way is bounded; there is nothing to guarantee that the error will monotonically decrease, and the performance of such a system may be completely unacceptable. This theorem is therefore similar to Theorem 3.1 for continuous systems. Both theorems provide necessary and sufficient conditions for convergence but do not state how well behaved this convergence is. Recall the constant c in Corollary 3.2 that permitted a huge increase of the error (in norm) from one trial to the next, and compare it to an example given in Longman (2000) where a discrete-time system is proven to converge but produces mean squared errors of over 10^{51} along the way.

A further convergence condition is one of monotonic convergence in norm. A sufficient condition for this is given by several papers including Sugie and Ono (1991), Norrlöf and Gunnarsson (2005), van de Wijdeven (2008) Hätönen (2004) and Norrlöf (2000b). Again, similar functional analysis-based results can be found in Edwards and Owens (1982).

Theorem 3.5. *Given an ILC system obeying equations 3.16 to 3.20, with the plant P of relative degree n and G_n given by equation 3.28, the error signal $e_k(\cdot) = y_{ref}(\cdot) - y_k(\cdot)$ converges monotonically in norm (i.e. $\|e_{k+1}\| < \|e_k\|$) if $\|G_n(Q(I - LP))G_n^T\| < 1$.*

The proof is based on viewing the system as a contraction mapping in the iteration direction. However, convergence of the error does not necessarily guarantee convergence of the input signal. For that the normed space must also be complete or it may not include the input signal that is required for perfect tracking.

Note that this theorem only provides a sufficient condition, however its use is widespread within the ILC literature, and it will appear later on in this thesis to derive some robustness results. An example where this condition is not met but the system still converges monotonically would be where the eigenvalues are all less than one and so the system converges, and the reference trajectory is such that it never excites the matrix in such a way as for the error to increase from one iteration to the next. Where $QL(y_{ref})$ is close enough to an eigenvector of $Q(I - LP)$ this could clearly happen.

In Norrlöf (1998) the roles of L and Q are fully discussed and a design procedure is suggested for choosing the matrices. The best performance is clearly obtained by setting $L = P^{-1}$ and $Q = I$, however this generally leads to a lack of robustness (this will be shown in Chapter 6), and so an appropriate trade-off is usually required. Here we will briefly discuss the convergence conditions when applying P- and D-type ILC.

3.3.4 Comparison of P-type and D-type ILC

Recall the condition for asymptotic stability: $|\lambda_i(G_nQ(I - LP)G_n^T)| < 1$ for all i. For a relative degree 0 plant it may be possible to achieve this criteria with a P-type algorithm (with $Q = I$) since

$$L_P P = \begin{pmatrix} \gamma & 0 & 0 & \cdots \\ 0 & \gamma & 0 & \cdots \\ 0 & 0 & \gamma & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} D & 0 & 0 & \cdots \\ CB & D & 0 & \cdots \\ CAB & CB & D & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} = \begin{pmatrix} \gamma D & 0 & 0 & \cdots \\ \gamma CB & \gamma D & 0 & \cdots \\ \gamma CAB & \gamma CB & \gamma D & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \tag{3.34}$$

Using this learning matrix, all that would be required to fulfil the condition would be that $|1 - \gamma D| < 1$. Therefore with a small enough γ of the correct sign the condition can always be satisfied. Since $Q = I$ this convergence would also be exact to the reference signal.

However, for a relative degree 1 plant the Ds in the above matrices drop out and we are left with a matrix $L_P P$ that cannot be full row rank no matter what value of γ we take. Therefore, with $Q = I$, the convergence criteria can not be met, since the smallest

possible value that the largest eigenvalue could take would be 1. Therefore P-type ILC cannot control a relative degree 1 plant (without the filter Q).

Returning to the D-type algorithm of equation 3.24, if we assimilate the constant $\frac{1}{T}$ into γ; in controlling a relative degree 1 plant we now have $L_D P$ as follows

$$
L_D P = \begin{pmatrix} -\gamma & \gamma & 0 & \cdots \\ 0 & -\gamma & \gamma & \cdots \\ 0 & 0 & -\gamma & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 & \cdots \\ CB & 0 & 0 & \cdots \\ CAB & CB & 0 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{3.35}
$$

$$
= \begin{pmatrix} \gamma CB & 0 & 0 & \cdots \\ \gamma CAB - \gamma CB & \gamma CB & 0 & \cdots \\ \gamma CA^2 B - \gamma CAB & \gamma CAB - \gamma CB & \gamma CB & \cdots \\ \vdots & & \vdots & & \vdots & \ddots \end{pmatrix}. \tag{3.36}
$$

The diagonal of γs one above the major diagonal in L_D has brought the CB terms up onto the major diagonal in $L_D P$ enabling the matrix $G_1 L_D P G_1^T$ to be full row rank and therefore making it possible for the convergence criteria to be met. Now we have the exact convergence condition of $|1 - \gamma CB| < 1$. Note that this condition is identical for both forms of the D-type learning law given above (equations 3.24 and 3.25).

This can therefore be likened to the continuous-time relative degree conditions given in Section 3.2, where the number of differentiations in the controller must equal the relative degree of the plant. Here, in discrete-time, the plant's relative degree must match the forward shifts in the controller.

Remark 3.6. Further in this thesis the G_n notation will be dropped as it is not strictly necessary. As Theorem 3.5 provides only a sufficient condition it will be replaced by the stronger requirement for monotonic convergence of $\|Q(I - LP)\| < 1$. Both of these can be regarded as essentially the same since the matrix Q can simply be chosen to have n rows of zeros at the bottom and n columns of zeros on the right. To show this first take a matrix Q. Another matrix Q_n can be denoted such that Q_n is identical to Q save for the last n columns and rows that are set to zero. It is easily shown that $\|Q_n(I - LP)\| = \|G_n(Q(I - LP))G_n^T\|$ and also that replacing Q with Q_n in the algorithm of equation 3.19 leads to the same sequence of signals at the output.

3.4   ILC Terminology

The following sections will detail some of the terminology used in the ILC literature, including discussions on different ILC algorithms and techniques. For more detail on different aspects of ILC see: the survey papers Ahn et al. (2007), Owens and Hätönen (2005) and Moore et al. (2006); a special issue of the International Journal of Control

on iterative learning control (see Moore and Xu, 2000, for the editorial); and an article in the IEEE Control Systems Magazine (Bristow et al., 2006).

3.4.1 Higher order algorithms

Higher order algorithms use more than just the previous iteration's error for the update. As many past iterations can be summed as deemed necessary. In equation 3.37, N previous errors are summed, each having a separate scalar factor β_n, allowing the learning matrix β to be weighted differently depending on which past error is being considered. This method may help to reduce the effects of disturbances and changing initial conditions as the errors are being averaged over a longer window.

$$u_{k+1} = u_k(t) + \sum_{n=0}^{N} \beta_n e_{k-n}(t) \tag{3.37}$$

Evidence has been produced arguing both for (Bien and Huh, 1989; Gunnarsson and Norrlöf, 2006; Hätönen et al., 2006; Sun and Wang, 2001; Chen et al., 1998) and against (Xu and Tan, 2001; Norrlöf, 2000a; Saab, 2006) higher order algorithms. Occam's razor springs to mind, higher order algorithms adding complications to the system and requiring more processing power, perhaps without any gain in performance.

In the lifted system representation an N^{th} order algorithm (Norrlöf, 2000b) can be written as

$$u_{k+1} = \sum_{j=k-N+1}^{k} \left(Q_{k-j+1}(u_j + L_{k-j+1}(e_j)) \right). \tag{3.38}$$

3.4.2 Adaptive ILC

Adaptive ILC (AILC) algorithms use adaptive techniques to change the controller, as opposed to the previous methods that have changed the input signal (Owens and Munde, 1996; Hätönen et al., 2004; French et al., 1999). Owens and Munde (2000) describes an AILC controller for a linear time-invariant (LTI), single-input-single-output (SISO), relative degree one, minimum phase plant with the update rule given by

$$u_{k+1}(t) = u_k(t) + \text{sgn}(CB)[(K_{k+1}e_{k+1}(t) + (F_{k+1}e_k)(t)], \quad 0 \le t \le T \tag{3.39}$$

where

$$K_{k+1} = K_k + c\|e_k\|^2, \quad c > 0. \tag{3.40}$$

In this algorithm the current error is fed back and weighted by the adapting K_{k+1} term. This means that the most recent state of the output appears at the input, unlike non-adaptive ILC, where the plant is effectively working open-loop for the duration of the iteration and the new open-loop input is calculated off-line between iterations. The

algorithm is based on high-gain stabilisation, with similarities to the universal adaptive stabilisation control law (see Ilchmann, 1991) — the idea being to create a control law capable of stabilising any plant within a certain plant structure (LTI, SISO, relative degree one and minimum phase).

Choi and Lee (2000) uses an adaptive technique with ILC, where "*uncertain parameters are estimated in the time domain whereas repetitive disturbances are identified and compensated in the iteration domain*", using the ILC to improve trial-to-trial performance and the adaptive control to estimate uncertain parameters.

3.4.3   Model-based ILC

As is apparent from previous sections of this thesis that the ideal input trajectory is $u(t) = P^{-1}y(t)$ (provided it lies within the plant's domain). Early ILC inverse model work was based on this approach (Lee et al., 2000), using a model of the plant inverse to update the next trajectory (Oh et al., 1988; Lucibello, 1992). If the plant is known explicitly then this method should work perfectly; however, a model of a plant will seldom be perfect and the plant inverse would likely be highly sensitive to inaccuracies. Robustness to model uncertainty is therefore a key issue in inverse model ILC (Harte et al., 2005). Also, as most plants are relative degree one or higher (in continuous time), the inverse would have to include a differentiator, greatly amplifying high frequency disturbances.

In Markusson et al. (2001) the inverse of a linearised model of the plant is used and weighted to achieve stability. As the inverse of a non-minimum phase system is unstable the paper also suggests the use of non-causal filtering techniques (this is possible as the filtering is carried out on data along the previous trial, see Section 3.4.5).

3.4.4   Norm optimal ILC

As the aim of ILC is to reduce the error signal, this can be deemed an optimisation problem where an attempt is made to minimise the error signal in terms of some norm. In essence, this is done by finding the input signal that produces the best output (Hätönen, 2004).

$$\min_{u_k \in \mathcal{U}} \|e_k\| \tag{3.41}$$

If the plant is well known and is invertible the solution is simply $u_k = P^{-1}y_{ref}$, where y_{ref} is the reference signal. Often the plant is not well known or not invertible, in which case an approximation of the plant can be used to find the solution iteratively.

For optimal ILC strategies a cost function is developed and an attempt is made to minimise it to give an optimal input for the next trial (Amann et al., 1996b,a; Amann,

1996; Verwoerd, 2005; Hätönen, 2004). This cost function can include different terms to allow the system to retain certain properties. As an example Amann et al. (1996b) and Amann et al. (1996a) both include a term that penalises large changes in the input signal from trial-to-trial in order to maintain an appropriate step size, with the cost function given by

$$J_{k+1}(u_{k+1}) = \|e_{k+1}\|^2 + \|u_{k+1} - u_k\|^2. \tag{3.42}$$

Gunnarsson and Norrlöf (1999) includes a term that penalises the energy of the input signal, constraining rather than penalising a change in error:

$$J_{k+1} = \|e_{k+1}\|^2 + \|u_{k+1}\|^2$$
$$\|u_{k+1} - u_k\|^2 \leq \delta. \tag{3.43}$$

Both of these also penalise the size of the error signal and propose weighting the different terms in the cost functions in order to provide a trade-off between the control objectives. Tao et al. (1994) includes a term that penalises large derivatives of the input signal; which has the effect of preferring a smoother input, similar to the filtering techniques explained earlier that increase robustness.

More recently, Chu and Owens (2009) provided a link between the norm optimal algorithm of equation 3.42 and a successive projection algorithm given in Owens and Jones (1978). In essence the latter shows that, for two closed convex sets K_1 and K_2 in a real Hilbert space H with a non-empty intersection $K_1 \cap K_2$, by iteratively projecting a point from one set to the other the point converges to the intersection. An example of this is depicted in Figure 3.4, where a point $k_0 \in H$ is first projected onto K_1, the result is then projected onto K_2 etc.
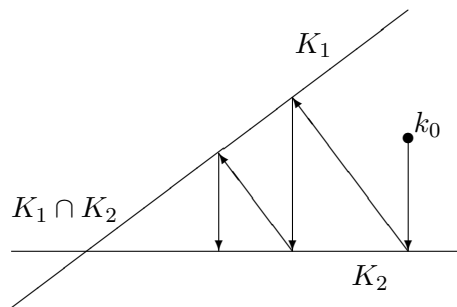
FIGURE 3.4: Example sequential projection

Chu and Owens (2009) shows that this corresponds to the norm optimal algorithm of equation 3.42 if one of the sets is taken to be the set of signals compatible with the plant

and the other as the set of signals with zero error:

$$K_1 = \{ (e, u) \in H \mid e = y_{ref} - Pu \}$$
$$K_2 = \{ (e, u) \in H \mid e = 0 \}. \tag{3.44}$$

Clearly the intersection of these two sets consists of the input u and error e such that both are compatible with the plant and the error is zero — the ideal solution. The norm optimal algorithm is shown to fit the successive projection theorem as demonstrated here with the help of Figure 3.4. Recall the norm optimal algorithm of equation 3.42:

$$J_{k+1}(u_{k+1}) = \|e_{k+1}\|^2 + \|u_{k+1} - u_k\|^2. \tag{3.45}$$

The point $(0, u_k)$ is in K_2, and e_{k+1} and u_{k+1} are compatible with the plant therefore $(e_{k+1}, u_{k+1}) \in K_1$. Minimising $J_{k+1}(u_{k+1})$ can clearly be viewed as a projection of $(0, u_k)$ onto K_1. The projection of the result back on to K_2 is achieved simply by setting $e_{k+1} = 0$. We now have the point $(0, u_{k+1}) \in K_2$. We can now repeat the process by projecting this back on to K_1 by minimising $J_{k+2}(u_{k+2})$. For a full rigorous mathematical explanation see Chu and Owens (2009).

The link between norm optimal and successive projection also allows the implementation of an accelerated algorithm from Owens and Jones (1978) for norm optimal ILC.

For more on norm-optimal ILC and explanations of various cost functions and their properties see Hätönen (2004).

3.4.5 Causality

In an iterative learning context, causality is a property of the two-dimensional nature described previously. At time t on iteration k the controller can have access to all information along t for any iteration $< k$. This has wide implications: data from an entire trial can be processed by an algorithm and filtered at will, and any delays caused can be removed before the new signal is applied to the system.

For ILC, non-causal is a label applied to algorithms that exploit data from *after* the current time along *previous* iterations.

Using the lifted system framework described in Section 3.3.1 this would be equivalent to an update rule where non-zero terms exist in L above the major diagonal. D-type ILC in discrete time is therefore non-causal as the learning matrix is empty save for the major diagonal and the one above it (or simply the diagonal directly above the major diagonal depending on which D-type representation is used — see Section 3.3.2). Within traditional feedback control of continuous systems derivative control is fundamentally non-proper due to a transfer function with more zeros than poles and so has to be

approximated. This problem can be resolved in ILC as a non-causal controller is able to differentiate an entire trajectory; although problems relating to noise amplification by differential operations still remain.

Several papers including Norrlöf and Gunnarsson (2005), Verwoerd et al. (2002) and Phan et al. (2000) show that (in general) non-causal ILC algorithms are required to fulfil monotone convergence criteria and enable the error to decay to zero. This is supported by results in Section 3.3.2, where it is shown that a non-causal algorithm would be required for a relative degree 1 system to to follow a reference trajectory exactly, although a causal algorithm could lead to convergence to a non-zero error by including the filter Q.

In Verwoerd (2005) it is shown that "*to each element in the set of causal iterations there corresponds a particular causal feedback map...which delivers the exact same performance as the ILC scheme, yet without the need to iterate.*" The used of non-causal ILC is advocated, particularly for the control of non-minimum phase systems and systems with relative degree higher than one. Goldsmith (2001, 2002) makes the strong statement that there is no reason to implement causal ILC, and that non-causal is the "*only viable route for ILC.*"

3.4.6 Other assorted ILC terminology

Terminal ILC considers the case where instead of attempting to follow a reference trajectory, the system is designed to simply reach a certain output by the end of the trial (see Chen et al., 1997; Gauthier and Boulet, 2005). Such tasks can be found in rapid thermal processing (Xu et al., 1999).

Set-point regulation is a type of ILC implementation where the output is attempting to remain constant over repeated trials in the face of a repeating disturbance. A building's central heating/air conditioning system could be fitted to this model: the outside air temperature could be regarded as a periodic disturbance, and the aim of the system would be to maintain a constant temperature inside.

Most of the algorithms considered so far have used the error along previous trials to determine an appropriate input signal for the next trial. 'Current iteration tracking error' (CITE) algorithms also make use of the error along the current trial when determining the input signal (Owens, 1992; Chen et al., 1996; Chin et al., 2004; Norrlöf and Gunnarsson, 2005; Ma et al., 1993; Amann et al., 1994). This could be considered as a standard ILC implementation with an added time-domain feedback controller. The adaptive ILC algorithm in Section 3.4.2 uses CITE together with high gain feedback for plant stabilisation.

Hybrid controllers involve implementing more than one controller on the same plant.

This usually consists of wrapping a feedback controller around a plant before adding the ILC loop. This could be done to reduce the effect of non-repeating disturbances or simply to change the properties of the system that the IL controller has to control. This can be also be used to stabilise a plant before adding an IL controller to minimise instability issues as explained later in Section 3.6.4. Examples of hybrid ILC systems include Tayebi and Islam (2006), where an adaptive ILC controller is implemented together with a PD controller on a robot manipulator; and Barton et al. (2000), which combines the plant in feedback with a PID controller and then applies the ILC around the entire closed-loop.

3.5 Relaxing of Restrictions

Some of the restrictions of ILC described in the introduction can be relaxed under certain circumstances.

3.5.1 Allowing some error

Sometimes it may not be possible for an ILC algorithm to completely reduce the error signal to zero. For example, if a system is requested to follow a step input the ILC algorithm would continue to increase the plant's input signal in order to reduce the error (a similar example is given in Amann et al., 1996b). This could lead to a spike in the control signal attempting to force the plant output as quickly as possible to the desired value. Although the output signal may well converge the input would be approaching a signal not in the plant's domain. In cases like this it would be more appropriate to allow a small error and maintain an acceptable control signal. This example is very fabricated but the principle is manifest. In this case filtering of the input signal could remove the spike, achieving an acceptable input but sacrificing the possibility of perfect tracking.

It may also not be possible to achieve the required robustness conditions and also deliver perfect tracking.

3.5.2 Resetting between iterations

Most ILC algorithms' stability proofs assume that the system's initial conditions are completely reset before commencing each iteration; in practice this may not occur and may lead to instability (Lee and Bien, 1991). Methods for combating this problem include accepting that the initial error may be present and developing algorithms that are robust to it (Jiang and Unbehauen, 1999; Yong et al., 2002), or measuring this error and using it to adjust the control signal (Chen et al., 1996, 1999; Lee and Bien, 1996).

3.5.3 Changing plant dynamics

Some ILC algorithms are designed around the problem of changing the controller as the plant changes rather than simply reducing the error. Examples would include mechanical systems where wear-and-tear would change the system properties over time. ILC could be used in this context to maintain a sensible output.

Some industrial applications require the use of very accurate robot motion, this requires expensive equipment and modelling of components in order to keep uncertainty to a minimum. A successful implementation of ILC may result in being able to use cheaper actuators, allowing the learning algorithm to compensate for the poor knowledge of the system characteristics. This would shift the costs away from expensive hardware to cheaper software.

3.6 Long Term Stability

ILC often suffers from long term instability issues (Huang and Longman, 1996; Longman, 2000; Barton et al., 2000): the error profile appearing to converge for a number of iterations and then suddenly diverging, driving the system unstable. Figure 3.5 (from Cai, 2009) shows the mean squared error during an ILC experiment on a gantry robot (see Chapter 6 for more detail on the robot). The red line shows the algorithm converging and then after approximately 120 iterations start to diverge. The experiment was terminated after 200 iteration to prevent damage to the apparatus.

Long term instability is a potential barrier to the use of ILC in many applications as once a controller is implemented it would be ideal to leave it running, confident that it would perform as expected. This is especially pertinent when considering slowly time varying systems, since it may not be possible to 'switch off' the learning algorithm once a system has converged satisfactorily. It would therefore be preferred to be in a position to explain this instability and be able to predict and hence prevent it. There is much discussion in the literature as to the possible causes of this phenomenon. The following sections will review some of these ideas and describe a few of the practical solutions that are used to avoid it.

3.6.1 Frequency domain

The instability could be a form of the waterbed effect: the ILC algorithm reducing unwanted low frequency components of the error but in doing so increasing the high frequencies. These high frequencies are negligible at first but as the low frequencies are reduced the high frequencies become more dominant (Huang and Longman, 1996). As ILC is based on improvement using the previous iteration (which in turn used the
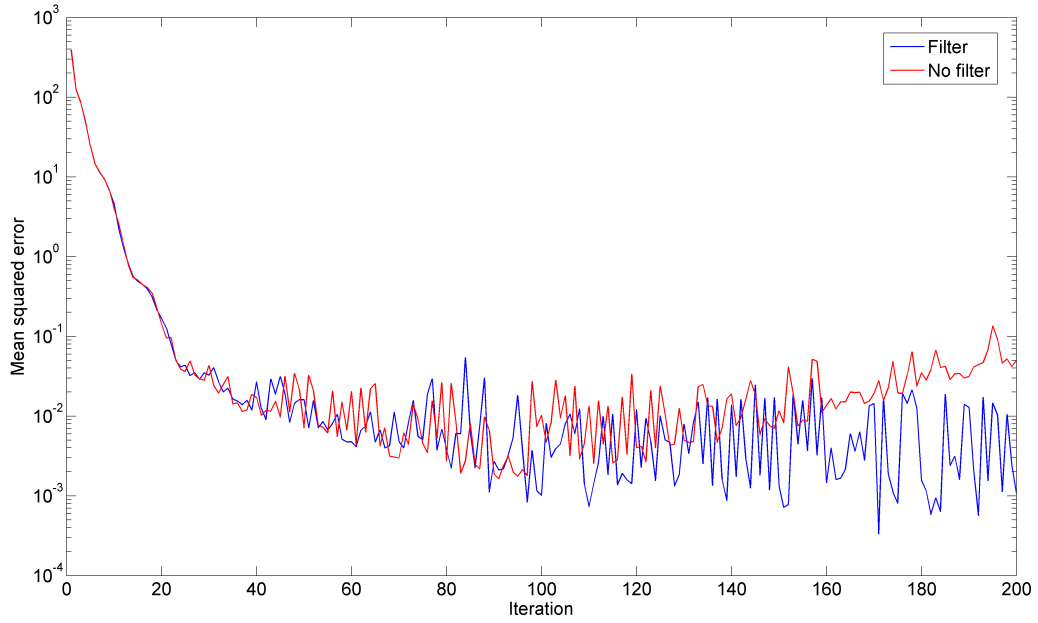
FIGURE 3.5: Long term stability simulation — error profile as a function of iteration number

iteration before that) the error is effectively summed over all the iterations meaning that these high frequency components may be compounded, causing the onset of instability. This is explained in Norrlöf (1998) in terms of the system's Nyquist diagram, where the high frequency components are found to lie outside the unit circle.

3.6.2 Time domain

Iterative learning controllers typically perform all the calculations on the whole signal off-line between iterations. With a badly chosen learning rate the ILC algorithm may improve the tracking at the start of the iteration at a cost to the error as a whole.

The basic ILC algorithms make use of the errors at each time instance and feed these forward for use in the next trial. The ILC algorithm is therefore basing its current plant input on information gathered in the previous iteration. This ignores any changes that have occurred in the current iteration. At the start of any trial the error is zero and so this effect is negligible; towards the end of a trial though, the rest of the trial will have evolved differently to the previous one and so the 'improvement' on the input signal at that point may not be as successful. (This process is also explained in Ratcliffe, 2005.)

The result of this is that the learning is concentrated toward the beginning of each cycle and the accuracy spreads further along with each iteration, as the smaller error is forced further along the trial. One corollary of this is the exponentially weighted norm in Arimoto et al. (1984) — explained in Section 3.2.

Figure 3.6 illustrates this phenomenon. The error toward the start of the 100th iteration is small, but the cost of this has been driven to the end of the iteration. If this system were allowed to continue the error would eventually be reduced over the whole domain; however it would take many iterations.
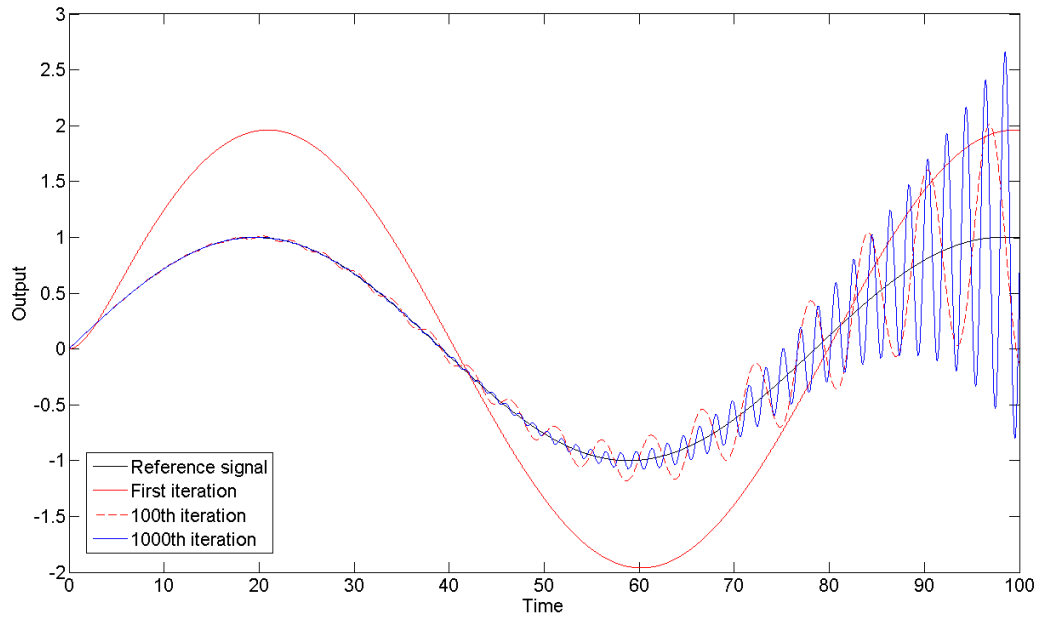
FIGURE 3.6: Long term stability simulation — graph showing an output after 1, 100 and 1000 iterations

This process is explained in several ILC papers and theses including Longman (2000) and Hätönen (2004).

3.6.3 Robustness

The instability could be caused by robustness issues. The controller may be long term stable on the model of the plant, but some of the unmodelled dynamics may then drive the real system to instability. Robustness could be considered to be an umbrella over the other possible causes of the long term stability problem since, once a nominal system is stable, the more robust it is the larger the errors it will be able to withstand without compromising stability.

This thesis will concentrate on this issue, using robustness analysis to prove the stability of plants within a neighbourhood of the model. Section 3.7 will therefore go into more detail on other work that has been undertaken on robustness for ILC.

3.6.4 Practical solutions

Several solutions have been developed to counter the problem of long term instability.

- Quantisation (Longman, 2003; Hsin et al., 1997).

  It has been observed that the error can decrease for many iterations and reach very low levels, however these low level signals accumulate and eventually emerge as long term instability. Quantisation can therefore be introduced into the control system to ensure that at high frequencies, where the Nyquist curve lies outside the unit circle, the errors are low enough that the are simply not captured as the signals are discretised. This does have adverse effects on the final convergence of the ILC system, since the introduction of quantisation will apply across all frequencies.

- Suspending the learning whilst the system's error is low (Barton et al., 2000).

  Once the error has reached a given value the learning could be switched off and the system allowed to run with no learning. Once the error has grown larger again it could be switched back on. The advantage of this would be for plants with changing dynamics as the learning would kick back in if the controller decided it was necessary. The disadvantage is that the estimate of this threshold may be too conservative; the algorithm may be capable of reducing the error further. In order to be able to switch off the learning the plant may also need an additional, non-ILC feedback loop wrapped around it in order to maintain stability.

- Resetting the system or stopping the learning once divergence is detected (Longman, 2000, 2003).

  Resetting the system back to some known previous iteration where the error was acceptable is the obvious solution but could undo some good learning that had been done.

  It could be assumed that one (or more) divergent signal(s) signifies the onset of instability and at a chosen point all learning could cease. As all systems possess unknown disturbances, small amounts of noise could cause the error profile to diverge and the learning to be prematurely cut short.

- Filtering of the input signal.

  If the instability arises out of high frequency oscillations, as observed on the gantry robot from Ratcliffe (2005), then one solution could be to add a low-pass filter. Figure 3.5 shows an ILC implementation with and without a low-pass filtering of the signal.

  In the algorithm given in Section 3.3.1, the Q is often implemented as a filter in order to remove some of the unwanted frequency components; however a reduction of Q (from $Q = I$) leads to larger tracking errors in the converged output signal.

Q can therefore be tuned to approximate the identity at low frequencies in order to provide a decent level of performance, and zero at high frequencies in order to reject high frequency disturbances. De Roover (1996), de Roover and Bosgra (2000) and Harte et al. (2005) suggest the use of Q as a low pass filter for this purpose. This is returned to in Section 6.2.1.

Filtering would add phase lag; however, this can be removed by instead using non-causal, zero-phase filtering between iterations — possible since ILC algorithms can have access to signals along the whole time series over previous iterations (Ratcliffe, 2005). Ratcliffe et al. (2005b) compares various filtering techniques on their ability to prevent long term instability, including band-stop, low-pass filter and zero-phase filtering. These are then compared to an aliasing technique, where the output of the controller is sampled at less than twice the plant's resonant frequency.

All of these practical solutions go some way to improving the situation but all require at least some amount of user input to define courses of action. This works fundamentally against the ILC aim of letting the machine do the learning. Also, these 'fixes' often have little or no mathematical grounding. This leads to the questions: how do they work? and why are they required?

A much better result would explain, predict and allow prevention of this instability mathematically, permitting the implementation of ILC algorithms knowing them to be stable. Providing an understanding of how and why the instability occurs could lead to more principled or efficient solutions.

3.7   Robustness for ILC

Robustness analysis for iterative learning control can be split into two groups:

- Robustness to model uncertainty.

  Applicable throughout control literature and not just ILC, as explained in Chapter 2, robustness to model uncertainty is required throughout in order to guarantee that the controller that stabilises the nominal model of a plant also stabilises the plant itself. This will be the central theme of this thesis.

- Robustness to initial state error.

  As explained in Section 3.5.2, a principal requirement of ILC is that the plant is returned to the same initial conditions for the start of each iteration. This may not be met in practice and so algorithms must be able to cope with these disturbances, to a certain degree. This subject will not be covered in any more detail in this thesis as we will concentrate on robustness to model uncertainty, for more detail and references to relevant papers refer back to Section 3.5.2.

A significant proportion of ILC literature that concerns robustness concentrates on multiplicative and additive perturbations. Algorithms are often proven to converge under nominal conditions and afterwards examined under multiplicative perturbations to prove that they are robust.

De Roover (1996) and de Roover and Bosgra (2000) both consider ILC in an $\mathcal{H}^\infty$ framework for the synthesis of a robust ILC controller. Using the ILC update equation $u_{k+1}(t) = Q(u_k(t) + Le_k(t))$ (identical to the update in equation 3.19) and the z-transform the monotonic convergence criteria is given by $\|Q(z)(I - L(z)P(z))\|_\infty < 1$ where $\| \cdot \|_\infty$ is the induced 2-norm. In order to achieve perfect tracking Q must equal the identity and so $L(z)$ must approximate $P(z)^{-1}$ to a certain extent (at the very least the shifts in L must match the relative degree of P, see Section 3.3.4). Since this may not be possible Q is used to *robustify* the algorithm with respect to uncertainty in P with a loss of perfect tracking (de Roover, 1996). The papers introduce a robust design procedure based on a multiplicative uncertainty model, where Q is chosen as a low-pass filter and L found using a μ-synthesis procedure.

Moon et al. (1998) and Tayebi and Zeremba (2003) examine ILC as a method to control a plant with an added non-iterative feedback controller. The plant and controller are implemented in a feedback loop and then an IL controller adjusts the plant's input to improve the trial-to-trial performance. Given a multiplicative uncertainty model, if the plant and non-iterative controller obey the robust performance condition, the papers state that the design of an additional IL controller around the pair, such that the whole system is robustly stable, is simple. In Tayebi and Zeremba (2003), for the plant given by the uncertainty set $P_1 = (1 + \Delta W_2)P$ in feedback with a controller C, where S and T are the sensitivity and complementary sensitivity functions, the robust performance condition can be written as

$$\| |W_1 S| + |W_2 T| \|_\infty < 1, \tag{3.46}$$

where W_1 and W_2 are known stable transfer functions. The ILC algorithm is given, with update rule

$$U_{k+1}(s) = W_1(U_k(s) + C(s)E_k(s)), \tag{3.47}$$

where U_k and E_k are the Laplace transforms of the input and error respectively, along trial k. The algorithm is proven to converge in L^2 as $k \to \infty$.

In van de Wijdeven and Bosgra (2007) a new ILC strategy is introduced that is robust to additive or multiplicative uncertainty. Some of the issues with other robustness work on ILC are discussed in the paper: the use of frequency domain techniques, which are inherently conservative due to their assumption of infinite-time horizons and ILC implementations being fundamentally of finite-time; $\mathcal{H}^\infty$ approaches restrict to only causal algorithms which, as discussed in Section 3.4.5, are often not appropriate; and the wide use of filtering in ILC that lacks a theoretical basis for the frequencies to choose.

The paper then suggests a method that is non-causal and utilises the finite-time aspect of ILC to find the learning function L that results in a system that is robust to additive and multiplicative uncertainty.

Donkers et al. (2008b,a) are concerned with the robustness of norm optimal ILC in terms of additive uncertainty. In Donkers et al. (2008b) a MIMO system is examined and sufficient conditions are given under which the algorithm converges robustly and monotonically. The framework makes use of the finite-time aspects of ILC and derives conditions using μ-synthesis techniques. Frequency domain results are also given for SISO systems. Donkers et al. (2008a) discusses the design of an ILC controller based on $\mathcal{H}^\infty$ methods; but one that permits a non-causal controller.

Hätönen et al. (2003) introduces a steepest-descent algorithm for discrete ILC. The algorithm can be made robust to multiplicative uncertainty provided the uncertainty satisfies a positivity condition. Using the update equation

$$u_{k+1} = u_k + \epsilon_{k+1} P^T e_k, \tag{3.48}$$

by a sufficiently small choice of $\epsilon_{k+1} > 0$ the algorithm is monotonically convergent. It is also shown that if the uncertainty is written $P_1 = PU$ (U denotes the multiplicative uncertainty), then if $U + U^T$ is a positive-definite matrix there again exists an $\epsilon_{k+1} > 0$ such that the system is stable. The paper also provides a method of determining ϵ_{k+1}.

Similarly, Harte et al. (2005) describes an inverse ILC algorithm given by

$$u_{k+1} = u_k + \beta P^{-1} e_k. \tag{3.49}$$

As expected, with an exact plant model, the algorithm would converge in a single iteration with $\beta = 1$. Where $P_1 = UP$ (P_1 is the real plant given by the modelled plant and some form of multiplicative uncertainty U), if $U + U^T$ is positive-definite then there exists β^* such that $0 < \beta < \beta^*$ ensures monotonic convergence. A sufficient condition for monotonic convergence is therefore that

$$\sup_{|z|=1} \left| \frac{1}{\beta} - U(z) \right| < \frac{1}{\beta}. \tag{3.50}$$

The paper then makes three remarks concerning the result:

i) To satisfy equation 3.50 it is sufficient that the Nyquist curve of $U(z)$ lies within a circle of radius $\frac{1}{\beta}$ centred at the point $(\frac{1}{\beta}, 0)$.

ii) Since this circle expands to cover the entire right half plane as $\beta \to 0$, for any strictly positive uncertainty $U(z)$, robust monotonic convergence can be satisfied by a sufficiently small β.

iii) The plant uncertainty must have a phase shift less that $\pm 90°$ over all frequencies.

Although a strong result the plant uncertainty being contained within a phase shift of ±90° is a restrictive assumption. As explained at the start of Chapter 2, it is physically impossible to measure the high frequency gain of a system by identification so it is quite possible that the condition would be violated. It would not be expected that a shift of 90° at high frequency with a sufficiently small gain would make a difference to the stability.

Aside from multiplicative and additive uncertainty representations, ILC is examined in Ahn et al. (2005) in terms of its robustness to parametric or interval perturbations to the system's Markov parameters. In the uncertainty representation, the plant is expressed in lifted system form and the Markov parameters that populate the lifted matrix are each given upper and lower bounds. The set of 'interval Markov matrices' is then defined as the set of matrices generated by these bounds, and the nominal matrix consists of the mid-points for each upper and lower bound. The paper then gives bounds on the radius of interval uncertainty around this nominal model such that that all plants contained in this radius are stable and monotonically convergent given a first order ILC algorithm of the following (lifted) form:

$$u_{k+1} = u_k + \gamma e_k. \tag{3.51}$$

It is interesting to note that the maximum robustness to interval uncertainty is found when $\gamma = 0$. It is clear that in this case no learning takes place, highlighting the trade-off that exists between robustness and performance. This will be returned to later in the thesis when examining ILC robustness using the gap.

Ahn et al. (2006) applies a similar technique but instead looks at interval uncertainty in the A matrix of the plant and examines bounds on the eigenvalues and eigenvectors of powers of A to calculate bounds on the interval uncertainty in the Markov matrix. These are then used to design an IL controller that is robust to this uncertainty and monotonically convergent.

More in line with the results of this thesis, ILC and the gap metric are brought together in French (2008), where a 2D gap metric is used to prove that there exists a non-zero stability margin for a class of adaptive ILC algorithms. The adaptive ILC algorithm is used to demonstrate some of the problems associated with the gap metric and ILC. The controller is defined as a gain that increases depending on the error measured on the previous iteration, and the plant is high-gain-stabilisible and relative degree one. If the measured error is high the gain is increased and so ultimately the error is reduced (by the high gain property). However; the stability margin depends on the gain, and so is reduced as the error forces the gain higher. If the given plant is initially unstable (in a classical sense along the trial) the traditional robust stability theory would predict zero stability margin. To overcome this problem it was shown that it is possible to measure the stability margin only from the point at which the system is first stabilised, enabling the algorithm to yield a non-zero stability margin. The paper also uses a 2D norm for

measuring the size of signals in ILC, a variant of which is used in this thesis.

3.8 ILC Implementations

Physical implementations of ILC can be separated into two broad categories. The first consists of work done on real industrial problems that can be solved with an ILC approach; and the second is where investigations are carried out on hardware to examine ILC issues such as the robustness or convergence rates of different ILC algorithms, with no specific industrial objectives defined. This section will detail a few examples from each category. For a bibliography of ILC applications (up to 2004) see Ahn et al. (2007).

3.8.1 Problem-oriented ILC

ILC has been implemented alongside functional electrical stimulation (FES) as an aid to stroke rehabilitation. A large number of individuals affected by strokes are left with an impairment in arm functionality. In order for a patient to attempt to regain control of the limb electrical stimulation can be applied to a muscle to move it. The aim is to encourage Hebbian learning to aid patient recovery: in simple terms, if the patient is attempting to move the arm and the arm is moving then this correlation should increase the association between synapses and lead to a stronger neurological pathway between brain and muscle. In order to achieve this for an impaired upper arm, a stroke patient's forearm is strapped to the end of a robotic arm and a task is given to track a trajectory on a horizontal plane. The robot assists the patient to follow the path and maintains a natural feel to the movement. FES is then applied to the patient's triceps in order to induce the arm to follow the given trajectory. ILC is used to develop the FES input as the patient repeats the same task. Papers published on the work include: Freeman et al. (2009c,d) for the design of the workstation; Freeman et al. (2009b,f) and Le et al. (2009, 2010) for identification and modelling of the muscle model when applying FES; Freeman et al. (2008) for experimental work with an adjoint ILC algorithm; Freeman et al. (2009a, 2010) and Hughes et al. (2009b) for experimental results on a set of unimpaired subjects; and Hughes et al. (2009a,c) for results when applied to stroke patients.

When performing lithographic printing of an IC pattern onto a silicon disk (known as a wafer) it is required to print several layers on top of each other with processing of the wafer in between exposures (Dijkstra, 2003). It is necessary to: maintain high accuracy when positioning to ensure that the various layers are correctly aligned; suppress movement once in position to prevent blurring during exposure; and move quickly between exposures to maximise throughput. Dijkstra (2003) is concerned with the application of ILC to this process on a Sire3 wafer-stage, where it is used to accurately control the position of the wafer. Several papers have also been published on the work: Dijkstra

and Bosgra (2003) on the stepping from one position to another with maximum speed and minimal settling time using ILC for input shaping; Dijkstra and Bosgra (2002c) on noise suppression during the process; Dijkstra and Bosgra (2002b) on an implementation of lifted system ILC; and Dijkstra and Bosgra (2002a) on a quadratic criterion ILC algorithm. There are also several publications on the use of ILC for rapid thermal processing, using ILC to control the temperature across a wafer during manufacture: Cho et al. (2003) on model development and identification, sensor location and optimal multivariable ILC; Chin et al. (2001) and Yang et al. (2003) on optimal ILC techniques; and Choi and Do (2001) for a neural network based algorithm.

Outside of these two examples, ILC has been applied to a variety of applications as diverse as: laser cutting using the controlled fracture technique (Tsai and Chen, 2004); dental implants (Huang et al., 2003); the motion of an electrostrictive actuator (Hu et al., 2004); the motion control of a camera tracking a target point in an image-plane (Liu et al., 2002); and the control of traffic levels on a freeway (Hou et al., 2008). See Ahn et al. (2007) for a comprehensive list of other ILC applications.

3.8.2   Experimental ILC

ILC is implemented on a non-minimum phase system in Freeman (2004). The thesis covers the construction of the facility and testing of various ILC strategies including P-, D- and delay-type ILC algorithms. The thesis also examines a phase-lead update law (also in Freeman et al., 2007) and the use of causal and non-causal filtering. Freeman et al. (2005a) concludes that the phase-lead algorithm outperforms P- and D-type ILC in this case although long term stability is not assured and so the paper advocates that the learning law be switched off once a suitable level of error has been achieved. A discrete implementation of a predictive optimal ILC algorithm is also implemented on the test-bed in Freeman et al. (2005b). A second thesis (Cai, 2009) introduces a 'reference shift' algorithm which is implemented on the same experimental facility and shows significant improvements over previous algorithms (Cai et al., 2008b).

Different ILC algorithms have been implemented and compared on a multi-axis gantry robot in two theses at the University of Southampton (Ratcliffe, 2005 and Cai, 2009). The first details the design and build of the test facility and then proceeds to examine different ILC methods. P-type ILC and a norm optimal IL controller are compared with traditional feedback control and shown to offer improved perfomance (also in Ratcliffe et al., 2005a, 2006a,b). Since the P-type controller is ill-suited to non-relative-degree zero systems a variation of the algorithm is implemented: a hybrid system consisting of the ILC with proportional feedback controller and an aliasing module. The aliasing module is compared with other filtering techniques in Ratcliffe et al. (2005b), where both aliasing and zero-phase filtering are found to offer improvement over traditional methods since they add no phase shift. Inverse model-based ILC is implemented in Hätönen et al.

(2007) and an adjoint algorithm in Ratcliffe et al. (2008) that can be made robust to positive multiplicative uncertainty. Cai (2009) uses the facility to investigate stochastic ILC algorithms (seen also in Cai et al., 2008a) and examines the along-the-trial dynamics using linear matrix inequalities (Hladowski et al., 2008, 2009). The robot is also used for research into methods of generating the plant's initial ILC input signal in Alsubaie et al. (2009) and Freeman et al. (in press); and the synchronisation of sub-systems whilst implementing ILC in Freeman et al. (2009e).

3.9   Summary

In this chapter the field of iterative learning control has been briefly introduced; a few of the different algorithms and terminology have been explained and some of the issues of stability have been raised; in particular, the often observed issue of ILC's long term instability, its effects and some of the possible causes.

The trial-to-trial aspect of ILC hinders analysis by introducing a two-dimensional element to the control structure: the time and iteration equations entwining to form a more complex whole. The classical ideas of stability in the time domain are no longer sufficient — for any linear system the output is always bounded over finite time if the input remains bounded — and therefore stability over successive iterations needs to be defined, and also depends on the way that signals are measured. Typically, stability concerns the convergence of the output signal to a given trajectory, but in practice it is nigh-on impossible to achieve tracking with zero error and so an acceptable level of error must be suffered. In order to measure this error a norm must be used, usually L^2 or l^2.

The crux of this thesis is the issue of long term instability. This has been discussed and some of the possible causes and cures mentioned. As it stands, the root of the problem is not well understood and so it is hoped that an investigation into robustness may shed some light onto it. Previous robustness results in the field have been given, mainly relating to multiplicative uncertainty descriptions; and a few implementations of ILC on hardware have been explained, both for the solving of real industrial problems and for the testing and benchmarking of ILC techniques.

The remainder of this thesis will introduce a notion of robust stability and apply it to ILC. Using this to provide a mathematical basis for the stability of ILC systems could go a long way to discovering the mechanisms that cause instability in ILC and hopefully provide some clearer answers about how systems can be developed to better cope with uncertainties.

For more on ILC there have been various survey papers published including Ahn et al. (2007), Owens and Hätönen (2005) and Moore et al. (2006); a special issue of the

International Journal of Control on iterative learning control (see Moore and Xu, 2000, for the editorial); and an article in the IEEE Control Systems Magazine (Bristow et al., 2006).

Chapter 4

The Gap Metric for Tracking

This chapter will build on the uncertainty modelling from Chapter 2 and present the gap metric in both linear and non-linear forms. Robust stability theorems will be given for both and the link between them will be discussed.

The non-linear gap metric will then be extended to incorporate systems with biases, where the standard non-linear gap fails to give appropriate robustness results. The associated 'biased robust stability theorem' was published without proof in Georgiou and Smith (1997b) and so one is given here.

In the latter part of the chapter the biased robust stability theorem is extended to incorporate a notion of tracking. Although the overall intention of this thesis is the stability of ILC systems, this is not directly addressed here. The work will instead take the form of an abstract modification of the gap metric applicable to a wide range of problems. This modification is directed towards ILC, however the link is only made explicit in Chapter 5.

4.1 Linear Gap Metric

The linear gap metric dates back to work in Newburgh (1951), which introduced a topology on which to measure the distance between two operators based on the Hausdorff distance. In Berkson (1963) this was compared to earlier work on 'the opening' from Krein and Krasnosel'skii (1947), showing that the opening and the Newburgh metric were equivalent.

The gap metric was developed into a form suitable for linear systems in Zames and El-Sakkary (1980) and in further detail in El-Sakkary (1985). The problem considered in these papers was robustness to plant uncertainty for plants that are unstable but stabilised via feedback.

For the analysis of stable plants that are continuous on some normed space it is possible to use the continuity principle (Zames and El-Sakkary, 1980; El-Sakkary, 1985) to relate the distance between open-loop plants and the responses of these plants closed-loop. For the unstable case however, bounded operators are no longer defined on the whole of a space but on a proper subspace of it, consisting of the bounded signals that produce bounded outputs. This prevents the direct use of the continuity principle.

To avoid this problem the gap metric describes the difference between two operators via a suitable notion of distance between their graphs. For a plant $P: \mathcal{U}_e \to \mathcal{Y}_e$ the graph is defined as the set of pairs of plant input-output signals that are bounded in norm:

$$\mathcal{M} = \mathcal{G}_P = \left\{ w \in \mathcal{U} \times \mathcal{Y} \ \middle| \ w = \begin{pmatrix} u_1 \\ Pu_1 \end{pmatrix} \right\}. \tag{4.1}$$

Similarly, the graph of the controller is given by

$$\mathcal{N} = \mathcal{G}_C = \left\{ w \in \mathcal{U} \times \mathcal{Y} \ \middle| \ w = \begin{pmatrix} Cy_2 \\ y_2 \end{pmatrix} \right\}. \tag{4.2}$$

Note that together, when under the feedback structure of Figure 4.1, these form the bounded signals that are compatible with the closed-loop system equations $[P, C]$; and also that the graph is a linear subspace when the plant is linear. It must be stressed that these graphs are subspaces (subsets in the non-linear case) of the *bounded* signal space $\mathcal{U} \times \mathcal{Y}$ which is, in turn, a subspace of the *extended* signal space $\mathcal{U}_e \times \mathcal{Y}_e$ on which the plant and controller's input and output are defined.
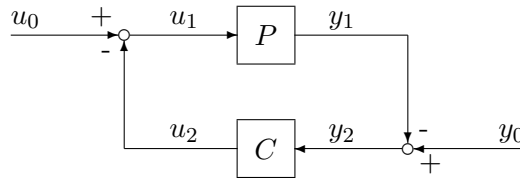
FIGURE 4.1: Feedback configuration $[P, C]$

For a plant P_1 the graph will be denoted $\mathcal{M}_1$. El-Sakkary (1985) then defines the gap between two closed operators as the gap between their graphs:

$$\delta_0(P, P_1) = \delta_0(\mathcal{M}, \mathcal{M}_1) \tag{4.3}$$

where the gap between two graphs is the maximum of the directed gap between the two graphs:

$$\delta_0(\mathcal{M}, \mathcal{M}_1) = \max \left\{ \vec{\delta}_0(\mathcal{M}, \mathcal{M}_1), \vec{\delta}_0(\mathcal{M}_1, \mathcal{M}) \right\}. \tag{4.4}$$

This is required when we wish to refer to the gap as a metric since it is required to fulfil the metric space axiom that states $\delta(x, y) = \delta(y, x)$. The directed gap is then given by:

$$\vec{\delta}_0(\mathcal{M}, \mathcal{M}_1) = \sup_{w \in \mathcal{M}, \ w \neq 0} \ \inf_{v \in \mathcal{M}_1} \frac{\|w - v\|}{(\|w\|)}. \tag{4.5}$$

Graphs are mainly characterised in terms of coprime factors. Recall the coprime factor results given in Section 2.2.3:

From equation 2.20, the set of perturbed plants is given by

$$\mathcal{P}_\Delta = \left\{ P_1 \in \mathcal{R} \ \middle| \ P_1 = (N + \Delta_N)(M + \Delta_M)^{-1}, \ \left\| \begin{bmatrix} \Delta_N \\ \Delta_M \end{bmatrix} \right\|_\infty < \frac{1}{\gamma} \right\}, \tag{4.6}$$

where $P = NM^{-1}$ forms a normalised right coprime factorisation of P. Then $[P_1, C]$ is stable for all $P_1 \in \mathcal{P}_\Delta$ provided $[P, C]$ is stable and $\|M^{-1}(I - CP)^{-1}[C \ -I]\|_\infty \leq \gamma$.
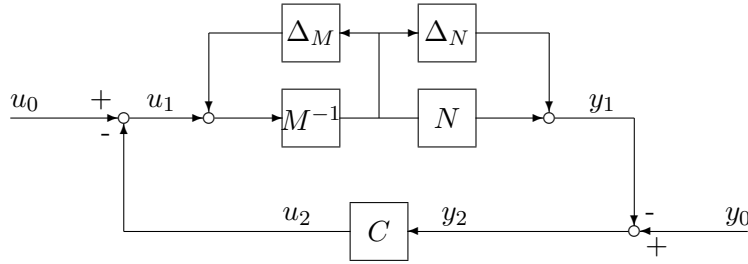
FIGURE 4.2: Coprime factor uncertainty block diagram representation

Theorem 4.1. (From Proposition 1.33 of Vinnicombe, 2001.) *Let* $G = [M, N]$ *where* $NM^{-1} = P$ *is a right coprime factorisation of the plant* $P \colon \mathcal{H}^2 \to \mathcal{H}^2$. *Define the graph in* $\mathcal{H}^2$ *as:*

$$\mathcal{G}_P = \left\{ w \in \mathcal{H}^2 \times \mathcal{H}^2 \ \middle| \ w = \begin{pmatrix} u_1 \\ Pu_1 \end{pmatrix} \right\}. \tag{4.7}$$

Then this graph is given by:

$$\mathcal{G}_P = \left\{ Gq \in \mathcal{H}^2 \times \mathcal{H}^2 \mid q \in \mathcal{H}^2 \times \mathcal{H}^2 \right\}. \tag{4.8}$$

Proof. Given any pair $(u, y) \in \mathcal{G}_P$ then $y = NM^{-1}u$. If we choose $q = M^{-1}u$ then $(u, y) = Gq$. Since N and M are right coprime there exists $X \in \mathcal{H}^\infty$ such that $XG = I$. We can therefore conclude that $q = X(u, y) \in \mathcal{H}^2$. Conversely, for any $q \in \mathcal{H}^2$ and $(u, y) = Gq$ it is clear that $y = NM^{-1}u$ with $u, y \in \mathcal{H}^2$. $\qquad\square$

Based on the equivalence of $L^2(\mathbb{R}+)$ and $\mathcal{H}^2$ (see Section 2.1.5) we can therefore conclude

that the graph of the above P in the time domain can be written as the inverse Laplace transform of the graph in the frequency domain, $\mathscr{L}^{-1}(\mathcal{G}_P)$. (We use $\mathcal{G}_P$ to denote the graph in either the time or frequency domain. The domain being considered will be clear from the context in which it is used.)

The full extent of the relationship between the coprime factorisation perturbation work and the gap metric appears in a large number of papers, notably in Georgiou and Smith (1990). It is stated that robustness optimisation in the gap metric is equivalent to robustness optimisation for coprime factor perturbations and they share the same ball of uncertainty:

Theorem 4.2. (From Theorem 4 of Georgiou and Smith, 1990.) *Consider a system P with normalised coprime fraction $P = NM^{-1}$ and a controller C which stabilises P. Take a real number $0 < b \leq 1$. Under the assumption that $[P_1, C]$ is well-posed, the following are equivalent:*

  a) *$[P_1, C]$ is stable for all P_1 with transfer function $P_1 = (N + \Delta_N)(M + \Delta_M)^{-1}$ where $\Delta_N, \Delta_M \in \mathcal{RH}^\infty$ and $\|(\Delta_N, \Delta_M)\|_\infty < b$.*

  b) *$[P_1, C]$ is stable for all P_1 with $\vec{\delta}_0(P, P_1) < b$*

Therefore b determines a neighbourhood in the given topology in which all plants in the uncertainty set are stable. The question then arises as to the choice of b in the above theorem. The distance between the two plants is denoted $\vec{\delta}_0(P, P_1)$, and b therefore becomes the stability margin. The stability margin applicable to this distance can be written in terms of the nominal plant and controller as follows (Georgiou and Smith, 1990; Foias et al., 1990; Georgiou and Smith, 1992):

$$b_{P,C} = \begin{cases} \left\| \begin{pmatrix} I \\ P \end{pmatrix} (I - PC)^{-1} (I, -C) \right\|_\infty^{-1} & \text{if } [P, C] \text{ is stable} \\ 0, & \text{otherwise} \end{cases} . \qquad (4.9)$$

To summarise in a single theorem (for example see Theorem 1 of Georgiou and Smith, 1992 or Proposition 1.1 of Vinnicombe, 1993):

Theorem 4.3. (From Theorem 1 of Georgiou and Smith, 1992.) *A globally well-posed system $[P_1, C]$ is stable (gain and BIBO stable, see Section 2.1.6) if the nominal system $[P, C]$ is stable and the distance between the two plants is less than the stability margin ($b_{P,C}$) of the nominal system:*

$$\vec{\delta}_0(P, P_1) < b_{P,C}. \qquad (4.10)$$

Vinnicombe (1993, 2001) describe the 'ν-gap' metric based on winding numbers and compare it to the gap metric; and Vidyasagar (1984) introduces the 'graph metric'

using a coprime factorisation technique that is also able to be applied to unstable plants, which is also compared to the gap in Vidyasagar (1984) and Zhu (1989). All of these papers are based on the same topology as the gap metric, being the coarsest topology on which feedback stability is maintained for small neighbourhoods around the plant, and the map from plant to closed-loop operators remains continuous (Pascoal et al., 1993; Cantoni and Vinnicombe, 2002). Neither of these alternate gaps will be discussed in this thesis as the work here will mainly be based on a non-linear adaption of the gap metric of Zames and El-Sakkary discussed above; although it should be noted that the ν-gap of Vinnicombe (1993) provides the tightest possible metric for linear systems using the graph topology and is often easiest to work with as it relates directly to the frequency response and to the familiar winding number conditions used in the Nyquist stability criterion. In a recent paper, Lanzon and Papageorgiou (2009), winding number conditions are used to define a distance between two plants that is easily applicable to a range of uncertainty representations.

4.2 Non-Linear Gap Metric

This section presents work from Georgiou and Smith (1997a) concerning the generalisation of the gap metric mentioned in the previous section to non-linear systems. The theorems are not specific to any signal space, such as L^p. All that is required is that the defined signal space has an applicable truncation definition, and so the domain of the signal space must be totally ordered. This in turn allows definitions of extended signal spaces, stabilisability and causality to be available.

The non-linear method from Georgiou and Smith (1997a) determines a distance between plants based on the geometry of their graphs in the same way as its linear counterpart. In the non-linear case, however, the graph is no longer a subspace and so appears in a more abstract fashion. Before outlining the non-linear gap a parallel projection operator will be introduced.

4.2.1 A parallel projection operator

In the previous section the stability margin $b_{P,C}$ was given in equation 4.9. It is remarked in Doyle et al. (1992) that the elements of this margin consist of the closed-loop transfer functions of the system: sensitivity, complementary sensitivity etc; and therefore in order to maximise robustness these quantities must be minimised. The paper then extends this to the non-linear case by generalising these closed-loop transfer functions to a pair of non-linear operators which are (for a stable system) parallel projections onto the graphs of the plant and controller. For an unstable system the operators will map to the signals in the extended space that are compatible with the plant and controller and therefore not the graphs, since the graphs only contain the compatible bounded signals.
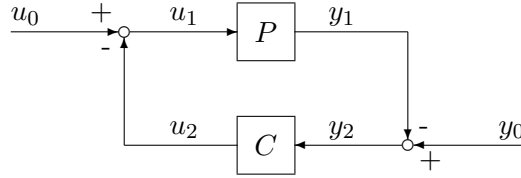
FIGURE 4.3: Feedback configuration $[P, C]$

Recall from equation 2.13 the map $H_{P,C}$ that maps the external signals to the internal signals. We can now define two parallel projections, $\Pi_{\mathcal{M}//\mathcal{N}}$ and $\Pi_{\mathcal{N}//\mathcal{M}}$, given by

$$\Pi_{\mathcal{M}//\mathcal{N}}: \mathcal{W} \to \mathcal{W}_e: \begin{pmatrix} u_0 \\ y_0 \end{pmatrix} \mapsto \begin{pmatrix} u_1 \\ y_1 \end{pmatrix}$$

$$\Pi_{\mathcal{N}//\mathcal{M}}: \mathcal{W} \to \mathcal{W}_e: \begin{pmatrix} u_0 \\ y_0 \end{pmatrix} \mapsto \begin{pmatrix} u_2 \\ y_2 \end{pmatrix}. \tag{4.11}$$

Note that from equations 2.12, 2.13 and 4.11 we can demonstrate that the stability of any one of the operators ($H_{P,C}$, $\Pi_{\mathcal{N}//\mathcal{M}}$ and $\Pi_{\mathcal{M}//\mathcal{N}}$) implies the stability of the others. Also, if the closed-loop system is stable then these projections map onto the graphs of the plant and controller respectively:

$$\Pi_{\mathcal{M}//\mathcal{N}}: \mathcal{W} \to \mathcal{G}_P \qquad\qquad \Pi_{\mathcal{N}//\mathcal{M}}: \mathcal{W} \to \mathcal{G}_C. \tag{4.12}$$

These parallel projection operators have some significant properties which we record here for completeness (from e.g. Doyle et al., 1992; Georgiou and Smith, 1997a, 2000)[1]:

1) $\Pi_{\mathcal{M}//\mathcal{N}} + \Pi_{\mathcal{N}//\mathcal{M}} = I$

2) $\Pi^2 = \Pi$

3) $H_{P,C} = (\Pi_{\mathcal{M}//\mathcal{N}}, \Pi_{\mathcal{N}//\mathcal{M}})$

4) $\Pi(\Pi w_1 + (I - \Pi)w_2) = \Pi w_1 \quad \forall w_1, w_2 \in \mathcal{W}$

5) $\Pi_{\mathcal{M}//\mathcal{N}}$ is the identity on $\mathcal{M}$ and respectively for $\Pi_{\mathcal{N}//\mathcal{M}}$ on $\mathcal{N}$ and therefore the norm of Π is greater than or equal to 1

6) Π induces a coordinatisation of $\mathcal{W}$: any $w \in \mathcal{W}$ has a unique additive decomposition $w = m + n$, where $m \in \Pi\mathcal{W}$ and $n \in (I - \Pi)\mathcal{W}$.

For the feedback system configuration in Figure 4.3 we can now use the operator $\Pi_{\mathcal{M}//\mathcal{N}}$ as the basis of the non-linear equivalent of the stability margin $b_{P,C}$ but in the time rather than frequency domain, defining he stability margin as $b_{P,C} = \|\Pi_{\mathcal{M}//\mathcal{N}}\|^{-1}$.

---

[1]Where Π is written here it can refer to either $\Pi_{\mathcal{M}//\mathcal{N}}$ or $\Pi_{\mathcal{N}//\mathcal{M}}$

To demonstrate the equivalence between time and frequency domain consider a linear plant $P\colon \mathcal{H}^2 \to \mathcal{H}^2$ and controller $C\colon \mathcal{H}^2 \to \mathcal{H}^2$ such that the closed-loop equations $[P, C]$ are stable. As per Figure 4.3 the signals $\hat{u}_0, \hat{y}_0, \hat{u}_1, \hat{y}_1 \in \mathcal{H}^2$. The map $\Pi_{\mathcal{M}//\mathcal{N}}$ can therefore be defined as follows:

$$\Pi_{\mathcal{M}//\mathcal{N}} = \mathscr{L}^{-1} \left(\begin{pmatrix} I \\ P \end{pmatrix} (I - PC)^{-1}(I, -C) \right) \mathscr{L}, \tag{4.13}$$

with $u_0, y_0, u_1, y_1 \in L^2(\mathbb{R}+)$ as the inverse Laplace transforms of the signals $\hat{u}_0, \hat{y}_0, \hat{u}_1, \hat{y}_1$. So for signals in $L^2(\mathbb{R}+)$ and $\mathcal{H}^2$ we have:

$$
\begin{array}{ccc}
\begin{pmatrix} I \\ P \end{pmatrix} (I - PC)^{-1}(I, -C) & : & \begin{pmatrix} \hat{u}_0 \\ \hat{y}_0 \end{pmatrix} \mapsto \begin{pmatrix} \hat{u}_1 \\ \hat{y}_1 \end{pmatrix} \\
\updownarrow & \mathscr{L}^{-1} \updownarrow \mathscr{L} & \\
\Pi_{\mathcal{M}//\mathcal{N}} & : & \begin{pmatrix} u_0 \\ y_0 \end{pmatrix} \mapsto \begin{pmatrix} u_1 \\ y_1 \end{pmatrix}
\end{array} \tag{4.14}
$$

The stability margins in both time and frequency domain are equal (see Section 2.1.5 concerning the equivalence of the $\mathcal{H}^\infty$- and L^2- norms):

$$b_{P,C} = \left\| \begin{pmatrix} I \\ P \end{pmatrix} (I - PC)^{-1}(I, -C) \right\|_\infty^{-1} = \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{L^2(\mathbb{R}+)}^{-1}. \tag{4.15}$$

These equivalences between the time and frequency domain signals apply only in $L^2(\mathbb{R}+)$ and $\mathcal{H}^2$, however this thesis considers the more general case of L^p where such parallels to the frequency domain are unavailable (see Section 2.1.5).

4.2.2 The non-linear gap

If we now let $\mathcal{M}_1$ denote the graph of the perturbed plant

$$\mathcal{M}_1 = \mathcal{G}_{P_1} = \left\{ w \in \mathcal{U} \times \mathcal{Y} \;\middle|\; w = \begin{pmatrix} u_1 \\ P_1 u_1 \end{pmatrix} \right\} \tag{4.16}$$

and recall from equation 4.1 that $\mathcal{M}$ denotes the graph of the nominal plant we can generalise the linear gap measure to the non-linear case.

In order to achieve this Georgiou and Smith (1997a) introduces a non-linear operator Φ that maps the graph of one plant onto the graph of the other. The distance of the

smallest such map from the identity is the directed gap between the two plants:

$$\vec{\delta}(\mathcal{M}, \mathcal{M}_1) = \begin{cases} \inf\{\|\Phi - I|_{\mathcal{M}}\| \colon \Phi \text{ is a causal,} \\ \qquad \text{surjective map from } \mathcal{M} \text{ to } \mathcal{M}_1 \\ \qquad \text{with } \Phi 0 = 0\}, \\ \infty \text{ if no such operator } \Phi \text{ exists} \end{cases} . \qquad (4.17)$$

The requirement for Φ to be surjective arises in Theorem 4.6, which analyses the stability of $[P_1, C]$ in terms of the signals in $\mathcal{G}_P$ mapped using Φ. Since all bounded, closed-loop signals compatible with $[P_1, C]$ must be written in terms of $[P, C]$, Φ must map onto all the possible signals. The condition that $\Phi 0 = 0$ follows naturally from the definition of induced norm however is included here explicitly as it will become pertinent later in the chapter and in Chapter 5 when a bias is included in the gap definition.

As well as the equivalence of the linear and non-linear stability margins, for linear systems in $\mathcal{RH}^\infty$, provided $\vec{\delta}_0(P, P_1) < 1$ the two measures of distance are also equal, $\vec{\delta}_0(P, P_1) = \vec{\delta}(P, P_1)$ (Georgiou and Smith, 1997a, Appendix).

Before presenting the main non-linear gap metric theorem a few definitions are required.

Definition 4.4. A mapping $Q \colon \mathcal{X} \to \mathcal{Y}$ is said to be **causal** if and only if

$$\forall x \in \mathcal{X} \quad \forall \tau \in \mathrm{dom}(x) \cap \mathrm{dom}(Qx) \quad : \quad T_\tau Q x = T_\tau Q T_\tau x. \qquad (4.18)$$

Definition 4.5. A causal plant $P \colon \mathcal{U}_e \to \mathcal{Y}_e$ is **stabilisable** if for all $T > 0$ and for all $(u, y) \in \mathcal{U}_e \times \mathcal{Y}_e$ satisfying $y = Pu$ there exists $(\tilde{u}, \tilde{y}) \in \mathcal{U} \times \mathcal{Y}$ such that $\tilde{y} = P\tilde{u}$ and $(\tilde{u}, \tilde{y})|_{[0,T]} = (u, y)|_{[0,T]}$.

The definition of stabilisability will be related to the stronger condition of behavioural controllability in Proposition 4.16.

Theorem 4.6. (From Theorem 1 of Georgiou and Smith, 1997a.) *Let* $\mathcal{U}$ *and* $\mathcal{Y}$ *be signal spaces and* $\mathcal{W} = \mathcal{U} \times \mathcal{Y}$. *Consider* $P_1 \colon \mathcal{U}_e \to \mathcal{Y}_e$, $P \colon \mathcal{U}_e \to \mathcal{Y}_e$ *and* $C \colon \mathcal{Y}_e \to \mathcal{U}_e$ *with* $P(0) = 0$ *and* $C(0) = 0$. *Suppose* $[P, C]$ *is gain stable on* $\mathcal{W}$, $[P_1, C]$ *is globally well-posed and* P_1 *is stabilisable. If*

$$\vec{\delta}(\mathcal{M}, \mathcal{M}_1) < \|\Pi_{\mathcal{M}//\mathcal{N}}\|^{-1} := b_{P,C} \qquad (4.19)$$

then the gain stability of $[P_1, C]$ *is assured on* $\mathcal{W}$ *and*

$$\|\Pi_{\mathcal{M}_1//\mathcal{N}}\| \le \|\Pi_{\mathcal{M}//\mathcal{N}}\| \frac{1 + \vec{\delta}(\mathcal{M}, \mathcal{M}_1)}{1 - \|\Pi_{\mathcal{M}//\mathcal{N}}\|\vec{\delta}(\mathcal{M}, \mathcal{M}_1)}. \qquad (4.20)$$

Note that this is only a sufficient condition for stability, the theorem says nothing about the stability of plants outside of this ball.

4.2.3 Other non-linear generalisations of the gap

Other than the non-linear gap of Georgiou and Smith (1997a), there are several other non-linear generalisations of the previous gap metric work.

Anderson et al. (1998), Sontag (1989) and Verma and Hunt (1993) discuss extensions of coprime factor robustness to cover non-linear systems. Paice and van der Schaft (1996) uses a Kernel representation to study the stability of plant-controller pairs and compares the results to those of coprime factorisation. The work on non-linear coprime factor robustness was then related to the non-linear gap in Bian and French (2005, 2003) and James et al. (2005).

Vinnicombe (1999) provides an approximate non-linear version of the ν-gap designed to collapse down to the standard ν-gap when applied to linear systems. This gap is also considered for non-linear operators in Anderson et al. (2002) where it is pointed out that the tightness of the linear ν-gap over the linear gap may not remain for the non-linear ν-gap.

Georgiou and Smith (1997a) and Bian and French (2005) also discuss the use of gain functions for robust stability analysis. For non-linear systems the existence of a gain function is a weaker requirement than the existence of a gain by an induced norm and so generalises the result. This thesis concentrates on the non-linear gap metric given in Georgiou and Smith (1997a) and so will not consider these other approaches to robust stability, although they are mentioned in Chapter 7 as further avenues to be explored for the iterative learning case.

4.3 Biased Norm

Georgiou and Smith (1997b) extends Georgiou and Smith (1997a) to the case where systems include some form of bias. Consider a closed-loop system which is stable in bounded-input-bounded-output terms (Definition 2.3) but where the map $\Pi_{\mathcal{M}//\mathcal{N}}$ does not map the zero trajectory to zero. Such a system may be perfectly acceptable but the induced norm of $\Pi_{\mathcal{M}//\mathcal{N}}$ would be infinite and so the system would not be gain stable (Definition 2.4). The previous robustness theorem would not give any robust stability guarantees as the robust stability margin defined by equation 4.19 would be zero.

For systems such as this it is possible to include a bias into the induced norm definition. This enables a measurement of the gain of an operator relative to a chosen signal. A biased norm on the map $A\colon \mathcal{X}_1 \to \mathcal{X}_2$, with $\mathcal{X}_1, \mathcal{X}_2 \subset \mathcal{W}$ and the bias $x_0 \in \mathcal{X}_1$, is now defined as

$$\|A\|_{\mathcal{W}, x_0} := \sup_{\substack{x_1 \in \mathrm{dom}(A),\ \tau > 0 \\ \|T_\tau(x_1 - x_0)\|_{\mathcal{W}} \neq 0}} \frac{\|T_\tau(Ax_1 - Ax_0)\|_{\mathcal{W}}}{\|T_\tau(x_1 - x_0)\|_{\mathcal{W}}}, \tag{4.21}$$

where $\| \cdot \|_{\mathcal{W}}$ is the norm on the vector space $\mathcal{W}$. Note: this biased norm relaxes to a non-biased norm if both $x_0 = 0$ and $A0 = 0$.

Consider the example above of a system with the property that $\Pi_{\mathcal{M}//\mathcal{N}}$ does not map the zero trajectory to zero, but where we are still interested in the system's stability relative to $x_0 = 0$. In this case the biased norm $\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}$ could be finite since the denominator in equation 4.21 approaching zero does not preclude the numerator approaching zero. Under the standard induced norm this would clearly not be possible.

The first four properties in the following proposition demonstrate that this biased norm obeys the axioms required for it to be labelled a norm. The final property shows that the norm is submultiplicative.

Proposition 4.7. *For a signal space $\mathcal{W}$, let $\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}_3 \subset \mathcal{W}$. Then for all causal operators $A, B\colon \mathcal{X}_1 \to \mathcal{X}_2$, $C\colon \mathcal{X}_1 \to \mathcal{X}_3$ and $D\colon \mathcal{X}_3 \to \mathcal{X}_2$ and for $x_0 \in \mathcal{W}$,*

  i) $\|A\|_{\mathcal{W},x_0} \geq 0$

  ii) $\|A\|_{\mathcal{W},x_0} = 0 \quad \Leftrightarrow \quad Ax = Ax_0 \quad \forall x$

  iii) $\|\lambda A\|_{\mathcal{W},x_0} = |\lambda|\|A\|_{\mathcal{W},x_0} \quad \forall \lambda \in \mathbb{R}$

  iv) $\|A + B\|_{\mathcal{W},x_0} \leq \|A\|_{\mathcal{W},x_0} + \|B\|_{\mathcal{W},x_0}$

  v) $\|DC\|_{\mathcal{W},x_0} \leq \|D\|_{\mathcal{W},Cx_0} \cdot \|C\|_{\mathcal{W},x_0}$

Proof. Only the proof of property v) is given here since the proofs for the remaining properties are trivial. By definition

$$\|DC\|_{\mathcal{W},x_0} := \sup_{\substack{x_1 \in \mathrm{dom}(C),\ \tau>0 \\ \|T_\tau(x_1-x_0)\|_{\mathcal{W}} \neq 0}} \frac{\|T_\tau(DCx_1 - DCx_0)\|_{\mathcal{W}}}{\|T_\tau(x_1 - x_0)\|_{\mathcal{W}}}. \tag{4.22}$$

In splitting up the operators C and D it is required to add an extra constraint into the supremum: that $\|T_\tau(Cx_1 - Cx_0)\|_{\mathcal{W}} \neq 0$. In order to do this it must therefore be shown that the inequality still holds if this condition is *not* met. This condition will therefore be added in the following way.

$$\|DC\|_{\mathcal{W},x_0} \leq \max \left\{ \sup_{\substack{x_1 \in \mathrm{dom}(C),\ \tau>0 \\ \|T_\tau(x_1-x_0)\|_{\mathcal{W}} \neq 0 \\ \|T_\tau(Cx_1-Cx_0)\|_{\mathcal{W}} \neq 0}} \frac{\|T_\tau(DCx_1 - DCx_0)\|_{\mathcal{W}}}{\|T_\tau(x_1 - x_0)\|_{\mathcal{W}}}, \right.$$

$$\left. \sup_{\substack{x_1 \in \mathrm{dom}(C),\ \tau>0 \\ \|T_\tau(x_1-x_0)\|_{\mathcal{W}} \neq 0 \\ \|T_\tau(Cx_1-Cx_0)\|_{\mathcal{W}} = 0}} \frac{\|T_\tau(DCx_1 - DCx_0)\|_{\mathcal{W}}}{\|T_\tau(x_1 - x_0)\|_{\mathcal{W}}} \right\}. \tag{4.23}$$

Suppose $\|T_\tau(Cx_1 - Cx_0)\|_\mathcal{W} = 0$, then it follows that $T_\tau Cx_1 = T_\tau Cx_0$. Consequently, by the causality of D,

$$
\begin{aligned}
\|T_\tau(DCx_1 - DCx_0)\|_\mathcal{W} &= \|T_\tau(DT_\tau Cx_1 - DT_\tau Cx_0)\|_\mathcal{W} \\
&= \|T_\tau(DT_\tau Cx_0 - DT_\tau Cx_0)\|_\mathcal{W} \\
&= 0.
\end{aligned}
\tag{4.24}
$$

Hence

$$
\sup_{\substack{x_1 \in \mathrm{dom}(C),\ \tau > 0 \\ \|T_\tau(x_1-x_0)\|_\mathcal{W} \neq 0 \\ \|T_\tau(Cx_1-Cx_0)\|_\mathcal{W} = 0}} \frac{\|T_\tau(DCx_1 - DCx_0)\|_\mathcal{W}}{\|T_\tau(x_1 - x_0)\|_\mathcal{W}} = 0 \leq \|D\|_{\mathcal{W},Cx_0} \cdot \|C\|_{\mathcal{W},x_0}.
\tag{4.25}
$$

On the other hand, if $\|T_\tau(Cx_1 - Cx_0)\|_\mathcal{W} \neq 0$, then

$$
\begin{aligned}
&\sup_{\substack{x_1 \in \mathrm{dom}(C),\ \tau > 0 \\ \|T_\tau(x_1-x_0)\|_\mathcal{W} \neq 0 \\ \|T_\tau(Cx_1-Cx_0)\|_\mathcal{W} \neq 0}} \frac{\|T_\tau(DCx_1 - DCx_0)\|_\mathcal{W}}{\|T_\tau(x_1 - x_0)\|_\mathcal{W}} \\[2mm]
&= \sup_{\substack{x_1 \in \mathrm{dom}(C),\ \tau > 0 \\ \|T_\tau(x_1-x_0)\|_\mathcal{W} \neq 0 \\ \|T_\tau(Cx_1-Cx_0)\|_\mathcal{W} \neq 0}} \left[\frac{\|T_\tau(DCx_1 - DCx_0)\|_\mathcal{W}}{\|T_\tau(Cx_1 - Cx_0)\|_\mathcal{W}} \cdot \frac{\|T_\tau(Cx_1 - Cx_0)\|_\mathcal{W}}{\|T_\tau(x_1 - x_0)\|_\mathcal{W}} \right] \\[2mm]
&\leq \left[\sup_{\substack{x_1 \in \mathrm{dom}(C),\ \tau > 0 \\ \|T_\tau(Cx_1-Cx_0)\|_\mathcal{W} \neq 0}} \frac{\|T_\tau(DCx_1 - DCx_0)\|_\mathcal{W}}{\|T_\tau(Cx_1 - Cx_0)\|_\mathcal{W}} \right] \left[\sup_{\substack{x_1 \in \mathrm{dom}(C),\ \tau > 0 \\ \|T_\tau(x_1-x_0)\|_\mathcal{W} \neq 0}} \frac{\|T_\tau(Cx_1 - Cx_0)\|_\mathcal{W}}{\|T_\tau(x_1 - x_0)\|_\mathcal{W}} \right] \\[2mm]
&\leq \left[\sup_{\substack{x_2 \in \mathrm{dom}(D),\ \tau > 0 \\ \|T_\tau(x_2-Cx_0)\|_\mathcal{W} \neq 0}} \frac{\|T_\tau(Dx_2 - DCx_0)\|_\mathcal{W}}{\|T_\tau(x_2 - Cx_0)\|_\mathcal{W}} \right] \left[\sup_{\substack{x_1 \in \mathrm{dom}(C),\ \tau > 0 \\ \|T_\tau(x_1-x_0)\|_\mathcal{W} \neq 0}} \frac{\|T_\tau(Cx_1 - Cx_0)\|_\mathcal{W}}{\|T_\tau(x_1 - x_0)\|_\mathcal{W}} \right] \\[2mm]
&= \|D\|_{\mathcal{W},Cx_0} \cdot \|C\|_{\mathcal{W},x_0}.
\end{aligned}
\tag{4.26}
$$

Hence by equations 4.23, 4.25 and 4.26 it follows that $\|DC\|_{\mathcal{W},x_0} \leq \|D\|_{\mathcal{W},Cx_0} \cdot \|C\|_{\mathcal{W},x_0}$. $\qquad\square$

The theorem in Georgiou and Smith (1997b) was published without proof and so one is given here.

Theorem 4.8. (From Theorem 1 of Georgiou and Smith, 1997b.) *Let $\mathcal{U}$ and $\mathcal{Y}$ be signal spaces and $\mathcal{W} = \mathcal{U} \times \mathcal{Y}$. Suppose $P \colon \mathcal{U}_e \to \mathcal{Y}_e$, $P_1 \colon \mathcal{U}_e \to \mathcal{Y}_e$ and $C \colon \mathcal{Y}_e \to \mathcal{U}_e$ are causal. Suppose that P_1 is stabilisable and $[P_1, C]$ is globally well-posed.*

Let $\mathcal{M} := \mathcal{G}_P$ and $\mathcal{M}_1 := \mathcal{G}_{P_1}$ and suppose that $\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0} < \infty$ for some $x_0 \in \mathcal{W}$.

Suppose the plant P_1 is such that there exists a causal surjective map $\Phi \colon \mathcal{M} \to \mathcal{M}_1$, which satisfies $\|(\Phi - I)|_\mathcal{M}\|_{\mathcal{W},x_1} < \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}^{-1}$, where $x_1 = \Pi_{\mathcal{M}//\mathcal{N}} x_0$.

Define $g_0 = (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$.

Then

$$\|\Pi_{\mathcal{M}_1//\mathcal{N}}\|_{\mathcal{W},g_0} \leq \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0} \frac{1 + \|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1}}{1 - \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}\|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1}}. \tag{4.27}$$

It should be noted that the above theorem does not explicitly use the gap $\vec{\delta}(P, P_1)$. However, with the definition

$$\vec{\delta}(P, P_1) = \left\{ \inf_{\Phi} \|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1} \;\middle|\; \Phi \colon \mathcal{M} \to \mathcal{M}_1 \text{ is causal and surjective} \right\}, \tag{4.28}$$

then if there does exists a Φ such that $\|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1} < \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}^{-1}$ it is clear that the gap $\vec{\delta}(P, P_1) \leq \|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1}$. The reason that Φ is utilised directly in the theorem statement as opposed to the gap is due to it being necessary to define the perturbed bias $g_0 = (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$.

Before the proof of Theorem 4.8 is given it is necessary to include the following lemma.

Lemma 4.9. *Let $\mathcal{X} \subseteq \mathcal{W}$ where $\mathcal{W}$ is a signal space and suppose $x_0 \in \mathcal{W}$. Suppose the operator $A \colon \mathcal{X} \to \mathcal{W}$ is causal and $\alpha := \|A\|_{\mathcal{W},x_0} < 1$.*

Let $\tau > 0$ and suppose $g, x, g_0 \in \mathcal{W}$ satisfy $T_\tau g = T_\tau(I + A)x$ and $T_\tau g_0 = T_\tau(I + A)x_0$.

Then $\|T_\tau(x - x_0)\|_{\mathcal{W}} \leq \frac{\|T_\tau(g - g_0)\|_{\mathcal{W}}}{1 - \alpha}$.

Proof. Since $T_\tau g = T_\tau(I + A)x$ and $T_\tau g_0 = T_\tau(I + A)x_0$ it follows that

$$T_\tau g - T_\tau g_0 = T_\tau(I + A)x - T_\tau(I + A)x_0, \tag{4.29}$$

and hence,

$$(T_\tau x - T_\tau x_0) - (T_\tau g - T_\tau g_0) = T_\tau A x_0 - T_\tau A x. \tag{4.30}$$

By the triangle inequality it can be shown that $\|T_\tau(y - z)\|_{\mathcal{W}} \geq \|T_\tau y\|_{\mathcal{W}} - \|T_\tau z\|_{\mathcal{W}}$. Hence from equation 4.30,

$$\begin{aligned}
\|T_\tau(Ax_0 - Ax)\|_{\mathcal{W}} &= \|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} \\
&= (1 - \alpha)\|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} + \alpha\|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} \\
&\geq (1 - \alpha)\|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} + \alpha\|T_\tau(x - x_0)\|_{\mathcal{W}} - \alpha\|T_\tau(g - g_0)\|_{\mathcal{W}}.
\end{aligned} \tag{4.31}$$

Since $\alpha := \|A\|_{\mathcal{W},x_0}$ it follows that $\|T_\tau(Ax_0 - Ax)\|_{\mathcal{W}} - \alpha\|T_\tau(x - x_0)\|_{\mathcal{W}} \leq 0$. Therefore, by equation 4.31,

$$0 \geq (1 - \alpha)\|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} - \alpha\|T_\tau(g - g_0)\|_{\mathcal{W}} \tag{4.32}$$

and hence,

$$\frac{\alpha \|T_\tau(g - g_0)\|_{\mathcal{W}}}{1 - \alpha} \geq \|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}}. \tag{4.33}$$

It follows that

$$\begin{aligned}
\|T_\tau(x - x_0)\|_{\mathcal{W}} &= \|T_\tau((x - x_0) - (g - g_0) + (g - g_0))\|_{\mathcal{W}} \\
&\leq \|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} + \|T_\tau((g - g_0))\|_{\mathcal{W}} \\
&\leq \frac{\alpha \|T_\tau(g - g_0)\|_{\mathcal{W}}}{1 - \alpha} + \|T_\tau(g - g_0)\|_{\mathcal{W}} \tag{4.34}
\end{aligned}$$

and therefore

$$\|T_\tau(x - x_0)\|_{\mathcal{W}} \leq \frac{1}{1 - \alpha} \|T_\tau(g - g_0)\|_{\mathcal{W}} \tag{4.35}$$

as required. $\qquad\qquad\square$

Proof of Theorem 4.8. By definition we have $x_1 = \Pi_{\mathcal{M}//\mathcal{N}} x_0$. Hence, from Proposition 4.7 property v), we have

$$\alpha := \|(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0} \leq \|(\Phi - I)\|_{\mathcal{W}, x_1} \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0} < 1. \tag{4.36}$$

Let $\tau > 0$. We first show that for any $g \in \mathcal{W}$ there exists x such that

$$\begin{aligned}
T_\tau g &= T_\tau (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x \\
&= T_\tau(\Pi_{\mathcal{N}//\mathcal{M}} + \Phi\Pi_{\mathcal{M}//\mathcal{N}})x. \tag{4.37}
\end{aligned}$$

Since $[P_1, C]$ is globally well-posed there exist $g_1, g_2 \in \mathcal{W}_e$ such that the system equations hold on $[0, \tau]$:

$$\begin{aligned}
y_1 &= P_1 u_1 & g_1 &= (u_1, y_1) \\
u_2 &= C y_2 & g_2 &= (u_2, y_2) \\
T_\tau g &= T_\tau g_1 + T_\tau g_2. & & \tag{4.38}
\end{aligned}$$

Since P_1 is stabilisable we can find $g_1'' \in \mathcal{M}_1$ such that $T_\tau g_1'' = T_\tau g_1$. By definition of $\mathcal{W}_e$ we can also find $g_2' \in \mathcal{W}$ such that $T_\tau g_2' = T_\tau g_2$.

As $\Phi \colon \mathcal{M} \to \mathcal{M}_1$ is surjective there exists $g_1' \in \mathcal{M}$ such that $\Phi g_1' = g_1''$ and therefore,

$$T_\tau \Phi g_1' = T_\tau g_1'' = T_\tau g_1. \tag{4.39}$$

We can now see that $x = g_1' + g_2'$ is a solution of equation 4.37. Since this holds for any $g \in \mathcal{W}$, then $T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}} g) = T_\tau(\Phi\Pi_{\mathcal{M}//\mathcal{N}} x)$.

In order to show that $T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}} g_0) = T_\tau(\Phi \Pi_{\mathcal{M}//\mathcal{N}} x_0)$ recall that by definition

$$
\begin{aligned}
T_\tau g_0 &= T_\tau(I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0 \\
&= T_\tau(I - \Pi_{\mathcal{M}//\mathcal{N}})x_0 + T_\tau \Phi \Pi_{\mathcal{M}//\mathcal{N}} x_0 \\
&= T_\tau \Pi_{\mathcal{N}//\mathcal{M}} x_0 + T_\tau \Phi \Pi_{\mathcal{M}//\mathcal{N}} x_0.
\end{aligned}
\tag{4.40}
$$

As $[P_1, C]$ and $[P, C]$ are globally well-posed, if m is an element in the graph of the plant and n an element in the graph of the controller then $\Pi_{\mathcal{N}//\mathcal{M}}(m + n) = \Pi_{\mathcal{N}//\mathcal{M}}(n) = n$. Therefore, as Φ maps its domain to elements in the graph of P_1 and $\Pi_{\mathcal{N}//\mathcal{M}}$ maps its domain to elements in the graph of C, we can state that.

$$
\begin{aligned}
T_\tau \Pi_{\mathcal{M}_1//\mathcal{N}} g_0 &= T_\tau \Pi_{\mathcal{M}_1//\mathcal{N}}(\Pi_{\mathcal{N}//\mathcal{M}} x_0 + T_\tau \Phi \Pi_{\mathcal{M}//\mathcal{N}} x_0) \\
&= T_\tau \Phi \Pi_{\mathcal{M}//\mathcal{N}} x_0
\end{aligned}
\tag{4.41}
$$

Combining both $T_\tau \Pi_{\mathcal{M}_1//\mathcal{N}} g = T_\tau \Phi \Pi_{\mathcal{M}//\mathcal{N}} x$ and $T_\tau \Pi_{\mathcal{M}_1//\mathcal{N}} g_0 = T_\tau \Phi \Pi_{\mathcal{M}//\mathcal{N}} x_0$ gives

$$
T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}} g - \Pi_{\mathcal{M}_1//\mathcal{N}} g_0) = T_\tau(\Phi \Pi_{\mathcal{M}//\mathcal{N}} x - \Phi \Pi_{\mathcal{M}//\mathcal{N}} x_0).
\tag{4.42}
$$

So, by causality and Lemma 4.9 with $A = (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}$:

$$
\begin{aligned}
\|T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}} g - \Pi_{\mathcal{M}_1//\mathcal{N}} g_0)\|_{\mathcal{W}} &= \|T_\tau(\Phi \Pi_{\mathcal{M}//\mathcal{N}} x - \Phi \Pi_{\mathcal{M}//\mathcal{N}} x_0)\|_{\mathcal{W}} \\
&\leq \|\Phi \Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0} \|T_\tau(x - x_0)\|_{\mathcal{W}} \\
&\leq \|\Phi \Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0} \frac{\|T_\tau(g - g_0)\|_{\mathcal{W}}}{1 - \alpha}.
\end{aligned}
\tag{4.43}
$$

So, since equation 4.43 holds for any $g \in \mathcal{W}$ and $\mathrm{dom}(\Pi_{\mathcal{M}_1//\mathcal{N}}) \subset \mathcal{W}$ then,

$$
\begin{aligned}
\|\Pi_{\mathcal{M}_1//\mathcal{N}}\|_{\mathcal{W}, g_0} &= \sup_{\substack{g \in \mathrm{dom}(\Pi_{\mathcal{M}_1//\mathcal{N}}),\ \tau > 0 \\ \|T_\tau(g - g_0)\|_{\mathcal{W}} \neq 0}} \frac{\|T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}} g - \Pi_{\mathcal{M}_1//\mathcal{N}} g_0)\|_{\mathcal{W}}}{\|T_\tau(g - g_0)\|_{\mathcal{W}}} \\
&\leq \frac{\|\Phi \Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0}}{1 - \alpha} \\
&\leq \frac{\|I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0}}{1 - \alpha} \\
&\leq \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0} \frac{1 + \|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W}, x_1}}{1 - \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0} \|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W}, x_1}}. \qquad \square
\end{aligned}
$$

4.4 A Biased Gap for Tracking

The previous section detailed the biased robust stability theorem from Georgiou and Smith (1997b). The original motivation for that paper was to extend earlier robust stability results to include systems that contained a bias. Stability, in this case, is relative to a bias signal. In this thesis the focus is on ILC and its robust stability whilst

engaged in tracking tasks. It is therefore wished to include the desired trajectory in the stability theorem as the trajectory to which the system is stable.

Previous results in this direction were published in Bradley and French (2009). The paper examined the robust stability of ILC systems using Theorem 4.8 with the signals lying in a 2D signal space, similar to one given in French (2008), with norm given by:

$$
\|u(\cdot,\cdot)\|_{l^p(\mathbb{N}\times[0,T])} = \begin{cases} \left(\sum\limits_{i=0}^{\infty} \sum\limits_{t=0}^{T} |u(i,t)|^p \right)^{\frac{1}{p}} & p < \infty \\ \sup\limits_{0\le i\le\infty} \sup\limits_{0\le t\le T} |u(i,t)| & p = \infty \end{cases}. \tag{4.44}
$$

This is equivalent to taking the $l^p[0,T]$-norm along each iteration, and applying the $l^p(\mathbb{N})$-norm to the resulting sequence of iteration 'sizes'. For a more in-depth explanation of the norm see Section 5.1.

A robust stability theorem for trajectory tracking ILC was given for signals in this 2D space under the biased norm, with the ILC's reference signal appearing in x_0. However, Theorem 4.8 includes a fundamental constraint that the reference trajectory must lie within the chosen signal space, $x_0 \in \mathcal{W}$, which causes problems with ILC.

This constraint is suitable for working within an l^∞-type signal space where the norm is based on the maximum amplitude of a signal, as any reference signal with bounded elements will be in the signal space. For other l^p spaces some difficulties arise. Consider a reference signal that is bounded in $l^p[0,T]$ but is non-zero. Once this signal is repeated over $\mathbb{N}$ it will be bounded in $l^p[0,T]$ along each iteration but will be infinite in the above norm as $i \to \infty$. The signal will therefore lie in $l^p_e(\mathbb{N} \times [0,T]) \setminus l^p(\mathbb{N} \times [0,T])$ and so is not an allowable x_0 (since Theorem 4.8 has constrained x_0 to lie in $l^p(\mathbb{N} \times [0,T])$). Since the very signals we wish to examine in this thesis are non-zero signals that repeat over an infinite number of iterations it is required to alter the theorem in order to include these.

In order to improve Theorem 4.8 the signal spaces will be adjusted to include the appropriate references and all domains and ranges of operators altered accordingly. The graphs of operators will therefore be defined relative to (unbounded) biases, and the norms of these operators will all be defined relative to these biases. This section will therefore: introduce a new set of notation in order to rigorously re-define the notation used previously relative to biases; prove the biased robust stability theorem under this notation; and show some of the properties of this new theorem when the analysis is restricted to linear systems.

The notation is split into three sections. The first covers the signals being examined, the second the norms and some new definitions of stability, and the third a series of definitions and propositions necessary for the main theorem.

4.4.1 Biased signal spaces

First it is needed to redefine the bias trajectories for the modelled and perturbed systems. For the feedback systems below these will be written as follows:
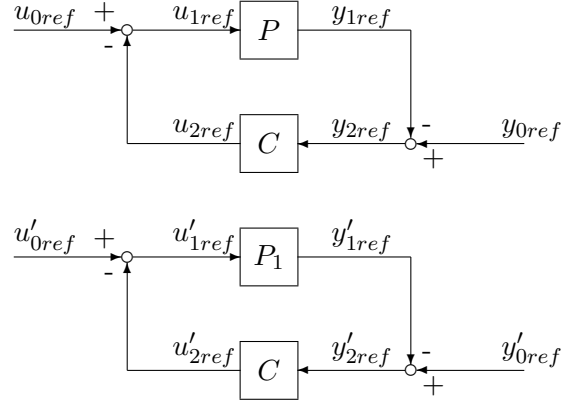
FIGURE 4.4: Feedback configuration denoting system biases

For the modelled plant

$$x_0 = \begin{pmatrix} u_{0ref} \\ y_{0ref} \end{pmatrix}$$

$$x_1 = \begin{pmatrix} u_{1ref} \\ y_{1ref} \end{pmatrix} = \Pi_{\mathcal{M}//\mathcal{N}} x_0$$

For the real plant

$$g_0 = \begin{pmatrix} u'_{0ref} \\ y'_{0ref} \end{pmatrix}$$

$$g_1 = \begin{pmatrix} u'_{1ref} \\ y'_{1ref} \end{pmatrix} = \Pi_{\mathcal{M}_1//\mathcal{N}} g_0$$

$$(4.45)$$

Within this thesis the signals x and g are generally used for the modelled and perturbed systems respectively. The subscript in these cases refers to where the biases lie — 0 lies around the closed-loop system, 1 around the plant and 2 the controller, in order to match the feedback notation used previously (see equation 2.12 and Figure 2.4).

In Chapter 5 this notation will be used in an ILC context. There we wish to examine the robust stability of an ILC system relative to a repeating reference trajectory y_{ref} defined on $l^p[0, T]$. This results in the 2D signal $y_{ref}(\cdot, \cdot)$, given by $y_{ref}(k, t) = y_{ref}(t)$ for all k. Note that unless $y_{ref}(\cdot) = 0$ or $p = \infty$ then $y_{ref}(\cdot, \cdot)$ will lie in $l_e^p(\mathbb{N} \times [0, T]) \backslash l^p(\mathbb{N} \times [0, T])$. This reference will be applied at y_{0ref} and therefore the external bias will be given by $x_0 = (0, y_{ref})$.

These signals are then used to define biased signal spaces. As before, define the signal

spaces $\mathcal{U}$ and $\mathcal{Y}$ and $\mathcal{W} = \mathcal{U} \times \mathcal{Y}$. Their biased counterparts are then given by

$$\mathcal{U}_{u_0} = \{\, u \in \mathcal{U}_e \mid u - u_0 \in \mathcal{U} \,\}$$
$$\mathcal{Y}_{y_0} = \{\, y \in \mathcal{Y}_e \mid y - y_0 \in \mathcal{Y}\}$$
$$\mathcal{W}_{\left(\begin{smallmatrix} u_0 \\ y_0 \end{smallmatrix}\right)} = \mathcal{U}_{u_0} \times \mathcal{Y}_{y_0}. \tag{4.46}$$

As we will generally regard signals in pairs the biased version of $\mathcal{W}$ given above will be denoted using the reference trajectory notation and so will become (when measured relative to the external biases) $\mathcal{W}_{x_0}$. Note that $\mathcal{U}_{u_0} \subset \mathcal{U}_e$, $\mathcal{Y}_{y_0} \subset \mathcal{Y}_e$ and $\mathcal{W}_{x_0} \subset \mathcal{W}_e$.

The graphs of the plants are now defined as subsets (affine sets in the linear case) of these biased signal spaces with the bias included. For a plant $P : \mathcal{U}_e \to \mathcal{Y}_e$ and controller $C : \mathcal{Y}_e \to \mathcal{U}_e$ the graphs are now defined by

$$\mathcal{M} = \mathcal{G}_{P,\left(\begin{smallmatrix} u_1 \\ y_1 \end{smallmatrix}\right)} = \left\{ \begin{pmatrix} u \\ Pu \end{pmatrix} \in \mathcal{U}_e \times \mathcal{Y}_e \;\middle|\; u \in \mathcal{U}_{u_1}, \; Pu \in \mathcal{Y}_{y_1} \right\} \tag{4.47}$$

$$\mathcal{N} = \mathcal{G}_{C,\left(\begin{smallmatrix} u_2 \\ y_2 \end{smallmatrix}\right)} = \left\{ \begin{pmatrix} Cy \\ y \end{pmatrix} \in \mathcal{U}_e \times \mathcal{Y}_e \;\middle|\; Cy \in \mathcal{U}_{u_2}, \; y \in \mathcal{Y}_{y_2} \right\}. \tag{4.48}$$

The operator $\Pi_{\mathcal{M}//\mathcal{N}}$ is now given by:

$$\Pi_{\mathcal{M}//\mathcal{N}} : \mathcal{W}_{x_0} \to \mathcal{W}_e. \tag{4.49}$$

Note that if $[P,C]$ is stable with respect to x_0 (see Definition 4.11 below) then

$$\Pi_{\mathcal{M}//\mathcal{N}} : \mathcal{W}_{x_0} \to \mathcal{G}_{P,x_1} \subset \mathcal{W}_{x_1}. \tag{4.50}$$

Recall the biased norm given in equation 4.21, which we use to measure the gain of operators with respect to biases in extended signal spaces. Given a signal space $\mathcal{W}$, a signal $z_0 \in \mathcal{W}_e$ and $A : \mathcal{W}_{z_0} \to \mathcal{W}_{Az_0}$ with $\mathcal{W}_{z_0}, \mathcal{W}_{Az_0} \subset \mathcal{W}_e$ then

$$\|A\|_{\mathcal{W},z_0} = \sup_{\substack{z \in \mathrm{dom}(A), \; \tau > 0 \\ \|T_\tau(z-z_0)\|_{\mathcal{W}} \neq 0}} \frac{\|T_\tau(Az - Az_0)\|_{\mathcal{W}}}{\|T_\tau(z - z_0)\|_{\mathcal{W}}}. \tag{4.51}$$

4.4.2 A series of definitions and propositions

The previous definitions of globally well-posedness, stability and stabilisability need to be altered in order to support the gap theorem in its biased norm state and so further proofs of properties and some new definitions are given here.

The first proposition extends the notion of well-posedness (given in Definition 2.2) to also include external signals (u_0, y_0) within the extended signal space. This is necessary

as we will now be considering signals that lie in extended spaces.

Proposition 4.10. *If a system is globally well-posed, then for any $(u_0, y_0) \in \mathcal{W}_e$ there exists a unique set of signals $((u_1, y_1), (u_2, y_2)) \in \mathcal{W}_e \times \mathcal{W}_e$ such that the set of equations $[P, C]$ hold on $[0, \tau]$ for any $\tau > 0$.*

Proof. First we prove the existence of the solution. Let $\tau > 0$. For any $(u_0, y_0) \in \mathcal{W}_e$ there exists $(u_0^\tau, y_0^\tau) \in \mathcal{W}$ such that $T_\tau(u_0^\tau, y_0^\tau) = T_\tau(u_0, y_0)$. As the set of equations $[P, C]$ is globally well-posed by Definition 2.2 then there exists a unique set of signals $((u_1^\tau, y_1^\tau), (u_2^\tau, y_2^\tau)) \in \mathcal{W}_e \times \mathcal{W}_e$ such that the set of equations $[P, C]$, with input (u_0^τ, y_0^τ), hold on $[0, \infty]$, and therefore $[P, C]$ also holds on $[0, \tau]$ for $((u_0, y_0), (u_1^\tau, y_1^\tau), (u_2^\tau, y_2^\tau))$.

It is also necessary to demonstrate that this same set of signals is generated when choosing a different truncation. We claim that if $\tau_1 > \tau_2$ and two sets of signals are formed using the above then

$$T_\tau \left[(u_1^{\tau_1}, y_1^{\tau_1}), (u_2^{\tau_1}, y_2^{\tau_1}) \right] = T_\tau \left[(u_1^{\tau_2}, y_1^{\tau_2}), (u_2^{\tau_2}, y_2^{\tau_2}) \right], \tag{4.52}$$

and therefore this *procedure* defines a unique set of signals $((u_1, y_1), (u_2, y_2)) \in \mathcal{W}_e \times \mathcal{W}_e$ as required. This claim is verified by the global well-posedness of $[P, C]$, the causality of P and C and the definition that

$$
\begin{aligned}
T_{\tau_2}(u_0^{\tau_1}, y_0^{\tau_1}) &= T_{\tau_2} T_{\tau_1}(u_0, y_0) \\
&= T_{\tau_2}(u_0, y_0) \\
&= T_{\tau_2}(u_0^{\tau_2}, y_0^{\tau_2}).
\end{aligned}
\tag{4.53}
$$

Lastly we show that the set of signals given by the above procedure is the only possible solution. Suppose there exist two sets of signals $w = ((u_1, y_1), (u_2, y_2)) \in \mathcal{W}_e \times \mathcal{W}_e$ and $\bar{w} = ((\bar{u}_1, \bar{y}_1), (\bar{u}_2, \bar{y}_2)) \in \mathcal{W}_e \times \mathcal{W}_e$ such that both w and $\bar{w}$ are compatible with (u_0, y_0) on the equations $[P, C]$ and $w \neq \bar{w}$ for all $t > 0$.

As $w \neq \bar{w}$ there exists $s > 0$ such that $T_s w \neq T_s \bar{w}$. Since both w and $\bar{w}$ are compatible with (u_0, y_0) on $[P, C]$ over $[0, s]$ then they must also be compatible with a some pair of signals $(\tilde{u}_0, \tilde{y}_0) \in \mathcal{W}$ such that $T_s(\tilde{u}_0, \tilde{y}_0) = T_s(u_0, y_0)$. However, global well-posedness and causality now imply that $T_s w = T_s \bar{w}$. This contradiction precludes the existence of a second solution. $\square$

Analogous to Definition 2.3 regarding closed-loop stability:

Definition 4.11. For a globally well-posed system given by the set of equations $[P, C]$, **stability with respect to a signal** x_0 is attained when for all $(u_0, y_0) \in \mathcal{W}_{x_0}$ then the internal signals $((u_1, y_1), (u_2, y_2)) \in \mathcal{W}_{x_1} \times \mathcal{W}_{x_2}$ (no signals lie outside of the biased signal spaces).

Analogous to Definition 2.4:

Definition 4.12. A system given by $[P, C]$ is **gain stable with respect to a signal** x_0 when the biased induced norms of $\Pi_{\mathcal{M}//\mathcal{N}} \colon \mathcal{W}_{x_0} \to \mathcal{W}_e$ and $\Pi_{\mathcal{N}//\mathcal{M}} \colon \mathcal{W}_{x_0} \to \mathcal{W}_e$ are finite:

$$\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0} < \infty$$
$$\|\Pi_{\mathcal{N}//\mathcal{M}}\|_{\mathcal{W}, x_0} < \infty \tag{4.54}$$

The following proposition uses the truncation operator to show that any signal lying in an extended signal space can be considered equal to a signal in the biased signal space up to any time τ.

Proposition 4.13. *For any $\tau > 0$ and any $z_0 \in \mathcal{W}_e$, then for all $x \in \mathcal{W}_e$ there exists $y \in \mathcal{W}_{z_0}$ such that $T_\tau y = T_\tau x$.*

Proof. Let $x \in \mathcal{W}_e$ and $\tau > 0$. Let $y = T_\tau x + (I - T_\tau)z_0$. Then $T_\tau y = T_\tau x$ and $y - z_0 = T_\tau x - T_\tau z_0 \in \mathcal{W}$ and so $y \in \mathcal{W}_{z_0}$. $\qquad\square$

The previous definition of stabilisability (Definition 4.5) is insufficient for the following theorem and so a new notion is required to relate stabilisability to the biased reference trajectory given.

Definition 4.14. A system is **stabilisable with respect to a given trajectory** $x_0 \in \mathcal{W}_e$ if for all $\tau > 0$ and for all $(u, y) \in \mathcal{U}_e \times \mathcal{Y}_e$ satisfying $y = Pu$ there exists $(\tilde{u}, \tilde{y}) \in \mathcal{W}_{x_0}$ such that $\tilde{y} = P\tilde{u}$ and $(\tilde{u}, \tilde{y})|_{[0,\tau]} = (u, y)|_{[0,\tau]}$.

The following definition of controllability is a stronger condition than either stabilisability or stabilisability with respect to a given trajectory. Since controllability is a more standard definition Proposition 4.16 is included to convey its relative strength and to show that controllability is a sufficient condition for stabilisability. However, since the stabilisability conditions are significantly weaker they will be included in the final theorem.

In order to fit with the material in this thesis we will here use a behavioural definition of controllability that can be found in Polderman and Willems (1998). This is the only point in this thesis where any behavioural notation is used and so only a sufficient amount of detail is given to enable controllability to be adequately defined. The 'behaviour' of a time-invariant dynamical system is denoted $\mathfrak{B}$, a set that contains all possible trajectories that are compatible with the system. Therefore, from the point of view of this thesis, the only detail required is that the graph and biased graph are both subsets of the behaviour — $\mathcal{G}_P, \mathcal{G}_{P, x_1} \subset \mathfrak{B}$ — although the behaviour may also include the signals compatible with the plant that are unbounded.

Definition 4.15. (From Definition 5.2.2 of Polderman and Willems, 1998.) Let $\mathfrak{B}$ be the behaviour of a time-invariant dynamical system. This system is called **controllable** if for any two trajectories $w_1, w_2 \in \mathfrak{B}$ there exists a $t_1 \geq 0$ and a trajectory $w \in \mathfrak{B}$ with the property

$$w(t) = \begin{cases} w_1(t) & t \leq 0 \\ w_2(t) & t \geq t_1 \end{cases}. \tag{4.55}$$

Note that any set of signals in the behaviour can be shifted arbitrarily in time and so the time 0 in the above definition can be placed anywhere within the domain of elements of the behaviour.

Proposition 4.16. *If a system is controllable then it is also stabilisable with respect to a given trajectory.*

Proof. For any $T > 0$, any $x_0 \in \mathcal{W}_e$ compatible with the plant, and all $(u, y) \in \mathcal{U}_e \times \mathcal{Y}_e$ satisfying $y = Pu$ it is necessary to find a pair of signals $(\tilde{u}, \tilde{y}) \in \mathcal{W}_{x_0}$ such that $\tilde{y} = P\tilde{u}$ and $(\tilde{u}, \tilde{y})|_{[0,T]} = (u, y)$.

Let $w = (u, y)$. By controllability there exists z and Δ such that

$$\tilde{w}(t) = \begin{cases} w(t) & t \in [0, T] \\ z(t) & t \in [T, T + \Delta] \\ x_0(t) & t \geq T + \Delta \end{cases} \tag{4.56}$$

with $\tilde{w} = (\tilde{u}, \tilde{y})$ satisfying $\tilde{y} = P\tilde{u}$. It is then clear that $\tilde{w} \in \mathcal{W}_{x_0}$ and $\tilde{w}|_{[0,T]} = w|_{[0,T]}$. $\square$

4.4.3 Biased graph robust stability theorem

The previous robust stability theorem (Theorem 4.8) will now be reestablished using the biased graph definition and the new notation. This theorem can therefore provide robust stability results that are applicable to systems with reference trajectories that lie in extended signal spaces. As before the theorem is not restricted to signals in 1D and so can be applied in an ILC setting in Chapter 5.

Theorem 4.17. *Let $\mathcal{U}$ and $\mathcal{Y}$ be signal spaces and $\mathcal{W} = \mathcal{U} \times \mathcal{Y}$. Suppose $P \colon \mathcal{U}_e \to \mathcal{Y}_e$, $P_1 \colon \mathcal{U}_e \to \mathcal{Y}_e$ and $C \colon \mathcal{Y}_e \to \mathcal{U}_e$ are causal and $[P_1, C]$ is globally well-posed.*

Let $\mathcal{M} = \mathcal{G}_{P,x_1}$ and $\mathcal{N} = \mathcal{G}_{C,x_2}$ where $x_1 = \Pi_{\mathcal{M}//\mathcal{N}} x_0$ and $x_2 = \Pi_{\mathcal{N}//\mathcal{M}} x_0$ for some $x_0 \in \mathcal{W}_e$. Suppose that $\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0} < \infty$.

Let $g_1 = \Phi \Pi_{\mathcal{M}//\mathcal{N}} x_0$ and suppose the plant P_1 is stabilisable with respect to g_1. Let $\mathcal{M}_1 := \mathcal{G}_{P_1,g_1}$ and suppose there exists a causal surjective map $\Phi \colon \mathcal{M} \to \mathcal{M}_1$ satisfying $\|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1} < \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}^{-1}$. Define $g_0 = (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$.

Then

$$\|\Pi_{\mathcal{M}_1//\mathcal{N}}\|_{\mathcal{W},g_0} \leq \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0} \frac{1 + \|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1}}{1 - \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}\|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1}}. \tag{4.57}$$

As before it should be noted that the above theorem does not explicitly use the gap $\vec{\delta}(P, P_1)$. However, since the gap is given by

$$\vec{\delta}(P, P_1) = \left\{ \inf_{\Phi} \|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1} \, \Big| \, \Phi \colon \mathcal{M} \to \mathcal{M}_1 \text{ is causal and surjective} \right\}, \tag{4.58}$$

if there does exists a Φ such that $\|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1} < \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}^{-1}$ then it is clear that the gap $\vec{\delta}(P, P_1) \leq \|(\Phi - I)|_{\mathcal{M}}\|_{\mathcal{W},x_1}$. The reason that Φ is included as opposed to the gap is due to it being necessary to define $g_0 = (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$.

The proof first requires the following two lemmata. Since the theorem is more involved than Theorem 4.8 some of the detail now appears in Lemma 4.18. Lemma 4.19 is similar to Lemma 4.9 but is now able to work with signals in extended spaces.

Lemma 4.18. *Let $\mathcal{U}$ and $\mathcal{Y}$ be signal spaces and $\mathcal{W} = \mathcal{U} \times \mathcal{Y}$. Consider $P \colon \mathcal{U}_e \to \mathcal{Y}_e$, $P_1 \colon \mathcal{U}_e \to \mathcal{Y}_e$ and $C \colon \mathcal{Y}_e \to \mathcal{U}_e$ to be causal and $[P_1, C]$ to be globally well-posed.*

Let $\mathcal{M} = \mathcal{G}_{P,x_1}$ and $\mathcal{N} = \mathcal{G}_{C,x_2}$ where $x_1 = \Pi_{\mathcal{M}//\mathcal{N}}x_0$ and $x_2 = \Pi_{\mathcal{N}//\mathcal{M}}x_0$ for some $x_0 \in \mathcal{W}_e$.

Let $g_1 = \Phi\Pi_{\mathcal{M}//\mathcal{N}}x_0$ and suppose the plant P_1 is stabilisable with respect to g_1. Let $\mathcal{M}_1 := \mathcal{G}_{P_1,g_1}$ and suppose there exists a causal surjective map $\Phi \colon \mathcal{M} \to \mathcal{M}_1$. Define $g_0 = (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$.

Then for any $g \in \mathcal{W}_{g_0}$ there exists $x \in \mathcal{W}_{x_0}$ such that for all $\tau > 0$:

$$\begin{aligned} T_\tau g &= T_\tau(I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x \\ &= T_\tau(\Pi_{\mathcal{N}//\mathcal{M}} + \Phi\Pi_{\mathcal{M}//\mathcal{N}})x, \end{aligned} \tag{4.59}$$

and

$$\begin{aligned} T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}}g) &= T_\tau(\Phi\Pi_{\mathcal{M}//\mathcal{N}}x) \\ T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}}g_0) &= T_\tau(\Phi\Pi_{\mathcal{M}//\mathcal{N}}x_0). \end{aligned} \tag{4.60}$$

Proof. Since $[P_1, C]$ is globally well-posed and by Proposition 4.10 it follows that for all $w_0 \in \mathcal{W}_{g_0}$ there exists $w_1, w_2 \in \mathcal{W}_e$ such that the system equations hold on $[0, \tau]$:

$$\begin{aligned} y_1 &= P_1 u_1 & w_1 &= (u_1, y_1) \\ u_2 &= C y_2 & w_2 &= (u_2, y_2) \\ T_\tau w_0 &= T_\tau w_1 + T_\tau w_2. \end{aligned} \tag{4.61}$$

Since P_1 is stabilisable with respect to g_1 we can find $w_1'' \in \mathcal{M}_1$ such that $T_\tau w_1'' = T_\tau w_1$. By Proposition 4.13 we can also find $w_2' \in \mathcal{W}_{x_2}$ such that $T_\tau w_2' = T_\tau w_2$.

As $\Phi \colon \mathcal{M} \to \mathcal{M}_1$ is surjective there exists $w_1' \in \mathcal{M}$ such that $\Phi w_1' = w_1''$ and therefore,

$$T_\tau \Phi w_1' = T_\tau w_1'' = T_\tau w_1. \tag{4.62}$$

We can now see that $x = w_1' + w_2'$ is a solution of equation 4.59. As this holds for any $w_0 \in \mathcal{W}_{g_0}$ we can substitute $w_0 = g$ and state that $T_\tau(\Pi_{\mathcal{M}_1 /\!/ \mathcal{N}} g) = T_\tau(\Phi \Pi_{\mathcal{M} /\!/ \mathcal{N}} x)$, and $T_\tau(\Pi_{\mathcal{M}_1 /\!/ \mathcal{N}} g_0) = T_\tau(\Phi \Pi_{\mathcal{M} /\!/ \mathcal{N}} x_0)$ from the definition of g_0. □

Lemma 4.19. *Let $\mathcal{W}$ be a signal space and suppose $x_0, g_0 \in \mathcal{W}_e$. Suppose the operator $A \colon \mathcal{W}_{x_0} \to \mathcal{W}_{x_1} \cup \mathcal{W}_{g_1}$ is causal with $T_\tau g_0 = T_\tau(I + A)x_0$ and $\alpha := \|A\|_{\mathcal{W}, x_0} < 1$.*

Suppose $g \in \mathcal{W}_{g_0}$, $x \in \mathcal{W}_{x_0}$ satisfy $T_\tau g = T_\tau(I + A)x$.

Then $\|T_\tau(x - x_0)\|_{\mathcal{W}} \leq \frac{\|T_\tau(g - g_0)\|_{\mathcal{W}}}{1 - \alpha}$.

Proof. Since $T_\tau g = T_\tau(I + A)x$ and $T_\tau g_0 = T_\tau(I + A)x_0$ it follows that

$$T_\tau g - T_\tau g_0 = T_\tau(I + A)x - T_\tau(I + A)x_0, \tag{4.63}$$

and hence,

$$(T_\tau x - T_\tau x_0) - (T_\tau g - T_\tau g_0) = T_\tau A x_0 - T_\tau A x. \tag{4.64}$$

By the triangle inequality it follows that $\|T_\tau(a - b)\|_{\mathcal{W}} \geq \|T_\tau a\|_{\mathcal{W}} - \|T_\tau b\|_{\mathcal{W}}$. Hence from equation 4.64,

$$
\begin{aligned}
\|T_\tau(Ax_0 - Ax)\|_{\mathcal{W}} &= \|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} \\
&= (1 - \alpha)\|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} + \alpha\|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} \\
&\geq (1 - \alpha)\|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} + \alpha\|T_\tau(x - x_0)\|_{\mathcal{W}} - \alpha\|T_\tau(g - g_0)\|_{\mathcal{W}}.
\end{aligned}
\tag{4.65}
$$

Now, $\|T_\tau(Ax_0 - Ax)\|_{\mathcal{W}} - \alpha\|T_\tau(x - x_0)\|_{\mathcal{W}} \leq 0$ by definition of $\alpha := \|A\|_{\mathcal{W}, x_0}$. Therefore

$$0 \geq (1 - \alpha)\|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}} - \alpha\|T_\tau(g - g_0)\|_{\mathcal{W}} \tag{4.66}$$

and hence,

$$\frac{\alpha\|T_\tau(g - g_0)\|_{\mathcal{W}}}{1 - \alpha} \geq \|T_\tau((x - x_0) - (g - g_0))\|_{\mathcal{W}}. \tag{4.67}$$

It follows that

$$
\begin{aligned}
\|T_\tau(x - x_0)\|_\mathcal{W} &= \|T_\tau((x - x_0) - (g - g_0) + (g - g_0))\|_\mathcal{W} \\
&\leq \|T_\tau((x - x_0) - (g - g_0))\|_\mathcal{W} + \|T_\tau((g - g_0))\|_\mathcal{W} \\
&\leq \frac{\alpha\|T_\tau(g - g_0)\|_\mathcal{W}}{1 - \alpha} + \|T_\tau(g - g_0)\|_\mathcal{W}
\end{aligned}
\tag{4.68}
$$

and therefore

$$
\|T_\tau(x - x_0)\|_\mathcal{W} \leq \frac{1}{1 - \alpha}\|T_\tau(g - g_0)\|_\mathcal{W}
\tag{4.69}
$$

as required. $\qquad\square$

We now give the proof of the main result:

Proof of Theorem 4.17. From Proposition 4.7, property v) and from the assumption that $\|(\Phi - I)|_{\mathcal{G}_{P,x_1}}\|_{\mathcal{W},x_1} < \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}^{-1}$ let

$$
\alpha := \|(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_o} \leq \|(\Phi - I)|_{\mathcal{G}_{P,x_1}}\|_{\mathcal{W},x_1}\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_o} < 1.
\tag{4.70}
$$

From Lemma 4.18 we can state that $T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}}g) = T_\tau(\Phi\Pi_{\mathcal{M}//\mathcal{N}}x)$ and also that $T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}}g_0) = T_\tau(\Phi\Pi_{\mathcal{M}//\mathcal{N}}x_0)$. Combining the two we arrive at

$$
T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}}g - \Pi_{\mathcal{M}_1//\mathcal{N}}g_0) = T_\tau(\Phi\Pi_{\mathcal{M}//\mathcal{N}}x - \Phi\Pi_{\mathcal{M}//\mathcal{N}}x_0).
\tag{4.71}
$$

So, by causality and Lemma 4.19 with $A = (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}$:

$$
\begin{aligned}
\|T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}}g - \Pi_{\mathcal{M}_1//\mathcal{N}}g_0)\|_\mathcal{W} &= \|T_\tau(\Phi\Pi_{\mathcal{M}//\mathcal{N}}x - \Phi\Pi_{\mathcal{M}//\mathcal{N}}x_0)\|_\mathcal{W} \\
&\leq \|\Phi\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}\|T_\tau(x - x_0)\|_\mathcal{W} \\
&\leq \|\Phi\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}\frac{\|T_\tau(g - g_0)\|_\mathcal{W}}{1 - \alpha}.
\end{aligned}
\tag{4.72}
$$

So, by equation 4.72:

$$
\begin{aligned}
\|\Pi_{\mathcal{M}_1//\mathcal{N}}\|_{\mathcal{W},g_0} &= \sup_{\substack{g\in\mathrm{dom}(\Pi_{\mathcal{M}_1//\mathcal{N}}),\ \tau>0 \\ \|T_\tau(g-g_0)\|_\mathcal{W}\neq 0}} \frac{\|T_\tau(\Pi_{\mathcal{M}_1//\mathcal{N}}g - \Pi_{\mathcal{M}_1//\mathcal{N}}g_0)\|_\mathcal{W}}{\|T_\tau(g - g_0)\|_\mathcal{W}} \\
&\leq \frac{\|\Phi\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}}{1 - \alpha} \\
&\leq \frac{\|I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}}{1 - \alpha} \\
&\leq \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}\frac{1 + \|(\Phi - I)|_{\mathcal{G}_{P,x_1}}\|_{\mathcal{W},x_1}}{1 - \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}\|(\Phi - I)|_{\mathcal{G}_{P,x_1}}\|_{\mathcal{W},x_1}}. \qquad\square
\end{aligned}
$$

4.4.4 Where do g_0 and g_1 lie?

Theorem 4.17 only shows that the gain of the system P_1 is bounded relative to a reference signal $g_0 = (I+(\Phi-I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$. It therefore becomes necessary to bound this reference signal to show that this calculation is worthwhile. The following proposition provides this bound for the general case and also a neater bound applicable when the set of plants is restricted to those where $(\Phi-I)\Pi_{\mathcal{M}//\mathcal{N}}0 = 0$. (In order to formally match the domain of the operators this condition will be written as $T_\tau(\Phi-I)\Pi_{\mathcal{M}//\mathcal{N}}(I-T_\tau)x_0 = 0$ for any $\tau > 0$.) This condition is satisfied by linear plants due to the property that $P0 = C0 = 0$. The theorem also provides an approximation for the case where this condition is not met.[2]

Proposition 4.20. *Let $\tau > 0$, $x_0 \in \mathcal{W}$ and $g_0 = (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$. Define $(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}$ to be causal where $\epsilon = \|(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}$. Then the magnitude of g_0 is bounded by*

$$\|T_\tau g_0\|_{\mathcal{W}} \le (1+\epsilon)\|T_\tau x_0\|_{\mathcal{W}} + \|T_\tau(\Phi-I)\Pi_{\mathcal{M}//\mathcal{N}}(I-T_\tau)x_0\|_{\mathcal{W}}. \tag{4.73}$$

With the added condition that $T_\tau(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}(I - T_\tau)x_0 = 0$ this bound reduces to

$$\|T_\tau g_0\|_{\mathcal{W}} \le (1+\epsilon)\|T_\tau x_0\|_{\mathcal{W}}. \tag{4.74}$$

Note that under the conditions of Theorem 4.17 $\epsilon = \|(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0} < 1$.

Proof. Let $x \in \mathcal{W}_{x_0}$ and $\tau > 0$ such that $T_\tau(x-x_0) \ne 0$. Since $g_0 = x_0+(\Phi-I)\Pi_{\mathcal{M}//\mathcal{N}}x_0$, we have

$$x_0 - g_0 + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x = (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x - (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_0 \tag{4.75}$$

and therefore

$$\frac{\|T_\tau(x_0 - g_0 + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x)\|_{\mathcal{W}}}{\|T_\tau(x - x_0)\|_{\mathcal{W}}} = \frac{\|T_\tau((\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x - (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_0)\|_{\mathcal{W}}}{\|T_\tau(x - x_0)\|_{\mathcal{W}}}$$

$$\le \epsilon, \tag{4.76}$$

since $\epsilon = \|(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}$.

Now let $x_\tau = (I - T_\tau)x_0$. Therefore $x_\tau \in \mathcal{W}_{x_0}$ and $\|T_\tau x_\tau\|_{\mathcal{W}} = 0$ (see Proposition 4.13).

---

[2]The condition here that $T_\tau(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}(I - T_\tau)x_0 = 0$ for any $\tau > 0$ can be easily constructed from the condition that $(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}0 = 0$. Since if $x_0 \in \mathcal{W}_e \setminus \mathcal{W}$ then the signal $0 \notin \text{dom}(\Pi_{\mathcal{M}//\mathcal{N}})$ we have instead chosen the signal $(I - T_\tau)x_0$, which is equal to 0 when truncated at τ. As both $\Pi_{\mathcal{M}//\mathcal{N}}$ and Φ are causal both conditions are therefore the same under truncation.

We can now write

$$\|T_\tau(x_0 - g_0 + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_\tau)\|_\mathcal{W} \leq \epsilon\|T_\tau(x_\tau - x_0)\|_\mathcal{W}$$
$$\leq \epsilon\|T_\tau x_0\|_\mathcal{W} + \epsilon\|T_\tau x_\tau\|_\mathcal{W}$$
$$\leq \epsilon\|T_\tau x_0\|_\mathcal{W}. \tag{4.77}$$

For the alternative case where $x_\tau \in \mathcal{W}_{x_0}$, $\tau > 0$ and $T_\tau(x_\tau - x_0) = 0$, from equation 4.75 and by the causality of $(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}$ we can state that:

$$\|T_\tau(x_0 - g_0 + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_\tau)\|_\mathcal{W} = 0$$
$$\leq \epsilon\|T_\tau x_0\|_\mathcal{W}, \tag{4.78}$$

and therefore for any $x_\tau = (I - T_\tau)x_0$, from equations 4.77 and 4.78 we have the bound:

$$\epsilon\|T_\tau x_0\|_\mathcal{W} \geq \|T_\tau(x_0 - g_0 + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_\tau)\|_\mathcal{W}$$
$$\geq \|T_\tau g_0\|_\mathcal{W} - \|T_\tau(x_0 + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_\tau)\|_\mathcal{W}, \tag{4.79}$$

and therefore

$$\|T_\tau g_0\|_\mathcal{W} \leq (1 + \epsilon)\|T_\tau x_0\|_\mathcal{W} + \|T_\tau(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_\tau\|_\mathcal{W}. \tag{4.80}$$

If we restrict the set of systems to those that fulfil $T_\tau(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_\tau = 0$ then this collapses to

$$\|T_\tau g_0\|_\mathcal{W} \leq (1 + \epsilon)\|T_\tau x_0\|_\mathcal{W} \qquad \square$$

This bound exhibits some unfortunate side effects when applied to any signal space not based on an l^∞-type norm, due to the inclusion of the truncation operator. Under l^∞ the truncation can be removed as the reference signals lie within the un-extended signal space. Clearly this results in the maximum magnitude of g_0 being less than twice the maximum magnitude of x_0.

Under other signal spaces the norm may not be as well behaved. Consider a constant x_0 and the l^1-norm. As τ increases the norm $\|T_\tau x_0\|_\mathcal{W}$ increases in a linear fashion. The bound on $\|T_\tau g_0\|_\mathcal{W}$ therefore increases. However, this does not imply that g_0 is a constant. The signal g_0 could remain small for $t < \tau - \delta$ and then 'blow up' in the range $[\tau - \delta, \tau]$ whilst still remaining within the bound. Figure 4.5 demonstrates the potential problem. The graph shows how $\|T_\tau g_0\|_\mathcal{W}$ increases as τ increases. The blue area is the region above the upper bound; the blue line shows the value that we would naturally expect $\|T_\tau g_0\|_\mathcal{W}$ to take given this bound; and the red line shows a possible value that $\|T_\tau g_0\|_\mathcal{W}$ could actually take. It is clear from the graph that (in the 'possible' case) g_0 has a sizeable magnitude for a short time but the low magnitude beforehand ensures the norm is below the given bound. This problem could be solved by placing additional

restrictions such as a bound on the differential of g_0.
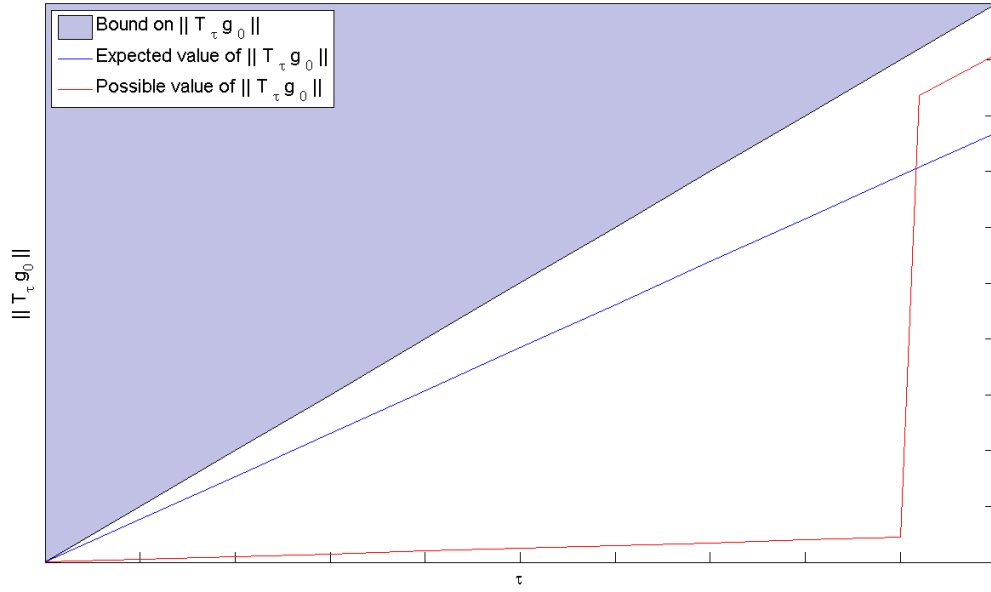
FIGURE 4.5: Demonstration of the bound on g_0

Recall that the bias around the perturbed plant is denoted g_1 and is defined by the equation $g_1 = \Phi\Pi_{\mathcal{M}//\mathcal{N}}x_0 = \Pi_{\mathcal{M}_1//\mathcal{N}}g_0$. This is therefore the signal that a stable $[P_1, C]$ will converge to, given the input g_0, and so it is of primary importance that this signal is appropriate. As all of the operators in Theorem 4.17 are measured using biased norms this is non-trivial, since the signal g_0 may be bounded appropriately but g_1 may not. As an example: the biased norm of a map $A\colon \mathcal{W}_{w_0} \to \mathcal{W}_{w_1}$ may equal zero, however this simply means that $Aw = w_1$ for all $w \in \mathcal{W}_{w_0}$; this does not mean that w_1 is bounded or convergent. It is therefore necessary to develop a bound on g_1 similar to the above bound on g_0. This time we will restrict to the case where $T_\tau(\Phi-I)\Pi_{\mathcal{M}//\mathcal{N}}(I-T_\tau)x_0 = 0$ from the outset (although a similar proof could be developed for the case where this is not met).

Proposition 4.21. *Let* $\tau > 0$, $x_0 \in \mathcal{W}$, $T_\tau(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}(I - T_\tau)x_0 = 0$ *and define* $\epsilon = \|(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}$. *Let* $g_1 = \Phi\Pi_{\mathcal{M}//\mathcal{N}}x_0$ *and let the maps* Φ *and* $\Pi_{\mathcal{M}//\mathcal{N}}$ *be causal. Then the magnitude of* g_1 *is bounded by*

$$\|T_\tau g_1\|_{\mathcal{W}} \leq \epsilon\|T_\tau x_0\|_{\mathcal{W}} + \|T_\tau\Pi_{\mathcal{M}//\mathcal{N}}x_0\|_{\mathcal{W}}. \tag{4.81}$$

Proof. Start with the definition:

$$\begin{aligned} g_1 &= \Phi\Pi_{\mathcal{M}//\mathcal{N}}x_0 \\ &= (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_0 - \Pi_{\mathcal{M}//\mathcal{N}}x_0. \end{aligned} \tag{4.82}$$

Suppose $T_\tau x_0 \neq 0$. Observe that taking $x_\tau = (I - T_\tau)x_0$ gives $T_\tau(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_\tau = 0$ by the causality of Φ and $\Pi_{\mathcal{M}//\mathcal{N}}$, and $T_\tau(x_\tau - x_0) = T_\tau((I - T_\tau)x_0 - x_0) = 0$. Therefore:

$$
\begin{aligned}
\epsilon &= \|(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0} \\
&= \sup_{\substack{x \in \mathcal{W}_{x_0},\ \tau' > 0 \\ \|T_{\tau'}(x-x_0)\|_{\mathcal{W}} \neq 0}} \frac{\|T_{\tau'}((\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x - (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_0)\|_{\mathcal{W}}}{\|T_{\tau'}(x - x_0)\|_{\mathcal{W}}} \\
&\geq \frac{\|T_\tau((\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_\tau - (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_0)\|_{\mathcal{W}}}{\|T_\tau(x_\tau - x_0)\|_{\mathcal{W}}} \\
&= \frac{\|T_\tau((\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_0)\|_{\mathcal{W}}}{\|T_\tau x_0\|_{\mathcal{W}}}, \tag{4.83}
\end{aligned}
$$

and so $\|T_\tau((\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_0)\|_{\mathcal{W}} \leq \epsilon\|T_\tau x_0\|_{\mathcal{W}}$. If $T_\tau x_0 = 0$ then this inequality holds trivially. Therefore:

$$
\begin{aligned}
\|T_\tau g_1\|_{\mathcal{W}} &= \|T_\tau((\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_0 - \Pi_{\mathcal{M}//\mathcal{N}}x_0)\|_{\mathcal{W}} \\
&\leq \|T_\tau((\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}x_0\|_{\mathcal{W}} + \|T_\tau \Pi_{\mathcal{M}//\mathcal{N}}x_0\|_{\mathcal{W}} \\
&\leq \epsilon\|T_\tau x_0\|_{\mathcal{W}} + \|T_\tau \Pi_{\mathcal{M}//\mathcal{N}}x_0\|_{\mathcal{W}}. \qquad\qquad \square
\end{aligned}
$$

Note that the bound given above suffers from the same problem as that given for g_0 when dealing with signal spaces other than l^∞ due to the truncation.

Note also that the second term in the final line contains the operator $\Pi_{\mathcal{M}//\mathcal{N}}$. Recall that this is the map from external signals around the closed-loop system to internal signals around the plant for the nominal system and therefore is known. This operator should also be well behaved since the nominal system is stable and its performance has been deemed acceptable. The exact nature of $\|\Pi_{\mathcal{M}//\mathcal{N}}\|$ will depend on the system being implemented, the domain in which it lies and the norms used to measure it. In the following chapter we will examine this in the ILC case. Proposition 5.13 in Section 5.5 will provide a bound on $\|\Pi_{\mathcal{M}//\mathcal{N}}\|$ in 2D for a monotonically convergent ILC system, where the monotonicity is used to provide a bound on $\|\Pi_{\mathcal{M}//\mathcal{N}}\|$.

4.4.5 A note on linearity

For the case of linear systems, Theorem 4.17 collapses onto the unbiased gap theorem from Georgiou and Smith (1997a). The gap is then equivalent to the linear gap. This section will present a simple proof to show this.

Lemma 4.22. *Let P and P_1 be linear and let $x_1 \in \mathcal{W}_e$ and $g_1 = \Phi x_1$. Then for every biased map $\Phi\colon \mathcal{G}_{P,x_1} \to \mathcal{G}_{P_1,g_1}$ there exists a non-biased map $\tilde{\Phi}\colon \mathcal{G}_P \to \mathcal{G}_{P_1}$ such that*

$$
\|(\Phi - I)|_{\mathcal{G}_{P,x_1}}\|_{\mathcal{W},x_1} = \|(\tilde{\Phi} - I)|_{\mathcal{G}_P}\|_{\mathcal{W}}. \tag{4.84}
$$

Proof. Let $y \in \mathcal{G}_P$. Define $\tilde{\Phi} \colon \mathcal{G}_P \to \mathcal{G}_{P_1}$ such that $\tilde{\Phi}(y) = \Phi(y + x_1) + \Phi(x_1)$. Let $x \in \mathcal{G}_{P,x_1}$.

Since $x \in \mathcal{G}_{P,x_1} \subset \mathcal{W}_{x_1}$ it follows that $x - x_1 \in \mathcal{W}$. Since both x and x_1 are compatible with the plant and since the plant is linear it follows that $(x - x_1)$ is also compatible with the plant, hence $x - x_1 \in \mathcal{G}_P$. Then $\tilde{\Phi}(x - x_1) = \Phi(x) - \Phi(x_1)$ and

$$
\begin{aligned}
\|(\Phi - I)|_{\mathcal{G}_{P,x_1}}\|_{\mathcal{W},x_1} &= \sup_{\substack{x \in \mathcal{G}_{P,x_1}, \ \tau > 0 \\ \|T_\tau(x - x_1)\|_{\mathcal{W}} \neq 0}} \frac{\|T_\tau\left[(\Phi - I)x - (\Phi - I)x_1\right]\|_{\mathcal{W}}}{\|T_\tau(x - x_1)\|_{\mathcal{W}}} \\
&= \sup_{\substack{x \in \mathcal{G}_{P,x_1}, \ \tau > 0 \\ \|T_\tau(x - x_1)\|_{\mathcal{W}} \neq 0}} \frac{\|T_\tau\left[\Phi x - \Phi x_1 - x + x_1\right]\|_{\mathcal{W}}}{\|T_\tau(x - x_1)\|_{\mathcal{W}}} \\
&= \sup_{\substack{(x - x_1) \in \mathcal{G}_P, \ \tau > 0 \\ \|T_\tau(x - x_1)\|_{\mathcal{W}} \neq 0}} \frac{\|T_\tau[\tilde{\Phi}(x - x_1) - x + x_1]\|_{\mathcal{W}}}{\|T_\tau(x - x_1)\|_{\mathcal{W}}} \\
&= \|(\tilde{\Phi} - I)|_{\mathcal{G}_P}\|_{\mathcal{W}} \tag{4.85}
\end{aligned}
$$

$\square$

4.5 Summary

This chapter has extended the material of Chapter 2 on robust stability. The subject of the linear and non-linear gap metrics have been introduced and some of their applications explained. Sufficient conditions for the stability of a system in the face of plant perturbations have then been proven in terms of the gain of a system.

The theory of robust stability with biases, which was outlined in Georgiou and Smith (1997b), has been given along with a detailed proof. This has then be extended using a notion of biased signal spaces, operators and norms in order to permit biases that lie in extended signal spaces. This led to the establishment of a biased norm robust stability theorem that can be applied to a variety of problems but was primarily developed to deal with the robust stability of trajectory tracking with iterative learning control.

The following chapter will apply this theorem in an ILC context.

Chapter 5

A Robust Stability Framework for ILC

The previous chapter introduced the gap metric in a biased fashion. As explained at the start of that chapter, the subject of ILC was not directly addressed and the theorems provided are therefore able to deal with a multitude of situations. This chapter will now demonstrate the applicability of the biased graph robust stability theorem (Theorem 4.17) to iterative learning control.

The chapter will begin by detailing some of the research carried out on bringing ILC and the gap metric together in French (2008) and Bradley and French (2009). Work from the latter will then be extended to incorporate the biased graph robust stability theorem from the previous chapter; providing a 2D robust stability margin and robustness results for ILC. The 2D gap used within this theorem will then be examined and related to the standard 1D gap, and a final theorem given that enables an ILC system to be examined with the traditional linear gap metric. This allows the 2D ILC problem to be provided with robustness guarantees based on 1D gap measurements. During these discussions surrounding issues of linearity, causality and surjectivity will be studied in order to provide a clear picture of the results in terms of their applicability, relation to previous work and required constraints.

5.1 An Overview of Previous ILC Gap Metric Work

As explained in Sections 3.7 and 4.4, ILC and the theory of the gap metric are brought together in French (2008), where a gap metric approach is used to prove that there exists a non-zero stability margin for a class of adaptive ILC algorithms. The paper concentrates on the simpler setting of disturbance attenuation rather than the more general issue of non-zero trajectory tracking, which will be our concern in this chapter.

As explained in Section 3.7, in French (2008) an adaptive ILC controller is defined via a feedback gain that increases depending on the error measured on the previous iteration. Robust stability guarantees are then given for high-gain-stabilisable plants of relative degree one. Whilst the error is high the gain increases and so reduces the error (by the high gain property). However, as has been relayed throughout this thesis, a high gain leads to a low stability margin.

In order to measure the size of signals in an ILC setting a 2D norm is defined in French (2008). The norm effectively wraps the real line up into segments of length T, forming a product space from the $L^p[0, T]$ norm and the natural numbers, $\mathbb{N}$:

$$\|u(\cdot, \cdot)\|_{L^p(\mathbb{N} \times [0,T])} = \begin{cases} \left(\sum_{k=0}^{\infty} \int_0^T |u(k,t)|^p \, dt \right)^{\frac{1}{p}} & p < \infty \\ \sup_{0 \leq k \leq \infty} \sup_{0 \leq t \leq T} |u(k,t)| & p = \infty \end{cases}. \tag{5.1}$$

The $L^p(\mathbb{N} \times [0, T])$ signal space is then defined as the set of all signals that are finite under this norm.

For the causality definition to hold (Definition 4.4) the domain of the space must be totally ordered and so it is required to define this ordering. For any $\tau_1, \tau_2 \in \mathbb{N} \times [0, T]$, $(k_1, t_1) = \tau_1 < \tau_2 = (k_2, t_2)$ if $k_1 < k_2$, or if $k_1 = k_2$ and $t_1 < t_2$. With the ordering defined a truncation operator is available and so there exist definitions of causality and stabilisability. (Figure 5.1 provides a graphical explanation of this ordering: the red showing the past and black the future, with respect to a time τ.)
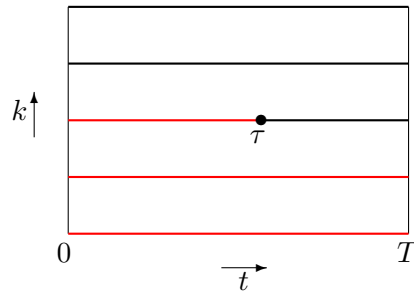
FIGURE 5.1: Signal space ordering

This provides a truncation definition as follows: where $\tau = (\tilde{k}, \tilde{t})$

$$T_\tau f(k,t) := \begin{cases} f(k,t) & k < \tilde{k} \\ f(k,t) & k = \tilde{k}, \ t \leq \tilde{t} \\ 0 & k = \tilde{k}, \ t > \tilde{t} \\ 0 & k > \tilde{k} \end{cases}. \tag{5.2}$$

From Definition 4.4, a map $Q \colon L^p(\mathbb{N} \times [0, T]) \to L^p(\mathbb{N} \times [0, T])$ is then said to be causal

if and only if for all $x \in L^p(\mathbb{N} \times [0,T])$ and all $\tau \in (\mathbb{N} \times [0,T])$ then

$$T_\tau Q x = T_\tau Q T_\tau x. \tag{5.3}$$

With causality defined in this manner it is possible to consider algorithms that are 'non-causal' in an ILC sense (see Section 3.4.5) since these algorithms use data from further forward in time along previous iterations.

This norm is used in Bradley and French (2009) to examine the 2D robust stability of ILC in a trajectory tracking setting. The paper provides a robust stability theorem based on the 2D norm given above and the biased robust stability theorem of Georgiou and Smith (1997b) (in this thesis it is shown as Theorem 4.8). The 2D gain of $\Pi_{\mathcal{M}//\mathcal{N}}$ is then calculated for ILC systems given in the lifted framework described in Section 3.3.1. The results show the effects on robustness of implementing model-based ILC and adding filtering.

In the previous chapter one of the problems with Theorem 4.8 was explained — that an ILC reference signal will not be permitted to lie within the signal space unless analysis is restricted to L^∞-type spaces. The work in this chapter will extend the work from Bradley and French (2009) to include these signals. This work will instead use Theorem 4.17 with a more general 2D norm to provide improved results.

5.2 A Robust Stability Margin for ILC

5.2.1 2D biased norm

The 2D norm of equation 5.1 is now generalised to allow different measurements along and up the trials. In this chapter we will be using two variants of norm: one continuous and one discrete. The continuous norm is defined as:

$$\|u(\cdot,\cdot)\|_{L^{p,q}(\mathbb{N} \times [0,T])} = \begin{cases} \left(\sum\limits_{k=0}^{\infty} \|u(k,\cdot)\|_{L^p}^q \right)^{\frac{1}{q}} & q < \infty \\ \sup\limits_{0 \le k \le \infty} \|u(k,\cdot)\|_{L^p}^q & q = \infty \end{cases}. \tag{5.4}$$

Here we have a continuous time L^p norm along each trial and the l^q norm of the results over all iterations. This generalisation allows, for example, the use of L^2 along the trial and l^∞ up the trials. Such a measurement would provide the maximum energy along all the trials. This is particularly useful when attempting to quantify specific properties such as along-the-trial performance and trial-to-trial convergence.

As with the other spaces, the signal space $L^{p,q}(\mathbb{N} \times [0,T])$ consists of all signals that are finite in this norm. The ordering of the space's domain is the same as that given

for the space with norm of equation 5.1 (and Figure 5.1): for any $\tau_1, \tau_2 \in \mathbb{N} \times [0, T]$, $(k_1, t_1) = \tau_1 < \tau_2 = (k_2, t_2)$ if $k_1 < k_2$, or if $k_1 = k_2$ and $t_1 < t_2$.

When examining ILC in a lifted system context we wish to use a discrete version of the above norm, detailed here:

$$\|u(\cdot, \cdot)\|_{l^{p,q}(\mathbb{N} \times [0,T])} = \begin{cases} \left(\sum_{k=0}^{\infty} \|u(k, \cdot)\|_{l^p[0,T]}^q \right)^{\frac{1}{q}} & q < \infty \\ \sup_{0 \le k \le \infty} \|u(k, \cdot)\|_{l^p[0,T]} & q = \infty \end{cases}. \qquad (5.5)$$

The norm is the same but with the 'along-the-trial' magnitude determined by a discrete time l^p norm. The domain ordering is formally the same. Note that $[0, T]$ now denotes the discrete interval[1]:

$$[0, T] = \{ s \in \mathbb{N} \mid 0 \le s \le T \}. \qquad (5.6)$$

The normed vector spaces associated with these norms will be defined in the standard form, consisting of all the functions for which the norm is finite. The extended spaces are also defined in the same way as previously. It should be noted that, although the internal norms in both cases *cannot* have extensions of their own since they are defined on a finite interval (see Section 2.1.3), the external norm sums to infinity and so the extended 2D space *does* exist.

This 2D space can be used with the biased signal space and biased norm definition from the previous chapter to provide a 2D biased norm. Take $\mathcal{W} = l^{p,q}(\mathbb{N} \times [0, T])$ and $x_0 = (u_0, y_0) \in \mathcal{W}_e$. Recall the biased signal space definition: for $\mathcal{W}_{x_0}$ this is given by

$$\begin{aligned} \mathcal{U}_{u_0} &= \{ u \in \mathcal{U}_e \mid u - u_0 \in \mathcal{U} \} \\ \mathcal{Y}_{y_0} &= \{ y \in \mathcal{Y}_e \mid y - y_0 \in \mathcal{Y} \} \\ \mathcal{W}_{x_0} &= \mathcal{U}_{u_0} \times \mathcal{Y}_{y_0}. \end{aligned} \qquad (5.7)$$

Also recall the biased norm definition: for $\mathcal{W}_{x_0}, \mathcal{W}_{Ax_0} \subset \mathcal{W}_e$ the norm of the map $A \colon \mathcal{W}_{x_0} \to \mathcal{W}_{Ax_0}$ is given by

$$\|A\|_{\mathcal{W}, x_0} := \sup_{\substack{x_1 \in \mathrm{dom}(A), \ \tau > 0 \\ \|T_\tau(x_1(\cdot, \cdot) - x_0(\cdot, \cdot))\|_{\mathcal{W}} \ne 0}} \frac{\|T_\tau(Ax_1(\cdot, \cdot) - Ax_0(\cdot, \cdot))\|_{\mathcal{W}}}{\|T_\tau(x_1(\cdot, \cdot) - x_0(\cdot, \cdot))\|_{\mathcal{W}}}. \qquad (5.8)$$

Theorem 4.17 was shown to hold for any signal space on which a truncation definition was defined and therefore the theorem holds for signals lying within appropriately biased versions of $l^{p,q}(\mathbb{N} \times [0, T])$. We now therefore have the robust stability theorem and a notion of signal sizes that we can apply to ILC.

---

[1]As previously mentioned, the set of natural numbers $\mathbb{N}$ here contains 0

5.2.2 ILC system

In order to implement the biased graph robust stability theorem (Theorem 4.17) it is required to define the plant and controller in terms of the 2D signal space given above. Start with a discrete-time plant $\tilde{P} \colon l_e^p(\mathbb{R}) \to l_e^p(\mathbb{R})$, of the form

$$
\begin{aligned}
x(t+1) &= Ax(t) + Bu_1(t) \\
y_1(t) &= Cx(t) + Du_1(t),
\end{aligned}
\tag{5.9}
$$

with $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times 1}$, $C \in \mathbb{R}^{1 \times n}$ and $D \in \mathbb{R}$ (using t as the discrete-time index).

Since we will only be examining signals on $[0, T]$ we will rewrite this plant as a map $l^p[0, T] \to l^p[0, T]$. Also, setting the initial conditions for the state variable $x(0) = 0$, the plant can be expanded as a matrix P mapping $u_1(\cdot)$ to $y_1(\cdot)$, matching the lifted system ILC given in Chapter 3.

$$
P =
\begin{pmatrix}
D & & & & \\
CB & D & & & \\
CAB & CB & D & & \\
CA^2B & CAB & CB & D & \\
\vdots & \vdots & \vdots & \vdots & \ddots
\end{pmatrix}
\tag{5.10}
$$

This matrix is lower-triangular and Toeplitz. Therefore the map $P \colon u_1(\cdot) \mapsto y_1(\cdot)$ is given by:

$$
\begin{pmatrix}
y_1(0) \\
y_1(1) \\
y_1(2) \\
y_1(3) \\
\vdots
\end{pmatrix}
= P
\begin{pmatrix}
u_1(0) \\
u_1(1) \\
u_1(2) \\
u_1(3) \\
\vdots
\end{pmatrix}.
\tag{5.11}
$$

This plant is now defined over $l^{p,q}(\mathbb{N} \times [0, T])$. This is the same plant but repeated over $\mathbb{N}$. Here $k \in \mathbb{N}$ represents the trial number.

$$
\begin{pmatrix}
y_1(k, 0) \\
y_1(k, 1) \\
y_1(k, 2) \\
\vdots \\
y_1(k, T)
\end{pmatrix}
=
\begin{pmatrix}
D & & & & \\
CB & D & & & \\
CAB & CB & D & & \\
\vdots & \vdots & \vdots & \ddots & \\
CA^{T-1}B & CA^{T-2}B & CA^{T-3}B & \cdots & D
\end{pmatrix}
\begin{pmatrix}
u_1(k, 0) \\
u_1(k, 1) \\
u_1(k, 2) \\
\vdots \\
u_1(k, T)
\end{pmatrix}
$$
$$
k \in \mathbb{N} \quad (5.12)
$$

For most of the analysis the signals will be examined as vectors along each trial and so the time index shall be dropped. Therefore $y_1(k)$ will denote the entire vector y_1 along

trial k:

$$y_1(k) = \begin{pmatrix} y_1(k,0) \\ y_1(k,1) \\ y_1(k,2) \\ \vdots \\ y_1(k,T) \end{pmatrix}. \tag{5.13}$$

Similarly, the other signals are written in the same way and, since $y_{ref}(k)$ is identical for all k, it will be written simply as y_{ref}. Using this notation the 2D plant is now denoted $\overline{P} : l^{p,q}(\mathbb{N} \times [0,T]) \to l^{p,q}(\mathbb{N} \times [0,T]) : u_1(\cdot) \mapsto y_1(\cdot)$ defined by $y_1(k) = Pu_1(k)$, $k \in \mathbb{N}$.

In order to fit the lifted system representation given in Section 3.3.1 with the system representation used in the gap work (given in Section 2.1.6) the structure in which the plant will be presented is given by Figure 5.2:
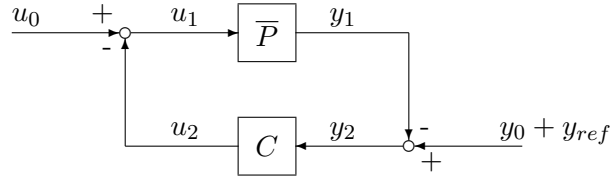
FIGURE 5.2: System diagram

The figure is shown with the reference trajectory y_{ref} inserted alongside the disturbance y_0. With this feedback structure we can match the lifted system update equation by implementing a controller of the form:

$$
\begin{aligned}
C \colon u_2(k+1, \cdot) &= Q(u_2(k, \cdot) - Ly_2(k, \cdot)) \\
u_2(0, \cdot) &= 0
\end{aligned}
\tag{5.14}
$$

This is the same ILC controller in Section 3.3.1 but with a sign change introduced to fit with the slightly different feedback structure. Here $L : l^p[0,T] \to l^p[0,T]$ is the learning gain and $Q : l^p[0,T] \to l^p[0,T]$ is usually a filter, as explained previously.

The plant $P_1 \colon l^p_e(\mathbb{R}) \to l^p_e(\mathbb{R})$ is written in a 2D form in the same way as the plant P. To summarise, the two plants and the controller are defined in 2D as:

$$
\begin{aligned}
\overline{P}(k,t) \colon l^{p,q}(\mathbb{N} \times [0,T]) &\to l^{p,q}_e(\mathbb{N} \times [0,T]) : u_1 \mapsto y_1, \ y_1(k, \cdot) = Pu_1(k, \cdot) \\
\overline{P}_1(k,t) \colon l^{p,q}(\mathbb{N} \times [0,T]) &\to l^{p,q}_e(\mathbb{N} \times [0,T]) : u_1 \mapsto y_1, \ y_1(k, \cdot) = P_1u_1(k, \cdot) \\
C(k,t) \colon l^{p,q}(\mathbb{N} \times [0,T]) &\to l^{p,q}_e(\mathbb{N} \times [0,T]) : y_2 \mapsto u_2, \\
u_2(k+1, \cdot) &= Q(u_2(k, \cdot) - Ly_2(k, \cdot)). \tag{5.15}
\end{aligned}
$$

5.2.3 A 2D biased robust stability margin

In this section the plants and controller given above will be examined with the biased robust stability theorem in 2D. The robust stability margin will be calculated for the feedback system $[\overline{P}, C]$, and then a 2D robust stability theorem will be stated, relating the stability margin of $[\overline{P}, C]$ to the stability of $[\overline{P}_1, C]$.

Recall the biased signal space definitions of Section 4.4.1:

$$
\begin{aligned}
\mathcal{U}_{u_0} &= \{\, u \in \mathcal{U}_e \mid u - u_0 \in \mathcal{U} \,\} \\
\mathcal{Y}_{y_0} &= \{\, y \in \mathcal{Y}_e \mid y - y_0 \in \mathcal{Y} \} \\
\mathcal{W}_{\left(\substack{u_0 \\ y_0}\right)} &= \mathcal{U}_{u_0} \times \mathcal{Y}_{y_0}.
\end{aligned}
\tag{5.16}
$$

In this case we wish the bias to take the form of the reference trajectory. The reference trajectory is inserted at y_0 so the bias is $(u_0, y_0) = (0, y_{ref})$.

This leaves the projection operator $\Pi_{\mathcal{M}//\mathcal{N}}$ defined as

$$
\Pi_{\mathcal{M}//\mathcal{N}} \colon \mathcal{W}_{\left(\substack{0 \\ y_{ref}}\right)} \to \mathcal{W}_e.
\tag{5.17}
$$

In Theorem 4.17 it is the biased induced gain of this projection that acts as the robust stability margin and so the following theorem will calculate this gain. The calculation is restricted to the case of $q \in \{1, \infty\}$ and $p \in [1, \infty]$.

Theorem 5.1. *Let $q \in \{1, \infty\}$, $p \in [1, \infty]$, $\mathcal{W} = l^{p,q}(\mathbb{N} \times [0, T])$ and $x_0 = \left(\substack{0 \\ y_{ref}}\right) \in \mathcal{W}_e$. Consider the closed-loop system given by Figure 5.2, and equations 5.11 and 5.14, with the matrix P given by equation 5.10. Suppose $\|Q(I - LP)\|_{l^p[0,T]} < 1$. Then with the biased norm of equation 5.8, the gain of $\Pi_{\mathcal{M}//\mathcal{N}}$ is bounded by*

$$
\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0} \le \|\left(\substack{I \\ P}\right)\|_{l^p[0,T]} \left(1 + \frac{\|QLP\|_{l^p[0,T]} + \|QL\|_{l^p[0,T]}}{1 - \|Q(I - LP)\|_{l^p[0,T]}}\right).
\tag{5.18}
$$

Note that the induced norms in $l^p[0, T]$ on the right hand side of equation 5.18 are always finite.

Proof. Let $q \in \{1, \infty\}$, $p \in [1, \infty]$, $\mathcal{W} = l^{p,q}(\mathbb{N} \times [0, T])$, $x_0 = \left(\substack{0 \\ y_{ref}}\right) \in \mathcal{W}_e$ and suppose $(u_0, y_0)^T \in l^{p,q}(\mathbb{N} \times [0, T])$. We can derive an expression for $u_2(k + 1)$ in terms of the disturbances and the previous control signal $u_2(k)$.

$$
\begin{aligned}
u_2(k + 1) &= Q(u_2(k) - L(y_0(k) + y_{ref} - P(u_0(k) - u_2(k)))) \\
&= Q(I - LP)u_2(k) - QL(y_0(k) + y_{ref} - Pu_0(k))
\end{aligned}
\tag{5.19}
$$

Since $u_2(0) = 0$ this can then be written as a recurrence relation to obtain

$$u_2(k) = \sum_{i=1}^{k} [Q(I - LP)]^{i-1} \left(-QL(y_0(k - i) + y_{ref} - Pu_0(k - i))\right). \tag{5.20}$$

To calculate the stability margin using the biased norm defined previously, with the bias as the reference signal y_{ref} we require $u_1(k)$ and $y_1(k)$. These are given by

$$u_1(k) = \sum_{i=1}^{k} [Q(I - LP)]^{i-1} \left(QL(y_0(k - i) + y_{ref} - Pu_0(k - i))\right) + u_0(k)$$

$$y_1(k) = Pu_1(k). \tag{5.21}$$

We can now substitute in the appropriate terms into the biased norm of equation 5.8 to find $\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{l^{p,q}(\mathbb{N}\times[0,T]),x_0}$:

$$x_0(k) = \begin{pmatrix} 0 \\ y_{ref} \end{pmatrix}$$

$$x_1(k) = \begin{pmatrix} u_0(k) \\ y_0(k) + y_{ref} \end{pmatrix}$$

$$\Pi_{\mathcal{M}//\mathcal{N}}x_0(k) = \begin{pmatrix} I \\ P \end{pmatrix} \sum_{i=1}^{k} [Q(I - LP)]^{i-1} \left(QLy_{ref}\right)$$

$$\Pi_{\mathcal{M}//\mathcal{N}}x_1(k) = \begin{pmatrix} I \\ P \end{pmatrix} \sum_{i=1}^{k} [Q(I - LP)]^{i-1} \left(QL(y_0(k - i) + y_{ref} - Pu_0(k - i))\right)$$

$$+ u_0(k). \quad (5.22)$$

For $q = 1$ we have

$$\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{l^{1,p}(\mathbb{N}\times[0,T]),x_0}$$

$$= \sup_{\substack{x_1\in\text{dom}(\Pi_{\mathcal{M}//\mathcal{N}}),\ \tau>0 \\ \|T_\tau(x_1(\cdot,\cdot)-x_0(\cdot,\cdot))\|_{l^{1,p}(\mathbb{N}\times[0,T]),x_0}\neq0}} \frac{\|T_\tau(\Pi_{\mathcal{M}//\mathcal{N}}x_1(\cdot,\cdot) - \Pi_{\mathcal{M}//\mathcal{N}}x_0(\cdot,\cdot))\|_{l^{1,p}(\mathbb{N}\times[0,T]),x_0}}{\|T_\tau(x_1(\cdot,\cdot) - x_0(\cdot,\cdot))\|_{l^{1,p}(\mathbb{N}\times[0,T]),x_0}}$$

$$= \sup_{\substack{\left(\begin{smallmatrix} u_0 \\ y_0 \end{smallmatrix}\right)\in l^{1,p}(\mathbb{N}\times[0,T]) \\ \left\|\left(\begin{smallmatrix} u_0 \\ y_0 \end{smallmatrix}\right)\right\|_{l^{1,p}(\mathbb{N}\times[0,T])}\neq0}}$$

$$\frac{\sum_{k=0}^{\infty} \left\|\begin{pmatrix} I \\ P \end{pmatrix} \sum_{i=1}^{k} [Q(I - LP)]^{i-1} \left(QL(y_0(k - i) - Pu_0(k - i))\right) + u_0(k)\right\|_{l^p[0,T]}}{\sum_{k=0}^{\infty} \left\|\begin{pmatrix} u_0(k) \\ y_0(k) \end{pmatrix}\right\|_{l^p[0,T]}}$$

$$(5.23)$$

and for $q = \infty$ we have

$$
\begin{aligned}
&\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{l^{\infty,p}(\mathbb{N}\times[0,T]),x_0} \\
&= \sup_{\substack{x_1\in\mathrm{dom}(\Pi_{\mathcal{M}//\mathcal{N}}),\ \tau>0 \\ \|T_\tau(x_1(\cdot,\cdot)-x_0(\cdot,\cdot))\|_{l^{\infty,p}(\mathbb{N}\times[0,T]),x_0}\neq 0}} \frac{\|T_\tau(\Pi_{\mathcal{M}//\mathcal{N}}x_1(\cdot,\cdot)-\Pi_{\mathcal{M}//\mathcal{N}}x_0(\cdot,\cdot))\|_{l^{\infty,p}(\mathbb{N}\times[0,T]),x_0}}{\|T_\tau(x_1(\cdot,\cdot)-x_0(\cdot,\cdot))\|_{l^{\infty,p}(\mathbb{N}\times[0,T]),x_0}} \\
&= \sup_{\substack{\left(\begin{smallmatrix}u_0\\y_0\end{smallmatrix}\right)\in l^{\infty,p}(\mathbb{N}\times[0,T]) \\ \|\left(\begin{smallmatrix}u_0\\y_0\end{smallmatrix}\right)\|_{l^{\infty,p}(\mathbb{N}\times[0,T])}\neq 0}} \\
&\quad\frac{\displaystyle\sup_{0\leq k\leq\infty}\left\|\begin{pmatrix}I\\P\end{pmatrix}\sum_{i=1}^{k}[Q(I-LP)]^{i-1}\left(QL(y_0(k-i)-Pu_0(k-i)))+u_0(k)\right)\right\|_{l^p[0,T]}}{\displaystyle\sup_{0\leq k\leq\infty}\left\|\begin{pmatrix}u_0(k)\\y_0(k)\end{pmatrix}\right\|_{l^p[0,T]}}.
\end{aligned}
$$

$$(5.24)$$

In both cases the gain of $\Pi_{\mathcal{M}//\mathcal{N}}$ can be bounded by examining the gains due to u_0 and y_0 separately and using the following property:

$$
\begin{aligned}
\frac{\|\Pi_{\mathcal{M}//\mathcal{N}}\left(\begin{smallmatrix}u_0\\y_0\end{smallmatrix}\right)\|}{\|\left(\begin{smallmatrix}u_0\\y_0\end{smallmatrix}\right)\|} &= \frac{\|\Pi_{\mathcal{M}//\mathcal{N}}\left(\begin{smallmatrix}u_0\\0\end{smallmatrix}\right)\|+\|\Pi_{\mathcal{M}//\mathcal{N}}\left(\begin{smallmatrix}0\\y_0\end{smallmatrix}\right)\|}{\|\left(\begin{smallmatrix}u_0\\y_0\end{smallmatrix}\right)\|} \\
&\leq \frac{\|\Pi_{\mathcal{M}//\mathcal{N}}|_{y_0=0}\|\|\left(\begin{smallmatrix}u_0\\0\end{smallmatrix}\right)\|+\|\Pi_{\mathcal{M}//\mathcal{N}}|_{u_0=0}\|\|\left(\begin{smallmatrix}0\\y_0\end{smallmatrix}\right)\|}{\|\left(\begin{smallmatrix}u_0\\y_0\end{smallmatrix}\right)\|} \\
&\leq \|\Pi_{\mathcal{M}//\mathcal{N}}|_{y_0=0}\|+\|\Pi_{\mathcal{M}//\mathcal{N}}|_{u_0=0}\|
\end{aligned}
$$

$$(5.25)$$

since $\|\left(\begin{smallmatrix}u_0\\0\end{smallmatrix}\right)\|\leq\|\left(\begin{smallmatrix}u_0\\y_0\end{smallmatrix}\right)\|$ and $\|\left(\begin{smallmatrix}0\\y_0\end{smallmatrix}\right)\|\leq\|\left(\begin{smallmatrix}u_0\\y_0\end{smallmatrix}\right)\|$.

Using this property we claim the following inequality can be obtained from equations 5.23 and 5.24:

$$
\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{l^{p,q}(\mathbb{N}\times[0,T]),x_0} \leq \|\left(\begin{smallmatrix}I\\P\end{smallmatrix}\right)\|_{l^p[0,T]}\left(1+\frac{\|QLP\|_{l^p[0,T]}+\|QL\|_{l^p[0,T]}}{1-\|Q(I-LP)\|_{l^p[0,T]}}\right).
$$

$$(5.26)$$

1. From equation 5.23:

Firstly set $u_0 = 0$. By various substitutions; triangle inequalities; and rearrangement of summations (Apostol, 1974, Definition 8.21, due to all terms being normed and therefore

positive):

$$\left\|\Pi_{\mathcal{M}//\mathcal{N}}\left(\begin{smallmatrix}0\\y_0\end{smallmatrix}\right)\right\|_{l^{1,p}(\mathbb{N}\times[0,T]),x_0}$$

$$= \sum_{k=0}^{\infty}\left\|\begin{pmatrix}I\\P\end{pmatrix}\sum_{i=1}^{k}\left[Q(I-LP)\right]^{i-1}QLy_0(k-i)\right\|_{l^p[0,T]}$$

$$\leq \left\|\begin{pmatrix}I\\P\end{pmatrix}\right\|\sum_{k=0}^{\infty}\sum_{n=0}^{k-1}\left\|\left[Q(I-LP)\right]^{k-n-1}QLy_0(n)\right\|_{l^p[0,T]}$$

$$\leq \left\|\begin{pmatrix}I\\P\end{pmatrix}\right\|\sum_{n=0}^{\infty}\sum_{k=n+1}^{\infty}\left\|\left[Q(I-LP)\right]^{k-n-1}QLy_0(n)\right\|_{l^p[0,T]}$$

$$\leq \left\|\begin{pmatrix}I\\P\end{pmatrix}\right\|\|QL\|\sum_{j=0}^{\infty}\left\|\left[Q(I-LP)\right]^j\right\|\sum_{n=0}^{\infty}\|y_0(n)\|_{l^p[0,T]}.$$

$$(5.27)$$

Hence, as $\|Q(I-LP)\|_{l^p[0,T]} < 1$ and $\|Q(I-LP)^j\|_{l^p[0,T]} \leq \|Q(I-LP)\|_{l^p[0,T]}^j$ for $0 \leq j \leq \infty$,

$$\left\|\Pi_{\mathcal{M}//\mathcal{N}}|_{u_0=0}\right\|_{l^{1,p}(\mathbb{N}\times[0,T]),x_0} \leq \frac{\left\|\left(\begin{smallmatrix}1\\P\end{smallmatrix}\right)\right\|_{l^p[0,T]}\|QL\|_{l^p[0,T]}}{1-\|Q(I-LP)\|_{l^p[0,T]}}. \qquad (5.28)$$

Setting $y_0 = 0$ we obtain:

$$\left\|\Pi_{\mathcal{M}//\mathcal{N}}\left(\begin{smallmatrix}u_0\\0\end{smallmatrix}\right)\right\|_{l^{1,p}(\mathbb{N}\times[0,T]),x_0}$$

$$= \sum_{k=0}^{\infty}\left\|\begin{pmatrix}I\\P\end{pmatrix}\left(u_0(k)+\sum_{i=1}^{k}\left[Q(I-LP)\right]^{i-1}QLPu_0(k-i)\right)\right\|_{l^p[0,T]}$$

$$\leq \sum_{k=0}^{\infty}\left\|\begin{pmatrix}I\\P\end{pmatrix}\sum_{i=1}^{k}\left[Q(I-LP)\right]^{i-1}QLPu_0(k-i)\right\|_{l^p[0,T]} + \sum_{k=0}^{\infty}\left\|\begin{pmatrix}I\\P\end{pmatrix}u_0(k)\right\|_{l^p[0,T]}.$$

$$(5.29)$$

The term on the left is bounded using the same substitutions and method as for the case above for $u_0 = 0$, and so the bound is given by

$$\left\|\Pi_{\mathcal{M}//\mathcal{N}}|_{y_0=0}\right\|_{l^{1,p}(\mathbb{N}\times[0,T]),x_0} \leq \frac{\left\|\left(\begin{smallmatrix}I\\P\end{smallmatrix}\right)\right\|_{l^p[0,T]}\|QLP\|_{l^p[0,T]}}{1-\|Q(I-LP)\|_{l^p[0,T]}} + \left\|\left(\begin{smallmatrix}I\\P\end{smallmatrix}\right)\right\|_{l^p[0,T]}. \qquad (5.30)$$

By equations 5.28, 5.30 and the property given by equation 5.25 we can derive equation 5.26.

2. Looking at equation 5.24:

As before we will start with $u_0 = 0$. By the triangle inequality:

$$\|\Pi_{\mathcal{M}//\mathcal{N}}\begin{pmatrix} 0 \\ y_0 \end{pmatrix}\|_{l^{\infty,p}(\mathbb{N}\times[0,T]),x_0}$$

$$= \sup_{0 \le k \le \infty} \left\| \begin{pmatrix} I \\ P \end{pmatrix} \sum_{i=1}^{k} [Q(I-LP)]^{i-1} QLy_0(k-i) \right\|_{l^p[0,T]}$$

$$\le \sup_{0 \le k \le \infty} \sum_{i=1}^{k} \left\| \begin{pmatrix} I \\ P \end{pmatrix} [Q(I-LP)]^{i-1} QLy_0(k-i) \right\|_{l^p[0,T]}$$

$$\le \sup_{0 \le k \le \infty} \sum_{i=1}^{k} \left(\left\| \begin{pmatrix} I \\ P \end{pmatrix} [Q(I-LP)]^{i-1} QL \right\|_{l^p[0,T]} \|y_0(k-i)\|_{l^p[0,T]} \right) \quad (5.31)$$

$$\le \left(\sup_{0 \le k \le \infty} \sum_{i=1}^{k} \left\| \begin{pmatrix} I \\ P \end{pmatrix} [Q(I-LP)]^{i-1} QL \right\|_{l^p[0,T]} \right) \sup_{0 \le k \le \infty} \|y_0(k)\|_{l^p[0,T]}$$

$$= \left(\lim_{k \to \infty} \sum_{i=1}^{k} \left\| \begin{pmatrix} I \\ P \end{pmatrix} [Q(I-LP)]^{i-1} QL \right\|_{l^p[0,T]} \right) \sup_{0 \le k \le \infty} \|y_0(k)\|_{l^p[0,T]}.$$

Hence, as $\|Q(I-LP)\|_{l^p[0,T]} < 1$ and $\|Q(I-LP)^j\|_{l^p[0,T]} \le \|Q(I-LP)\|^j_{l^p[0,T]}$ for $0 \le j \le \infty$,

$$\|\Pi_{\mathcal{M}//\mathcal{N}}|_{u_0=0}\|_{l^{\infty,p}(\mathbb{N}\times[0,T]),x_0} \le \frac{\|\begin{pmatrix} I \\ P \end{pmatrix}\|_{l^p[0,T]} \|QL\|_{l^p[0,T]}}{1 - \|Q(I-LP)\|_{l^p[0,T]}}. \quad (5.32)$$

Now moving on to u_0 and so setting $y_0 = 0$ we obtain

$$\|\Pi_{\mathcal{M}//\mathcal{N}}\begin{pmatrix} u_0 \\ 0 \end{pmatrix}\|_{l^{\infty,p}(\mathbb{N}\times[0,T]),x_0}$$

$$= \sup_{0 \le k \le \infty} \left\| \begin{pmatrix} I \\ P \end{pmatrix} \left(u_0(k) + \sum_{i=1}^{k} [Q(I-LP)]^{i-1} QLPu_0(k-i) \right) \right\|_{l^p[0,T]}$$

$$\le \sup_{0 \le k \le \infty} \left\| \begin{pmatrix} I \\ P \end{pmatrix} \sum_{i=1}^{k} [Q(I-LP)]^{i-1} QLPu_0(k-i) \right\|_{l^p[0,T]}$$

$$+ \sup_{0 \le k \le \infty} \left\| \begin{pmatrix} I \\ P \end{pmatrix} u_0(k) \right\|_{l^p[0,T]}. \quad (5.33)$$

The term on the left is bounded using the same substitutions and method as for the case above for $u_0 = 0$, and so the bound is given by

$$\|\Pi_{\mathcal{M}//\mathcal{N}}|_{y_0=0}\|_{l^{\infty,p}(\mathbb{N}\times[0,T]),x_0} \le \frac{\|\begin{pmatrix} I \\ P \end{pmatrix}\|_{l^p[0,T]} \|QLP\|_{l^p[0,T]}}{1 - \|Q(I-LP)\|_{l^p[0,T]}} + \|\begin{pmatrix} I \\ P \end{pmatrix}\|_{l^p[0,T]}. \quad (5.34)$$

By equations 5.32, 5.34 and the property given by equation 5.25 we can derive equation 5.26. $\quad\square$

In order to bound the summations in inequalities 5.28 and 5.32 we imposed the sufficient condition that $\|Q(I - LP)\|_{l^p[0,T]} < 1$. A similar condition arises in a host of papers including Hätönen (2004), Norrlöf and Gunnarsson (2002), Hätönen et al. (2004), Norrlöf (2000b) and Harte et al. (2005); where this is supplied as the sufficient condition for monotonic convergence of this algorithm (see Theorem 3.5).

Therefore, we now have a bound on $b_{P,C} = \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}^{-1}$ for $\mathcal{W} = l^{p,q}(\mathbb{N} \times [0,T])$ so we can apply the robust stability theorem (Theorem 4.17).

Theorem 5.2. *Consider the closed-loop system $[\overline{P}, C]$ given by Figure 5.2, equations 5.11 and 5.14, and with the matrix P given by equation 5.10. Let $q \in \{1, \infty\}$ and $p \in [1, \infty]$ Suppose $\|Q(I - LP)\|_{l^p[0,T]} < 1$. Suppose $[\overline{P}_1, C]$ is globally well-posed and let $\mathcal{W} = l^{p,q}(\mathbb{N} \times [0,T])$ and $x_0 = \begin{pmatrix} 0 \\ y_{ref} \end{pmatrix} \in \mathcal{W}_e$.*

Define

$$b_{\overline{P},C} = \left(\|\begin{pmatrix} I \\ P \end{pmatrix}\|_{l^p[0,T]} \left(1 + \frac{\|QLP\|_{l^p[0,T]} + \|QL\|_{l^p[0,T]}}{1 - \|Q(I - LP)\|_{l^p[0,T]}} \right) \right)^{-1}. \tag{5.35}$$

Suppose the plant $\overline{P}_1$, with $\mathcal{M}_1 := \mathcal{G}_{\overline{P}_1, g_1}$ is such that there exists a causal surjective $\Phi: \mathcal{M} \to \mathcal{M}_1$, where $\mathcal{M} := \mathcal{G}_{\overline{P}, x_1}$ with the plant satisfying $\|(\Phi - I)|_{\mathcal{G}_{\overline{P}, x_1}}\|_{\mathcal{W}, x_1} < b_{\overline{P}, C}$, where $x_1 = \Pi_{\mathcal{M}//\mathcal{N}} x_0$.

Then the stability of $[\overline{P}_1, C]$ on $\mathcal{W}$ with respect to $g_0 = (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$ is assured and

$$\|\Pi_{\mathcal{M}_1//\mathcal{N}}\|_{\mathcal{W}, g_0} \leq \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0} \frac{1 + \|(\Phi - I)|_{\mathcal{G}_{\overline{P}, x_1}}\|_{\mathcal{W}, x_1}}{1 - \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W}, x_0} \|(\Phi - I)|_{\mathcal{G}_{\overline{P}, x_1}}\|_{\mathcal{W}, x_1}}. \tag{5.36}$$

Proof. The proof follows directly from Theorems 4.8 and 5.1. □

This theorem provides a 2D robust stability guarantee for a set of closed-loop linear plants engaged in trajectory tracking within the lifted system framework, providing a set of plants for which the controller is sufficient to guarantee stability.

It should be noted that the above theorem does not explicitly use the gap $\vec{\delta}(\overline{P}, \overline{P}_1)$. However, since the gap is given by

$$\vec{\delta}(\overline{P}, \overline{P}_1) = \left\{ \inf_{\Phi} \|(\Phi - I)|_{\mathcal{G}_{\overline{P}, x_1}}\|_{\mathcal{W}, x_1} \; \middle| \; \Phi: \mathcal{M} \to \mathcal{M}_1 \text{ is causal and surjective} \right\}, \tag{5.37}$$

if there does exists a Φ such that $\|(\Phi - I)|_{\mathcal{G}_{\overline{P}, x_1}}\|_{\mathcal{W}, x_1} < b_{P,C}$ then it is clear that $\vec{\delta}(\overline{P}, \overline{P}_1) \leq \|(\Phi - I)|_{\mathcal{G}_{\overline{P}, x_1}}\|_{\mathcal{W}, x_1}$. The reason that Φ is included as opposed to the gap is due to it being necessary to define $g_0 = (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$.

5.3 The Relationship Between 1D and 2D Gaps

This section will show how this 2D gap can be related back to the standard 1D gap metric in the linear case. The reason for this restriction will be made clear.

Before a theorem can be given showing this relation it is necessary to introduce some new notation. As the theorem provides a link between the gap metric on different domains: $l^{p,q}(\mathbb{N} \times [0,T])$, $l^p[0,T]$ and $l^p[\mathbb{R}+]$; and also links the biased and non-biased gaps, the notation is sizeable but essential. Where possible, conventions are used: for example, the signal space in question will always appear as a superscript in graph and operator definitions.

We recall the graph of an operator $P \colon \mathcal{U}_e \to \mathcal{Y}_e$ relative to a biased pair of signals $x \in \mathcal{W}_e$

$$\mathcal{G}_{P,x} = \left\{ w \in \mathcal{W}_e \ \middle|\ w = \begin{pmatrix} u_1 \\ Pu_1 \end{pmatrix} \in \mathcal{W}_x \right\}, \tag{5.38}$$

and the graph of the operator $P \colon \mathcal{U}_e \to \mathcal{Y}_e$ without bias

$$\mathcal{G}_P = \left\{ w \in \mathcal{W}_e \ \middle|\ w = \begin{pmatrix} u_1 \\ Pu_1 \end{pmatrix} \in \mathcal{W} \right\}. \tag{5.39}$$

As the graphs only contain bounded signals (in the appropriate sense) the biased and non-biased graphs are likely to be different sets unless the bias lies in the non-extended signal space.

We now define two sets of operators:

$$\mathcal{O}_{P,P_1}^{\mathcal{W}} := \left\{ \Phi \colon \mathcal{G}_P^{\mathcal{W}} \to \mathcal{G}_{P_1}^{\mathcal{W}} \ \middle|\ \Phi \text{ is causal, surjective and } \Phi(0) = 0 \right\}, \tag{5.40}$$

and

$$\mathcal{Q}_{P,x_1,P_1,g_1}^{\mathcal{W}} := \left\{ \Psi \colon \mathcal{G}_{P,x_1}^{\mathcal{W}} \to \mathcal{G}_{P_1,g_1}^{\mathcal{W}} \ \middle|\ \exists \Phi \in \mathcal{O}_{P,P_1}^{\mathcal{W}} \text{ such that } \forall x \in \mathcal{G}_{P,x_1}^{\mathcal{W}} \right.$$
$$\Psi(x) = \Phi(x - x_1) + \Psi(x_1) \tag{5.41}$$
$$\left. \text{and } \Psi \text{ is surjective} \right\}.$$

The biased gap is then defined as

$$\vec{\delta}_{x_1,g_1}^{\mathcal{W}}(P,P_1) = \begin{cases} \displaystyle \inf_{\Psi \in \mathcal{Q}_{P,x_1,P_1,g_1}^{\mathcal{W}}} \left\| (\Psi - I)|_{\mathcal{G}_{P,x_1}^{\mathcal{W}}} \right\|_{\mathcal{W},x_1} \\[2mm] \infty \text{ if no such operator } \Psi \text{ exists} \end{cases} \tag{5.42}$$

and the non-biased gap as

$$\vec{\delta}^{\mathcal{W}}(P, P_1) = \begin{cases} \inf_{\Phi \in \mathcal{O}^{\mathcal{W}}_{P,P_1}} \| (\Phi - I)|_{\mathcal{G}^{\mathcal{W}}_P} \|_{\mathcal{W}} \\ \\ \infty \text{ if not such operator } \Phi \text{ exists} \end{cases} . \tag{5.43}$$

5.3.1 Causality

Before the main theorem it is necessary to show that the causality of all operators within $\mathcal{Q}^{\mathcal{W}}_{P,x_1,P_1,g_1}$ is inherited directly from the causality (by definition) of all operators within $\mathcal{O}^{\mathcal{W}}_{P,P_1}$.

Lemma 5.3. *All operators in the set $\mathcal{Q}^{\mathcal{W}}_{P,x_1,P_1,g_1}$ are causal.*

Proof. Let $\Psi \in \mathcal{Q}^{\mathcal{W}}_{P,x_1,P_1,g_1}$ and $\Phi \in \mathcal{O}^{\mathcal{W}}_{P,P_1}$ such that $\Psi(x) = \Phi(x - x_1) + \Psi(x_1)$. Let x, y be such that $T_\tau x = T_\tau y$ for any $\tau > 0$. Then it follows from the causality of Φ (Definition 4.4) that for all $\tau > 0$:

$$\begin{aligned} T_\tau \Psi x &= T_\tau \left[\Phi(x - x_1) + \Psi(x_1) \right] \\ &= T_\tau \Phi(x - x_1) + T_\tau \Psi(x_1) \\ &= T_\tau \Phi(T_\tau x - T_\tau x_1) + T_\tau \Psi(x_1) \\ &= T_\tau \Phi(T_\tau y - T_\tau x_1) + T_\tau \Psi(x_1) \\ &= T_\tau \Phi(y - x_1) + T_\tau \Psi(x_1) \\ &= T_\tau \left[\Phi(y - x_1) + \Psi(x_1) \right] \\ &= T_\tau \Psi y, \end{aligned} \tag{5.44}$$

and therefore Ψ is causal for all $\Psi \in \mathcal{Q}^{\mathcal{W}}_{P,x_1,P_1,g_1}$. $\qquad\square$

5.3.2 Surjectivity

It was shown in Lemma 5.3 that the causality of operators in the unbiased set $\mathcal{O}^{\mathcal{W}}_{P,P_1}$ implies the causality of all operators within the biased set $\mathcal{Q}^{\mathcal{W}}_{P,x_1,P_1,g_1}$. Unfortunately the same currently remains unclear for the inheritance of surjectivity. In general it is best to remove as many unnecessary conditions on the sets of operators and therefore removing the surjectivity condition on the set of biased operators is desirable. Also, in Lemma 5.6 we will require that there exists an operator in the biased set $\mathcal{Q}$ based on the existence of an operator in $\mathcal{O}$. Without the inheritance of surjectivity from operators in $\mathcal{O}$ to operators in $\mathcal{Q}$ this may not be possible since the resulting operator may not be surjective and therefore may be excluded from $\mathcal{Q}$.

Proposition 5.5 proves that the surjectivity (by definition) of operators within $\mathcal{O}^{\mathcal{W}}_{P,P_1}$ implies the surjectivity of operators within $\mathcal{Q}^{\mathcal{W}}_{P,x_1,P_1,g_1}$ if the plants P and P_1 are linear. It

is hoped that a proof could be developed to complete the theorem for nonlinear systems also, since this is currently the only area within the theorem that requires linearity. Before this can be given, the following lemma is required to relate the biased and non-biased graphs. The non-biased graph is a subspace and so possesses some additional properties. The biased graph is not however, since for example if $x_0 \in \mathcal{W}_e \setminus \mathcal{W}$ then $2x_0 \notin \mathcal{W}_{x_0}$. In most instances the biased and non-biased graphs do not overlap, therefore in order to compare the two it is necessary to prove the existence of elements within one graph based on the existence of elements within the other.

In the following lemma and proposition the superscript $\mathcal{W}$ above the graph symbol $\mathcal{G}$ has been dropped for convenience.

Lemma 5.4. *Given a linear operator $P: \mathcal{U}_e \to \mathcal{Y}_e$, then for any $w = (u_1, Pu_1) \in \mathcal{W}_e$ the following three properties hold:*

i) For any $x \in \mathcal{G}_P$ there exists $y \in \mathcal{G}_{P,w}$ such that $y = x + w$.

ii) For any $y \in \mathcal{G}_{P,w}$ there exists $x \in \mathcal{G}_P$ such that $x = y - w$.

iii) For any $y \in \mathcal{G}_{P,w}$ then $w \in \mathcal{G}_{P,y}$.

iv) For any $y \in \mathcal{G}_{P,w}$ then $\mathcal{G}_{P,y} = \mathcal{G}_{P,w}$

Proof. Let $w = \left(\begin{smallmatrix} w_1 \\ w_2 \end{smallmatrix}\right)$, $x = \left(\begin{smallmatrix} x_1 \\ x_2 \end{smallmatrix}\right)$ and $y = \left(\begin{smallmatrix} y_1 \\ y_2 \end{smallmatrix}\right)$. We therefore have $w_2 = Pw_1$ with $w_2 \in \mathcal{Y}_e$ and $w_1 \in \mathcal{U}_e$.

For the proof of property i), let $x \in \mathcal{G}_P$ therefore $x_2 = Px_1$ with $x_2 \in \mathcal{Y}$ and $x_1 \in \mathcal{U}$.

By the linearity of P we can define $y_2 = Py_1$ such that $y_2 = (x_2 + w_2) \in \mathcal{Y}_{w_2}$ and $y_1 = (x_1 + w_1) \in \mathcal{U}_{w_1}$.

Therefore $y \in \mathcal{G}_{P,w}$. The proof of property ii) is trivially similar.

To prove property iii) first take any $y \in \mathcal{G}_{P,w}$, recalling the biased graph and biased signal space definitions:

$$
\mathcal{G}_{P,\left(\begin{smallmatrix} w_1 \\ w_2 \end{smallmatrix}\right)} = \left\{ \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} \in \mathcal{U}_e \times \mathcal{Y}_e \; \middle| \; y_1 \in \mathcal{U}_{w_1}, \; y_2 \in \mathcal{Y}_{w_2}, \; y_2 = Py_1 \right\}.
$$
$$
\mathcal{U}_{w_1} = \{\, y_1 \in \mathcal{U}_e \;\mid\; y_1 - w_1 \in \mathcal{U} \,\}
$$
$$
\mathcal{Y}_{w_2} = \{\, y_2 \in \mathcal{Y}_e \;\mid\; y_2 - w_2 \in \mathcal{Y} \,\} \tag{5.45}
$$

Since $\mathcal{U}$ and $\mathcal{Y}$ are subspaces, if $y_1 - w_1 \in \mathcal{U}$ then $w_1 - y_1 \in \mathcal{U}$, and if $y_2 - w_2 \in \mathcal{Y}$ then $w_2 - y_2 \in \mathcal{Y}$. Therefore if $y \in \mathcal{W}_w$ then $w \in \mathcal{W}_y$. Also, since w consists of pairs of signals compatible with the plant by definition, then $w \in \mathcal{G}_{P,y}$. This holds for any $y \in \mathcal{G}_{P,w}$.

Property iv) is demonstrated by considering that for all $y \in \mathcal{G}_{P,w}$ then for any $x \in \mathcal{G}_{P,y}$ we can state that $x \in \mathcal{G}_{P,w}$ (and conversely that for all $w \in \mathcal{G}_{P,y}$ then for any $x \in \mathcal{G}_{P,w}$

then $x \in \mathcal{G}_{P,y}$). To show this first recall from property ii) that for any $x \in \mathcal{G}_{P,y}$ there exists $u \in \mathcal{G}_P$ such that $u = x - y$. Similarly, since $y \in \mathcal{G}_{P,w}$ there exists $v \in \mathcal{G}_P$ such that $v = y - w$. As P is linear $\mathcal{G}_P$ forms a subspace. Therefore, since $u, v \in \mathcal{G}_P$ we can state that $(u - v) \in \mathcal{G}_P$ and therefore $(x - w) \in \mathcal{G}_P$ and so $x \in \mathcal{G}_{P,w}$. As this holds for any $x \in \mathcal{G}_{P,y}$, for any $y \in \mathcal{G}_{P,w}$ and this argument is symmetrical we can state that $\mathcal{G}_{P,y} = \mathcal{G}_{P,w}$. $\qquad\square$

The following proposition now establishes that, when confining the analysis to linear plants, the restriction to surjective operators can be dropped in the definition of the set $\mathcal{Q}_{P,x_1,P_1,g_1}^{\mathcal{W}}$ given in equation 5.41. This is due to the surjectivity condition being inherited automatically from the set $\mathcal{O}_{P,P_1}^{\mathcal{W}}$.

Proposition 5.5. *Suppose $P\colon \mathcal{U}_e \to \mathcal{Y}_e$ and $P_1\colon \mathcal{U}_e \to \mathcal{Y}_e$ are linear. Let $\Psi \in \mathcal{Q}_{P,x_1,P_1,g_1}^{\mathcal{W}}$, where*

$$\mathcal{Q}_{P,x_1,P_1,g_1}^{\mathcal{W}} := \big\{ \, \Psi \colon \mathcal{G}_{P,x_1}^{\mathcal{W}} \to \mathcal{G}_{P_1,g_1}^{\mathcal{W}} \ \mid \ \exists \Phi \in \mathcal{O}_{P,P_1}^{\mathcal{W}} \ \text{ such that } \forall x \in \mathcal{G}_{P,x_1}^{\mathcal{W}}$$
$$\Psi(x) = \Phi(x - x_1) + \Psi(x_1) \, \big\}. \qquad (5.46)$$

All operators $\Psi \in \mathcal{Q}_{P,x_1,P_1,g_1}^{\mathcal{W}}$ are surjective.

Proof. In order to prove the surjectivity of all operators $\Psi \in \mathcal{Q}_{P,x_1,P_1,g_1}^{\mathcal{W}}$ it must be shown that for all $z \in \mathcal{G}_{P_1,g_1}$ there exists $x \in \mathcal{G}_{P,x_1}$ such that $z = \Psi x$.

By definition, recall that for any $\Psi \in \mathcal{Q}_{P,x_1,P_1,g_1}^{\mathcal{W}}$ there exists $\Phi \in \mathcal{O}_{P,P_1}^{\mathcal{W}}$ such that for all $x \in \mathcal{G}_{P,x_1}$

$$\Psi(x) = \Phi(x - x_1) + \Psi(x_1). \qquad (5.47)$$

Now, since $\Phi\colon \mathcal{G}_P \to \mathcal{G}_{P_1}$ is surjective, for any $w \in \mathcal{G}_{P_1}$ there exists $z \in \mathcal{G}_P$ such that $w = \Phi z$. Let $x = z + x_1$, then $x - x_1 \in \mathcal{G}_P$ and $w = \Phi(x - x_1)$.

Since x_1 is fixed and since P is linear, from property i) of Lemma 5.4 we see that since $(x - x_1) \in \mathcal{G}_P$ then $x \in \mathcal{G}_{P,x_1}$.

We have therefore proven that for all $w \in \mathcal{G}_{P_1}$ there exists $x \in \mathcal{G}_{P,x_1}$ such that

$$w = \Psi(x) - \Psi(x_1). \qquad (5.48)$$

All that therefore remains to be shown is that for any $z \in \mathcal{G}_{P_1,g_1}$ there exists $w \in \mathcal{G}_{P_1}$ such that $w = z - \Psi(x_1)$ for then, by equation 5.48, it follows that $z = \Psi(x)$ as required. In order to show this let $y = \Psi(x_1)$. By definition $y \in \mathcal{G}_{P_1,g_1}$. Then, by property ii), for any $z \in \mathcal{G}_{P_1,y}$ we can find a $w \in \mathcal{G}_{P_1}$ such that $w = z - y$. Finally, by property iv) of Lemma 5.4 we can state that $\mathcal{G}_{P_1,g_1} = \mathcal{G}_{P_1,y}$. $\qquad\square$

To clarify: The 2D gap given obeys all of the conditions given in Theorem 4.17 and so is permissible gap measure; as is the 1D gap measure. The above lemma shows that, for the case of linear systems, surjectivity is not compromised in the movement between gaps in Theorem 5.10 (given in the following section) and so the condition of surjectivity is not required in equation 5.41. If this were proven for nonlinear systems then the result would allow the use of the standard nonlinear gap measure on nonlinear 2D systems.

5.3.3 Relating the 1D and 2D gaps

The following three lemmata will detail the relationship between the biased 2D gap and the equivalent unbiased 1D gap. Some of the details within the proofs are adapted from French (2008), where they instead appeared for continuous time systems using the norm of equation 5.1. Here we use the discrete version of the 2D norm, given in equation 5.5.

The first lemma relates the biased and unbiased 2D gaps to each other, and is the only one of the three that requires the plants to be linear.

Lemma 5.6. *For $p, q \in [1, \infty]$ and linear plants $\overline{P}, \overline{P}_1 \colon l_e^{p,q}(\mathbb{N} \times [0, T]) \to l_e^{p,q}(\mathbb{N} \times [0, T])$*

$$\vec{\delta}_{x_1, g_1}^{l^{p,q}(\mathbb{N} \times [0,T])}(\overline{P}, \overline{P}_1) = \vec{\delta}^{l^{p,q}(\mathbb{N} \times [0,T])}(\overline{P}, \overline{P}_1). \tag{5.49}$$

Proof. Starting with the definition:

$$\vec{\delta}_{x_1, g_1}^{l^{p,q}(\mathbb{N} \times [0,T])}(\overline{P}, \overline{P}_1)$$

$$= \inf_{\Psi \in \mathcal{Q}_{\overline{P}, x_1, \overline{P}_1, g_1}^{l^{p,q}(\mathbb{N} \times [0,T])}} \left\| (\Psi - I)\big|_{\mathcal{G}_{\overline{P}, x_1}^{l^{p,q}(\mathbb{N} \times [0,T])}} \right\|_{l^{p,q}(\mathbb{N} \times [0,T]), x_1}$$

$$= \inf_{\Psi \in \mathcal{Q}_{\overline{P}, x_1, \overline{P}_1, g_1}^{l^{p,q}(\mathbb{N} \times [0,T])}} \sup_{\substack{x \in \mathcal{G}_{\overline{P}, x_1}^{l^{p,q}(\mathbb{N} \times [0,T])}, \ \tau > 0 \\ \|T_\tau(x - x_1)\|_{l^{p,q}(\mathbb{N} \times [0,T])} \neq 0}} \frac{\|T_\tau((\Psi - I)x - (\Psi - I)x_1)\|_{l^{p,q}(\mathbb{N} \times [0,T])}}{\|T_\tau(x - x_1)\|_{l^{p,q}(\mathbb{N} \times [0,T])}}. \tag{5.50}$$

Recalling that $\Psi(x) = \Phi(x - x_1) + \Psi(x_1)$ we claim that

$$\vec{\delta}_{x_1, g_1}^{l^{p,q}(\mathbb{N} \times [0,T])}(\overline{P}, \overline{P}_1)$$

$$= \inf_{\Phi \in \mathcal{O}_{\overline{P}, \overline{P}_1}^{l^{p,q}(\mathbb{N} \times [0,T])}} \sup_{\substack{(x - x_1) \in \mathcal{G}_{\overline{P}}^{l^{p,q}(\mathbb{N} \times [0,T])}, \ \tau > 0 \\ \|T_\tau(x - x_1)\|_{l^{p,q}(\mathbb{N} \times [0,T])} \neq 0}} \frac{\|T_\tau(\Phi - I)(x - x_1)\|_{l^{p,q}(\mathbb{N} \times [0,T])}}{\|T_\tau(x - x_1)\|_{l^{p,q}(\mathbb{N} \times [0,T])}}$$

$$= \vec{\delta}^{l^{p,q}(\mathbb{N} \times [0,T])}(\overline{P}, \overline{P}_1). \tag{5.51}$$

It remains to show that this claim is correct. In order to do this consider the following two statements:

For all $\Psi \in \mathcal{Q}_{\overline{P},x_1,\overline{P}_1,g_1}^{lp,q(\mathbb{N}\times[0,T])}$ there exists $\Phi_\Psi \in \mathcal{O}_{\overline{P},\overline{P}_1}^{lp,q(\mathbb{N}\times[0,T])}$ such that

$$\|(\Psi - I)x - (\Psi - I)x_1\|_{lp,q(\mathbb{N}\times[0,T])} = \|(\Phi_\Psi - I)(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])}; \qquad (5.52)$$

and conversely, for all $\Phi \in \mathcal{O}_{\overline{P},\overline{P}_1}^{lp,q(\mathbb{N}\times[0,T])}$ there exists $\Psi_\Phi \in \mathcal{Q}_{\overline{P},x_1,\overline{P}_1,g_1}^{lp,q(\mathbb{N}\times[0,T])}$ such that

$$\|(\Phi - I)(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])} = \|(\Psi_\Phi - I)x - (\Psi_\Phi - I)x_1\|_{lp,q(\mathbb{N}\times[0,T])}. \qquad (5.53)$$

Both these follow from the definition of $\Psi \in \mathcal{Q}_{\overline{P},x_1,\overline{P}_1,g_1}^{\mathcal{W}}$ given in equation 5.41, with the second also requiring the surjectivity of operators in $\mathcal{Q}_{\overline{P},x_1,\overline{P}_1,g_1}^{lp,q(\mathbb{N}\times[0,T])}$ to be inherited from $\mathcal{O}_{\overline{P},\overline{P}_1}^{lp,q(\mathbb{N}\times[0,T])}$. This is established by Proposition 5.5 due to the linearity of $\overline{P}$ and $\overline{P}_1$. (See Remark 5.7 for an explanation of this requirement.)

From these two statements it follows that

$$\begin{aligned}
&\inf_{\substack{\Psi \in \mathcal{Q}_{\overline{P},x_1,\overline{P}_1,g_1}^{lp,q(\mathbb{N}\times[0,T])}}} \sup_{\substack{x \in \mathcal{G}_{\overline{P},x_1}^{lp,q(\mathbb{N}\times[0,T])},\ \tau>0 \\ \|T_\tau(x-x_1)\|_{lp,q(\mathbb{N}\times[0,T])}\neq 0}} \frac{\|T_\tau((\Psi - I)x - (\Psi - I)x_1)\|_{lp,q(\mathbb{N}\times[0,T])}}{\|T_\tau(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])}} \\
&= \inf_{\substack{\Psi \in \mathcal{Q}_{\overline{P},x_1,\overline{P}_1,g_1}^{lp,q(\mathbb{N}\times[0,T])}}} \sup_{\substack{x \in \mathcal{G}_{\overline{P},x_1}^{lp,q(\mathbb{N}\times[0,T])},\ \tau>0 \\ \|T_\tau(x-x_1)\|_{lp,q(\mathbb{N}\times[0,T])}\neq 0}} \frac{\|T_\tau(\Phi_\Psi - I)(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])}}{\|T_\tau(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])}} \\
&\geq \inf_{\substack{\Phi \in \mathcal{O}_{\overline{P},\overline{P}_1}^{lp,q(\mathbb{N}\times[0,T])}}} \sup_{\substack{(x-x_1) \in \mathcal{G}_{\overline{P}}^{lp,q(\mathbb{N}\times[0,T])},\ \tau>0 \\ \|T_\tau(x-x_1)\|_{lp,q(\mathbb{N}\times[0,T])}\neq 0}} \frac{\|T_\tau(\Phi - I)(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])}}{\|T_\tau(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])}}.
\end{aligned} \qquad (5.54)$$

However, if the process is taken in the opposite direction:

$$\begin{aligned}
&\inf_{\substack{\Phi \in \mathcal{O}_{\overline{P},\overline{P}_1}^{lp,q(\mathbb{N}\times[0,T])}}} \sup_{\substack{(x-x_1) \in \mathcal{G}_{\overline{P}}^{lp,q(\mathbb{N}\times[0,T])},\ \tau>0 \\ \|T_\tau(x-x_1)\|_{lp,q(\mathbb{N}\times[0,T])}\neq 0}} \frac{\|T_\tau(\Phi - I)(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])}}{\|T_\tau(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])}} \\
&= \inf_{\substack{\Phi \in \mathcal{O}_{\overline{P},\overline{P}_1}^{lp,q(\mathbb{N}\times[0,T])}}} \sup_{\substack{(x-x_1) \in \mathcal{G}_{\overline{P}}^{lp,q(\mathbb{N}\times[0,T])},\ \tau>0 \\ \|T_\tau(x-x_1)\|_{lp,q(\mathbb{N}\times[0,T])}\neq 0}} \frac{\|T_\tau((\Psi_\Phi - I)x - (\Psi_\Phi - I)x_1)\|_{lp,q(\mathbb{N}\times[0,T])}}{\|T_\tau(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])}} \\
&\geq \inf_{\substack{\Psi \in \mathcal{Q}_{\overline{P},x_1,\overline{P}_1,g_1}^{lp,q(\mathbb{N}\times[0,T])}}} \sup_{\substack{x \in \mathcal{G}_{\overline{P},x_1}^{lp,q(\mathbb{N}\times[0,T])},\ \tau>0 \\ \|T_\tau(x-x_1)\|_{lp,q(\mathbb{N}\times[0,T])}\neq 0}} \frac{\|T_\tau((\Psi - I)x - (\Psi - I)x_1)\|_{lp,q(\mathbb{N}\times[0,T])}}{\|T_\tau(x - x_1)\|_{lp,q(\mathbb{N}\times[0,T])}}.
\end{aligned}$$

$$(5.55)$$

It is clear that the two inequalities are in different directions and therefore both infimums must be identical in order for the two inequalities to hold.

Note that in both of the above processes (equations 5.54 and 5.55) it is required that $(x - x_1) \in \mathcal{G}_{\overline{P}}^{L^p(\mathbb{N}\times[0,T])}$ is equivalent to $x \in \mathcal{G}_{\overline{P},x_1}^{L^p(\mathbb{N}\times[0,T])}$. This is shown in Lemma 5.4,

properties i) and ii). □

Remark 5.7. The requirement for the plants to be linear was necessary in the above lemma to ensure that the statement made in equation 5.53 was legitimate. Without the linearity restriction there may be no such operator $\Psi_\Phi \in \mathcal{Q}^{l^{p,q}(\mathbb{N}\times[0,T])}_{\overline{P},x_1,\overline{P}_1,g_1}$ in equation 5.53 such that

$$\|(\Phi - I)(x - x_1)\|_{l^{p,q}(\mathbb{N}\times[0,T])} = \|(\Psi_\Phi - I)x - (\Psi_\Phi - I)x_1\|_{l^{p,q}(\mathbb{N}\times[0,T])}, \qquad (5.56)$$

since the operator that obeys the above equation may not be surjective and so wouldn't lie in $\mathcal{Q}^{l^{p,q}(\mathbb{N}\times[0,T])}_{\overline{P},x_1,\overline{P}_1,g_1}$. It is therefore hoped that the inheritance of surjectivity given in Proposition 5.5 can be expanded to nonlinear systems to enable the removal of the surjectivity condition from equation 5.41 and allow the biased and non-biased gaps to be compared for nonlinear systems.

Using the continuous version of the proof (with appropriately defined linear plants $\overline{P}$ and $\overline{P}_1$) it can also be shown that $\vec{\delta}^{L^{p,q}(\mathbb{N}\times[0,T])}_{x_1,g_1}(\overline{P},\overline{P}_1) = \vec{\delta}^{L^{p,q}(\mathbb{N}\times[0,T])}(\overline{P},\overline{P}_1)$.

The next lemma provides an inequality relating the 2D gap with the 1D gap on a finite time interval (the largest gap between plants along one trial).

Lemma 5.8. *For $p,q \in [1,\infty]$ and plants $P, P_1 : l^p_e[\mathbb{R}+] \to l^p_e[\mathbb{R}+]$ with*

$$\overline{P} : l^{p,q}_e(\mathbb{N}\times[0,T]) \to l^{p,q}_e(\mathbb{N}\times[0,T]) \quad : \quad u_1 \mapsto y_1, \quad y_1(k,\cdot) = Pu_1(k,\cdot)$$
$$\overline{P}_1 : l^{p,q}_e(\mathbb{N}\times[0,T]) \to l^{p,q}_e(\mathbb{N}\times[0,T]) \quad : \quad u_1 \mapsto y_1, \quad y_1(k,\cdot) = P_1u_1(k,\cdot) \qquad (5.57)$$

then
$$\vec{\delta}^{l^{p,q}(\mathbb{N}\times[0,T])}(\overline{P},\overline{P}_1) \leq \vec{\delta}^{l^p[0,T]}(P,P_1). \qquad (5.58)$$

Proof. Firstly, let $\Phi \in \mathcal{O}^{l^p[0,T]}_{P,P_1}$. We can clearly find an operator $\Psi \in \mathcal{O}^{l^{p,q}(\mathbb{N}\times[0,T])}_{\overline{P},\overline{P}_1}$ such that

$$(\Psi w)(k,t) = (\Phi w(k,\cdot))(t). \qquad (5.59)$$

We shall first detail the proof for $q < \infty$, which is similar to Theorem 3.1 of French (2008). We can state that there exists a Ψ obeying equation 5.59 such that the following

holds:

$$\vec{\delta}^{l^{p,q}(\mathbb{N}\times[0,T])}(\overline{P},\overline{P}_1)$$

$$\leq \sup_{\substack{w\in\mathcal{G}_{\overline{P}}^{l^{p,q}(\mathbb{N}\times[0,T])},\ \tau>0 \\ \|T_\tau w\|_{l^{p,q}(\mathbb{N}\times[0,T])}\neq 0}} \frac{\|T_\tau(\Psi-I)w\|_{l^{p,q}(\mathbb{N}\times[0,T])}}{\|T_\tau w\|_{l^{p,q}(\mathbb{N}\times[0,T])}}$$

$$= \sup_{\substack{w\in\mathcal{G}_{\overline{P}}^{l^{p,q}(\mathbb{N}\times[0,T])},\ \tau=(k,t)>0 \\ \|T_\tau w\|_{l^{p,q}(\mathbb{N}\times[0,T])}\neq 0}} \left(\frac{\displaystyle\sum_{i=0}^{k-1}\|(\Phi-I)w(i,\cdot)\|_{l^p[0,T]}^q + \|T_t(\Phi-I)w(k,\cdot)\|_{l^p[0,T]}^q}{\displaystyle\sum_{i=0}^{k-1}\|w(i,\cdot)\|_{l^p[0,T]}^q + \|T_t w(k,\cdot)\|_{l^p[0,T]}^q}\right)^{\frac{1}{q}}$$

$$\leq \sup_{\substack{w\in\mathcal{G}_{\overline{P}}^{l^{p,q}(\mathbb{N}\times[0,T])},\ \tau=(k,t)>0 \\ \|T_\tau w\|_{l^{p,q}(\mathbb{N}\times[0,T])}\neq 0 \\ 0<\tau'\leq T}} \left(\frac{\|T_{\tau'}(\Phi-I)|_{\mathcal{G}_P^{l^p[0,T]}}\|_{l^p[0,T]}^q\|T_\tau w\|_{l^{p,q}(\mathbb{N}\times[0,T])}^q}{\|T_\tau w\|_{l^{p,q}(\mathbb{N}\times[0,T])}^q}\right)^{\frac{1}{q}}$$

$$= \sup_{0<\tau\leq T}\|T_\tau(\Phi-I)|_{\mathcal{G}_P^{l^p[0,T]}}\|_{l^p[0,T]}. \tag{5.60}$$

Since the above holds for any $\Phi\in\mathcal{O}_{P,P_1}^{l^p[0,T]}$ we can therefore state that for $q<\infty$

$$\vec{\delta}^{l^{p,q}(\mathbb{N}\times[0,T])}(\overline{P},\overline{P}_1)\leq\vec{\delta}^{l^p[0,T]}(P,P_1). \tag{5.61}$$

Similarly, for $q=\infty$ there exists a Ψ obeying equation 5.59 such that the following holds:

$$\vec{\delta}^{l^{p,\infty}(\mathbb{N}\times[0,T])}(\overline{P},\overline{P}_1)$$

$$\leq \sup_{\substack{w\in\mathcal{G}_{\overline{P}}^{l^{p,\infty}(\mathbb{N}\times[0,T])},\ \tau>0 \\ \|T_\tau w\|_{l^{p,\infty}(\mathbb{N}\times[0,T])}\neq 0}} \frac{\|T_\tau(\Psi-I)w\|_{l^{p,\infty}(\mathbb{N}\times[0,T])}}{\|T_\tau w\|_{l^{p,\infty}(\mathbb{N}\times[0,T])}}$$

$$= \sup_{\substack{w\in\mathcal{G}_{\overline{P}}^{l^{p,\infty}(\mathbb{N}\times[0,T])},\ \tau=(k,t)>0 \\ \|T_\tau w\|_{l^{p,\infty}(\mathbb{N}\times[0,T])}\neq 0}} \frac{\displaystyle\max_{0<i\leq k-1}\left\{\|(\Phi-I)w(i,\cdot)\|_{l^p[0,T]}, \|T_t(\Phi-I)w(k,\cdot)\|_{l^p[0,T]}\right\}}{\displaystyle\max_{0<i\leq k-1}\left\{\|w(i,\cdot)\|_{l^p[0,T]}, \|T_t w(k,\cdot)\|_{l^p[0,T]}\right\}}$$

$$\leq \sup_{\substack{w\in\mathcal{G}_{\overline{P}}^{l^{p,\infty}(\mathbb{N}\times[0,T])},\ \tau=(k,t)>0 \\ \|T_\tau w\|_{l^{p,\infty}(\mathbb{N}\times[0,T])}\neq 0 \\ 0<\tau'\leq T}} \frac{\|T_{\tau'}(\Phi-I)|_{\mathcal{G}_P^{l^p[0,T]}}\|_{l^p[0,T]}\|T_\tau w\|_{l^{p,\infty}(\mathbb{N}\times[0,T])}}{\|T_\tau w\|_{l^{p,\infty}(\mathbb{N}\times[0,T])}}$$

$$= \sup_{0<\tau\leq T}\|T_\tau(\Phi-I)|_{\mathcal{G}_P^{l^p[0,T]}}\|_{l^p[0,T]}. \tag{5.62}$$

Again, the above holds for any $\Phi\in\mathcal{O}_{P,P_1}^{l^p[0,T]}$ therefore for $q=\infty$

$$\vec{\delta}^{l^{p,q}(\mathbb{N}\times[0,T])}(\overline{P},\overline{P}_1)\leq\vec{\delta}^{l^p[0,T]}(P,P_1). \qquad \square$$

As before, the lemma also holds in the continuous case (see French, 2008) and therefore we can also state that: $\vec{\delta}^{L^{p,q}(\mathbb{N}\times[0,T])}(\overline{P},\overline{P}_1) \leq \vec{\delta}^{L^p[0,T]}(P,P_1)$.

The following lemma states that the 1D gap on a finite time interval is less than or equal to the standard nonlinear gap on $\mathbb{R}+$.

Lemma 5.9. *For $p \in [1,\infty]$ and plants $P, P_1: l_e^p[\mathbb{R}+] \to l_e^p[\mathbb{R}+]$ where P is stabilisable:*

$$\vec{\delta}^{l^p[0,T]}(P,P_1) \leq \vec{\delta}^{l^p[\mathbb{R}+]}(P,P_1). \tag{5.63}$$

Proof. Again following French (2008), since P is stabilisable then for all $u \in \mathcal{G}_P^{l^p[0,T]}$ there exists $v \in \mathcal{G}_P^{l^p[\mathbb{R}+]}$ such that $v|_{[0,T]} = u$. Additionally, for all $u \in \mathcal{G}_P^{l^p[\mathbb{R}+]}$ there exists $v \in \mathcal{G}_P^{l^p[0,T]}$ such that $v|_{[0,T]} = u$. Therefore we can conclude that

$$\mathcal{G}_P^{l^p[\mathbb{R}+]}\Big|_{[0,T]} = \mathcal{G}_P^{l^p[0,T]}. \tag{5.64}$$

Since Φ is causal, for any $x, y \in \mathcal{G}_P^{l^p[\mathbb{R}+]}$ such that $x|_{[0,T]} = y|_{[0,T]}$ then $T_T\Phi x = T_T\Phi y$. We can therefore define $\Phi_T: l^p[\mathbb{R}+] \to l^p[0,T]$ such that

$$\Phi_T x = (\Phi x)|_{[0,T]} \tag{5.65}$$

It then follows that

$$
\begin{aligned}
\vec{\delta}^{l^p[\mathbb{R}+]}(P,P_1) &= \inf_{\Phi \in \mathcal{O}_{P,P_1}^{l^p[\mathbb{R}+]}} \sup_{\substack{x \in \mathcal{G}_P^{l^p[\mathbb{R}+]},\ t>0 \\ \|T_t x\|_{l^p[\mathbb{R}+]} \neq 0}} \left(\frac{\|T_t(\Phi - I)x\|_{l^p[\mathbb{R}+]}}{\|T_t x\|_{l^p[\mathbb{R}+]}} \right) \\
&\geq \inf_{\Phi \in \mathcal{O}_{P,P_1}^{l^p[\mathbb{R}+]}} \sup_{\substack{x \in \mathcal{G}_P^{l^p[\mathbb{R}+]}, \\ \|x|_{[0,T]}\|_{l^p[0,T]} \neq 0}} \left(\frac{\|(\Phi_T x - x|_{[0,T]}\|_{l^p[0,T]}}{\|x|_{[0,T]}\|_{l^p[0,T]}} \right) \\
&= \inf_{\Upsilon \in \mathcal{O}_{P,P_1}^{l^p[0,T]}} \sup_{\substack{x \in \mathcal{G}_P^{l^p[0,T]}, \\ \|x\|_{l^p[0,T]} \neq 0}} \left(\frac{\|(\Upsilon - I)x\|_{l^p[0,T]}}{\|x\|_{l^p[0,T]}} \right) \\
&= \vec{\delta}^{l^p[0,T]}(P,P_1). \qquad \square
\end{aligned}
$$

Theorem 5.10. *For $p \in [1,\infty]$, $q \in \{1,\infty\}$ and linear plants $P, P_1: l_e^p[\mathbb{R}+] \to l_e^p[\mathbb{R}+]$ where P is stabilisable and*

$$
\begin{aligned}
\overline{P}: l_e^{p,q}(\mathbb{N} \times [0,T]) &\to l_e^{p,q}(\mathbb{N} \times [0,T]) &&:&& u_1 \mapsto y_1, && y_1(k,\cdot) = Pu_1(k,\cdot) \\
\overline{P}_1: l_e^{p,q}(\mathbb{N} \times [0,T]) &\to l_e^{p,q}(\mathbb{N} \times [0,T]) &&:&& u_1 \mapsto y_1, && y_1(k,\cdot) = P_1 u_1(k,\cdot) \tag{5.66}
\end{aligned}
$$

we can state that

$$\vec{\delta}_{x_1,g_1}^{l^{p,q}(\mathbb{N}\times[0,T])}(\overline{P},\overline{P}_1) = \vec{\delta}^{l^{p,q}(\mathbb{N}\times[0,T])}(\overline{P},\overline{P}_1) \leq \vec{\delta}^{l^p[0,T]}(P,P_1) \leq \vec{\delta}^{l^p[\mathbb{R}+]}(P,P_1). \tag{5.67}$$

Proof. The proof is evident from Lemmata 5.6 to 5.9. □

5.4   A Biased Robust Stability Theorem for ILC

The thesis will now culminate in this section with a theorem on robust stability for iterative learning control. The theorem is based on the biased graph robust stability theorem (Theorem 4.17). The stability margin used in this result is taken from Theorem 5.1. The 1D gap is utilised, with Theorem 5.10 providing the required comparison to the 2D gap.

Theorem 5.11. *Let $q \in \{1, \infty\}$ and $p \in [1, \infty]$. Consider the linear plants given by $P, P_1 \colon l_e^p[\mathbb{R}+] \to l_e^p[\mathbb{R}+]$ and their respective 2D analogues:*

$$\overline{P} \colon l_e^{p,q}(\mathbb{N} \times [0,T]) \to l_e^{p,q}(\mathbb{N} \times [0,T]) \quad : \quad u_1 \mapsto y_1, \quad y_1(k, \cdot) = P u_1(k, \cdot)$$
$$\overline{P}_1 \colon l_e^{p,q}(\mathbb{N} \times [0,T]) \to l_e^{p,q}(\mathbb{N} \times [0,T]) \quad : \quad u_1 \mapsto y_1, \quad y_1(k, \cdot) = P_1 u_1(k, \cdot). \quad (5.68)$$

Consider the linear closed-loop system $[\overline{P}, C]$ given by Figure 5.2, and equations 5.11 and 5.14. Suppose $\|Q(I - LP)\|_{l^p[0,T]} < 1$. Let $\mathcal{W} = l^{p,q}(\mathbb{N} \times [0,T])$ and suppose $[\overline{P}_1, C]$ is globally well-posed. Define $x_0 = \left(\begin{smallmatrix} 0 \\ y_{ref} \end{smallmatrix} \right) \in \mathcal{W}_e$ and

$$b_{\overline{P},C} = \left(\left\| \left(\begin{smallmatrix} I \\ P \end{smallmatrix} \right) \right\|_{l^p[0,T]} \left(1 + \frac{\|QLP\|_{l^p[0,T]} + \|QL\|_{l^p[0,T]}}{1 - \|Q(I - LP)\|_{l^p[0,T]}} \right) \right)^{-1}. \quad (5.69)$$

Let $\mathcal{M} := \mathcal{G}_P^{l^p[\mathbb{R}+]}$ and $\mathcal{M}_1 := \mathcal{G}_{P_1}^{l^p[\mathbb{R}+]}$ denote the graphs of P and P_1 respectively, and let

$$\mathcal{O}_{P,P_1}^{l^p[\mathbb{R}+]} := \left\{ \Phi \colon \mathcal{G}_P^{l^p[\mathbb{R}+]} \to \mathcal{G}_{P_1}^{l^p[\mathbb{R}+]} \mid \Phi \text{ is causal, surjective and } \Phi(0) = 0 \right\}. \quad (5.70)$$

If there exists a $\Phi \in \mathcal{O}_{P,P_1}^{l^p[\mathbb{R}+]}$ such that $\|(\Phi - I)|_{\mathcal{M}}\|_{l^p[\mathbb{R}+]} < b_{\overline{P},C}$, then the stability of $[\overline{P}_1, C]$ on $\mathcal{W}$ with respect to $g_0 = (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$ is assured.

Proof. The proof follows from Theorems 5.2 and 5.10. □

5.5   The Convergence and Boundedness of g_1

Recall Section 4.4.4 in the previous chapter regarding bounds on the biases g_0 and g_1, where g_0 is the external bias around the perturbed closed-loop system and g_1 is the bias around the perturbed plant. These are given by $g_0 = (I + (\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}})x_0$ and $g_1 = \Phi\Pi\mathcal{M}//\mathcal{N}x_0 = \Pi\mathcal{M}_1//\mathcal{N}g_0$. In Section 4.4.4 these bounds were only given for the general case. Here we will exploit the structure of the ILC system described in this

chapter to extend this and provide a proof that the bias surrounding the perturbed plant converges and is bounded.

5.5.1 Convergence

First we shall consider the convergence of the signal g_1.

Proposition 5.12. *Let $p, q \in [1, \infty]$. Consider the linear plants $P, P_1 \colon l_e^p[\mathbb{R}+] \to l_e^p[\mathbb{R}+]$ and their respective 2D analogues:*

$$\overline{P} \colon l_e^{p,q}(\mathbb{N} \times [0, T]) \to l_e^{p,q}(\mathbb{N} \times [0, T]) \quad : \quad u_1 \mapsto y_1, \quad y_1(k, \cdot) = P u_1(k, \cdot)$$

$$\overline{P}_1 \colon l_e^{p,q}(\mathbb{N} \times [0, T]) \to l_e^{p,q}(\mathbb{N} \times [0, T]) \quad : \quad u_1 \mapsto y_1, \quad y_1(k, \cdot) = P_1 u_1(k, \cdot). \quad (5.71)$$

Consider the linear closed-loop system $[\overline{P}, C]$ given by Figure 5.2, and equations 5.11 and 5.14. Suppose $\|Q(I - LP)\|_{l^p[0,T]} < 1$ and $\Phi \in \mathcal{O}_{P,P_1}$ is chosen to be continuous. If $x_0(k, \cdot) \in l^p[0, T]$ is constant from trial-to-trial, i.e. $x_0(k_1, \cdot) = x_0(k_2, \cdot)$ for all $k \in \mathbb{N}$, then there exists a signal $g_1^* \in l^p[0, T]$ such that

$$g_1^*(\cdot) = \lim_{k \to \infty} (\Phi \Pi_{\mathcal{M}//\mathcal{N}} x_0)(k, \cdot) \tag{5.72}$$

Proof. First recall that the nominal system $[\overline{P}, C]$ is monotonically convergent since $\|Q(I - LP)\|_{l^p[0,T]} < 1$ and therefore also obeys the eigenvalue condition of convergence. Therefore there exists an $x_1^*(\cdot) \in l^p[0, T]$ such that:

$$x_1^*(\cdot) = \lim_{k \to \infty} (\Pi_{\mathcal{M}//\mathcal{N}} x_0)(k, \cdot) \tag{5.73}$$

Therefore, since Φ is continuous there clearly exists a $g_1^* \in l^p[0, T]$ such that:

$$g_1^*(\cdot) = \lim_{k \to \infty} (\Phi x_1(k, \cdot)). \qquad \qquad \square$$

Note that forcing Φ to be continuous is not particularly restrictive. Recall that the plants in question are linear and therefore in l^2 the map Φ is linear (see section 4.2.2 and also the appendix of Georgiou and Smith, 1997a). An investigation would be required to determine how restrictive this is under other spaces.

5.5.2 Boundedness

In order to examine the boundedness of the signals g_0 and g_1 we will restrict our analysis to $l^{p,\infty}(\mathbb{N} \times [0, T])$ and linear systems. This enables the use of Propositions 4.20 and

4.21 to give bounds on g_0 and g_1 by the following:

$$\|g_0\|_{\mathcal{W}} \leq (1 + \epsilon)\|x_0\|_{\mathcal{W}}$$
$$\|g_1\|_{\mathcal{W}} \leq \epsilon\|x_0\|_{\mathcal{W}} + \|\Pi_{\mathcal{M}//\mathcal{N}}x_0\|_{\mathcal{W}}. \tag{5.74}$$

Note that since we are now restricting to l^∞ the truncations have been removed (see equations 4.74 and 4.81). This also enables the ILC reference signal in the following proposition to be written as a signal in $\mathcal{W}$ rather than $\mathcal{W}_e$ and the operator $\Pi_{\mathcal{M}//\mathcal{N}}: \mathcal{W} \to \mathcal{W}$.

The bound given for g_0 in equation 5.74 is adequate since it only depends on ϵ and x_0. The following proposition will determine a bound on the internal signals g_1 using a property of $\Pi_{\mathcal{M}//\mathcal{N}}$ for a monotonically convergent ILC system..

Proposition 5.13. *Let $p \in [1, \infty]$ and $q = \infty$. Consider the linear plants given by $P, P_1: l_e^p[\mathbb{R}+] \to l_e^p[\mathbb{R}+]$ and their respective 2D analogues:*

$$\overline{P}: l_e^{p,q}(\mathbb{N} \times [0,T]) \to l_e^{p,q}(\mathbb{N} \times [0,T]) \quad : \quad u_1 \mapsto y_1, \quad y_1(k, \cdot) = Pu_1(k, \cdot)$$
$$\overline{P}_1: l_e^{p,q}(\mathbb{N} \times [0,T]) \to l_e^{p,q}(\mathbb{N} \times [0,T]) \quad : \quad u_1 \mapsto y_1, \quad y_1(k, \cdot) = P_1u_1(k, \cdot). \tag{5.75}$$

Consider the linear closed-loop system $[\overline{P}, C]$ given by Figure 5.2, and equations 5.11 and 5.14. Let $\mathcal{W} = l^{p,q}(\mathbb{N} \times [0,T])$. Suppose $\|Q(I - LP)\|_{l^p[0,T]} < 1$ and $x_0 = \begin{pmatrix} 0 \\ y_{ref} \end{pmatrix} \in \mathcal{W}$.

Let $\Pi_{\mathcal{M}//\mathcal{N}}: \mathcal{W} \to \mathcal{W}$ such that $\Pi_{\mathcal{M}//\mathcal{N}}: \begin{pmatrix} u_0 \\ y_0 \end{pmatrix} \mapsto \begin{pmatrix} u_1 \\ y_1 \end{pmatrix}$, where $\begin{pmatrix} u_0 \\ y_0 \end{pmatrix}$ and $\begin{pmatrix} u_1 \\ y_1 \end{pmatrix}$ obey $[\overline{P}, C]$, define $g_1 = \Phi\Pi_{\mathcal{M}//\mathcal{N}}x_0$ and let $\epsilon = \|(\Phi - I)\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}$.

Then
$$\|g_1\|_{\mathcal{W}} \leq \epsilon\|x_0\|_{\mathcal{W}} + 2\left\|\begin{matrix} I \\ P^{-1} \end{matrix}\right\|_{l^p[0,T]} \|y_{ref}\|_{l^p[0,T]}. \tag{5.76}$$

Proof. Let $x_0 = \begin{pmatrix} 0 \\ y_{ref} \end{pmatrix} \in \mathcal{W}$. From the monotonic convergence criteria, and since $u_0(0) = 0$, we can state that for all $k \in \mathbb{N}$:

$$\|y_1(k+1) - y_{ref}\|_{l^p[0,T]} \leq \|y_1(k) - y_{ref}\|_{l^p[0,T]}$$
$$\leq \|y_{ref}\|_{l^p[0,T]}. \tag{5.77}$$

Therefore $\|y_1(k)\|_{l^p[0,T]} \leq 2\|y_{ref}\|_{l^p[0,T]}$ for all $k \in \mathbb{N}$. Also, in the same fashion:

$$\|u_1(k+1) - P^{-1}y_{ref}\|_{l^p[0,T]} = \|P^{-1}y_1(k) - P^{-1}y_{ref}\|_{l^p[0,T]}$$
$$= \|P^{-1}\|_{l^p[0,T]}\|y_1(k) - y_{ref}\|_{l^p[0,T]}$$
$$\leq \|P^{-1}\|\|y_{ref}\|_{l^p[0,T]}. \tag{5.78}$$

Therefore $\|u_1(k)\| \leq 2\|P^{-1}\|\|y_{ref}\|$ for all $k \in \mathbb{N}$.

Since $\Pi_{\mathcal{M}//\mathcal{N}}x_0 = \begin{pmatrix} u_1 \\ y_1 \end{pmatrix}$ then for all $k \in \mathbb{N}$:

$$\|[\Pi_{\mathcal{M}//\mathcal{N}}x_0](k)\|_{l^p[0,T]} \le 2 \left\| \begin{matrix} I \\ P^{-1} \end{matrix} \right\| \|y_{ref}\|_{l^p[0,T]}, \tag{5.79}$$

and therefore:

$$\|\Pi_{\mathcal{M}//\mathcal{N}}x_0\|_{l^{p,\infty}(\mathbb{N}\times[0,T])} \le 2 \left\| \begin{matrix} I \\ P^{-1} \end{matrix} \right\| \|y_{ref}\|_{l^p[0,T]}. \tag{5.80}$$

From Proposition 4.21 we have

$$\|g_1\|_{\mathcal{W}} \le \epsilon\|x_0\|_{\mathcal{W}} + \|\Pi_{\mathcal{M}//\mathcal{N}}x_0\|_{\mathcal{W}}$$

$$\le \epsilon\|x_0\|_{\mathcal{W}} + 2 \left\| \begin{matrix} I \\ P^{-1} \end{matrix} \right\| \|y_{ref}\|. \qquad \square$$

Note that the requirement that the analysis is conducted in l^∞ is not considerably restrictive. Firstly, the signal spaces given by l^p are subsets of l^∞ and from this it is trivial to also show that $l^{p,q} \subset l^{p,\infty}$. Secondly, $l^{p,\infty}$ makes the ideal choice for examining the size of the reference trajectories since it provides the maximum l^p-norm of every trajectory; whereas norms such as $l^{p,2}$ would increase over successive iterations, blurring the bounds on g_0 and g_1 (see Figure 4.5). Using l^∞ prevents the manifestation of long term instability, since the divergence of signals to norms of 10^{51} observed in Longman (2000) would be revealed by the supremum in $l^{p,\infty}$.

The downside to restricting analysis to $l^{p,\infty}$ is that it may be desirable to determine robust stability results using Theorem 4.17 in other spaces, such as $l^{p,2}$. It would then be necessary to examine the boundedness of the reference trajectories in $l^{p,\infty}$. The robust stability theorem would provide a bound on $\epsilon = \|(\Phi - I)\|\|\Pi_{\mathcal{M}//\mathcal{N}}\| < 1$ in $l^{p,2}$ due to the requirement that the gap is less than the stability margin but the bound calculations would then require a bound on ϵ in $l^{p,\infty}$.

5.6   Summary

This chapter has extended work from Bradley and French (2009) concerning the application of the gap metric to iterative learning control. The existence of a non-zero stability margin for a broad class of ILC algorithms has also been shown.

The biased norm introduced in Section 4.3 has been used with a two-dimensional signal space to give a bound on the stability margin of an ILC system. The set of plants and controller for which this is applicable are those that can be explained by the lifted system representation given in Chapter 3. This has enabled the construction of a 2D biased graph robust stability theorem for iterative learning control.

A theorem was also given showing that the one-dimensional gap measure can be used in lieu of the two-dimensional measure when used with the 2D measure of stability margin, but only when considering linear systems. Since the 2D measure of gap would likely be a very difficult result to calculate, relating it back to the 1D gap is a powerful result enabling the theorem to be used with the well-established traditional gap measure.

The justification for the restriction to linear systems lies in the necessity of the gap operators to be surjective. This restriction to surjective operators could simply be defined in equation 5.41, but this may then invalidate the statement in Lemma 5.6 that there exists a $\Psi_\Phi \in \mathcal{Q}^{l^{p,q}(\mathbb{N}\times[0,T])}_{P,x_1,\overline{P}_1,g_1}$ such that

$$\|(\Phi - I)(x - x_1)\|_{l^{p,q}(\mathbb{N}\times[0,T])} = \|(\Psi_\Phi - I)x - (\Psi_\Phi - I)x_1\|_{l^{p,q}(\mathbb{N}\times[0,T])}, \tag{5.81}$$

since the operator Ψ_Φ that obeys the above equation may not be surjective. It is therefore hoped that a proof can be developed such that Proposition 5.5 holds for nonlinear systems. This would enable the 1D gap to be used with the 2D robust stability margin without restricting to linearity.

Finally, Section 5.5 provided a proof that the reference signal for the perturbed system converges provided the map Φ is uniformly continuous and gave bounds on both g_0 and g_1 in $l^{p,\infty}(\mathbb{N} \times [0,T])$.

Chapter 6

Interpretation

This chapter will apply the results established in the previous chapter to some ILC problems to highlight some issues within the field. This will provide a link between real ILC implementations and the abstract nature that makes up the majority of this thesis. The bulk of the chapter will take the form of analyses of two different aspects of iterative learning controllers with a view on robustness. The first topic covered will be the role of the filter Q for improving robustness of lifted system ILC and the trade-off between robustness and performance. The second shall be the use of inverse model ILC. Following this, a section will detail the relationship between the robust stability margin and the eigenvalue and monotonic convergence criteria.

Examples are drawn from other papers that are concerned with ILC robustness. Some simulations are also given in order to provide some quantitative results to back up the qualitative reasoning. These simulations will be based on data from the gantry robot mentioned in Section 3.8.2 (see the theses Ratcliffe, 2005 and Cai, 2009).

Before commencing the work on robustness the following section will introduce the gantry robot and discuss the development of twenty-first, first and seventh order models of it.

6.1   An Introduction to the Gantry Robot

The system examined in this chapter is a multi-axis gantry robot. The robot features in a large quantity of applied ILC work — in the PhD theses Ratcliffe (2005) and Cai (2009), and papers detailed in Section 3.8.2.

The robot's three axes allow an end-effector to move anywhere within a cuboid environment; however, this thesis will only examine the x-axis. It was manufactured by Aerotech Inc. USA and is the type of gantry robot commonly used in industry and

therefore a good choice for the analysis of ILC algorithms. The end-effector consists of an electromagnet to enable the machine to 'pick-and-place' ferromagnetic materials. The robot is shown in Figure 6.1.

FIGURE 6.1: A photograph of the gantry robot

The x-axis has an effective travel of 640mm controlled by a motor. This is powered by an amplifier providing a ± 10 Volt control signal which determines the velocity and direction of the motor. The plant's output is considered to be the displacement (in metres) of the end-effector and the input is the velocity control signal (in Volts).

The plant is first modelled and simplified using data from Ratcliffe (2005). It will only be used as an example of a physical ILC system in order to apply some of the work developed earlier. As such the control mechanism and specific features of the plant are ignored; all that is examined is its frequency response data.

This section will start by detailing the method of modelling the system and its conversion into state space form. It will then be described in the lifted system form to allow it to be related to the robustness work from earlier chapters. Three different models of the gantry robot will be examined. The first is a twenty-first order model that is a very close approximation of the robot. The second is a first order model that is a simplification of the twenty-first order model that loses some of the accuracy at high frequencies. The final model was derived from the plant after it was moved to a different location. This model (seventh order) therefore describes the same plant but differs more substantially from the other two. The inclusion of the seventh order model provides a cogent set of results with regards to robustness, since algorithms can be developed based on measured data from the other models and then implemented on the seventh order model. Since

they all describe the same underlying plant a robust algorithm should be able to control all three.

6.1.1 Twenty-first order model

The plant is treated as a 'black box' structure — the internal dynamics of the plant are ignored and the only concern is the relationship between inputs and outputs. It is also considered to be linear, permitting the use of frequency domain analysis using transfer functions.

In order to model the system sine waves of differing frequencies are applied at the input and the gain and phase shift for each of these is measured at the output. The results can then be used to construct a Bode plot which is then used to develop a transfer function description of the plant. This was first approximated manually and then refined using a least-mean-square optimisation (Ratcliffe, 2005). Figure 6.2 shows the Bode plot for the x-axis, as measured by frequency response tests, along with the Bode plot for a 21st order transfer function that approximates it, given in equation 6.1. (The figure also shows the plant model in discrete-time, which will be detailed shortly.)
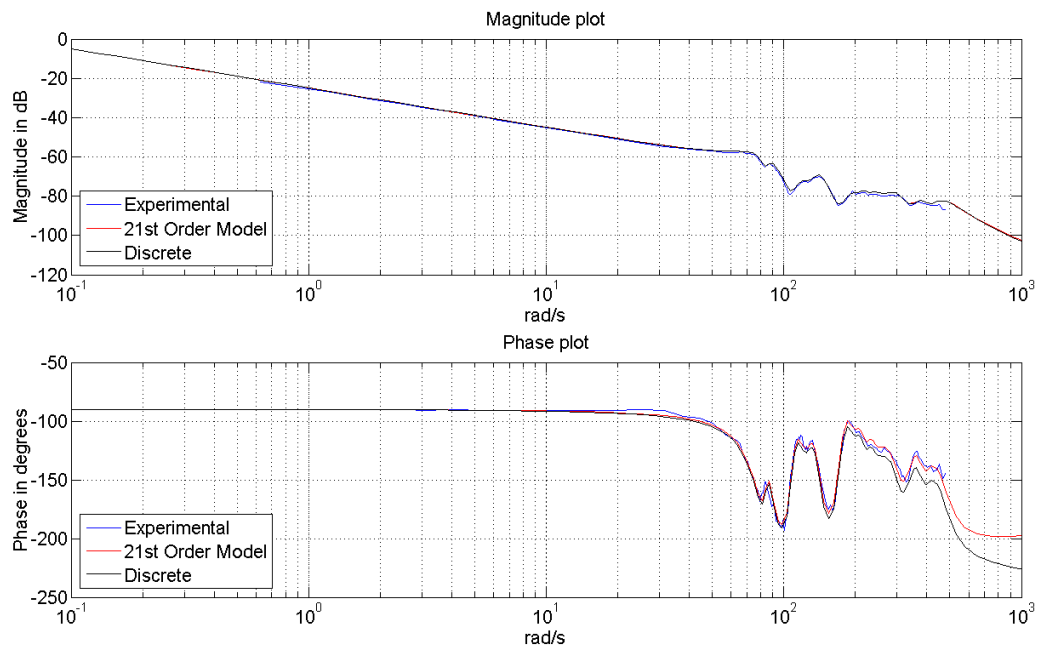

FIGURE 6.2: 21st order Bode plot for the x-axis

$$G(s)^{21} = \frac{0.0000157274(s + 3.64 \times 10^5)(s + 500.21)(s + 35.88 \pm i400.76)}{s(s + 69.73 \pm i459.64)(s + 30.53 \pm i378.98)} \cdots$$

$$\cdots \frac{(s + 30.46 \pm i336.23)(s + 27.47 \pm i249.8)(s + 10.78 \pm i223.56)}{(s + 32.59 \pm i297.99)(s + 21.14 \pm i239.36)(s + 10.64 \pm i220.02)} \cdots$$

$$\cdots \frac{(s + 14.06 \pm i195.13)(s + 10.59 \pm i169.44)(s + 8.83 \pm i124.71)}{(s + 10.67 \pm i192.24)(s + 10.45 \pm i141.63)(s + 8.51 \pm i119.75)} \cdots$$

$$\cdots \frac{(s + 5.33 \pm i106.87)(s + 3.36 \pm i83.93)}{(s + 6.05 \pm i86.78)(s + 12.02 \pm i79.09)} \quad (6.1)$$

As is clear from the Bode plot, the frequency response of the 21st order model matches the plant well in the frequency range 10^{-1} to 10^3 rad/s. This should therefore provide an accurate model. However, a 21st order model is cumbersome and so the model order is reduced to aid the design of suitable controllers. Further details of the modelling of the plant are given in Ratcliffe (2005).

6.1.2   First order model

The 21st order model is now reduced to a 1st order one. This model is taken from Ratcliffe (2005) and is formulated to approximate the low frequency gain of the system. As such some of the high frequency effects are lost in the reduction. A Bode plot for the 1st order model is given in Figure 6.3 (along with its discrete-time representation) and the transfer function is given in equation 6.2.
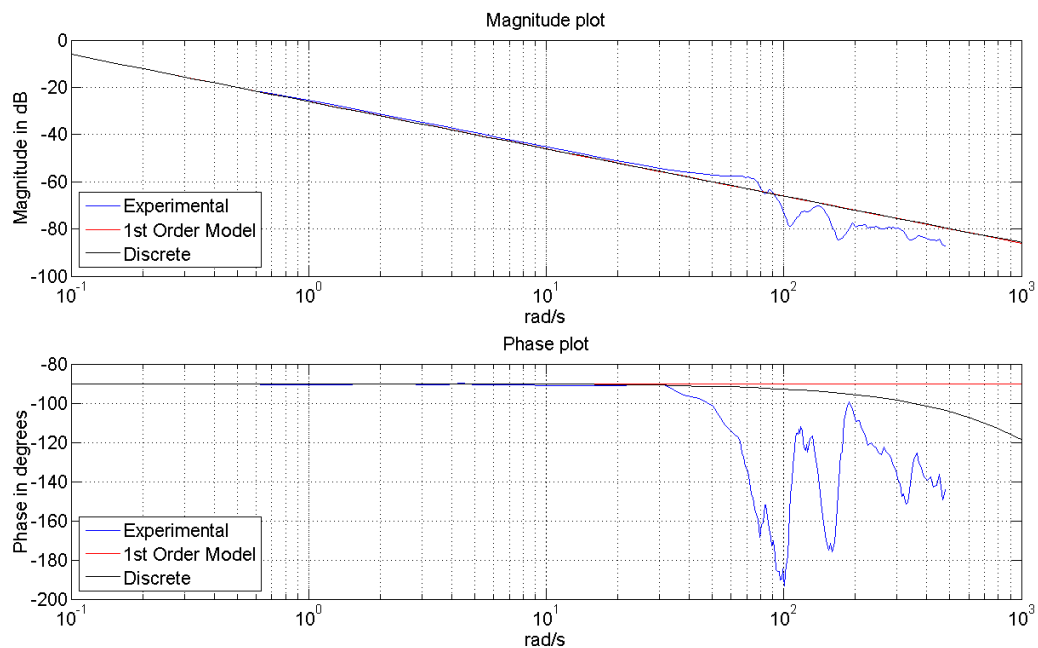

FIGURE 6.3: 1st order Bode plot for the x-axis

$$G(s)^1 = \frac{0.05}{s} \tag{6.2}$$

This 21$^{\text{st}}$ order description will be regarded as the actual plant and a 1$^{\text{st}}$ order approximation to it will be regarded as the nominal plant model. Since the 1$^{\text{st}}$ order model is a good approximation of the 21$^{\text{st}}$ order system it would be expected that a controller able to control one would also be able to control the other. We can therefore examine robustness properties by designing a controller for the 1$^{\text{st}}$ order model and then testing it on the 21$^{\text{st}}$ order model.

The gap between the 1$^{\text{st}}$ and 21$^{\text{st}}$ order models is calculated as 0.0634. (This was calculated using MATLAB's *μ-Analysis and Synthesis Toolbox* — see Balas et al., 1998.)

6.1.3   Seventh order model

The gantry robot has presented a useful opportunity for the examination of the robustness of ILC. Between the theses Ratcliffe (2005) and Cai (2009) the entire apparatus was moved to a different location and remodelled. This has therefore provided two different sets of data for the same plant, the original 21$^{\text{st}}$ and 1$^{\text{st}}$ order models given above (from Ratcliffe, 2005) and a new 7$^{\text{th}}$ order model (from Cai, 2009). This therefore provides an opportunity to examine the performance of algorithms developed for the first model and implemented on the new one.

A Bode plot of the 7$^{\text{th}}$ order model of Cai (2009) along with the experimental measurements is given in Figure 6.4, and the transfer function is given by:

$$G(s)^7 = \frac{13077183.4436(s + 113.4)}{s(s^2 + 61.57s + 1.125 \times 10^4)} \cdots$$
$$\cdots \frac{(s^2 + 30.28s + 2.13 \times 10^4)}{(s^2 + 227.9s + 5.647 \times 10^4)(s^2 + 466.1s + 6.142 \times 10^5)}. \tag{6.3}$$

The gap between the 1$^{\text{st}}$ and 7$^{\text{th}}$ order models is calculated as 0.2364. It is therefore clear that the 1$^{\text{st}}$ order model is a much better description of the plant as given by the old model than the new. This is not surprising, since it was based on a reduction of the old model. It would therefore be expected that a controller developed using the 1$^{\text{st}}$ order model would fare better on the 21$^{\text{st}}$ order model than it would on the 7$^{\text{th}}$ order model.
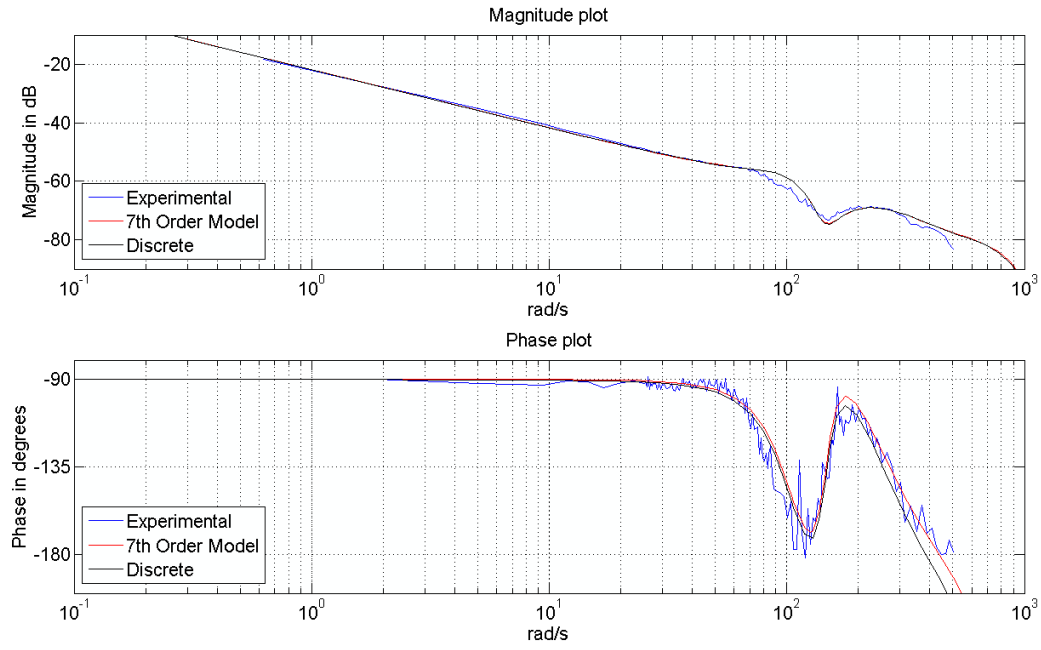
FIGURE 6.4: 7$^{\text{th}}$ order Bode plot for the x-axis

6.1.4 Lifted plant models

As ILC requires the use of data from (at least) a previous trial the use of memory is required for storage. This ultimately leads to a discretisation at some stage. Generally the entire ILC algorithm is operated in discrete-time leaving the plant as the only continuous-time object being analysed. It therefore becomes appropriate to regard the plant as a discrete mapping and transfer it into the lifted system formulation described in Section 3.3.1. In order to do this the plant is first discretised at a sampling rate of 1kHz using a zero-order hold. The 21$^{\text{st}}$ order model now becomes a 21$^{\text{st}}$ order discrete-time relative degree one model in state-space form which can be easily transferred into the lifted system representation by computing the Markov parameters and populating the lifted system matrix:

$$
P = \begin{pmatrix}
0 & 0 & 0 & 0 & \cdots \\
CB & 0 & 0 & 0 & \cdots \\
CAB & CB & 0 & 0 & \cdots \\
CA^2B & CAB & CB & 0 & \cdots \\
\vdots & \vdots & \vdots & \vdots & \ddots
\end{pmatrix}.
\tag{6.4}
$$

This 21$^{\text{st}}$ order lifted model is denoted P^{21}, with the first few columns and rows of the matrix shown in equation 6.5.

$$P^{21} = 1 \times 10^{-4} \begin{pmatrix} 0 & 0 & 0 & 0 & \cdots \\ 0.0318 & 0 & 0 & 0 & \cdots \\ 0.1034 & 0.0318 & 0 & 0 & \cdots \\ 0.1787 & 0.1034 & 0.0318 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{6.5}$$

The 1$^{\text{st}}$ order discrete-time model is also computed from the continuous time 1$^{\text{st}}$ order model. The above matrix P is then populated as before, this time with all elements $CA^iB = 0.5 \times 10^{-4}$ for all $i \geq 0$. This 1$^{\text{st}}$ order plant model matrix is denoted P^1:

$$P^1 = 1 \times 10^{-4} \begin{pmatrix} 0 & 0 & 0 & 0 & \cdots \\ 0.5 & 0 & 0 & 0 & \cdots \\ 0.5 & 0.5 & 0 & 0 & \cdots \\ 0.5 & 0.5 & 0.5 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{6.6}$$

Similarly, the 7$^{\text{th}}$ order model is given by:

$$P^7 = 1 \times 10^{-4} \begin{pmatrix} 0 & 0 & 0 & 0 & \cdots \\ 0.0047 & 0 & 0 & 0 & \cdots \\ 0.0588 & 0.0047 & 0 & 0 & \cdots \\ 0.1985 & 0.0588 & 0.0047 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{6.7}$$

The discrete-time systems for the 21$^{\text{st}}$, 1$^{\text{st}}$ and 7$^{\text{th}}$ order models are also depicted in Figures 6.2, 6.3 and 6.4 respectively.

The reference trajectory that will be examined is a single pick and place task that is 2 seconds long (shown in Figure 6.5). (The trajectory is the same as that used in Ratcliffe, 2005.) At the 1kHz sample rate P becomes a 2000×2000 matrix.

6.2   The Role of the Filter Q

This section covers the robustness and performance issues relating to the filter Q. It is often not possible to achieve the performance criteria and also maintain an appropriate level of robustness with only the learning matrix in the lifted system ILC formulation. In order to make the algorithm more robust exact convergence may have to be sacrificed. An increase in robustness can be accomplished by implementing a filter into the learning loop.
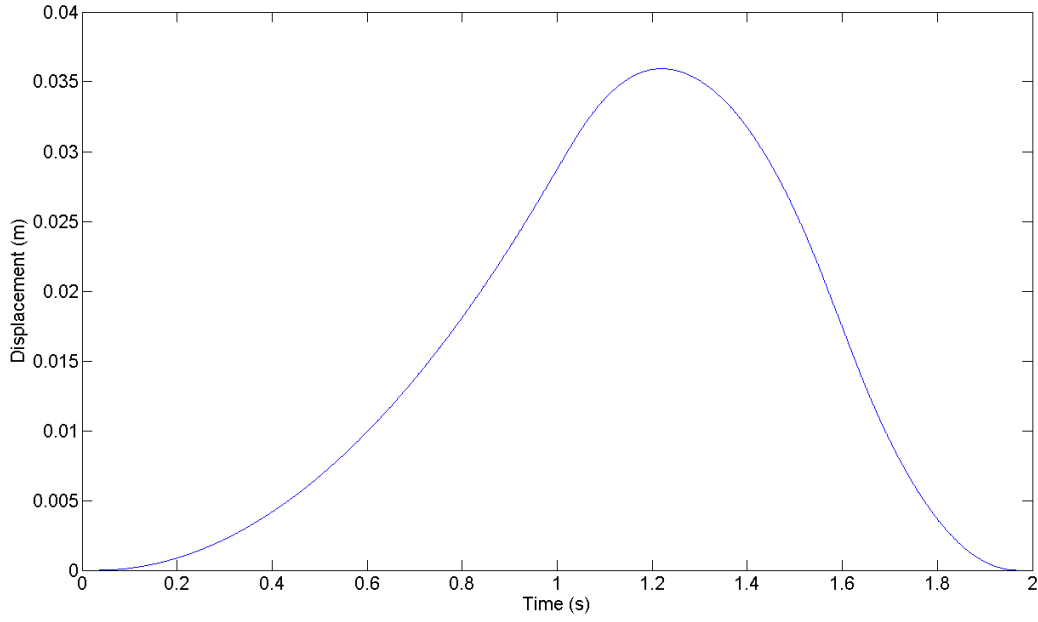
FIGURE 6.5: Reference trajectory for the x-axis

This section will first demonstrate the effect that the filter has on the final convergence of the ILC system and on the stability margin. An example will then be drawn from the gantry robot. This example will show an IL controller that is able to stabilise the simple 1st order model but is not capable of stabilising the 21st order model. A filter will then be introduced into the IL controller to illustrate it can extend the stability margin around the 1st order plant and controller such that the 21st order system is convergent. The same will also be shown for the 7th order model using identical controllers.

6.2.1 The trade-off between robustness and performance

Restricting ourselves to the single-input-single-output case, a plant is given by

$$
\begin{aligned}
x(t+1) &= Ax(t) + Bu(t) \\
y(t) &= Cx(t) + Du(t),
\end{aligned}
\tag{6.8}
$$

with $u(t), y(t) \in \mathbb{R}$, $x(t) \in \mathbb{R}^m$ and A, B, C and D of appropriate dimensions. We again re-write this plant as a matrix mapping $u(\cdot)$ to $y(\cdot)$, on a finite time-span $[0, N]$:

$$
\begin{pmatrix}
y(0) \\
y(1) \\
y(2) \\
\vdots \\
y(N)
\end{pmatrix}
=
\begin{pmatrix}
D & 0 & 0 & \cdots & 0 \\
CB & D & 0 & \cdots & 0 \\
CAB & CB & D & \cdots & 0 \\
\vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\
CA^{N-1}B & CA^{N-2}B & CA^{N-3}B & \cdots & D
\end{pmatrix}
\begin{pmatrix}
u(0) \\
u(1) \\
u(2) \\
\vdots \\
u(N)
\end{pmatrix}.
\tag{6.9}
$$

By writing out the desired reference trajectory as $y_{ref}(\cdot)$ we establish our IL controller as

$$u_{k+1}(\cdot) = Q(u_k(\cdot) + Le_k(\cdot)), \tag{6.10}$$

where $e_k(\cdot) = y_{ref}(\cdot) - y_k(\cdot)$. As explained in Section 3.3.1, the learning matrix L determines how the error is fed back to produce the following input and the filter Q is used to tune the asymptotic properties of the system.

We will now calculate the plant's output trajectory as the number of iterations increase. By combining the plant and controller equations we arrive at

$$u_{k+1}(\cdot) = Q(I - LP)u_k(\cdot) + QLy_{ref}(\cdot). \tag{6.11}$$

From this we can derive an expression for the limit of the plant's output trajectory as the iteration number tends to infinity in the face of zero disturbances:

$$y_\infty = \lim_{k \to \infty} P \sum_{i=0}^{k} [Q(I - LP)]^i (QLy_{ref}) \tag{6.12}$$

This limit can then be calculated to obtain the trajectory to which the plant converges. Let $S_k = \sum_{i=0}^{k} [Q(I - LP)]^i$. Then

$$
\begin{aligned}
(I - Q(I - LP))S_k &= \sum_{i=0}^{k} [Q(I - LP)]^i - Q(I - LP) \sum_{i=0}^{k} [Q(I - LP)]^i \\
&= \sum_{i=0}^{k} \left([Q(I - LP)]^i - [Q(I - LP)]^{i+1} \right) \\
&= I - [Q(I - LP)]^{k+1}
\end{aligned}
\tag{6.13}
$$

Hence:

$$S_k = (I - Q(I - LP))^{-1}(I - [Q(I - LP)]^{k+1}) \tag{6.14}$$

If we examine an asymptotically stable system where $|\lambda_i(Q(I - LP))| < 1$, where $\{\lambda_0, \ldots, \lambda_{N-n}\}$ is the set of all eigenvalues of $Q(I - LP)$ (see Theorem 3.4) it follows that $[Q(I - LP)]^{k+1} \to 0$ as $k \to \infty$. Substituting this into the equation for y_∞ gives

$$y_\infty = P(I - Q(I - LP))^{-1}(QLy_{ref}). \tag{6.15}$$

It can be clearly seen from this that setting Q to I allows the plant's output to converge exactly to the reference signal ($y_\infty = y_{ref}$). Using a non-zero Q can lead to a degradation in performance since the final output trajectory may not equal the desired trajectory. (It should be noted that, since we have only used the eigenvalue condition here, the matrix $[Q(I - LP)]^{k+1}$ may become very large for $k < \infty$ as explained in Section 3.3.3.)

There are two drawbacks to setting Q to the identity in the controller. Firstly, assuming that it is desired for the closed-loop system to obey the stronger condition of monotonic convergence (Theorem 3.5), it is required that $\|(Q(I - LP))\|_{l^p[0,T]} < 1$. This places some strong constraints on the choice of L. Secondly, Q has a significant bearing on robustness, as can be gathered from equation 5.69 regarding the robust stability margin.

It is remarked in Ahn et al. (2005) that, in terms of interval uncertainties, the maximum level of robustness is achieved with a learning gain of zero (this was briefly discussed in Section 3.7). The same result can be established from the approach of this thesis. Consider the robust stability margin:

$$b_{P,C} = \left(\left\| \begin{pmatrix} I \\ P \end{pmatrix} \right\|_{l^p[0,T]} \left(1 + \frac{\|QLP\|_{l^p[0,T]} + \|QL\|_{l^p[0,T]}}{1 - \|Q(I - LP)\|_{l^p[0,T]}} \right) \right)^{-1}. \tag{6.16}$$

In order to maximise $b_{P,C}$ the fraction in the above equation must be minimised. It is straightforward to observe that $b_{P,C}$ is minimised with respect to Q by setting $Q = 0$. It is also clear that this results in no learning taking place and, assuming a first-iteration plant input of zero, the plant output will always be zero. This can be deduced from equation 6.15.

These results regarding Q also echo work in Section 3.7, where de Roover (1996) and de Roover and Bosgra (2000) were discussed with regards to robustness in the $\mathcal{H}^\infty$ framework. The condition is given that Q must equal I in order for exact convergence, and the advocation of Q as a filter to increase robustness.

The role of Q has therefore been clearly expounded: with Q equal to the identity exact convergence to the reference signal is possible, and with Q small we can gain an improvement in robustness (see also Norrlöf and Gunnarsson, 2002, Theorem 9). An example of how to choose Q would therefore be to set it to attenuate high frequency error components. In this case Q would approximate I at low frequencies and fall-off at higher frequencies. This would allow the desired convergence to zero error at low frequencies and a higher robustness at high frequencies, where the model would be more prone to inaccuracies. The use of filtering for this purpose was discussed in Section 3.6.4.

6.2.2 Robustifying the D-type algorithm on the gantry robot

The assertion in the previous section will now be demonstrated more concretely using the gantry robot as an example. Consider the plant to be given by the 1st order model P^1. The controller will be in the lifted system form of

$$u_{k+1}(\cdot) = Q(u_k(\cdot) + Le_k(\cdot)), \tag{6.17}$$

with $e_k(\cdot) = y_{ref}(\cdot) - y_k(\cdot)$. Since it represents a relative degree one plant all elements of P^1 on and above the major diagonal are zero. The appropriate controller to test first

is therefore the D-type algorithm with the learning matrix:

$$L = 10 \begin{pmatrix} 0 & 1 & 0 & \cdots \\ 0 & 0 & 1 & \cdots \\ 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{6.18}$$

(based on equation 3.25 given earlier). A learning rate of 10 was chosen by experimentation. At first we will also set Q to the identity. Since we are examining a trajectory of 2s at 1000kHz the matrices P^1, L and Q are all 2000×2000.

At first we examine the nominal plant model P^1. The following two figures show this algorithm in operation: Figure 6.6 shows the output trajectory along iterations 5, 10 and 1000 along with the reference trajectory; and Figure 6.7 shows the 2-norm of the error along each iteration.

The algorithm clearly converges (at least for the first 1000 iterations) and so appears to be a good candidate to implement on the 21st order plant. In this case, however, the result is not as satisfactory. Figure 6.8 shows the output along iterations 5, 10 and 100 — the same iterations as Figure 6.6 for the 1st order case. For iterations 5 and 10 the 21st order system appears to boast roughly the same output as the 1st order but then carries a substantial divergence at higher iterations. What could be regarded a small perturbation between the two models leads to a considerable difference in result.

In order to counter this effect it is therefore necessary to make the algorithm more robust. As has been noted earlier, Q can be used to accomplish this. We will therefore maintain the same learning matrix and introduce a Q not equal to the identity. In order to clearly demonstrate the convergence to a non-zero error here it is simply set to the identity multiplied by a scalar gain.

Figure 6.10 shows the output trajectories at 5, 10 and 1000 iterations for the 21st order plant with a D-type controller with Q reduced to 0.99. The algorithm now appears to converge. Figure 6.11 shows the error norm per iteration, no longer exhibiting the divergence found in the case where $Q = 1$.

The drawback is that the algorithm no longer converges to zero error. The robustness has been increased such that the 21st order plant is stable based on a controller designed for the 1st order plant, but this has come with a reduction in asymptotic performance.

The same simulations are also given for the 7th order model: Figure 6.12 shows the plant's output along trials 5, 10 and 1000 for the 7th order model with the D-type algorithm and Figure 6.13 shows the error norm per iteration; and Figures 6.14 and 6.15 show the same but with $Q = 0.99$.

In order to achieve near-perfect tracking and also an acceptable level of robustness it is
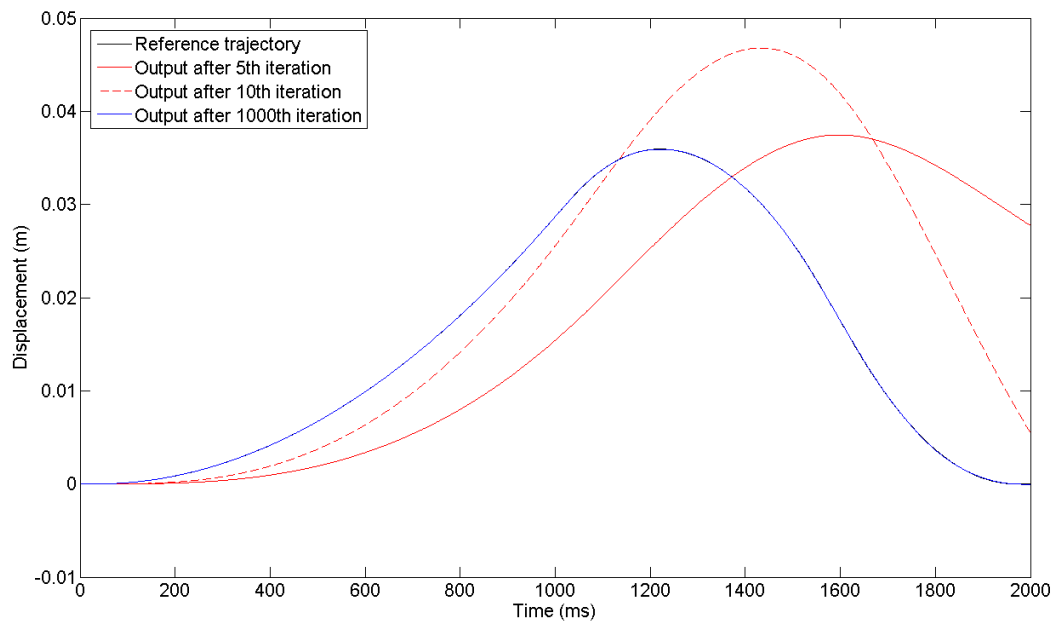
FIGURE 6.6: Output trajectory along trials 5, 10 and 1000 for the 1st order model with a D-type algorithm. (The reference and 1000th iteration output coincide.)

FIGURE 6.7: Output error per iteration for the 1st order model with a D-type algorithm

FIGURE 6.8: Output trajectory along trials 5, 10 and 1000 for the 21^{st} order model with a D-type algorithm

FIGURE 6.9: Output error per iteration for the 21^{st} order model with a D-type algorithm

FIGURE 6.10: Output trajectory along trials 5, 10 and 1000 for the 21^{st} order model with a D-type algorithm with $Q = 0.99$

FIGURE 6.11: Output error per iteration for the 21^{st} order model with a D-type algorithm with $Q = 0.99$

FIGURE 6.12: Output trajectory along trials 5, 10 and 1000 for the 7^{th} order model with a D-type algorithm
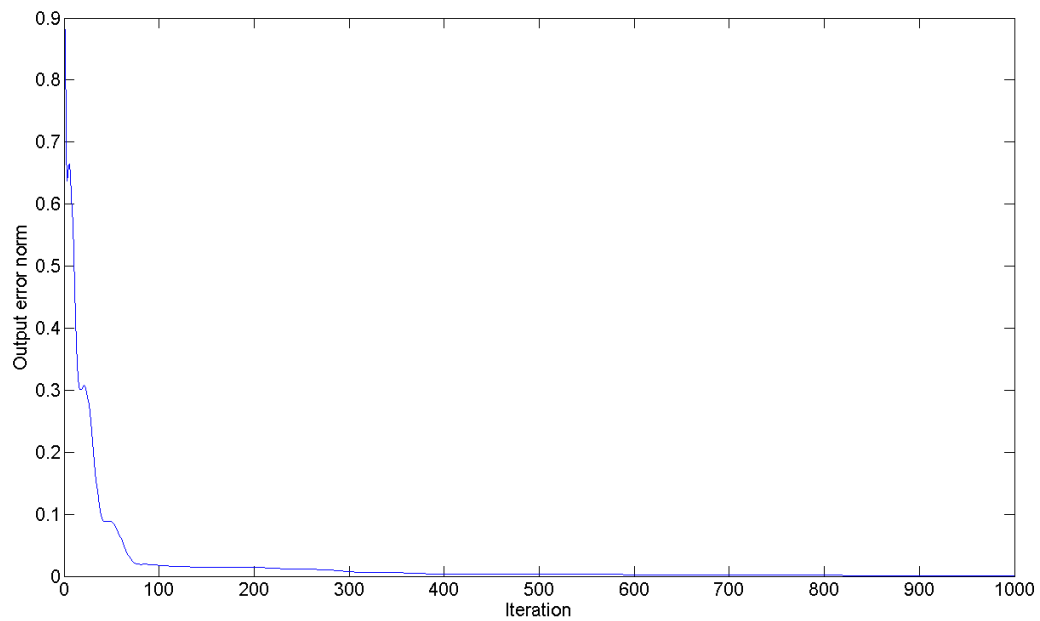

FIGURE 6.13: Output error per iteration for the 7^{th} order model with a D-type algorithm

FIGURE 6.14: Output trajectory along trials 5, 10 and 1000 for the 7^{th} order model with a D-type algorithm with $Q = 0.99$
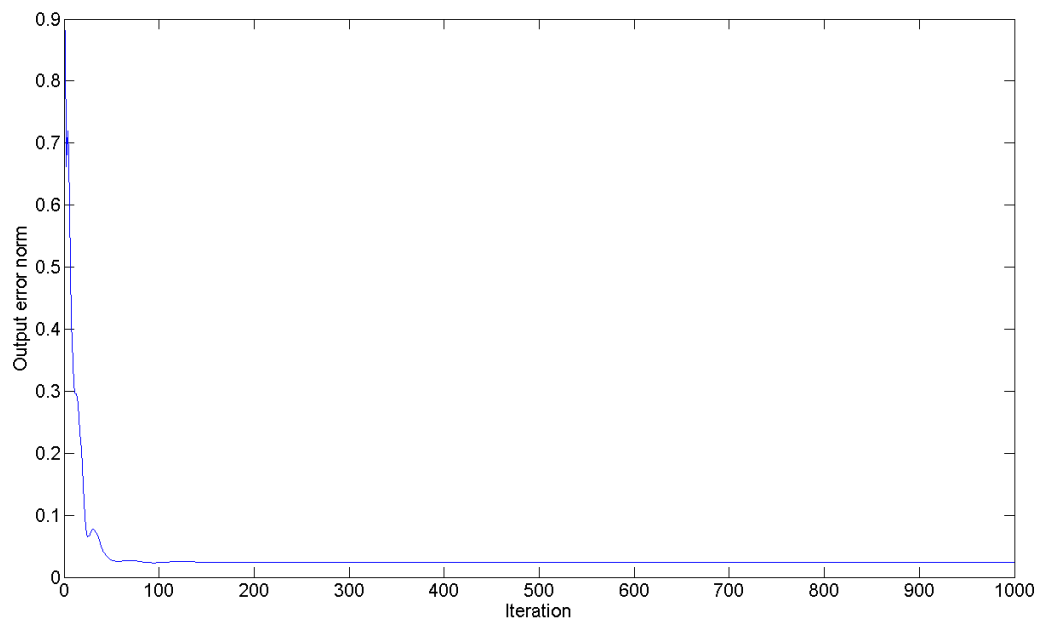
FIGURE 6.15: Output error per iteration for the 7^{th} order model with a D-type algorithm with $Q = 0.99$

here necessary to implement Q as a filter rather than a constant. In this case a low-pass filter would enable the algorithm to maintain tracking at low frequencies and increase the robustness at higher frequencies. In order not to add phase shift 'non-causal' filtering can be used. This is possible with ILC applications since the entire signal along the previous trial is available. (See Ratcliffe et al., 2005b.)

Here we will examine the case where Q is based on a second order Butterworth filter. In order to cancel the shifts in phase this filter will be operated twice on each iteration. The first will run forwards in the standard fashion; and the second will be reversed in order to cancel the shift in phase. Here a cut off frequency of 80rad/s has been chosen by examining the Bode plot of Figure 6.3, since at this frequency the phase difference between the 1st order model and the experimental results is approaching 180°. A Bode plot showing the response of this second order Butterworth filter is shown in Figure 6.16.
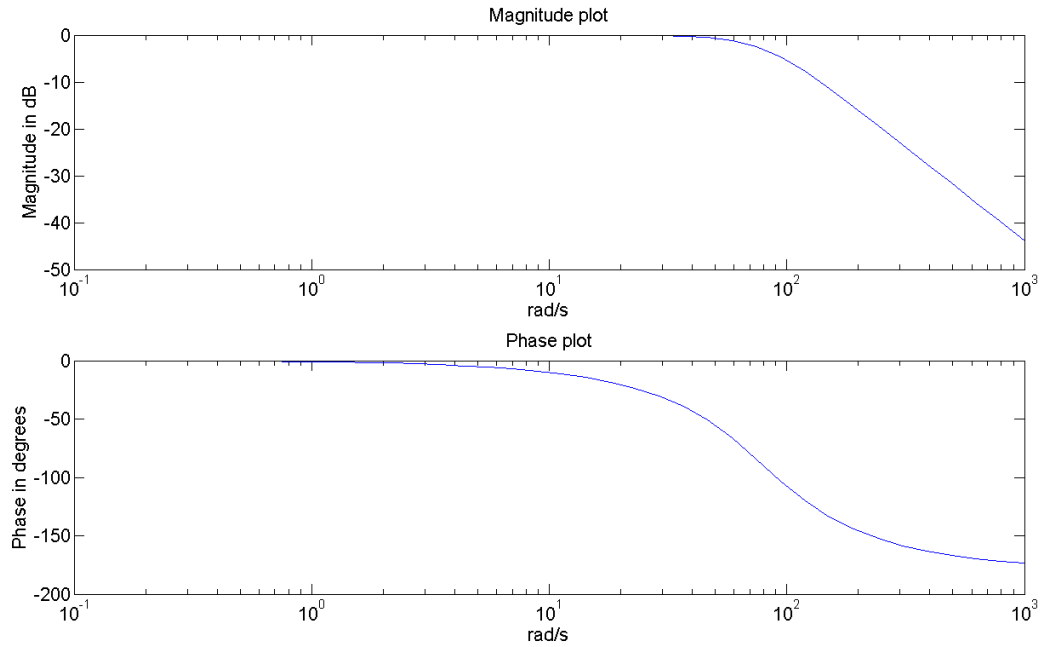
FIGURE 6.16: Butterworth filter Bode plot

This filter is first discretised and then a lifted system matrix $\bar{Q}$ is generated from its Markov parameters. A matrix I_r is also defined that is an identity matrix that has been reflected horizontally. This matrix serves to reverse a vector that it is applied to. The matrix Q is then given by $Q = I_r \bar{Q} I_r \bar{Q}$. This matrix Q can therefore be likened to two second order Butterworth filters with zero phase shift.

Figure 6.17 shows a simulation of the 21st order plant model with a the D-type algorithm, using this filter Q and the same learning gain of 10 used previously. Figure 6.18 shows the error norm per iteration. Figures 6.19 and 6.20 show the same for the 7th order plant model.
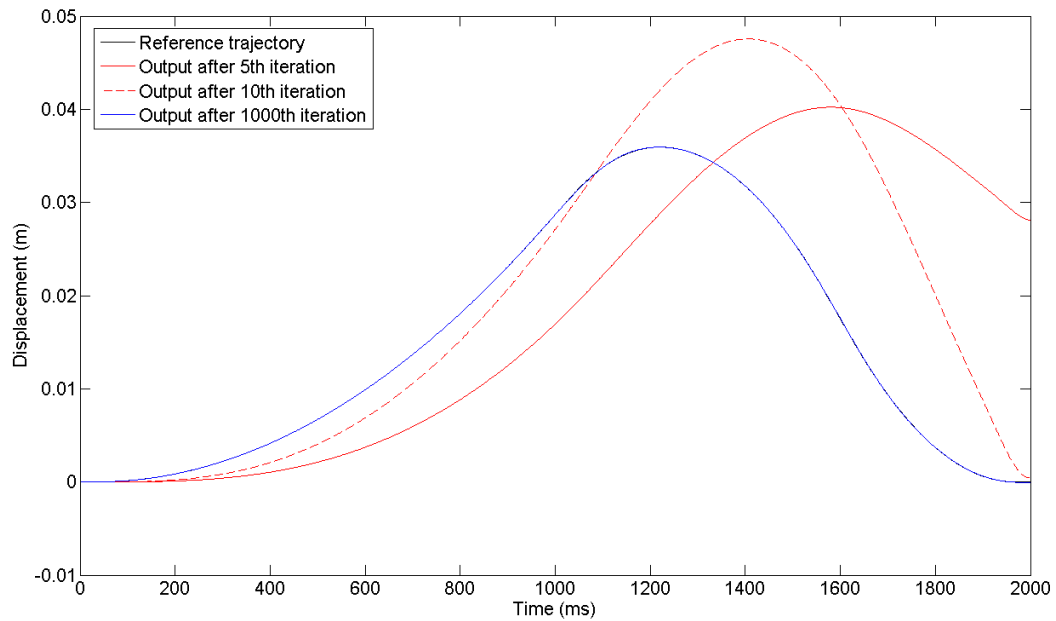
FIGURE 6.17: Output trajectory along trials 5, 10 and 1000 for the 21^{st} order model with a D-type algorithm with Q as a zero-phase-shift Butterworth filter. (The reference and 1000^{th} iteration output coincide.)
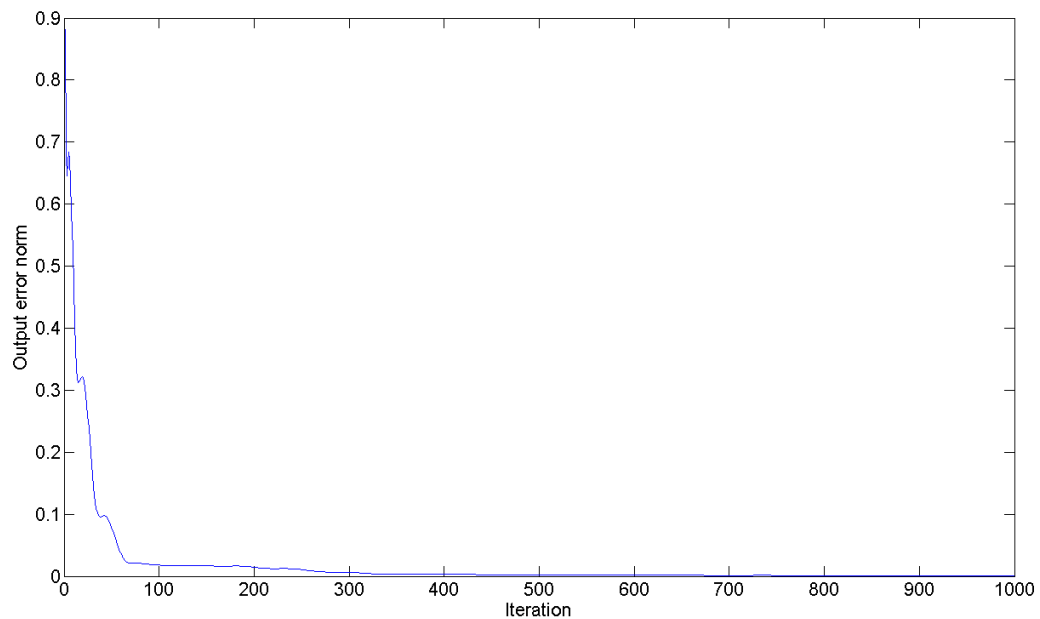
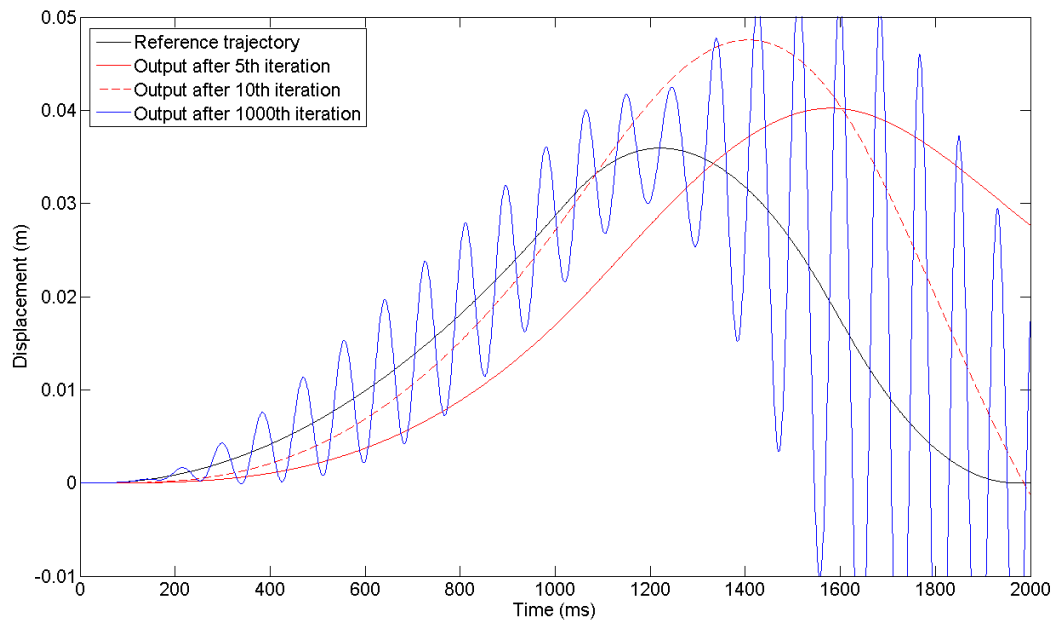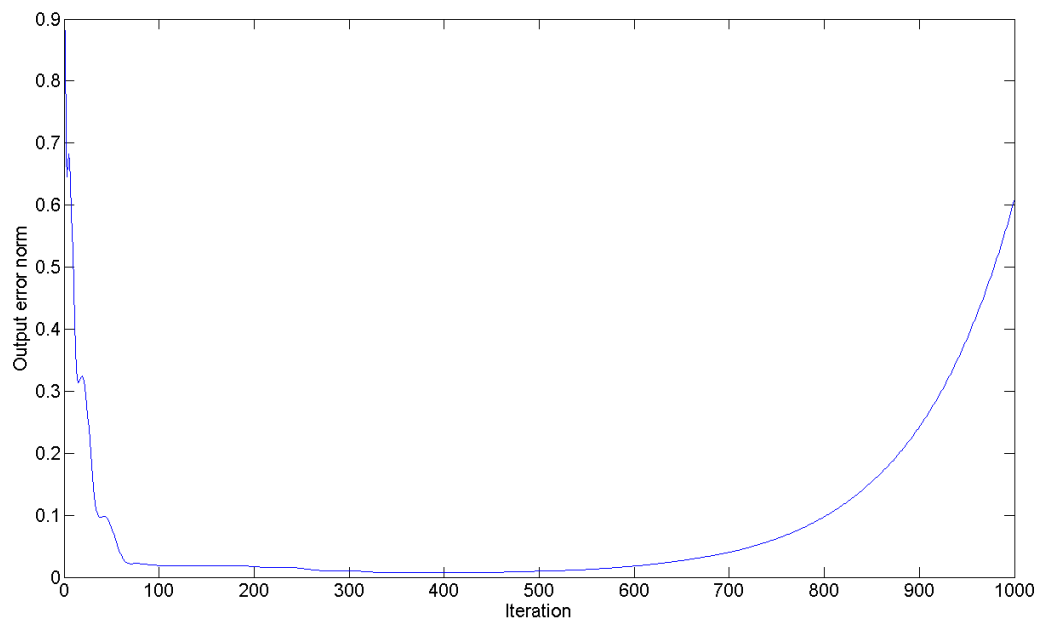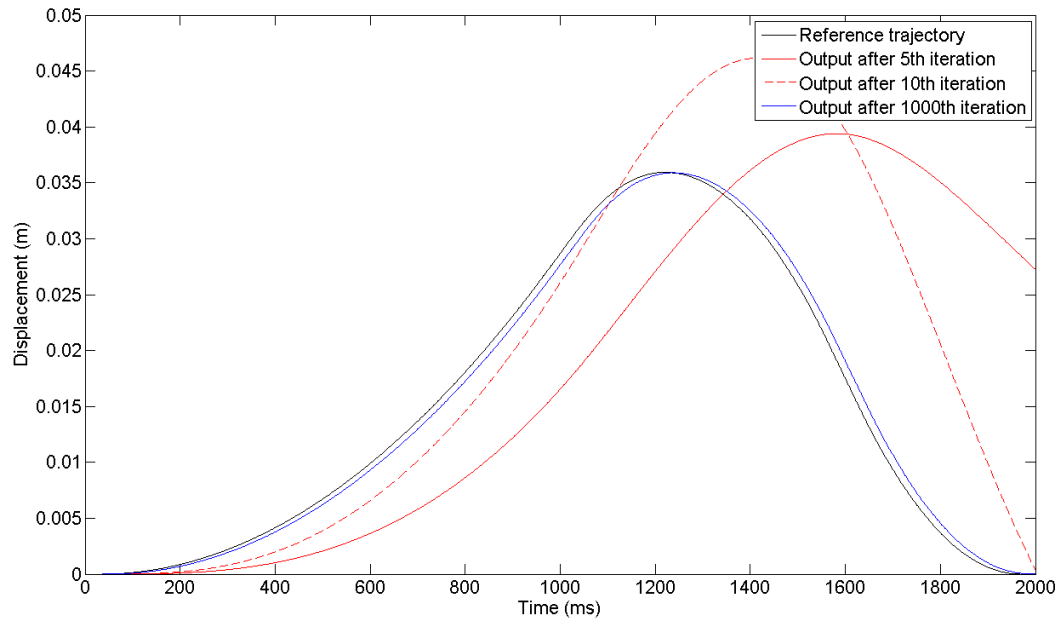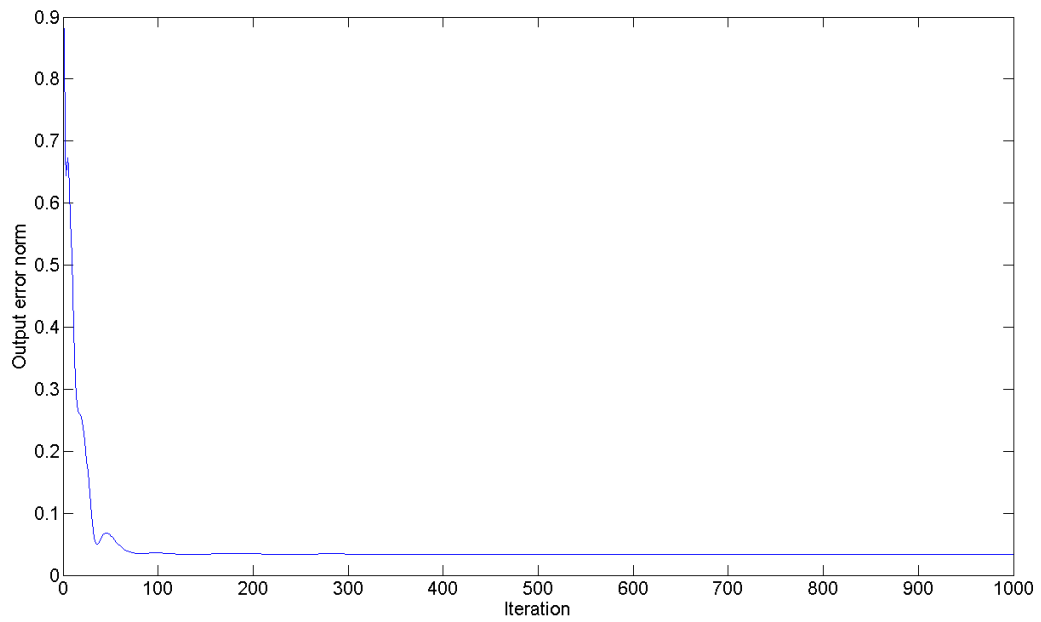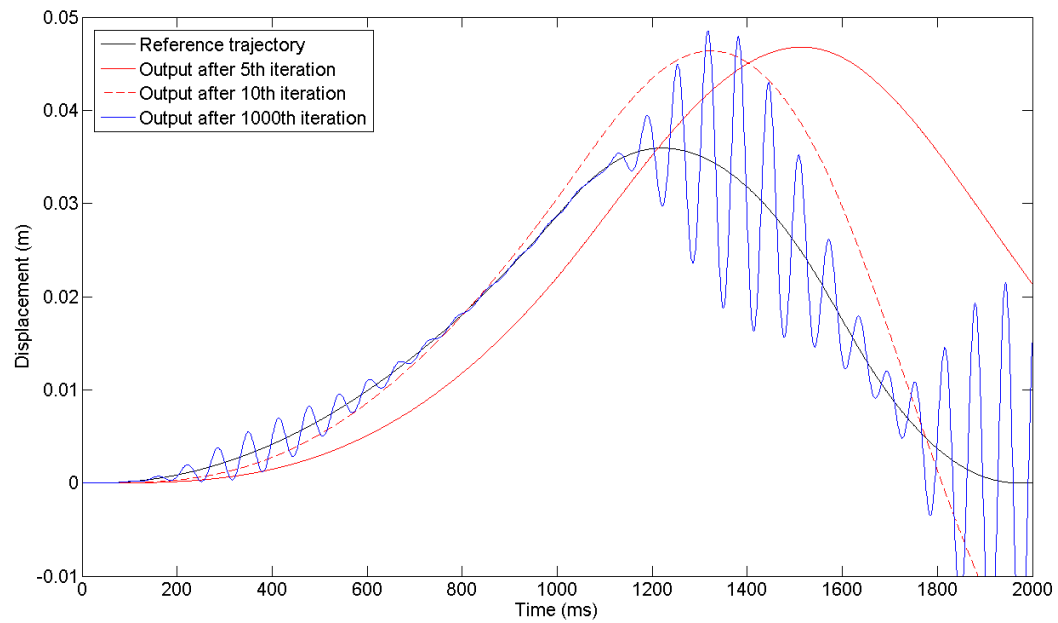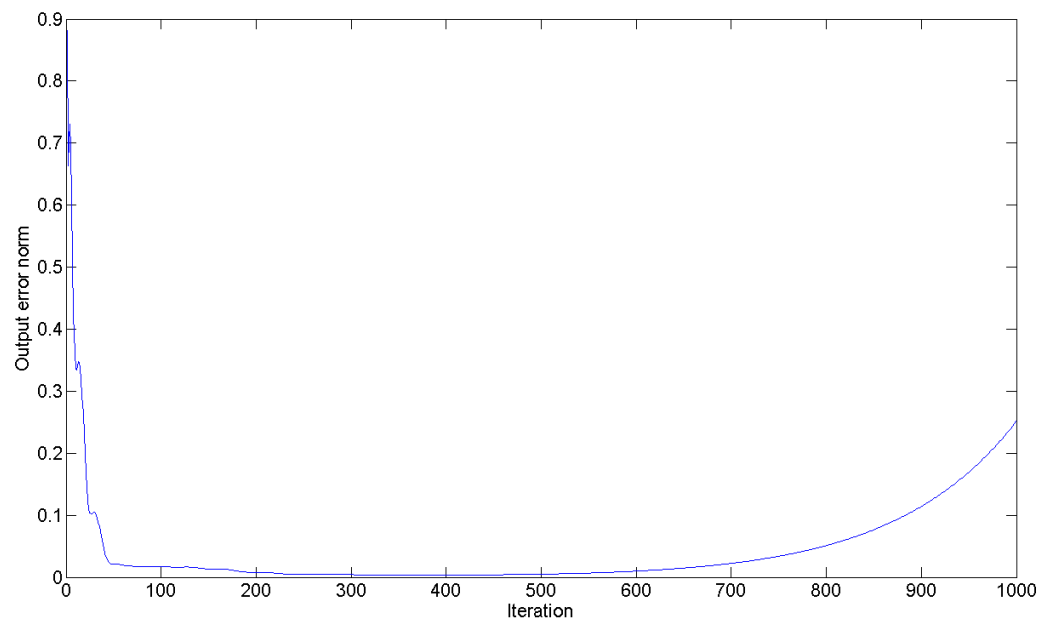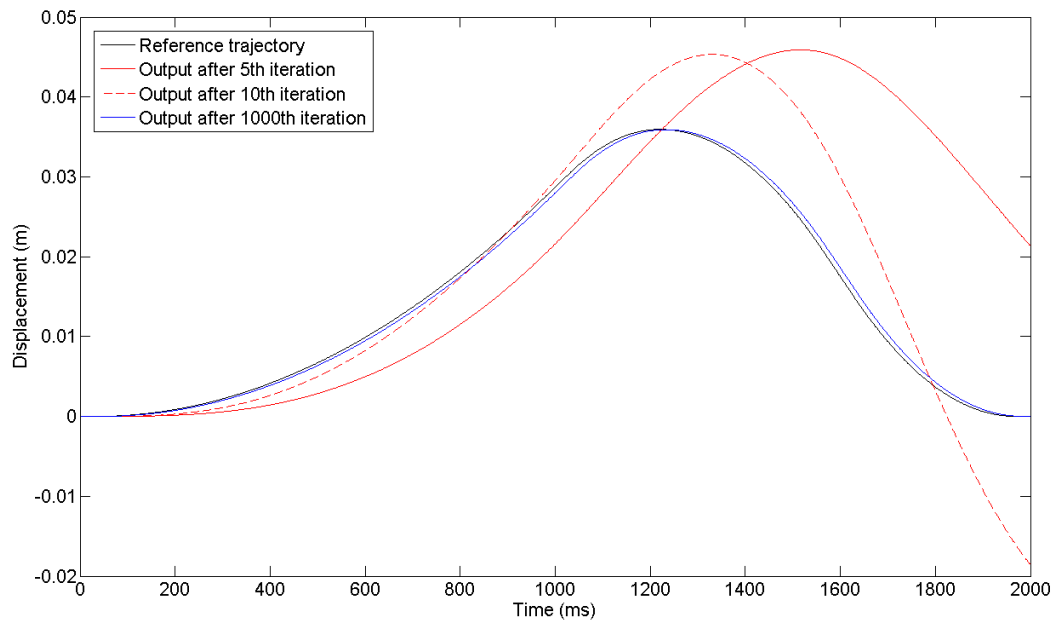FIGURE 6.18: Output error per iteration for the 21^{st} order model with a D-type algorithm with Q as a zero-phase-shift Butterworth filter
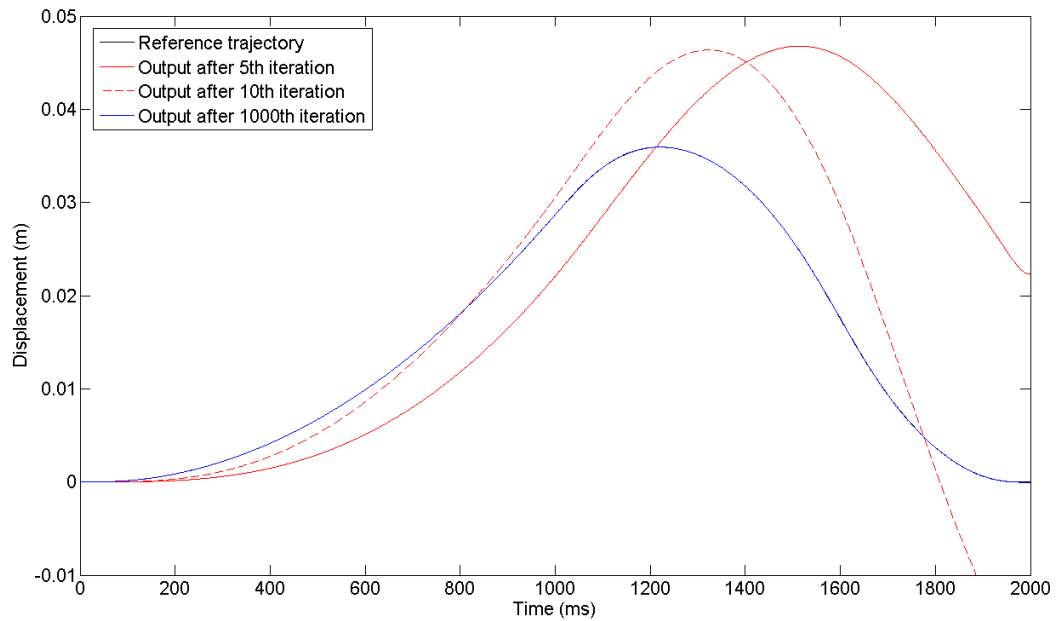
FIGURE 6.19: Output trajectory along trials 5, 10 and 1000 for the 7^{th} order model with a D-type algorithm with Q as a zero-phase-shift Butterworth filter. (The reference and 1000^{th} iteration output coincide.)
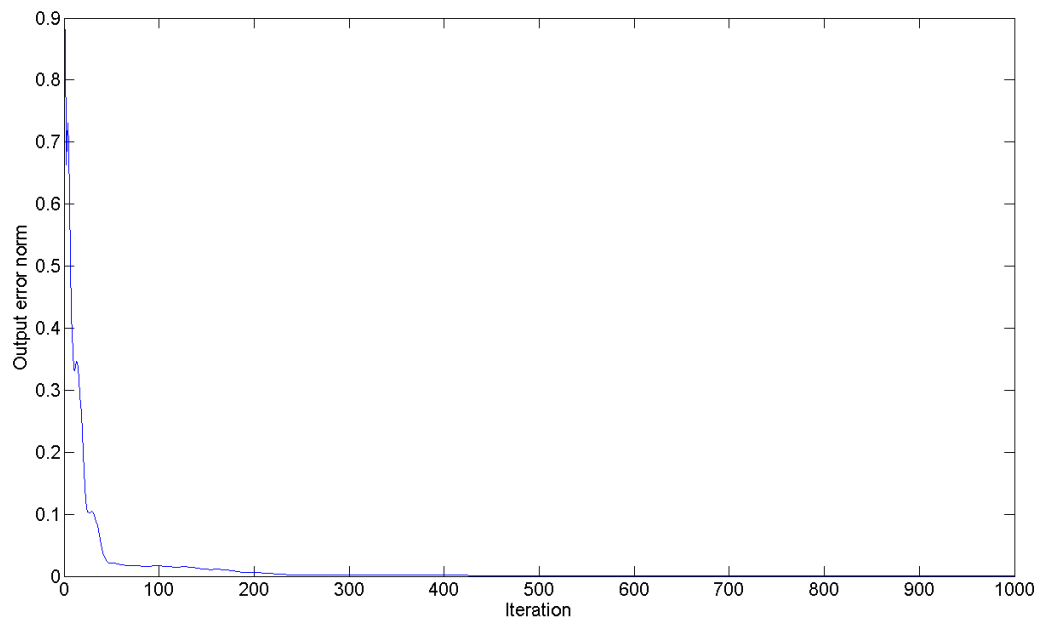
FIGURE 6.20: Output error per iteration for the 7^{th} order model with a D-type algorithm with Q as a zero-phase-shift Butterworth filter

Implementing Q as a Butterworth filter clearly provides convergence to a smaller error than for $Q = 0.99$ and is more robust to plant uncertainty than the case where $Q = I$.

6.3 Inverse Model ILC

As detailed in Section 3.4.3, model-based ILC uses an implementation of the plant's inverse in the controller in order to update the plant's next input trajectory, since the ideal input trajectory for the plant to achieve the output y_{ref} is given by $P^{-1}y_{ref}$. As explained previously however, this does lead to a degradation in robustness.

Looking now at the lifted system representation again. If the learning matrix L is given by P^{-1} and Q by the identity then very good results can be achieved as $Q(I - LP)$ would be zero. This clearly meets the monotonic convergence criteria $\|Q(I - LP)\|_{l^p[0,T]} < 1$. If we assume an input of zero along the first trial, producing zero output, the plant's input along the second trial would be given by $u_2 = Q(u_1 + Le_1) = P^{-1}y_{ref}$ — a convergence in a single iteration.

A major weakness in this approach is demonstrated by examination of the definition of $b_{P,C}$ (given again in equation 6.19). With Q as the identity, the numerator of the fraction contains two terms: $\|LP\|_{l^p[0,T]}$ and $\|L\|_{l^p[0,T]}$. Since we have set $L = P^{-1}$ exactly, the first term will equal one. However, the second term only contains P^{-1}, and it is a reasonable assumption that, in most instances, the norm of the inverse of P would be large. This would lead to a very small stability margin.

$$b_{P,C} = \left(\|\left(\begin{smallmatrix} I \\ P \end{smallmatrix}\right)\|_{l^p[0,T]} \left(1 + \frac{\|QLP\|_{l^p[0,T]} + \|QL\|_{l^p[0,T]}}{1 - \|Q(I - LP)\|_{l^p[0,T]}} \right) \right)^{-1} \qquad (6.19)$$

Clearly, matching the controller too tightly to the plant can result in a severe degradation of robustness to plant uncertainty. This leads to precisely the same conclusion that appears when considering filtering — attaining perfect tracking (this time in a single iteration) leads to a reduction in robustness.

Again similar results are found in other ILC robustness papers. De Roover (1996) and de Roover and Bosgra (2000) suggest a design procedure where a frequency ω is selected, below which the plant is accurately modelled. The learning matrix is then chosen to be the inverse of the plant up to ω and Q as a low pass filter, acting as the identity up to ω and approximating zero beyond. The papers point out that ω should be chosen to be as large as possible to guarantee tracking for the maximum frequency range.

Harte et al. (2005) discusses an inverse model ILC strategy, where the learning matrix is the inverse of the plant model and $Q = I$. The paper then provides a condition that the plant uncertainty must be contained within a phase shift of $\pm 90°$ over all frequencies

for the actual plant to be stabilised by this controller. This has been explained in more detail in Section 3.7.

6.3.1 Inverse model ILC on the gantry robot

To demonstrate the lack of robustness that can be inherent in inverse model ILC this section will again present an example based on the gantry robot. It is assumed that the 1st order model is the modelled system and the 21st order model is akin to the actual plant. The controller will therefore be based on the 1st order nominal model.

As explained in Section 6.1.4 the plant is relative degree one and so direct computation of the inverse $(P^1)^{-1}$ is not possible. In order to develop the inverse of the model matrix we therefore remove the top line and right-most column of the plant matrix P^1 and then find the inverse[1]. All elements are then shifted up one row within the matrix (by multiplication of the D-type learning matrix of equation 3.25) to produce the inverse model learning matrix L, the first few rows and columns shown here:

$$L = \begin{pmatrix} -20000 & 20000 & 0 & 0 & \cdots \\ 0 & -20000 & 20000 & 0 & \cdots \\ 0 & 0 & -20000 & 20000 & \cdots \\ 0 & 0 & 0 & -20000 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \tag{6.20}$$

This matrix will be regarded as the inverse of P^1 since $G_1 L P G_1^T = I$, with G_n given in equation 3.28.

Similarly to the analysis of Q in the previous section we will first look at the nominal case. The ILC system is therefore given by the plant P^1 and the lifted system controller of equation 6.10 with L given above and Q as the identity. A simulation of this algorithm produces the expected result that exact convergence to the reference trajectory is achieved in a single iteration (graph not given).

The monotonic convergence criteria is clearly met since $\|Q(I - LP)\|_{l^p[0,T]} = 0$. The robust stability margin of equation 6.19 however reveals a problem. As explained in the previous section, in this case the term $\|QL\|_{l^p[0,T]}$ in the denominator has a considerable effect on the robustness. The l^2-norm of L is given as 4×10^4. This results in a tiny stability margin that can be demonstrated by implementing the same algorithm on the 21st order plant.

Figure 6.21 shows several output trajectories for the 21st order plant and inverse model controller (still based on the inverse of the 1st order plant). Although the output after

---

[1]In order to produce a matrix L that was 2000×2000 as required a 2001×2001 matrix of the plant P^1 was used.

iteration 5 appears to have already converged, by the 50^{th} iteration the divergence of the output is obvious. Figure 6.22 shows how the error evolves for the first 100 iterations. The result is a clear example of the long term instability mentioned in Section 3.6 — for the first few iterations convergence seems to be certain but after many trials where the error is close to zero the system suddenly diverges.

It is apparent from the Bode plot of Figures 6.3 and 6.2 that the discrete 1^{st} and 21^{st} order models differ by more than 90° and so the condition of Harte et al. (2005) that the shift is within ±90° over all frequencies is clearly not satisfied.

The same simulation is now carried out on the 7^{th} order model. Figure 6.23 shows the output for iterations 1, 5 and 15 and Figure 6.24 shows the error norm for the first 25 iterations. As can be clearly seen in both figures the error divergence occurs much sooner for the 7^{th} order model and is also much faster. This is not surprising since the 1^{st} order model on which the controller is based was derived from the 21^{st} order model, which differs significantly from the 7^{th} order. Even though all the models are based on precisely the same plant they differ substantially, and the model-based controller is simply not robust enough to stabilise the perturbed plants.

To demonstrate the lack of robustness afforded by this inverse model controller compare the stability margin to the gap between the 1^{st} and 21^{st} order models. The gap between the plants was given in Section 6.1.2 as 0.0634. With the gain of order 10^4, as in this case, to guarantee stability using the theorem provided in this thesis would require Q to be of the order of 10^{-3} (if Q were the identity multiplied by a constant).

In order to improve this we here return to the second order Butterworth filter of Section 6.2.2, with a cut-off of 80rad/s. Implementing the same inverse model algorithm, but now with Q as the zero-phase-shift Butterworth filter provides improved results for both the 21^{st} and 7^{th} order plants. Figure 6.25 shows the output trajectories along trials 5, 10 and 1000 for this algorithm. Without the filter Q the trajectories were only given for the first 50 iterations due to the algorithms divergence. In this case the algorithm has effectively converged by the 5^{th} iteration and remains converged beyond the 1000^{th} (the error norm per iteration is shown in Figure 6.26).

Figures 6.27 and 6.28 show the same algorithm's operation on the 7^{th} order model. Although the convergence is not as fast it still produces small errors by the 10^{th} iteration which remain small beyond the 1000^{th}. The filter has therefore managed to dramatically improve the robustness of the inverse model algorithm, especially considering the gap between the 7^{th} and 1^{st} order models.

Compared to the D-type algorithm, the inverse model approach exhibits much faster convergence; however, without the addition of the filter the cost of this is a large reduction of robustness. The introduction of the filter appears to provide fast convergence for the perturbed plants, with this convergence also close to the desired output.
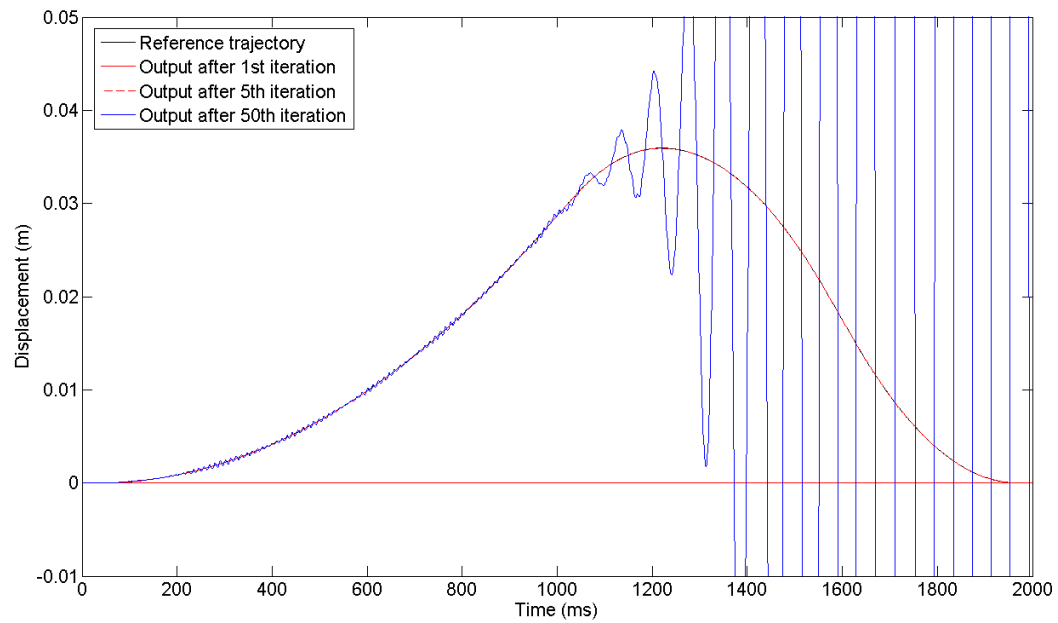
FIGURE 6.21: Output trajectory along trials 1, 5 and 50 for the 21st order model with an inverse model algorithm
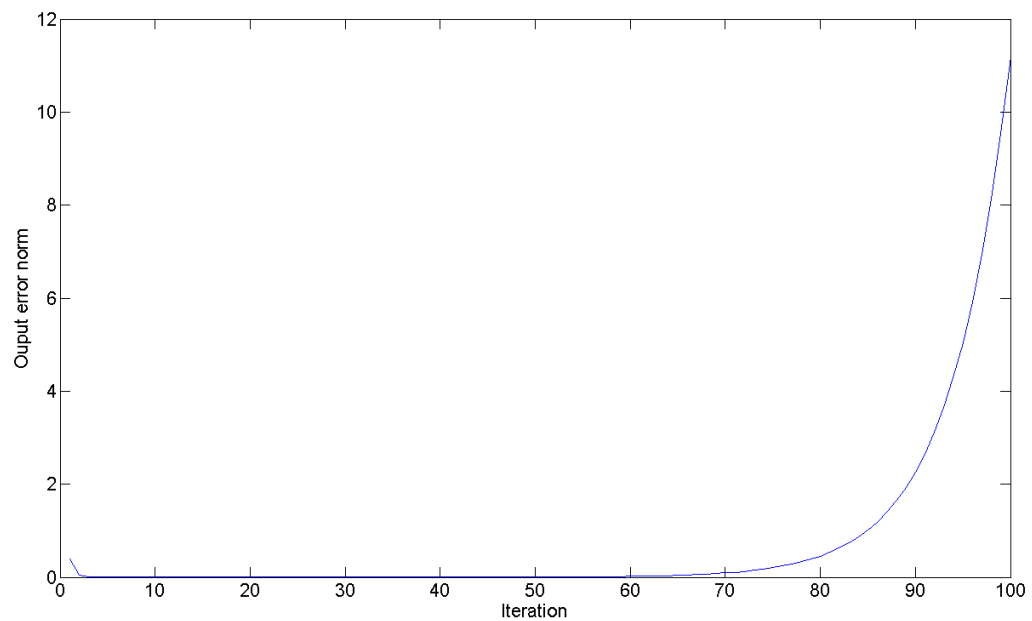
FIGURE 6.22: Output error per iteration for the 21st order model with an inverse model algorithm
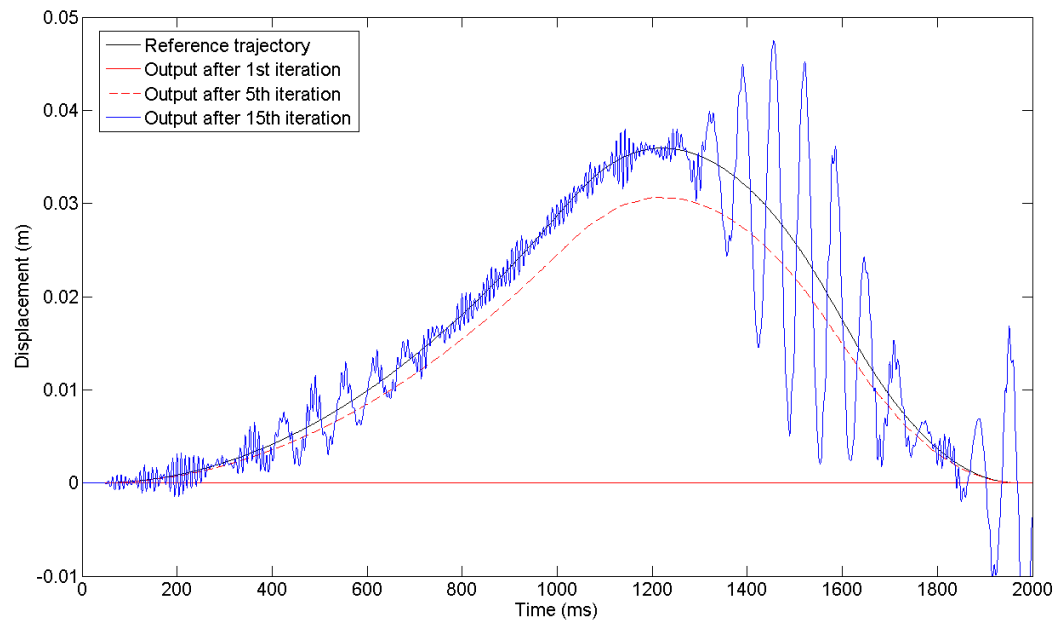
FIGURE 6.23: Output trajectory along trials 1, 5 and 15 for the 7^{th} order model with an inverse model algorithm
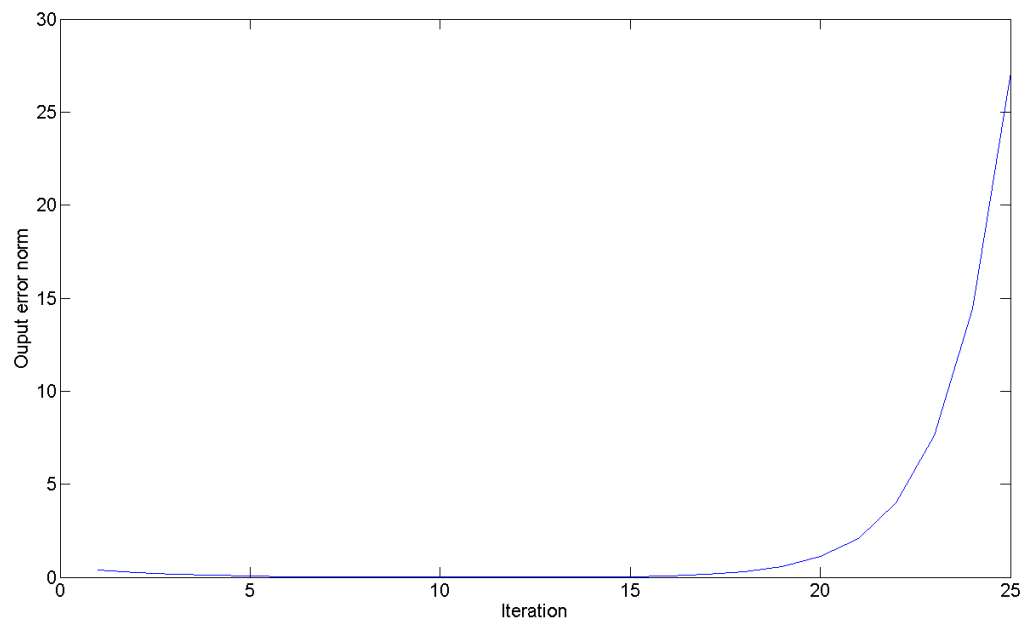
FIGURE 6.24: Output error per iteration for the 7^{th} order model with an inverse model algorithm
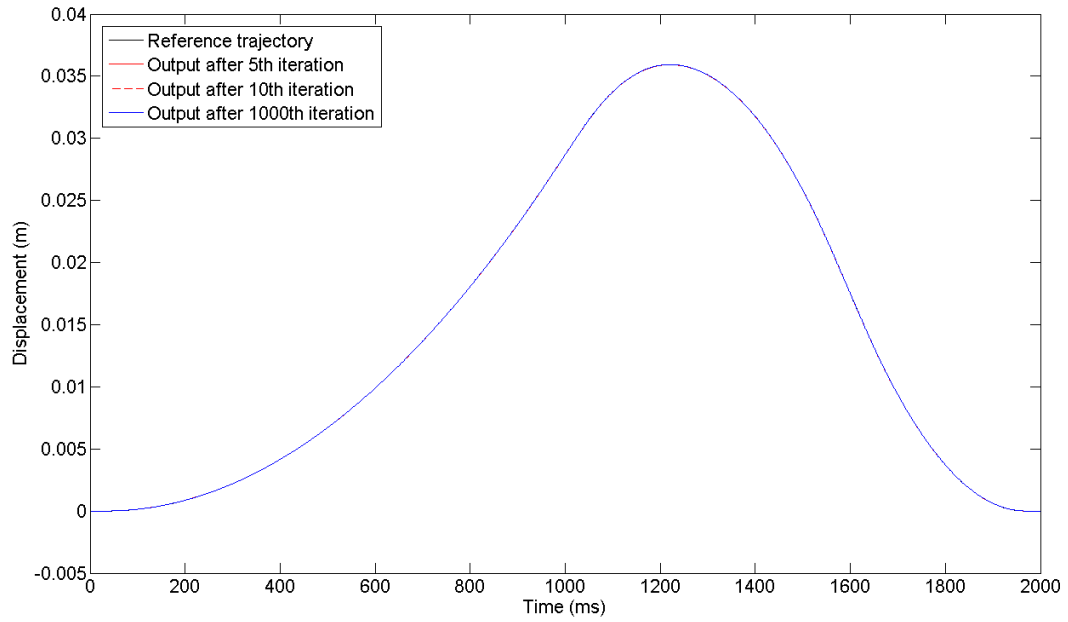
FIGURE 6.25: Output trajectory along trials 5, 10 and 1000 for the 21^{st} order model with an inverse model algorithm and Q as a zero-phase-shift Butterworth filter. (The reference and 5^{th}, 10^{th} and 1000^{th} iteration outputs coincide.)
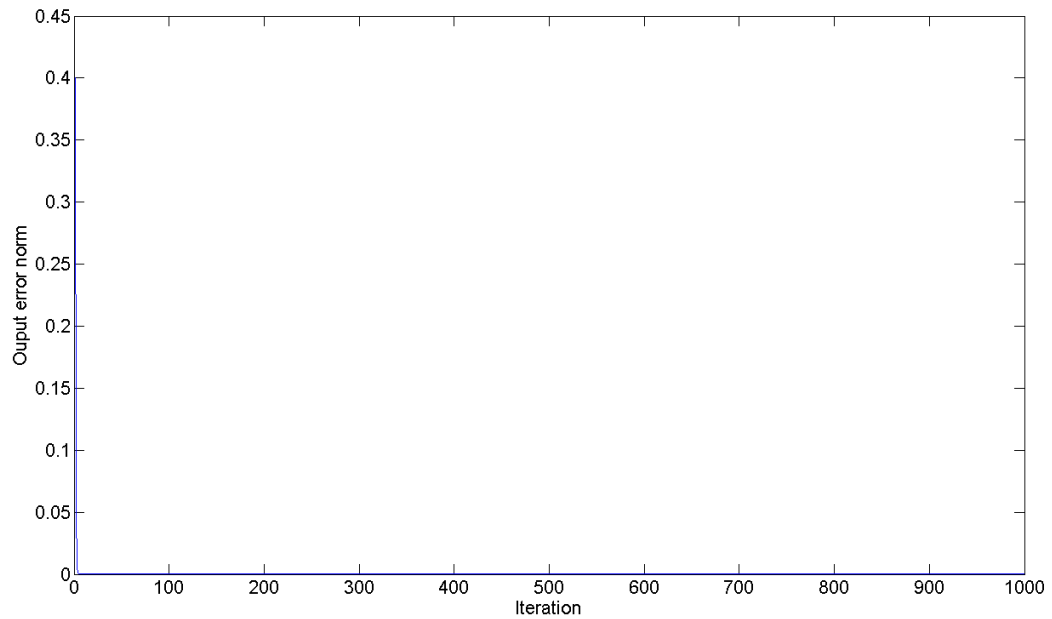
FIGURE 6.26: Output error per iteration for the 21^{st} order model with an inverse model algorithm and Q as a zero-phase-shift Butterworth filter
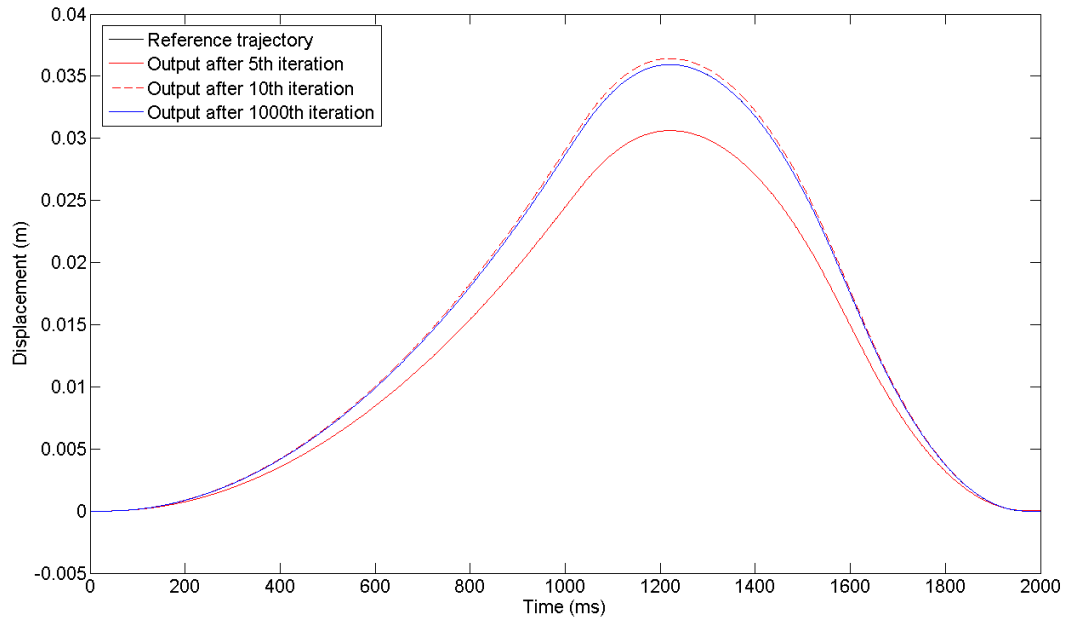
FIGURE 6.27: Output trajectory along trials 5, 10 and 1000 for the 7^{th} order model with an inverse model algorithm and Q as a zero-phase-shift Butterworth filter. (The reference and 1000^{th} iteration output coincide.)
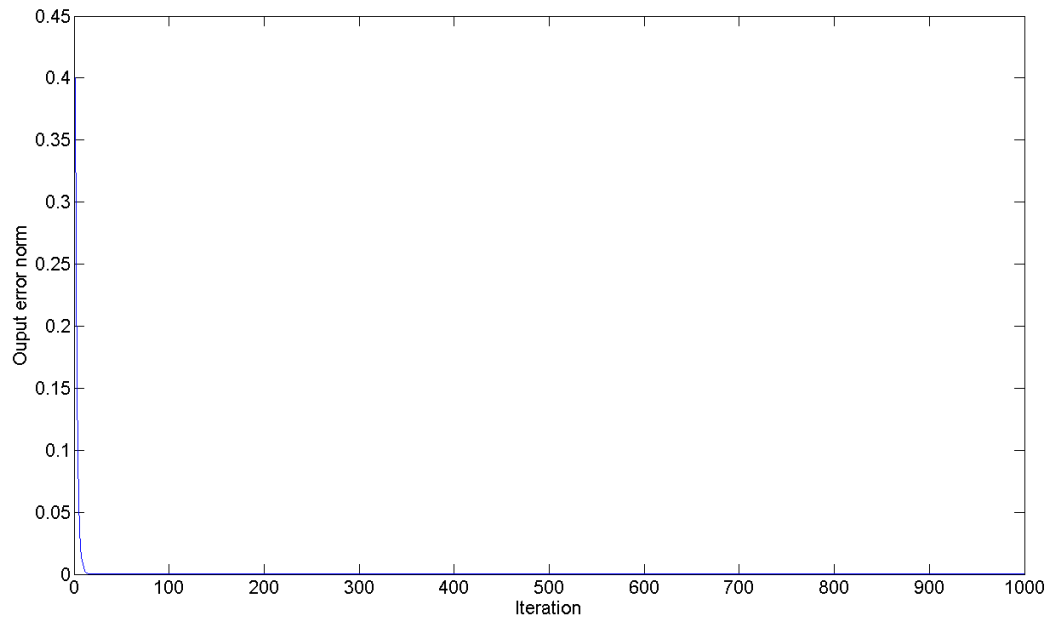
FIGURE 6.28: Output error per iteration for the 7^{th} order model with an inverse model algorithm and Q as a zero-phase-shift Butterworth filter

It was stated above that the norm of QL was calculated as 4×10^4 in the pure inverse model case, with Q as the identity. With Q as the above filter instead this reduces to 2.15×10^3. This is still a high gain, and therefore would result in a very small stability margin using this thesis' robust stability theorem. However, the simple introduction of a filter has reduced the norm by a factor of 18, without any significant loss of convergence.

6.4 Eigenvalues, Monotonic Convergence and Gain Stability

This section will discuss the relationship between the robust stability margin and the following two conditions on ILC systems:

- The eigenvalue convergence condition
  Detailed in Theorem 3.4, the eigenvalue condition states that a system is asymptotically stable if and only if

$$|\lambda_i(G_n Q(I - LP)G_n^T)| < 1, \tag{6.21}$$

  where $\{\lambda_0, \ldots, \lambda_{N-n}\}$ is the set of all eigenvalues of $G_n Q(I - LP)G_n^T$.

- The monotonic convergence criteria
  The monotonic convergence criteria of Theorem 3.5 states that a sufficient condition for the monotonic convergence of ILC systems expressed in lifted system form is that
$$\|G_n(Q(I - LP))G_n^T\| < 1. \tag{6.22}$$

In Chapter 5 a bound on the stability margin of lifted system ILC systems was calculated. In order to do this we required the strong condition of monotonic convergence to hold. This enabled the calculation of a simple bound for the summation $\sum_{j=0}^{\infty} \|[Q(I - LP)]^j\|_{l^p[0,T]}$. This condition is not required in general though, since a system may not fulfil this condition and yet still be gain stable, with a finite $\Pi_{\mathcal{M}//\mathcal{N}}$.

The following three examples will explain the role of the robust stability margin and detail how it fits with the two convergence conditions. This will be achieved by considering three nominal systems that follow different performance and convergence conditions.

i) A nominal system obeying both the monotonic and eigenvalue convergence criteria.

   Consider a nominal system that obeys the stronger condition of monotonic convergence. Using the robust stability theorem in this thesis we could find a robust stability margin for this system. While the theorem may not be able to guarantee that the perturbed system also converges monotonically, it will be possible to

ensure that its closed-loop performance is within acceptable limits since its gain would be bounded by the following:

$$\|\Pi_{\mathcal{M}_1//\mathcal{N}}\|_{\mathcal{W},g_0} \leq \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0} \frac{1 + \|(\Phi - I)|_{\mathcal{G}_{\overline{P},x_1}}\|_{\mathcal{W},x_1}}{1 - \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}\|(\Phi - I)|_{\mathcal{G}_{\overline{P},x_1}}\|_{\mathcal{W},x_1}}. \quad (6.23)$$

Recall also the condition given within the robust stability theorem that states that $\|(\Phi - I)|_{\mathcal{G}_{\overline{P},x_1}}\|_{\mathcal{W},x_1} < \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}^{-1}$. If this condition is tight then the denominator in equation 6.23 would approach zero and so the gain of the perturbed system could be large (although it would always remain finite). However, if we instead take a perturbed system that obeys $\|(\Phi - I)|_{\mathcal{G}_{\overline{P},x_1}}\|_{\mathcal{W},x_1} < \frac{1}{2}\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}^{-1}$ — and so the perturbed plant is well within the robust stability margin of the nominal system — then the bound on the gain of the perturbed system can be simplified to

$$\|\Pi_{\mathcal{M}_1//\mathcal{N}}\|_{\mathcal{W},g_0} \leq 2\|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0} + 1. \quad (6.24)$$

This would therefore prevent the perturbed system from exhibiting any long term instability effects (since we can also bound g_0 and g_1 as per Propositions 4.20 and 5.13). This reveals that, by minimising the uncertainty between plants and maximising the stability margin, we can improve the gain bound on the perturbed system. Note that this is only a bound. This means that even the largest possible gain of the most badly perturbed system may be well within this, depending on the conservatism that is introduced by the robust stability theorem.

ii) A nominal system obeying the eigenvalue convergence criteria but demonstrating long term instability.

As described above, a system may not obey the strong condition of monotonic convergence and still obey the eigenvalue condition and possess a bounded induced norm. (Note that this thesis has not provided a way of calculating the gain of a system that is not monotonic.) However, as can be the case with systems that are convergent but don't obey the monotonic condition, the transient performance may be completely unacceptable. Consider again the example given in Longman (2000) where a system is proven to converge but produces mean squared errors of over 10^{51} on the way. This system may be stable but the performance is unacceptable. This begs the question: *is there any validity in a robust stability theorem that allows this system to be considered stable and guarantees the stability of another system that is sufficiently close to it?*

Looking more closely at the process occurring in the robust stability theorem reveals the answer and also a better understanding as to the theorem's power. Analysis of the system above that produces the unacceptable transient behaviour would provide a stability margin close to zero: since the gain would be of the order of 10^{51} its inverse (and hence stability margin) would be of the order of 10^{-51}. Another plant would have to be within this distance (in the gap) in order

to guarantee stability and similar closed-loop behaviour. But why would this calculation even be considered if the nominal system is already known to have unacceptable performance? All that would be gained from such a result would be a guarantee that permits the perturbed system to also exhibit unacceptable performance.

iii) A system that has acceptable long term stability and obeys the eigenvalue condition but is non-monotonic.

Finally, consider the nominal system in-between the previous two, that only just fails to meet the monotonic convergence criteria, but has been considered appropriate for its task. Although this thesis would not be able to derive a robust stability margin for this plant directly (since $\|Q(I - LP)\|_{l^p[0,T]} \not< 1$) it would be possible to derive a robust stability margin by introducing a second 'nominal' system close to it that does obey the condition and then using that to derive a bound on $\|\Pi_{\mathcal{M}//\mathcal{N}}\|$ using equation 6.23. This could then be used to attain a gain bound for the perturbed system. An alternative would also be to find another method of bounding $\|\Pi_{\mathcal{M}//\mathcal{N}}\|$ that doesn't require $\|Q(I - LP)\|_{l^p[0,T]} < 1$.

The crux of this section is therefore that the robust stability result of this thesis provides a flexibility that complements the convergence criteria. The theorem allows a possibly non-monotonic perturbed system but provides a bound on the gain that enables a designer to ascertain whether the system's performance is acceptable. This means that the restrictive condition of monotonic convergence, that is often required to ensure long term stability, can be relaxed.

6.5 Summary

This chapter has provided a link between the robustness results of previous chapters and two topics within ILC. Analysis was conducted on the use of the filter Q and inverse model ILC, in terms of robustness. In both cases the robust stability margin given in this thesis was shown to back up previous results within the field. Various algorithms were also simulated on a plant model derived from a gantry robot, showing the applicability of these results to real situations.

The role of Q in lifted system designs was demonstrated: Q as the identity enabling convergence to zero error; and Q less than the identity increasing the robustness. It was then shown that implementing Q as a low-pass filter with a D-type algorithm allowed accurate convergence and was also capable of stabilising the perturbed plants. This complements work in papers such as de Roover (1996) and de Roover and Bosgra (2000), where Q is chosen as a low pass filter to maximise convergence at the desired low fre-

quencies and maximise robustness at high frequencies, often the location of unmodelled plant dynamics.

The robustness of inverse model ILC was examined using the 2D robust stability margin, showing that taking an exact inverse of a plant can lead to poor robustness results. This is due to the fact that the gain of P^{-1} is likely to be very large and is present in the inverse of the robust stability margin. The same low-pass filter Q was then used with the inverse model-based controller, demonstrating that the filter can dramatically increase the robustness and enable the controller to stabilise the perturbed plants.

Lastly, a comparison was given between the robust stability margin and the eigenvalue and monotonic conditions of convergence for ILC algorithms. It was shown that the robust stability theorem only guarantees stability in terms of the final convergence of the ILC algorithm. This means that a system proven to be stable by the theorem could exhibit some of the long term instability effects described in Longman (2000). It was also shown that the availability of a bound on the perturbed system's $\Pi_{\mathcal{M}//\mathcal{N}}$ allows a designer to guarantee that, although not necessarily monotonically convergent, a perturbed system has suitable performance.

Chapter 7

Conclusion

7.1 Summary

The purpose of this thesis has been the analysis of the robust stability of iterative learning control engaged in trajectory tracking. This has been achieved by developing a new robustness framework based on the nonlinear robust stability theory of Georgiou and Smith (1997a).

The generalisations made to the robust stability theorem, although primarily motivated by ILC, are ultimately applicable to a wider range of problems. Since Theorem 4.17 places no restrictions on the signal spaces used, beyond the existence of a definition of causality, it is possible to apply it to single- or multi-dimensional problems outside the field of ILC. The theorem is not restricted to linear systems and stability is defined relative to a reference trajectory so it is also capable of the examination of controllers performing trajectory tracking. This conclusion will therefore be split into two sections: the first shall summarise the developments this thesis has made to the gap in a nonlinear and biased norm setting, and the second will detail the results obtained when this approach is applied to ILC.

7.1.1 The gap metric for tracking

This thesis began with a demonstration as to the need for controllers to be formulated such that they are robust. The standard robustness problem stems from the lack of a perfect mathematical description of the plant that is required to be controlled. When it is known that there is a degree of uncertainty in the plant description it must be established that the controller designed for the modelled (nominal) system is also capable of stabilising the real (perturbed) one.

Structured uncertainty models restrict the difference between nominal and perturbed

systems to a highly uniform framework. As an example consider parametric uncertainty, where both models may be written as transfer functions sharing the same number of poles, the locations of which are allowed to fall anywhere within given ranges. However, these types of model uncertainty fail to capture the majority of possible high frequency perturbations since, by identification, it may not even be feasible to ascertain the *number* of poles at high frequency, let alone their location.

Unstructured models were introduced in Chapter 2 as a method of improving the description of uncertainty. The chapter culminated in an account of the coprime factor uncertainty representation, which allows the plant uncertainty to be captured in a much more general form. A stability theorem was also discussed that was based on the uncertainty measured using coprime factorisation, proving stability for a perturbed plant provided it lies close enough to the nominal plant — a open-loop condition for closed-loop stability. In this case the qualifier 'close' refers to the distance between the plants on a given topology. This is the graph topology, the same topology that is defined by the gap metric.

Chapter 4 detailed the linear gap and stability margin; and the associated 'classical' robust stability theorem (Theorem 4.3) was related to the coprime factor robust stability theorem. The linear gap defined the distance between linear plants as the distance between their graphs — linear subspaces of the signal space that contain all the signals compatible with the plant equation. This was then extended to nonlinear systems — where the graph is no longer a subspace and instead appears as an abstract subset of the signal space, again compatible with the plant equations — to define the nonlinear gap metric of Georgiou and Smith (1997a). A robust stability theorem based on this nonlinear gap was then given (Theorem 4.6 in this thesis and Theorem 1 of Georgiou and Smith, 1997a).

Since the aim of this thesis was to examine the tracking problem, and the nonlinear gap metric robust stability theorem considers stability relative to zero, a biased gap theorem was then given. This new 'biased' measurement and stability theorem was based on one published in Georgiou and Smith (1997b); however it was published there without proof and so a proof was developed and given here. This was done in such a way that the only restriction on the choice of signal space is that a definition of causality is available, therefore allowing a multi-dimensional signal space to be employed (in this thesis we consider a 2D space applicable to ILC). A robust stability theorem was then given concerning the robust stability of systems that contain a bias. The theorem collapses to the standard robust stability result of Georgiou and Smith (1997a) when considering systems where a zero input produces a zero output.

The theorem of Georgiou and Smith (1997b) restricted biases to lie within the chosen signal space, denoted $\mathcal{W}$. For spaces given by l^∞-type norms this is not a problem, but when considering other spaces difficulties arise. Consider the l^2-space. Any non-zero,

repeating signal (like that used as a reference within ILC) fails to lie within l^2 despite all elements being bounded. This can be demonstrated by truncating the signal and taking the norm: as the length of the truncation increases the norm increases *ad infinitum*. The biased norm theorem therefore needed to be developed further in order to ameliorate this.

The solution to this issue adopted in this thesis was to instead consider reference signals to lie within extended signal spaces $\mathcal{W}_e$. The biased robust stability theorem was then broadened to include these signals and re-proven to fit with new definitions of stability and stabilisability with respect to given trajectories, and graphs with biases. Chapter 4 culminated in a presentation of this 'biased graph robust stability theorem', which is therefore capable of examining the robust stability of trajectory tracking control in a wide variety of situations using single- or multi-dimensional signal spaces — ideally suited to an examination of iterative learning control.

7.1.2 Robust stability for ILC

Chapter 3 provided an introduction to iterative learning control, describing the motivation of ILC algorithms and the properties that distinguish ILC from other control methods. The idea is that when repeating tasks it should be possible to improve their execution based on previous 'experience'. For iterative learning control, from a mathematical standpoint, tasks that occur for a finite period of time are repeated for an infinite number of trials and an algorithm attempts to improve the process by learning along each.

The basic ILC algorithms were reviewed in Chapter 3 together with the lifted system or supervector description, where a plant is written down as a matrix mapping an input vector to an output vector; possible since ILC is usually operated in discrete-time and each trial is of a finite duration, after which it is reset before starting again. In this form it is easy to express the concepts of stability, performance and robustness that were given in Section 2.1.8:

i) Nominal stability
   The necessary and sufficient condition for the stability of an ILC algorithm is given by the eigenvalue condition that $|\lambda_i(Q(I - LP))| < 1$.

ii) Nominal performance
    As mentioned in Longman (2000), an ILC system fulfilling the eigenvalue condition may not provide an acceptable level of performance. However, the monotonic convergence criteria (given by $\|Q(I - LP)\| < 1$) guarantees that the ILC system's error decays monotonically in a given norm.

iii) Robust stability and performance

The stability and performance of a system containing a perturbed plant is then the concern of this thesis. It is wished to guarantee that when the plant is perturbed the closed-loop ILC system remains stable and that its performance is close to that of the nominal system.

In order to provide a measure of performance Chapter 5 introduced a 2D norm and signal space applicable to ILC. This led to a 2D stability margin result for ILC systems expressed in the lifted system form (given again in equation 7.1). Once the 2D space is equipped with a definition of causality, this 2D stability margin is used with the biased gap definition of Chapter 4 in 2D to form a full robust stability theorem applicable to ILC systems engaged in trajectory tracking.

It was shown in Section 5.3 that the 2D gap definition that arose from the biased graph gap definition with a 2D signal space can be related back to the 1D nonlinear gap. However, the relationship between biased and unbiased gaps is currently only applicable to linear ILC systems. This is due to the requirement that the operator defining the size of the gap is surjective (see section 5.3.2).

The relationship between 1D and 2D gaps for linear plants was then used to develop the main result of this thesis: Theorem 5.11, which implemented the biased graph robust stability theorem of Chapter 4 with a 2D measure of stability margin and the standard 1D gap measure. The result demonstrates that if an ILC system is nominally monotonically convergent and the plant is perturbed by a sufficiently small amount then the performance of the perturbed closed-loop system will be close to that of the nominal.

The gains of operators are all measured relative to reference signals. Proposition 5.12 proved that these references converge and Proposition 5.13 provided them with a bound. This ensures that the gain on the perturbed system is relative to an appropriate signal.

The stability margin calculated in Chapter 5 was then discussed in Chapter 6, drawing examples through simulations on a mathematical model of a gantry robot. The chapter demonstrated the trade-off between robustness and convergence found when varying the filter Q in lifted system ILC and expounded the benefits and costs of model-based ILC.

For ILC systems expressed in the lifted system representation setting Q to the identity was shown to allow exact convergence to the reference trajectory. However, the stability margin is maximised by minimising Q, as can be seen by examining the stability margin given in equation 7.1.

$$b_{P,C} = \left(\|\left(\begin{smallmatrix} I \\ P \end{smallmatrix} \right)\|_{l^p[0,T]} \left(1 + \frac{\|QLP\|_{l^p[0,T]} + \|QL\|_{l^p[0,T]}}{1 - \|Q(I-LP)\|_{l^p[0,T]}} \right) \right)^{-1} \qquad (7.1)$$

Also recall the equation denoting the trajectory that the plant converges to:

$$y_\infty = P(I - Q(I - LP))^{-1}(QLy_{ref}). \tag{7.2}$$

From these two equations the suggestion given in Section 6.2.1 is to implement Q as a low-pass filter (also advocated by several ILC publications including Norrlöf and Gunnarsson, 2002; de Roover, 1996; de Roover and Bosgra, 2000; Harte et al., 2005). In this way it would resemble the identity at low frequencies, where good tracking is required, and zero at high frequencies, where uncertainty is more likely; providing a method of trading-off robustness against convergence. Simulations to this effect were carried out in Section 6.3.1, where Q was designed as a second order low-pass Butterworth filter that operates forwards and backwards to cancel the phase shift it introduces. The results showed that the filter improved the robustness of a pure D-type controller and exhibited superior convergence to the D-type algorithm where Q is a constant less than one.

It was shown that model-based ILC can exhibit a very low stability margin, causing a severe lack of robustness. When L is given by P^{-1} convergence is achieved in a single iteration, but the appearance of L in the numerator of equation 7.1 demonstrates that the cost of this is borne by the stability margin. It is likely that the inverse of the plant would be sizeable in norm. This was indicated by the inverse of the gantry robot's 1^{st} order model being as high as 4×10^4. This leads to the same conclusion as found when examining Q: that robustness can be severely compromised when attempting to increase performance. This situation was then improved by again including Q as a low-pass filter. In this case the convergence was to a very low error and was achieved at a much faster rate than with a D-type controller.

Section 6.4 returned to the stability and performance conditions for ILC described at the start of this section, explaining how the derived robust stability theorem fits these conditions. It was explained that, in this thesis, the bound on the stability margin required the system to fulfil the monotonic convergence criteria, however this is restrictive since a non-monotonic system may posses gain stability. It was also clarified that the robust stability theorem does *not* guarantee that the perturbed system is monotonically convergent, though it can be used to bound the gain of $\Pi_{\mathcal{M}_1//\mathcal{N}}$ by equation 7.3.

$$\|\Pi_{\mathcal{M}_1//\mathcal{N}}\|_{\mathcal{W},g_0} \leq \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0} \frac{1 + \|(\Phi - I)|_{\mathcal{G}_{\overline{P},x_1}}\|_{\mathcal{W},x_1}}{1 - \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}\|(\Phi - I)|_{\mathcal{G}_{\overline{P},x_1}}\|_{\mathcal{W},x_1}} \tag{7.3}$$

This equation can therefore be used to demonstrate the long term stability of the perturbed plant-controller pair since, with a small gap, the *performance* of the perturbed system is *close* to that of the nominal system, and g_0 and $\Pi_{\mathcal{M}_1//\mathcal{N}}g_0$ are bounded by Propositions 4.20 and 5.13.

If the condition $\|(\Phi - I)|_{\mathcal{G}_{\overline{P},x_1}}\|_{\mathcal{W},x_1} < \|\Pi_{\mathcal{M}//\mathcal{N}}\|_{\mathcal{W},x_0}^{-1}$ is tight then this can lead to a very small denominator in equation 7.3. This can therefore allow a large gain for the

perturbed plant, such as the 10^{51} mentioned in Longman (2000), although this gain would be finite. However, if the nominal system possesses a large stability margin and the gap between plants is small then then gains of both nominal and perturbed systems will be similar.

The biased graph robust stability theorem in this thesis is a highly generalised result. The analysis is not based on frequency domain techniques so is exact rather than an approximation, since results based on Fourier transforms inherently assume an infinite trial length in order to find the transforms of signals. However, the gap can be related to frequency domain results when using l^2 and considering linear systems. This allows the consideration of the classical frequency domain perturbations.

The requirement that a plant and controller be causal are defined in a traditional sense and so 'non-causal' ILC algorithms (where information ahead of the current time on previous iterations) are able to be examined. This is due to the ordering that has been defined on the 2D space. Robustness is also calculated here with a more general uncertainty model, not restricted to additive, multiplicative, parametric or interval uncertainties. This allows a greater range of uncertainties to be captured; for example additive uncertainty does not allow stable and unstable models to be compared; and parametric does not permit changes on model order. (For references to published material on ILC's robust stability see Section 3.7.)

7.2   Future Work

This thesis leaves a number of avenues down which the work can be extended. An important direction would be to simultaneously continue research experimentally and theoretically: extending it to different ILC algorithms; using it to achieve a greater understanding of the causes of long term instability; and drawing the theory alongside non-ILC theorems to provide a reliable benchmark for comparing the capabilities of different algorithms both within and outside of ILC.

It is also anticipated that the analysis established here can be utilised to provide a base for quantitative design procedures. An example of this would be the design of the filter Q.

Some of the specific areas due for consideration are explained here. As before, this has been split into work on extending both the biased graph robust stability theorem and the robust stability result for ILC.

7.2.1 Furthering the biased graph robust stability theorem

- In this thesis we have considered the closed-loop stability with respect to disturbances of all magnitudes in a given space. In many situations this may not be possible to achieve and a robust stability theorem based on a bounded set of possible signals may be more appropriate. Georgiou and Smith (1997a) provides an example of an unstable plant with input saturation. In order to provide a stability theorem for such a plant the set of signals considered could be modified to those within a set of radius r, defined in Georgiou and Smith (1997a) by

$$\mathcal{S}_r := \left\{ w \in \mathcal{W} \; \middle| \; \sup_{\tau > 0} \|T_\tau w\|_{\mathcal{W}} \right\}. \tag{7.4}$$

This then leads to the gap between two plants being based on the difference between their graphs restricted to signals in $\mathcal{S}_r$:

$$\vec{\delta}_{\mathcal{S}_r}(\mathcal{M}, \mathcal{M}_1) = \begin{cases} \inf\{\|\Phi - I|_{\mathcal{M} \cap \mathcal{S}_r}\| : \Phi \text{ is causal,} \\ \qquad \text{maps } \mathcal{M} \cap \mathcal{S}_r \text{ into } \mathcal{M}_1 \\ \qquad \text{with } \Phi 0 = 0 \text{ such that} \\ \qquad T_\tau(\Phi - I)T_\tau \text{ is compact } \forall \tau > 0\} \\ \infty \text{ if no such operator } \Phi \text{ exists} \end{cases}, \tag{7.5}$$

and the robust stability margin also restricted to these signals:

$$b_{P,C}(\mathcal{S}_r) := \|\Pi_{\mathcal{M}//\mathcal{N}}|_{\mathcal{S}_r}\|_{\mathcal{W}}^{-1}. \tag{7.6}$$

(Note that the gain stability of the perturbed system will be defined for signals in some subset of $\mathcal{M}_1$, here denoted $\mathcal{M}_1 \cap \mathcal{S}_\epsilon$. In order to ensure that all the necessary signals are included in this subset it may be necessary to make the set $\mathcal{M} \cap \mathcal{S}_r$ quite large in order that for all signals $w \in \mathcal{M}_1 \cap \mathcal{S}_\epsilon$ there exists an $x \in \mathcal{M} \cap \mathcal{S}_r$ such that $w = \Phi x$. In order for this to be guaranteed r and ϵ may differ substantially.)

Theorem 4 of Georgiou and Smith (1997a) then provides a robust stability result on bounded sets similar to Theorem 4.6 in this thesis (Theorem 1 of Georgiou and Smith, 1997a).

This is also used in French (2008) to establish the bound on the stability margin of an adaptive ILC system (see Section 3.7). Since the system's adaption is based on the *size* of the error rather than the error itself, the stability margin has a heavy dependance on the size of disturbances.

An expansion of this thesis' biased graph robust stability theorem could therefore follow in this direction, relaxing slightly the condition of stability over the biased signal space to one of stability to signals sufficiently close to the bias. Algorithms such as the adaptive one mentioned above would require this in order for the analysis to proceed. It is also expected that the examination of other nonlinear

ILC strategies that contain norms explicitly in their algorithms would call for this work since the stability of the algorithm would likely depend on the disturbance level.

- As explained in Section 2.1.4, operators with superlinear growth would not possess a bounded induced norm and so are unable to fit with the induced norm-based robust stability theorem. However, for systems where $H_{P,C}$ is superlinear it is possible to use a gain function. Theorem 6 of Georgiou and Smith (1997a) is based on this approach. Recall the gain function given by

$$g[G](\alpha) := \sup\{\|Gx\| : \|x\| \le \alpha\}. \tag{7.7}$$

Define the set of functions $\mathcal{K}_\infty$ where all $g \in \mathcal{K}_\infty$ are continuous, strictly increasing functions $g\colon [0,\infty) \to [0,\infty)$ such that $g(0) = 0$ and $g(\infty) = \infty$. If $H_{P,C}$ is gain function stable, then the gain function stability of $H_{P_1,C}$ is guaranteed if there exists a function $\epsilon \in \mathcal{K}_\infty$ such that

$$g[I - \Phi] \circ g[\Pi_{\mathcal{M}//\mathcal{N}}](\alpha) \le (1 - \epsilon)^{-1}(\alpha). \tag{7.8}$$

Gain function stability is also examined in Bian and French (2005), where induced gain robust stability results are extended to weaker gain function stability. The topology on which the gain function stability theorem rests is also discussed and related to the graph topology since, as no metric is given, the gain function method does not explicitly define any topology on the uncertainty.

As with bounded set robust stability mentioned above, the work of this thesis could be continued towards some gain function stability results. If this was to proceed the gain function stability theorem of Georgiou and Smith (1997a) would have to be altered to accommodate the bias present in this thesis' biased graph robust stability theorem.

- Chapter 4 briefly mentioned the ν-gap in both linear (Section 4.1) and nonlinear (Section 4.2.3) settings. The linear ν-gap of Vinnicombe (1993, 2001), although inducing the same 'graph topology', provides a tighter bound than the gap metric. It may therefore be worthwhile to generate similar results to this thesis using the ν-gap, at least for linear systems.

For the nonlinear case the bound given by the ν-gap may not remain tighter than the nonlinear gap, as explained in Anderson et al. (2002). However, the nonlinear approximation of the ν-gap given in Vinnicombe (1999) remains a good candidate for furthering this thesis' work. An expansion of the relationship between the 2D biased gap and the standard nonlinear gap could be made to include the ν-gap; making the gap calculation accessible to an audience more comfortable with the ν-gap.

7.2.2 Furthering the robustness for ILC

- Recall that Theorem 5.10 related the 2D gap between plants to the standard 1D gap. In order to achieve this Lemma 5.5 proved the inheritance of the surjectivity property directly from the unbiased to the biased sets of maps $P \to P_1$, but only when restricted to linear systems. For the case of nonlinear systems there is therefore a weakness in Theorem 5.10.

  In the movement between gaps there is a natural choice of relationship between operators that $\Psi(x) = \Phi(x - x_1) + \Psi(x_1)$, where Φ is the unbiased operator and Ψ is the biased operator (see equation 5.41). Since both operators must be restricted to lying in sets of surjective operators, if surjectivity is not inherited from one set to the other then it must be imposed. For a chosen operator Φ, there may therefore be no such operator Ψ that is both surjective and obeys $\Psi(x) = \Phi(x - x_1) + \Psi(x_1)$.

  If it were proven that surjectivity *is* inherited from one set of operators to the other then Theorem 5.10 will also hold for nonlinear systems, since this is the only point at which linearity is required.

- The ILC algorithms discussed in this thesis are those expressed in the lifted system form. This can be thought of as a generalisation of P- and D-type ILC. The theorems here could therefore be extended to incorporate other algorithms such as those that employ current iteration feedback (see Sections 3.4.6 and 3.4.2 and also Owens and Munde, 2000). This would prompt a return to some of the adaptive work in French (2008), requiring the biased norm gap theorem to be reproduced for the case of bounded signals, as in the first point in the previous section.

  If the theorems were generalised to cover different types of ILC systems it may be possible to use them as a benchmark for robust stability. This would then allow the comparison of different algorithms side-by-side to determine the properties of different controllers in different situations, enabling a more precise method of choosing the correct controller for a given plant and task.

- As explained in Section 3.4.3 robustness is an important aspect of model-based ILC. This is backed up by the robustness results given in Section 6.3, showing the excellent performance but unacceptable robustness of a pure inverse model approach. Further investigation could therefore be conducted on model-based ILC methods.

  Here the only model-based approach that was discussed was the simple lifted system algorithm where $L = P^{-1}$. However, the inverse of a non-minimum phase system is unstable and so other techniques may have to be used. Markusson et al. (2001) considers finding the inverse through non-causal filtering techniques and Ratcliffe et al. (2008) instead makes use of an adjoint algorithm based on a steepest-descent method. An in-depth investigation of these and other algorithms

and a comparison in terms of robustness would be a solid candidate for furthering the work of this thesis.

- Theorem 5.1, concerning the stability margin of an ILC system, only permits the use of $L^{p,q}(\mathbb{N} \times [0, T])$ for $q = 1$ and $q = \infty$. This is due to the use of the triangle inequality to generate the bounds. As an example see equation 5.27: in order that the summation is moved outside the norm we use the triangle inequality. This is possible for both $q = 1$ and $q = \infty$ since the following two inequalities hold:

$$\sum_{N=0}^{\infty} \left\| \sum_{n=0}^{N} a_n \right\|^1 \leq \sum_{N=0}^{\infty} \sum_{n=0}^{N} \|a_n\|^1$$

$$\sup_{0 \leq N \leq \infty} \left\| \sum_{n=0}^{N} a_n \right\| \leq \sup_{0 \leq N \leq \infty} \sum_{n=0}^{N} \|a_n\|. \tag{7.9}$$

However, the proof for any value of q requires more work due to the difficult nature of exponents within norm calculations. For $1 < q < \infty$ the above inequalities become:

$$\sum_{N=0}^{\infty} \left\| \sum_{n=0}^{N} a_n \right\|^q \leq \sum_{N=0}^{\infty} f(N, q) \sum_{n=0}^{N} \|a_n\|^q, \tag{7.10}$$

where the value given by $f(N, q)$ increases as N and q increase. Since in this case the external summation sums for $N \to \infty$ this would lead to an infinite measure of gain regardless of the system in question. To demonstrate consider the simple case of summing two elements with $q = 2$:

$$\|A + B\|^2 \leq 3\|A\|^2 + 3\|B\|^2. \tag{7.11}$$

For larger values of q the situation worsens. Now consider the bound again when splitting three terms:

$$\|A + B + C\|^2 \leq 5\|A\|^2 + 5\|B\|^2 + 5\|C\|^2. \tag{7.12}$$

In order to combat this a slightly different approach may be called for or a different definition of norm used.

- Theorem 5.1 also requires the plant to be expressed in the lifted system framework detailed previously. This theorem does have a downside, that is the need to express the plant as a map P — mapping input to output — containing the Markov parameters of the system. This means that if each trial is 200 time-steps long P becomes a 200×200 matrix and if an extra time-step is added the gain alters. It would be much more prudent to produce a measure that did not require such large matrices and was invariant to the length of a given trial.

One possibility would be to fully exploit the structure of the matrices. For a linear time invariant plant the matrix P is Toeplitz and so eigenvalues and eigenvectors

obey certain properties (Gray, 1971). It is fairly intuitive — given the fact that P only contains the matrices A,B,C and D and only in a very repetitive fashion — that this kind of matrix should contain a large amount of redundancy. Further investigation could reveal the possibility of making some considerable simplifications.

On the existence of an alternative method that is invariant to trial length, there is some evidence to suggest that this may not be feasible. Section 3.6.2 detailed that the learning process is sometimes concentrated towards the start of each iteration and the performance degrades as time passes, due to the algorithm evolving differently to the previous iteration (this should not be as much of a concern with model-based methods since this would be compensated for automatically by the controller). This can produce outputs of the form shown in Figure 3.6. This clearly shows the dependance that trial length would have on the gain calculation — the longer the trials, the longer it would take the algorithm to converge and the stronger the possible divergence along the way.

- Testing of all these margins on a real-life system would provide some insight into the conservativeness of the stability bounds for specific cases and would demonstrate the practical use of the gap metric for ILC. In a similar vein to the experimental work from the theses Freeman (2004), Ratcliffe (2005), Cai (2009) and Norrlöf (1998, 2000b), different algorithms could be examined on physical systems and the robustness results compared to those derived from work in this thesis.

  Chapter 6 discussed the gantry robot (of Freeman, 2004 and Cai, 2009) in terms of the filter Q and a simple model-based algorithm. This could easily be extended by testing the performance and robustness of these algorithms on plant models and then implementing them on hardware. The non minimum phase testbed of Freeman (2004) could also be examined. This would prove useful for the experimentation of model-based algorithms, since the standard inverse of the plant would be unstable.

  Other ILC strategies could also be examined such as hybrid ILC, where a non-iterative controller is implemented around the plant together with an IL controller. Moon et al. (1998) and Tayebi and Zeremba (2003) consider the robustness of ILC using this approach (see Section 3.7). Tayebi and Zeremba (2003) states that, under multiplicative uncertainties, if the plant and non-iterative controller obey the robust performance condition then the design of a convergent IL controller is straightforward. An application of this thesis' results in this direction may yield some interesting conclusions. If the non-iterative controller and plant are considered a separate robustness problem, under the gap-based robust stability margin it would be ideal to minimise their associated map $\Pi_{\mathcal{M}//\mathcal{N}}$. Clearly, minimising this map should have a positive effect on minimising the equivalent map around the iterative system. This thesis therefore appears to support the view given in

Tayebi and Zeremba (2003), although more analysis in this direction is required.

- In Chapter 6 the role of the filter Q was briefly expounded, however the simulations given only implemented Q in thee forms: the identity, a constant multiplied by the identity and a simple Butterworth filter. A more useful result would be to examine the effect of designing Q using different low-pass filters and also compare these to the aliasing method described in Section 3.6.4. Further investigation could also be conducted on a design based on the approach of de Roover (1996) and de Roover and Bosgra (2000), where a frequency ω is selected, up to which an inverse model controller is accurate. Then Q can be chosen as a low-pass filter removing the uncertainty at high frequency. Since these papers have advocated this approach both for robustness and convergence it would also be worthwhile to fully ascertain the effect this would have on the stability margin $b_{P,C}$.

- An investigation could be carried to to discover just how conservative the results are, providing proof as to how useful they would be in implementation. This does not just concern the bound on stability margin but also the conservatism present in the inequality relating 1D and 2D, biased and non-biased gaps.

- Finally, this thesis has detailed a new robust stability theorem, applied it to ILC and provided some qualitative interpretation of the result. Following this up with some quantitative research would be a prime concern. This would involve calculating the stability margins for ILC systems, comparing robustnesses of ILC and non-ILC systems side-by-side, measuring the distance between two plants in 2D and comparing it to their distance in 1D, and giving some concrete evidence of the costs, benefits and trade-offs of various ILC systems.

Bibliography

H.-S. Ahn, YQ. Chen, and K. L. Moore. Iterative learning control: Brief survey and categorization. *IEEE Transactions on Systems, Man and Cybernetics—Part C: Applications and Reviews*, 37(6):1099–1121, November 2007.

H.-S. Ahn, K. L. Moore, and Y. Chen. Schur stability radius bounds for robust iterative learning controller design. In *Proceedings of the American Control Conference*, pages 178–183, Portland, OR, USA, June 2005.

H.-S. Ahn, K. L. Moore, and Y. Chen. Monotonic convergent iterative leaning controller design based on interval model conversion. *IEEE Transactions on Automatic Control*, 51(2):366–371, February 2006.

M. A. Alsubaie, C. Freeman, Z. Cai, P. Lewin, and E. Rogers. ILC initial input selection with experimental verification. In *Symposium on Learning Control at IEEE CDC*, Shanghai, China, December 2009.

N. Amann. *Optimal Algorithms for Iterative Learning Control*. PhD thesis, Faculty of Engineering, University of Exeter, 1996.

N. Amann, D. H. Owens, and E. Rogers. New results in iterative learning control. In *International Conference on Control*, volume 1, pages 640–645, March 1994.

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

N. Amann, D. H. Owens, and E. Rogers. Iterative learning control using optimal feedback and feedforward actions. *International Journal of Control*, 65(2):277–293, September 1996b.

B. D. O. Anderson, T. S. Brinsmead, and F. De Bruyne. The Vinnicombe metric for nonlinear operators. *IEEE Transactions on Automatic Control*, 47(9):1450–1465, September 2002.

B. D. O. Anderson, M. R. James, and D. J. N. Limebeer. Robust stabilization of nonlinear systems via normalized coprime factor representations. *Automatica*, 34(12): 1593–1599, 1998.

T. M. Apostol. *Mathematical Analysis*. Addison-Wesley, 1974.

S. Arimoto, S. Kawamura, and F. Miyazaki. Bettering operation of robots by learning. *Journal of Robotic Systems*, 1(2):123–140, 1984.

S. Arimoto, S. Kawamura, F. Miyazaki, and S. Tamaki. Learning control for dynamical systems. In *Proceedings of the 24th IEEE Conference on Decision and Control*, pages 1375–1380, Fort Lauderdale, Florida, USA, December 1985.

K. J. Åström and B. Wittenmark. *Adaptive Control*. Addison-Wesley, 1989.

G. J. Balas, J. C. Doyle, K. Glover, A. Packard, and R. Smith. *μ-Analysis and Synthesis Toolbox: For Use with MATLAB*. The MathWorks, 3rd edition, October 1998.

A. D. Barton, P. L. Lewin, and D. J. Brown. Practical implementation of a real-time iterative learning position controller. *International Journal of Control*, 73(10):992–999, July 2000.

E. Berkson. Some metrics on the subspaces of a banach space. *Pacific Journal of Mathematics*, 13(1):7–22, 1963.

W. Bian and M. French. Coprime factorisation and gap metric for nonlinear systems. In *Proceedings of the 42th IEEE Conference on Decision and Control*, pages 4694–4699, Maui, Hawaii, USA, December 2003.

W. Bian and M. French. Graph topologies, gap metrics, and robust stability for nonlinear systems. *SIAM Journal of Control and Optimization*, 44(2):418–443, 2005.

Z. Bien and K. M. Huh. Higher-order iterative learning control algorithm. *IEE Proceedings — Control Theory and Applications*, 136(3):105–112, May 1989.

R. Bradley and M. French. Robust stability for iterative learning control. In *Proceedings of the American Control Conference*, pages 955–960, St. Louis, Missouri, USA, June 2009.

D. A. Bristow, M. Tharayil, and A. G. Alleyne. A survey of iterative learning control. *IEEE Control Systems Magazine*, 26(3):96–114, June 2006.

Z. Cai. *Iterative Learning Control: Algorithm Development and Experimental Benchmarking*. PhD thesis, Electronics and Computer Science, University of Southampton, 2009.

Z. Cai, C. Freeman, P. Lewin, and E. Rogers. Experimental comparison of stochastic iterative learning control algorithms. In *Proceedings of the American Control Conference*, pages 4548–4553, Seattle, Washington, USA, June 2008a.

Z. Cai, C. T. Freeman, P. L. Lewin, and E. Rogers. Iterative learning control for a non-minimum phase plant based on a reference shift algorithm. *Control Engineering Practice*, 16(6):633–643, 2008b.

M. Cantoni and G. Vinnicombe. Linear feedback systems and the graph topology. *IEEE Transactions on Automatic Control*, 47(5):710–719, May 2002.

G. Casalino and G. Bartolini. A learning procedure for the control of movements of robotic manipulators. In *IASTED Symposium on Robotics and Automation*, pages 108–111, Amsterdam, The Netherlands, 1984.

Y. Chen, Z. Gong, and C. Wen. Analysis of a high-order iterative learning control algorithm for uncertain nonlinear systems with state delays. *Automatica*, 34(3):345–353, 1998.

Y. Chen, C. Wen, Z. Gong, and M. Sun. An iterative learning controller with initial state learning. *IEEE Transactions on Automatic Control*, 44(2):371–376, February 1999.

Y. Chen, J.-X. Xu, and T. H. Lee. An iterative learning controller using current iteration tracking error information and initial state learning. In *Proceedings of the 35th IEEE Conference on Decision and Control*, volume 3, pages 3064–3069, December 1996.

Y. Chen, J.-X. Xu, and C. Wen. A high-order terminal iterative learning control scheme. In *Proceedings of the 36th IEEE Conference on Decision and Control*, pages 3771–3772, San Diego, CA, USA, December 1997.

I. Chin, S. J. Qin, K. S. Lee, and M. Cho. A two-stage iterative learning control technique combined with real-time feedback for independent disturbance rejection. *Automatica*, 40(11):1913–1922, 2004.

I. S. Chin, J. Lee, H. Ahn, S. Joo, K. S. Lee, and D. Yang. Optimal iterative learning control of wafer temperature uniformity in rapid thermal processing. In *Proceedings of the IEEE International Symposium on Industrial Electronics*, volume 2, pages 1225–1230, 2001.

M. Cho, Y. Lee, S. Joo, and K. S. Lee. Sensor location, identification, and multivariable iterative learning control of an RTP process for maximum uniformity of wafer temperature distribution. In *Proceedings of the IEEE International Conference on Advanced Thermal Processing of Semiconductors*, pages 177–184, September 2003.

J. Y. Choi and H. M. Do. A learning approach of wafer temperature control in a rapid thermal processing system. *IEEE Transactions on Semiconductor Manufacturing*, 14 (1):1–10, February 2001.

J. Y. Choi and J. S. Lee. Adaptive iterative learning control of uncertain robotic systems. *IEE Proceedings — Control Theory and Applications*, 147(2):217–223, March 2000.

B. Chu and D. H. Owens. Accelerated norm-optimal iterative learning control algorithms using successive projection. *International Journal of Control*, 82(8):1469–1484, August 2009.

J. J. Craig. Adaptive control of manipulators through repeated trials. In *Proceedings of the American Control Conference*, pages 1566–1573, San Diego, USA, 1984.

D. de Roover. Synthesis of a robust iterative learning controller using an H_∞ approach. In *Proceedings of the 35th IEEE Conference on Decision and Control*, volume 3, pages 3044–3049, Kobe, Japan, December 1996.

D. de Roover and O. H. Bosgra. Synthesis of robust multivariable iterative learning controllers with application to a wafer stage motion system. *International Journal of Control*, 73(10):968–979, July 2000.

B. G. Dijkstra. *Iterative Learning Control with Applications to a Wafer Stage*. PhD thesis, Delft University of Technology, 2003.

B. G. Dijkstra and O. H. Bosgra. Convergence design considerations of low order Q-ILC for closed loop systems, implemented on a high precision wafer stage. In *Proceedings of the 41st IEEE Conference on Decision and Control*, volume 3, pages 2494–2499, December 2002a.

B. G. Dijkstra and O. H. Bosgra. Extrapolation of optimal lifted system ILC solution, with application to a waferstage. In *Proceedings of the American Control Conference*, volume 4, pages 2595–2600, May 2002b.

B. G. Dijkstra and O. H. Bosgra. Noise suppression in buffer-state iterative learning control, applied to a high precision wafer stage. In *Proceedings of the International Conference on Control Applications*, volume 2, pages 998–1003, September 2002c.

B. G. Dijkstra and O. H. Bosgra. Exploiting iterative learning control for input shaping, with application to a wafer stage. In *Proceedings of the American Control Conference*, volume 6, pages 4811–4815, June 2003.

T. Donkers, J. van de Wijdeven, and O. Bosgra. A design approach for noncausal robust iterative learning control using worst case disturbance optimisation. In *Proceedings of the American Control Conference*, pages 4567–4572, Seattle, Washington, USA, 2008a.

T. Donkers, J. van de Wijdeven, and O. Bosgra. Robustness against model uncertainties of norm optimal iterative learning control. In *Proceedings of the American Control Conference*, pages 4561–4566, Seattle, Washington, USA, 2008b.

J. Doyle, B. Francis, and A. Tannenbaum. *Feedback Control Theory*. Macmillan, 1990.

J. C. Doyle, T. T. Georgiou, and M. C. Smith. The parallel projection operators of a nonlinear feedback system. In *Proceedings of the 31st IEEE Conference on Decision and Control*, pages 1050–1054, Tucson, Arizona, December 1992.

J. B. Edwards and D. Owens. *Analysis and Control of Multipass processes*. Research Studies Press, 1982.

A. K. El-Sakkary. The gap metric: robustness of stabilisation of feedback systems. *IEEE Transactions on Automatic Control*, 30(3):240–247, March 1985.

C. Foias, T. T. Georgiou, and M. C. Smith. Geometric techniques for robust stabilization of linear time-varying systems. In *Proceedings of the 29th IEEE Conference on Decision and Control*, pages 2868–2873, Honolulu, Hawaii, December 1990.

C. Freeman, M. Alsubaie, Z. Cai, P. Lewin, and E. Rogers. Initial input selection for iterative learning control. *ASME Journal of Dynamic Systems, Measurement and Control*, in press.

C. Freeman, A. M. Hughes, J. Burridge, P. Chappell, P. Lewin, and E. Rogers. Iterative learning control of FES applied to the upper extremity for rehabilitation. *Control Engineering Practice*, 17(3):368–381, 2009a.

C. Freeman, A. M. Hughes, J. Burridge, P. Chappell, P. Lewin, and E. Rogers. A model of the upper extremity using FES for stroke rehabilitation. *ASME Journal of Biomechanical Engineering*, 131(3), 2009b.

C. Freeman, A. M. Hughes, J. Burridge, P. Chappell, P. Lewin, and E. Rogers. A robotic workstation for stroke rehabilitation of the upper extremity using FES. *Medical Engineering and Physics*, 31(3):364–373, 2009c.

C. Freeman, A. M. Hughes, J. Burridge, P. H. Chappell, P. L. Lewin, and E. Rogers. Design and control of an upper arm FES workstation for rehabilitation. In *Proceedings of the 11th International Conference on Rehabilitation Robotics*, pages 66–72, Kyoto, Japan, June 2009d.

C. Freeman, P. Lewin, and E. Rogers. Experimental evaluation of iterative learning control algorithms for non-minimum phase plants. *International Journal of Control*, 78(11):826–846, 2005a.

C. Freeman, P. Lewin, E. Rogers, and J. Ratcliffe. Synchronisation of multi-axis automation processes using iterative learning control. In *Proceedings of the European Control Conference*, pages 1529–1534, Budapest, Hungary, August 2009e.

C. T. Freeman. *Experimental Evaluation of Iterative Learning Control Performance for Non-Minimum Phase Plants*. PhD thesis, Electronics and Computer Science, University of Southampton, 2004.

C. T. Freeman, I. L. Davies, P. L. Lewin, and E. Rogers. Iterative learning control of upper limb reaching using functional electrical stimulation. In *Proceedings of the 17th International Federation of Automatic Control World Congress*, pages 13444–13449, Seoul, Korea, July 2008.

C. T. Freeman, A. M. Hughes, J. H. Burridge, P. H. Chappell, P. L. Lewin, and E. Rogers. An upper limb model using FES for stroke rehabilitation. In *Proceedings of the European Control Conference*, pages 3208–3213, Budapest, Hungary, August 2009f.

C. T. Freeman, P. L. Lewin, and E. Rogers. Discrete predictive optimal ILC implemented on a non-minimum phase experimental testbed. In *Proceedings of the American Control Conference*, pages 282–287, Portland, Oregan, USA, June 2005b.

C. T. Freeman, P. L. Lewin, and E. Rogers. Experimental evaluation of iterative learning control algorithms for non-minimum phase plants. *International Journal of Control*, 78(11):826–846, July 2005c.

C. T. Freeman, P. L. Lewin, and E. Rogers. Further results on the experimental evaluation of iterative learning control algorithms for non-minimum phase plants. *International Journal of Control*, 80(4):569–582, 2007.

C. T. Freeman, P. L. Lewin, E. Rogers, and D. Owens. Iterative learning control of the redundant upper limb for rehabilitation. In *Proceedings of the American Control Conference*, Baltimore, Maryland, USA, July 2010.

M. French. Robust stability of iterative learning control schemes. *International Journal of Robust and Nonlinear Control*, 18(10):1018–1033, July 2008.

M. French, G. Munde, E. Rogers, and D. H. Owens. Recent developments in adaptive iterative learning control. In *Proceedings of the 38th IEEE Conference on Decision and Control*, volume 1, pages 264–269, Phoenix, AZ, USA, December 1999.

G. Gauthier and B. Boulet. Convergence analysis of terminal ILC in the z domain. In *Proceedings of the American Control Conference*, pages 184–189, Portland, OR, USA, June 2005.

T. T. Georgiou and M. C. Smith. Optimal robustness in the gap metric. *IEEE Transactions on Automatic Control*, 35(6):673–686, June 1990.

T. T. Georgiou and M. C. Smith. Robust stabilization in the gap metric: controller design for distributed plants. *IEEE Transactions on Automatic Control*, 37(8):1133–1143, August 1992.

T. T. Georgiou and M. C. Smith. Robustness analysis of nonlinear feedback systems: an input output approach. *IEEE Transactions on Automatic Control*, 42(9):1200–1221, September 1997a.

T. T. Georgiou and M. C. Smith. Robustness of a relaxation oscillator. *International Journal of Robust and Nonlinear Control*, 10:1005–1024, 2000.

T.T. Georgiou and M.C. Smith. Biased norms and robustness analysis for nonlinear feedback systems. In *Proceedings of the 36th IEEE Conference on Decision and Control*, pages 642–643, 1997b.

P. B. Goldsmith. The fallacy of causal iterative learning control. In *Proceedings of the 40th IEEE Conference on Decision and Control*, pages 4475–4480, Orlando, Florida, USA, December 2001.

P. B. Goldsmith. On the equivalence of causal LTI iterative learning control and feedback control. *Automatica*, 38:703–708, 2002.

R. M. Gray. Toeplitz and circulant matrices: a review. Technical report, Information Systems Laboratory, Stanford University, 1971.

S. Gunnarsson and M. Norrlöf. Some aspects of an optimization approach to iterative learning control. In *Proceedings of the 38th IEEE Conference on Decision and Control*, volume 2, pages 1581–1586, December 1999.

S. Gunnarsson and M. Norrlöf. On the disturbance properties of high order iterative learning control algorithms. *Automatica*, 42(11):2031–2034, 2006.

W. B. J. Hakvoort, R. G. K. M. Aarts, J. van Dijk, and J.B. Jonker. Lifted system iterative learning control applied to an industrial robot. *Control Engineering Practice*, 16(4):377–391, 2008.

T. J. Harte, J. Hätönen, and D. H. Owens. Discrete-time inverse model-based iterative learning control: stability, monotonicity and robustness. *International Journal of Control*, 78(8):577–586, May 2005.

J. Hätönen. *Issues of Algebra and Optimality in Iterative Learning Control*. PhD thesis, Department of Process and Environmental Engineering, University of Oulu, 2004.

J. Hätönen, D. H. Owens, and K. Feng. Basis functions and parameter optimisation in high-order iterative learning control. *Automatica*, 42(2):287–294, 2006.

J. J. Hätönen, T. J. Harte, D. H. Owens, J. D. Ratcliffe, P. L. Lewin, and E. Rogers. A new robust iterative learning control algorithm for application on a gantry robot. In *9th IEEE International Conference on Emerging Technologies and Factory Automation*, volume 2, pages 305–312, 2003.

J. J. Hätönen, D. H. Owens, and K. L. Moore. An algebraic approach to iterative learning control. *International Journal of Control*, 77(1):45–54, January 2004.

J. J. Hätönen, D. H. Owens, J. D. Ratcliffe, P. L. Lewin, and E. Rogers. Experimentally verified robustness properties of a class of model inverse algorithms. In *11th International Federation of Automatic Control Workshop on Adaptation and Learning in Control and Signal Processing*, St Petersburg, Russia, August 2007.

L. Hladowski, Z. Cai, K. Galkowski, E. Rogers, C. Freeman, P. Lewin, and W. Paszke. Repetitive process based iterative learning control designed by LMIs and experimentally verified on a gantry robot. In *Proceedings of the American Control Conference*, pages 949–954, St. Louis, Missouri, USA, June 2009.

L. Hladowski, K. Galkowski, Z. Cai, E. Rogers, C. Freeman, and P. Lewin. A 2D systems approach to iterative learning control with experimental verification. In *Proceedings*

of the 17th International Federation of Automatic Control World Congress, pages 2832–2837, Seoul, Korea, July 2008.

Z. Hou, J.-X. Xu, and J. Yan. An iterative learning approach for density control of freeway traffic flow via ramp metering. *Transportation Research Part C: Emerging Technologies*, 16(1):71–97, 2008.

Y.-P. Hsin, R. W. Longman, E. J. Solcz, and J. deJong. Stabilization due to finite word length in repetitive and learning control. *Advances in the Astronautical Sciences*, 97: 817–836, 1997.

M. Hu, H. Du, S.-F. Ling, Z. Zhou, and Y. Li. Motion control of an electrostrictive actuator. *Mechatronics*, 14(2):153–161, 2004.

Y. Huang and R. W. Longman. The source of the often observed property of initial convergance followed by divergance in learning and repetitive control. *Advances in the Astronautical Sciences*, 90:555–572, 1996.

Y.-C. Huang, M. Chan, Y.-P. Hsin, and C.-C. Ko. Use of PID and iterative learning controls on improving intra-oral hydraulic loading system of dental implants. *JSME International Journal Series C*, 46(4):1449–1455, 2003.

A. M. Hughes, C. Freeman, J. Burridge, P. Chappell, P. Lewin, and E. Rogers. Feasibility of iterative learning control mediated by functional electrical stimulation for reaching after stroke. *Journal of Neurorehabilitation and Neural Repair*, 23(6):559–568, 2009a.

A. M. Hughes, C. T. Freeman, J. Burridge, P. Chappell, P. Lewin, R. Pickering, and E. Rogers. Shoulder and elbow muscle activity during fully supported trajectory tracking in neurologically intact older people. *Journal of Electromyography and Kinesiology*, 19(6):1025–1034, 2009b.

A. M. Hughes, C. T. Freeman, J. Burridge, P. Chappell, P. Lewin, E. Rogers, B. Dibb, and M. Donovan-Hall. Efficacy of iterative learning control for stroke rehabilitation. *Progress in Neurology and Psychiatry*, 13(5):16–20, 2009c.

A. Ilchmann. Non-identifier-based adaptive control of dynamical systems. *IMA Journal of Mathematical Control and Information*, 8:321–366, 1991.

M. R. James, M. C. Smith, and G. Vinnicombe. Gap metrics, representations, and nonlinear robust stability. *SIAM Journal of Control and Optimization*, 43(5):1535–1582, 2005.

T.-J. Jang, H.-S. Ahn, and C.-H. Choi. Iterative learning control for discrete-time nonlinear systems. *International Journal of Systems Science*, 25(7):1179–1189, July 1994.

P. Jiang and R. Unbehauen. An iterative learning control scheme with deadzone. In *Proceedings of the 38th IEEE Conference on Decision and Control*, pages 3816–3817, Phoenix, Arizona, USA, December 1999.

S. Kawamura, F. Miyazaki, and S. Arimoto. Applications of learning method for dynamic control of robot manipulators. In *Proceedings of the 24th IEEE Conference on Decision and Control*, pages 1381–1387, Fort Lauderdale, Florida, USA, December 1985.

S. Kawamura, F. Miyazaki, and S. Arimoto. Realization of robot motion based on a learning method. *IEEE Transactions on Systems, Man and Cybernetics*, 18(1):126–134, January/February 1988.

M. G. Krein and M. A. Krasnosel'skii. Fundamental theorems on the extension of Hermitian operators and certain of their applications to the theory of orthogonal polynomials and the problem of moments. *Uspekhi Matematicheskikh Nauk*, 2(3): 60–106, 1947. (in Russian).

A. Lanzon and G. Papageorgiou. Distance measures for uncertain linear systems: A general theory. *IEEE Transactions on Automatic Control*, 54(7):1532–1547, July 2009.

F. Le, I. Markovski, C. Freeman, and E. Rogers. Identification of electrically stimulated muscle after stroke. In *Proceedings of the European Control Conference*, pages 1576–1581, Budapest, Hungary, August 2009.

F. Le, I. Markovsky, C. T. Freeman, and E. Rogers. Identification of electrically stimulated muscle models of stroke patients. *Control Engineering Practice*, 18:396–407, 2010.

H. Lee and Z. Bien. Study on robustness of iterative learning control with non-zero initial conditions. *International Journal of Control*, 64(3):345–359, 1996.

J. H. Lee, K. S. Lee, and W. C. Kim. Model-based iterative learning control with a quadratic criterion for time-varying linear systems. *Automatica*, 36(5):641–657, 2000.

K. H. Lee and Z. Bien. Initial condition problem of learning control. *IEE Proceedings — Part D*, 138(6):525–528, November 1991.

T. Lin. *Newton Method Based Iterative learning control for nonlinear systems*. PhD thesis, The Department of Automatic Control and Systems Engineering, The University of Sheffield, 2006.

D. Liu, L.-C. Fu, S.-H. Hsu, and T.-K. Kuo. Analysis on an on-line iterative correction control law for visual tracking. In *4th Asian Control Conference*, pages 2130–2133, 2002.

R. W. Longman. Iterative learning control and repetitive control for engineering practice. *International Journal of Control*, 73(10):930–954, 2000.

R. W. Longman. On the interaction between theory, experiments, and simulation in developing practical learning control algorithms. *International Journal of Applied Mathematics and Computer Science*, 13(1):101–111, March 2003.

P. Lucibello. Learning control of linear systems. In *Proceedings of the American Control Conference*, pages 1888–1892, Chicago, USA, June 1992.

L. Y. X. Ma, T. S. Low, and S. K. Tso. Discrete iterative learning controller. *Electronics Letters*, 29(12):1046–1048, June 1993.

O. Markusson, H. Hjalmarsson, and M. Norrlöf. Iterative learning control of nonlinear non-minimum phase systems and its application to system and model inversion. In *Proceedings of the 40th IEEE Conference on Decision and Control*, volume 5, pages 4481–4482, December 2001.

J.-H. Moon, T.-Y. Doh, and M. J. Chung. A robust approach to iterative learning control design for uncertain systems. *Automatica*, 34(8):1001–1004, 1998.

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

K. L. Moore. An observation about monotonic convergence in discrete-time, P-type iterative learning control. In *2001 IEEE International Symposium on Intelligent Control*, pages 45–49, Mexico City, Mexico, September 2001.

K. L. Moore, Y. Chen, and H.-S. Ahn. Iterative learning control: A tutorial and big picture view. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 2352–2357, 2006.

K. L. Moore and J. Xu. Special issue on iterative learning control. *International Journal of Control*, 73(10):819–823, 2000.

J. D. Newburgh. A topology for closed operators. *Annals of mathematics*, 53(2):250–255, March 1951.

M. Norrlöf. *On Analysis and Implementation of Iterative Learning Control*. PhD thesis, Linköping Studies in Science and Technology, 1998.

M. Norrlöf. Comparative study on first and second order ILC — frequency domain analysis and experiments. In *Proceedings of the 39th IEEE Conference on Decision and Control*, pages 3415–3420, Sydney, Australia, December 2000a.

M. Norrlöf. *Iterative Learning Control Analysis, Design, and Experiments*. PhD thesis, Linköping Studies in Science and Technology, 2000b.

M. Norrlöf and S. Gunnarsson. Time and frequency domain convergence properties in iterative learning control. *International Journal of Control*, 75(14):1114–1126, 2002.

M. Norrlöf and S. Gunnarsson. A note on causal and CITE iterative learning control algorithms. *Automatica*, 41:345–350, 2005.

S.-R. Oh, Z. Bien, and I.H. Suh. An iterative learning control method with application to robot manipulators. *IEEE Journal of Robotics and Automation*, 4(5):508–514, October 1988.

D. H. Owens. Iterative learning control — convergence using high gain feedback. In *Proceedings of the 31st IEEE Conference on Decision and Control*, volume 3, pages 2545–2546, December 1992.

D. H. Owens and K. Feng. Parameter optimization in iterative learning control. *International Journal of Control*, 76(11):1059–1069, 2003.

D. H. Owens and J. Hätönen. Iterative learning control — an optimization paradigm. *Annual Reviews in Control*, 29(1):57–70, 2005.

D. H. Owens and R. P. Jones. Iterative solution of constrained differential/algebraic systems. *International Journal of Control*, 27(6):957–974, 1978.

D. H. Owens and G. Munde. Adaptive iterative learning control. In *IEE Colloquium on Adaptive Control (Digest No: 1996/139)*, pages 6/1–6/4, June 1996.

D. H. Owens and G. Munde. Error convergance in an adaptive iterative learning controller. *International Journal of Control*, 73(10):851–857, 2000.

A. D. B. Paice and A. J. van der Schaft. The class of stabilizing nonlinear plant controller pairs. *IEEE Transactions on Automatic Control*, 41(5):634–645, May 1996.

A. M. Pascoal, R. Ravi, and P. P. Khargonekar. The graph topology for linear plants with applications to nonlinear robust stabilization. *IEEE Transactions on Automatic Control*, 38(2):298–302, February 1993.

M. Q. Phan, R. W. Longman, and K. L. Moore. Unified formulation of iterative learning control. *Advances in the Astronautical Sciences*, 105:93–111, 2000.

J. W. Polderman and J. C. Willems. *Introduction to Mathematical Systems Theory: A Behavioral Approach*. Springer, 1998.

J. Ratcliffe, L. V. Duinkerken, P. Lewin, E. Rogers, J. Hätönen, T. Harte, and D. Owens. Fast norm-optimal iterative learning control for industrial applications. In *Proceedings of the American Control Conference*, pages 1951–1956, Portland, Oregan, USA, June 2005a.

J. D. Ratcliffe. *Iterative Learning Control Implemented on a Multi-Axis System*. PhD thesis, Electronics and Computer Science, University of Southampton, 2005.

J. D. Ratcliffe, J. J. Hätönen, P. L. Lewin, E. Rogers, T. J. Harte, and D. H. Owens. P-type iterative learning control for systems that contain resonance. *International Journal of Adaptive Control and Signal Processing*, 19:769–796, 2005b.

J. D. Ratcliffe, J. J. Hätönen, P. L. Lewin, E. Rogers, and D. H. Owens. Robustness analysis of an adjoint optimal iterative learning controller with experimental verification. *International Journal of Robust and Nonlinear Control*, 18(10):1089–1113, 2008.

J. D. Ratcliffe, P. L. Lewin, and E. Rogers. Comparing the performance of two iterative learning controllers with optimal feedback control. In *International Symposium on Intelligent Control*, pages 838–843, Munich, Germany, October 2006a.

J. D. Ratcliffe, P. L. Lewin, E. Rogers, J. J. Hätönen, and D. H. Owens. Norm-optimal iterative learning control applied to gantry robots for automation applications. *IEEE Transactions on Robotics*, 22(6):1303–1307, 2006b.

W. Rudin. *Real and Complex Analysis*. McGraw-Hill, international edition edition, 1987.

S. S. Saab. On a discrete-time stochastic learning control algorithm. *IEEE Transactions on Automatic Control*, 46(8):1333–1336, August 2001.

S. S. Saab. Optimality of first-order ILC among higher order ILC. *IEEE Transactions on Automatic Control*, 51(8):1332–1335, August 2006.

E. D. Sontag. Smooth stabilization implies coprime factorization. *IEEE Transactions on Automatic Control*, 34(4):435–443, April 1989.

T. Sugie and T. Ono. An iterative control law for dynamical systems. *Automatica*, 27 (4):729–732, 1991.

M. Sun and D. Wang. Robust discrete-time iterative learning control: initial shift problem. In *Proceedings of the 40th IEEE Conference on Decision and Control*, volume 2, pages 1211–1216, December 2001.

K. M. Tao, R. L. Kosut, and G. Aral. Learning feedforward control. In *Proceedings of the American Control Conference*, volume 3, pages 2575–2579, June 1994.

A. Tayebi and S. Islam. Adaptive iterative learning control for robot manipulators: Experimental results. *Control Engineering Practice*, 14(7):843–851, 2006.

A. Tayebi and M. B. Zeremba. Robust iterative learning control design is straightforward for uncertain LTI systems satisfying the robust performance condition. *IEEE Transactions on Automatic Control*, 48(1):101–106, 2003.

C.-H. Tsai and C.-J. Chen. Application of iterative path revision technique for laser cutting with controlled fracture. *Optics and Lasers in Engineering*, 41(1):189–204, 2004.

M. Uchyama. Formulation of high-speed motion pattern of a mechanical arm by trial. *Transactions of the Society for Instrumentation and Control Engineers*, 14:706–712, 1978. (in Japanese).

J. van de Wijdeven. *Iterative Learning Control Design for Uncertain and Time-Windowed Systems*. PhD thesis, Technische Universiteit Eindhoven, 2008.

J. van de Wijdeven and O. Bosgra. Noncausal finite-time robust iterative learning control. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 258–263, New Orleans, LA, USA, December 2007.

M. S. Verma and L. R. Hunt. Right coprime factorizations and stabilization for nonlinear systems. *IEEE Transactions on Automatic Control*, 38(2):222–231, February 1993.

M. H. A. Verwoerd. *Iterative Learning Control: A Critical Review*. PhD thesis, University of Twente, 2005.

M. H. A. Verwoerd, G. Meinsma, and T. J. A. de Vries. One the use of noncausal LTI operators in iterative learning control. In *Proceedings of the 41st IEEE Conference on Decision and Control*, pages 3362–3366, Las Vegas, Nevada, USA, December 2002.

M. Vidyasagar. The graph metric for unstable plants and robustness esitimates for feedback stability. *IEEE Transactions on Automatic Control*, AC-29(5):403–418, May 1984.

G. Vinnicombe. Frequency domain uncertainty and the graph topology. *IEEE Transactions on Automatic Control*, 38(9):1371–1383, September 1993.

G. Vinnicombe. A ν-gap distance for uncertain and nonlinear systems. In *Proceedings of the 38th IEEE Conference on Decision and Control*, pages 2557–2562, Phoenix, Arizona, USA, December 1999.

G. Vinnicombe. *Uncertainty and Feedback: $\mathcal{H}_\infty$ loop-shaping and the ν-gap metric*. Imperial College Press, 2001.

D. Wang. On anticipatory iterative learning control designs for continuous time nonlinear dynamic systems. In *Proceedings of the 38th IEEE Conference on Decision and Control*, volume 2, pages 1605–1610, December 1999.

D. Wang. On D-type and P-type ILC designs and anticipatory approach. *International Journal of Control*, 73(10):890–901, July 2000.

C. Xie. *Nonlinear Output Feedback Control: An Analysis of Performance and Robustness*. PhD thesis, Electronics and Computer Science, University of Southampton, May 2004.

J. Xu and Y. Tan. On the converence speed of a class of higher-order ILC schemes. In *Proceedings of the 40th IEEE Conference on Decision and Control*, pages 4932–4937, Orlando, Florida, USA, December 2001.

J.-X. Xu, Y. Chen, T. H. Lee, and S. Yamamoto. Terminal iterative learning control with an application to RTPCVD thickness control. *Automatica*, 35(9):1535–1542, 1999.

D. R. Yang, K. S. Lee, H. J. Ahn, and J. H. Lee. Experimental application of a quadratic optimal iterative learning control method for control of wafer temperature uniformity in rapid thermal processing. *EEE Transactions on Semiconductor Manufacturing*, 16 (1):36–44, February 2003.

F. Yong, Y. C. Soh, and G. G. Feng. Congergence analysis of iterative learning control with uncertain initial condition. In *4th World Congress on Intelligent Control and Automation*, Shanghai, June 2002.

G. Zames. On the input-output stability of time-varying nonlinear feedback systems — Part 1: Conditions derived using concepts of loop gain, conicity, and positivity. *IEEE Transactions on Automatic Control*, 11(2):228–238, 1966.

G. Zames and A. El-Sakkary. Unstable systems and feedback: the gap metric. In *Proceedings of the Allerton Conference*, pages 380–385, 1980.

K. Zhou, J. C. Doyle, and K. Glover. *Robust and Optimal Control*. Prentice Hall, 1996.

S. Q. Zhu. Graph topology and gap topology for unstable systems. *IEEE Transactions on Automatic Control*, 34(8):848–855, 1989.