

Radiotherapy Cancer Treatment: Investigating
Real-Time Position and Dose Control, the
Sensor-Delayed Plant Output Estimation Problem,
and the Nonovershooting Step Response Problem

by

James Michael Patrick Stewart

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Master of Applied Science
in
Electrical and Computer Engineering

Waterloo, Ontario, Canada, 2006
© James Michael Patrick Stewart 2006

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

For over a century, physicians have prescribed x-ray radiation to destroy or impede the growth of cancerous tumours. Modern radiation therapy machines shape the radiation beam to balance the competing goals of maximizing irradiation of cancerous tissue and minimizing irradiation of healthy tissue, an objective complicated by tumour motion during the treatment and errors positioning the patient to align the tumour with the radiation beam. Recent medical imaging advances have motivated interest in using feedback during radiation therapy to track the tumour in real time and mitigate these complications. This thesis investigates how real-time feedback control can be used to track the tumour and focus the radiation beam tightly around the tumour. Improving on these results, a feedback control system is proposed for intensity modulated radiation therapy which allows a *non-uniform* radiation dose to be applied to the tumour. Motivated by the results of the proposed control systems, this thesis also examines two theoretical control problems: estimating the output of an unknown system when a sensor delay prevents its direct measurement, and designing a controller to provide an arbitrarily fast nonovershooting step response.

Acknowledgments

I am indebted to my supervisor, Professor Daniel Davison, for his boundless insight and support. I would also like to thank Professor Daniel Miller, Julie Vale, Wilfred Kwok, and Darrell Gaudette for their constant willingness to help me through stubborn problems. I also valued the support of Adam, Alan, Dana, Kristina, Lena, Noodles, and Ollie while I wrote this thesis.

I would like to acknowledge the government of Ontario and the Department of Electrical and Computer Engineering for providing the funding to make this work possible.

Finally, I owe a deep debt of gratitude to my parents for their steadfast support and encouragement.

Contents

1	Introduction	1
1.1	Radiotherapy Cancer Treatment	1
1.2	Objective	3
1.3	Outline	4
2	Background: Jaw-Based Collimator Controller	6
2.1	High-Level Approach	6
2.2	Multirate Observers	7
2.3	Jaw-Based Collimator Controller	12
3	Multileaf Collimator Tumour Tracking Control	15
3.1	Observer Gain Optimization	15
3.2	Multileaf Collimator Controller	18
3.3	Simulation Results	23
4	Tumour Dose Control for Intensity-Modulated Radiotherapy	29
4.1	Motivation	30
4.2	Proposed Dose Controller	31
4.2.1	Control Structure	31
4.2.2	Desired Irradiation Area Calculation	32
4.3	Simulation Results	34
4.4	Weaknesses of the Proposed Control System	38

5	Adaptive Estimation in the Presence of Sensor Delays	40
5.1	Problem Formation	41
5.2	Adaptive Estimation Technique	44
5.3	Simulation Results	50
5.3.1	Convergence Rate Sensitivity to Plant Input	54
5.3.2	Improving Convergence Rates Through Scaling	55
5.3.3	Tracking Time-Varying Parameters	57
5.4	Applicability to Tumour Position Estimation	58
5.5	Relationship to Adaptive Prediction	59
6	The Nonovershooting Step Response Problem	61
6.1	Background and Definitions	62
6.2	Case where $H(s)$ has Two NMP Zeros	67
6.3	Case where $H(s)$ has One NMP Zero	70
6.4	Simulation Results	72
6.5	Extensions to Discrete Time	75
6.5.1	Background and Definitions	76
6.5.2	Case where $H^d(z)$ has Two NMP Zeros	80
6.5.3	Case where $H^d(z)$ has One NMP Zero	82
7	Summary and Future Work	84

List of Figures

1.1	Collimation with jaws and multileaf collimation	2
1.2	External beam radiotherapy machine and MLC	3
1.3	Radiotherapy using a multileaf collimator	4
2.1	Proposed radiotherapy scheme	7
2.2	Basic multirate observer setup	8
2.3	More general multirate observer setup	10
2.4	Radiotherapy with a jaw collimator	12
2.5	Proposed jaw-based collimator controller	13
3.1	Optimized multirate tumour position observer	18
3.2	Radiotherapy using a MLC	19
3.3	Proposed MLC controller	20
3.4	Computing MLC leaf offset	21
3.5	Sample leaf reference signal	22
3.6	CTV and PTV used in the simulations	24
3.7	Radiotherapy fraction simulations	25
3.8	Accumulated dose	26
3.9	Radiotherapy fraction simulations with tumour bias change	27
3.10	Accumulated dose with tumour bias change	28
4.1	Proposed MLC controller fraction simulation at 0.2 s	30
4.2	Proposed IMRT dose controller	32

4.3	Calculation of the desired irradiation area	34
4.4	Reference profile used in the simulations	35
4.5	Dose controller fraction simulation	36
4.6	Corrective action of the proposed dose controller	37
4.7	Dose controller accumulated dose	38
5.1	Plant output estimation problem	42
5.2	Geometric interpretation of the projection estimation algorithm	45
5.3	Example estimation problem	51
5.4	Estimator performance over the first 40 s	53
5.5	Estimator performance over the first 1000 s	54
5.6	Prediction error with different breathing flow rate signals	55
5.7	Estimator performance with the scaling improvement	57
5.8	Estimator performance with a tumour bias change	58
5.9	Aliasing effects of the multirate tumour position estimation problem	59
6.1	Control configurations under investigation	63
6.2	Step response definitions	64
6.3	Representative step response of $f(t)$	67
6.4	Factoring $H(s)$	69
6.5	Representative plot of $f(t)$	69
6.6	Lower bound on the step response overshoot	72
6.7	Closed-loop step responses exhibiting overshoot	74
6.8	Closed-loop step responses not exhibiting overshoot	75
6.9	Discrete-time control configuration under investigation	77
6.10	Factoring $H^d(z)$	81

Chapter 1

Introduction

1.1 Radiotherapy Cancer Treatment

Radiation therapy, or radiotherapy, has been used in the treatment of various diseases for over a century. The first use of radiation for therapeutic purposes followed only a few years after the categorization of x-rays by Röntgen in 1895 [1]. Out of this beginning came the field of radiation oncology, which, in general, refers to the use of radiation in the treatment of cancer.

There are three common ways to apply radiation to a cancerous tumour in current day radiotherapy: brachytherapy (implanting a radioactive source in or near the tumour), unsealed source radiotherapy (injecting or ingesting a soluble radioactive substance) and, the most common form of radiotherapy, external beam radiotherapy (applying a radiation beam to the tumour from a source external to the patient) [2]. This work focuses exclusively on external beam radiotherapy and future references to radiotherapy refer to this particular treatment. For medical reasons, the prescribed radiation dose is not delivered in one treatment, but over several increments called *fractions*, with treatment typically progressing at the rate of one fraction per day [3].

The inherent challenge of radiotherapy is to deliver a prescribed radiation dose to a tumour while sparing radiation to healthy tissue surrounding the tumour. To this end, radiotherapy machines collimate the radiation beam in an attempt to shield healthy tissue from radiation. Previous generations of radiotherapy machines used two pairs of metal blocks called “jaws” to decrease radiation of healthy tissue, as shown in Figure 1.1(a). Modern radiotherapy machines employ multiple opposing metal slats to reduce this healthy tissue irradiation, as illustrated in Figure 1.1(b).




Figure 1.1: (a) Collimation with two pairs of jaws; (b) multileaf collimation.

Each slat is called a leaf and the set of leaves is referred to as a multileaf collimator (MLC). Figure 1.2 shows a radiotherapy machine and a multileaf collimator with 80 leaves. The radiation beam is collimated in the treatment head of the radiotherapy machine.

Radiotherapy using a MLC is an example of *conformal* radiotherapy, which, as its name suggests, refers to radiotherapy in which the radiation beam is shaped to match the geometry of the tumour. Traditionally, the collimator leaves are static during the course of a fraction, but motivated by medical and computing advances, interest in *intensity-modulated* radiotherapy (IMRT) has grown over the past decade. IMRT actuates the collimator leaves during the fraction so the post-fraction radiation dose has a non-uniform intensity. By appropriately choosing the radiation profile for each fraction, the aggregate (from all fractions) radiation profile can be made to conform more closely to the three-dimensional geometry of the tumour¹.

To position the collimator leaves for a fraction, a radiation oncologist first defines three volumes related to the tumour geometry: the gross tumour volume (GTV), the clinical target volume (CTV) and the planning target volume (PTV) [5]. The GTV is the palpable volume of malignant tumour growth. The CTV contains the GTV and a tissue volume surrounding the GTV that is suspected to contain subclinical cancerous tissue. Finally, the PTV contains the CTV plus another margin to ensure the prescribed dose is absorbed by the CTV when tumour motion (caused by biological processes such as breathing and the beating of the heart, for example) and errors in aligning the patient with the radiation beam (called “patient setup errors”) are taken into account. These volumes are illustrated in

¹For a comprehensive overview of IMRT see [1].

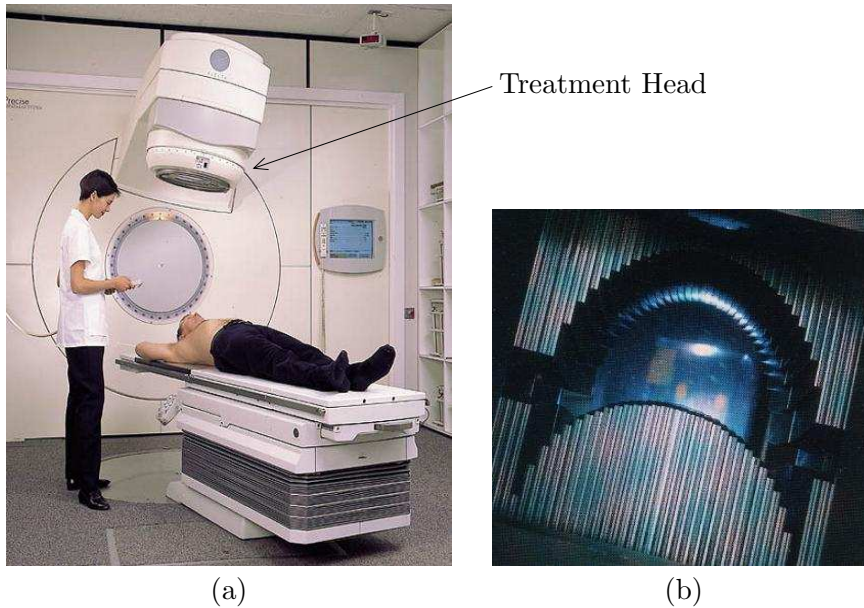


Figure 1.2: (a) External beam radiotherapy machine; (b) multileaf collimator, from [4].

Figure 1.3(a), which shows tumour motion during a fraction. While portions of the PTV are blocked by the collimator leaves as the tumour moves, at no time is any portion of the CTV blocked by the leaves. In essence, the PTV includes a volume of healthy tissue which is sacrificed to ensure radiation delivery to the CTV. This healthy tissue irradiation typically leads to undesirable side effects and constrains the overall efficacy of the fraction.

1.2 Objective

Extending [6] and [7], this work seeks to reduce or eliminate the PTV by tracking the tumour as the fraction progresses and adjusting the collimator leaves appropriately. As illustrated in Figure 1.3(b), the proposed technique will keep the radiation beam tightly focused around the tumour, reducing side effects caused by healthy tissue irradiation. While [6] and [7] focus on a tumour tracking control scheme for the jaw-based collimator, this work develops analogous control schemes for the MLC and IMRT.

The central objective of accounting for tumour position uncertainty is not new. It has been widely accepted in the radiotherapy community that patient setup errors

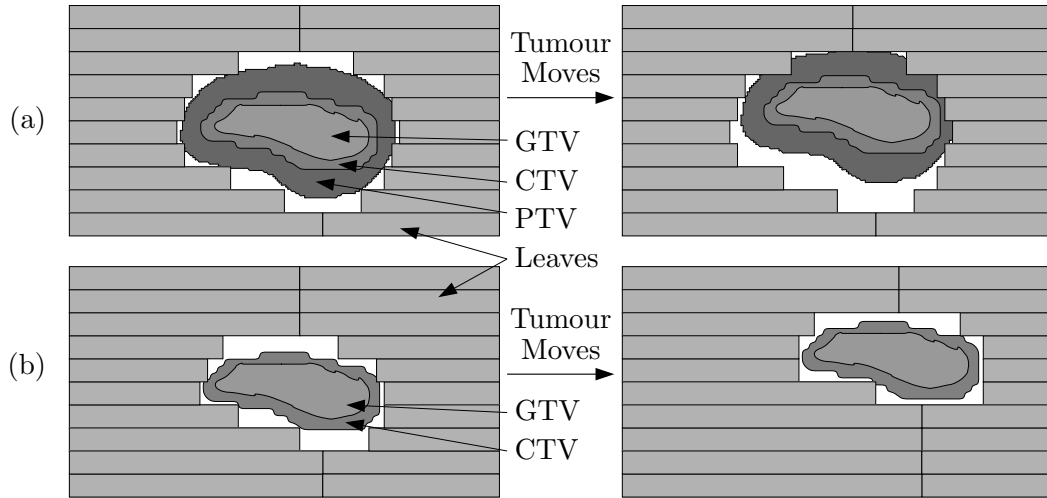


Figure 1.3: *Radiotherapy using a MLC: (a) traditional scheme with static leaves; (b) proposed scheme where the leaves move to track the tumour motion.*

and intrafractional tumour motion are limiting the potential of radiotherapy (e.g., see [9], [10], [11]). Moreover, other researchers have advocated the use of image feedback, but they use the feedback in a limited way, typically just to correct patient setup errors (e.g., see [12]) or to verify tumour position at the conclusion of a fraction (e.g., see [13]). More ambitious approaches monitor the tumour in real time and toggle the radiation beam on and off depending on the tumour location (e.g., see [14]). The novelty of the approach outlined in this work is that *both* patient setup errors and intrafractional tumour motion are accommodated.

1.3 Outline

The next chapter summarizes the work of [6] and [7] in the development of a tumour tracking controller for a jaw-based collimator. This previous work serves as the foundation for the development of a tumour tracking controller for the multileaf collimator in Chapter 3. Chapter 4 extends these results to design a controller for IMRT that incorporates tumour tracking capabilities. This controller provides more accurate delivery of radiation than traditional IMRT.

Chapters 5 and 6 examine theoretical control problems motivated by the work in Chapters 3 and 4. Chapter 5 considers the problem of estimating the output of a system when a known sensor delay prevents its direct measurement. An algorithm is developed that provides asymptotic estimation of the system output under

very mild assumptions on the system. Finally, motivated by the need to prevent opposing collimator leaves from colliding, Chapter 6 explores the nonovershooting step response problem. The chapter will show that for particular systems it is impossible to design a linear time-invariant controller that provides an arbitrarily fast nonovershooting step response.

Chapter 2

Background: Jaw-Based Collimator Controller

This chapter summarizes the work of [6] and [7] in the development of a tumour tracking jaw-based collimator controller. The goal of this controller is to position the collimator jaws to focus the radiation beam tightly around the tumour. The control strategy summarized in this chapter serves as the foundation of the proposed control schemes for the MLC and IMRT in Chapters 3 and 4, respectively.

2.1 High-Level Approach

In order to track the tumour for the jaw-based collimator controller, it is proposed that the therapeutic radiation beam be used as an imaging beam to determine the tumour position¹. One immediate drawback of this approach is the rate at which the images are acquired and processed is significantly slower than the rate of intrafractional tumour motion. To overcome this hindrance, it is also proposed that the breathing flow rate of the patient (which can be measured relatively quickly) be used in conjunction with the image feedback to *estimate* the tumour position. Based on the estimated tumour position, the collimator jaws are actuated so the radiation beam tracks the tumour. A high level diagram of the proposed system is shown in Figure 2.1.

¹Having said that, the control schemes presented in this work are general enough to accommodate any imaging setup, such as the separate radiation beam and x-ray imager employed in [8].

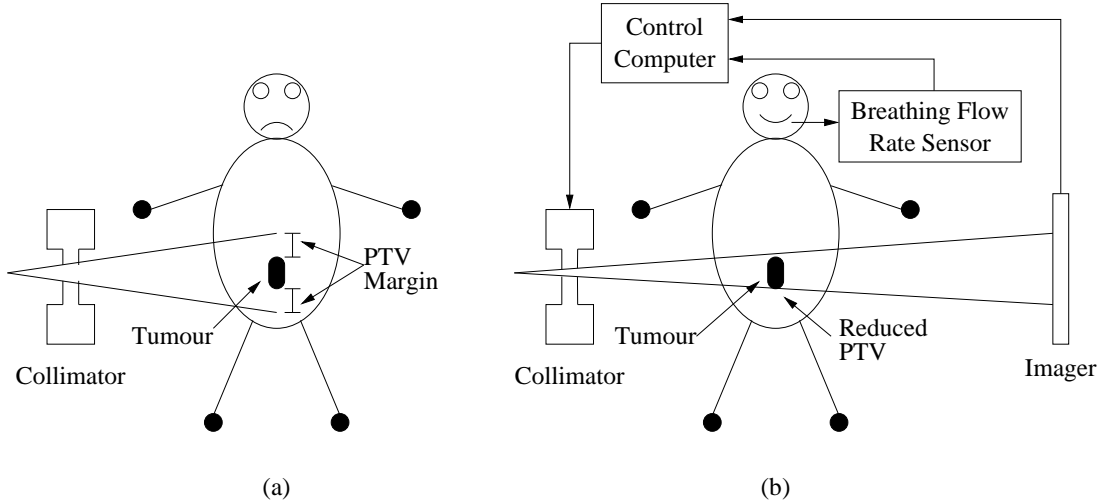


Figure 2.1: (a) Traditional radiotherapy; (b) proposed scheme from [6] and [7] using image feedback and breathing flow rate measurement.

2.2 Multirate Observers

Before describing the jaw-based collimator control scheme in the next section, this section outlines the design of a multirate observer. This observer will be used to estimate the tumour and collimator jaw positions based on relatively fast measurements of the patient’s breathing flow rate and relatively slow measurements from the x-ray imager. Although it may be possible to formulate the jaw-based collimator controller design problem in a general multirate framework, a simpler, more intuitive approach is taken in [6] and [7]. To accommodate the fast measurements of the patient breathing flow rate and the slow measurements of the image feedback, a “sampling timing signal” is used, as explained below.

Consider first the general discrete-time multirate observer setup of Figure 2.2. The linear time-invariant (LTI) system, $P(z)$, is modelled by the state space equations

$$x[k+1] = Ax[k] + Bu[k] \quad (2.1)$$

$$y[k] = Cx[k] \quad (2.2)$$

and is assumed to be known and operating with an h second sampling period. Although the input signal, $u[k]$, is measurable at every time instant, the output signal is only measurable at time instants $k = 0, N, 2N, 3N, \dots$, where $N \geq 1$ is a constant integer. In the figure, the measured signal is denoted by $y_s[k]$ where the

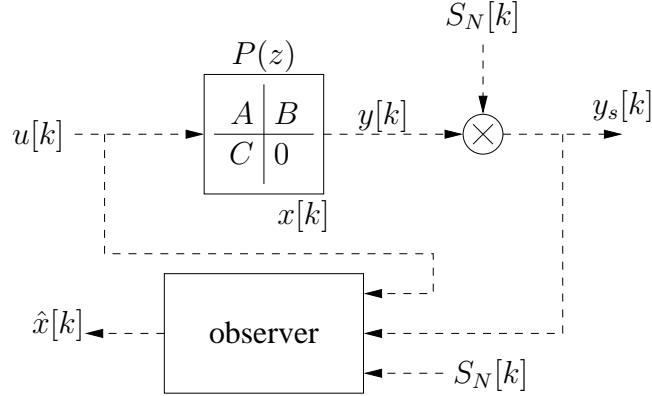


Figure 2.2: Basic multirate observer setup, from [6].

subscript “s” denotes “sampled”. $y_s[k]$ is defined to be zero at time instants when $y[k]$ is not measurable. Thus,

$$y_s[k] = \begin{cases} y[k] & \text{if } k = 0, N, 2N, 3N, \dots \\ 0 & \text{otherwise.} \end{cases} \quad (2.3)$$

The “sampling timing signal”, $S_N[k]$, is I (the identity matrix) at time instants when a measurement is available and 0 otherwise. For the uniform sampling of (2.3),

$$S_N[k] = \begin{cases} I & \text{if } k = 0, N, 2N, 3N, \dots \\ 0 & \text{otherwise} \end{cases} \quad (2.4)$$

and the sampled plant output is

$$y_s[k] = S_N[k] \cdot y[k]. \quad (2.5)$$

This sampling technique has the added benefit of providing for more complicated output sampling arrangements. For example, missing a sample can easily be indicated by setting the appropriate element of $S_N[k]$ to 0. (In the radiotherapy application, missing a sample could be caused by an image processing failure, for example.)

The observer problem of Figure 2.2 is to estimate the state $x[k]$ based on measurements of $u[k]$ and $y_s[k]$. Denote the estimate of $x[k]$ by $\hat{x}[k]$ and consider the observer

$$\hat{x}[k+1] = A\hat{x}[k] + Bu[k] + H(y_s[k] - \hat{y}_s[k]) \quad (2.6)$$

$$\hat{y}_s[k] := S_N[k] \cdot \hat{y}[k] \quad (2.7)$$

$$= S_N[k] \cdot C\hat{x}[k], \quad (2.8)$$

where H is the to-be-determined observer gain. This observer structure is identical to the standard discrete-time observer (e.g., see [21]) with the exception that the state equation (2.6) runs without a corrective term when no output measurements are available (i.e., when $S_N[k] = 0$). As the next theorem shows, this observer provides asymptotic estimation of the system state for an appropriate choice of the observer gain H .

Theorem 2.2 [6]: Consider the plant (2.1)–(2.5), the observer (2.6)–(2.8), and the sampling timing signal (2.4). Assume that the pair

$$(CA^{N-1}, A^N) \tag{2.9}$$

is detectable. Then the matrix H can be chosen such that $A^N - HCA^{N-1}$ has all its eigenvalues in the unit circle. Any such choice of H results in the state estimation error, $e[k] := x[k] - \hat{x}[k]$, tending to zero as $k \rightarrow \infty$ (for arbitrary initial conditions on $x[0]$ and $\hat{x}[0]$ and for arbitrary signal $u[\cdot]$).

Proof [6]: Define $\bar{A}_k := A - HC \cdot S_N[k]$ for $k \geq 0$. Now use (2.1)–(2.5) and (2.6)–(2.8) to obtain

$$\begin{aligned} e[k+1] &= (A - HC \cdot S_N[k]) \cdot e[k] \\ &= \bar{A}_k \cdot e[k], \quad k \geq 0. \end{aligned} \tag{2.10}$$

For the special case of uniform sampling reflected in (2.4) it follows that, for $j \geq 0$,

$$\begin{aligned} e[(j+1)N+1] &= \bar{A}_N \bar{A}_{N-1} \cdots \bar{A}_2 \bar{A}_1 \cdot e[jN+1] \\ &= (A^N - HCA^{N-1}) \cdot e[jN+1]. \end{aligned} \tag{2.11}$$

If (CA^{N-1}, A^N) is detectable, then, by definition, H can be chosen to make $A^N - HCA^{N-1}$ stable. Assume that H is chosen as such, then from (2.11), $e[jN+1] \rightarrow 0$ as $j \rightarrow \infty$. Now observe that

$$e[jN+2] = \bar{A}_{jN+1} \cdot e[jN+1] = A \cdot e[jN+1].$$

Since $e[jN+1] \rightarrow 0$ as $j \rightarrow \infty$, it is also true that $e[jN+2] \rightarrow 0$ as $j \rightarrow \infty$. Similarly,

$$e[jN+3] = \bar{A}_{jN+2} \cdot e[jN+2] = A \cdot e[jN+2],$$

giving $e[jN + 3] \rightarrow 0$ as $j \rightarrow \infty$. Repeat this argument until one obtains $e[jN + N] \rightarrow 0$ as $j \rightarrow \infty$. This proves that $e[k] \rightarrow 0$ as $k \rightarrow \infty$. The result holds for any $e[0]$, and is independent of the input $u[k]$. \square

The observer technique outlined above can be applied to more complicated observer arrangements. Consider the observer setup of Figure 2.3, in which there is an output disturbance, $w[k]$, and stable sensor dynamics modelled by the filter $F(z)$. The observer measures the plant input, $u[k]$, and the sampled output of $F(z)$, denoted by $\bar{y}_{fs}[k]$ (the subscripts “f” and “s” denote that the signal has been filtered by $F(z)$ and sampled by the sampling timing signal $S_N[k]$, respectively). Based on these measurements, the observer goal is to estimate the signal $\bar{y}[k]$. The motivation for this goal is the need to estimate the tumour and collimator jaw positions; this will become apparent in Section 2.3.

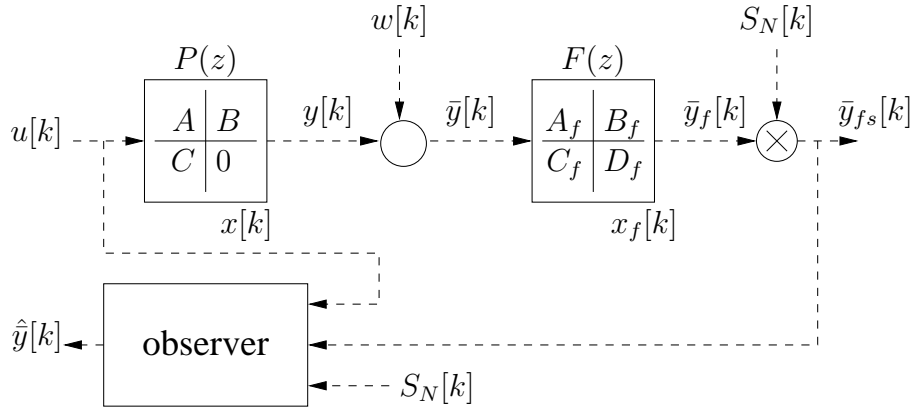


Figure 2.3: *More general multirate observer setup, from [6].*

The observer for this general multirate setup can be designed using the results of Theorem 2.2 by first forming an augmented plant, composed of $P(z)$, $F(z)$, and the disturbance dynamics and then constructing an observer of the form (2.6)-(2.8). Specifically, first assume that the disturbance is generated by the exosystem

$$w[k + 1] = A_w w[k], \quad (2.12)$$

and using the augmented state

$$x_{aug}[k] := [x[k]^T \ x_f[k]^T \ w[k]^T]^T$$

(where superscript T denotes the transpose operator), construct the augmented state-space model

$$x_{aug}[k+1] = A_{aug}x_{aug}[k] + B_{aug}u[k] \quad (2.13)$$

$$\bar{y}_f[k] = C_{aug}x_{aug}[k] \quad (2.14)$$

$$\bar{y}_{fs}[k] = S_N[k] \cdot \bar{y}_f[k], \quad (2.15)$$

where

$$A_{aug} := \begin{bmatrix} A & 0 & 0 \\ B_f C & A_f & B_f \\ 0 & 0 & A_w \end{bmatrix}, \quad B_{aug} := \begin{bmatrix} B \\ 0 \\ 0 \end{bmatrix}, \quad \text{and} \quad C_{aug} := [D_f C \quad C_f \quad D_f].$$

The observer for this augmented plant is

$$\hat{x}_{aug}[k+1] = A_{aug}\hat{x}_{aug}[k] + B_{aug}u[k] + H_{aug}(\bar{y}_{fs}[k] - \hat{\bar{y}}_{fs}[k]) \quad (2.16)$$

$$\begin{aligned} \hat{\bar{y}}_{fs}[k] &:= S_N[k] \cdot \hat{\bar{y}}_f[k] \\ &:= S_N[k] \cdot C_{aug}\hat{x}_{aug}[k], \end{aligned} \quad (2.17)$$

where H_{aug} is the to-be-chosen observer gain. The desired estimate is

$$\hat{y}[k] := \hat{y}[k] + \hat{w}[k] = C_y \hat{x}_{aug}[k], \quad (2.18)$$

where

$$C_y := [C \quad 0 \quad I]. \quad (2.19)$$

Applying Theorem 2.2 to the plant (2.13)-(2.15) with the observer (2.16)-(2.17) and the sampling timing signal (2.4) yields that, at long as $(C_{aug}A_{aug}^{N-1}, A_{aug}^N)$ is detectable, there exists an H_{aug} to place the eigenvalues of $A_{aug}^N - H_{aug}C_{aug}A_{aug}^{N-1}$ inside the unit circle and drive the state estimation error to zero. Consequently,

$$\hat{y}[k] - \bar{y}[k] = C_y \cdot (\hat{x}_{aug}[k] - x_{aug}[k]) \rightarrow 0$$

as $k \rightarrow \infty$.

The next section, which summarizes the proposed jaw-based collimator controller of [6] and [7], will demonstrate the importance of the observer setup of Figure 2.3. Recall that this controller serves as the basis for the MLC and IMRT controllers of Chapters 3 and 4, respectively.

2.3 Jaw-Based Collimator Controller

The goal of the jaw-based collimator controller of [6] and [7] is to use image feedback to track the tumour as the fraction progresses and adjust the radiation beam accordingly (recall Figure 2.1). Figure 2.4 compares traditional radiotherapy in part (a) with the the proposed scheme in part (b). Traditional radiotherapy requires a relatively large PTV to ensure that the CTV is irradiated throughout the fraction in spite of patient setup errors and intrafractional tumour motion. On the other hand, by tracking the tumour during the fraction, the proposed jaw-based collimator controller focuses the radiation beam tightly around the tumour, allowing for a reduction in the PTV.

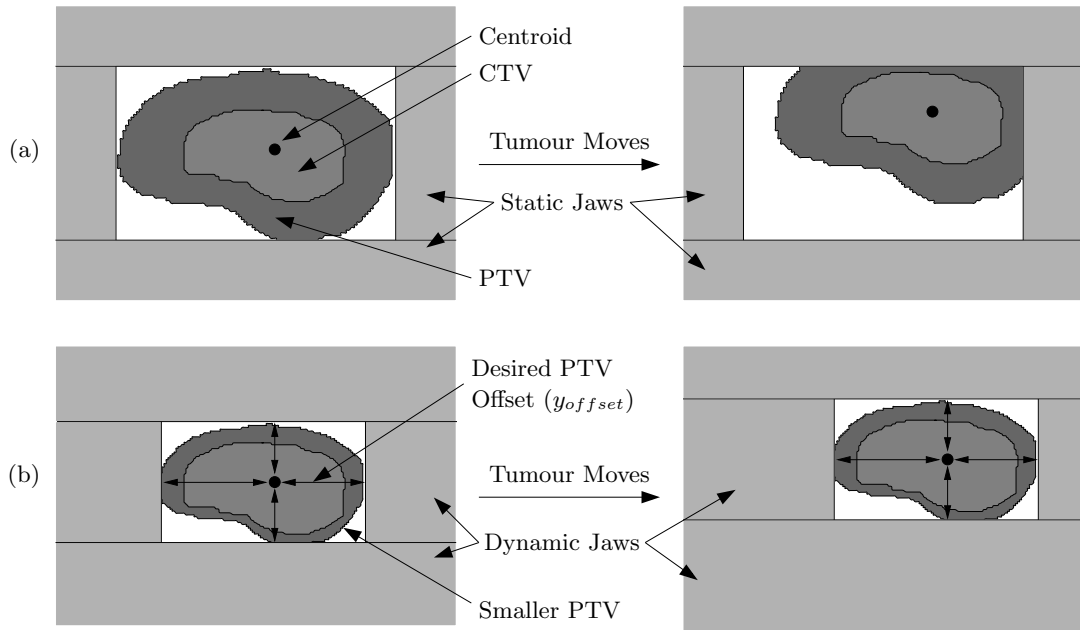


Figure 2.4: (a) Traditional radiotherapy with a jaw collimator; (b) radiotherapy scheme in [6] and [7], with tumour tracking and dynamic jaws.

Although the problem is inherently multivariable, for clarity the controller described in this section considers tumour motion along only one degree of freedom and the control of only one collimator jaw. Also, only rigid translational tumour motion is considered; rotational, scaling and warping effects are ignored. Consequently, the centroid of the tumour uniquely defines the tumour position. It is worth noting that this assumption effectively decouples the inherently multivariable controller; the desired position of each collimator jaw depends on tumour motion

along *one* degree of freedom only. Finally, note from Figure 2.4(b) that each collimator jaw should be positioned a *constant* distance from the tumour centroid. This distance is referred to as the desired PTV offset and can be readily determined from the relationship between the tumour centroid and tumour geometry.

The proposed jaw-based collimator control structure is shown in Figure 2.5. There are two central components in the control setup: estimating the position of the tumour centroid (shown above the bold dash-dot line) and controlling the collimator jaw position, based on the estimated tumour position (below the bold dash-dot line).

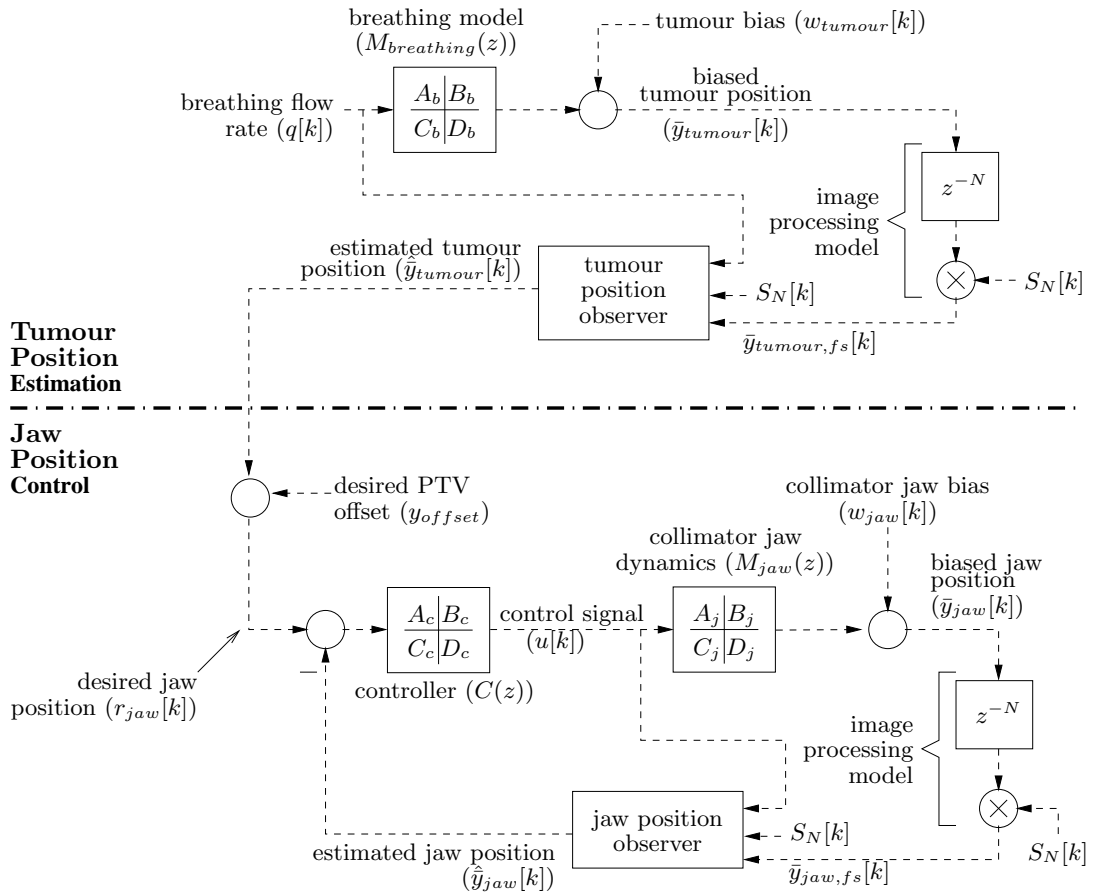


Figure 2.5: Overview of the proposed jaw-based collimator controller, from [6] and [7].

Focus first on the problem of estimating the tumour position. Recall that the control structure uses image feedback to determine the position of the tumour; it is assumed that this image processing imposes a pure delay in determining the tumour position². Thus, it is only possible to directly measure the past position of the tumour, motivating the need to estimate its current position. It is proposed that the breathing flow rate of the patient, $q[k]$, and the delayed and sampled tumour position, $\bar{y}_{tumour,fs}[k]$, be used to estimate the tumour position³. It is assumed that a model relating the breathing flow rate to the tumour position, $\bar{y}_{tumour}[k]$, is known; this model consists of the LTI component, $M_{breathing}(z)$, and an unknown bias, $w_{tumour}[k]$. As previously mentioned, the controller cannot access $\bar{y}_{tumour}[k]$ in real time due to the delay imposed by image processing; it can only access a delayed, slowly sampled version of this signal. To overcome this obstruction, a multirate observer of the same structure as Figure 2.3 is used to estimate $\bar{y}_{tumour}[k]$; this estimate is denoted by $\hat{y}_{tumour}[k]$. Note that the signals $q[k]$, $w_{tumour}[k]$, $\bar{y}_{tumour}[k]$, and $\hat{y}_{tumour}[k]$ operate at the “fast” sampling rate, only the signal used by the observer, $\bar{y}_{tumour,fs}[k]$, operates at the “slow” sampling rate. All sampling is assumed to be uniform, with the slow sampling governed by the signal $S_N[k]$, according to (2.4).

The bottom half of Figure 2.5 shows the feedback loop used to control the collimator jaw position. The desired jaw position, $r_{jaw}[k]$, is simply an offset (the desired PTV offset) from the tumour position (recall Figure 2.4(b)). The “plant” in the feedback loop is the collimator jaw mechanism, assumed to be modelled by the discrete-time transfer function, $M_{jaw}(z)$, and an unknown bias term, $w_{jaw}[k]$. By treating $w_{jaw}[k]$ as an *unknown* constant, it is possible to gain some robustness by guaranteeing that the controller can accommodate translational calibration errors in the jaw. The plant output, $\bar{y}_{jaw}[k]$, is the collimator jaw position in the x-ray imager frame of reference. Note that the scaling and translation factors required to transform from the collimator frame of reference to the x-ray imager frame of reference are included in $M_{jaw}(z)$ and $w_{jaw}[k]$, respectively. Since $\bar{y}_{jaw}[k]$ cannot be measured in real time, it is proposed that another multirate observer of the type described in Figure 2.3 be used to estimate the jaw position. Finally, the controller, $C(z)$, uses this estimate to appropriately position the jaw as the fraction progresses.

The next chapter will describe in detail a control scheme for the MLC that builds upon the jaw-based collimator controller described in this chapter. Note that only a summary of the results of [6] and [7] has been presented in this chapter; refer to these works for complete details.

²Note that the focus of this work is not on image processing and it is assumed throughout that image processing imposes a delay in the control schemes. Image processing algorithms for the radiotherapy project have previously been explored in [7].

³For a good overview on tumour motion caused by respiration, see [11].

Chapter 3

Multileaf Collimator Tumour Tracking Control

Based on the results of [6] and [7] (as summarized in the previous chapter), this chapter develops a tumour tracking control scheme for the MLC¹. Similar to the jaw-based collimator controller, the goal of this scheme is to use the MLC to position the radiation beam tightly around the tumour by tracking the tumour during the fraction. As will be described later in the chapter, the control of the MLC poses some unique challenges compared to the jaw-based collimator controller.

Before proceeding with the design of the MLC controller in Section 3.2, the next section describes a notable optimization to the transient performance of the leaf and tumour position observers, the pivotal components of both the jaw-based collimator and MLC controllers. Simulations detailing the efficacy of the proposed MLC controller are provided in Section 3.3.

3.1 Observer Gain Optimization

Recall from the development of the multirate observer in Section 2.2 (especially Theorem 2.2) that the state estimation error asymptotically tends to zero for any suitable choice of the observer gain H . However, no claims were made about the transient characteristics of the estimation error. Indeed, in [6] and [7], the placement of the observer eigenvalues inside the unit circle was arbitrary and transient performance was largely ignored. Poor transient performance can have devastating

¹This control scheme has been previously proposed in [15].

consequences in the radiotherapy application at hand; even an error in the tumour position of a centimeter can have detrimental effect on the success of a fraction (consider, for example, a tumour near the spinal cord or gall bladder).

To improve the observer transient performance, the calculation of the observer gain, H , will be formulated as a linear quadratic regulator (LQR) optimization problem. The optimization outlined in this section is similar to the LQR optimization of a standard state observer (e.g., see [16]). First, define the observer error at every N^{th} sample as

$$e_N[k] := x_{aug}[kN + 1] - \hat{x}_{aug}[kN + 1].$$

Recall from (2.11) that this error signal satisfies

$$e_N[k + 1] = (A^N - HCA^{N-1})e_N[k]. \quad (3.1)$$

Ideally, the observer gain, H , should be chosen to minimize the cost index

$$J_e = \sum_{k=0}^{\infty} (e_N[k])^T Q_e e_N[k], \quad (3.2)$$

where $Q_e > 0$ is a design parameter. However, to restate the problem in a standard LQR setup, introduce a system whose dynamic matrix is the transpose of that in (3.1),

$$a_N[k + 1] = ((A^N)^T - (A^{N-1})^T C^T H^T) a_N[k], \quad (3.3)$$

where the superscript T denotes the matrix transpose. Note that (3.1) and (3.3) have the same poles; however, (3.3) can be viewed as the “closed-loop” system of the following artificial plant and controller:

$$a_N[k + 1] = (A^N)^T a_N[k] + (A^{N-1})^T C^T v[k] \quad (3.4)$$

$$v[k] = -H^T a_N[k]. \quad (3.5)$$

Viewing (3.4) as the “plant” and (3.5) as the “controller”, LQR optimal control methods can be used to determine the “controller gain” H^T that minimizes the cost index

$$J = \sum_{k=0}^{\infty} [(a_N[k])^T Q a_N[k] + (v[k])^T R v[k]], \quad (3.6)$$

where $Q \geq 0$ and $R > 0$ are design parameters. Since the observer is used to estimate $\bar{y}[k]$, Q is chosen to be $Q = C_y^T \bar{Q} C_y$, where \bar{Q} is a design parameter and

C_y is defined in (2.19). Note that when $Q_e = C_y^T \bar{Q} C_y$ in (3.2), J_e becomes

$$\begin{aligned} J_e &= \sum_{k=0}^{\infty} (e_N[k])^T (C_y)^T \bar{Q} C_y e_N[k] \\ &= \sum_{k=0}^{\infty} (x_{aug}[kN+1] - \hat{x}_{aug}[kN+1])^T (C_y)^T \bar{Q} C_y (x_{aug}[kN+1] - \hat{x}_{aug}[kN+1]) \\ &= \sum_{k=0}^{\infty} (\bar{y}[kN+1] - \hat{y}[kN+1])^T \bar{Q} (\bar{y}[kN+1] - \hat{y}[kN+1]), \end{aligned}$$

and estimation errors of $\bar{y}[k]$ are explicitly penalized. Although choosing $Q = C_y^T \bar{Q} C_y$ in (3.6) does not explicitly penalize estimation errors of $\bar{y}[k]$ in this manner, this choice of Q seems reasonable.

To demonstrate the improved transient performance that is attainable using this optimization approach, consider the tumour position observer in the top half of Figure 2.5. As in [6] and [7], the dynamics of the model relating the breathing flow rate to the tumour position is approximately an integrator:

$$M_{breathing}(z) = \frac{0.0149}{z - 0.990}$$

(for a sampling period of $h = 0.1$ s). The image processing delay, N , is 5 (i.e., the image processing sampling is 5 times slower than the breathing flow rate sampling) and the tumour bias is constant, $A_w = 1$. Based on these models, there are 7 states in the augmented plant and, therefore, 7 observer poles to assign. The observer of [6] and [7] used pole-placement techniques to arbitrarily place the observer poles at $\{0, 0, 0, 0, 0, 0.2, 0.3\}$. Using the optimization technique outlined above, with $\bar{Q} = 10^5$ and $R = 1$, the observer poles are optimally placed at $\{0, 0, 0, 0, 0, 0, 0.951\}$.

Figure 3.1 compares the performance of the two observers over the first 10 s (100 samples using a 0.1 s sampling period). The breathing flow rate, $q[k]$, is assumed to be sinusoidal with a frequency of $1/6$ Hz and is shown in part (a). Part (b) shows the tumour position, $\bar{y}_{tumour}[k]$, and the observer output, $\hat{y}_{tumour}[k]$, for the original observer proposed in [6] and [7] in response to an initial transient and a change in the tumour bias (caused by the patient shifting on the treatment table, for example) of 2.0 cm at $t = 5.1$ s. The estimation error of the original observer is shown in part (c). Finally, parts (d) and (e) show the markedly improved performance of the optimal observer in response to the same initial transient and tumour bias change. Note that the lag between the tumour bias change and the update of $\hat{y}_{tumour}[k]$ is due to the delay and slow sampling of the image processing.




Figure 3.1: *Multirate tumour position observer performance for the observer from [6] and [7] (parts (b) and (c)) and the optimized observer (parts (d) and (e)).*

3.2 Multileaf Collimator Controller

Using the observer gain improvement of the previous section and the original jaw-based collimator controller of [6] and [7], this section describes the design of a tumour-tracking controller for the MLC. Recall from Figure 1.1 that the purpose of using a MLC in radiotherapy, instead of a jaw collimator, is to provide better conformity of the radiation beam to the tumour geometry. The high-level approach is similar to that taken for the jaw-based collimator controller: image feedback is used to track the tumour as the fraction progresses, and the collimator leaves are positioned to keep the radiation beam tightly focused around the PTV (recall Figure 2.1). As shown in Figure 3.2(b), this approach permits a smaller margin between the CTV and the PTV, reducing the amount of healthy tissue irradiation compared to the traditional (fixed collimation) approach of Figure 3.2(a).




Figure 3.2: *Radiotherapy using a MLC: (a) the traditional scheme, where the leaves are fixed; (b) the proposed scheme, where the leaves move in real time to track the tumour motion, allowing a smaller PTV to be used.*

A block diagram of the proposed MLC controller is shown in Figure 3.3; for simplicity, the control of only one leaf is shown in the diagram. The tumour is assumed to move with two degrees of freedom, where the “horizontal” direction is defined to be the direction of travel of the collimator leaf and the “vertical” direction perpendicular to it (see Figure 3.2(b)). Where appropriate, horizontal and vertical signals are denoted with a superscript “h” or “v”, respectively.

At the top of Figure 3.3 are two tumour position observer blocks used to estimate the position of the tumour centroid in the horizontal and vertical directions. (As with the jaw-based collimator controller, only translational and rigid tumour motion is considered.) Each block is a consolidation of the breathing model, image processing delay, and observer of the “tumour position estimation” portion of Figure 2.5. Unlike the jaw-based collimation scheme, the estimated tumour position in *both* directions is required to determine the reference position for each leaf. In each observer block, the breathing model is assumed to be²

$$M_{breathing}(z) = \frac{0.0149}{z - 0.990}$$

²The models and constants used in this section are the same as those chosen in [6].

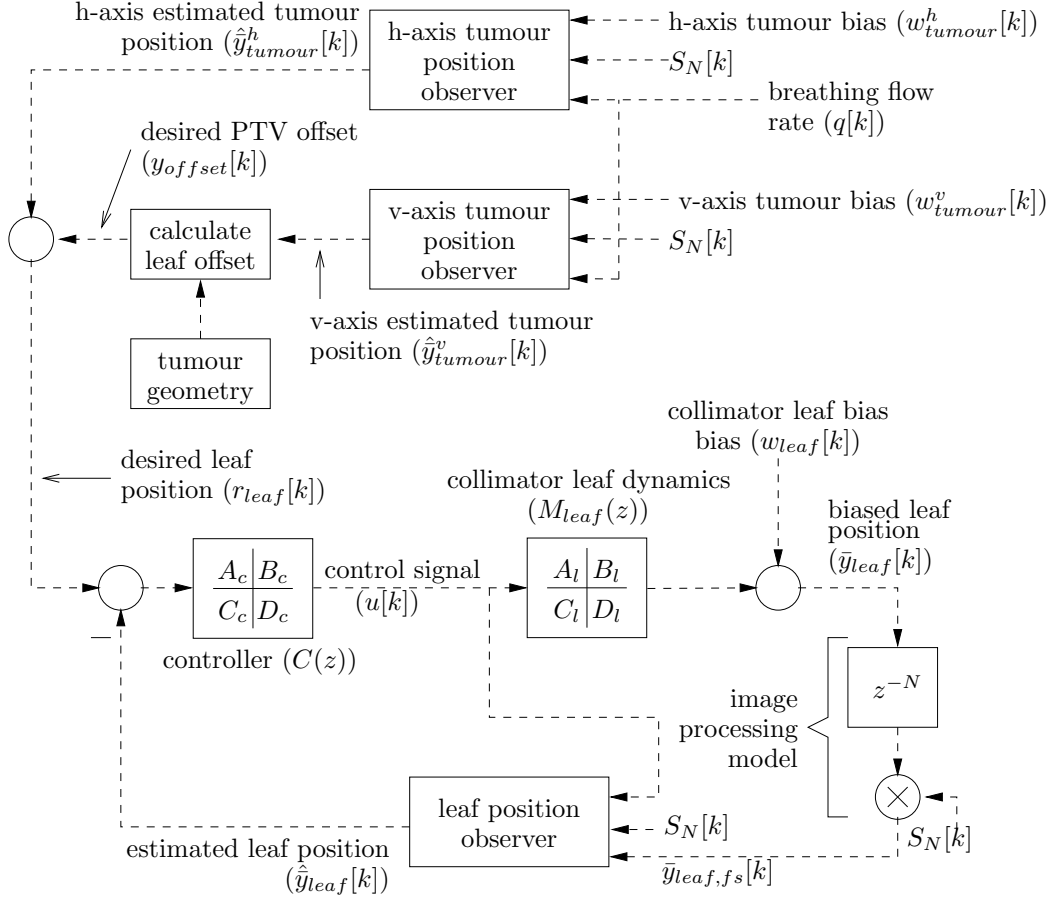


Figure 3.3: Overview of the proposed MLC controller for a single leaf.

with a sampling period of $h = 0.1$ s. The image processing delay, N , is 5 and the tumour bias is constant, $A_w = 1$. Using the LQR optimization technique outlined in the previous section, with $\bar{Q} = 10^5$ and $R = 1$, the observer poles in both tumour position observer blocks are optimally placed at $\{0, 0, 0, 0, 0, 0, 0.951\}$.

The most interesting difference between the jaw-based collimator controller and the MLC controller is the calculation of the desired PTV offset, $y_{offset}[k]$. In the jaw-based collimator controller, this offset is *constant*, reflecting the constant distance between the tumour centroid and the desired jaw position (recall Figure 2.4(b)). However, in the MLC controller, this signal is *time-varying* and depends on the v-axis estimated tumour position and the geometric relationship between the tumour centroid and the PTV. Figure 3.4 illustrates the calculation of $y_{offset}[k]$ in two cases: when the path of the leaf intersects the PTV, in part (a), and when it




Figure 3.4: *Computing leaf offsets: (a) leaf is beside PTV; (b) leaf is below PTV.*

does not intersect the PTV, in part (b) (this determination is made using the v-axis estimated tumour position and the tumour geometry). The leaves marked “A” and “B” in Figure 3.2(b) are examples of the former case and those marked “C” and “D” are examples of the latter. In the part (a) case, the offset is calculated as the maximum (absolute) horizontal distance between the horizontal tumour centroid position and the portion of the PTV that is estimated to be along vertical edge of the leaf. In the part (b) case, it is desired that the left and right opposing leaves touch to completely block the radiation. This is accomplished by setting the offset such that the leaf is in the “middle” of the PTV. When the leaves are below the PTV, they are set to the middle of the PTV when approaching it from the bottom. Similarly, when the leaves are above the PTV, they are set to the middle of the PTV when approaching it from above. The leaves are positioned in this fashion to improve transient leaf performance when the tumour moves into or out of the plane of the collimator leaves (e.g., see the leaf marked “E” in Figure 3.2(b)).

The leaf controller feedback loop in the lower part of Figure 3.3 is almost identical to the jaw position control feedback loop of Figure 2.5; the only difference is the replacement of the jaw dynamics model with the leaf dynamics model, $M_{leaf}(z)$. The collimator leaf is modelled as a first-order lag system with a time constant of 0.1 s and a DC gain of 3. With the “fast” sampling rate of $h = 0.1$ s, the leaf model is

$$M_{leaf}(z) = \frac{1.896}{z - 0.3679}.$$

The LQR observer optimization technique of the previous section was also applied to the leaf position observer. Using $\bar{Q} = 10^5$ and $R = 1$, the leaf position observer eigenvalues were optimally placed at $\{0, 0, 0, 0, 0, 0, 0.0681\}$.

The final component to be designed is the leaf controller, $C(z)$. Fortunately, the separation principle applies to the feedback loop (see [17]) and the leaf controller can be designed independently of the observer. The leaf controller is required to

1. stabilize the loop;
2. reject constant disturbances (i.e., the collimator leaf bias, $w_{leaf}[k]$);
3. provide asymptotic tracking for a periodic reference signal with discrete-time frequency θ ;
4. provide asymptotic tracking for a step reference signal.

The third and fourth requirements are imposed by the nature of the tumour motion: the tumour moves in a roughly sinusoidal pattern with frequency θ , the same frequency as the breathing rate of the patient. The step changes in the reference signal occur when the tumour enters or exits the path of the leaf (e.g., see the leaf marked “E” in Figure 3.2(b)). Figure 3.5 shows a typical leaf reference signal, determined by simulation, which justifies these controller requirements. The final leaf controller structure is

$$C(z) = K \cdot \frac{z - \alpha}{z - 1} \cdot \frac{z^2 - 2(\cos \theta)\omega_n z + \omega_n^2}{z^2 - 2(\cos \theta)z + 1}. \quad (3.7)$$

The pole at $z = 1$ guarantees that requirements 2 and 4 are met, while the poles on the unit circle ensure that requirement 3 is met. The gains and zeros were chosen to provide “good” step response transients, in addition to closed loop stability (requirement 1). Root locus methods were employed to obtain $K = 0.6$, $\omega_n = 0.9$ and $\alpha = 0.4$.




Figure 3.5: *Sample leaf reference signal, $r_{leaf}[k]$.*

3.3 Simulation Results

This section compares representative simulations of the proposed MLC controller to equivalent simulations of traditional multileaf collimation. Figure 3.6(a) shows the CTV used in the simulations; the CTV has a height of 5.1 *cm* and a width of 6.0 *cm*. From a medical perspective, the ideal outcome of the fraction would be the irradiation of the CTV only, without irradiating any of the healthy tissue surrounding the CTV. The breathing flow rate, $q[k]$, is sinusoidal and causes translational sinusoidal motion of the CTV with amplitude 0.56 *cm* and frequency 1/6 *Hz* in both the horizontal and vertical directions for a total translational motion of $\sqrt{(0.56 \text{ cm})^2 + (0.56 \text{ cm})^2} \approx 1.58 \text{ cm}$. The magnitude of this motion is consistent with the 1.6 *cm* respiration-induced movement of the kidney reported in [18]. Moreover, the magnitude of this motion is not unreasonable as intrafractional tumour movement of up to 3.5 *cm* has been documented in abdominal tumours [19]. A PTV margin (distance between the edges of the CTV and PTV) of 0.63 *cm* was used for the traditional multileaf collimation simulations (see Figure 3.6(b)). This margin is just large enough to ensure radiation delivery to the CTV in spite of the sinusoidal tumour motion (assuming constant h-axis and v-axis tumour biases). The PTV margin for the proposed MLC controller was set to 0.0 *cm* (the ideal case where the PTV and CTV perfectly overlap).

The fraction was simulated for 10 *s*, or 100 samples at a sampling period of 0.1 *s*. The radiation dose rate was assumed constant and scaled such that the CTV should receive 10 *s* of radiation. During each sampling interval, any tissue covered by leaves was assumed to receive no radiation, and tissue not covered by leaves was assumed to receive full radiation; effects such as scattering and penumbra (a geometric effect arising from the positioning of the radiation source, MLC and patient) were ignored. The MLC consists of 50 leaves, 25 on each side of the collimator, each with a vertical height of 3.924 *mm*. This choice is consistent with modern MLCs, such as the MLC with a 4.0 *mm* leaf height reported in [20].




Figure 3.6: (a) *CTV* used in the simulations; (b) *PTV* used in the simulation of traditional multileaf collimation.

The proposed and traditional MLC controller are compared in Figure 3.7, which shows each simulated fraction at 0.0 s , 1.4 s , and 4.5 s from the start of the fraction. The leaves in the traditional controller are statically placed so the radiation beam is shaped around the *PTV* at the start of the fraction. Although the traditional controller irradiates the *CTV*, there is significant irradiation of healthy tissue surrounding the *CTV* (shown in white between the *CTV* and collimator leaves). By using image feedback, the proposed MLC controller tracks the tumour and adjusts the collimator leaves to focus the radiation beam tightly around the *CTV*. At the start of the fraction (0.0 s), the leaves of the MLC are fully retracted to obtain a wide image of the cancerous region and determine the tumour position. The vertical blue lines in the proposed MLC controller simulation indicate reference positions for each collimator leaf ($r_{leaf}[k]$).




Figure 3.7: *Radiotherapy fraction simulations: (a) traditional MLC controller; (b) proposed MLC controller.*

The dose delivered to the region in and around the CTV is shown Figure 3.8 for both the traditional (part (a)) and proposed (part (b)) MLC controllers. The first column shows the dose delivered to the patient, with intensity proportional to the delivered radiation dose. The second column shows the magnitude of the dose error, where red indicates tissue that has received too much radiation and blue indicates tissue that has received too little radiation; colour intensity is proportional to the magnitude of the dose error. (In this simulation, very little tissue has been under irradiated and it may be difficult to see these regions.) As expected, the traditional controller succeeds in irradiating the CTV, but at the cost of significant healthy tissue irradiation. As shown in Figure 3.8(b), the proposed MLC controller irradiates the CTV just as well as the traditional controller, but with much less healthy tissue irradiation. The vertical “tails” at the top and bottom of the delivered dose are caused by the leaves not closing quickly enough when the fraction begins (see the $t = 1.4$ s screenshot in Figure 3.7(b)).




Figure 3.8: *Accumulated dose: (a) traditional MLC controller; (b) proposed MLC controller.*

The effectiveness of the proposed MLC controller is even more pronounced in the simulation shown in Figure 3.9. This simulation is identical to the previous simulation, except a change in the vertical tumour bias, $w_{tumour}^v[k]$ (caused by the patient shifting on the treatment table, or a sudden movement of intestinal gas, for example), of -1.0 cm occurs at 5.0 s . This causes the CTV to be partially covered by the collimator leaves in both the traditional and proposed MLC controllers. However, after the bias change propagates through the image processing delay, the proposed MLC controller repositions the collimator leaves so that the radiation beam is again focused around the CTV. Lacking the image feedback of the proposed controller, the traditional MLC controller cannot accommodate this small change in the tumour bias and some of the CTV is covered by the collimator leaves later in the fraction (compare the $t = 5.7\text{ s}$ screenshots in Figure 3.9).




Figure 3.9: *Radiotherapy fraction simulations: (a) traditional MLC control with tumour bias change; (b) proposed MLC controller with tumour bias change.*

Figure 3.10 shows the result of the tumour bias change in the dose delivered to the patient. Most troublesome is the under irradiated region (shown in blue) of the CTV for the traditional MLC controller resulting from the CTV moving beneath the collimator leaves. Due to the image feedback, the proposed MLC controller accounts for the bias change and the results are nearly identical to the previous simulation of Figure 3.8(b).




Figure 3.10: Accumulated dose: (a) traditional MLC controller with bias change; (b) proposed MLC controller with bias change.

Chapter 4

Tumour Dose Control for Intensity-Modulated Radiotherapy

Based on the MLC tumour tracking controller developed in the previous chapter, this chapter proposes a dose controller for IMRT. Recall that in IMRT the collimator leaves are dynamically positioned during a fraction so the post-fraction delivered radiation dose has a non-uniform intensity. By appropriately defining the radiation dose profile for each fraction, the aggregate (from all fractions) applied radiation dose conforms more closely to the actual geometry of the tumour. In addition, denser or “more malignant” portions of the tumour can receive more radiation than surrounding cancerous tissue.

Despite its effectiveness, IMRT suffers from the same problems as standard radiotherapy: intrafractional tumour motion and errors in aligning the patient with the radiation beam at the start of a fraction (patient setup errors) lead to unnecessary healthy tissue irradiation. These, and other problems, are explored in the next section before the proposed dose controller is outlined in Section 4.2¹. In Section 4.3, simulations comparing the IMRT dose controller with standard IMRT illuminate the advantage of the proposed controller. Finally, Section 4.4 describes two important weaknesses of the proposed dose controller; these weaknesses motivate the theoretical control problems of Chapters 5 and 6.

¹This controller has previously been proposed in [22].

4.1 Motivation

Recent studies have confirmed the improved dose coverage to cancerous tissue and decreased irradiation of healthy tissue with IMRT (e.g., see [23], [24]). In spite of this improvement, other studies have cautioned that tumour motion and patient setup errors can significantly degrade the planned delivered radiation dose (e.g., see [25], [26]). Furthermore, leaf and tumour transient motion can adversely affect the delivered radiation dose. For example, Figure 4.1 shows the MLC controller fraction simulation previously described in Section 3.3 at 0.2 s; the left and right leaves have overshoot their reference positions. The overshooting leaves beside the tumour inadvertently block radiation and the leaves above and below the tumour are overlapping, as indicated by the hatched region.

Building on the MLC tumour-tracking controller of the previous chapter, the dose controller developed in this chapter aims to overcome these deficiencies. Using the same image feedback mechanism proposed in Chapters 2 and 3, tumour motion and patient setup errors can be accommodated. In addition, by estimating the applied radiation dose as the fraction progresses, the dose controller can compensate for “cold spots” (under-irradiated tissue) caused by leaf and tumour motion transients. (The more worrisome issue of overlapping leaves is considered in Chapter 6.)




Figure 4.1: *Proposed MLC controller fraction simulation at 0.2 s.*

4.2 Proposed Dose Controller

The high-level approach of the dose controller design is similar to the MLC controller design of Chapter 3: image feedback is used to estimate the positions of the tumour and collimator leaves in real time and the collimator leaves are positioned so the post-fraction dose profile is as close as possible to a reference dose profile. For the purposes of the dose controller, the nomenclature “desired irradiation area” is used to distinguish the desired geometry of the radiation beam from the actual tumour. By appropriately defining the desired irradiation area as the fraction progresses, the delivered dose profile can be varied as the fraction progresses; this is the high-level mechanism employed by the dose controller to create a delivered dose profile that is close to a reference dose profile.

The reference dose profile is assumed to be specified as an $M \times M$ array, where each element of the array is called a voxel. (Although “voxels” typically refer to volume elements, the term is used here to emphasize that the reference dose profile is a two-dimensional projection of an overall three-dimensional prescribed dose profile.) Each element in the array specifies a desired absorbed dose for that voxel, in units called grays (Gy). A gray is the SI unit for absorbed dose and is defined as $1 \text{ Gy} = 1 \text{ J/kg}$.

4.2.1 Control Structure

Figure 4.2 shows a block diagram of the proposed IMRT dose controller; the control of only one collimator leaf is shown for clarity. The two tumour centroid position observers (at the top of the figure) and the leaf controller feedback loop (at the bottom right of the figure) are identical to those of the MLC controller of Figure 3.3; consult Section 3.2 for the design of these blocks.

To compensate for under-irradiated tissue caused by leaf and tumour transient motion, it is proposed that the dose controller estimate the accumulated dose profile as the fraction progresses. Estimating the dose profile is the function of the “calculate hit matrix” and “accumulator” blocks at the bottom left of Figure 4.2. Using the estimated tumour centroid and leaf positions, the “calculate hit matrix” block estimates which voxels were hit with radiation at the current sample time. For simplicity, the estimated tumour and leaf positions are rounded to the nearest voxel to avoid partially covered voxels. The output of this block is a signal of width M^2 where each element is either 1 or 0, indicating whether the corresponding voxel received or did not receive radiation. For simplicity, any voxels covered by leaves receive no radiation, and voxels not covered by leaves receive full radiation at a

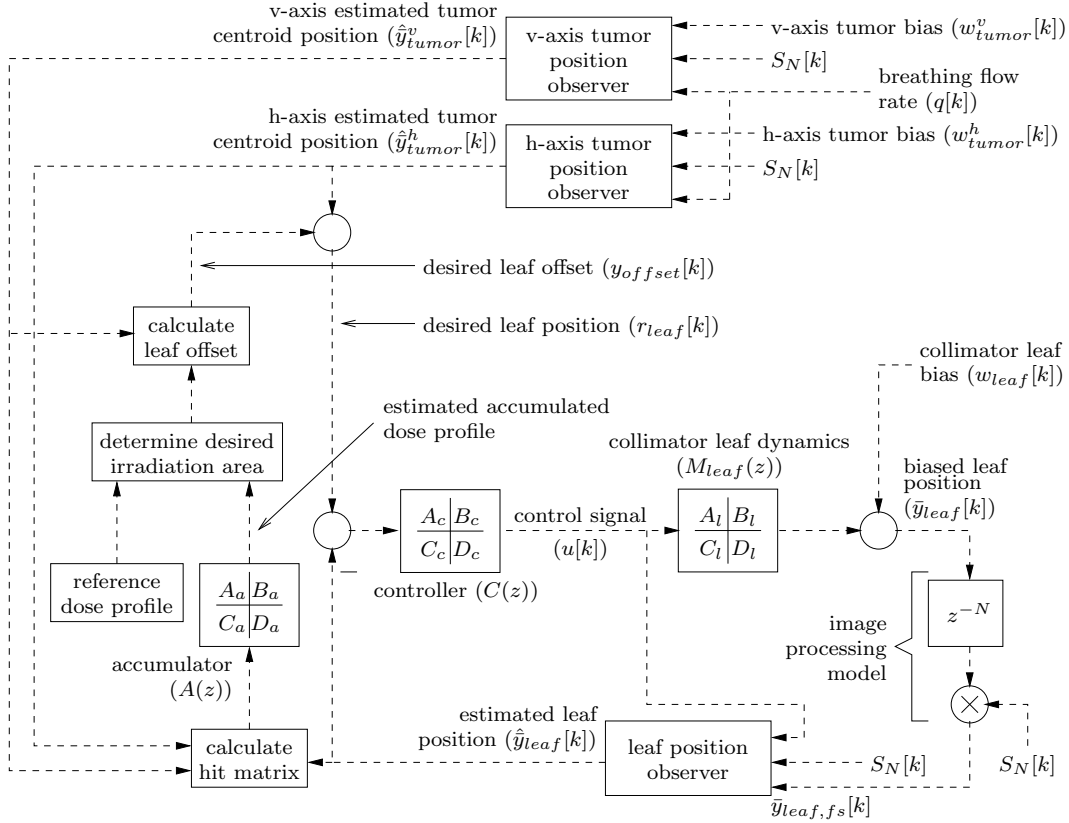


Figure 4.2: Overview of the proposed IMRT dose controller for a single leaf.

constant dose rate of G Gy/sample; as with the MLC controller, effects such as scattering and penumbra are ignored. The “accumulator” block sums the dose received by each voxel as the fraction progresses, i.e., $A(z)$ is a fully decoupled system and each diagonal component has the transfer function $G/(z - 1)$.

4.2.2 Desired Irradiation Area Calculation

The most interesting difference between the MLC and IMRT dose controllers is the calculation of the desired irradiation area. In the former controller, the geometric relationship between the tumour centroid and the desired irradiation area is *constant*, reflecting the goal of positioning the collimator leaves tightly around the tumour. However, in the dose controller design, this relationship is *time-varying* so the dose profile can be varied as the fraction progresses.

The desired irradiation area is defined at each sample time such that the post-

fraction dose profile is as close as possible to the reference dose profile. The high-level technique is to use the reference dose profile and the estimated accumulated dose profile to determine areas that have not received their prescribed dose. Specifically, at the start of the fraction, the dose controller chooses the voxel from the reference dose profile that requires the highest radiation dose; if this voxel is not unique, any voxel from the set can be chosen. This voxel is referred to as the “seed” voxel. Then, at each sample time, the following steps are performed:

1. The estimated dose error profile is calculated by subtracting the estimated accumulated dose profile from the reference dose profile. Voxels still requiring radiation are indicated by a value greater than zero in this profile.
2. Using flood-fill operations, the region of voxels connected to the seed voxel that still require radiation is determined.
3. The total dose (in Gy) still required by the region determined in Step 2 is calculated, and one of the following actions is taken:
 - (a) If the dose in this area is less than a threshold value (indicating that the remaining dose required by the region is negligible), a new seed voxel is chosen. This seed is chosen as the voxel that requires the highest radiation dose from the estimated dose error profile. If this voxel is not unique, any voxel from the set is chosen.
 - i. If a new seed voxel is determined in this step, the algorithm resumes at Step 2.
 - ii. Otherwise, there are no regions remaining that require significant radiation and the fraction is complete.
 - (b) If this area is greater than a threshold value, define the desired irradiation area as the region calculated in Step 2.

The purpose of the threshold value in the above algorithm is to prevent the dose controller from focusing on small regions that require radiation. Without the threshold, it is possible for the controller to focus the radiation beam on a single voxel, likely resulting in over irradiation of tissue surrounding the voxel. In a sense, the threshold value indicates the tradeoff between reducing under-irradiated tissue and reducing over-irradiated tissue.

The calculation of the desired irradiation area is illustrated by the simple example in Figure 4.3. (Note that, to preserve the simplicity of the example, effects

caused by tumour and leaf motion transients are ignored.) Part (a) shows the estimated dose error profile at a particular sample time in the fraction; the elliptical area on the left requires an additional 2 Gy/voxel of radiation to reach its reference dose and the area on the right requires an additional 1 Gy/voxel. A voxel within the ellipse on the left is initially chosen as the seed voxel. All voxels connected to this seed voxel are then determined, as indicated by the white region in part (b). This region becomes the desired irradiation area and will remain so until it receives its full 2 Gy/voxel dose. Afterward, the estimated dose error profile will resemble that shown in part (c) and a new voxel is chosen as the seed voxel. The white region of part (d) shows the new desired irradiation area; this area will be targeted by the radiation beam until it receives its prescribed 1 Gy/voxel dose. Once this region receives its prescribed dose, there are no voxels remaining that require radiation and the fraction is terminated.

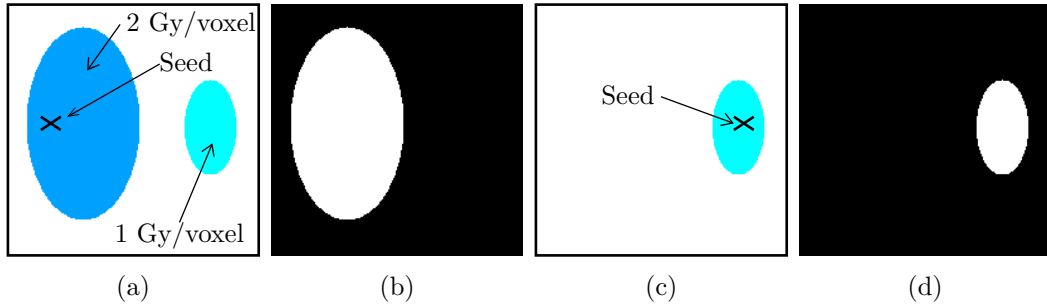


Figure 4.3: *Illustration of the calculation of the desired irradiation area.*

4.3 Simulation Results

This section compares a representative simulation of the IMRT dose controller described in the previous section to that of traditional IMRT. The reference dose profile used in the simulations is a matrix of 545×545 voxels² and is shown in Figure 4.4. Note that the profile consists of two distinct sections; such a profile may be required if a tumour surrounds a critical structure such as the spine or the liver. The prescribed dose to each voxel in the reference profile is indicated by the colour intensity at that voxel, with brighter reds denoting a higher prescribed dose. The maximum prescribed dose of all voxels in the reference dose profile is 2 Gy/voxel.

²These matrix dimensions are used for compatibility with a proof-of-concept experiment.




Figure 4.4: *Reference profile used in the simulations.*

In the simulation, the tumour is perfectly aligned with the radiation beam at the start of the fraction (i.e., there are no patient setup errors) and exhibits sinusoidal translational motion with amplitude 0.56 cm and frequency $1/6\text{ Hz}$ in both the horizontal and vertical directions. (It is worth repeating that intrafraction tumour motion of up to 3.5 cm has been reported in [19].) The MLC consists of 50 leaves, 25 on each side of the collimator, each with a vertical height of 3.924 mm . At each sampling interval, any voxels covered by the collimator leaves were assumed to receive no radiation, and voxels not covered by leaves were assumed to receive radiation at the constant rate of 0.02 Gy/sample (0.2 Gy/s at the sample rate of 0.1 s). Finally, the threshold value in determining the desired irradiation area is 4 Gy . That is, if a connected region still needing radiation requires less than 4 Gy of radiation (as determined by summing the radiation required by each voxel in this region from the estimated dose error profile), it is ignored.

The proposed dose controller is compared with traditional IMRT in the simulation screenshots of Figure 4.5. The screenshots show the collimator leaves and the reference profile in the imager frame of reference. Due to intrafractional tumour motion, traditional IMRT erroneously focuses the radiation beam during the fraction (note especially the regions of healthy tissue in white that receive radiation). Using image feedback, the proposed dose controller estimates the tumour position during the fraction and the radiation beam is focused appropriately. The desired irradiation area is indicated by the green outline in Figure 4.5(b).




Figure 4.5: *Radiotherapy fraction simulations: (a) traditional IMRT controller; (b) proposed dose controller.*

Note that the first two screenshots of both methods in Figure 4.5 show their respective simulation at 1.0 s and 6.8 s from the start of the fraction; the third screenshot is at 11.4 s for traditional IMRT and 17.4 s for the proposed dose controller. In general, due to the feedback mechanism used to correct for under-irradiated tissue, a fraction for the proposed dose controller will be longer than a comparable fraction for traditional IMRT. In the simulations shown in Figure 4.5, the fraction for traditional IMRT requires 15.2 s to complete, while the fraction for the dose controller completes in 25.1 s.

The efficacy of the proposed dose controller is illustrated in Figure 4.6 which shows the simulation of the controller at 0.2 s (part (a)) and 14.4 s (part (b)) from the start of the fraction. At 0.2 s, the right collimator leaves have overshoot their reference positions, causing some voxels to be inadvertently covered by the leaves. By estimating the accumulated dose profile as the fraction progresses, the proposed dose controller compensates for this transient effect and later in the fraction, the




Figure 4.6: *Simulation of the proposed dose controller at: (a) 0.2 seconds; (b) 14.4 seconds.*

region that was covered is targeted by the radiation beam, as shown in Figure 4.6(b). As in Figure 4.1, the collimator leaves above and below the profile have overshoot their reference positions and are overlapping. This issue is considered in detail in Chapter 6.

The dose delivered to the region in and around the tumour is shown in Figure 4.7 for both traditional IMRT (part (a)) and the proposed dose controller (part (b)). The first column shows the dose delivered to the patient, with intensity proportional to the delivered radiation dose. The second column shows the magnitude of the dose error, where red indicates regions that have received too much radiation and blue indicates regions that has received too little; colour intensity is proportional to the magnitude of the error. Due to the tumour motion and the lack of corrective feedback, traditional IMRT results in significant errors between the desired and delivered dose profiles. Especially disconcerting are the regions of tissue which do not receive their full prescribed dose (note that these regions may be difficult to see in Figure 4.7(a)); this under-irradiation may severely affect the success of the fraction. Using image feedback and estimation of the accumulated dose profile during the fraction, the proposed dose controller significantly reduces the number of under-irradiated voxels and provides a tighter bound between the reference and actual dose profiles.




Figure 4.7: *Accumulated dose: (a) traditional IMRT; (b) proposed dose controller.*

4.4 Weaknesses of the Proposed Control System

Despite the demonstrable benefits of the proposed IMRT dose controller, this control scheme (and the MLC controller of Chapter 3) has two significant deficiencies. The first is the assumption that a model relating the breathing flow rate of the patient and the tumour position is known (recall the tumour position estimation portion of Figure 2.5). Although this assumption reduces the complexity of the controller design, it is unrealistic. The model will vary between patients and tumours and will even vary in the same patient from fraction to fraction (e.g., the amount of food in the stomach affects abdominal tumour motion). One approach to overcome this problem is to identify this model and estimate the tumour position as the fraction progresses; such an approach is the focus of the next chapter.

The other notable weakness of the proposed control scheme is the possibility of opposing collimator leaves colliding with each other. The leaf controller, $C(z)$, was not designed to handle this requirement and, in some of the simulations, op-

posing leaves do in fact collide. One method to prevent opposing leaf collisions is to design $C(z)$ so each collimator leaf does not overshoot its reference position; however, it is not clear how to design such a controller, or even if such a controller exists. Chapter 6 explores the problem of designing a controller that provides a fast nonovershooting step response.

Chapter 5

Adaptive Estimation in the Presence of Sensor Delays

To estimate the current tumour position, the control schemes of Chapters 2 through 4 assume the existence of a *known* and *constant* model relating the breathing flow rate of the patient to the tumour position (recall the transfer function $M_{breathing}(z)$ of Figure 2.5). While this assumption makes the control system design more tractable, it is decidedly unrealistic; the model will differ between patients and even over time in the same patient (e.g., changes in bowel volume between fractions can affect the relationship between patient breathing and tumour motion). Although it may be possible to identify this model using standard system identification techniques (e.g., see [27]) before a fraction begins, the increased treatment complexity would not be desirable. This chapter explores a technique to identify this model on-line and estimate the tumour position as the fraction progresses.

Although this chapter focuses on the problem of estimating the tumour position, the results of this chapter are applicable to many diverse estimation problems where a sensor delay prevents direct measurement of the system output. For example, the results of this chapter could be applied to [28], which describes the difficulty of designing an air-fuel ratio controller for gasoline engines due to an air-flow sensor delay, or [29], which details the challenges of fermenting ethanol using baker's yeast with a sensor delay in ethanol measurement.

The next section defines the general plant output estimation problem investigated in this chapter; this problem is the single-rate version of the multirate tumour position estimation problem in the control schemes of Chapters 2 through 4.

5.1 Problem Formation

The topology of the plant output estimation problem is shown below in Figure 5.1(b); this setup is similar to the tumour position estimation setup duplicated in Figure 5.1(a) with the exception of sampling. To reduce the complexity of the estimation problem, the plant output estimation setup assumes *single rate* sampling. (The applicability of the results of this chapter to the tumour position estimation problem is further discussed in Section 5.4.) In Figure 5.1(b), a sensor delay (modelled as a pure delay of N steps) prevents direct measurement of the plant output, $\bar{y}[k]$. The technique developed in this chapter uses the plant input, $u[k]$, and the delayed plant output, $\bar{y}[k - N]$, to produce an estimate of the plant output, denoted by $\hat{y}[k]$, such that

$$\lim_{k \rightarrow \infty} (\bar{y}[k] - \hat{y}[k]) = 0.$$

The plant is assumed to consist of an *unknown* linear, time-invariant and proper system, $P(z)$, and an *unknown* constant bias, w . The only other assumption on the plant is that an upper bound on the number of poles and zeros of $P(z)$ is known.

Since $P(z)$ is linear, time-invariant and proper, it can be written in the difference equation form

$$y[k] = \sum_{i=1}^n a_i y[k - i] + \sum_{i=0}^m b_i u[k - i], \quad k \geq t_0 \quad (5.1)$$

where the coefficients a_i and b_i are *unknown* constants. Note that if an upper bound on the number of poles and zeros is known, the plant model can be written in such a form. Now, from (5.1), $y[k - 1]$ can be written as a linear combination of

$$\{y[k - 2], y[k - 3], \dots, y[k - n - 1]\}$$

and

$$\{u[k - 1], u[k - 2], \dots, u[k - m - 1]\}.$$

In general, $y[k - i]$ can be written as a linear combination of

$$\{y[k - i - 1], y[k - i - 2], \dots, y[k - i - n]\}$$

and

$$\{u[k - i], u[k - i - 1], \dots, u[k - i - m]\}$$

for

$$i = \{1, 2, \dots, N - 1\}.$$

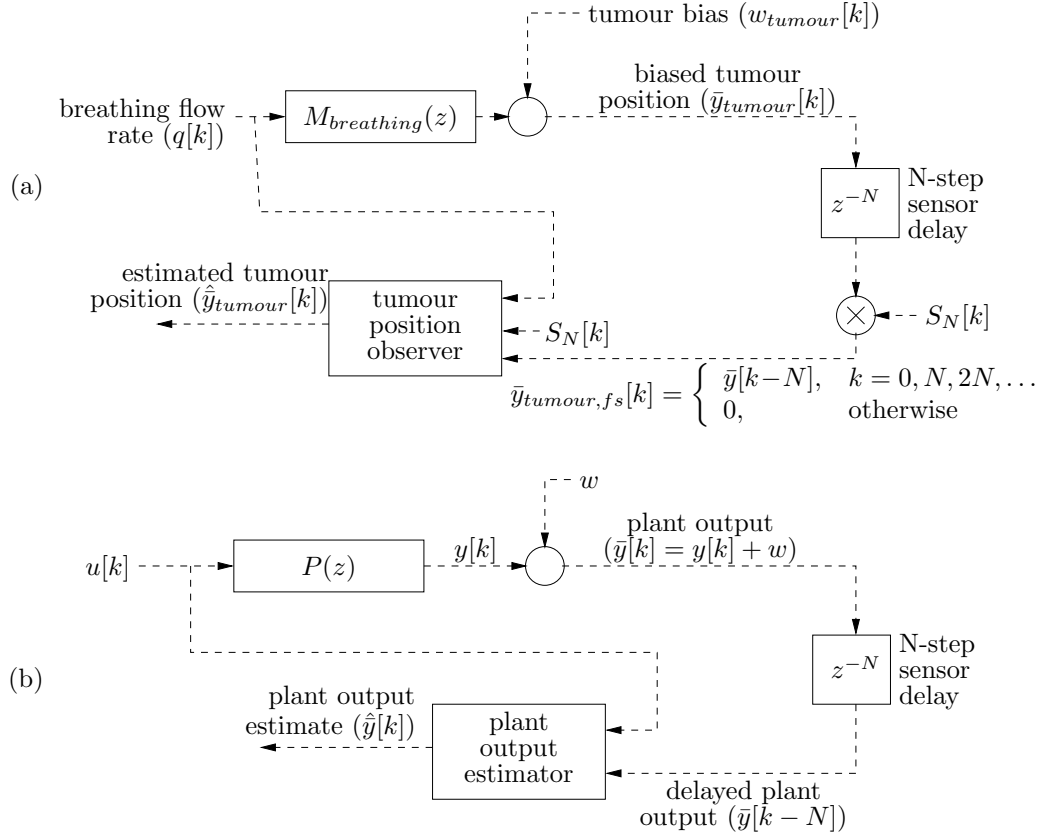


Figure 5.1: (a) The multirate tumour position estimation problem; (b) simplified single rate plant output estimation problem considered in this chapter.

Therefore, (5.1) can be rewritten as

$$y[k] = \sum_{i=0}^{\bar{n}} \bar{a}_i y[k - N - i] + \sum_{i=0}^{\bar{m}} \bar{b}_i u[k - i], \quad k \geq \bar{t}_0,$$

where the constants, \bar{n} , \bar{m} , \bar{a}_i , \bar{b}_i and \bar{t}_0 are functions of n , m , a_i , b_i and t_0 . The biased plant output is then given by

$$\begin{aligned}
\bar{y}[k] &= y[k] + w \\
&= \sum_{i=0}^{\bar{n}} \bar{a}_i y[k - N - i] + \sum_{i=0}^{\bar{m}} \bar{b}_i u[k - i] + w \\
&= \sum_{i=0}^{\bar{n}} \bar{a}_i (\bar{y}[k - N - i] - w) + \sum_{i=0}^{\bar{m}} \bar{b}_i u[k - i] + w \\
&= \sum_{i=0}^{\bar{n}} \bar{a}_i \bar{y}[k - N - i] + \sum_{i=0}^{\bar{m}} \bar{b}_i u[k - i] + \underbrace{w - w \cdot \sum_{i=0}^{\bar{n}} \bar{a}_i}_c \\
&= \underbrace{[\bar{y}[k - N] \quad \cdots \quad \bar{y}[k - N - \bar{n}] \quad u[k] \quad \cdots \quad u[k - \bar{m}] \quad 1]}_{\phi[k]^T} \underbrace{\begin{bmatrix} \bar{a}_0 \\ \vdots \\ \bar{a}_{\bar{n}} \\ \bar{b}_0 \\ \vdots \\ \bar{b}_{\bar{m}} \\ c \end{bmatrix}}_{\theta^*} \\
&= \phi[k]^T \theta^*, \quad k \geq \bar{t}_0.
\end{aligned} \tag{5.2}$$

Note that $\phi[k]$ consists of measurable signals and θ^* consists of unknown constants; this is the standard form used in many on-line estimation techniques (a good overview of standard on-line estimation results can be found in [30]). The basic premise behind the estimation technique outlined in this chapter is to estimate θ^* by measuring $\phi[k]$ at every time instant and from this estimate, calculate the estimate of the plant output, $\hat{y}[k]$ ¹. Specifically, denoting the current estimate of

¹ Although the plant is assumed to consist of a linear time-invariant component, $P(z)$, the results of this chapter are applicable to a more general class of systems. Specifically, the estimation technique developed in this chapter is applicable to any system where $\bar{y}[k]$ can be written as

$$\bar{y}[k] = \phi[k]^T \theta^*,$$

where θ^* is a vector of constant parameters and $\phi[k]$ is a linear or nonlinear combination of the data

$$D[k] := \{\bar{y}[k - N], \bar{y}[k - N - 1], \dots, u[k], u[k - 1], \dots\}.$$

θ^* by $\hat{\theta}[k]$, define the *plant output estimate* as²

$$\hat{y}[k] := \phi[k]^T \hat{\theta}[k - 1].$$

Also, define the *parameter estimate error* as

$$\tilde{\theta}[k] := \hat{\theta}[k] - \theta^*,$$

and the *prediction error* as

$$\begin{aligned} e[k] &:= \bar{y}[k] - \hat{y}[k] \\ &= \phi[k]^T \theta^* - \phi[k]^T \hat{\theta}[k - 1] \\ &= -\phi[k]^T \tilde{\theta}[k - 1]. \end{aligned}$$

The next section outlines a projection algorithm technique which provides asymptotic plant output estimation, i.e.,

$$\lim_{k \rightarrow \infty} e[k] = 0.$$

5.2 Adaptive Estimation Technique

Based on the plant model (5.2), this section develops a projection algorithm technique to estimate θ^* based on the standard projection algorithm for on-line parameter estimation (e.g., see [31]). The lemma and theorem proofs in this section are derived from analogous results for the standard projection algorithm in [32] and [33]. The basic intuition behind the algorithm is as follows: given a new delayed output, $\bar{y}[k - N]$ and a vector of input/output data, $\phi[k - N]$, all possible values of θ^* satisfying the model (5.2) lie on the hypersurface

$$H := \{ \theta : \bar{y}[k - N] = \phi[k - N]^T \theta \}. \quad (5.3)$$

Among these candidate values for θ^* , the new estimate, denoted by $\hat{\theta}[k]$, is chosen as the one closest to $\hat{\theta}[k - N - 1]$. Hence, $\hat{\theta}[k]$ is chosen as the value of θ which minimizes

$$J(\theta) = \frac{1}{2} \left\| \theta - \hat{\theta}[k - N - 1] \right\|_2^2 \quad (5.4)$$

subject to $\theta \in H$ where (5.4) uses the standard Holder 2-norm. A geometric interpretation of this technique is illustrated for a second-order case in Figure 5.2. The following lemma shows the minimizing θ is unique and easy to express.

²Although it may seem more natural to use $\hat{\theta}[k]$ in this equation, developments later in the chapter require $\hat{\theta}[k - 1]$.




Figure 5.2: *Geometric interpretation of the projection estimation algorithm.*

Lemma 5.2:

$$\arg \min_{\theta \in H} J(\theta) = \begin{cases} \hat{\theta}[k-N-1], & \phi[k-N] = 0 \\ \hat{\theta}[k-N-1] + \frac{\phi[k-N]}{\phi[k-N]^T \phi[k-N]} \left(\bar{y}[k-N] - \phi[k-N]^T \hat{\theta}[k-N-1] \right), & \text{otherwise.} \end{cases}$$

Proof: Since H is closed and convex,

$$\hat{\theta}[k] := \arg \min_{\theta \in H} J(\theta)$$

exists and is unique.

If $\phi[k-N] = 0$, then $\bar{y}[k-N]$ must be zero as well, so from (5.3), H is the entire space and

$$\theta = \hat{\theta}[k-N-1]$$

minimizes (5.4).

Now, suppose that $\phi[k-N] \neq 0$. Introduce a Lagrange multiplier for the constraint

$\theta \in H$ and define

$$\begin{aligned} J_l(\hat{\theta}[k], \lambda) &:= \frac{1}{2} \left\| \hat{\theta}[k] - \hat{\theta}[k - N - 1] \right\|_2^2 + \lambda \left(\bar{y}[k - N] - \phi[k - N]^T \hat{\theta}[k] \right) \\ &= \frac{1}{2} \left(\hat{\theta}[k] - \hat{\theta}[k - N - 1] \right)^T \left(\hat{\theta}[k] - \hat{\theta}[k - N - 1] \right) \\ &\quad + \lambda \left(\bar{y}[k - N] - \phi[k - N]^T \hat{\theta}[k] \right). \end{aligned}$$

The necessary conditions for a minimum are

$$\frac{\partial J_l(\hat{\theta}[k], \lambda)}{\partial \hat{\theta}[k]} = 0 \quad \& \quad \frac{\partial J_l(\hat{\theta}[k], \lambda)}{\partial \lambda} = 0.$$

Now,

$$\begin{aligned} J_l(\hat{\theta}[k], \lambda) &= \frac{1}{2} \sum_{i=1}^p \left(\hat{\theta}_i[k] - \hat{\theta}_i[k - N - 1] \right)^2 + \\ &\quad \lambda \left(\bar{y}[k - N] - \sum_{i=1}^p \phi_i[k - N] \hat{\theta}_i[k] \right), \end{aligned}$$

so

$$\frac{\partial J_l(\hat{\theta}[k], \lambda)}{\partial \hat{\theta}_i[k]} = \hat{\theta}_i[k] - \hat{\theta}_i[k - N - 1] - \lambda \phi_i[k - N] = 0,$$

which implies

$$\frac{\partial J_l(\hat{\theta}[k], \lambda)}{\partial \hat{\theta}[k]} = \hat{\theta}[k] - \hat{\theta}[k - N - 1] - \lambda \phi[k - N] = 0. \quad (5.5)$$

Also,

$$\frac{\partial J_l(\hat{\theta}[k], \lambda)}{\partial \lambda} = \bar{y}[k - N] - \phi[k - N]^T \hat{\theta}[k] = 0. \quad (5.6)$$

Substituting (5.5) into (5.6) and rearranging gives

$$\lambda = \frac{\bar{y}[k - N] - \phi[k - N]^T \hat{\theta}[k - N - 1]}{\phi[k - N]^T \phi[k - N]}.$$

Substituting this back into (5.5) gives

$$\hat{\theta}[k] = \hat{\theta}[k - N - 1] + \frac{\phi[k - N]}{\phi[k - N]^T \phi[k - N]} \left(\bar{y}[k - N] - \phi[k - N]^T \hat{\theta}[k - N - 1] \right). \quad \square$$

A potential problem with the above algorithm is that $\phi[k - N]$ may get close (or equal) to zero, potentially causing numerical problems³. To avoid this problem, a small constant, c , is added to the denominator of the algorithm leading to the following slightly modified version of the algorithm.

Projection Algorithm for Plants with Sensor Delays. *With $0 < a < 2$, $c > 0$, let*

$$\hat{\theta}[k] = \hat{\theta}[k - N - 1] + \frac{a\phi[k - N]}{c + \phi[k - N]^T\phi[k - N]} \left(\bar{y}[k - N] - \phi[k - N]^T\hat{\theta}[k - N - 1] \right),$$

$$k \geq \bar{t}_0 + N + 1 \quad (5.7)$$

with $\hat{\theta}[\bar{t}_0] = \hat{\theta}[\bar{t}_0 + 1] = \dots = \hat{\theta}[\bar{t}_0 + N] = \theta_o$.

The next theorem, the main result of this chapter, shows that this projection algorithm provides asymptotic estimation of the plant output for bounded input, bounded output (BIBO) plants. In other words, under the mild assumptions that the *unknown* plant is linear with an offset, time-invariant, stable, and an upper bound on the number of poles and zeros of the plant is known, this projection algorithm can asymptotically estimate the output of the plant in spite of a known sensor delay.

Theorem 5.2: Suppose that the projection algorithm for plants with sensor delays is applied to the system of Figure 5.1(b). Then

(i) $\left\| \hat{\theta}[k] - \theta^* \right\|_2 \leq \left\| \hat{\theta}[k - N - 1] - \theta^* \right\|_2 \leq \dots \leq \|\theta_0 - \theta^*\|_2, \quad k \geq \bar{t}_0 + N + 1$
(i.e., every $(N + 1)^{\text{th}}$ parameter estimate error is nonincreasing);

(ii) If $P(z)$ is BIBO stable and the input to the system, $u[k]$, is bounded then,

$$\lim_{k \rightarrow \infty} e[k] = \lim_{k \rightarrow \infty} (\bar{y}[k] - \hat{y}[k]) = 0$$

(i.e., the algorithm provides asymptotic estimation of the plant output).

³Having said that, recall from (5.2) that the last element of $\phi[k - N]$ is the constant 1; thus, $\phi[k - N]^T\phi[k - N] \geq 1$. However, in other formulations of this estimation problem (e.g., the basic estimation problem of Figure 5.1(b) with the bias term, w , removed), it is possible for $\phi[k - N] = 0$.

Proof: Subtracting θ^* from both sides of the update equation (5.7) gives

$$\begin{aligned}\tilde{\theta}[k] &= \tilde{\theta}[k - N - 1] + \frac{a\phi[k - N]}{c + \phi[k - N]^T\phi[k - N]} \times \\ &\quad \left(\bar{y}[k - N] - \phi[k - N]^T\hat{\theta}[k - N - 1] \right) \\ &= \tilde{\theta}[k - N - 1] + \frac{a\phi[k - N]}{c + \phi[k - N]^T\phi[k - N]} e[k - N].\end{aligned}$$

Now,

$$\begin{aligned}\|\tilde{\theta}[k]\|_2^2 &= \|\tilde{\theta}[k - N - 1]\|_2^2 + 2a \frac{\phi[k - N]^T\tilde{\theta}[k - N - 1]}{c + \phi[k - N]^T\phi[k - N]} e[k - N] + \\ &\quad a^2 \frac{\phi[k - N]^T\phi[k - N]}{(c + \phi[k - N]^T\phi[k - N])^2} e[k - N]^2 \\ &= \|\tilde{\theta}[k - N - 1]\|_2^2 - 2a \frac{1}{c + \phi[k - N]^T\phi[k - N]} e[k - N]^2 + \\ &\quad a^2 \frac{\phi[k - N]^T\phi[k - N]}{(c + \phi[k - N]^T\phi[k - N])^2} e[k - N]^2 \\ &= \|\tilde{\theta}[k - N - 1]\|_2^2 + \\ &\quad \underbrace{a \left(\overbrace{\frac{\phi[k - N]^T\phi[k - N]}{c + \phi[k - N]^T\phi[k - N]} - 2}^{\in[0,a]} \right)}_{<0} \underbrace{\left(\frac{e[k - N]^2}{c + \phi[k - N]^T\phi[k - N]} \right)}_{\geq 0} \quad (5.8) \\ \|\tilde{\theta}[k]\|_2^2 &\leq \|\tilde{\theta}[k - N - 1]\|_2^2,\end{aligned}$$

which implies (i).

Next, for $k \geq \bar{t}_0 + N + 1$, denote k by

$$k = \bar{t}_0 + \gamma(N + 1) + \tau,$$

where $\gamma \in \{1, 2, 3, \dots\}$ and $\tau \in \{0, 1, 2, \dots, N\}$ and, for convenience, define

$$\bar{\phi}[k] := \phi[k]^T\phi[k] \geq 0.$$

Then, iterating (5.8) gives

$$\begin{aligned}
\left\| \tilde{\theta}[k] \right\|_2^2 &= \left\| \tilde{\theta}[\bar{t}_0 + \gamma(N+1) + \tau] \right\|_2^2 \\
&= \left\| \tilde{\theta}[\bar{t}_0 + (\gamma-1)(N+1) + \tau] \right\|_2^2 + a \left(a \frac{\bar{\phi}[\bar{t}_0 + (\gamma-1)(N+1) + \tau + 1]}{c + \bar{\phi}[\bar{t}_0 + (\gamma-1)(N+1) + \tau + 1]} - 2 \right) \times \\
&\quad \left(\frac{e[\bar{t}_0 + (\gamma-1)(N+1) + \tau + 1]^2}{c + \bar{\phi}[\bar{t}_0 + (\gamma-1)(N+1) + \tau + 1]} \right) \\
&= \dots \\
&= \left\| \tilde{\theta}[\bar{t}_0 + \tau] \right\|_2^2 + \sum_{i=0}^{\gamma-1} a \left(\frac{a\bar{\phi}[\bar{t}_0 + i(N+1) + \tau + 1]}{c + \bar{\phi}[\bar{t}_0 + i(N+1) + \tau + 1]} - 2 \right) \times \\
&\quad \left(\frac{e[\bar{t}_0 + i(N+1) + \tau + 1]^2}{c + \bar{\phi}[\bar{t}_0 + i(N+1) + \tau + 1]} \right) \\
&= \left\| \tilde{\theta}[\bar{t}_0] \right\|_2^2 + \sum_{i=0}^{\gamma-1} a \underbrace{\left(\frac{a\bar{\phi}[\bar{t}_0 + i(N+1) + \tau + 1]}{c + \bar{\phi}[\bar{t}_0 + i(N+1) + \tau + 1]} - 2 \right)}_{\in [-2, a-2]} \times \\
&\quad \left(\frac{e[\bar{t}_0 + i(N+1) + \tau + 1]^2}{c + \bar{\phi}[\bar{t}_0 + i(N+1) + \tau + 1]} \right), \tag{5.9}
\end{aligned}$$

where the last equality arises from the initial conditions of the algorithm in (5.7). Now, from (5.9), the following inequality results:

$$0 \leq \left\| \tilde{\theta}[k] \right\|_2^2 \leq \left\| \tilde{\theta}[\bar{t}_0] \right\|_2^2 + a(a-2) \sum_{i=0}^{\gamma-1} \frac{e[\bar{t}_0 + i(N+1) + \tau + 1]^2}{c + \bar{\phi}[\bar{t}_0 + i(N+1) + \tau + 1]}.$$

Hence, $\forall \gamma$,

$$\sum_{i=0}^{\gamma-1} \frac{e[\bar{t}_0 + i(N+1) + \tau + 1]^2}{c + \bar{\phi}[\bar{t}_0 + i(N+1) + \tau + 1]} \leq \frac{1}{a(2-a)} \left\| \tilde{\theta}[\bar{t}_0] \right\|_2^2,$$

which implies

$$\lim_{\gamma \rightarrow \infty} \frac{e[\bar{t}_0 + \gamma(N+1) + \tau + 1]}{(c + \bar{\phi}[\bar{t}_0 + \gamma(N+1) + \tau + 1])^{1/2}} = 0. \tag{5.10}$$

Now, assume that $P(z)$ is BIBO stable and the system input, $u[k]$, is bounded. Then, $\phi[k]$ is bounded (recall from (5.2) that the elements of $\phi[k]$ are constants, plant inputs and plant outputs) and the denominator of (5.10) is finite. Therefore,

$$\lim_{k \rightarrow \infty} e[k] = 0,$$

and (ii) holds. □

Note that Theorem 5.2 makes no claims on the convergence of the projection algorithm and, in general, $\hat{\theta}[k]$ will not converge to θ^* (i.e., the plant is not fully identified)⁴. Some insight into this behaviour can be gained by considering a high order plant with a constant frequency sinusoidal input; this input does not yield “rich” enough data to fully identify the plant. In the standard projection algorithm, $\hat{\theta}[k]$ will converge to θ^* if the input, $u[k]$, meets a persistent excitation criterion [35]; it is unknown if a similar result applies to the projection algorithm for plants with sensor delays. However, such a result is not the focus of this chapter.

5.3 Simulation Results

This section explores the modified tumour position estimation problem of Figure 5.3; this problem is identical to the tumour position estimation problem of Chapters 2 through 4, except that single rate sampling is assumed. The biased tumour position, $\bar{y}_{tumour}[k]$, is related to the breathing flow rate of the patient, $q[k]$, by a dynamic component, $M_{breathing}(z)$, and a bias term, $w_{tumour}[k]$. $M_{breathing}(z)$ is assumed to be a stable first-order lag system of the form

$$M_{breathing}(z) = \frac{\beta}{z - \alpha},$$

where $\beta > 0$ and $|\alpha| < 1$ are *unknown* constants. A sensor (image processing) delay of $N = 5$ steps prevents direct measurement of the tumour position.

⁴Moreover, in the standard projection algorithm, it is possible for $\hat{\theta}[k]$ to not converge at all [34].

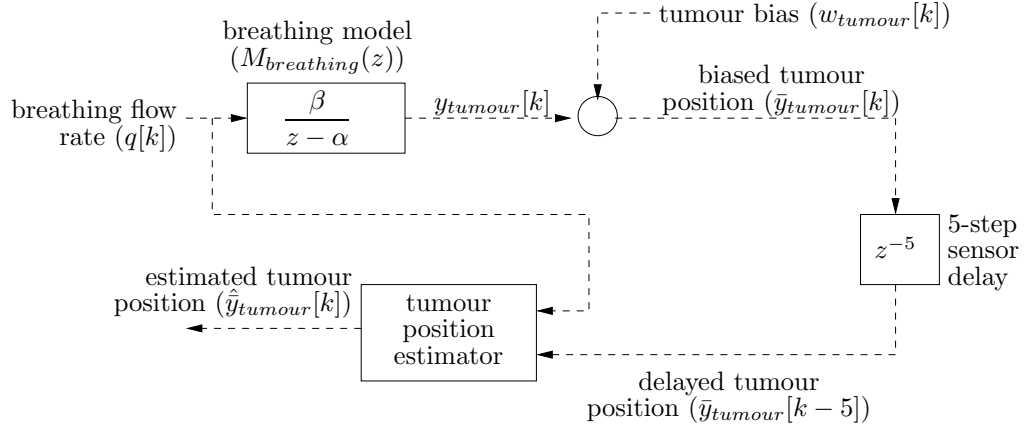


Figure 5.3: *Example estimation problem.*

First, note that $\bar{y}_{tumour}[k]$ can be written in the form of (5.2):

$$\begin{aligned}
 y_{tumour}[k] &= \alpha \cdot y_{tumour}[k-1] + \beta \cdot q[k-1] \\
 &= \alpha (\alpha \cdot y_{tumour}[k-2] + \beta \cdot q[k-2]) + \beta \cdot q[k-1] \\
 &= \alpha^2 y_{tumour}[k-2] + \alpha\beta \cdot q[k-2] + \beta \cdot q[k-1] \\
 &= \dots \\
 &= \alpha^5 y_{tumour}[k-5] + \sum_{i=1}^5 \beta \alpha^{i-1} q[k-i],
 \end{aligned}$$

giving

$$\begin{aligned}
\bar{y}_{tumour}[k] &= y_{tumour}[k] + w_{tumour}[k] \\
&= \alpha^5 y_{tumour}[k-5] + \sum_{i=1}^5 \beta \alpha^{i-1} q[k-i] + w_{tumour}[k] \\
&= \alpha^5 (\bar{y}_{tumour}[k-5] - w_{tumour}[k]) + \sum_{i=1}^5 \beta \alpha^{i-1} q[k-i] + w_{tumour}[k] \\
&= \alpha^5 \bar{y}_{tumour}[k-5] + \sum_{i=1}^5 \beta \alpha^{i-1} q[k-i] + (1 - \alpha^5) w_{tumour}[k] \\
&= \underbrace{\begin{bmatrix} \bar{y}_{tumour}[k-5] & q[k-1] & q[k-2] & \cdots & q[k-5] & 1 \end{bmatrix}}_{\phi[k]^T} \times \\
&\quad \underbrace{\begin{bmatrix} \alpha^5 \\ \beta \\ \alpha\beta \\ \vdots \\ \alpha^4\beta \\ (1 - \alpha^5)w_{tumour}[k] \end{bmatrix}}_{\theta^*} \\
&= \phi[k]^T \theta^*. \tag{5.11}
\end{aligned}$$

The projection algorithm for plants with sensor delays will be used to compute an estimate of θ^* and calculate an estimate of the plant output,

$$\hat{\bar{y}}_{tumour}[k] = \phi[k]^T \hat{\theta}[k-1].$$

As long as $w_{tumour}[k]$ is constant, Theorem 5.2 guarantees that this estimate will satisfy

$$\lim_{k \rightarrow \infty} (\bar{y}_{tumour}[k] - \hat{\bar{y}}_{tumour}[k]) = 0.$$

In the simulation, the sampling period is 0.1 s, the breathing flow rate, $q[k]$, is assumed to be sinusoidal with frequency 1/6 Hz and amplitude 250 cm³/s and the tumour bias is constant, $w_{tumour}[k] = 1.8$ mm. The breathing model is

$$M_{breathing}(z) = \frac{0.0149}{z - 0.99}$$

(i.e., $\alpha = 0.99$ and $\beta = 0.0149$) and the estimator uses the parameters $a = 1$ and $c = 1$, and an initial parameter estimate of $\theta_0 = 0$.

Figure 5.4 shows the performance of the estimator over the first 40 s. Note that the prediction error converges to zero impractically slowly in this example; even after 40 s, the magnitude of the prediction error is roughly 3 mm. Figure 5.5, which shows the performance of the estimator over 1000 s, confirms that the prediction error converges to 0. Also, note that the normed parameter estimate error converges to approximately 0.4 (i.e., the parameter estimate, $\hat{\theta}[k]$, does not converge to θ^*).

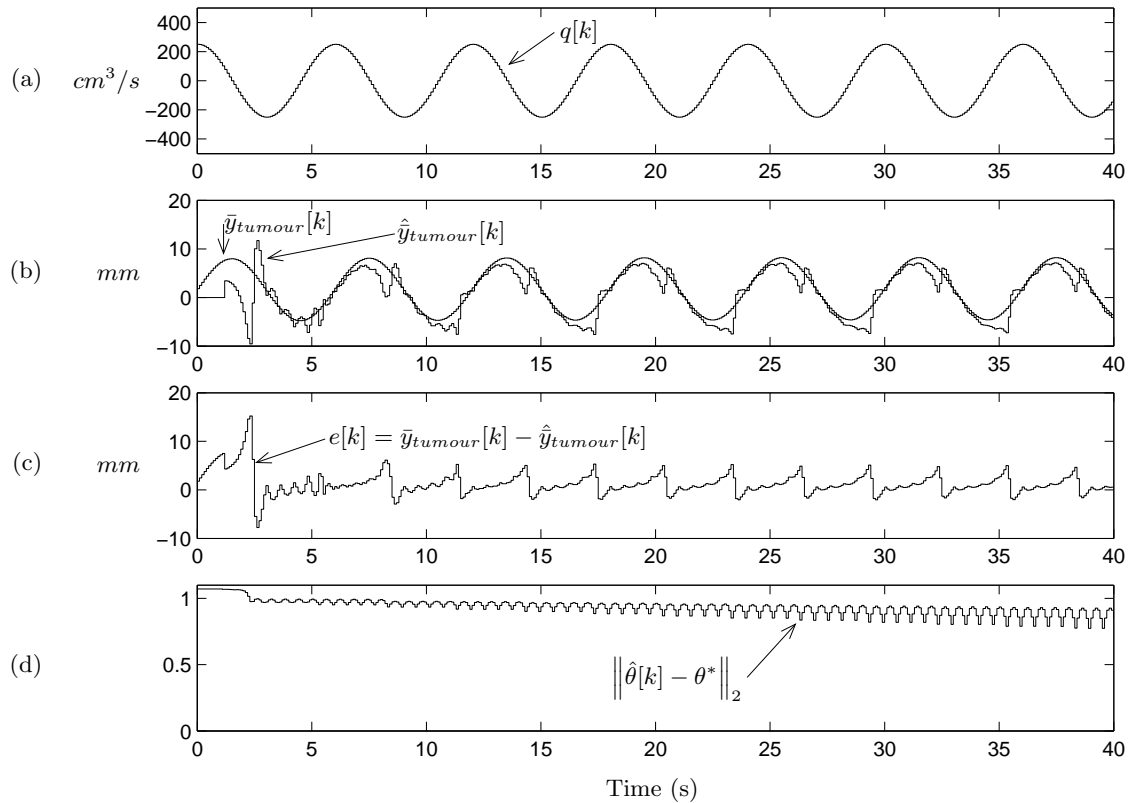


Figure 5.4: Estimator performance over the first 40 s: (a) breathing flow rate; (b) actual and estimated tumour positions; (c) prediction error; (d) normed parameter estimate error.




Figure 5.5: *Estimator performance over the first 1000 s: (a) prediction error; (b) normed parameter estimate error.*

5.3.1 Convergence Rate Sensitivity to Plant Input

Intuitively, the performance of the estimator is largely a function of the “richness” of the plant output which is itself a function of the plant input. Different plant input signals should yield different responses from the estimator; however, it is surprising just how profound an impact such differences have. For example, Figure 5.6 compares the prediction error of the estimator of Figure 5.3 for three different breathing flow rate signals. All of the signals are sinusoidal with amplitude $250 \text{ cm}^3/\text{s}$, they only differ in frequency. Even a relative difference of 1 % in the frequency results in a drastically different prediction error convergence. Note that the frequencies are not realistic for a breathing flow rate, they are only used for their illustrative value.




Figure 5.6: *Prediction error with different breathing flow rate signals*

5.3.2 Improving Convergence Rates Through Scaling

Although the prediction error converges to zero in the preceding simulations, the slow rate of convergence (for most breathing flow rate signals) would not be practical for the radiotherapy application at hand. Recognizing that the magnitude of the individual elements of θ^* vary greatly (the minimum and maximum values are approximately 0.014 and 0.95), it is natural to wonder whether scaling the elements of θ^* so they have comparable magnitudes will improve the rate of convergence. Of course, since the elements of θ^* are unknown, it is not generally possible to scale the elements so they are equal; however, scaling can be applied if some information on the relative magnitudes of the elements is known a priori. For example, assume that the parameters α , β and $w_{tumour}[k]$ are known to be bounded:

$$\begin{aligned} 0.9 &\leq \alpha \leq 1 \\ 0.01 &\leq \beta \leq 0.05 \\ 0 &\leq w_{tumour}[k] \leq 4. \end{aligned}$$

Using the median value for each parameter, the tumour position can be scaled as follows:

$$\bar{y}_{tumour}[k] = \underbrace{[0.95^5 \bar{y}_{tumour}[k-5] \ 0.03q[k-1] \cdots (0.03 \cdot 0.95^4)q[k-5] \ 2(1-0.95^5)]}_{\phi[k]^T} \times \underbrace{\begin{bmatrix} \alpha^5/0.95^5 \\ \beta/0.03 \\ \vdots \\ \alpha^4\beta/(0.03 \cdot 0.95^4) \\ (1 - \alpha^5)w_{tumour}[k]/(2(1 - 0.95^5)) \end{bmatrix}}_{\theta^*}. \quad (5.12)$$

Note that when $\alpha = 0.95$, $\beta = 0.03$ and $w_{tumour}[k] = 2$, θ^* of (5.12) will be a vector of ones. With $\alpha = 0.99$, $\beta = 0.0149$ and $w_{tumour}[k] = 1.8$ as in the previous simulation, (5.12) becomes

$$\bar{y}_{tumour}[k] \approx \underbrace{[0.95^5 \bar{y}_{tumour}[k-5] \ 0.03q[k-1] \cdots (0.03 \cdot 0.95^4)q[k-5] \ 2(1-0.95^5)]}_{\phi[k]^T} \times \underbrace{\begin{bmatrix} 1.229 \\ 0.497 \\ 0.518 \\ 0.539 \\ 0.562 \\ 0.586 \\ 0.195 \end{bmatrix}}_{\theta^*},$$

and the elements of θ^* are now roughly within one order of magnitude of each other. Figure 5.7 demonstrates the improved speed of convergence of the prediction error with the above scaling improvement; the output estimate error decays to less than 0.5 mm after only 5 s.

Although the scaling improvement outlined above has intuitive appeal, there is no guidance on how close together the elements of θ^* must be to achieve a practical rate of convergence of the prediction error to zero. Efforts to deterministically determine a rate of convergence have proven unsuccessful and other authors have noted the difficulty in determining a bound on the convergence rate of the standard (non-sensor delayed) projection algorithm [36], [37].




Figure 5.7: *Estimator performance with the scaling improvement: a) breathing flow rate; (b) actual and estimated tumour positions; (c) prediction error; (d) normed parameter estimate error.*

5.3.3 Tracking Time-Varying Parameters

Until this point in the simulations, it has been assumed that the tumour bias, $w_{tumour}[k]$, is constant. Recall from (5.11) that this assumption makes θ^* constant and therefore Theorem 5.2 applies. However, the projection algorithm for plants with sensor delays can be used to track some time-varying parameters. Figure 5.8 shows the estimator performance using the scaling improvement outlined above with a change in the tumour bias from 1.8 mm to 11.8 mm at $t = 7.0$ s. Although Theorem 5.2 no longer applies, the prediction error still converges to zero. Also note that property (i) of Theorem 5.2 does not hold in this simulation; however, this is not surprising since θ^* is now time-varying.




Figure 5.8: *Estimator performance with a tumour bias change: a) breathing flow rate; (b) actual and estimated tumour positions; (c) prediction error; (d) normed parameter estimate error.*

5.4 Applicability to Tumour Position Estimation

The results of this chapter, although motivated by the problem of estimating the tumour position, are not directly applicable to this problem. While the problem formulation of this chapter assumes single rate sampling, the tumour position estimation problem in Chapters 2 through 4 uses a simplified multirate sampling scheme (recall Section 2.2) and efforts to apply the work of this chapter to the tumour position estimation problem have remained unfruitful⁵. Despite this, two intuitive results shed some light on the difficulty of transitioning from the single rate to the simplified multirate framework. First, aliasing effects can “mask” the true tumour position. For example, Figure 5.9 shows sinusoidal tumour motion

⁵However, if the sampling rate of the breathing flow rate signal is slowed to match the sampling rate of the image processor, then the results of this chapter can be applied.




Figure 5.9: *Aliasing effects of the multirate tumour position estimation problem.*

at $1/2$ Hz with a sampling period of $h = 0.1$ s. If the image processing delay is 20 steps, i.e. $N = 20$ (recall Figure 5.1(a)), then the delayed, sampled tumour position, $\bar{y}_{tumour,fs}[k]$ is a constant 20.

The second, related, problem is the loss of information due to sampling. In Figure 5.9, since the tumour position is sampled every 20 steps, only $1/20^{\text{th}}$ of the tumour position “information” is accessible by the tumour position estimator. In general, for an N step delay, $1/N^{\text{th}}$ of the “information” is available. Does this loss present a limitation in the design of an estimator? Is it always possible to design an asymptotic estimator for the multirate tumour position estimation scheme for any value of N (assuming no aliasing effects)? To date, these remain open questions.

5.5 Relationship to Adaptive Prediction

Before leaving this chapter, it should be noted that after the work of this chapter was completed, it was discovered that the main motivation and result is a special case of adaptive prediction. Conceptually, the problem of this chapter was to *estimate* the signal $\bar{y}[k]$ (recall Figure 5.1(b)). Adaptive prediction, on the other hand, would focus on the mathematically identical problem of *predicting* where the signal

$\bar{y}[k - N]$ would be N steps from the current time. Mathematically, the projection algorithm commonly used in adaptive prediction minimizes the cost index

$$J(\theta) = \frac{1}{2} \left\| \theta - \hat{\theta}[k - 1] \right\|_2^2,$$

whereas the results in this chapter focused on the cost index

$$J(\theta) = \frac{1}{2} \left\| \theta - \hat{\theta}[k - N - 1] \right\|_2^2$$

(recall (5.4)). Otherwise, the results of this chapter and the standard projection algorithm are derived similarly. A good overview of adaptive prediction results can be found in [38].

Chapter 6

The Nonovershooting Step Response Problem

The design of the collimator leaf controller, $C(z)$, in Section 3.2 focused on meeting stability, tracking and disturbance rejection requirements; however, in retrospect, one critical requirement was overlooked. Simulations revealed that it is possible for opposing collimator leaves to collide with each other. If this were to occur on a real collimator, serious damage to the collimator (or worse) could result. One possible method to prevent opposing leaf collisions is to design $C(z)$ so that each collimator leaf does not overshoot its reference position; however, it is not immediately clear how to design $C(z)$ to meet this requirement, or if it is even possible. In addition, $C(z)$ should have a “fast” response to reduce errors caused by leaf transient motion. Focusing, in particular, on the step response, this chapter will show that it is impossible to design a stabilizing controller to provide an arbitrarily fast nonovershooting response if the plant has two real nonminimum phase (NMP) zeros. This result has been published in a different form in [39].

The work of this chapter is motivated not only by the control of the collimator leaves; many control problems require a fast response with no (or minimal) overshoot. Robot force control, for example, typically requires minimizing overshoot to prevent damage to the environment (e.g., see [40]). Also, [41] describes the design of a missile autopilot controller with a fast response and minimal overshoot.

The next section provides some background and related work on the design of compensators to provide a nonovershooting response and definitions used in this chapter. For simplicity, this chapter focuses on continuous-time systems; analogous single rate discrete-time results are summarized in Section 6.5.

6.1 Background and Definitions

The focus of this chapter can be summarized in the following questions: given a linear time-invariant (LTI) plant, is it always possible to design a stabilizing LTI controller to provide a nonovershooting step response? If so, can a stabilizing LTI controller provide an *arbitrarily fast* nonovershooting step response? The answer to the former question is yes. In [42], the authors describe a stabilizing two degree-of-freedom compensator to render the step response nonovershooting for stabilizable plants without a zero at the origin. However, the authors note that their technique “is equivalent to sufficiently slowing down the speed of the response”. In [43] and [44], nonovershooting controllers are considered where there are no constraints on the speed of the response, but the results only apply if the plant is minimum phase. In particular, [44] presents a controller design procedure for minimum-phase plants to achieve independently specified overshoot and rise time. With these results in mind, the only open question is whether a stabilizing controller can be designed to provide an arbitrarily fast nonovershooting step response for *nonminimum phase* (NMP) plants, i.e., plants with zeros in the open right half of the complex plane.

It is well established that NMP zeros limit the achievable closed-loop performance; representative works that discuss such performance limitations include [45]-[50] and good overviews of these and related results include [51] and [52]. Perhaps the most familiar characteristic of real NMP zeros is that they lead to *undershoot* in the step response. The purpose of this chapter is to show that real NMP zeros can also lead to *overshoot* in the step response.

Throughout this chapter, let $P(s)$ denote the transfer function of the known plant, let $r(t)$ denote the reference signal and let $y(t)$ denote the plant output. The control topology is assumed to be the general two degrees-of-freedom (2-DOF) control configuration shown in Figure 6.1(a). Note that classical 1-DOF control (Figure 6.1(b)), open-loop control (Figure 6.1(c)) and standard control schemes based on state-space methods (such as the observer-based state-feedback control scheme with integral action of Figure 6.1(d)) are special cases of Figure 6.1(a). Denote the closed-loop transfer function from $R(s)$ to $Y(s)$ in Figure 6.1(a) by $H(s)$. Assuming closed-loop stability, the zeros of $H(s)$ include all the NMP zeros of $P(s)$ and assuming that $P(s)$ is strictly proper, then $H(s)$ is also strictly proper. Hence, throughout this chapter, $H(s)$ is assumed to be a strictly proper NMP transfer function with input $R(s)$ and output $Y(s)$. Since this chapter investigates the unit step response of $H(s)$, $R(s) = 1/s$ throughout. The following characteristics of the unit step response, $y(t)$, will be used (see Figure 6.2):

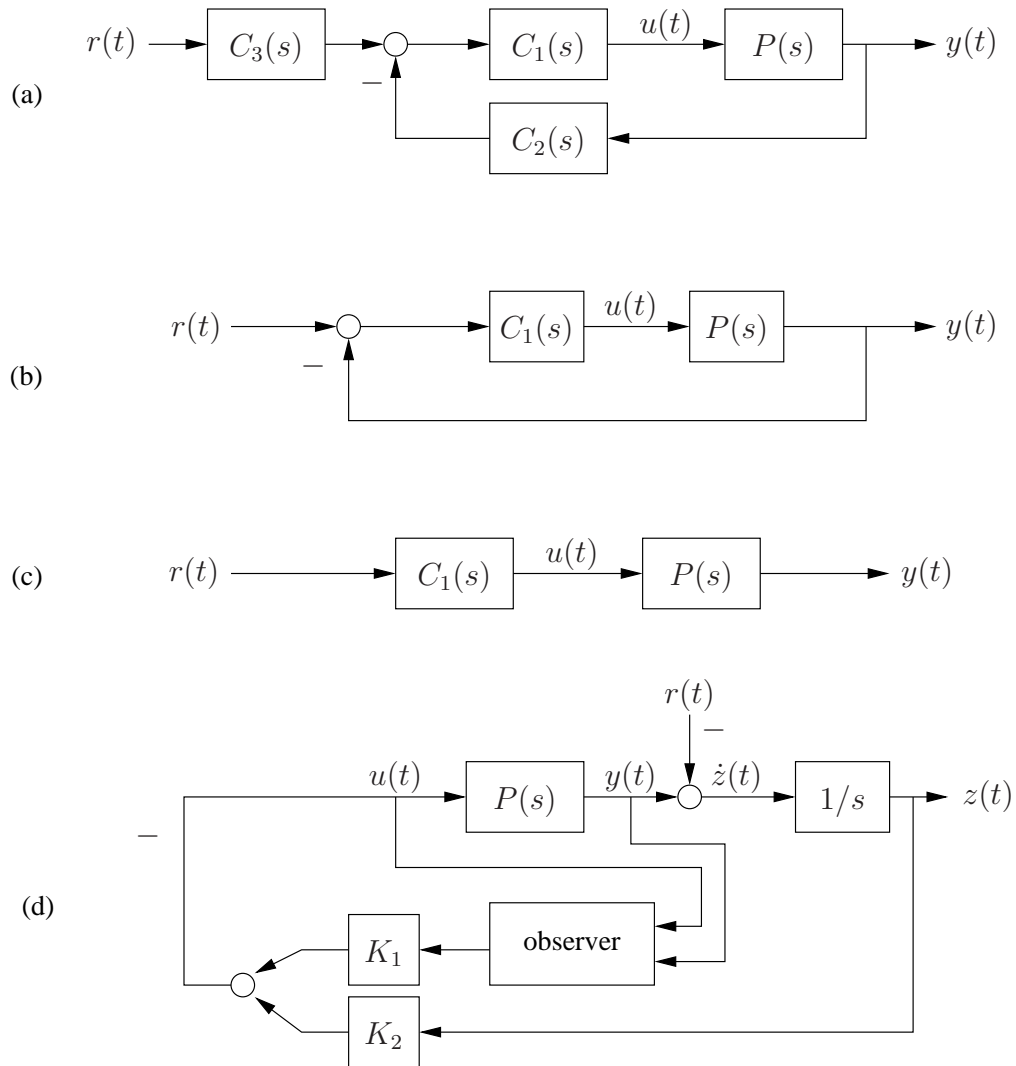


Figure 6.1: Control configurations considered in this chapter: (a) a general two degree-of-freedom feedback scheme; (b) the classical unity feedback one degree-of-freedom feedback scheme; (c) an open-loop scheme; (d) an observer-based state-feedback control scheme with integral action.




Figure 6.2: A representative step response of $H(s)$ used to show the definitions of steady-state value (y_∞), undershoot (y_{us}), overshoot (y_{os}), and settling time (t_s).

- The *steady-state value* of $y(t)$, denoted y_∞ , is

$$y_\infty = \lim_{t \rightarrow \infty} y(t).$$

The Final Value Theorem immediately gives $y_\infty = H(0)$. For simplicity, it is assumed $y_\infty > 0$, although similar results can be derived if $y_\infty < 0$.

- The *undershoot* of the step response, denoted y_{us} , is the smallest nonnegative number such that

$$y(t) \geq -y_{us} \quad \forall t \geq 0.$$

$y(t)$ is said to “exhibit undershoot” if $y_{us} > 0$ ¹.

- The *overshoot* of $y(t)$, denoted y_{os} , is the smallest nonnegative number such that

$$y(t) \leq (y_\infty + y_{os}) \quad \forall t \geq 0.$$

$y(t)$ is said to “exhibit overshoot” if $y_{os} > 0$.

- The *settling time* of $y(t)$, denoted t_s , is the smallest time such that

$$|y(t) - y_\infty| \leq 0.02y_\infty \quad \forall t \geq t_s.$$

¹ Note that alternate definitions of undershoot have been used. For example, in [46] the term undershoot is used in the more limited situation where there exists a $T > 0$ such that $y(t) < 0$ for $0 < t < T$.

The following theorem, which shows that real NMP zeros lead to undershoot, is essential for the main result of this chapter. The proof is included for completeness.

Theorem 6.1 [47]: Let $H(s)$ be a stable strictly proper transfer function with a real NMP zero at $s = z$ and with any number of additional zeros located anywhere in the complex plane. Assume $y_\infty > 0$. Then the step response of $H(s)$ exhibits undershoot that, as a fraction of the steady-state value, is bounded by

$$\frac{y_{us}}{y_\infty} \geq \frac{0.98}{e^{zt_s} - 1} > 0. \quad (6.1)$$

As $t_s \rightarrow 0$, the bound (6.1) tends to ∞ .

Proof: Since z is a zero of $H(s)$, it is true that $Y(z) = H(z)R(z) = H(z)/z = 0$. Now, the Laplace transform of $y(t)$ is

$$Y(s) = \int_0^\infty y(t)e^{-st} dt,$$

which (due to the stability of $H(s)$) converges for all s with positive real part, including the point $s = z$. Hence,

$$Y(z) = \int_0^\infty y(t)e^{-zt} dt = 0. \quad (6.2)$$

Since $y_\infty > 0$, (6.2) implies that time intervals where $y(t)$ is positive must be balanced by time intervals where $y(t)$ is negative. Thus, the step response of $H(s)$ exhibits undershoot. To derive the undershoot bound, first use (6.2) and the fact that $-y(t) \leq y_{us}$ to obtain

$$\begin{aligned} \int_{t_s}^\infty y(t)e^{-zt} dt &= \int_0^{t_s} -y(t)e^{-zt} dt \\ &\leq \int_0^{t_s} y_{us}e^{-zt} dt \\ &= \frac{y_{us}}{z} (1 - e^{-zt_s}). \end{aligned} \quad (6.3)$$

Second, use the fact that $y(t) \geq 0.98y_\infty$ for $t \in [t_s, \infty)$ to compute that

$$\begin{aligned} \int_{t_s}^\infty y(t)e^{-zt} dt &\geq \int_{t_s}^\infty 0.98y_\infty e^{-zt} dt \\ &= \frac{0.98y_\infty}{z} e^{-zt_s}. \end{aligned} \quad (6.4)$$

Combine (6.3) and (6.4) to obtain (6.1). □

The following lemma is also crucial in the investigation of step response overshoot in the next section. It shows that augmenting a stable strictly proper transfer function with a real NMP zero can never decrease the settling time of the step response.

Lemma 6.1: Let $G(s)$ be a stable, strictly proper transfer function and

$$H(s) = G(s) \left(1 - \frac{1}{z}s \right),$$

where $z > 0$ is a real number. Denote the step responses of $H(s)$ and $G(s)$ by $y(t)$ and $f(t)$, respectively. Then the settling time of $y(t)$ is not less than the settling time of $f(t)$.

Proof: By the lemma assumptions,

$$y(t) = f(t) - \frac{1}{z} \frac{d}{dt} f(t) \tag{6.5}$$

holds for all $t > 0$. Since $G(0) = H(0)$, the signal $f(t)$ tends to a positive number as $t \rightarrow \infty$, namely y_∞ . Denote the settling time of $f(t)$ by t_{sf} . There are two cases to consider, as shown in Figure 6.3. In the first case, $f(t_{sf}) = 1.02y_\infty$, and therefore $\frac{d}{dt}f(t_{sf}) \leq 0$. Equation (6.5) yields

$$\begin{aligned} y(t_{sf}) &= f(t_{sf}) - \frac{1}{z} \frac{d}{dt} f(t_{sf}) \\ &= 1.02y_\infty - \frac{1}{z} \frac{d}{dt} f(t_{sf}) \\ &\geq 1.02y_\infty. \end{aligned} \tag{6.6}$$

In the second case, $f(t_{sf}) = 0.98y_\infty$, and therefore $\frac{d}{dt}f(t_{sf}) \geq 0$ and

$$\begin{aligned} y(t_{sf}) &= f(t_{sf}) - \frac{1}{z} \frac{d}{dt} f(t_{sf}) \\ &= 0.98y_\infty - \frac{1}{z} \frac{d}{dt} f(t_{sf}) \\ &\leq 0.98y_\infty. \end{aligned} \tag{6.7}$$




Figure 6.3: A representative step response of $f(t)$, including a zoom-in near $t = t_{sf}$, showing the two possible cases discussed in the proof.

Inequalities (6.6) and (6.7) imply that $|y(t_{sf}) - y_{\infty}| \geq 0.02y_{\infty}$; in other words, signal $y(t)$ must have a settling time at least as long as the settling time of the signal $f(t)$. \square

The next two sections investigate step response overshoot when $H(s)$ has two real NMP zeros (the following section) and one real NMP zero (Section 6.3).

6.2 Case where $H(s)$ has Two NMP Zeros

Theorem 6.1 establishes that the presence of at least one real NMP zero in $H(s)$ is enough to guarantee undershoot, and that the undershoot grows without bound as the settling time tends to zero. The following theorem, the main result of this chapter, shows that a similar result holds for step response overshoot, assuming that $H(s)$ has at least *two* real NMP zeros. While two real NMP zeros do not

necessarily guarantee overshoot, overshoot will occur if the settling time is small enough and the overshoot grows without bound as the settling time tends to zero.

Theorem 6.2: Let $H(s)$ be a stable strictly proper transfer function with real NMP zeros at $s = z_1$ and $s = z_2$ and with any number of additional zeros located anywhere in the complex plane. Assume $y_\infty > 0$. Then

$$\frac{y_{os}}{y_\infty} > \left(\frac{0.98}{e^{z_2 t_s} - 1} \right) \left(\frac{1 - z_1 t_s}{z_1 t_s} \right) - 1. \quad (6.8)$$

For sufficiently small t_s , the bound in (6.8) is positive, so the step response necessarily exhibits overshoot. As $t_s \rightarrow 0$, the bound (6.8) tends to ∞ .

Proof: The intuition behind the proof is that one zero leads to undershoot while the other zero introduces negative derivation action (see (6.10) below); the effect of the negative derivative action is to increase $y(t)$ and the increase gets more severe as the settling time decreases.

Begin by factoring $H(s)$ as

$$H(s) = G(s) \left(1 - \frac{1}{z_1} s \right), \quad (6.9)$$

where $G(s)$ is stable and strictly proper with a NMP zero at $s = z_2$. Let $f(t)$ denote the step response of $G(s)$ and, as usual, let $y(t)$ denote the step response of $H(s)$ (see Figure 6.4). By (6.9),

$$y(t) = f(t) - \frac{1}{z_1} \frac{d}{dt} f(t) \quad (6.10)$$

holds for all $t > 0$. The signal $f(t)$ tends to a positive number as $t \rightarrow \infty$, namely y_∞ , since $H(0) = G(0)$. Figure 6.5 shows a representative plot of $f(t)$. Since $G(s)$ is NMP, Theorem 6.1 implies that $f(t)$ exhibits undershoot; more precisely, denoting the minimum value of $f(t)$ by $-f_{us}$, Theorem 6.1 states

$$f_{us} \geq \frac{0.98 y_\infty}{e^{z_2 t_{sf}} - 1}, \quad (6.11)$$

where t_{sf} is the settling time of $f(t)$ (again, see Figure 6.5). From Lemma 6.1, signal $y(t)$ must have a longer settling time than signal $f(t)$:

$$t_s \geq t_{sf}. \quad (6.12)$$




Figure 6.4: Factoring $H(s)$ as in (6.9).

From (6.12) and (6.11), the following bound on f_{us} results:

$$f_{us} \geq \frac{0.98y_\infty}{e^{z_2 t_s} - 1}. \quad (6.13)$$

Next, choose t_2 to be any value in the (necessarily non-empty) set $\{t : f(t) = -f_{us}\}$ and define $t_1 := \max\{t : f(t) = 0, 0 \leq t < t_2\}$, as indicated in Figure 6.5. By the Mean Value Theorem, there exists a t^* in the interval (t_1, t_2) such that

$$\frac{d}{dt}f(t^*) = \frac{f(t_2) - f(t_1)}{t_2 - t_1} = \frac{-f_{us}}{t_2 - t_1}. \quad (6.14)$$




Figure 6.5: Representative plot of $f(t)$.

The machinery now exists to determine a bound on the peak value of $y(t)$. In particular, use (6.10), (6.14), the fact that $f(t^*) \geq -f_{us}$, the fact that $t_s > t_2 - t_1$, and (6.13) to obtain

$$\begin{aligned}
\max_{t \in [0, \infty)} y(t) &\geq y(t^*) \\
&= f(t^*) - \frac{1}{z_1} \frac{d}{dt} f(t^*) \\
&= f(t^*) + \frac{f_{us}}{z_1} \frac{1}{t_2 - t_1} \\
&\geq -f_{us} + \frac{f_{us}}{z_1} \frac{1}{t_2 - t_1} \\
&> -f_{us} + \frac{f_{us}}{z_1} \frac{1}{t_s} \\
&= f_{us} \left(\frac{1}{z_1 t_s} - 1 \right) \\
&\geq \frac{0.98 y_\infty}{e^{z_2 t_s} - 1} \left(\frac{1}{z_1 t_s} - 1 \right). \tag{6.15}
\end{aligned}$$

Finally, use (6.15) and the relationship $\max_{t \in [0, \infty)} y(t) = y_{os} + y_\infty$ to obtain (6.8). \square

Note that the bound in Theorem 6.2 is equally valid with z_1 and z_2 exchanged, but neither of the two arrangements of z_1 and z_2 yields a consistently tighter bound. Also note that, although the bound may be conservative in particular examples, it does prove that the step response of a transfer function with at least two real nonminimum phase zeros exhibits overshoot if the settling time is short enough.

6.3 Case where $H(s)$ has One NMP Zero

While Theorem 6.2 establishes that step response overshoot results from the presence of two real NMP zeros and sufficiently small settling time, an analogous result does not hold if $H(s)$ has only one NMP zero:

Theorem 6.3: Let $H(s)$ be a stable strictly proper transfer function with exactly one NMP zero, at $s = z$. Then there does not exist a non-trivial lower bound (possibly dependent on t_s , y_∞ , and z) on the step response overshoot of $H(s)$.

Proof: Consider the special case where all the poles of $H(s)$ are real and distinct:

$$H(s) = \frac{-K(s - z)}{\prod_{i=1}^{n_p} (s + p_i)}. \quad (6.16)$$

In (6.16), $K > 0$, $z > 0$, and $p_{n_p} > p_{n_p-1} > \dots > p_1 > 0$ with $n_p \geq 2$. As usual, let $y(t)$ denote the step response of $H(s)$. Then, $y(t)$ exhibits the following characteristics:

1. $y(0) = 0$ (since $H(s)$ is strictly proper);
2. $y_\infty > 0$ (by direct calculation of $H(0)$);
3. $y(t)$ exhibits undershoot (using Theorem 6.1); and
4. $y(t)$ has exactly one extrema for $t > 0$ (using the results in [53]).

These four properties imply that $y(t)$ does not exhibit overshoot, no matter how small t_s is. Consequently, there does not exist a non-trivial lower bound on the overshoot. \square

Theorem 6.3 does not imply that there is never closed-loop step response overshoot for a plant that has a single NMP zero; indeed, it is not difficult to generate counter-examples where overshoot is exhibited because of the positions of other poles and zeros. However, the theorem does establish that a single NMP zero does not *necessarily* lead to step response overshoot.

6.4 Simulation Results

Consider the plant

$$P(s) = \frac{(s-1)(s-2)}{(s+1)^3}.$$

Assume that any of the control configurations shown in Figure 6.1 is used, and, for simplicity, assume that perfect steady-state step tracking is required, i.e., $y_\infty = 1$. Using $z_1 = 1$ and $z_2 = 2$, the bound on overshoot in (6.8) was computed as a function of the settling time, and is plotted in Figure 6.6 as a solid curve. Note that, for $t_s > 0.4237$ s, the bound is negative, and therefore provides no meaningful information. However, for $t_s < 0.4237$ s, the bound is positive and therefore there *must* be overshoot for every possible stabilizing linear time-invariant control scheme. As expected, the overshoot tends to infinity as t_s tends to zero.




Figure 6.6: Lower bound on the step response overshoot. The dash line marks where the bound becomes positive.

To demonstrate more explicitly the overshoot trend for this example, consider an observer-based state-feedback controller with integral action, as shown in Figure 6.1(d). Let $\dot{x} = Ax + Bu$, $y = Cx$ be any minimal state-space realization of $P(s)$. To describe the controller in Figure 6.1(d), introduce

$$\tilde{A} := \begin{bmatrix} A & 0 \\ C & 0 \end{bmatrix} \quad \text{and} \quad \tilde{B} := \begin{bmatrix} B \\ 0 \end{bmatrix}.$$

The controller equations are then

$$\begin{aligned} \dot{z} &= y - r \\ \dot{\hat{x}} &= (A - BK_1 - HC)\hat{x} - BK_2z + Hy \\ u &= -K_1\hat{x} - K_2z, \end{aligned}$$

where the two controller gains, $K := [K_1 \ K_2]$ and H , are chosen such that the eigenvalues of $(\tilde{A} - \tilde{B}K)$ and $(A - HC)$ are all in the open left-half complex plane. (The closed-loop poles are, of course, the union of the eigenvalues of $(\tilde{A} - \tilde{B}K)$ and those of $(A - HC)$.) The gain K was computed so that the eigenvalues of $(\tilde{A} - \tilde{B}K)$ are placed at $\{-\alpha, -2\alpha, -3\alpha, -4\alpha\}$ (where $\alpha > 0$ is a parameter), and H was computed so that the eigenvalues of $(A - HC)$ are placed at $\{-5\alpha, -6\alpha, -7\alpha\}$. Figure 6.7 shows the final closed-loop step responses for three values of α . In all cases, zero initial conditions are assumed. Observe that, due to the integrator in the feedback loop, $y_\infty = 1$ holds. Also, note the presence of overshoot in the plots; the overshoot is much larger than the bound (6.8), indicating that Theorem 6.2 can be conservative. Finally, it is interesting that the overshoot is significantly larger in magnitude than the undershoot, implying that overshoot is, at least for some examples, a greater cost associated with multiple real NMP zeros than is undershoot.




Figure 6.7: Closed-loop step responses using three different values of α .

Now, consider the same controller constructed for the plant

$$P(s) = \frac{(s - 1)}{(s + 1)^3}.$$

Note that $P(s)$ is the same as before, except one of the NMP zeros has been removed. The closed-loop transfer function from $R(s)$ to $Y(s)$ is in the form of (6.16) in the proof of Theorem 6.3 and the step response will not exhibit overshoot, independent of the value of α . Figure 6.8 shows the final closed-loop step responses for three values of α .




Figure 6.8: *Nonovershooting closed-loop step responses using three different values of α .*

6.5 Extensions to Discrete Time

Although this chapter has focused on a continuous-time performance limitation associated with NMP zeros, the radiotherapy control setups of Chapters 2 through 4 operate in discrete time. This section returns to the discrete-time framework and investigates whether a discrete-time stabilizing LTI controller can be designed to provide an arbitrarily fast nonovershooting step response for any stabilizable LTI discrete-time plant. It turns out that the continuous-time results presented earlier in this chapter have analogous discrete-time counterparts; however, while Section 6.2 showed that there is always overshoot for continuous-time plants with two real NMP zeros if the speed of response is made fast enough, the corresponding discrete-time result (see Theorem 6.5.2 below) is not as strong. In particular, since the speed of response is lower bounded by the system delay (relative degree) of the discrete-time plant, the speed of response cannot be made arbitrarily fast and overshoot cannot be guaranteed. On the other hand, continuous-time plants do

not, in general, have such a limit and Theorem 6.2 guarantees overshoot as the settling time approaches zero.

6.5.1 Background and Definitions

Section 6.1 noted that, given a stabilizable continuous-time LTI plant without a zero at the origin, it is always possible to design a continuous-time LTI controller to render the step response nonovershooting. A similar result holds for discrete-time plants: there always exists a discrete-time LTI controller that results in a nonovershooting step response for stabilizable discrete-time LTI plants with no zeros on the unit circle [54]. Although proving the *existence* of a nonovershooting controller, the authors of [54] make no claims on the speed of response. This section will shed some insight into the design of a discrete-time controller to provide an arbitrarily fast nonovershooting step response by developing discrete-time counterparts to the theorems and lemmas of Sections 6.1 through 6.3.

Throughout this section, let $P^d(z)$ denote the transfer function of the known plant, let $r^d[k]$ denote the reference signal and let $y^d[k]$ denote the plant output. (A superscript “d” is used to distinguish the discrete-time definitions from the continuous-time definitions used earlier in the chapter.) Denote the z-transform of the reference signal and plant output by $R^d(z)$ and $Y^d(z)$, respectively. The control topology is assumed to be the general two degrees-of-freedom (2-DOF) control configuration shown in Figure 6.9. Denote the closed-loop transfer function from $R^d(z)$ to $Y^d(z)$ in Figure 6.9 by $H^d(z)$. Assuming closed-loop stability, the zeros of $H^d(z)$ include all the NMP zeros of $P^d(z)$ and assuming $P^d(z)$ is strictly proper, $H^d(z)$ is also strictly proper. Therefore, $H^d(z)$ is assumed to be a stable strictly proper NMP transfer function with input $R^d(z)$ and output $Y^d(z)$. Since this section explores the unit step response of $H^d(z)$, $R^d(z) = z/(z - 1)$. The following characteristics of the unit step response, $y^d[k]$, will be used:




Figure 6.9: General two degree-of-freedom discrete-time control configuration considered in this section.

- The *steady-state value* of $y^d[k]$, denoted y_∞^d , is $y_\infty^d = \lim_{k \rightarrow \infty} y^d[k]$. The Final Value Theorem immediately gives $y_\infty^d = H^d(1)$. It is assumed $y_\infty^d > 0$ throughout this section.
- The *undershoot* of the step response, denoted y_{us}^d , is the smallest nonnegative number such that

$$y^d[k] \geq -y_{us}^d \quad \forall k \geq 0.$$

$y^d[k]$ is said to “exhibit undershoot” if $y_{us}^d > 0$.

- The *overshoot* of $y^d[k]$, denoted y_{os}^d , is the smallest nonnegative number such that

$$y^d[k] \leq (y_\infty^d + y_{os}^d) \quad \forall k \geq 0.$$

$y^d[k]$ is said to “exhibit overshoot” if $y_{os}^d > 0$.

- The *settling time* of $y^d[k]$, denoted t_s^d , is the smallest time such that

$$|y^d[k] - y_\infty^d| \leq 0.02y_\infty^d \quad \forall k \geq t_s^d.$$

The following theorem, the discrete-time counterpart of Theorem 6.1, is a key component of the main development in this section and is important in its own right since it shows that real positive NMP zeros lead to step response undershoot.

Theorem 6.5.1 [55]: Let $H^d(z)$ be a stable strictly proper transfer function with a real positive NMP zero at $z = a > 1$ and with any number of additional zeros located anywhere in the complex plane. Assume $y_\infty^d > 0$. Then the step response of $H^d(z)$ exhibits undershoot that, as a fraction of the steady-state value, is bounded by

$$\frac{y_{us}^d}{y_\infty^d} \geq \frac{0.98}{a^{t_s^d} - 1} > 0. \quad (6.17)$$

Proof: Since a is a zero of $H^d(z)$, it is true that $Y^d(a) = H^d(a)R^d(a) = H^d(a) \times \frac{a}{a-1} = 0$. Now, due to the stability of $H^d(z)$, the z-transform of $y^d[k]$ converges for all

$$\{z \in \Re : z > 1\},$$

including the point $z = a$. Hence,

$$Y^d(a) = \sum_{k=0}^{\infty} y^d[k]a^{-k} = 0. \quad (6.18)$$

Since $y_{\infty}^d > 0$, (6.18) implies that time intervals where $y^d[k]$ is positive must be balanced by time intervals where $y^d[k]$ is negative. Therefore, the step response of $H^d(z)$ exhibits undershoot. To derive (6.17), use (6.18) and the fact that $-y^d[k] \leq y_{us}^d$ to obtain

$$\begin{aligned} \sum_{k=t_s^d}^{\infty} y^d[k]a^{-k} &= \sum_{k=0}^{t_s^d-1} -y^d[k]a^{-k} \\ &\leq y_{us}^d \sum_{k=0}^{t_s^d-1} a^{-k} \\ &= y_{us}^d \left(\frac{1 - a^{-t_s^d}}{1 - 1/a} \right). \end{aligned} \quad (6.19)$$

Now, use the fact that $y^d[k] \geq 0.98y_{\infty}^d$ for $k \in [t_s^d, \infty)$ to compute that

$$\begin{aligned} \sum_{k=t_s^d}^{\infty} y^d[k]a^{-k} &\geq 0.98y_{\infty}^d \sum_{k=t_s^d}^{\infty} a^{-k} \\ &= 0.98y_{\infty}^d \left(\frac{a^{-t_s^d}}{1 - 1/a} \right). \end{aligned} \quad (6.20)$$

Combine (6.19) and (6.20) to obtain (6.17). \square

The following lemma, the discrete-time equivalent of Lemma 6.1, is also required in the main development of this section. It shows that augmenting a stable strictly proper transfer function with a positive real NMP zero can never decrease the settling time of the step response.

Lemma 6.5.1: Let $G^d(z)$ be a stable strictly proper transfer function and

$$H^d(z) = G^d(z) \left(\frac{a}{a-1} - \frac{1}{a-1}z \right),$$

where $a > 1$ is a real number. (Note that $H^d(z)$ has been factored in this form so the DC gain of $H^d(z)$ and $G^d(z)$ are equal.) Denote the step responses of $H^d(z)$ and $G^d(z)$ by $y^d[k]$ and $f^d[k]$, respectively. Then the settling time of $y^d[k]$ is at least as long as the settling time of $f^d[k]$.

Proof: By the lemma assumptions,

$$y^d[k-1] = \frac{a}{a-1}f^d[k-1] - \frac{1}{a-1}f^d[k] \quad (6.21)$$

holds for all $k \geq 1$. Since $G^d(1) = H^d(1)$, the signal $f^d[k]$ tends to a positive number as $k \rightarrow \infty$, namely y_∞^d . Denote the settling time of $f^d[k]$ by t_{sf}^d . There are two cases to consider. In the first case, $f[t_{sf}^d - 1] > 1.02y_\infty^d$ and $|f^d[t_{sf}^d] - y_\infty^d| \leq 0.02y_\infty^d$ (i.e., the signal $f^d[k]$ enters the “settling band” from above), and therefore (6.21) yields

$$\begin{aligned} y^d[t_{sf}^d - 1] &= \frac{a}{a-1}f^d[t_{sf}^d - 1] - \frac{1}{a-1}f^d[t_{sf}^d] \\ &> \frac{a}{a-1}(1.02y_\infty^d) - \frac{1}{a-1}(1.02y_\infty^d) \\ &= 1.02y_\infty^d \end{aligned} \quad (6.22)$$

In the second case, $f^d[t_{sf}^d - 1] < 0.98y_\infty^d$ and $|f^d[t_{sf}^d] - y_\infty^d| \leq 0.02y_\infty^d$ (i.e., the signal $f^d[k]$ enters the “settling band” from below), and therefore (6.21) yields

$$\begin{aligned} y^d[t_{sf}^d - 1] &= \frac{a}{a-1}f^d[t_{sf}^d - 1] - \frac{1}{a-1}f^d[t_{sf}^d] \\ &< \frac{a}{a-1}(0.98y_\infty^d) - \frac{1}{a-1}(0.98y_\infty^d) \\ &= 0.98y_\infty^d \end{aligned} \quad (6.23)$$

Inequalities (6.22) and (6.23) imply that $|y^d[t_{sf}^d - 1] - y_\infty^d| > 0.02y_\infty^d$; in other words, signal $y^d[k]$ must have a settling time no less than the settling time of signal $f^d[k]$. \square

The next two subsections investigate step response overshoot in the discrete-time framework.

6.5.2 Case where $H^d(z)$ has Two NMP Zeros

Theorem 6.2 confirms that the presence of two real NMP zeros in a (stable) continuous-time transfer function results in step response overshoot if the settling time is made short enough. The following theorem shows that a similar result holds for discrete-time transfer functions.

Theorem 6.5.2: Let $H^d(z)$ be a stable, strictly proper transfer function with real, positive NMP zeros at $z = a_1 > 1$ and $z = a_2 > 1$ and with any number of additional zeros located anywhere in the complex plane. Assume $y_\infty^d > 0$. Then

$$\frac{y_{os}^d}{y_\infty^d} > \left(\frac{0.98}{a_2^{t_s^d} - 1} \right) \left(\frac{1}{a_1 - 1} \right) - 1. \quad (6.24)$$

Proof: The intuition behind the proof is similar to that of Theorem 6.2: one zero leads to undershoot, while the other zero introduces a negative scaling action.

Begin by factoring $H^d(z)$ as

$$H^d(z) = G^d(z) \left(\frac{a_1}{a_1 - 1} - \frac{1}{a_1 - 1} z \right), \quad (6.25)$$

where $G^d(z)$ is stable and strictly proper with a NMP zero at $z = a_2$. (As before, $H^d(z)$ has been factored in this form to make the DC gain of $H^d(z)$ and $G^d(z)$ equal.) Let $f^d[k]$ and $y^d[k]$ denote the step responses of $G^d(z)$ and $H^d(z)$ respectively (see Figure 6.10). By (6.25),

$$y^d[k] = \frac{a_1}{a_1 - 1} f^d[k] - \frac{1}{a_1 - 1} f^d[k + 1] \quad (6.26)$$

holds for all $k \geq 0$. The signal $f^d[k]$ tends to a positive number as $k \rightarrow \infty$, namely y_∞^d , since $H^d(1) = G^d(1)$. Since $G^d(z)$ has a positive real NMP zero, Theorem 6.5.1 implies that $f^d[k]$ exhibits undershoot. Denoting the minimum value of $f^d[k]$ by $-f_{us}^d$, Theorem 6.5.1 states

$$f_{us}^d \geq \frac{0.98 y_\infty^d}{a_2^{t_{sf}^d} - 1}, \quad (6.27)$$




Figure 6.10: Factoring $H^d(z)$ as in (6.25).

where t_{sf}^d is the settling time of $f^d[k]$. From Lemma 6.5.1, signal $y^d[k]$ must have a longer settling time than signal $f^d[k]$:

$$t_s^d \geq t_{sf}^d. \quad (6.28)$$

From (6.28) and (6.27), the following bound on f_{us}^d results:

$$f_{us}^d \geq \frac{0.98y_\infty^d}{a_2^{t_s^d} - 1}. \quad (6.29)$$

Next, choose k^* to be any value in the necessarily non-empty set $\{k : f^d[k^*] \geq 0, f^d[k^* + 1] < 0\}$. Now, use (6.26), the fact that $f^d[k^* + 1] \geq -f_{us}^d$ and (6.29) to obtain

$$\begin{aligned} \max_{k \geq 0} y^d[k] &\geq y^d[k^*] \\ &= \frac{a_1}{a_1-1} f^d[k^*] - \frac{1}{a_1-1} f^d[k^* + 1] \\ &\geq -\frac{1}{a_1-1} f^d[k^* + 1] \\ &\geq \frac{1}{a_1-1} f_{us}^d \\ &\geq \left(\frac{1}{a_1-1} \right) \left(\frac{0.98y_\infty^d}{a_2^{t_s^d} - 1} \right) \end{aligned} \quad (6.30)$$

Finally, use (6.30) and the relationship

$$\max_{k \geq 0} y^d[k] = y_{os}^d + y_\infty^d$$

to obtain (6.24). □

Note that the bound in Theorem 6.5.2 is equally valid with a_1 and a_2 exchanged, but neither of the two arrangements yields a consistently tighter bound. Also note from the control configuration of Figure 6.9 that $H^d(z)$ (the closed-loop transfer function from $R^d(z)$ to $Y^d(z)$) will have a system delay (relative degree) at least equal to the system delay of $P^d(z)$. Hence, the settling time of $y^d[k]$ can be no smaller than the system delay of $P^d(z)$. If $P^d(z)$ is strictly proper, then $t_s^d \geq 1$ and the bound in Theorem 6.5.2 is positive only for certain arrangements of a_1 and a_2 . In contrast, the analogous continuous-time result of Theorem 6.2 does not have a similar lower limit and as the settling time approaches zero, there will be step response overshoot irrespective of the positions of the real NMP zeros.

6.5.3 Case where $H^d(z)$ has One NMP Zero

While Theorem 6.5.2 establishes that step response overshoot can result from the presence of two positive real NMP zeros and sufficiently small settling time, a similar result does not hold for one positive real NMP zero. The following theorem, the discrete-time counterpart to Theorem 6.3, shows that one positive real NMP zero does not necessarily lead to overshoot, regardless of the speed of response.

Theorem 6.5.3: Let $H^d(z)$ be a stable strictly proper transfer function with exactly one positive real NMP zero at $z = a > 1$. Then there does not exist a non-trivial lower bound on the step response overshoot of $H^d(z)$.

Proof: Consider the special case where all the poles are real and distinct:

$$H^d(z) = \frac{-K(z - a)}{\prod_{i=1}^{n_p} (z - p_i)}. \quad (6.31)$$

In (6.31), $K > 0$, $a > 1$, and $1 > p_{n_p} > p_{n_p-1} > \dots > p_1 > 0$ with $n_p \geq 2$. As usual, let $y^d[k]$ denote the step response of $H^d(z)$. Then, $y^d[k]$ exhibits the following characteristics:

1. $y^d[0] = 0$ (since $H^d(z)$ is strictly proper);
2. $y_\infty^d > 0$ (by direct calculation of $H^d(1)$);

3. $y^d[k]$ exhibits undershoot (using Theorem 6.5.1); and
4. $y^d[k]$ has exactly one extrema for $k > 0$ (using the results in [56]).

These four properties imply that $y^d[k]$ does not exhibit overshoot, independent of the zero location or the value of t_s^d . Consequently, there does not exist a non-trivial lower bound on the overshoot. \square

As with Theorem 6.3, Theorem 6.5.3 does not imply that there is never closed-loop step response overshoot for a plant with a single positive NMP zero. However, both theorems establish that a single NMP zero does not *necessarily* lead to step response overshoot.

Chapter 7

Summary and Future Work

This work proposed control schemes to improve the effectiveness of conformal and intensity-modulated radiotherapy cancer treatment by using image feedback to track the tumour during treatment. By tracking the tumour, a smaller PTV can be used in conformal radiotherapy, reducing the harm to healthy tissue surrounding the tumour. The proposed control scheme for intensity-modulated radiotherapy, in addition to tracking the tumour, also estimates the dose profile as the treatment progresses. This provides more accurate radiation delivery and a post-treatment dose profile that is closer to a reference profile than that achieved with traditional intensity-modulated radiotherapy.

Motivated by the need to determine the current tumour position, this work also investigated estimating the output of an *unknown* affine time-invariant system when a sensor delay prevents its direct measurement. An algorithm was developed for a discrete single rate framework that provides asymptotic estimation of the output under mild assumptions on the system.

Finally, this work examined the design of a controller to provide an arbitrarily fast nonovershooting step response, a requirement arising from the need to prevent opposing collimator leaf collisions. A new performance limitation was derived that proves it is impossible to design such a controller when the plant contains two or more real nonminimum phase zeros.

Avenues for future work on the radiotherapy control system include:

1. Developing a technique to estimate the tumour and collimator leaf positions when the breathing model is unknown. Recall that this estimation problem operates in a *multirate* framework, whereas the algorithm developed in Chapter 5 applies to *single rate* topologies.

2. Preventing opposing collimator leaf collisions (by, for example, designing the leaf controller to provide a nonovershooting response).
3. Developing an observer design technique that explicitly penalizes errors in the tumour and leaf position estimates. Recall that the LQR optimized observer described in Section 3.1 penalized a related, but not identical, cost index.
4. Devising a method to toggle the radiation beam on and off depending on the certainty of the tumour and leaf position estimates.
5. Quantifying the ability to tolerate random image processing failures (caused by a failure to determine the tumour location in an image, for example).
6. Accounting for deformable and rotating tumours.
7. Using a more realistic collimator leaf model (e.g., a stepper motor model) for the leaf controller design.
8. Designing a tumour and leaf position estimator to accommodate image processing noise (caused by quantization error in the image processing algorithm, for example).

Portions of this thesis have been previously published and have been used with permission from the IEEE. The specific works are as follows:

J. Stewart and D.E. Davison, “Conformal radiotherapy cancer treatment with multileaf collimators: improving performance with real-time feedback,” *Proceedings of the IEEE Conference on Control Applications*, Aug. 2005, pp. 125–130. © 2005 IEEE.

J. Stewart and D.E. Davison, “Dose control in radiotherapy cancer treatment: improving dose coverage with estimation and feedback,” *Proceedings of the American Control Conference*, June 2006, pp. 4806–4811. © 2006 IEEE.

J. Stewart and D.E. Davison, “On overshoot and nonminimum phase zeros,” *IEEE Trans. Auto. Control*, vol. 51, pp. 1378–1382, Aug. 2006. © 2006 IEEE.

Bibliography

- [1] F.M. Khan, *The Physics of Radiation Therapy*. Philadelphia: Lippincott Williams & Wilkins, 2003.
- [2] R.L. Souhami, J. Moxham, *Textbook of Medicine*. Philadelphia: Elsevier Health Sciences, 2004.
- [3] H.K. Awwad, *Radiation Oncology: Radiobiological and Physiological Perspectives*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1990.
- [4] M. Sharpe, L. Dawson, “Introduction to Radiation Oncology,” *Guest Lecture*, Toronto, November 2004.
- [5] International Commission on Radiation Units and Measurements, *ICRU Report No. 50. Prescribing, recording and reporting photon beam therapy*. Bethesda, Maryland, 1993.
- [6] D.E. Davison and E.S. Hwang, “Automating radiotherapy cancer treatment: use of multirate observer-based control,” *Proceedings of the American Control Conference*, June 2003, pp. 1194–1199.
- [7] X. Li, “Application of feedback control and image processing in radiotherapy,” master’s thesis, Department of Electrical and Computer Engineering, University of Waterloo, 2003.
- [8] D. Jaffray, J. Siewerdsen, J. Wong, et al., “Flat panel cone-beam computed tomography for image-guided radiation therapy,” *Int. J. Radiation Oncology Biol. Phys.*, vol. 53, pp. 1337–1349, 2002.
- [9] J. Balter, G. Chen, C.A. Pelizzari, et al., “Online repositioning during treatment of the prostate: A study of potential limits and gains,” *Int. J. Radiation Oncology Biol. Phys.*, vol. 27, pp. 137–143, 1993.

- [10] K. Kitamura, H. Shirato, Y. Seppenwoolde, et al., “Tumor location, cirrhosis, and surgical history contribute to tumour movement in the liver, as measured during stereotactic irradiation using a real-time tumor-tracking radiotherapy system,” *Int. J. Radiation Oncology Biol. Phys.*, vol. 56, pp. 221–228, 2003.
- [11] C. Ozhasoglu and M.J. Murphy, “Issues in respiratory motion compensation during external-beam radiotherapy,” *Int. J. Radiation Oncology Biol. Phys.*, vol. 52, pp. 1389–1399, 2002.
- [12] C.F. Serago, S.J. Chungbin, S.J. Buskirk, et al., “Initial experience with ultrasound localization for positioning prostate cancer patients for external beam radiotherapy,” *Int. J. Radiation Oncology Biol. Phys.*, vol. 53, pp. 1130–1138, 2002.
- [13] M.A. Mosleh-Shirazi, P.M. Evans, W. Swindell, et al., “A cone-beam megavoltage CT scanner for treatment verification in conformal radiotherapy,” *Radiotherapy and Oncology*, vol. 48, pp. 319–328, 1998.
- [14] H.D. Kubo and B.C. Hill, “Respiration gated radiotherapy treatment: a technical study,” *Phys. Med. Biol.*, vol. 41, pp. 83–91, 1996.
- [15] J. Stewart and D.E. Davison, “Conformal radiotherapy cancer treatment with multileaf collimators: improving performance with real-time feedback,” *Proceedings of the IEEE Conference on Control Applications*, Aug. 2005, pp. 125–130.
- [16] B.D.O. Anderson and J.B. Moore, *Optimal Control: Linear Quadratic Methods*. Englewood Cliffs, New Jersey: Prentice Hall, 1990.
- [17] D.E. Davison, E.S. Hwang and X. Li, “Generalization of the separation principle beyond constant-gain state-feedback control,” *Proceedings of the American Control Conference*, June 2003, pp. 3228–3233.
- [18] S.C. Davies, A.L. Hill, R.B. Holmes, et al., “Ultrasound quantification of respiratory organ motion in the upper abdomen,” *British J. of Radiology*, vol. 67, pp. 1096–1102, 1994.
- [19] P.J. Bryan, S. Custar, J.R. Haaga and V. Balsara, “Respiratory movement of the pancreas: An ultrasonic study,” *J. of Ultrasound in Medicine*, vol. 3, pp. 317–320, 1984.

- [20] A.S. Shiu, H.M. Kooy and J.R. Ewton, “Comparison of miniature multileaf collimation (MMLC) with circular collimation for stereotactic treatment,” *Int. J. Radiation Oncology Biol. Phys.*, vol. 37, pp. 679–688, 1997.
- [21] V. Kučera, *Analysis and Design of Discrete Linear Control Systems*. Systems and Control Engineering, Hertfordshire, England: Prentice Hall International, Ltd., 1991, pp. 177–183.
- [22] J. Stewart and D.E. Davison, “Dose control in radiotherapy cancer treatment: improving dose coverage with estimation and feedback,” *Proceedings of the American Control Conference*, June 2006, pp. 4806–4811.
- [23] A. Pirzkall, M. Carol, F. Lohr, et al., “Comparison of intensity-modulated radiotherapy with conventional conformal radiotherapy for complex-shaped tumors,” *Int. J. Radiation Oncology Biol. Phys.*, vol. 48, pp. 1371–1380, 2000.
- [24] E. Huang, B.S. Teh, D. Strother, et al., “Intensity-modulated radiation therapy for pediatric medulloblastoma: early report on the reduction of toxicity,” *Int. J. Radiation Oncology Biol. Phys.*, vol. 52, pp. 599–605, 2002.
- [25] D.P. Gierga, G.T.Y. Chen, J.H. Kung, et al., “Quantification of respiration-induced abdominal tumor motion and its impact on IMRT dose distributions,” *Int. J. Radiation Oncology Biol. Phys.*, vol. 58, pp. 1584–1595, 2004.
- [26] D.J. Little, L. Dong, L.B. Levy, et al., “Use of portal images and BAT ultrasonography to measure setup error and organ motion for prostate IMRT: implications for treatment margins,” *Int. J. Radiation Oncology Biol. Phys.*, vol. 56, pp. 1218–1224, 2003.
- [27] T. Soderstrom and P. Stoica, *System Identification*. Systems and Control Engineering, Hertfordshire, England: Prentice Hall International, Ltd., 1989.
- [28] S. Takahashi and T. Sekozawa, “Air-fuel ratio control in gasoline engines based on state estimation and prediction using dynamic models,” *Proceedings of the IEEE Conference on Industrial Electronics, Control, and Instrumentation*, November 1995, pp. 217–222.
- [29] C. F. Mandenius, B. Mattiasson, J. P. Axelsson, et al., “Control of an ethanol fermentation carried out with alginate entrapped *Saccharomyces cerevisiae*,” *Biotechnology and Bioengineering*, vol. 29, pp. 941–949, 1997.

- [30] L. Ljung and T. Söderström, *Theory and Practice of Recursive Identification*. MIT Press Series in Signal Processing, Optimization and Control, Cambridge, Massachusetts: MIT Press, 1986.
- [31] I. Mareels and J. W. Polderman, *Adaptive Systems, an Introduction*. Boston, Massachusetts: Birkhäuser, 1996, pp. 63–67.
- [32] G.C. Goodwin and K.S. Sin, *Adaptive Filtering Prediction and Control*. Information and Systems Science Series, Englewood Cliffs, New Jersey: Prentice Hall, Inc., 1984, pp. 49–54.
- [33] D. Miller, *Adaptive Control Course Lecture Notes*. Department of Electrical and Computer Engineering, Univ. of Waterloo, 2006.
- [34] E. Lefeber, J.W. Polderman, “On the possible divergence of the projection algorithm,” *IEEE Trans. Auto. Control*, vol. 40, pp. 495–497, 1995.
- [35] I. Mareels and J.W. Polderman, *Adaptive Systems, an Introduction*. Boston, Massachusetts: Birkhäuser, 1996, pp. 80–82.
- [36] B. Delyon, R. Izmailov and A. Juditsky, “On the projection algorithm and delay of peaking in adaptive control,” *IEEE Trans. Auto. Control*, vol. 38, pp. 581–584, 1993.
- [37] B.D.O. Anderson, P.R. Bitmead, C.R. Johnson, et al., *Stability of Adaptive Systems: Passivity and Averaging Analysis*. MIT Press Series in Signal Processing, Optimization and Control, Cambridge, Massachusetts: MIT Press, 1986.
- [38] G.C. Goodwin and K.S. Sin, *Adaptive Filtering Prediction and Control*. Information and Systems Science Series, Englewood Cliffs, New Jersey: Prentice Hall, Inc., 1984, pp. 106–117.
- [39] J. Stewart and D.E. Davison, “On overshoot and nonminimum phase zeros,” *IEEE Trans. Auto. Control*, vol. 51, pp. 1378–1382, Aug. 2006.
- [40] J. De Schutter, “Improved force control laws for advanced tracking applications,” *Proceedings of the IEEE Conference on Robotics and Automation*, April 1988, pp. 1497–1502.
- [41] D.M. McDowell, G.W. Irwin and G. McConnell, “A neural autopilot for bank-to-turn control,” *Proceedings of the IEE Colloquium on Multivariable Methods for Flight Control Applications*, May 1994, pp. 8/1–8/5.

- [42] S. Darbha and S.P. Bhattacharyya, "On the synthesis of controllers for a nonovershooting step response," *IEEE Trans. Auto. Control*, vol. 48, pp. 797–800, 2003.
- [43] S. Jayasuriya and J.-W. Song, "On the synthesis of compensators for nonovershooting step response," *ASME J. of Dyna. Syst., Measure., and Control*, vol. 118, pp. 757–763, 1996.
- [44] Y.C. Kim, L.H. Keel, and S.P. Bhattacharyya, "Transient response control via characteristic ratio assignment," *IEEE Trans. Auto. Control*, vol. 48, pp. 2238–2244, 2003.
- [45] J.S. Freudenberg and D.P. Looze, "Right half-plane poles and zeros and design tradeoffs in feedback systems," *IEEE Trans. Auto. Control*, vol. 30, pp. 555–565, 1985.
- [46] M. Vidyasagar, "On undershoot and nonminimum phase zeros," *IEEE Trans. Auto. Control*, vol. 31, pp. 440-440, 1986.
- [47] R.H. Middleton, "Trade-offs in linear control system design," *Automatica*, vol. 27, pp. 281-292, 1991.
- [48] L. Qiu and E. Davison, "Performance limitations of nonminimum phase systems in the servomechanism problem," *Automatica*, vol. 29, pp. 337-349, 1993.
- [49] J. Chen, L. Qiu and O. Toker, "Limitations on maximal tracking accuracy," *IEEE Trans. Auto. Control*, vol. 45, pp. 326-331, 2000.
- [50] K. Lau, R. Middleton and J. Braslavsky, "Undershoot and settling time trade-offs for nonminimum phase systems," *IEEE Trans. Auto. Control*, vol. 48, pp. 1389-1393, Aug. 2003.
- [51] M.M. Seron, J.H. Braslavsky and G.C. Goodwin, *Fundamental Limitations in Filtering and Control*. Communications and Control Engineering, London: Springer-Verlag, 1997.
- [52] J. Freudenberg, R. Middleton and A. Stefanopoulou, "A survey of inherent design limitations," *Proceedings of the American Control Conference*, June 2000, pp. 2987–3001.
- [53] M. El-Khoury, O.D. Crisalle and R. Longchamp, "Influence of zero locations on the number of step-response extrema," *Automatica*, vol. 29, pp. 1571–1574, 1993.

- [54] G. Deodhare and M. Vidyasagar, “Design of non-overshooting feedback control systems,” *Proceedings of the Conference on Decision and Control*, December 1990, pp. 1827–1834.
- [55] B.L. De La Barra, M. El-Khoury and M. Fernandez, “On undershoot in scalar discrete-time systems,” *Automatica*, vol. 32, pp. 255–259, 1996.
- [56] M. El-Khoury, O.D. Crisalle and R. Longchamp, “Discrete transfer-function zeros and step response extrema,” *Proceedings of the IFAC World Conference*, July 1993, pp. 81–86.