On-line control of d-Tubocurarine induced muscle relaxation

A thesis presented in fulfilment of the requirements for the degree of Doctor of Philosophy in the Department of Electrical and Electronic Engineering at the University of Cape Town by Larry Bertram Rametti.

- January 1985

DECLARATION

I declare that this thesis represents my original work both in concept and in execution except where specific reference has been made to the contribution of others.

signature removed

Larry Bertram Rametti

January 1985

ACKNOWLEDGEMENT

I wish to express my gratitude to those persons in the Department of Electrical and Electronic Engineering and in the Department of Anaesthetics at the University of Cape Town who assisted me in this research.

I feel that I must mention by name my supervisor Professor H.S. Bradlow and thank him for his support - especially for his unfailing enthusiasm.

I realise that many people contribute to work of this nature - sometimes beyond the call of duty - although only one receives credit for doing it. I therefore acknowledge my debt to these nameless people and hope to reciprocate their good actions in so far as I am able.

ABSTRACT

A microcomputer based control system designed around the 8088 Central Processing Unit and 8087 Numerics Data Processor has been developed to induce and then maintain a desired level of muscle relaxation in patients who undergo surgical proceedures. The neuromuscular blocking effect is measured by obtaining the evoked twitch response to Ulnar nerve stimulation and the computer is interfaced to a motorized syringe pump which administers the relaxant. The drug d-Tubocurarine was chosen as it is one of the longest acting of all commonly used relaxants and is consequently believed to be one of the most difficult to control. The methodology developed in the thesis can therefore easily be applied to other shorter acting drugs.

In the off-line analysis of the patient response, a non-linear model is fitted directly to the patient blocking effect (i.e. plasma concentration data are not required) using a non-linear Least Squares method. Stochastic effects are accounted for and the resulting model is shown to simulate the patient response realistically for control purposes. The variability in the effect response to d-Tubocurarine is established, based on the data of a total of 44 patients.

It was shown by simulation that simple Proportional plus Integral plus Derivative control was suitable for maintaining a previously induced level of paralysis but was not suitable for the transitional control phase. Nor could the self-tuning controlier of Clarke and Gawthrop deal successfully with this phase. The reason is shown to be the positive input nature of the patient which disallows negative control commands. A novel, computationally intensive control algorithm was therefore designed for this application. In it two patient model parameters are estimated on-line by a non-linear Least Squares method which uses Steepest Descent optimization and the patient states are approximated by a

simple linear estimator. The patient twitch response, the drug commands issued by the controller and the parameters estimated on-line are logged onto cassette tape during each control operation and are transferred subsequently to a UNIVAC 11/Ø8 mainframe computer for off-line analysis of the patient response and the controller performance.

Initial controller settings were obtained by testing the algorithm on simulated patients. These settings concerned patient model parameters for response to the drug, bandwidth settings for the State Estimator, the desired length of the transitional control phase, the spacing of the calculated drug commands and the patient circulatory dead-time.

The clinical controller trials provided data on 34 patients, all adults of ASA grading 1 or 2 who underwent various surgical proceedures. The controller settings were refined as experience was gained in controlling these real patients.

In the off-line analysis of the controller performance, the correct functioning of the on-line parameter estimation routine was verified, the efficacy of on-line parameter estimation was established and the performance of the State Estimator investigated. On average the on-line Least Squares method resulted in a residual sum of squares which was only 3% inferior to that calculated off-line with a sophisticated optimization package. It was found that the two parameters estimated on-line resulted in a significant reduction in the time taken to induce relaxation in the clinical trials performed over the case where no parameters are estimated - which is a major advantage - and further that the control was not far inferior to the hypothetical ideal which could be achieved if all the patient model parameters were known a priori. Finally the best bandwidth of the State Estimator was found to be that which was suggested by the controller trials on simulated patients performed prior to the clinical trials.

The off-line analysis showed however that the shortening of the induction period could have been achieved without parameter estimation had a better selection of the controller settings been made thus indicating that the variability in patient response to d-Tubocurarine encountered was insufficient to justify on-line parameter estimation with the control algorithm in use although it did render the selection of the settings which pertain to the patient model parameters less critical.

The controller was shown to operate successfully on the entire range of patient responses encountered. On average the degree of relaxation which was necessary to depress either the single twitch response or the train-of-four ratio by 80% was induced within 12,1 ± 6,2 minutes with an initial overshoot of 5,9 ± 3,4 %. Thereafter the relaxation was maintained with an off-set from set-point of 2,0 ± 1,6 % and standard deviation of 3,4 ± 1,6 % about this value. 80% depression of the single twitch response was found to be adequate for both lower and upper abdominal surgery whereas control to a 80% depression of the train-of-four ratio was found to be inadequate for either.

The contributions of this work are that it develops the most detailed model of the patient response to d-Tubocurarine yet reported thus enabling realistic patient simulations for controller evaluation. It is the first controller to demonstrate the benefit of on-line parameter estimation in the induction phase of control. Finally it is the first to report on clinical controller trials which use d-Tubocurarine and to demonstrate that control is successful on a wide range of patient responses.

CONTENTS OF THE THESIS

DEFINITION OF TERMS

Control period - the spacing between successive drug commands issued by the controller (typically 1 minute).

Data buffer - 256 bytes of RAM in the microcomputer used as tempory storage for patient data obtained from clinical trials. This data is written to/read from the cassette tape unit.

Deterministic model - that which produces an output in response to deterministic inputs.

Extreme Parameter Vector - vector of 4 parameters which defines a deterministic model that simulates a conservative patient response. It is used in calculating the sequence of injections issued during the initial phase of control.

Induction period - the total time spent in the initial and relaxation phases of control.

Initial dead-time - the time between the administration of the loading dose and onset of change in the effect produced. The threshold of onset is taken to be Ø,85 of the twitch ratio that is monitored or controlled.

Initial injection sequence - the drug commands issued by the Initial Phase Controller.

Initial Parameter Vector - vector of 4 parameters which defines a deterministic model with average patient response. It is used in the relaxation and regulation phases of control.

Initial phase of control - this extends throughout the initial dead-time of the patient i.e. it begins when the relaxant is first administered and ends when the twitch ratio being controlled drops below Ø,85.

Loading dose - the initial, and usually large, bolus dose of drug which begins the relaxation process.

Noise model - that which contributes to the total model output but is independent of deterministic inputs.

Recovery phase - this begins at the termination of on-line control and continues until the patient has adequately recovered from the effect of the relaxant. A reversal drug may be administered during this phase.

Regulation phase of control - this begins when the patient first reaches the set-point and continues until the end of control. During this period it is necessary to maintain the previously induced level of relaxation and possibly to follow changes in set-point.

Relaxation phase of control - this extends from the time that the Ø,85 twitch ratio threshold is crossed until the set-point is reached.

Residual	- error found at a particular time when comparing a process and its model.
Sampling period	- the spacing between successive measurements of the patient twitch response (typically 10 seconds).
Sensitivity of the patient	- the bolus dose of relaxant expressed in micrograms per kilogram of body mass which will depress the twitch ratio of the initially unrelaxed patient by 80%.
Stochastic model	- see Noise model.
Time to maximum effect	- the time between the administration of a single bolus dose of relaxant to an unrelaxed patient and the instant that the induced relaxation reaches maximum effect.
Train-of-four (T4) ratio	- the ratio of the fourth twitch amplitude to the first in response to a train-of-four pulse stimulation.
Train-of-two (T2) ratio	- the ratio of the second twitch amplitude to the first in response to a train-of-two (or train-of-four) pulse stimulation.
Twitch control height	- the amplitude of the evoked twitch response just prior to administration of the relaxant.
T1 ratio	- the ratio of the amplitude of a single evoked twitch or of the first twitch in a train-of-four response to the twitch control height.

EXPLANATION OF THE PRINCIPLE ABBRIVIATIONS

AR - Autoregressive

ARMA - Autoregressive moving average

CTIC - Cassette tape interface card

dTC - d-Tubocurarine (curare)

EE - Equation Error

EMG - Electromyogram

LS - Least Squares

MA - Moving average

ML - Maximum Likelihood

NDP - Numerics Data Processor

OE - Output Error

SABUS - South African Standard Bus

T1 - Single evoked twitch response

T2 - Train-of-two twitch response

T4 - Train-of-four twitch response

EXPLANATION OF THE PRINCIPLE SYMBOLS

$A(z^{-1})$ - polynomial in the time shift operator z^{-1} used in the patient deterministic model.

a_i - parameters of the polynomial $A(z^{-1})$.

$B(z^{-1})$ - polynomial in the time shift operator z^{-1} used in the patient deterministic model.

BT4 - parameter of the logistic relationship pharmacodynamic equation.

$C(z^{-1})$ ... $D(z^{-1})$ - polynomials in the time shift operator z^{-1} used in the patient noise model.

d_i - parameters of the polynomial $D(z^{-1})$.

DSTEP - constant used in the numerical approximation to the gradient vector and in the linear search for the Steepest Descent method.

$e(k)$ ... $E(z^{-1})$ - discrete time white noise source which drives the patient noise model.

$f(k)$ ... $F(z^{-1})$ - output of a patient noise model.

g - gain parameter of the patient model.

G - gain vector of the Kalman Filter or State Estimator.

G_i - elements of the gain vector of the Kalman Filter or State Estimator.

H - Output vector of the deterministic model in the State description of the patient.

H_n - Output vector of the combined deterministic and stochastic models in the State description of the patient.

i - a counter used for transitional control in the relaxation phase. (Equation 4.2)

k - discrete time at sampling period 'k'.

K - parameter of the Hill equation.

ℓ - patient circulatory delay expressed in sampling periods.

L - intended length of the relaxation phase.

P_c - corrected covariance matrix of the Kalman filter.

P_p - predicted covariance matrix of the Kalman filter.

p_1, p_2 - poles of the two-compartment patient model.

r - residual formed in the comparison of process and its model.

r_c - correlated residuals derived from fitting the deterministic model alone.

R_e - variance of the process measurement noise.

R_w - variance of the process driving noise.

SI - 8086/8 CPU register which is used as an index for array elements in the Short-real format. SI = 4.k

SP - set-point.

SP_{des} - desired set-point. $SP_{des} = SP_{lin} - BT4$

SP_{lin} - linearised set-point. $SP_{lin} = \ln[(1-SP)/SP]$

SP_t - intermediate set-point chosen for transitional control in the relaxation phase.

t - continuous time

T - fractional blocking effect on the range (0,1).

T_c - control period.

T_{lin} - linearized patient twitch response. $T_{lin} = \ln[(1-T)/T]$

T_s - sampling period

$u(k)$ / $U(z^{-1})$ — drug input to the patient model or drug command calculated by the controller.

V — loss function (residual sum of squared errors).

w — white process driving noise source.

W — weighing sequence with values 1 or Ø.

x — output of a deterministic patient model ; or the patient states used in the State formulation of the patient deterministic model.

$X(z^{-1})$ — output of the deterministic patient model.

$\underline{X}$ — state vector of the deterministic patient model. (Equation 4.1)

x_{max} — the greatest output of the patient deterministic model possible for given initial conditions and excitation.

$\underline{X}_n$ — state vector of the combined deterministic and stochastic patient models (Equation 4.3) that is used in the Kalman filter.

$\overline{\underline{X}}_n$ — predicted estimate of $\underline{X}_n$

$\hat{\underline{X}}_n$ — corrected estimate of $\underline{X}_n$

x_T — states of $\underline{X}_T$

$\underline{X}_T$ - true patient state vector. (Unknown to the controller).

$y(k)$, $Y(z^{-1})$ - output of the combined deterministic and stochastic patient models.

$y_F(k)$, $Y_F(z^{-1})$ - filtered output of the combined deterministic and stochastic patient models. $Y_F(z^{-1}) = D(z^{-1}).Y(z^{-1})$

z^{-1} - backward time shift operator.

γ - parameter of the Hill equation pharmacodynamic relationship.

Γ - input vector of the deterministic model in the State description of the patient.

Γ_n - input vector of the combined deterministic and stochastic models in the State description of the patient.

$\underline{\xi}$ - unit vector in the direction of steepest descent used in the on-line parameter estimation algorithm.

λ - the distance in the direction of steepest descent calculated by the linear search.

ϕ - Transition matrix of the deterministic model in the State description of the patient.

ϕ_n - Transition matrix of the combined deterministic and stochastic models in the State description of the patient.

$\underline{\theta}$ - true parameter vector of the patient model (unknown to the controller). This comprises $\underline{\theta}_D$ and $\underline{\theta}_S$.

$\underline{\theta}_D$ - parameter vector of the deterministic patient model.

$\underline{\theta}_E$ - Extreme Parameter Vector used in the Initial Phase Controller.

θ_i - elements of $\underline{\theta}$

$\underline{\theta}_S$ - parameter vector of the stochastic patient model.

CHAPTER 1

Introduction

1.1 The elements of anaesthesia

In modern anaesthesia, the anaesthetist is required to monitor a number of physiological parameters in patients who undergo surgical proceedures and to control some of them. Blood pressure, pulse rate and ECG response are routinely monitored throughout the operation. The depth of sleep and degree of analgesia can be varied although not readily controlled because they are difficult to estimate.

Muscle relaxation forms part of anaesthesia. It enables the surgeon to reach and operate on inaccessible organs without interference from muscle action.

1.2 How muscle relaxation is effected at present

The anaesthetist and surgeon can be thought of as comprising a crude control loop as shown in figure 1.1. In this context the surgeon is the transducer. He advises the anaesthetist of any undesirable patient movement which may indicate that the level of relaxation is inadequate. This may be the only feedback signal. The administration of the relaxant is based on the anaesthetist's experience of the drug and such movement.

More recently anaesthetists have started occasionally using a Nerve Stimulator to assist them in assessing the degree of blockade particularily

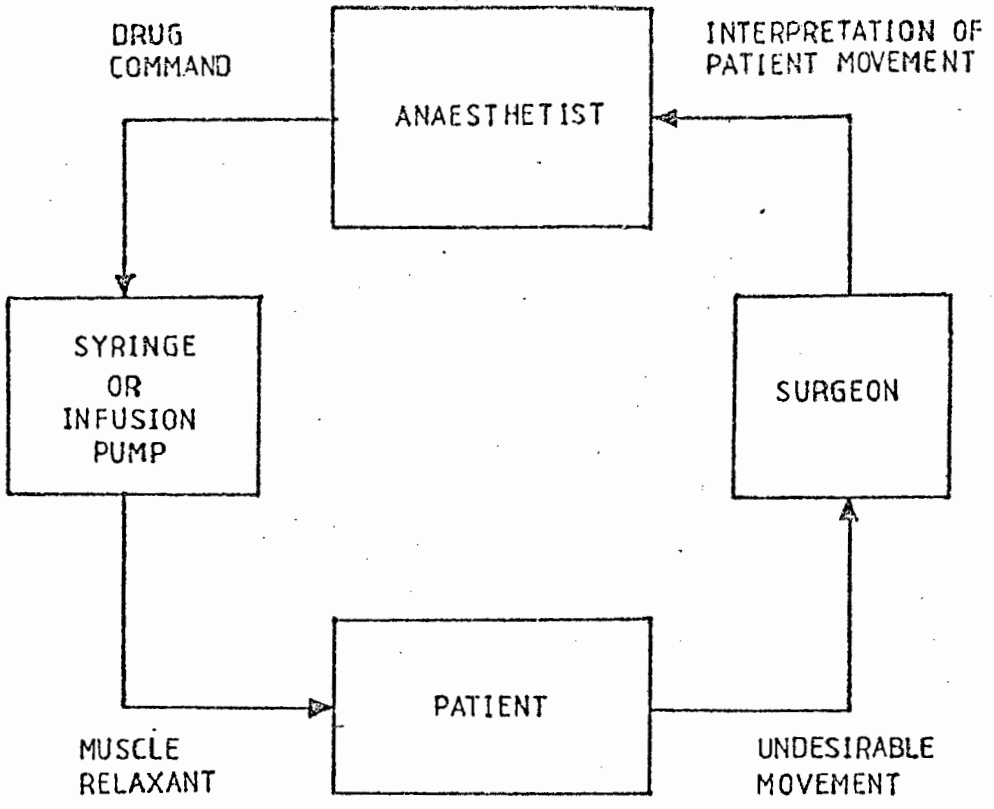

Figure 1.1 : The usual control loop for muscle relaxation in the operating theatre.

during the recovery phase of the operation. The patient is typically stimulated at the median or ulnar nerve. The anaesthetist may simply judge the strength of the evoked twitch response in the hand visually or he may measure it. This provides an objective norm by which to judge the level of relaxation although unfortunately not at the site that the surgeon requires it.

There are two techniques in use for measuring the evoked response :

1. the force of thumb adduction measured by a suitable transducer and amplified to produce a proportional voltage

2. a measured electromyogram (EMG) signal is processed (usually amplified, rectified and integrated over the period that a response to the stimulus is expected) to produce a voltage that represents the

electrical activity of the relevant part of the muscle.

This work considers the first of these methods only. It was chosen because it is less susceptible to surgical diathermy interference than the EMG response (Brown et al.,1980) and because of its simplicity. It was used specifically to measure the response to train-of-four stimulation. Thus there are available three useful measures of relaxation (Ali et al.,1971), namely the Train-of-four (T4), Train-of-two (T2) and Single Twitch response (T1) ratios. Each of these three measures were considered in this research. The EMG response on the other hand requires more sophisticated electronic processing and technical difficulties associated with it preclude its routine clinical use (Ali and Savarese,1976).

1.3 Motivation for automatic control

Automatic control of muscle relaxation would obviously relieve the anaesthetist from this task. This is especially desirable for short acting drugs such as Succinyldicholine (Scoline), ORG NC45 and Atracurium which are tedious to control manually. There is however an important reason for controlling the degree of paralysis that is related to patient safety.

The work of Ham et al.,(1979) on d-Tubocurarine (curare or dTC), shows that it is desirable to maintain blocking at a level not exceeding 90% depression of the Single Twitch response because there is then no delay in the patient recovering to a level where antagonism with anticholinesterase agents is effective. While it is possible for the anaesthetist to achieve this clinically, and at the same time induce and maintain an adequate level of relaxation by carefully controlling the relaxant dosage schedule, this can make undue demands on his time and skills. The use of an automatic control system is

thus suggested. In addition it is desirable from a practical point of view to induce relaxation as quickly as possible. This can be achieved by administering a large bolus dose, as is usually done for muscle relaxants, but the resulting overshoot of the set-point due to the excessive dose used is undesirable for the reason given by Ham et al. and because the extent of the side effects attributable to muscle relaxants is determined by the peak plasma concentration (Uys et al.,1984). The possibility of an extended time to return to a reversible level of relaxation is particularily undesirable if for some reason the patient needs to be reversed shortly after the relaxant has been administered or in respect of patients who have kidney diseases and whose ability to eliminate the drug is consequently impaired. Excessive overshoot also increases the danger of paralysis reoccuring in the recovery room. This is known as 'residual curarization' (Viby-Mogensen et al., 1979).

The postulates of an automatic control system are then that it be capable of bringing an initially unrelaxed patient to a selectable set-point in as short a time as possible with minimum overshoot and that it be able to maintain that level within an acceptible band thereafter with possible set-point changes.

1.4 The difficulties associated with the control problem

The control of d-Tubocurarine induced muscle relaxation has a number of difficulties which characterise the problem :

1. Curare has a long duration of action (typically tens of minutes).

2. There is inter-patient variability in response to the drug. This was noted by Bradlow et al.,(1983) and the extent of this variability is established more accurately in the present work.

3. The pharmacodynamic relationship is non-linear (Ham et al.,1979). Certain patients will show no visible block for a given loading dose while double this dose will depress the train-of-four response completely.

4. The controller acts on what is essentially a 'positive input system' and is thereby constrained to issue commands of a single polarity (the drug cannot be removed from the patient at will once administered). It was found that this restriction must be considered if the initial transitional control is to be accomplished with safety. It is shown in chapter 4 however that the restriction is not particularily serious if regulation is the only aim. Acceptable regulation is easier to achieve than transitional control which is required in this work.

Positive linear systems and the control of positive input linear systems have been studied at a theoretical level - see for example Luenberger,(1978) and Brammer,(1972) respectively. Their treatment is however beyond the scope of this research.

1.5 The contribution of various researchers

It is only recently that attempts have been made to control muscle relaxation automatically using a feedback control system.

Cass et al.,(1976) describe computer control of four non-depolarizing relaxants (Gallamine, dTC, Alcuronium and Pancuronium) in the sheep. The EMG response was used and the relaxant was administered by a motor driven syringe pump in the on/off mode rather than by varying the rate of administration. Their method however resultes in a fairly long induction period because of the

need to measure the response to an initial test dose in each case.

In a later development to the above work, Brown et al.,(1981) use a microprocessor to control muscle relaxation in sheep and in humans. They also use the evoked EMG response and have a fixed parameter Proportional plus Integral controller but only consider the regulation phase - a bolus dose is administered manually to take the patient through the induction period. When under computer control, the relaxant is similarily administered by a motor driven syringe pump in the on/off mode.

Most recently Berger and Brown, (1984) have developed an adaptive controller for dTC induced muscle relaxation in the sheep. This also uses the evoked EMG response and a motor driven syringe pump to administer the drug. The controller however is not operative until onset of blockade (a standard dosage schedule is used during this period), furthermore adaptation only begins some 13 minutes after the start of control thus effectively excluding the transitional control phase. Their work also does not discuss the important question of inter-patient variability to the drug and how their controller is able to cope with it.

They identify the patient model explicitly with a recursive Least Squares estimator and use this information in a pole-assignment controller. They state that the restriction imposed by the positive input nature of the patient is a particular problem and deal with it by taking a number of precautions one of which is the introduction of an interim set-point which is brought gradually to the desired set-point.

Sheppard et al.,(1982) announce their intention to control muscle relaxation in humans with the short acting depolarising agent Scoline. They advocate microcomputer control of both the induction and regulation phases. The evoked EMG response is to be used to measure the level of relaxation and the

drug administered by infusion pump. The parameters of the controller will not allow adaptation to individual patient responses.

In common with the above researchers, Linkens and co-workers also use the evoked EMG response and administer the relaxant by infusion pump (Linkens et al.,1982). They consider the short acting relaxant NC45 and the longer acting Pancuronium. They have reported on trials undertaken on both dogs and humans using fixed parameter Proportional plus Integral control. These have concentrated on the regulation phase however and it is still necessary to administer a bolus dose manually before switching to on-line control. Simulation results are given for a fixed parameter PID controller which uses a Smith Predictor to deal with the circulatory dead-time. Simulation results for a pole-assignment self-tuning controller are also briefly reported. The simulation trials cover both the induction and regulation phases of control.

As much of the work done by the researchers cited above has dealt solely with the regulation phase of control, the restriction on the controller to issue non-negative commands is not particularily serious. Indeed it was because the regulation phase was of primary concern that the identification studies of Linkens et al.,(1982) have concentrated on a patient model linearized about a set-point. Again the experience of Linkens et al. of the variability in patient responses has led them to consider a self-tuning controller.

In contrast to the work done by these researchers, the control system discussed in this thesis is intended for all phases of control - from administration of the first dose. It specifically includes the non-linearity in the patient pharmacodynamic relationship in its design. It proposes to account for inter-patient variability by on-line parameter estimation (including that of the non-linear element) for use by the controller. The design of the control algorithm specifically allows for operation on a positive input system and is based on dTC which is among the longest acting of all the commonly used

neuromuscular blocking agents (Bevan,1983) and is consequently believed to be the most difficult to control. The methodology developed in this thesis will therefore be directly applicable to the shorter acting drugs.

Besides the muscle relaxation problem, the use of computers in closed loop control has found other applications in areas related to patient care - depth of anaesthesia (Evans, Lampard - referenced in Jacobs,(1982)), Demand analgesia (Jacobs and Reasbeck,1982), Blood-sugar (Albisser,1979), Arrhythmias (Collins and Arzbaecher,1979), Blood pressure (Koivo,1980,1981; Auer and Rodler,1981; Kaufman et al.,1982; Walker et al.,1982; Slate and Sheppard,1982; Arnsparger et al.,1983), Body fluids (Parkin,Feiring,Sheppard - referenced in Jacobs,(1982)) and halothane concentration in anaesthetised patients (Morris et al.,1983). A review of some of these is given by Jelliffe,(1983).

It is instructive to consider briefly the blood pressure problem which has been treated by a number of researchers because the approaches made there immediately spring to mind for use in the control of muscle relaxation. Every research group cited above except Auer and Rodler propose a form of adaptive control to deal with the inter-patient variability in response to the drug. Where clinical trials are reported, all have used drug infusion. All who propose process models have assumed them to be linear with the exception of Slate and Sheppard. There are however different approaches to the control strategy ranging from the simple ad hoc proceedure of Auer, to Model reference techniques (Kaufman) to non-linear multiple-mode control (Slate and Sheppard) to optimal control with a quadratic performance index (of which minimum variance control is a special case). Finally, whilst the biological process being controlled is also a 'positive input system', none of the researchers have seen this as a severe restriction although some have memtioned it in passing.

The control of blood pressure is however simpler than that of muscle relaxation in two respects. Firstly the patient model has a single compartment

in the former case (for Sodium nitroprusside which is commonly used) whereas the response is at least bi-exponential in the latter and then the duration of action of nitroprusside (2 mins according to Koivo,(1980)) is much shorter than dTC which is in excess of 10 mins.

1.6 Outline of the thesis

The modelling of the patient response to dTC is an important step towards on-line control of muscle relaxation because it assists in the design of the control algorithm and enables the controller to be evaluated on simulated patients prior to clinical trials. This is covered in chapter 3.

A FORTRAN programme was written to model off-line the blocking effect of dTC. The effect is modelled directly i.e. plasma concentration data is not required. The programme uses non-linear Least Squares. The deterministic model has 4 parameters (including the non-linearity) and the stochastic model has a maximum of 5.

Chapter 4 deals with the design of the controller. The widely used PID controller and the self-tuning controller of Clarke and Gawthrop,(1975) are obvious candidates for the muscle relaxation problem. These were investigated but found to be unsuitable. A novel, computationally intensive control algorithm was then developed for the specific application. Three versions were produced, two for evaluation on simulated patients, the third for use in the operating theatre. The parameter estimation within the controller is explicit i.e. patient model parameters are estimated, not those of the control law. The model is non-linear. The Least Squares method used incorporates a Steepest Descent algorithm to minimise the residual sum of squares.

Chapter 2 deals with the hardware and software of the control system. Recent advances in microcomputer technology have made the optimization of control laws less critical than in the past. The control algorithm is implemented on a microcomputer which uses a 8088 CPU and a 8087 Numerics Data Processor (NDP). As a result the microcomputer has powerful arithmetic capabilities. It also incorporates a cassette tape unit which is used to log data during clinical trials for subsequent transfer to a mainframe computer and off-line analysis. Programming was done exclusively in assembly language. The microcomputer is interfaced to a Twitch Measurement System which stimulates the patient and measures the response. It is also interfaced to a syringe driven injection pump which administers the relaxant. The Twitch Measurement System was adapted for this research, the injection pump was specifically designed for this purpose. The control system is portable.

Software was written in FORTRAN for off-line analysis of the patient data files created in the operating theatre and also to compute the parameters needed to initialize the control algorithm.

CHAPTER 2

The hardware and software of the control system

2.1 Introduction

This chapter describes the design and performance of a microcomputer based control system which enables the control algorithm developed in the thesis to be implemented on-line in the operating theatre. Specifically

1. it presents the composition of the control system in terms of its logical elements and discusses each.

2. it describes the software of the control system and that written for the off-line analysis of the controller's performance on the UNIVAC 11/Ø8 computer. However the description of the control algorithm itself is deferred to Chapter 5 where a detailed treatment is given.

3. it discusses the performance of the hardware and software of the controller as determined from the clinical trials.

A block diagram of the control system is shown in figure 2.1. and a photograph of the same in figure 2.2. It has three logical sections:

1. The muscle relaxation measurement system. This comprises a nerve stimulator, transducer, twitch measurement system and chart recorder.

2. The microcomputer system. This includes a tape drive for data logging and a keyboard/display unit.

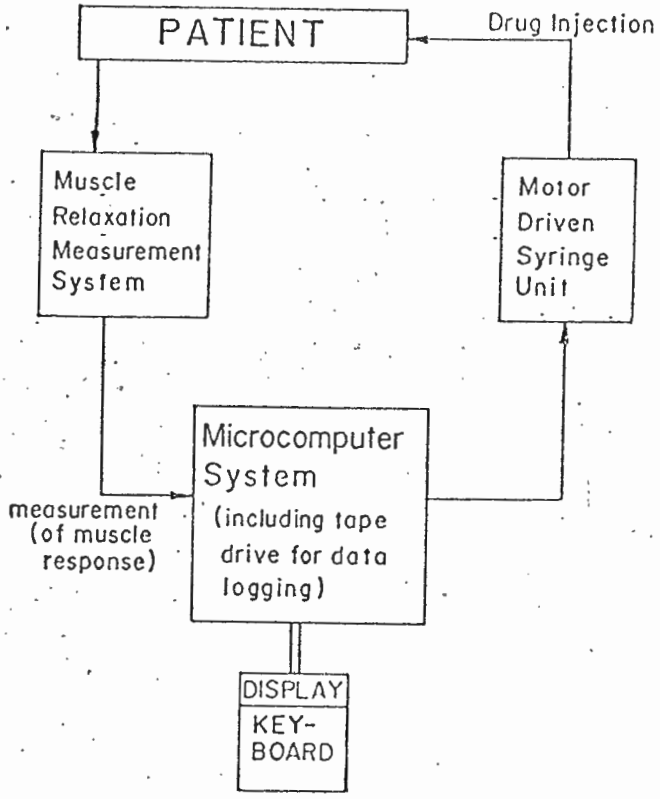

Figure 2.1 : Block diagram of the control system showing its logical sections.

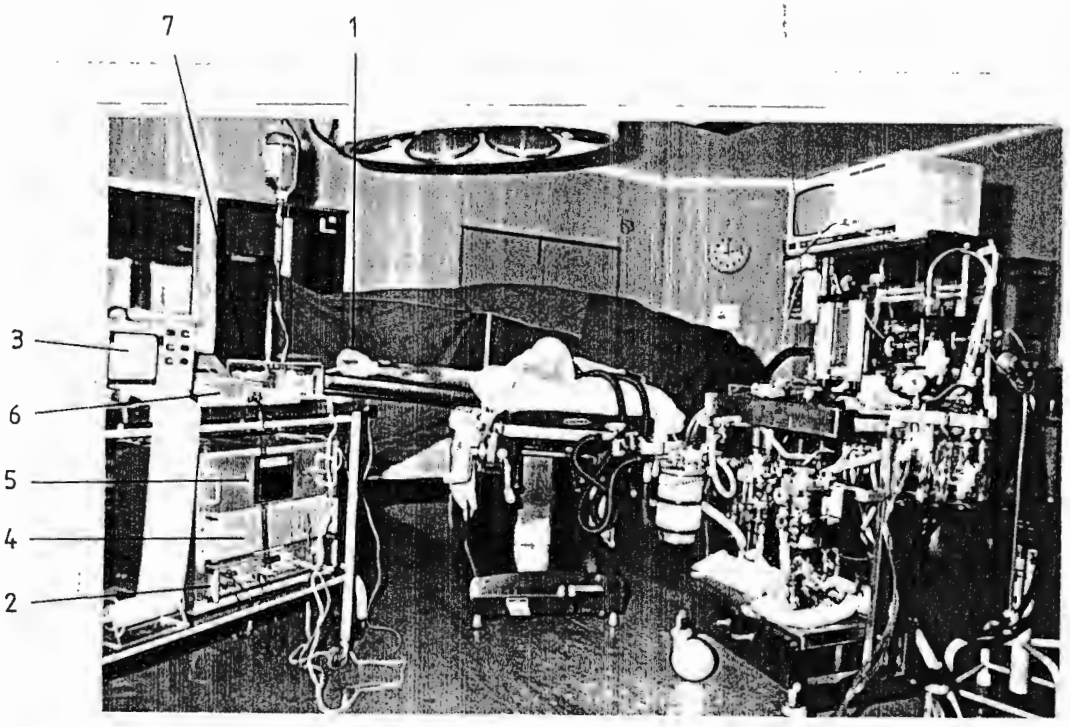

Figure 2.2 : Photograph of the control system connected to a patient in the operating theatre. The components are 1. The transducer. 2. The muscle relaxation measurement system. 3. The chart recorder. 4. The microcomputer. 5. The tape drive. 6. The keyboard and display. 7. The motor driven syringe unit.

3. A motor driven syringe unit for administering the relaxant.

The patient is stimulated at the ulnar nerve through surface or needle electrodes and the force of thumb adduction in response to this is measured by a strain-gauge type transducer. The Twitch Measurement System measures the peak amplitude of each twitch and digitises it so that the microcomputer can read it directly.

There is a monitor programme which runs in a loop on the microcomputer and successively calls subroutines to perform the major functions required of the controller unit namely

- to calculate the twitch ratio from the data read from the muscle relaxation measurement system (T4, T2 and T1 ratios are possible)

- to calculate the drug commands

- to log data on cassette tape

- to read the keyboard and present data to the display.

The motorized syringe unit or Drug Pump is designed to administer doses as boli. It uses disposable syringes and is typically connected to the drip feed in the patient's forearm via a 3 way tap valve through which the syringe can be primed. The anaesthetist can break the control loop by switching the pump OFF LINE. He can then administer the drug himself from the doses calculated by the computer and presented to the display.

The Drug Pump, Nerve Stimulator and transducer are all electrically isolated from the patient for safety.

The control system also logs on cassette tape for the duration of control the twitch response, the drug commands and certain parameters calculated by the control algorithm. This data is used off-line to model the patient's response to the drug and to check the efficacy of the control system.

2.2 The muscle relaxation measurement system

The Nerve Stimulator and Twitch Measurement System were adapted from an existing unit for the purpose of this research (Bradlow et al., 1984).

2.2.1 The Nerve Stimulator

The stimulator has three modes of operation selectable from the front panel:

- Train-of-four pulse
- Single Pulse
- Tetanic pulse

In the Train-of-four mode, a sequence of 4 voltage pulses spaced 0,5 seconds apart is repeated every 10 seconds. At the stimulator output, the pulses are approximately rectangular and have a duration of between 140 and 190 μS depending on the impedance between the electrodes and the current delivered. In all cases the duration of the stimulation pulse is less than the 200 μS limit necessary to avoid repetitive nerve firing (Ali and Saverese, 1976; Viby-Mogensen, 1982).

The "Single Pulse" mode causes stimuli to be output at Ø,3 Hz. These single pulses are to facilitate setting the supramaximal level of stimulation.

In the Tetanic mode, the patient is stimulated continuously at 5Ø Hz.

The amplitude of the output pulses is variable between Ø and 23ØV by means of a front panel control. This is used to set the supramaximal level of stimulation by raising the output voltage until the twitch response reaches its maximum strength then increasing it 1Ø% beyond this point (Ali et al., 1971).

A simplified circuit diagram of the Nerve Stimulator is shown in figure 2.3. (The complete circuit is given in Appendix E). It has three functional sections.

1. THE POWER SUPPLY. This provides two DC output voltages - 12V to power the CMOS logic circuits for the pulse generation and 25ØV used in the pulse shaping section to form the stimulation pulse. The power transformer is specially wound to ensure that leakage from its secondary winding to the primary and to the core is minimal. The Nerve Stimulator unit is also isolated with respect to earth and all voltage power rails float with respect to earth. The leakage currents for this unit are discussed in section 2.6.

2. THE PULSE GENERATOR. This uses a 5Ø Hz signal derived from the power supply transformer for its timing. The Schmitt triggers square the wave and the series of counters which follow produce a 2 Hz wave and a Ø,1 Hz wave of mark/space ratio 2:1Ø. The logical AND of these waves produces 4 pulses of 1ØØ mSec duration each at Ø,5 second intervals every 1Ø seconds. This signal has three uses - it drives the stimulation pulse shaping circuit described below, it lights a L.E.D. on the front panel which gives a visual indication when stimulation is

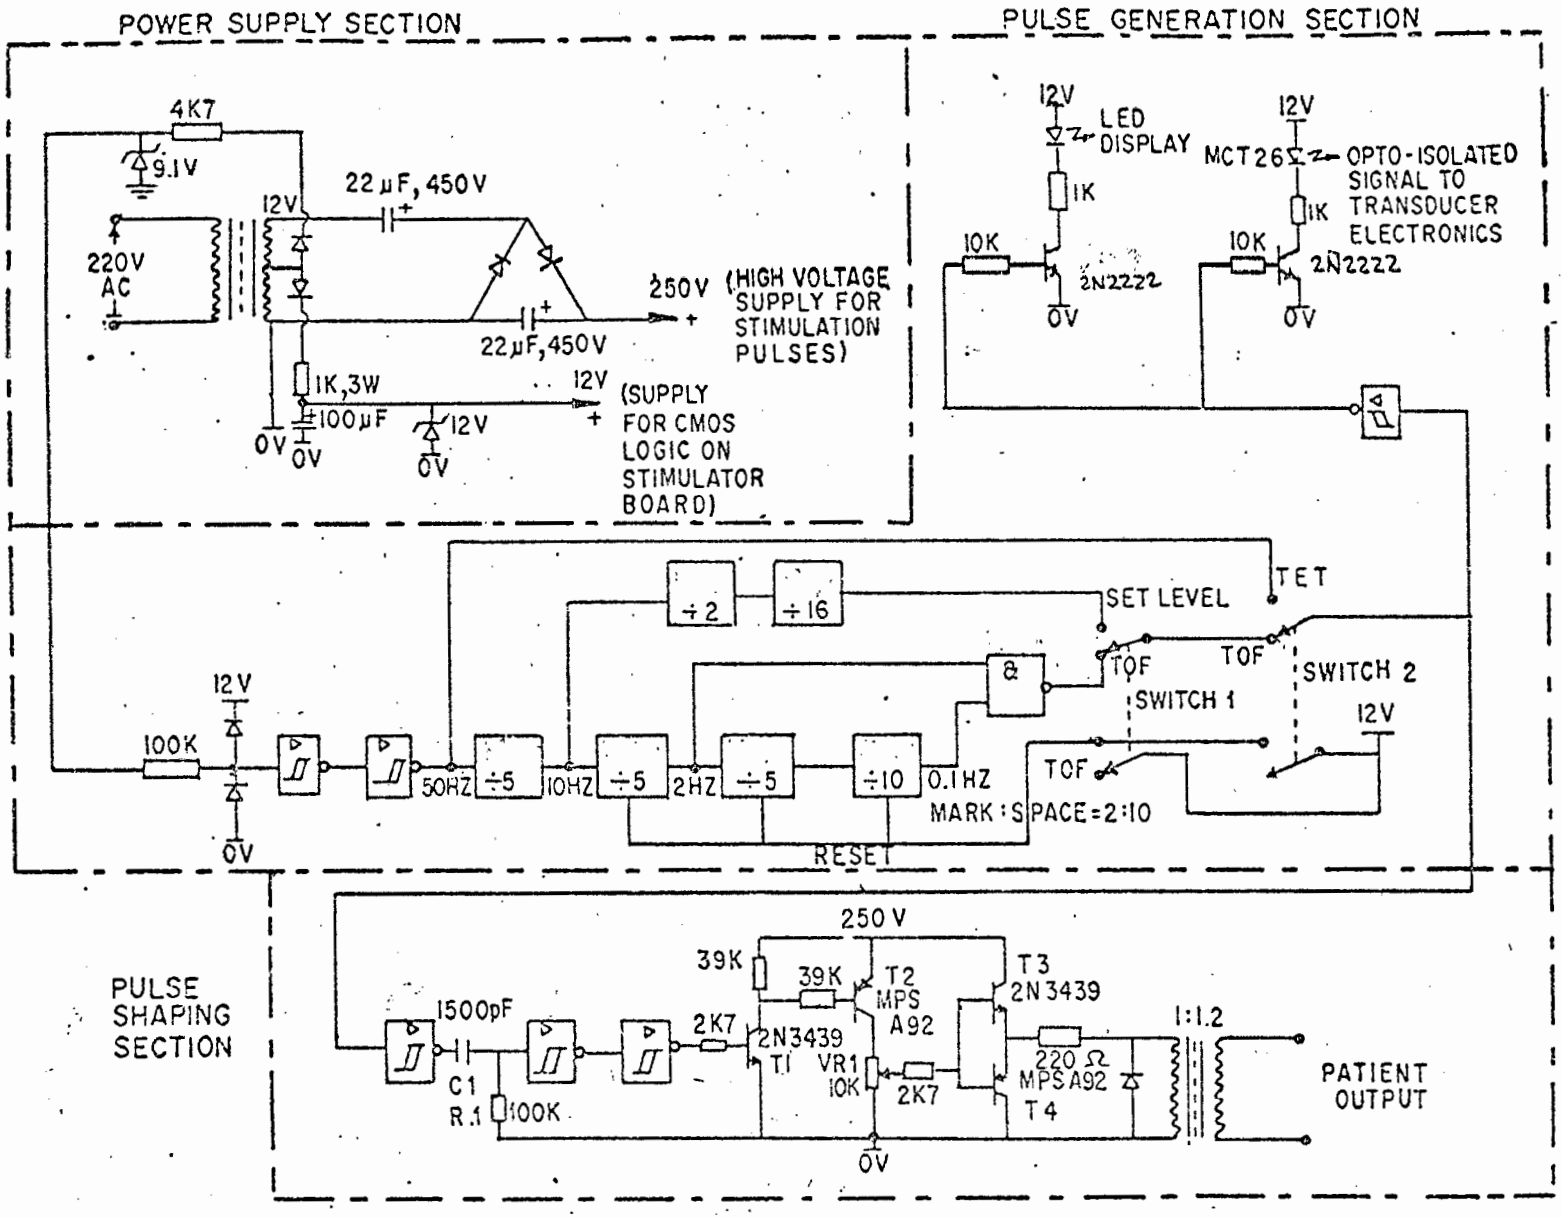

Figure 2.3 : Simplified circuit diagram of the Nerve Stimulator.

occuring and it is passed to the Twitch Measurement System via an opto-isolator where it is used to select only valid stimulation responses.

3. THE PULSE SHAPING CIRCUIT. The time constant of the differentiator consisting of capacitor C1 and resistor R1 near the input of this section limits the duration of the logic pulse to below 200 μS. The pulse is squared by the Schmitt triggers thereafter and the remainder of the circuit converts the logic pulse to a stimulation pulse whose amplitude can be set by VR1 to between 0 and 230 volts.

The output of the Nerve Stimulator is passed through a final isolation transformer and from there to the wrist electrodes which were of the silver/silver chloride paediatric ECG type. The transformer serves three purposes:

1. It boosts the maximum pulse height from 230V to approximately 275V. It has a turn ratio of 1:1,2.

2. It provides another level of isolation between the patient and the mains or earth.

3. It ensures that no DC voltage can be passed to the patient in the event of a failure in the pulse shaping circuit.

2.2.2 The Twitch Measurement System

The first stage of the Twitch Measurement System is an Instrumentation Amplifier, the "strain gauge amplifier", which is driven by the transducer. It has 6 discrete gains settings from 100 to 5000. Its output is passed to the

signal processing circuit described below and also to a microammeter on the front panel which is calibrated in grams force. This assists the anaesthetist in setting a suitable preload resting tension (Viby Mogensen, 1982) when positioning the transducer. Once this has been done, the microammeter is nulled by means of a potentiometer on the front panel which rebalances the amplifier.

A simplified circuit diagram of the signal processing circuit is shown in figure 2.4. (The complete circuit is given in Appendix E). It captures the peak value of each of the four twitch responses, digitises it and outputs each to the microcomputer together with the pulse number. It ensures that the output of the strain gauge amplifier is only digitised when a response to a stimulation pulse is expected. The synchronization is achieved by means of a 1ØØ mSec pulse the beginning of which is coincident with that of each stimulation pulse. This 1ØØ mSec pulse is produced by the Nerve Stimulator and opto-coupled into the signal processing circuit as shown in the figure.

The strain gauge amplifier output is passed through a DC level remover in order to remove any offset due to drift or inaccurate balancing of the amplification circuit. Thereafter its maximum value is obtained in a peak hold and discharge circuit. Peak holding takes place when transistor TR1 is off and this only occurs for 1ØØ mSec after the start of each stimulation pulse as stated above. The peak hold circuit is thereby prevented from storing spurious signals.

The A/D conversion of the output of the peak hold circuit is delayed for 1ØØ mSec after the stimulation pulse. This delay ensures that the maximum stimulation response is achieved. (The average time to maximum twitch strength for 1Ø patients selected at random from those who underwent clinical trials was found to be 77 $\pm$ 1Ø mSec. In no case did it exceed 1ØØ mSec). The START CONVERSION pulse of the A/D is derived from the trailing edge of the 1ØØ mSec pulses from the Nerve Stimulator.

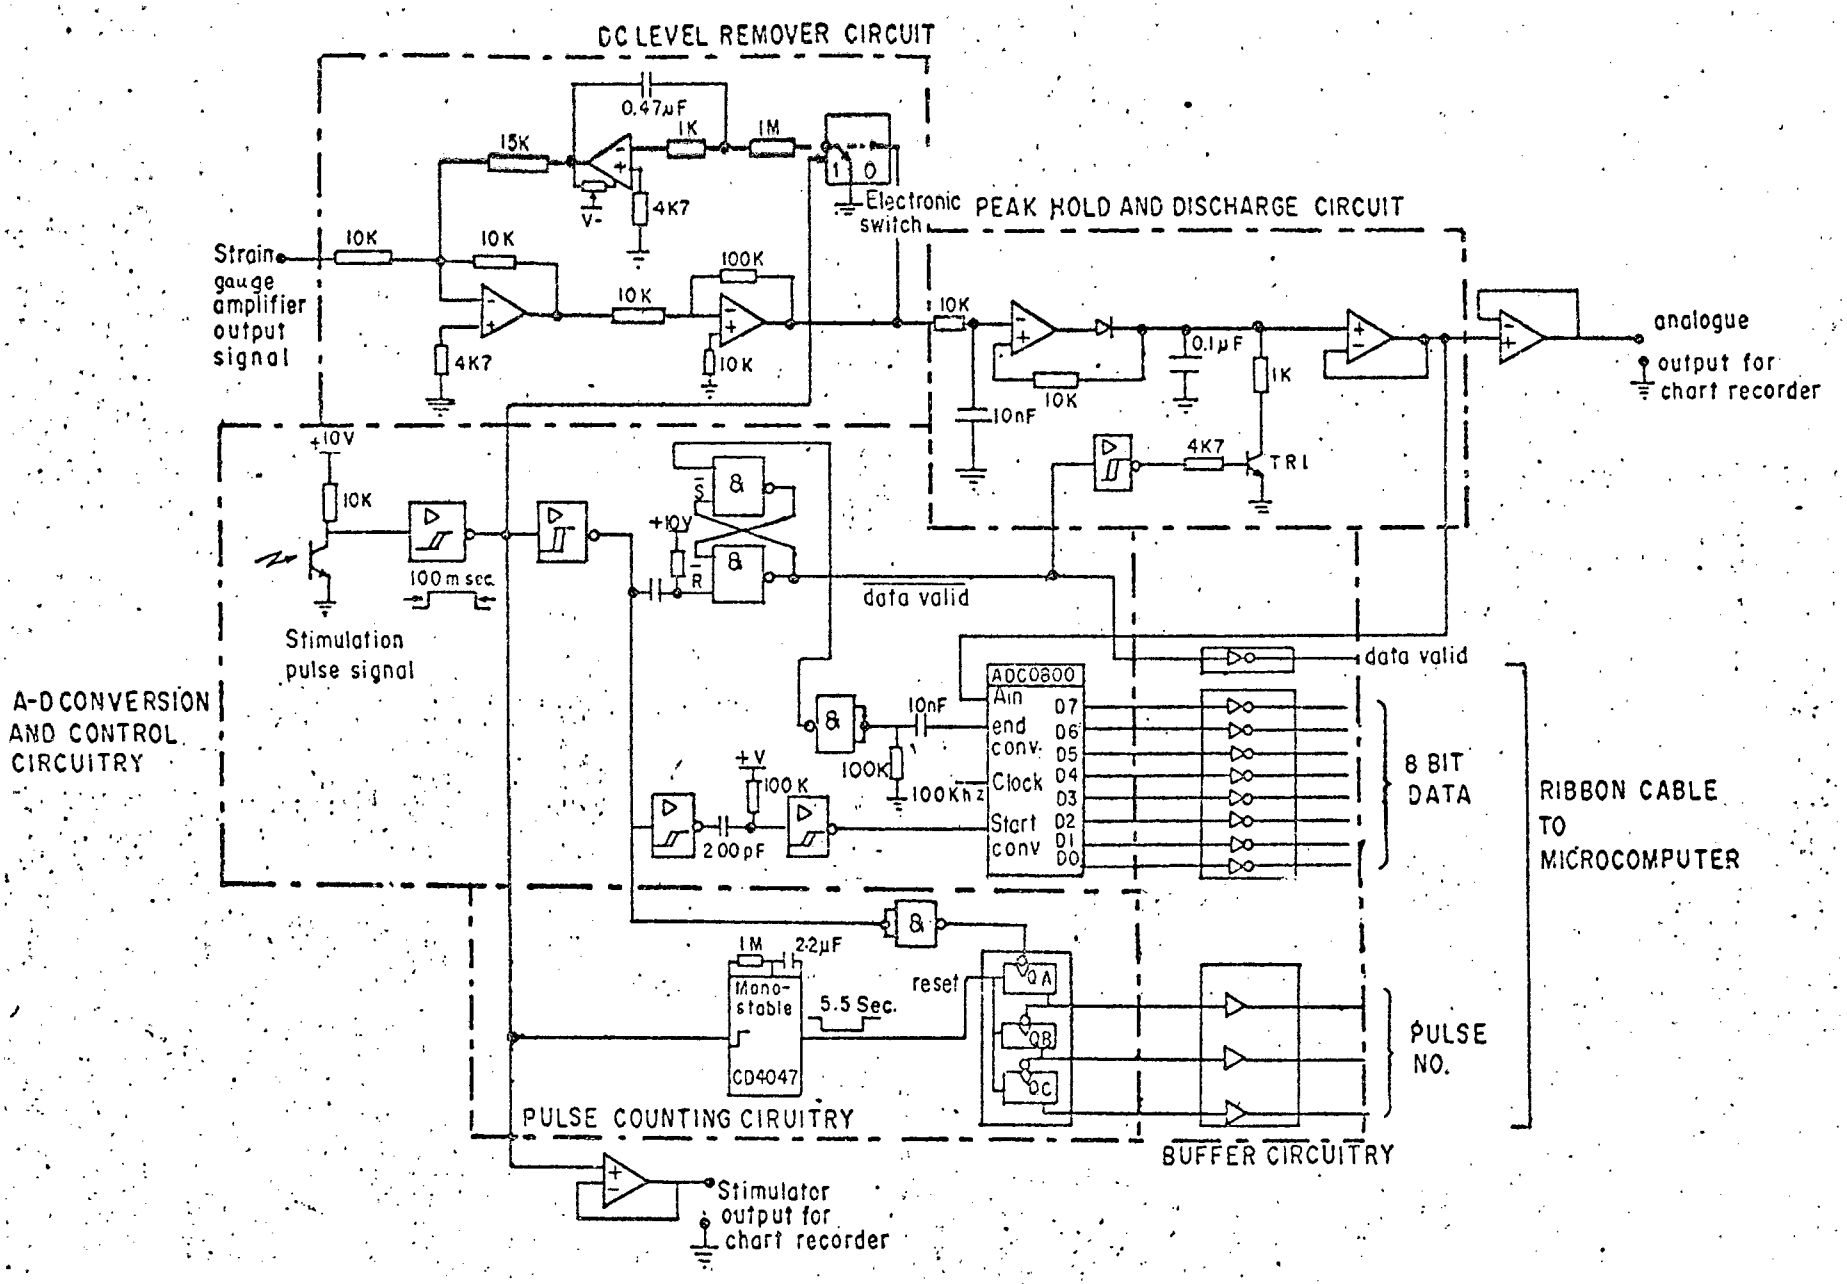

Figure 2.4 : Simplified circuit diagram of the signal processing circuit for the strain gauge amplifier response.

The stimulator pulses are also used to reset an R-S flip flop given by two cross connected NAND gates which indicates in the RESET state that the data is invalid. When the A/D has completed its conversion cycle (which takes 400 μS with a 100 kHz clock), its END OF CONVERSION pulse will set this flip flop which indicates to the computer that the data at the A/D output is valid.

The pulse number, which ranges from 1 to 4, is read by the computer from a binary counter where the stimulator pulses activate the CLOCK input. Each stimulator pulse causes the counter to be incremented by 1 on its trailing edge. The RESET pins of the counter chip are controlled by the output of a retriggerable monostable. This monostable is triggered on the leading edge of the first stimulator pulse in the train-of-four and remains low for approximately 5,5 seconds. When the RESET inputs are low, the counter is active to count, and increments after each of the four stimulator pulses. After 5,5 seconds the monostable goes high causing the counter to reset to zero. In this way the counter remains synchronised to the Nerve Stimulator because it is only reset between successive T4 trains.

The A/D output, pulse number and DATA VALID signal are buffered to convert the CMOS signal levels to TTL levels for compatibility with the microcomputer and to enable these digital outputs of the signal processing system to be interfaced to the microcomputer's bus.

2.2.3 The Chart Recorder

As an aid to the research, provision was made for the analogue output of the signal processing circuit of the Twitch Measurement System to be recorded on a chart recorder. (A MFE recorder, model M22C was used). This was used to display the patient's response continuously thus enabling any anomalies such as

spurious noise signals to be detected. Before control was begun it was used to verify that the stimulus was supramaximal and to check for adequate recovery from a previously administered relaxant (for those patients who were relaxed during intubation).

2.2.4 The Transducer

The transducer converts the force of contraction of the Adductor Policis muscle to stimulation of the Ulnar nerve into a resistance imbalance of a Wheatstone bridge. This forms the input of the Twitch Measurement System (section 2.2.2).

The transducer consists of a perspex plunger mounted in a precision linear bearing fitted inside a perspex housing. The thumb rests on a stainless steel cradle on one end of the plunger. The opposite end of the plunger applies the load induced by the thumb onto a cantilever on which is fitted four strain gauges - two active and two dummy - connected in a bridge in the usual way. The system has a working range of Ø - 1ØØN which is adequate for the intended purpose. The plunger, because of its perspex shaft and housing, provides electrical isolation between the patient and the electronic system to which the transducer is connected.

The transducer can either be mounted on an armboard to which the patient's arm and hand are strapped, as shown in figure 2.5 or it can be strapped into the palm of the hand as shown in figure 2.6. Both methods were employed in this research with the use of the armboard being generally preferred as it was less susceptible to interference from the patient's pulse. In terms of convienience of setting up the transducer there is little to choose between the two methods, both being relatively simple to use. A disadvantage found with the armboard

Figure 2.5 : Transducer unit mounted on the armboard.

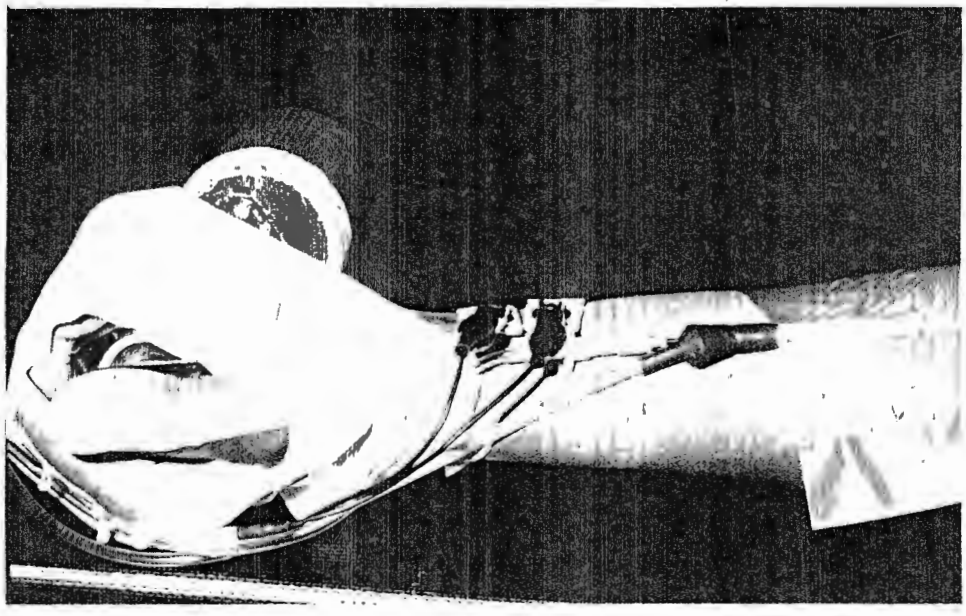

Figure 2.6 : Transducer unit mounted in the patient's hand.

arrangement was that on occasion it did pick up mechanically induced noises from the activity of the surgeons. This was however a relatively minor problem.

2.3 The Microcomputer

The microcomputer system was based on a S64 South African Standard Bus back plane (SABUS Committee, 1981). The various components of the computer system were implemented as printed circuit boards or wire wrap cards and slotted into the backplane. This system was chosen largely owing to its flexibility but also as it is low cost and commonly used in the Electrical Engineering laboratories at the university. Thus interchange of circuit boards for diagnostic purposes is a simple matter. A block diagram of the microcomputer system is shown in figure 2.7. Six boards are connected to the SABUS. These are discussed below.

Board 1. CPU card

The CPU card was designed and developed in the Electrical Engineering Department of the university (Sherlock, 1981). It allows the full co-processing features of the 8088 CPU and 8087 Numeric Data Processor (NDP).

Board 2. Keyboard controller

This card was also designed locally (Smith, 1979 and Putman, 1980) and uses an Intel 8279 keyboard/display interface chip. In addition it has provision for 4K Byte of ROM (Intel 2708 chips) and 2K Byte of RAM (Intel 2114 chips).

The keyboard and display consist of a 16 key keypad and a one-line, 8 digit, 7 segment display (Page, 1980). There is a bell which can be switched in

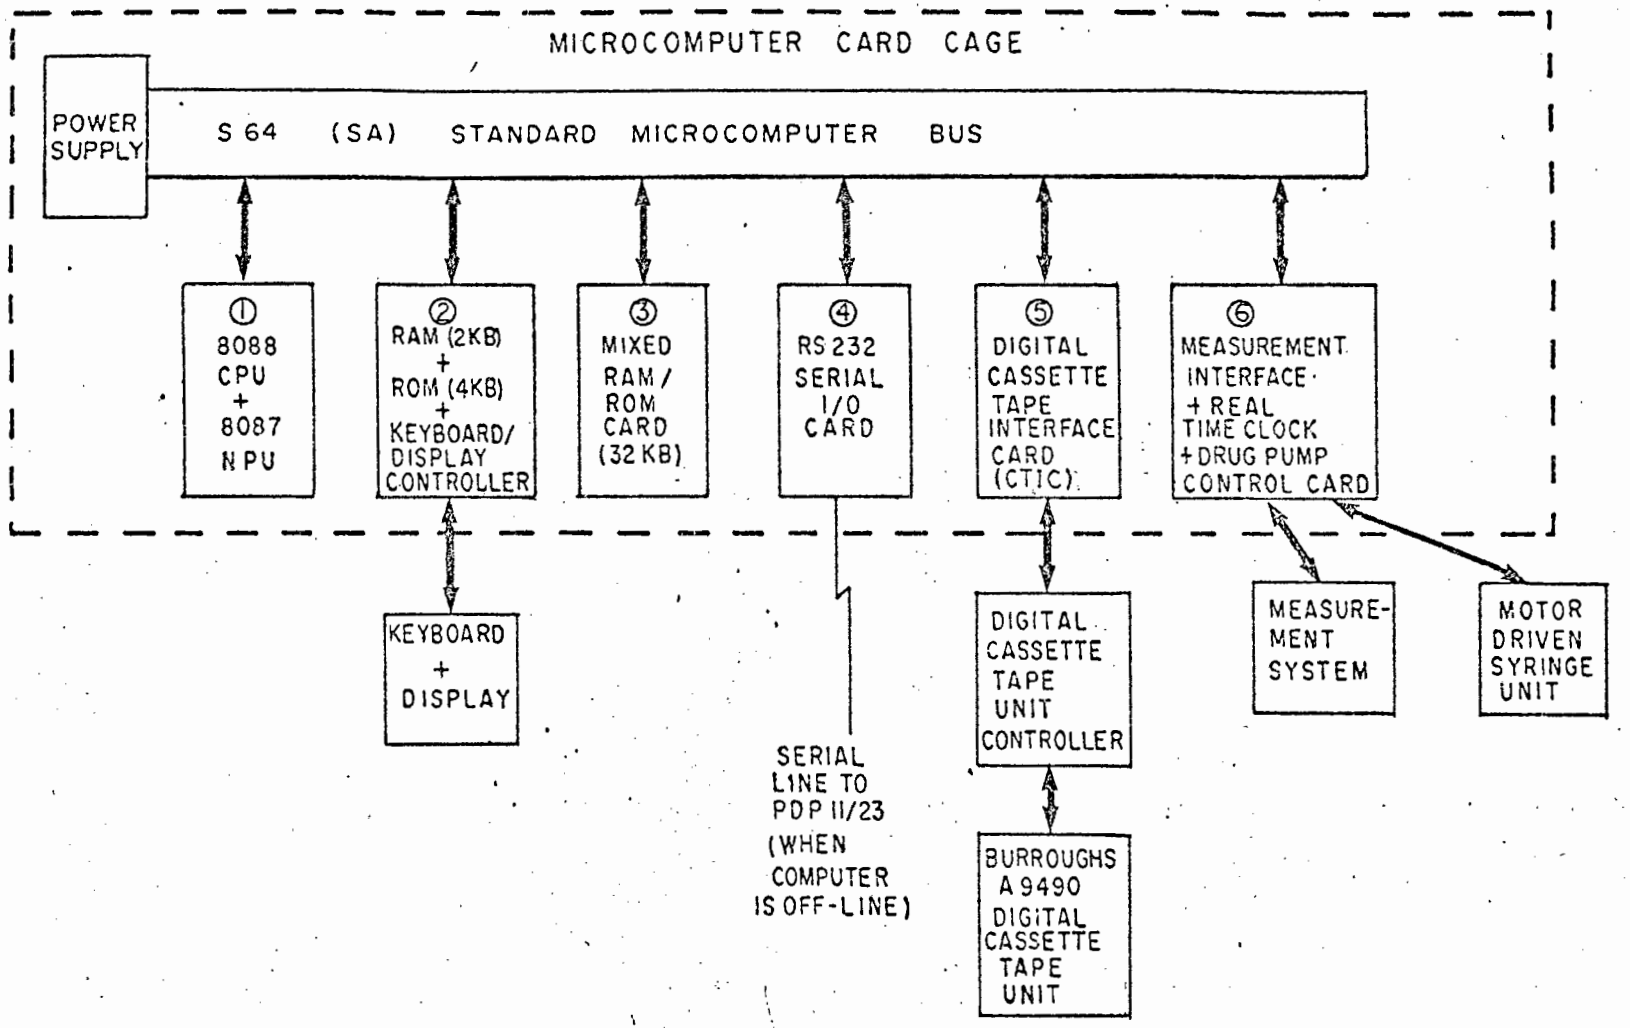

Figure 2.7 : Block diagram of the microcomputer system.

to ring whenever a drug command is output by the controller.

Board 3. Mixed RAM/ROM memory card

This card (ISCOR, 1981) contains 16 sockets that will accommodate either an Intel 2716 EPROM or Hitachi 6116 RAM chip. Both memory chips have a 2K Byte capacity and any mix is possible in the 16 sockets.

Board 4. Serial I/O card

This card was designed locally (Day, 1981) and employs an Intel 8251 USART to achieve parallel to serial conversion and vice versa. This card is only required when the computer is used "off-line" to play back recorded data.

Board 5. Digital Cassette Tape Interface Card (CTIC)

This card, (Strumpher, 1982) is implemented as a wire wrap circuit board and together with the "Digital Cassette Tape Unit Controller" (Eva, 1980), allows a digital cassette tape unit to be used as a standard peripheral in a SABUS microcomputer system.

The CTIC incorporates an Intel 8085 microprocessor in its design. This intelligence allows it to deal autonomously with tape operations and thus it does not load the system CPU (the Intel 8088). It also provides sophisticated error checking of tape read and write operations.

A block diagram of the Interface Card is shown in figure 2.8. The system CPU initiates tape operations by writing to the CTIC Control Register and confirms the results of these operations by reading the CTIC Status Register. Data transfer to or from the CTIC is done via a 2K Byte RAM memory buffer contained on the CTIC. This RAM is accessible to both the System Bus and the

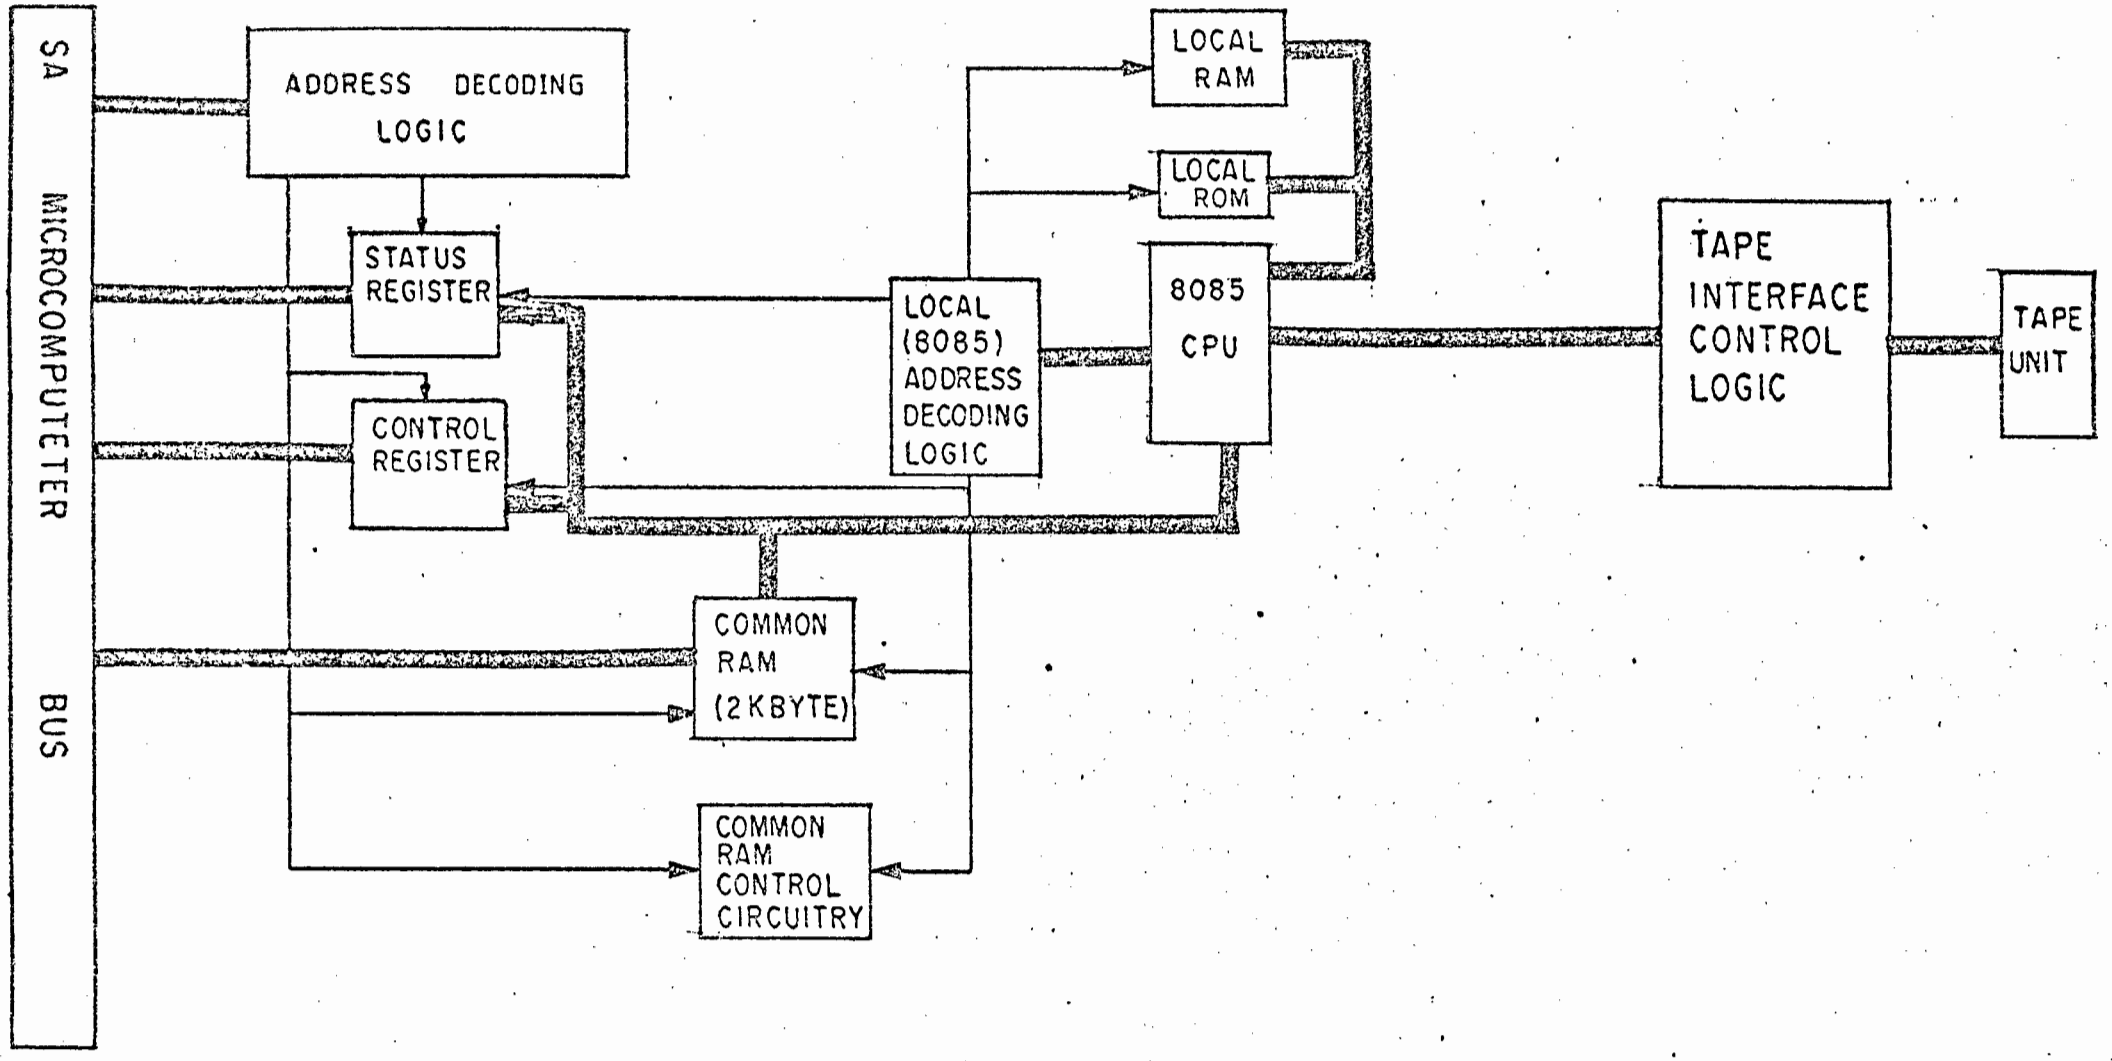

Figure 2.8 : Block diagram of the Cassette Tape Interface Card (CTIC).

local CPU on the CTIC. Thus any data transfers to or from the cassette tape do not delay the system CPU longer than normal memory transfer despite the fact that the actual tape operations may be of seconds duration.

The particular cassette tape unit used (a Burroughs A9490) is highly desirable for logging data in harsh electrical environments owing to its reliability which is achieved by recording clock and data information concurrently on the two tracks of the tape. It was accordingly chosen for this research as the use of diathermy in the operating theatre can result in severe electrical noise.

Board 6. Microcomputer interface to the Twitch Measurement System and Drug Pump

This card is the only one that had to be designed specifically for this research. It interfaces the microcomputer to the Twitch Measurement System (section 2.2.2) and the Drug Pump. It also provides a simple real-time clock circuit (timing derived from the 50 Hz mains supply). This enables the computer to read the real-time in seconds since control was initiated. A block diagram of the card is shown in figure 2.9.

The interface to the Twitch Measurement System (Page, 1980) consists merely of a set of tri-state buffers which, when selected by the card's address decoding logic, allows the computer system bus to read the digital data output from this unit. In all 12 lines are read and these are derived from the signal processing circuit of figure 2.4. The interface also contains a 5 Hz clock signal which is read by the system CPU and facilitates error checking of the measurement hardware by allowing the computer to place a time window around the valid data pulses.

The interface to the Drug Pump uses an Intel 8254 Interval Timer chip, the output of which determines whether the pump motor is on or off. The duration

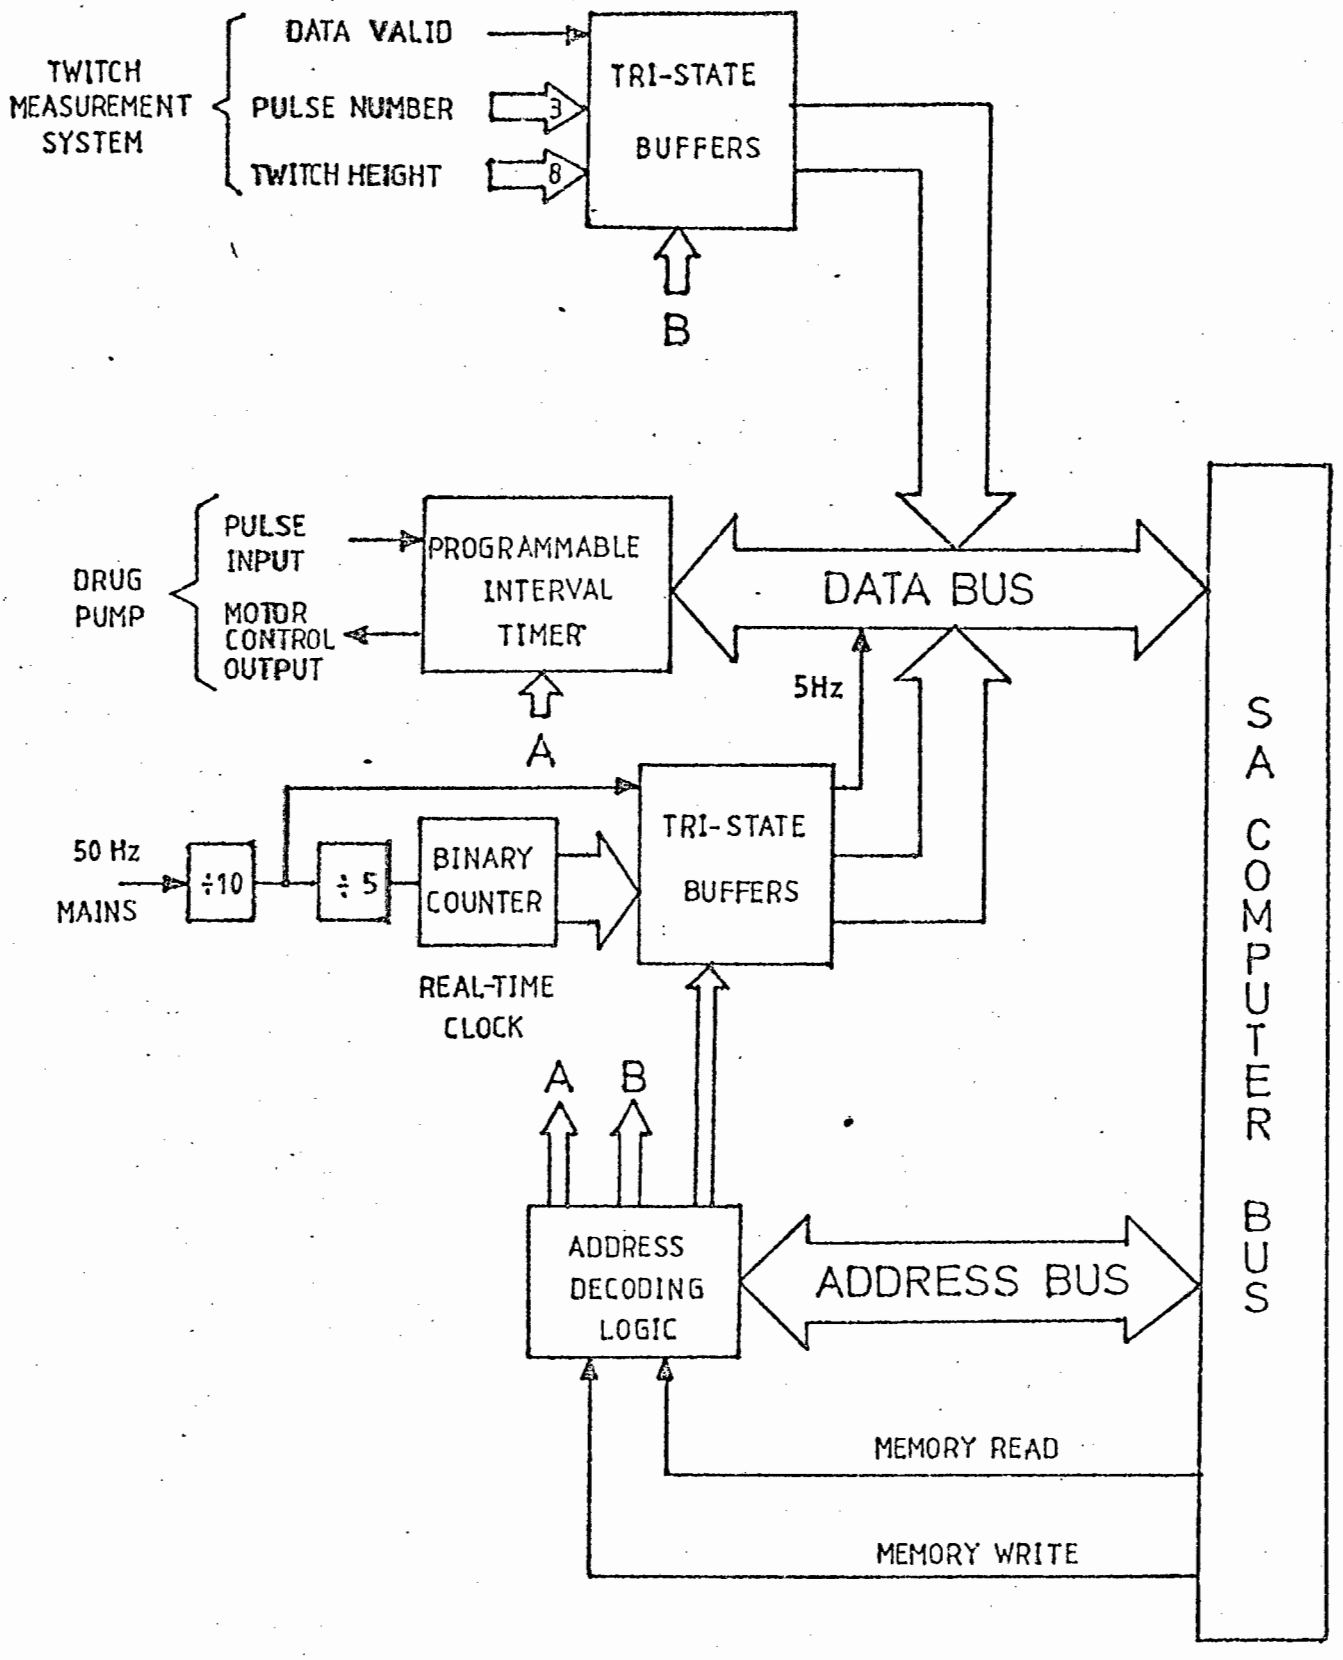

Figure 2.9 : Block diagram of the microcomputer interface to the Twitch Measurement System and Drug Pump.

for which the pump is to remain on is written into the 8254 by the computer. Thereafter pulses from the pump unit, each of which indicates the delivery of Ø,Ø36 millilitre of drug, decrement the counter until it reaches zero and switches the pump off.

The circuit diagram of this card is given in Appendix F.

2.4 The Drug Pump

Control is achieved by regulating the rate at which muscle relaxant is administered. This can be achieved either by controlling the rate of drug infusion or alternatively by administering bolus injections at regular intervals (thereby achieving an approximation to impulse inputs). The latter technique was chosen for this work for four reasons :

1. For a slow acting drug like curare, once the level of relaxation has reached the set-point, the rate of drug infusion required to maintain that set-point is very small (on average it was found to be as low as Ø,ØØ2 mg/Kg/min). However a maximum rate of about Ø,9 mg/Kg/min is desirable in the induction period in order not to prolong the interval unnecessarily. This wide range is hard to achieve accurately with simple equipment.

2. The drift of the measurement away from set-point between boli is acceptably small provided the boli are administered frequently enough.

3. The actual administration of the drug boli can be achieved using the very simple motor driven apparatus described below which uses the same syringes as would normally be used clinically to administer the drug.

4. As bolus injection is the most direct clinical method of administering muscle relaxants, it is likely that a bolus injection system would achieve more ready clinical acceptance than an infusion one. This is reinforced by the manner in which the drug administration has been implemented in the control system.

The Drug Pump delivers bolus doses to the patient as commanded by the controller. A photograph of the mechanical arrangement is shown in figure 2.1Ø.

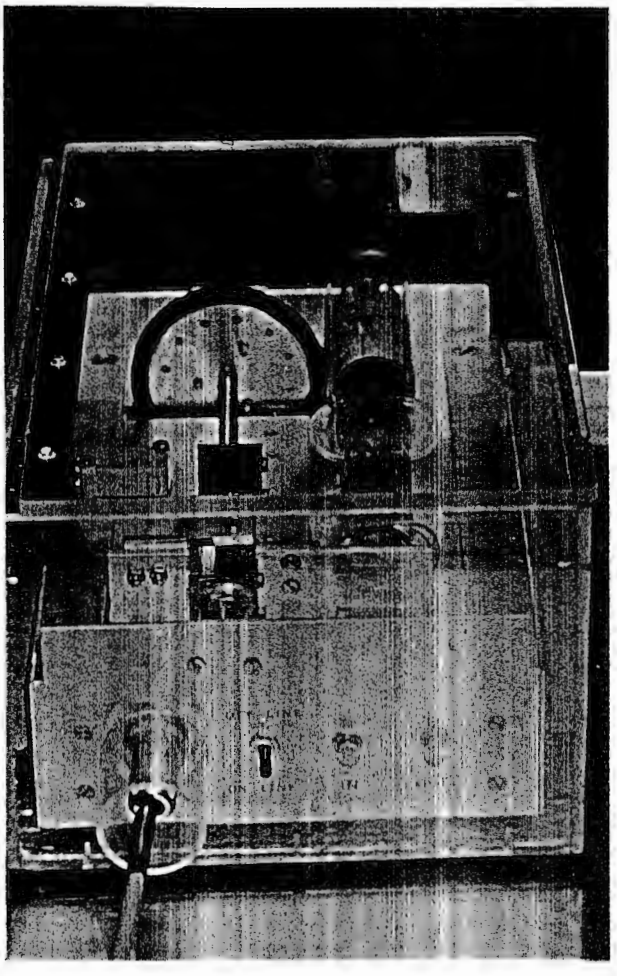

Figure 2.1Ø : The Drug Pump.

The unit is loaded with a disposable 2Ø ml syringe the plunger of which is depressed by a worm drive. A flywheel with holes spaced 3Ø° apart around the

circumference is fixed to the worm screw and driven by a motor. A lamp/photocell combination measures the rotation of the screw by counting the number of holes which pass it. The sequence of pulses so formed is the feedback signal to the computer. Thirty degrees of rotation corresponds to Ø,Ø36 ml of fluid for the particular syringe used but syringes of slightly different bore may also be used and the necessary conversion factor entered into the computer from the keyboard.

The computer decides the on/off status of the motor. Rapid decelleration is achieved by reversing the polarity of the motor supply with a DPCO relay. This is energised immediately that the "motor off" command is received for a time period which can be set from Ø to 5ØØ mSec. Thereafter the motor drive is removed. The timer was tuned for negligible run-on.

The Drug Pump is battery driven and opto-isolated from the computer circuitry to ensure patient safety. It is totally enclosed in a clear perspex case with no accessible conducting parts. The circuit diagrams of the pump and opto-isolated interface are given in Appendix F.

Use of the pump is straight forward. The user enters two settings at the keyboard before commencing control. One informs the computer of the drug concentration in the syringe, the other is the number of pulses which corresponds to 1 cm travel of the worm carriage. The second of these settings was included only because the mechanical arrangement of the pump was not finalized at the time the software was written. The pump is switched OFF-LINE whereupon the position of the worm carriage is controlled by push buttons on the pump itself. The anaesthetist loads the syringe as if he were going to administer the drug himself. He then attaches it to an intravenous infusion line, typically through a 3 way tap valve with which the system may be cleared of air bubbles, inserts it into the Drug Pump and switches it ON LINE. To

change a syringe, the unit is merely switched OFF LINE and the proceedure repeated. The syringe can be changed in the course of administering a bolus since when switched to OFF LINE the pulse transmitting circuit from the pump is disabled. When switched back to ON LINE, the remainder of the dose will be administered.

The pump is capable of injecting fluid at a rate of 1 ml/sec. This corresponds to 1 mg/sec for the drug concentration used in the clinical trials. Thus the largest dose likely to be commanded (about 15mg) will be administered in a time one order of magnitude less than the average of the smaller time constant of dTC (Bradlow et al.,1983). The doses are therefore acceptable approximations to impulses.

2.5 Twitch Response Simulator

Whilst developing the hardware and software of the control system it was at times necessary to check the operation of the intermediate stages whilst the equipment was in the laboratory. As no patient was then available, it was necessary to simulate the human twitch response in a way which is synchronised with the output of the Nerve Stimulator.

A twitch response simulator was constructed in the form of an electro-mechanical actuator device to which the transducer could be attached. This was driven by an electronic circuit which was timed by the Nerve Stimulator and produced a train-of-four voltage pulses in synchronism with it. The rise time of these pulses could be varied from 10 to approximately 100 mSec and their height set independantly. In this way it was possible to simulate various valid and invalid train patterns. The circuit diagram of the unit is given in Appendix G.

2.6 Electrical saftey tests

A necessary consideration for any electronic equipment used in medical applications is patient safety. All the elements of the control system were specifically designed to avoid micro and macro shock hazard and thus the necessary measurements to determine current leakage levels (Olson, 1978 and I.E.C., 1977) were carried out. The equipment was found to be safe for use. Details of the tests performed and their results are given in Appendix H.

2.7 The significance of noise in the measurement of the twitch ratio

In 6 of the 42 clinical trials performed, the controller failed because the measurement signal-to-noise ratio was too small. In 1 of these 6 cases, this failure was due to interference from the patient's pulse which was corrected by use of the armboard described in section 2.2.4. For the other 5 cases, failure was due to there being a certain degree of noise caused by the inevitable electrical and mechanical interference. Thus 2 of the 19 T4 trials failed and 3 of the 4 T2 trials but none of the T1 trials.

The plot of the T2 and T4 ratios against the T1 ratio taken from the clinical trial of a typical patient (no. 32) and shown in figure 2.11 helps to explain the reason for these failures.

Given a ratio of say 30%, it is possible to deduce the heights of twitches 1,2 and 4 for this patient from figure 2.11 and to compare them as done in figure 2.12. If one assumes that the noise signal is independant of the size of the measurement, the signal-to-noise ratio is dependant on the signal strength. S/N will be best for the T1 ratio because both numerator and denominator pulses

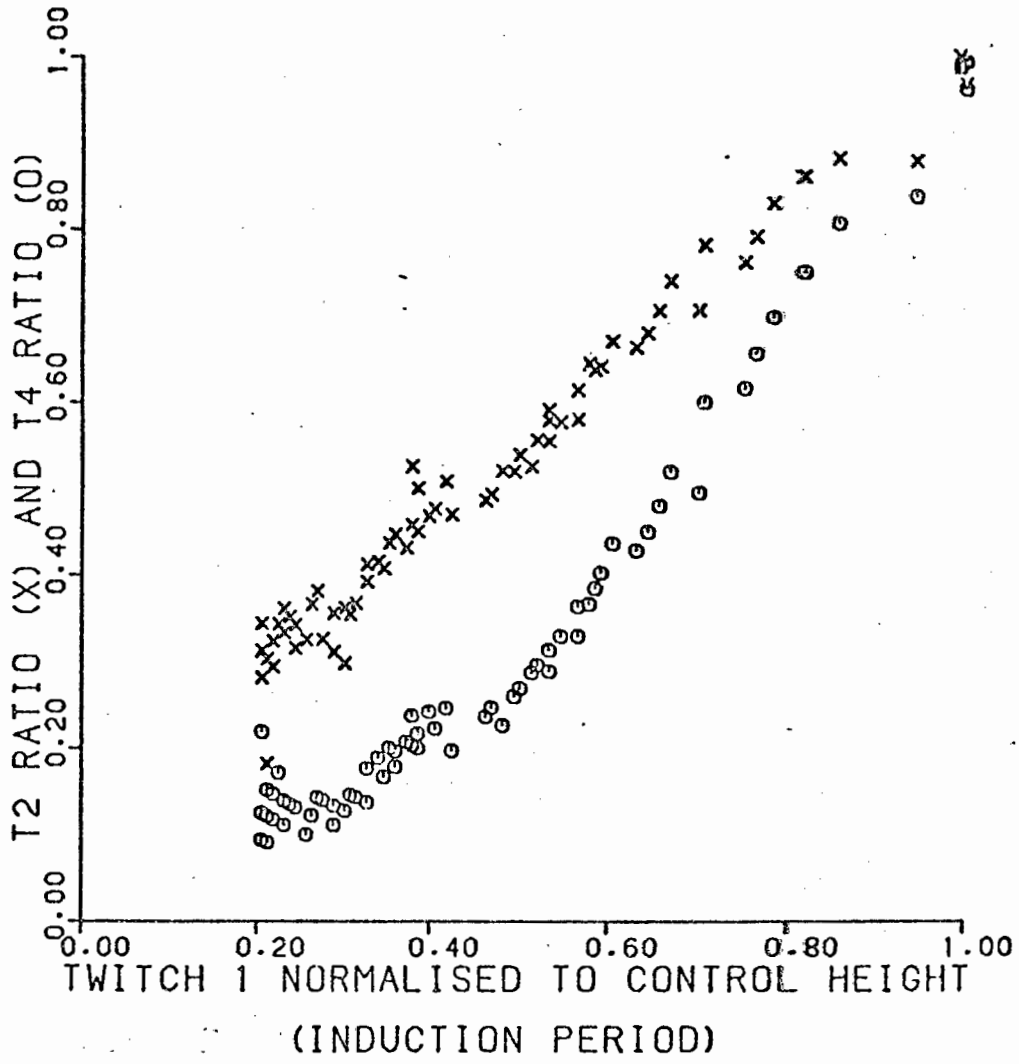

Figure 2.11 : Plot of the T2 and T4 ratios against the T1 ratio taken from the clinical trial of patient 32.

have the greatest magnitude. The T2 ratio will have the worst S/N because twitch 1 is depressed most. This is the reason that control of the T2 ratio was unsuccessful. S/N of the T4 ratio is somewhat better than for T2 but inferior to T1. This is verified by a comparison of the average noise variances of the estimated patient models derived from the T4 and T1 clinical trials in Table 3.2. They are $8,48.10^{-4}$ and $2,46.10^{-4}$ respectively. Finally, the signal-to-noise ratio decreases as the twitch ratio becomes smaller in all three cases.

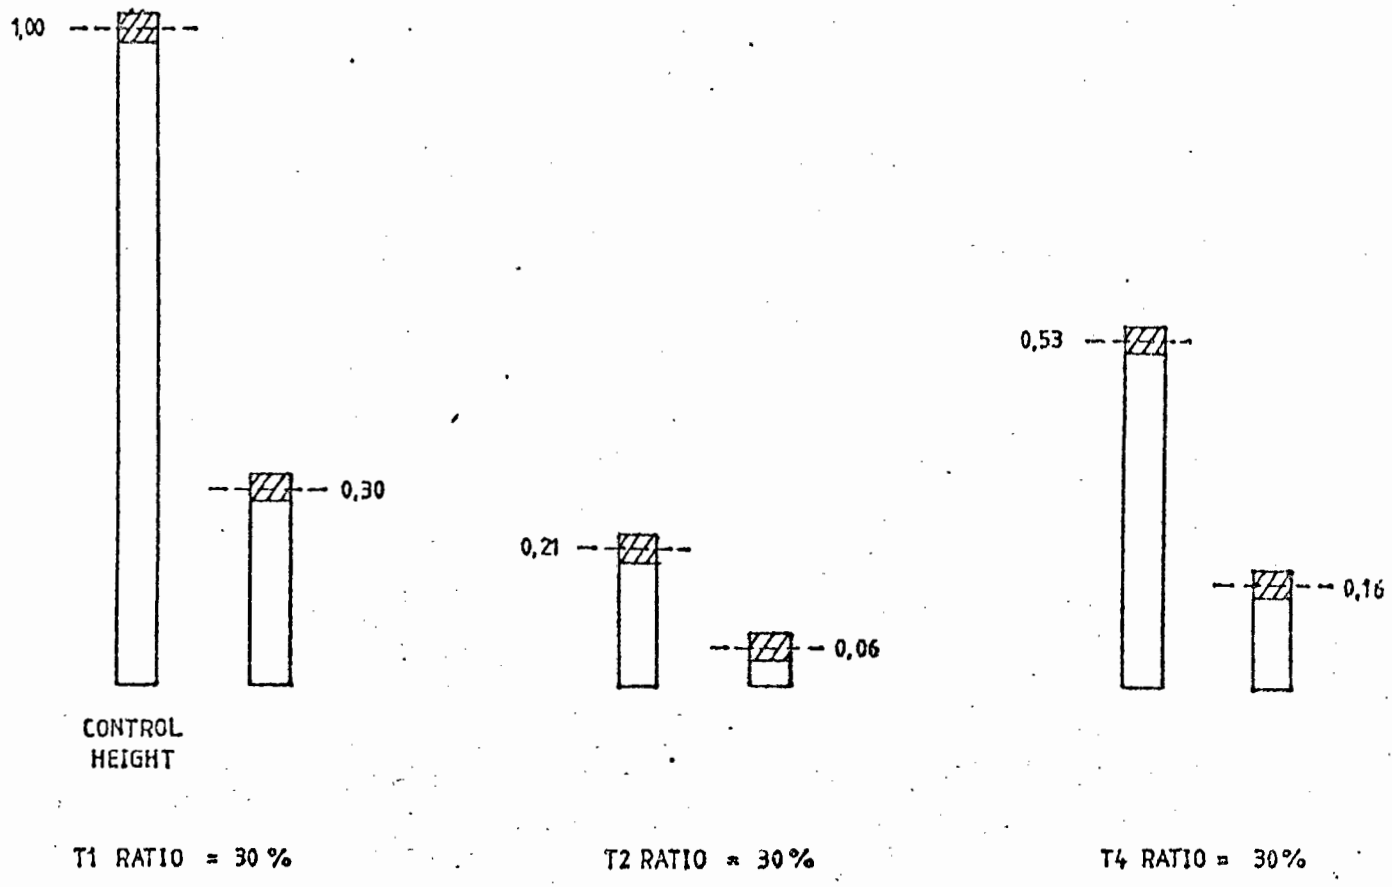

Figure 2.12 : Comparison of twitch amplitudes at T1, T2 and T4 ratios of 30% for patient 32. The standard deviation of the noise component was approximately 0,02 for this patient.

2.8 Software of the microcomputer

All the software for the microcomputer was written in the Intel 8086/88 assembly language. Assembly language programming was chosen mainly because at the time this software product was initiated the departmental laboratories were unable to support a high level language compiler that generated code for the 8087 Numerics Coprocessor. In addition use of assembly language results in more efficient code which is desirable in a real time application which makes extensive use of numerical functions.

The software is implemented as a single task running on the native computer system without the use of any operating system. The software was designed to

run without employing interrupts, as the application does not necessitate their use, and this is believed to improve the reliability of the system. The memory requirements of this software are 10K Bytes of EPROM and 22K Bytes of RAM.

The software is controlled by commands entered at the keyboard. These commands are made in response to prompts written to the display. They enable the user to:

1. Set the mode of operation.

   The ON LINE mode is used for real-time muscle relaxation control. The user may select whether or not to log patient and controller data on cassette tape. In the OFF-LINE mode data files are read from the cassette tape and transferred down a serial line to a PDP 11/23 minicomputer for storage on an Industry standard 1600 bpi half inch magnetic tape.

2. Select which of the T1,T2 or T4 ratios is to be controlled.

3. Enter the controller settings or select default values.

4. Interrupt control in order to change the set-point, override a drug command issued by the controller or to change the syringe on the Drug Pump. It is also possible to bypass the controller but continue to log data. This is useful on completion of control to record the recovery phase of the operation.

5. Initiate and terminate the control/logging operation.

Error checking facilities are provided by the software. These pertain to read/write tape operations, the link between the microcomputer and PDP

minicomputer for the transfer of data files, erroneous data arising either from a malfunction in the measurement interface or from an invalid train-of-four fade pattern and the control algorithm. A listing of the error checks which are made is given in Appendix C.

Figure 2.13 is a block diagram which shows the interconnection of the major software modules of the control system. (Each module may contain one or more independant sections of code). The functioning of the software can be deduced from the description given below. Full details are contained in Appendix C.

Modules BUFFER (the cassette tape interface controller),(Strumpher,1982) and PDPLNK (interface to the PDP11 minicomputer) written by Bradlow, 1982 are not discussed. The listings are however included in the appendix as these modules were subsequently modified.

Module DDOSE

This module contains the supervisory programme and a simplified flowchart of its logic is given in figure 2.14. The two major branches pertain to the ON LINE (control/logging) and OFF LINE (playback data files) modes of operation.

In the ON LINE mode the user first selects whether the tape unit is to be used to log data or not. If so the computer prompts for the file name and the user responds with a 6 character header. A sequence of prompts then follows for each of the controller settings to which the user may respond appropriately or select the default for each by pressing the <RETURN> key. The programme next initializes certain constants within the control algorithm and clears arrays (module RESET). After this point the programme enters a loop in which it successively reads and processes the patient's latest twitch response (module DATA), calculates the drug command for that sampling period (module BOLUS4) and updates the accumulated patient dose for display. The data which is written to

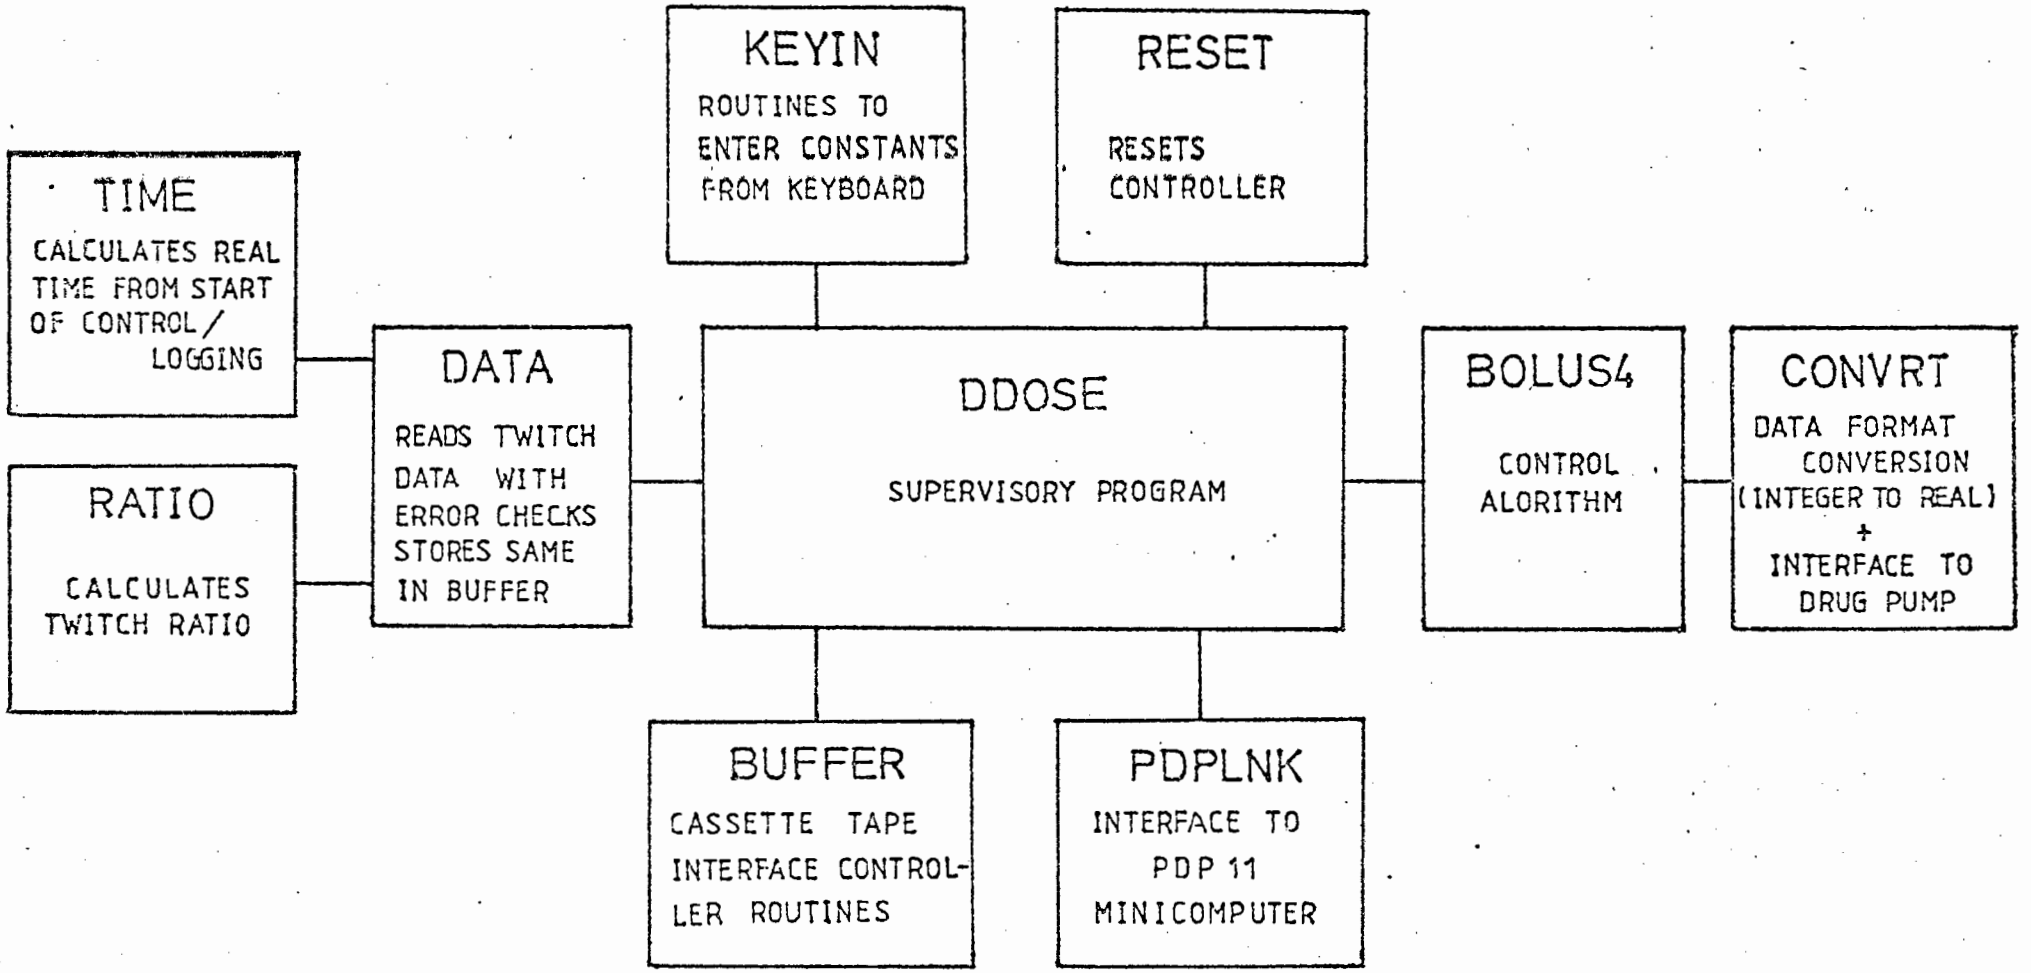

Figure 2.13 : Interconnection of the major software modules of the control system.

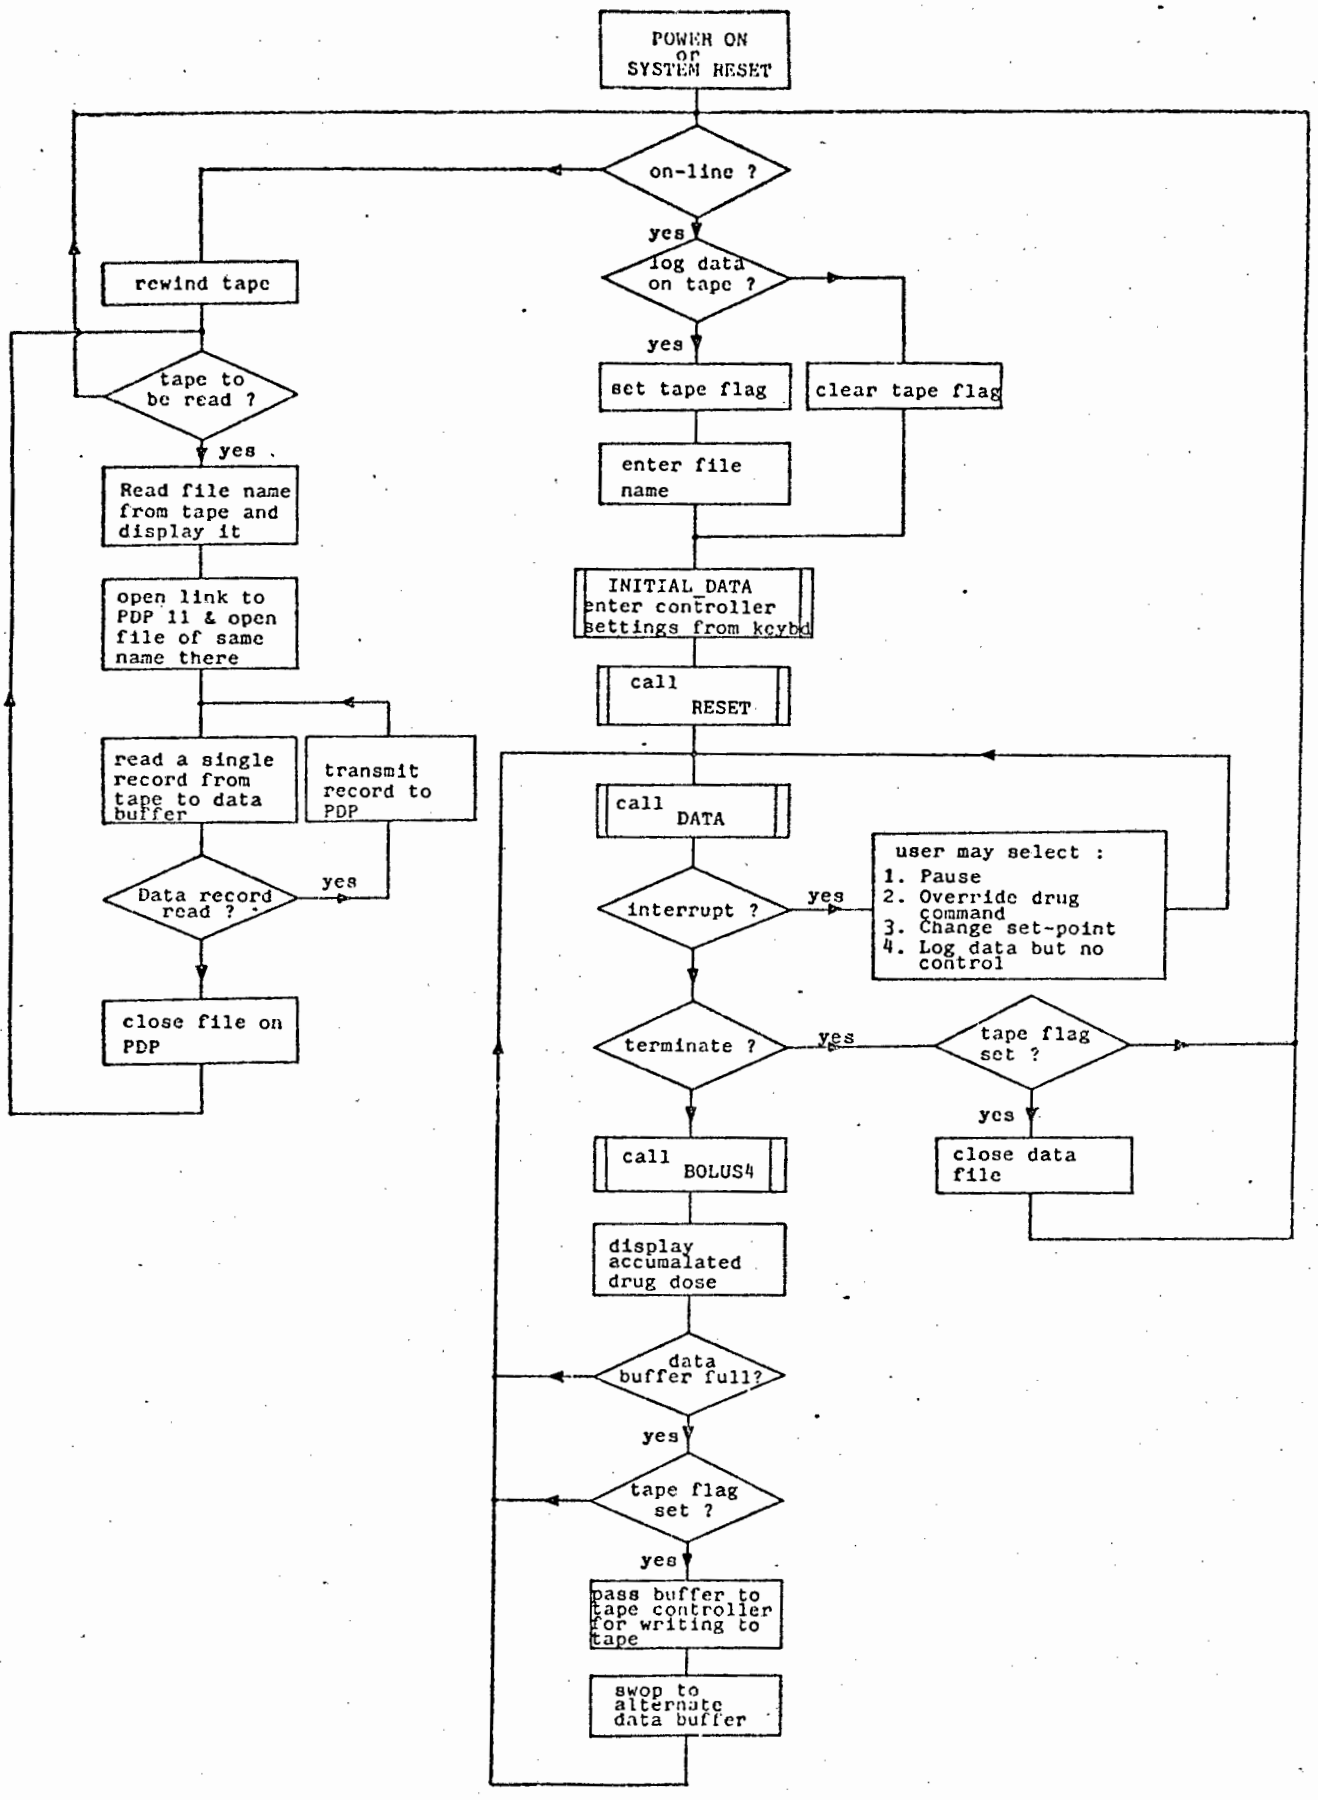

Figure 2.14 : Simplified flowchart of the supervisory programme.

tape is stored alternatively in either of two data buffers. These buffers are 256 Bytes each and take 2 minutes to fill. A full buffer is written to the Cassette Tape Interface Card. The time taken is of the order of a normal series of memory writes i.e. of the order of tens of microseconds. The programme then continues using the "empty" buffer. The Cassette Tape Interface Card, under control of the 8Ø85 CPU, then autonomously writes the data to tape. This is completed long before the next buffer is filled. The loop may be broken temporally by pressing the <PAUSE> key. This enables the user to perform the functions shown in figure 2.14.

To exit from the loop the user presses the <RETURN> key. This returns the programme to its original "POWER ON" or "SYSTEM RESET" state after closing the tape file if the tape is being written to.

In the OFF LINE mode the tape is initially rewound. The programme then enters a loop and one data file is read per pass. The programme halts at the beginning of each pass and the user may exit at this point. Files are transferred in the following way. The file name is read from tape and displayed. A link is then opened to the PDP minicomputer and a file of the same name created there. Each file record is read to the data buffer from tape and transmitted to the PDP. (Alternating buffers are not used.) When a Trailer record is encountered, the file on the PDP is closed and transfer thus completed.

MODULE DATA

This module contains subroutine DATA which is called from the main programme DDOSE. It reads and checks the data received from the interface to the Twitch Measurement System (Section 2.2.2).

It polls the measurement interface and waits for it to indicate that a particular pulse has been received and that the measurement is valid. In this way it reads each of the twitch responses in the train-of-four and their assosiated pulse number. On first entry the real-time clock and sample number counter are reset. The routine then checks that the amplitude of twitch 1 (the largest in the train) is in range and displays the appropriate overflow/underflow message if not. If the twitch 1 data is out of range, the programme delays 3 seconds and returns to the entry point of the subroutine. It continues in this way until the error has been rectified. An overflow message usually indicates that the gain setting of the strain gauge amplifier (section 2.2.2) is too high and an underflow that the patient is not being stimulated. Errors which are detected subsequently by the DATA subroutine are merely signalled by replacing the invalid data with the value FF Hex. When this data is passed to the control algorithm, it responds by rejecting the measured twitch ratio at these sampling periods and replaces them with the most recent valid data.

If T1 control has been selected, the DATA routine calculates the control pulse height by averaging the magnitude of Twitch 1 for the first 3 sampling periods. In the unlikely event that the twitch response is out of range at the second or third sampling period, the invalid data is replaced by the previous valid measurement. The average is then stored and used subsequently to calculate the T1 ratio.

If either T2 or T4 control is selected, the routine checks that the amplitude of the train-of-four response decreases monotonically and signals an error as described above if not.

In all three modes of control the spacing between successive twitch measurements in the train-of-four response (nominally 500 mSec) is checked. A time window of 400 - 700 mSec is set and the data is accepted only if it falls

within this band. The use of diathermy during operations makes this precaution necessary for if electrical interface advances the pulse counter and sets the DATA VALID flip flop within the signal processing circuit of the Twitch Measurement System (section 2.2.2) the DATA subroutine would otherwise interpret this condition as a valid twitch response.

DATA calls subroutine TIME at each sampling period to determine the real-time from the start of control in seconds and in sampling periods.

Subroutine RATIO is called to calculate the required twitch ratio from measurements made.

Finally the DATA routine checks whether either of the <PAUSE> or <RETURN> keys have been pressed. If so it returns with the appropriate indication to the main programme. This provides a convenient way of allowing the operator to interrupt the programme cycle without requiring the use of hardware interrupts.

Subroutine DATA loads 9 Bytes of information into the data buffer per sampling period - the number of the present sampling period, the real-time, each of the 4 twitch values and the twitch ratio being controlled.

MODULE KEYIN

This module contains four subroutines - INITIAL_DATA, TRANSMIT_WORD, TRANSMIT_DWRD and NUMBER. These are used to enter various settings of the control system. As the system is a research prototype every parameter of interest was included. The majority of these could however be fixed for routine control proceedures.

Subroutine INITIAL_DATA is called by the main programme DDOSE. It prompts the user to enter the mode of control (T1,T2 or T4), the 2 parameters which

pertain to the drug pump (section 2.4) and 8 settings of the control algorithm. (These are explained in section 4.7.)

The user has the option of entering any parameter at the keyboard or selecting a default value which is stored in EPROM. INITIAL_DATA calls either of subroutines TRANSMIT_WORD or TRANSMIT_DWRD to transfer whichever parameter was selected from the stack top of the 8087 to a region in the Extra Segment (ES) of memory for tempory storage. TRANSMIT_WORD is used if the parameter is treated as a Word Integer (16 bit) and the other if it is treated as a Short-real (32 bit). If the default value is selected, it is loaded on the 8087 stack top with a simple LOAD instruction; if the keyboard is used, subroutine NUMBER calculates the value of the parameter from the digits entered with the aid of the 8087 and places it on its stack.

MODULE RESET

This module has two subroutines - BOLUS4_RESET and RECEIVE_DATA. BOLUS4_RESET is called by the main programme DDOSE immediately that it exits subroutine INITIAL_DATA discussed above. BOLUS4_RESET in turn calls RECEIVE_DATA to read the parameters which were temporally stored in the Extra Segment (module KEYIN). These two subroutines together perform the following functions:

1. Define the 8 controller and 2 drug pump settings.

2. Reset the counters used by the control algorithm.

3. Clear the controller mode flag so as to select Initial Phase control (section 4.5).

4. Reset the sampling period counter.

5. Clear the flag which signals that a drug command is to be over-written by the user.

6. Initialize the twitch ratio value at one sampling period behind real-time.

7. Clear the four arrays of 1000 elements each used by the control algorithm.

MODULE CONVRT

There are five subroutines in this module, two of which are utility routines used to present data to the 7 segment display. The other three are:

1. T4_CONVERT: This converts a Word Integer to the Short-real format and is called by the control algorithm BOLUS4 so to transform the latest twitch ratio. If it encounters the data invalid value of FF Hex, it operates on the last valid ratio received.

2. U_CONVERT: This subroutine is called by the control algorithm BOLUS4. It rounds the drug command calculated there to the nearest 0,1 mg., displays this value and commands the Drug Pump to administer this dose by setting the 8254 pump controller appropriately. It also codes the latest drug dose and stores it temporally in the Extra Segment of memory from where it is written to the data buffer.

3. TOTAL_DRUG: This determines the accumalated drug dose administered to the patient from the start of the operation to the time that the subroutine is called by the supervising programme DDOSE. The total dose is rounded to the nearest milligram and displayed.

MODULE TIME

This uses the real-time clock to calculate the real-time from the start of the control operation in seconds and in sampling periods.

MODULE RATIO

This calculates one of the T1, T2 or T4 ratios from the train-of-four response derived from the Twitch Measurement System.

MODULE BOLUS4

This contains the control algorithm and is discussed in Chapter 5.

2.9 Software for off-line analysis of the patient data files

The software of the microcomputer controlled system described in section 2.8 allows the user to transfer patient data files from cassette tape to disk storage on the PDP 11 minicomputer. These files are transfered to the UNIVAC 11/Ø8 via an Industry Standard magnetic tape. They are then analysed with the aid of programme CONTRL.ANLYSE described in appendix D.

This programme prints the patient's data file in a suitable format and additionally it allows the user to plot any combination of the following graphs:

1. The twitch ratio being controlled against time.

2. The best control possible assuming that the estimated patient parameters are available (from programme IDNTFY.VAØ5A - section 3.7).

3. Parameters "GAIN" and "BT4" as they were estimated on-line by the controller.

4. The T4 ratio against the T1 ratio and the T2 ratio against the T1 ratio for the induction period of control.

5. The T4 ratio against the T1 ratio and the T2 ratio against the T1 ratio for the recovery phase of the operation.

When patient data is written to cassette tape, all numbers are decomposed into bytes, the Word Integer and Short-real types being 2 and 4 bytes respectively. Each byte is then expressed as 3 octal digits and written to tape in Ascii.

CONTRL.ANLYSE decodes the data file then calls separate subroutines for each of the plots required.

The patient data files have two further purposes. These are in modelling the patient response (programme IDNTFY.VAØ5A, section 3.7) and in checking the performance of the State Estimator used within the control algorithm (programme CONTRL.STEPLT - section 4.5.3).

CHAPTER 3

The Patient Model

3.1 Introduction

The aim of the present work is to design a system capable of controlling the level of induced muscle relaxation in humans. It would greatly assist the design to have a model of the process being controlled. To be specific there are three applications of system modelling in this thesis.

1. to obtain a model structure on which to base the design of the controller

2. to obtain patient parameters with which to initialize the control algorithm

3. to enable realistic simulations of patients to be made for evaluation of the candidate controllers.

The concept of system modelling has been put onto a rigorous foundation in the field of System Identification. In their survey paper on this subject, Astrom and Eykhoff,(1971) point out that certain principles underlie the many identification techniques which have evolved. They are

1. the model structure chosen to represent the process being identified

2. the method used to compare the process and its model

3. the functional which is calculated from the comparison in 2. above and

which is processed to obtain the model which is "best" in the required sense.

4. the means of system excitation (if this can be chosen freely).

The term 'identification' usually refers to estimation of parameters of a linear model. It is shown in this chapter however that the patient response is non-linear and hence the phrase 'parameter estimation' will be used in place of 'identification' henceforth.

In the following discussion, the model of the human response (to be referred to as the "patient model") will be considered to be subdivided into two parts as shown in figure 3.1

1. the deterministic model (that which produces an output in response to deterministic inputs)

2. the noise or stochastic model (that which contributes to the output but is independant of the deterministic inputs; its origin is in the process or measurement noise).

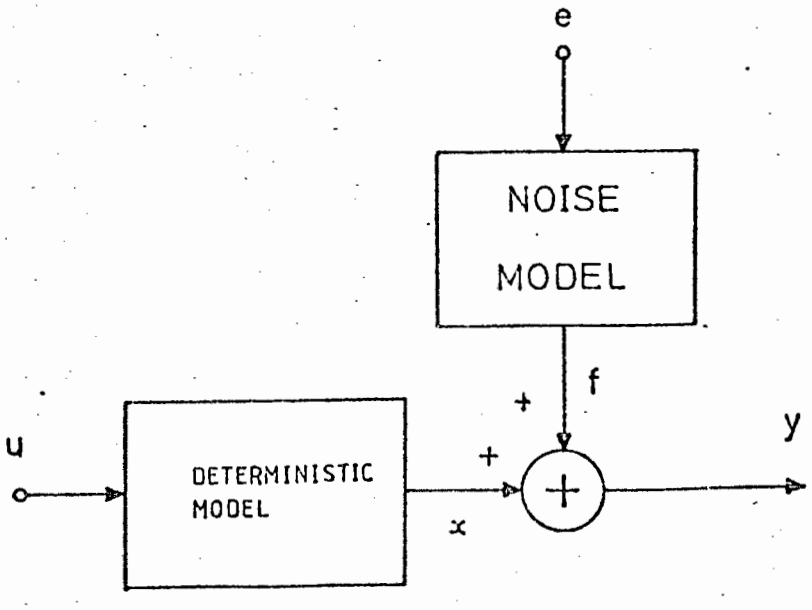

Figure 3.1 : The human blocking effect (y) is synthesised as the output (x) of a deterministic model excited by the muscle relaxant (u) and the output (f) of a noise model driven by a white noise source (e).

3.2 The Structure of the Deterministic Model

The aim here is to model the human pharmacological effect in response to the administered relaxant. This requires knowledge of the pharmacokinetics and pharmacodynamics of the drug in use. By pharmacokinetics is meant the mathematical description of the processes and rates of drug movement within the body from the sites of administration into the blood, distribution into the tissues and elimination by metabolism or excretion. Pharmacodynamics relates the drug concentration at a particular site to the pharmacological effect produced (Stanski and Watkins,1982).

It seems that 3 approaches have been made to model the pharmacological effect. The earliest of these is the two stage method used by the medical researchers. See for example Ham et al.,(1979), Stanski et al.,(1979) for dTC, Hull et al.,(1980) for Fazadinium and Pancuronium. They first determine the pharmacokinetics of the drug by modeling the plasma concentration from a number of blood samples taken over a period of time. They then fit a non-linear pharmacodynamic model to the concentration/ effect data. They do this because actual knowledge of the drug concentration is of importance to them. However, if the only concern is to model the level of relaxation for control purposes, which is the case in the other 2 approaches, one can either fit a composite non-linear model to the effect data directly (Bradlow et al., 1983) or linearize about the level at which one intends to regulate (Linkens et al.,1982).

The work of the medical researchers in modeling the pharacokinetics of dTC and related non-depolarisers has been based on Compartmental analysis. See for example the book by Stanski and Watkins,(1982). This method has intuitive appeal. One postulates a number of connected compartments or areas of action within the body on semi-physiological grounds and then proceeds to describe the transfer of drug between these compartments using differential equations. The impulse response of the plasma concentration in the compartment of interest is

then a sum of decaying exponentials.

Both two and three compartment models have been used to describe the pharmacokinetics of dTC in man. The three compartment model (Gibaldi et al., 1972, Wingard and Cook, 1976 and Meijer et al., 1979) has an impulse response

$$x(t) = ae^{-\alpha t} + be^{-\beta t} + ce^{-\gamma t} \quad \ldots \quad 3.1$$

where a,b,c, α, β and γ are real numbers. These three research groups were concerned with the elimination kinetics of dTC after intravenous injection. Gibaldi et al. found that the long term elimination pattern (14 hours) indicated the use of a three compartment model. Wingard and Cook considered both two and three compartment models and found that the latter gave better agreement with measured concentration data at 3 hours or longer after administration of the dose. Meijer et al. state that their three compartment model gave the best fit to the experimental data (elimination for 10 hours).

Stanski et al.,(1979) however found by statistical comparison with a three compartment model that the two compartment model with impulse response

$$x(t) = ae^{-\alpha t} + be^{-\beta t} \quad \ldots \quad 3.2$$

is adequate for surgical proceedures of up to 4 hours.

Two models have been proposed for the pharmacodynamic expression. These are the logistic relationship (Ham et al., 1979)

$$T(t) = \frac{1}{1 + e^{AT4x(t) + BT4}} \quad \ldots \quad 3.3$$

and Hill equation (Sheiner et al, 1979)

$$T(t) = \frac{1}{1 + Kx(t)^{\gamma}} \quad \ldots \quad 3.4$$

which relate the fractional blocking effect T to the plasma concentration x in the compartment of interest. It is assumed that AT4, BT4, K and γ are all real constants. Both expressions are non-linear sigmoidal forms.

Ham et al. obtained the complete pharmacological response from the composite model of equation 3.1 and equation 3.3. Sheiner et al. however were unable to use the simpler two compartment model directly for this purpose because they found that the time response of the effect lags behind that of the drug concentration. They therefore introduced a third compartment, a hypothetical "effect" compartment which receives a negligible mass of drug and thus has no bearing on plasma concentration. They consequently modelled the human pharmacological effect of dTC with a 3 compartment structure and the Hill equation 3.4.

When one is concerned with modeling the blocking effect for the purpose of on-line control, it is impractical to measure the plasma concentration for this involves taking blood samples during the control period and analysing them immediately. It is also unnecessary to do so because the pharmacokinetic model and pharmacodynamic relationship can be combined into a composite model. Such a model is non-linear and is more appropriate for transitional control than a linearized version, (which may however be suitable for regulation). It is more suitable because best results of identification can be expected when the unknown system is investigated in conditions under which it will be controlled (Balakrishnan and Peterka, 1969).

Bradlow et al. found that a two compartment model was adequate to fit the Effect data directly. It had an impulse response

$$x(t) = a(e^{-\alpha t} - e^{-\beta t}) \quad \ldots \quad 3.5$$

The non-linear part of their composite model was the inverse logistic relationship derived from equation 3.3

$$x(t) = \ln\left[\frac{1 - T(t)}{T(t)}\right] - BT4$$

where the constant AT4 has been incorporated into 'a' of equation 3.5. As he considers the short term action of the drug (up to about 2 hours) and models the effect data exclusively, Bradlow's findings are significant for this work and the model structure which he has derived is used henceforth in this research.

Equation 3.5 is a time domain expression. It is more convenient to work in the discrete time domain however because the envisaged controller is digital and although the patient can be considered as a continuous time system, measurements are made only at integral multiples of a fixed sampling period. Straight forward algebraec manipulation shows that the Z transform of equation 3.5 is equation 3.6

$$X(z^{-1}) = \frac{gz^{-1}}{(1 - p_1 z^{-1})(1 - p_2 z^{-1})} U(z^{-1}) \quad \ldots \quad 3.6$$

where

$$g = a(e^{-\alpha \cdot TSTEP} - e^{-\beta \cdot TSTEP})$$
$$p_1 = e^{-\alpha \cdot TSTEP}$$
$$p_2 = e^{-\beta \cdot TSTEP}$$

If one is concerned only with modelling the blocking effect, there is in fact only 1 parameter to be determined in equations 3.3 or 3.4 because the constants AT4 or K can be incorporated into the coefficient 'a' of equation 3.5. Therefore without loss of generality and in discrete time equations 3.3 and 3.4 can be written as

$$T(k) = \frac{1}{1 + e^{x(k) + BT4}} \quad \ldots \quad 3.7$$

$$T(k) = \frac{1}{1 + x(k)^{\gamma}} \quad \ldots \quad 3.8$$

In fitting the non-linear element, there are 4 possible permutations corresponding to equations 3.7, 3.8 and their inverses 3.9 and 3.10 respectively

$$x(k) = \ln\left[\frac{1 - T(k)}{T(k)}\right] - BT4 \quad \ldots \quad 3.9$$

$$x(k) = \left[\frac{1 - T(k)}{T(k)}\right]^{\frac{1}{\gamma}} \quad \ldots \quad 3.10$$

As the equations are non-linear, the forward and inverse relationships cannot be expected to yield the same estimated parameters.

The patient model of equation 3.6 assumes a unit delay between input and output. It should be modified to cater for the dead-time which always exists owing to the circulatory delay in the blood-stream. If this delay is 'ℓ' sampling periods the model is

$$X(z^{-1}) = \frac{z^{-\ell} g}{(1 - p_1 z^{-1})(1 - p_2 z^{-1})} U(z^{-1}) \quad \ldots \quad 3.11$$

The delay is a discrete variable and therefore cannot be fitted in the same way as the other model parameters. One can however re-iterate the parameter estimation process and find the best value by trial and error. Thus the delay can be regarded as a pseudo-parameter which to some extent compensates for modelling errors. This approach was used instead of fixing the delay by visual inspection of the patient response for it gives a better model as judged by the criterion of fit used.

3.3 Comparing the process and its model.

Astrom and Eykhoff,(1971) discuss 3 ways of comparing the process being identified and its model. These are known colloquilly as Input Error, Output Error (OE) and Equation Error (EE). Input Error is applicable to situations where there is uncertainty in the measurement of the process input signal. This method is not considered because it is assumed that the measurements are known exactly. The assumption is justified for the first half of the patients modelled because the relaxant was administered manually and with great care. In the remaining cases the relaxant was administered by a syringe pump and the quantisation error was insignificant except for the very smallest doses administered (see section 2.4). The advantage of OE over EE is that the deterministic model may fitted independently of the noise model (Moore, 1982). This is of value because

1. the patient model on which the controller is to be based is deterministic and this can be found without stochastic considerations and inaccuracies from potential errors in the noise model

2. it allows one to separate the estimation of the deterministic and noise model parameters thereby simplifying the estimation process (Moore, 1982).

The advantage of EE over OE is that the residuals are a linear function of the model parameters. If the process being modelled were known to be linear this property implies for the method of Least Squares that the parameter estimates converge to a unique point in parameter space. However as the process is known to be non-linear it is not clear whether EE maintains this advantage over OE. The point is beyond the scope of this thesis and will not be persued here. Suffice it to say that visual inspection of curve fits, estimated parameters and off-line control of simulated patients appeared to be reasonable in every case.

3.4 The functional

Two of the most commonly used methods for processing the residuals 'r' are Least Squares (LS) and Maximum Likelihood (ML). ML requires that the probability density function and variance of the noise be known, LS does not. Assuming that this information is available and one is prepared to use it, the computational effort associated with ML exceeds that of LS (Eykhoff, 1982). As LS is the simpler of the two methods and most of the researchers cited in this chapter have used it, it will be considered exclusively henceforth.

The LS method adjusts the deterministic parameter vector $\underline{\theta}_D$ to minimize the loss function V of equation 3.12 for N received data.

$$V(\underline{\theta}_D) = \sum_{k=1}^{N} r^2(k).W(k) \quad \ldots \quad 3.12$$

W is a weighing sequence and in this work is allowed to take the values 1 or Ø so as to accept or reject data. Rejection occurs during periods that measurement of the twitch response is interrupted or the twitch ratio is outside the range of validity of the pharmacodynamic equation. The twitch ratio is declared valid on the range Ø,1 to Ø,8. The reason for this choice is discussed in section 4.5.4 which deals with parameter estimation within the controller.

There is no known closed form solution to the LS problem for a system which has the non-linearity described. Numerical minimization on the other hand is straight forward. The non-linearity and weighing sequence can be included easily. For the OE formulation, there are four residuals possible which are derivable from equations 3.7 to 3.1Ø. They are

$$r_c(k) = \left[T(k) - \frac{1}{1 + e^{X(k) + BT4}} \right] \cdot W(k) \quad \ldots \quad 3.13$$

(logistic relationship)

$$r_c(k) = \left[\ln\left(\frac{1 - T(k)}{T(k)}\right) - X(k) - BT4\right] \cdot W(k) \quad \ldots \quad 3.14$$

(inverse logistic relationship)

$$r_c(k) = \left[T(k) - \frac{1}{1 + X(k)^{\gamma}}\right] \cdot W(k) \quad \ldots \quad 3.15$$

(Hill equation)

$$r_c(k) = \left[\left(\frac{1 - T(k)}{T(k)}\right)^{\frac{1}{\gamma}} - X(k)\right] \cdot W(k) \quad \ldots \quad 3.16$$

(inverse Hill equation)

Equations 3.13 and 3.15 give rise to the parameter estimation structure shown in figure 3.2(a). As the Output Error uses the patient's twitch ratio directly, this structure is appropriate for patient models intended to simulate the twitch response. However the patient model is non-linear in these cases and it is less convienient to design a controller for such a model than for the structure of figure 3.2 (b) which is derived from equations 3.14 and 3.16. This structure is therefore used for control purposes in Chapter 4.

If it is assumed that the noise on the patient output is symmetrically distributed about zero mean, then strictly speaking this cannot be true when passed through the non-linearity of figure 3.2(b). However the non-linear element can be regarded as piecewize linear if the noise which perturbs about the operating point (the true level of relaxation) has a small enough variance. This complication will therefore be ignored.

3.5 The Input Signal

The type of input signal has a bearing on the accuracy of an identification experiment. In particular for the method of Least Squares the input has to satisfy the condition of persistent excitation if the identified model

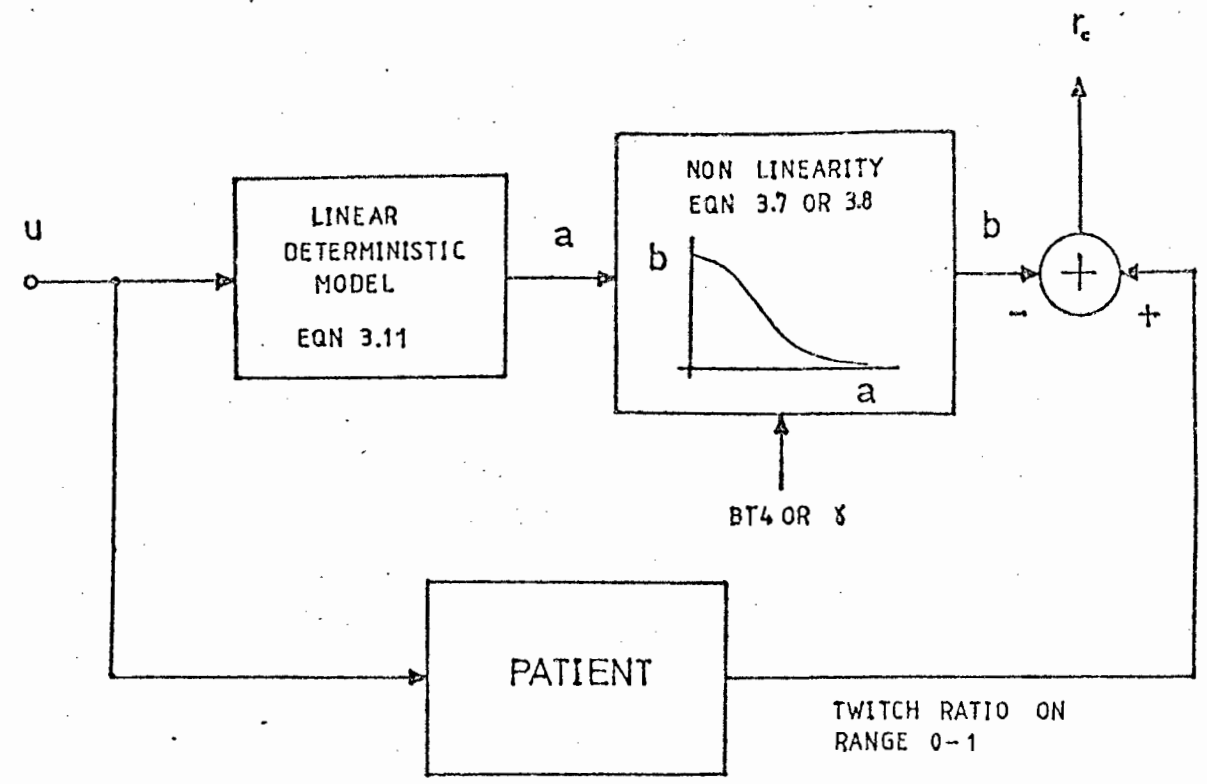

(a)

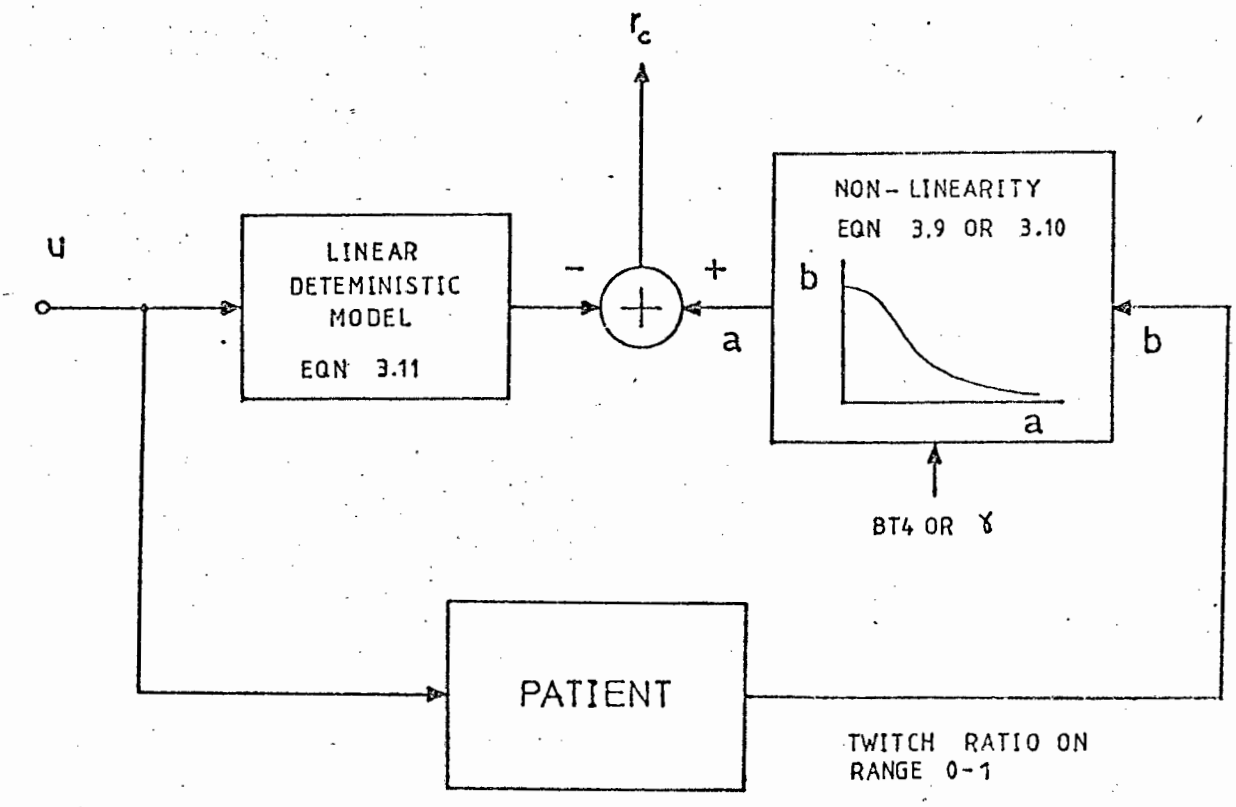

(b)

Figure 3.2 : Formation of the Output Error when using the forward and inverse versions of the non-linear pharmacodynamic relationship. These are shown in Figures 3.2(a) and 3.2(b) respectively.

parameters are to be consistent (the model parameters tend to the process parameters as the number of data tend to infinity) (Astrom and Eykhoff, 1971). A simple definition of persistent excitation for controllable systems has been given by Young,(1970). According to this the input signal must have 2 properties

1. it must activate the process throughout the observation interval

2. the number of distinct frequency components present in the signal must equal or exceed 'd' where

$$d = \begin{cases} \text{no. parameters}/2 & ; \text{ no. parameters is even} \\ (\text{no. parameters} + 1)/2 & ; \text{ no. parameters is odd} \end{cases}$$

For the intended application, the alternative approaches are to use a specially designed test input which satisfies the above conditions or to use the normal operating records i.e. the response to the usual dosage protocol administered during routine operations (or from that administered by candidate controllers). The second of these is desirable because it is more convenient but more especially because it allows off-line simulations on the very patient undergoing clinical trials. There is however no guarantee that the input is persistently exciting.

Perhaps the most common form of persistent excitation is a discrete binary noise type input such as a Pseudo Random Binary Sequence (PRBS). This was used by Linkens et al.,(1982), but is not without its disadvantages. In its application the PRBS would be applied to the patient once some adequate degree of relaxation had been induced. It would then perturb about this operating

point. But this would not estimate the patient parameters under the conditions he would be controlled. Further with a drug whose duration of action is as long as dTC, the bit interval would need to be set to a number of normal sampling periods which are typically 10 seconds each. (Linkens et al. used bit intervals of 33,3 and 100 seconds for pancuronium). One then loses information on the intersample noise correlations. Thus it is not clear that PRBS excitation would give better estimation of either the deterministic or noise model parameters.

3.6 The noise model

When identification is done on a process in which the noise is not white, stochastic models are introduced to account for the self-correlation in the noise. The most general form of such a model is the autoregressive moving average (ARMA) representation expressed by equation 3.17 in discrete time as

$$F(z^{-1}) = \frac{C(z^{-1})}{D(z^{-1})} E(z^{-1}) \quad \ldots \quad 3.17$$

where E is the Z domain description of a white noise source and C and D are polynomials.

$$C(z^{-1}) = 1 + c_1 z^{-1} + \ldots + c_n z^{-n}$$
$$D(z^{-1}) = 1 + d_1 z^{-1} + \ldots + d_m z^{-m}$$

The zeroes of C may lie on or within the unit circle but the zeroes of D must lie within it (Astrom and Soderstrom,1974). The two important simplifications of the ARMA model are the autoregressive (AR) and moving average (MA) models in which the coefficients of the C and D polynomials are taken as zero respectively. All 3 models have been investigated in the context of Least

Squares with the Equation Error criterion. See for example Soderstrom, Ljung and Gustavsson,(1978) for AR and MA and Talmon and van den Boom,(1973) for all three formulations. It was found that consistency of the parameter estimates is not automatically guaranteed for any of these. The ARMA representation requires the smallest number of parameters for an adequate noise model description (Talmon and van den Boom, 1973).

The AR version of the noise model in this work was used with the Output Error formulation for the deterministic model (Bradlow and Rametti, 1982) because it is simple to implement.

Given a sequence of N correlated residuals $r_c(k)$, $k = 1,\ldots,N$ derived, say, from the deterministic model equations 3.13 - 3.16 it follows from equation 3.17 that an AR model can be fitted to them by minimizing

$$V(\underline{\theta}_s) = \sum_{k=1}^{N} r^2(k).W(k) \quad \ldots \quad 3.19$$

where

$$r(k) = r_c(k) + d_1 r_c(k-1) + \ldots + d_n r_c(k-n)$$

The number of significant parameters in the stochastic polynomial can be determined by applying the F-test (Mandel, 1964) to the loss function V. This method of parameter estimation has intuitive appeal. It strives to reduce the sequence N to white noise which condition is necessary when all information has been extracted from the data.

The use of Output Error simplifies parameter estimation as noted previously because the fitting of the deterministic and noise models is then separated. There are three steps to be followed :

1. obtain the deterministic model using the LS Output Error method and

hence the vector of deterministic parameters $\underline{\theta}_D$

2. subtract the best fitting deterministic curve from the original data to form a sequence of correlated residuals

3. obtain the noise model and hence the vector of stochastic parameters $\underline{\theta}_S$.

The subtraction needed to calculate the correlated residuals is performed in both of the ways shown in figure 3.2. These give rise to the patient models shown in block diagram form in figure 3.3.

The model of figure 3.3(a) is intended for patient simulation. The noise model is placed as shown and not between the deterministic model and non-linear element for two reasons:

1. Clinical experience has shown that the level of the noise is approximately independant of the magnitude of the twitch ratio.

2. With the noise-model in the alternative position, one cannot separate the estimation of deterministic and noise model parameters in the manner proposed because the variance of the noise measured at the model output is a function of the operating point on the non-linearity which is determined by the output of the deterministic model.

The patient model shown in figure 3.3(b) is linear and for this reason is intended for use in the controller design as explained in section 3.4.

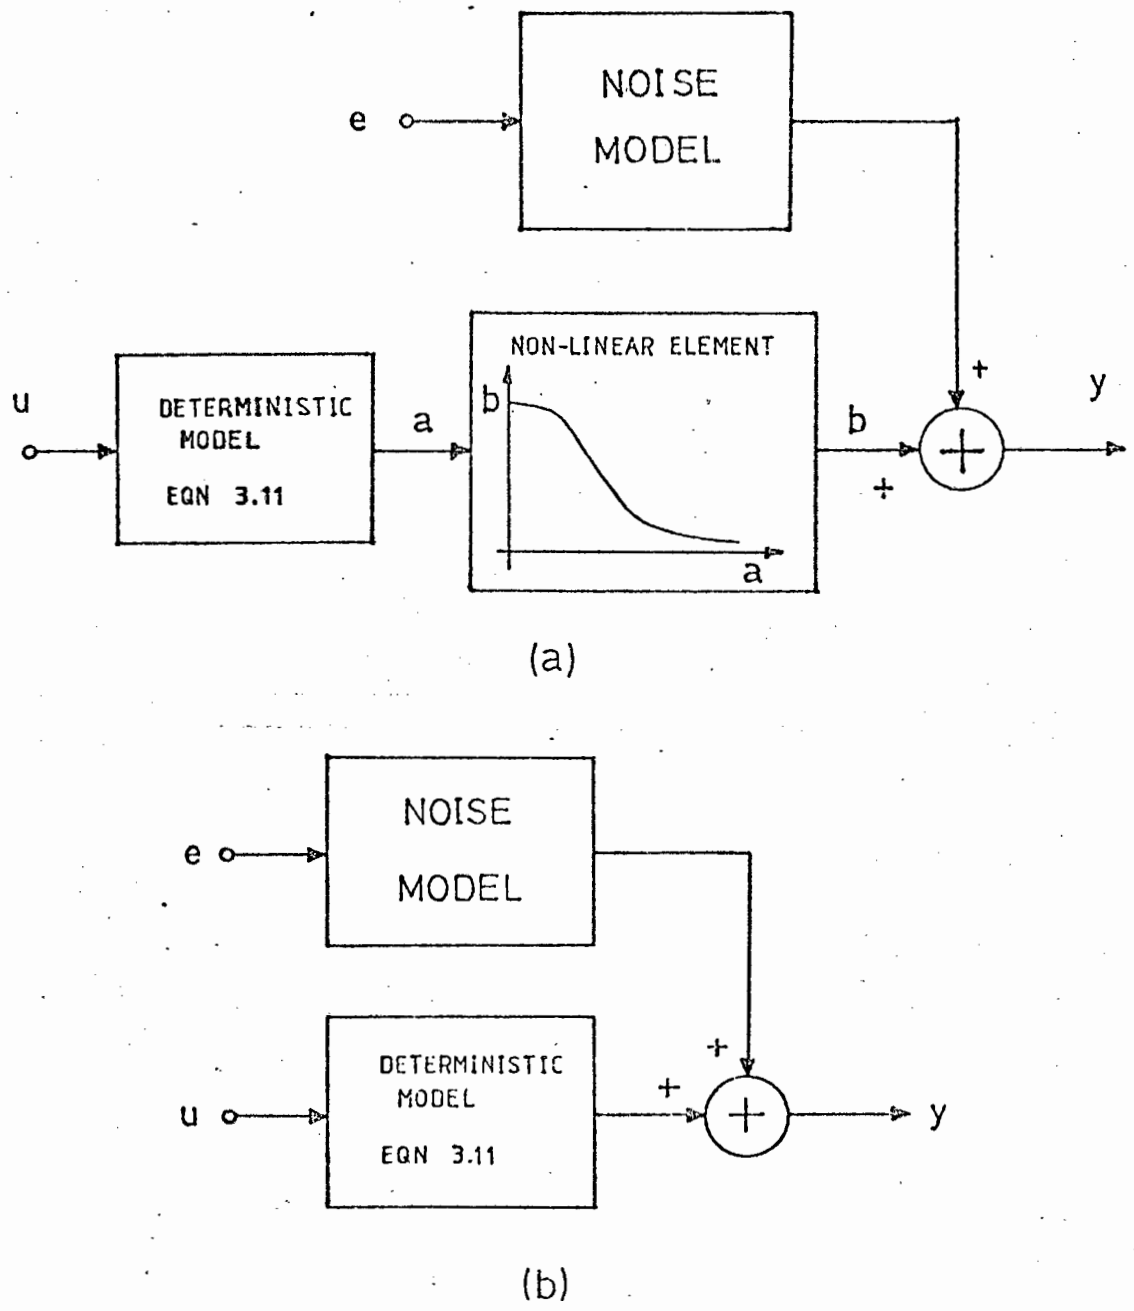

Figure 3.3 : Block diagram of two structures for the patient model -
(a) that used to simulate the patient twitch response,
(b) for use in the controller design.

3.7 Software written to model the patient effect response

A program IDNTFY.VA05A was written in FORTRAN to model the pharmacological effect response of patients to muscle relaxants and is used in the analysis of section 3.8. It fits a discrete time deterministic model by the LS method using the Output Error criterion. The loss function is minimized using subroutine VA05A of the Harwell library (Hopper, 1974). This uses a combination of Steepest Descent, Newton's and Marquardt's methods.

The user selects the number of poles and zeroes and the number of parameters fitted in the Autoregressive noise model. Any of the four non-linear elements discussed in section 3.4 may be used. The circulatory dead time is not fitted and must be supplied by the user.

The programme reads the data files created by the on-line controller (Chapter 2). In addition to determining the required model parameters, it calculates the autocorrelation function of the residuals after the noise model has been fitted.

This programme was used in the modelling reported in Bradlow and Rametti, (1982). The listing and further details are given in Appendix D.

3.8 Analysis of patient data

This section concerns off-line parameter estimation analysis performed on 44 real patients. Ten of these were reported in Bradlow and Rametti,(1982) and 34 in Rametti et al.,(1984). These were performed at different times and the patients were prepared under different medical conditions.

3.8.1 Preparation of the patients

The patients entered into the trials were all adults, ASA Grade I or II (i.e. normal healthy patients or those with mild system disease) in whom disturbance at the neuromuscular junction or muscle pathology was not suspected.

The first group of 1Ø (patients 1 - 1Ø) in which the twitch response was monitored but not controlled on-line underwent various surgical proceedures. They did not receive any aminoglycoside antibiotics. After premedication at the anaesthetist's discretion, anaesthesia was induced with thiopentone sodium and maintained with 4l of N_2O, 2l of O_2 and Ø,5% halothane. The end tidal CO_2 was maintained at 4%. After topical spray of 4 ml 4% lignocane to the upper respiratory tract, the trachea was intubated under vision with a cuffed endotracheal tube. Reversal of relaxation at the end of the operation was achieved with 2,5 mg neostigmine and 1,2 mg atropine. The transducer was placed in the palm of the patient's hand as described in Chapter 2. Stimulation of the ulnar and median nerves was obtained with 21G needle electrodes placed at the wrist 25 mm apart parallel to these nerves. The period between the train of 4 stimuli was 9,Ø5 seconds and the width of each stimulus 3ØØ μ Sec.

The second group of 16 (patients 11-26) in which the T4 ratio was monitored and controlled on-line were premedicated either with benzodiazepine or narcotic. They were induced with thiopentone 3-5mg/Kg, intubated with 1-1,5mg/ Kg succinyldicholine, maintained with 7Ø% N_2O in O_2 and supplemented by 1-Ø,5% halothane with narcotic as required during the control phase.

The third group of 18 (patients 27-44) in which the T1 ratio was monitored and controlled on-line were premedicated with diazepam 1Ø mg orally one hour pre-operatively. Induction of anaesthesia included fentanyl 1mg/Kg and thiopentone 3-5 mg/Kg. Intubation was achieved with the use of lignocane 1,5 - 2 mg/Kg sprayed topically onto the vocal cords, with halothene 1%. Maintenance

of anaesthesia included 70% N_2O in O_2 with halothane 0.5%, fentanyl 50 mg being added at half hourly intervals.

The latter two groups of patients were stimulated through surface electrodes (paediatric ECG type). The stimulation pulse width was about 150 μSec. The period between the train-of-four stimuli was 10,0 seconds. The transducer was either mounted in the palm of the hand or on the arm-board as described in Chapter 2. Administration of dTC only commenced once the response to ulnar nerve stimulation had stabilized (halothane concentration remaining at 0,5%). Patients were ventilated to maintain normocapnia, arterial pressure and electrocardiogram being monitored and recorded. All drugs known to potentiate dTC were withheld during the period of on-line control. On completion of surgery neostigmine 2,5 mg and glycopyrrolate 0,4 mg were given to both groups of patients and on-line control terminated. The T4 ratio was however further monitored until it exceeded 70% prior to extubation and return of the patients to recovery facilities.

In all 3 groups of patients the stimulation voltage was adjusted until supramaximal response was obtained prior to administration of dTC.

3.8.2 Results

The relative suitability of the logistic relationship, Hill equation and their inverses was first studied. The original analysis concerned the batch of patients 1-10 only. Loss functions were calculated corresponding to each of the four residuals defined in equations 3.13 to 3.16. The pharmacodynamic parameters γ and BT4 were then varied over a range of values. The plots of the four loss functions so obtained are shown in figure 3.4.

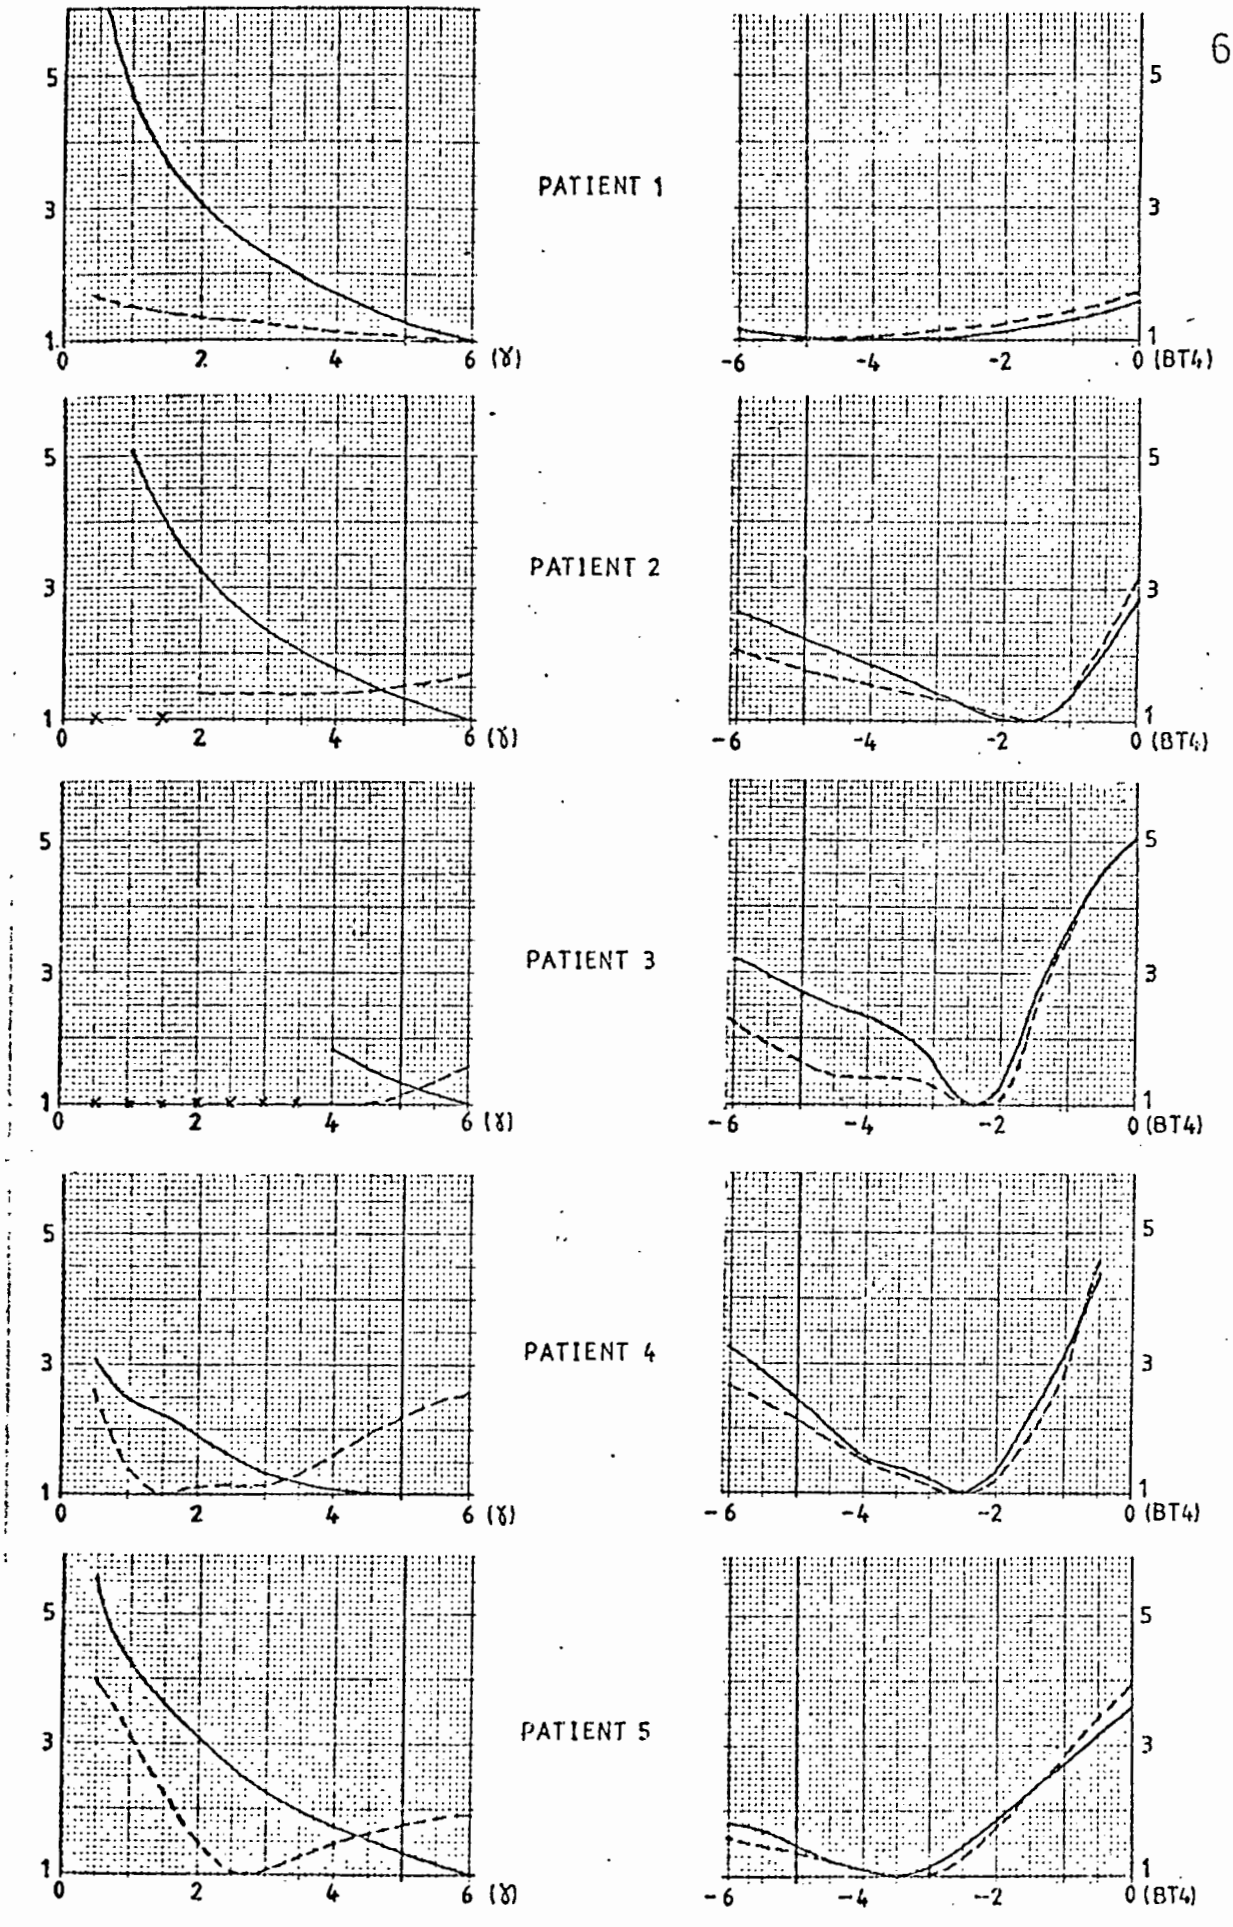

Figure 3.4 : The sensitivity of the loss function to the parameter of the Hill equation (left hand column) and BT4 of the logistic relationship (right hand column). The dashed curves are the forward relationship and the solid curves the inverses. The Y co-ordinate has been normalised to the minimum value of the loss function in each case.
(x = poles of model were outside the unit circle in the Z plane).

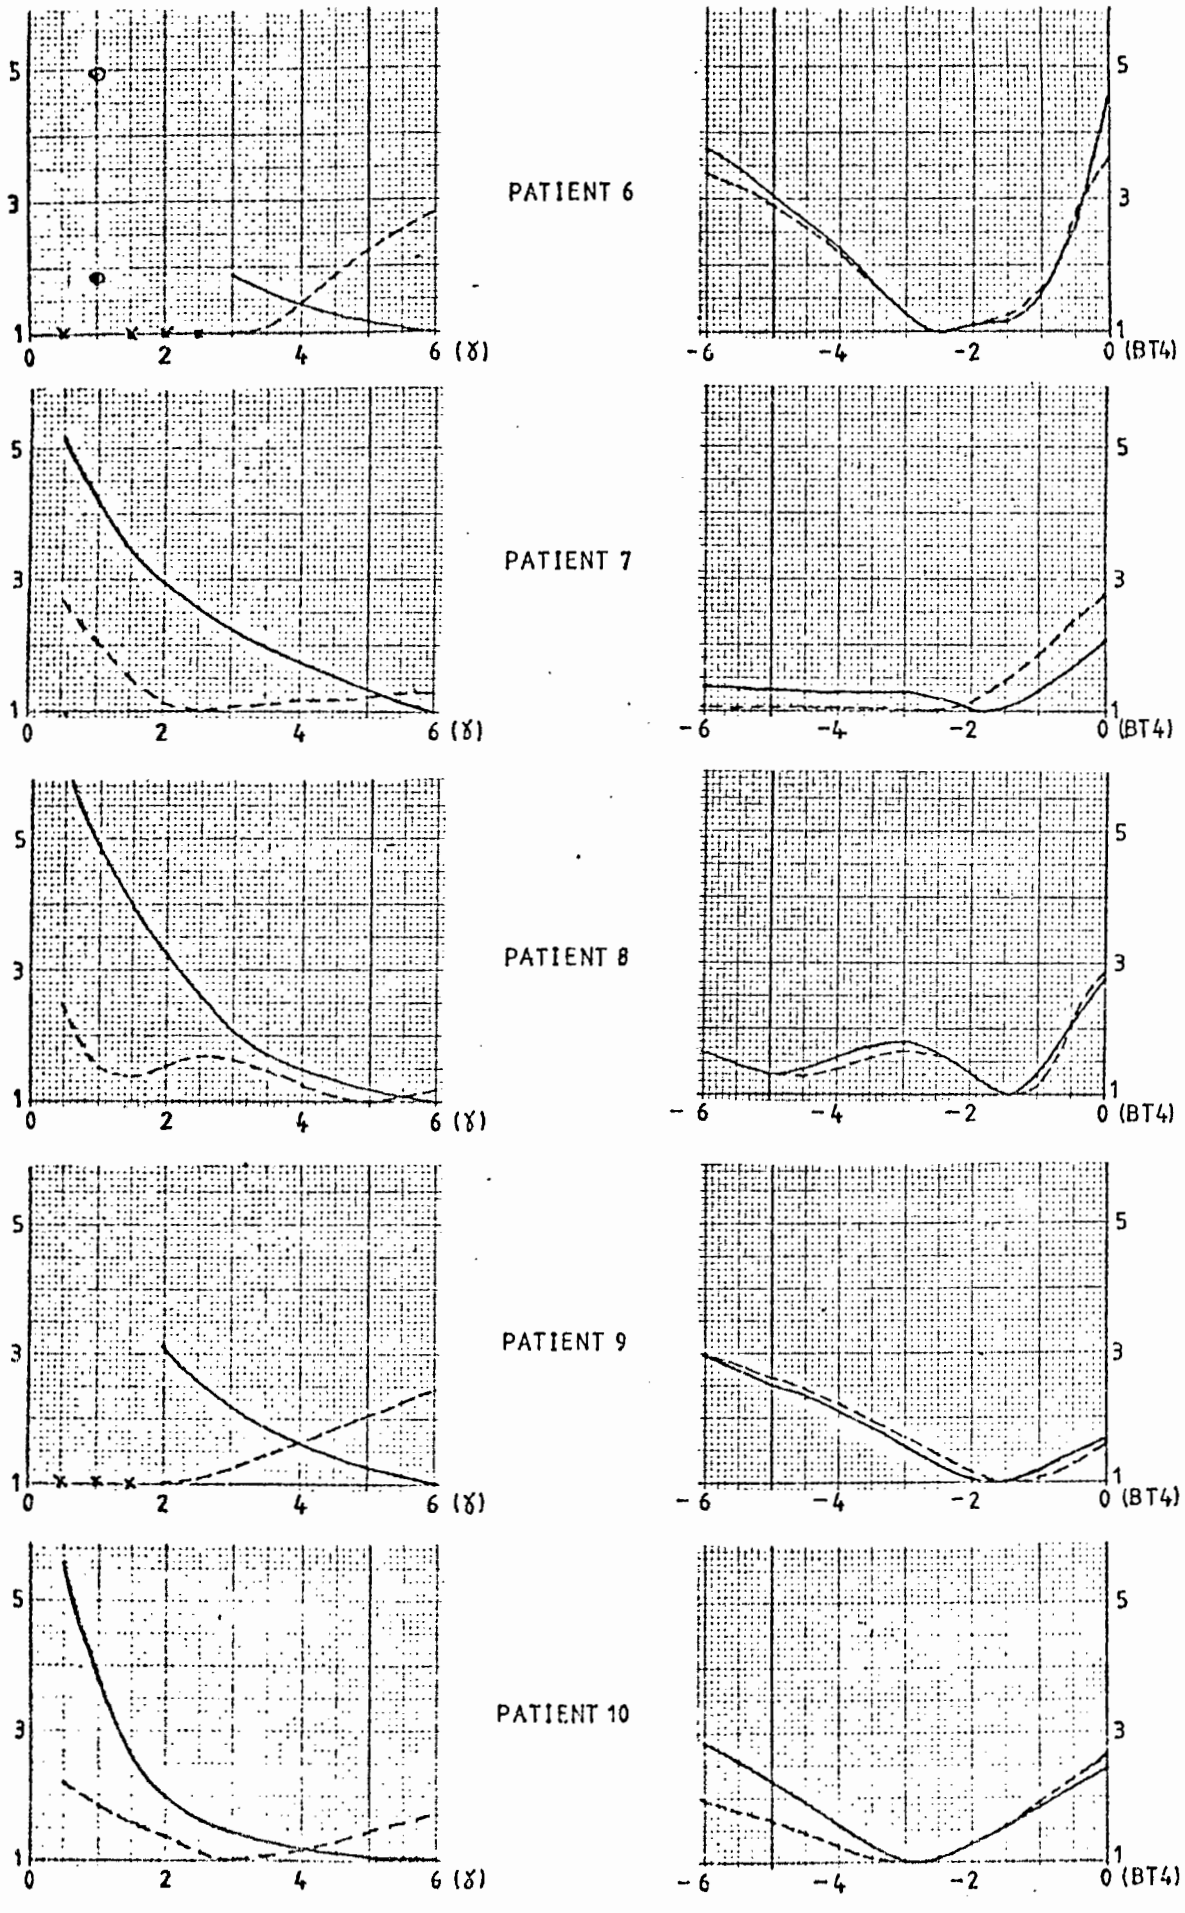

figure 3.4 (contd.)

Use of the Hill equation resulted in unstable models (poles outside the unit circle in the Z plane) in 4 cases for a certain range of γ which was close to the optimum value.

For the inverse Hill equation, no best value of γ could be found over the range Ø - 6 which was searched and the corresponding patient parameters gave a loss function between actual twitch ratio and model output which was considerably inferior to that obtained for the forward equation. The Hill equation and its inverse were therefore excluded from further analysis.

The logistic relationship was however well behaved and gave approximately the same model parameters and loss functions for both forward and inverse equations.

Part of the analysis was then extended to all 44 patients. The model parameters were estimated under slightly different conditions. The value of circulatory delay which resulted in the smallest loss function was used in each case. In the former work on patients 1 - 1Ø it had been determined by visual inspection of the twitch response. The limits of validity for the twitch ratio was Ø,1 to Ø,8. Formerly it had been Ø,1 to Ø,9.

The loss function between actual twitch ratio and best fitting curve was then calculated for each of the 44 patients with models which had used the logistic relationship and its inverse. These are given in Table 3.1. The average difference between the two columns is only 4,9 $\pm$ 5,3%.

It was stated in the introduction to this chapter that there are three applications of system modelling in this work namely to obtain a structure on which to base the design of the controller (sections 3.4 and 3.6), to model real patients for later simulated control experiments and to obtain parameters with which to initialize the control algorithm.

PATIENT	LOGISTIC RELATIONSHIP	INVERSE LOGISTIC RELATIONSHIP
1	0,734	0,739
2	0,081	0,081
3	0,025	0,027
4	0,060	0,061
5	0,137	0,146
6	0,079	0,081
7	0,736	0,904
8	0,254	0,264
9	0,179	0,198
10	0,574	0,589
11	0,181	0,182
12	0,094	0,098
13	0,304	0,313
14	0,222	0,253
15	0,840	0,887
16	0,115	0,115
17	*	0,668
18	0,692	0,739
19	0,895	0,908
20	0,265	0,271
21	0,256	0,269
22	0,300	0,303
23	0,203	0,205
24	0,203	0,208
25	0,123	0,133
26	0,075	0,077
27	0,204	0,214
28	0,183	0,208
29	0,093	0,096
30	0,221	0,225
31	0,277	0,289
32	0,168	0,170
33	0,021	0,023
34	0,678	0,767
35	0,069	0,069
36	0,114	0,117
37	0,260	0,264
38	0,251	0,264
39	0,443	0,454
40	0,415	0,446
41	0,122	0,124
42	0,077	0,091
43	0,168	0,174
44	0,927	0,996

Table 3.1 : Residual sum of squares between actual twitch data and best fitting curve for Output Error Least Squares fits using the logistic relationship and its inverse. The logistic relationship was found to be inappropriate for the patient 17.

In Table 3.2 is given the parameters of the 44 real patients who were modelled in the course of this work. The Output Error LS method was used as was the logistic relationship. The dead time quoted is that which resulted in the best loss function and was determined by trial and error. In estimating the noise model, the F-test was applied to determine the number of significant parameters and is as shown in the table.

The parameter estimation performed within the controller (Chapter 4) uses the inverse logistic relationship. As the circulatory dead-time is not determined on-line, it is fixed at the average value of 5 sampling periods. This value was determined by fitting the dead-time for each patient by trial and error hence finding the optimum value for each which minimized the loss functions in each case. The average for the 44 patients was 45,5 seconds. When rounded this is 5 sampling periods for measurement spacings of both 9,Ø5 and 1Ø,Ø secs. The model parameters of the 44 patients were therefore re-estimated at a fixed lag of 5 sampling periods and using the inverse logistic relationship. The parameters of the models thus obtained are given in Table 3.3 to which reference will be made in later chapters.

The disadvantage of the AR noise model formulation is that a relatively large number of parameters may be required (Talmon and van den Boom, 1973). Hence it was found that 7 patients of the 44 in Table 3.2 and 6 in Table 3.3 required more than 5 parameters. As the F-test quantity was not far from the threshold of significance in most of these cases, no more than 5 parameters were fitted. The noise model is thus somewhat cumbersome but this is of no eventual consequence since the final version of the controller design does not incorporate a noise model at all.

PATIENT	$POLE_1$	$POLE_2$	GAIN	BT4	LAG	d_1	d_2	d_3	d_4	d_5	d_6	R_e
1	0,99863	0,840	0,1154	-4,950	1	-0,401	-0,210	-0,080	-0,094			$9,42 \cdot 10^{-4}$
2	0,99764	0,937	0,0347	-2,736	1	-0,828						$1,13 \cdot 10^{-4}$
3	0,99788	0,946	0,0124	-3,969	5	-0,432	-0,211	-0,181				$5,31 \cdot 10^{-5}$
4	0,99505	0,912	0,0376	-1,210	0	-0,838						$1,14 \cdot 10^{-4}$
5	0,99899	0,970	0,0168	-3,580	1	-0,770						$2,86 \cdot 10^{-4}$
6	0,99855	0,906	0,0257	-2,439	6	0,0						$4,64 \cdot 10^{-4}$
7	0,99824	0,980	0,0228	-5,819	5	-0,400	-0,101	-0,216				$3,32 \cdot 10^{-3}$
8	0,99592	0,987	0,0086	-1,515	2	-0,702	-0,235					$1,71 \cdot 10^{-4}$
9	0,99871	0,805	0,0482	-1,368	4	-0,700	-0,099	0,070	-0,186			$1,66 \cdot 10^{-4}$
10	0,99765	0,830	0,0974	-2,904	7	-0,869						$1,33 \cdot 10^{-3}$
11	0,99801	0,862	0,0659	-1,163	8	-0,196	-0,205					$5,52 \cdot 10^{-4}$
12	0,99831	0,947	0,0255	-1,501	1	-0,699						$2,68 \cdot 10^{-4}$
13	0,99739	0,958	0,0281	-2,432	3	-0,231	-0,212	-0,114	-0,228	-0,055	*	$3,54 \cdot 10^{-4}$
14	0,99970	0,780	0,0890	-0,811	5	-0,253	-0,019	-0,091	-0,102	-0,229		$1,26 \cdot 10^{-3}$
15	0,99146	0,988	0,0104	-1,783	4†	-0,460	-0,119	-0,187	-0,063	-0,109		$5,63 \cdot 10^{-4}$
16	0,99827	0,848	0,0892	-2,665	3	-0,442	-0,191					$5,08 \cdot 10^{-4}$
17 °	0,99856	0,845	0,1250	-2,645	7†	-0,529						$1,08 \cdot 10^{-3}$
18	0,99602	0,969	0,0108	-0,680	11	-0,261						$3,13 \cdot 10^{-3}$
19	0,99581	0,955	0,0296	-1,815	2	-0,427						$1,94 \cdot 10^{-3}$
20	0,99803	0,972	0,00924	-0,362	5	-0,276	-0,222					$4,60 \cdot 10^{-4}$
21	0,98567	0,98559	0,0221	-2,910	5	-0,384	-0,192					$8,81 \cdot 10^{-4}$
22	0,99366	0,952	0,0186	-0,843	9	-0,228	-0,320					$7,07 \cdot 10^{-4}$
23	0,99809	0,777	0,0717	-1,215	8	-0,140	-0,168	-0,097	-0,144			$6,94 \cdot 10^{-4}$
24	0,99613	0,933	0,0261	-0,681	6	-0,328	-0,111					$4,30 \cdot 10^{-4}$
25	0,99577	0,918	0,0338	-0,736	7	-0,326	-0,103	-0,161	0,091	-0,152	*	$5,11 \cdot 10^{-4}$
26	0,99453	0,969	0,0402	-2,766	2	-0,493	-0,169					$2,32 \cdot 10^{-4}$
27	0,99757	0,716	0,0593	-2,371	4	-0,572	-0,177	-0,116	0,002	-0,088	*	$1,12 \cdot 10^{-4}$
28	0,99568	0,941	0,0162	-0,442	7	-0,429	-0,209	0,008	-0,047	-0,101		$2,65 \cdot 10^{-4}$
29	0,99738	0,824	0,0573	-2,501	4	-0,331	-0,242	-0,131	-0,057	-0,101		$1,07 \cdot 10^{-4}$
30	0,99591	0,914	0,0276	-2,085	2	-0,288	-0,200					$7,18 \cdot 10^{-4}$
31	0,99646	0,976	0,00911	-2,433	3	-0,279	-0,236	-0,173	-0,148			$2,92 \cdot 10^{-4}$
32	0,99783	0,958	0,00796	-1,055	10	-0,484	-0,182	-0,177	-0,089			$7,66 \cdot 10^{-5}$
33	0,99917	0,891	0,0247	-5,631	4	-0,353	-0,047	-0,033	-0,065	-0,184		$7,46 \cdot 10^{-5}$
34	0,99895	0,933	0,0199	-2,958	4†	-0,648	-0,148	-0,127	-0,020	-0,002	*	$2,31 \cdot 10^{-4}$
35	0,99818	0,896	0,0281	-2,009	2	-0,268	-0,179	-0,140	-0,107	-0,118		$9,17 \cdot 10^{-5}$
36	0,99803	0,972	0,00429	-0,707	9	-0,650	-0,120					$1,14 \cdot 10^{-4}$
37	0,99491	0,882	0,0237	-0,487	6	-0,504	-0,222	-0,014	-0,181	-0,100	*	$1,23 \cdot 10^{-4}$
38	0,99681	0,976	0,00532	-0,549	7	-0,517	-0,265	0,002	-0,146			$1,67 \cdot 10^{-4}$
39	0,99794	0,964	0,00952	-1,429	1	-0,277	-0,291	-0,070	-0,066	-0,108	*	$4,68 \cdot 10^{-4}$
40	0,99872	0,929	0,00900	-1,593	12	-0,369	-0,268	-0,113	-0,060	-0,139	*	$1,63 \cdot 10^{-4}$
41	0,99872	0,953	0,0102	-3,672	2	-0,161	-0,137	-0,093	-0,085	-0,175		$2,07 \cdot 10^{-4}$
42	0,99511	0,901	0,0398	-0,840	6	-0,464	-0,085	-0,236				$1,87 \cdot 10^{-4}$
43	0,99583	0,957	0,00778	-1,368	10	-0,594	-0,326					$9,78 \cdot 10^{-5}$
44	0,99625	0,964	0,00495	0,070	10	-0,378	-0,141	-0,134	-0,036	-0,131		$9,39 \cdot 10^{-4}$

Table 3.2 : Estimated model parameters taken from 44 real patients. The sampling period is 9,05 seconds for patients 1 – 10 and 10 seconds for the remainder.

\+ fitted by visual inspection of patient data

o inverse logistic relationship used to fit deterministic model. Fit not possible with logistic relat.

* F-test showed a significant reduction in the loss function in fitting 6 parameters to the Noise model

PATIENT	$POLE_1$	$POLE_2$	GAIN	BT4	d_1	d_2	d_3	d_4	d_5	d_6	R_e
1	0,99853	0,815	0,0733	-2,226	-0,462	-0,197	-0,081	-0,053	-0,100		$2,59 \cdot 10^{-2}$
2	0,99770	0,959	0,0165	-1,603	-0,818						$3,58 \cdot 10^{-3}$
3	0,99739	0,965	0,00802	-3,497	-0,436	-0,095	-0,208				$2,76 \cdot 10^{-3}$
4	0,99706	0,944	0,0213	-2,824	-0,963	-0,152	0,252				$2,45 \cdot 10^{-3}$
5	0,99914	0,961	0,0217	-3,724	-0,694	-0,011	-0,166				$9,90 \cdot 10^{-3}$
6	0,99855	0,916	0,0227	-2,391	0,0						$1,64 \cdot 10^{-2}$
7	0,99900	0,964	0,0425	-7,784	-0,461	-0,116	-0,231				$9,87 \cdot 10^{-2}$
8	0,99616	0,984	0,00973	-1,504	-0,623	-0,161	-0,144				$5,21 \cdot 10^{-3}$
9	0,99854	0,749	0,0700	-1,628	-0,676	-0,143	0,106	-0,220			$5,88 \cdot 10^{-3}$
10	0,99751	0,880	0,0680	-2,782	-0,907						$3,56 \cdot 10^{-2}$
11	0,99808	0,896	0,0624	-1,861	-0,257	-0,225					$2,71 \cdot 10^{-2}$
12	0,99717	0,963	0,0155	-0,846	-0,584						$9,09 \cdot 10^{-3}$
13	0,99750	0,956	0,0267	-2,173	-0,229	-0,211	-0,118	-0,219	-0,038	*	$1,29 \cdot 10^{-2}$
14	0,99813	0,910	0,0309	-0,193	-0,289						$5,11 \cdot 10^{-2}$
15	0,99088	0,990	0,00891	-1,584	-0,444	-0,112	-0,175	-0,052	-0,137		$1,84 \cdot 10^{-2}$
16	0,99839	0,800	0,1043	-2,237	-0,455	0,174					$2,15 \cdot 10^{-2}$
17	0,99856	0,864	0,1269	-3,296	-0,558						$4,30 \cdot 10^{-2}$
18	0,99627	0,971	0,0105	-0,816	-0,332						$1,10 \cdot 10^{-1}$
19	0,99580	0,954	0,0277	-1,467	-0,416						$8,02 \cdot 10^{-2}$
20	0,99809	0,967	0,0119	-0,563	-0,256	-0,235					$2,76 \cdot 10^{-2}$
21	0,98603	0,98564	0,0206	-2,683	-0,560						$3,34 \cdot 10^{-2}$
22	0,99340	0,960	0,0182	-1,133	-0,163	-0,289	-0,113	-0,076	-0,039	*	$2,19 \cdot 10^{-2}$
23	0,99792	0,879	0,0410	-1,276	-0,146	-0,211					$3,30 \cdot 10^{-2}$
24	0,99593	0,948	0,0186	-0,444	-0,240	-0,126					$2,10 \cdot 10^{-2}$
25	0,99484	0,944	0,0258	-0,780	-0,322	-0,185					$2,61 \cdot 10^{-2}$
26	0,99477	0,966	0,0381	-2,243	-0,461	-0,186					$8,66 \cdot 10^{-3}$
27	0,99697	0,946	0,00703	-0,657	-0,620	-0,218	-0,105				$5,66 \cdot 10^{-3}$
28	0,99554	0,953	0,0120	-0,280	-0,418	-0,238					$1,32 \cdot 10^{-2}$
29	0,99745	0,786	0,0680	-2,456	-0,386	-0,286	-0,184				$4,54 \cdot 10^{-3}$
30	0,99606	0,903	0,0271	-1,651	-0,347	-0,199					$2,90 \cdot 10^{-2}$
31	0,99641	0,977	0,00831	-2,187	-0,241	-0,219	-0,176	-0,161			$1,25 \cdot 10^{-2}$
32	0,99781	0,963	0,00781	-1,295	-0,456	-0,179	-0,186	-0,127			$3,20 \cdot 10^{-3}$
33	0,99919	0,888	0,0235	-5,084	-0,360	-0,094	-0,070	-0,185			$5,43 \cdot 10^{-3}$
34	0,99907	0,938	0,0134	-1,865	-0,610	-0,172	-0,145				$9,03 \cdot 10^{-3}$
35	0,99826	0,872	0,0277	-1,312	-0,357	-0,193	-0,162	-0,098			$4,79 \cdot 10^{-3}$
36	0,99797	0,975	0,00416	-0,902	-0,566	-0,149					$4,23 \cdot 10^{-3}$
37	0,99467	0,900	0,0223	-0,633	-0,438	-0,221	-0,034	-0,200	-0,112	*	$5,97 \cdot 10^{-3}$
38	0,99633	0,981	0,00452	-0,525	-0,482	-0,294	-0,010	-0,174			$7,16 \cdot 10^{-3}$
39	0,99762	0,969	0,00729	-0,983	-0,256	-0,339	-0,094	-0,068	-0,087	*	$1,83 \cdot 10^{-2}$
40	0,99875	0,948	0,00657	-1,584	-0,324	-0,243	-0,122	-0,101	-0,165	*	$6,74 \cdot 10^{-3}$
41	0,99874	0,943	0,0121	-3,633	-0,166	-0,148	-0,100	-0,100	-0,172		$8,72 \cdot 10^{-3}$
42	0,99499	0,921	0,0320	-0,830	-0,444	-0,115	-0,201				$1,02 \cdot 10^{-2}$
43	0,99556	0,968	0,00626	-1,439	-0,552	-0,362					$5,42 \cdot 10^{-3}$
44	0,99614	0,972	0,00339	0,254	-0,354	-0,132	-0,153	-0,044	-0,148	*	$3,88 \cdot 10^{-2}$

Table 3.3 : Estimated model parameters taken from 44 real patients. The inverse logistic relationship is used. The patient lag is fixed at 5 sampling periods. The sampling period is 9,05 seconds for patients 1 - 10 and 10 seconds for the remainder.

* F-test showed a significant reduction in fitting 6 parameters to the Noise model

CHAPTER 4

Controller Design

4.1 Introduction

The purpose of this chapter is to design a control algorithm capable of inducing muscle relaxation to a desired level and then maintaining it with possible set-point changes. This should be done with a view to the performance criteria given in chapter 1 viz rapid induction rate, minimum overshoot and acceptable regulation.

As stated in section 1.6 the PID controller and the self-tuning controller of Clarke and Gawthrop are obvious candidates for the muscle relaxation problem. It turns out however that these two are unsuitable choices (at least for transient control) when used on positive input systems with time constants of the length encountered when using curare.

A novel dual-mode controller which embodies explicit parameter estimation is consequently developed.

The algorithm that was implemented on the microcomputer was chosen from a number of variants of this type of controller as it was the simplist. These other controllers are evaluated together with it in chapter 6 to determine the degradation in performance in arriving at the version selected for implementation.

4.2 Linearization of the twitch measurement

It was shown in Chapter 3 that the effect response to dTC is non-linear and it is constrained to the range (Ø.1). The resultant pharmacological model thus consists of a linear system followed by a non-linear element. However the controller design can be simplified by linearizing the twitch measurement using the relationship

$$T_{lin}(k) = \ln\left[\frac{1 - T(k)}{T(k)}\right]$$

which is independant of the model parameters. The set-point is similarily transformed

$$SP_{lin}(k) = \ln\left[\frac{1 - SP(k)}{SP(k)}\right]$$

where SP is the set-point on the range (Ø,1). A comparison with the inverse logistic relationship, equation 3.9, on which these transformations are based shows that an off-set BT4 is introduced into the linearized twitch data. However in all other respects the controller design may proceed as for a linear system and is consequently simpler than that of the otherwise non-linear controller. In the ensuing discussion linear models will be considered exclusively. A block diagram of the control loop which shows the above transformations is given in figure 4.1.

It was shown in section 3.8.2 that use of the logistic relationship or its inverse resulted in nearly the same loss function between the twitch data and model output. The use of the inverse relationship here then should not adversely affect the parameters of a patient model estimated on-line within the controller.

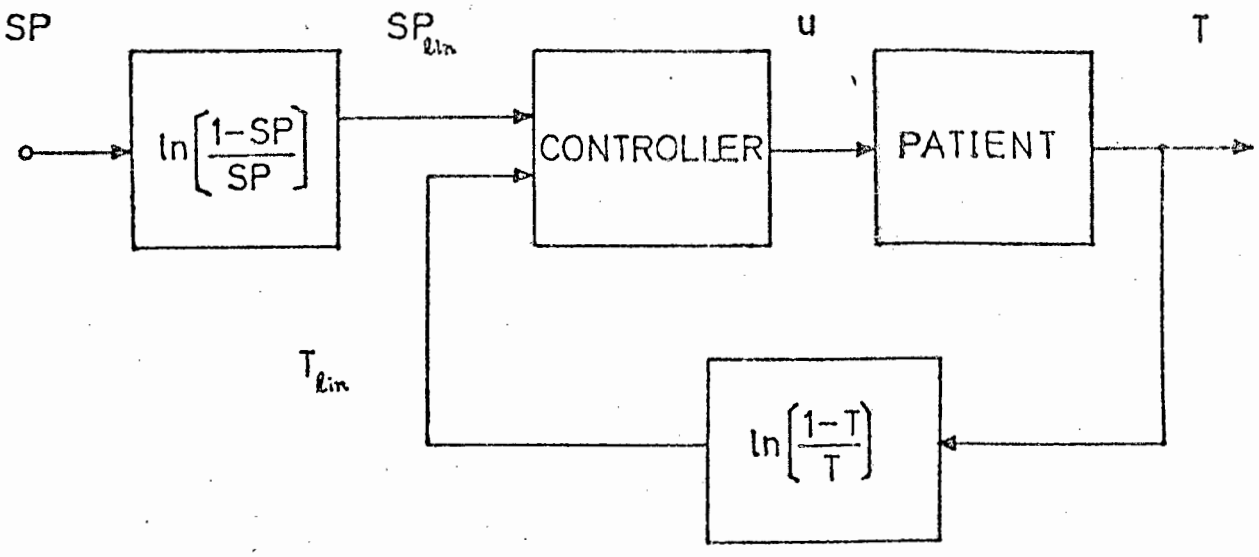

Figure 4.1 : Block diagram of the control loop which shows the linearization of the twitch measurement and set-point.

4.3 Investigation of the PID and Clarke/Gawthrop controllers

Owing to its simplicity, the PID controller was the first candidate investigated. The discrete version considered here as well as the Ziegler-Nichols transient response method that was used to tune it are described in many texts, see for example Astrom and Wittenmark,(1984). The simplest form which does not filter the derivative part of the controller was selected and the resulting algorithm is described in Appendix A. A second PID controller which disallows negative commands but which is otherwise identical is also described there.

The self-tuning k-step ahead optimal controller of Clarke and Gawthrop,(1975) is an implicit algorithm which tunes the parameters of the controller directly. The authors point out that it has certain advantages over the earlier self-tuning regulator of Astrom and Wittenmark,(1973). Two of these are that weighting of control is allowed for and set-point variation may be

optimally followed. A FORTRAN programme to implement this algorithm and another to implement a modified version which is constrained to disallow negative commands but is otherwise identical are given in Appendix A.

The PID and Clarke/Gawthrop controllers were investigated to determine how each would cope with the variability in patient response and the constraint that controllers may not issue negative drug commands. The added complication of the unknown patient dead-time was therefore ignored. Instead it was fixed at the known value of one sampling period in all the simulated patients for the purposes of this investigation.

Both the PID and Clarke/Gawthrop controllers are however sensitive to the process dead-time. Optimal controllers loose control when the dead-time exceeds the value assumed in the controller (Wellstead et al, 1979). Appreciable process dead-time will cause deterioration in PID control and Linkens et al.,(1982) used a Smith-predictor (Astrom and Wittenmark,1984) to compensate for this. (In a separate investigation on controlling simulated patients it was found that instability of the PID controller was likely if the patient's dead-time exceeded 4 sampling periods).

The four controllers described above were evaluated by simulation on the batch of original patients 1 - 1Ø. The parameters of patient 9 who had an average sensitivity to dTC (i.e. required an average bolus dose per unit body mass to depress the twitch ratio by 8Ø%) were selected from this group to initialize the self-tuning controllers and to tune the PID algorithms. In the latter case the Ziegler-Nichols method for systems with appreciable dead-time was used and the three constants so derived were:

$K = Ø,833$ (proportional) ; $T_I = 97,4$ (integral) ; $T_D = 24,4$ (derivative)

The operation of the PID controller does not require the parameter BT4 to be

known. The Clarke and Gawthrop controller makes provision for BT4 by tuning an off-set parameter.

The regulation ability of the PID controllers was first investigated as this is the least demanding type of control. The average patient number 9 was simulated and relaxation induced to a set-point of T4 ratio = Ø,2 by administering the appropriate bolus dose. When this had taken maximum effect, the PID controllers were switched on. In noise free simulations both functioned well. However the inability of the constrained version to regulate efficiently became apparant when the noise sourse was included in the patient model whereas the performance of the other did not change significantly. The graphs in figure 4.2 show the performance of the constrained PID controller for the noise-free and noisy simulations of patient 9.

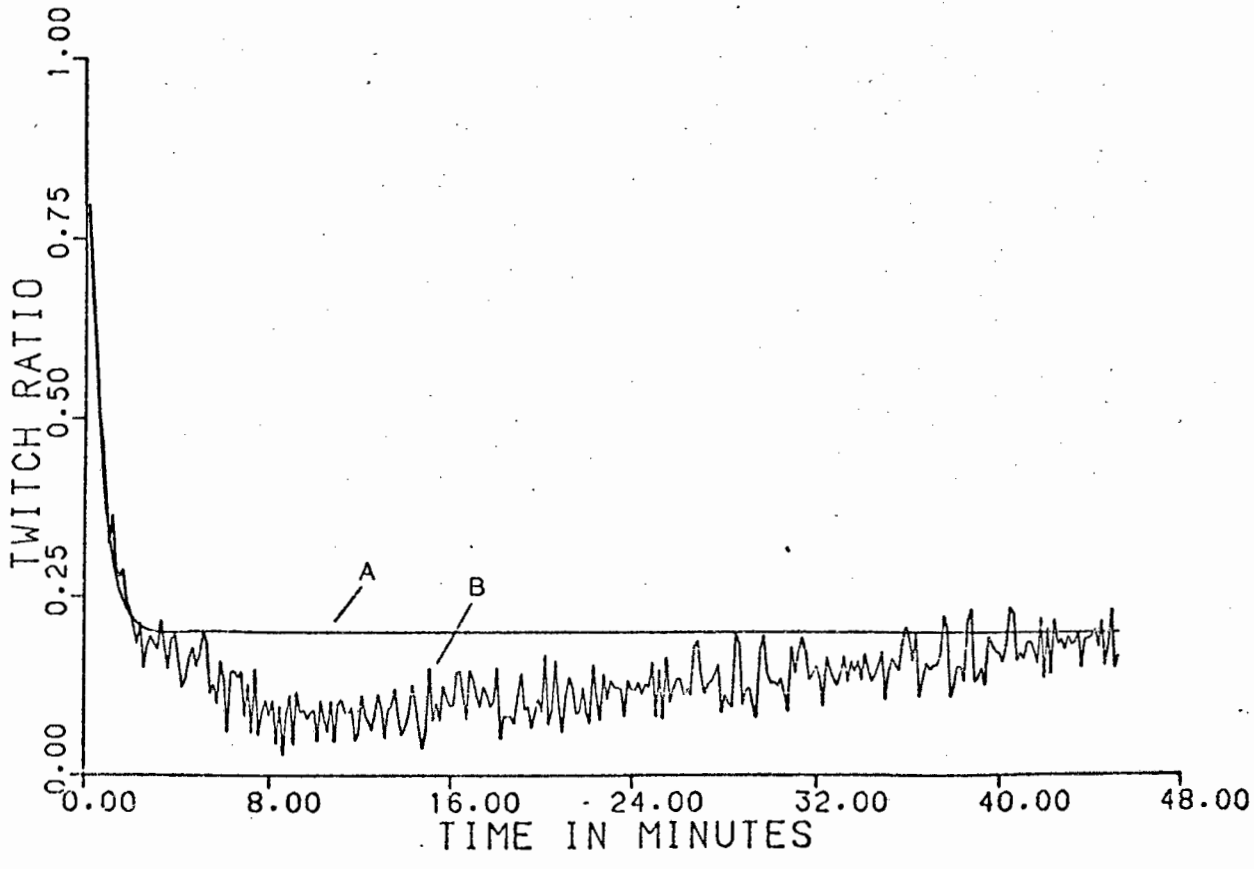

Figure 4.2 : PID controller constrained to disallow negative commands and used on a simulated average patient in the regulation phase. The required level of relaxation was induced by a calculated initial bolus dose injection. Regulation is good in the noise free case (curve A) but deteriorates quite significantly when noise is present (curve B).

The unconstrained controller was then tuned to the average patient as above but required to regulate all ten. The average setting was found to be suitable in each case.

Control during the induction period was investigated next. In all cases a loading dose of 6 mg dTC was administered and the controller switched on once this had taken maximum effect. This dose was considered to be reasonable because the most sensitive patient of the group required 7,2 mg to induce a T4 ratio of Ø,2. The loading dose was given to assist in overcoming the initial dead-time in the patient response.

The PID controllers tuned to the average patient were made to control the most sensitive patient no. 4, the average no. 9 and the most insensitive no. 3. The resulting control in these cases is shown in figure 4.3. The effect of de-tuning the unconstrained controller is apparant. The other is unreliable. In particular its failure to control the average patient - to whom the algorithm had been exactly tuned - suggests that simple PID control will not be suitable for inducing relaxation in real patients.

The above investigation into control during the induction period was repeated on the two self-tuning algorithms. The corresponding results are shown in figure 4.4. The beneficial effect of self-tuning is obvious in the case of the unconstrained controller but the constrained version gives unacceptable results.

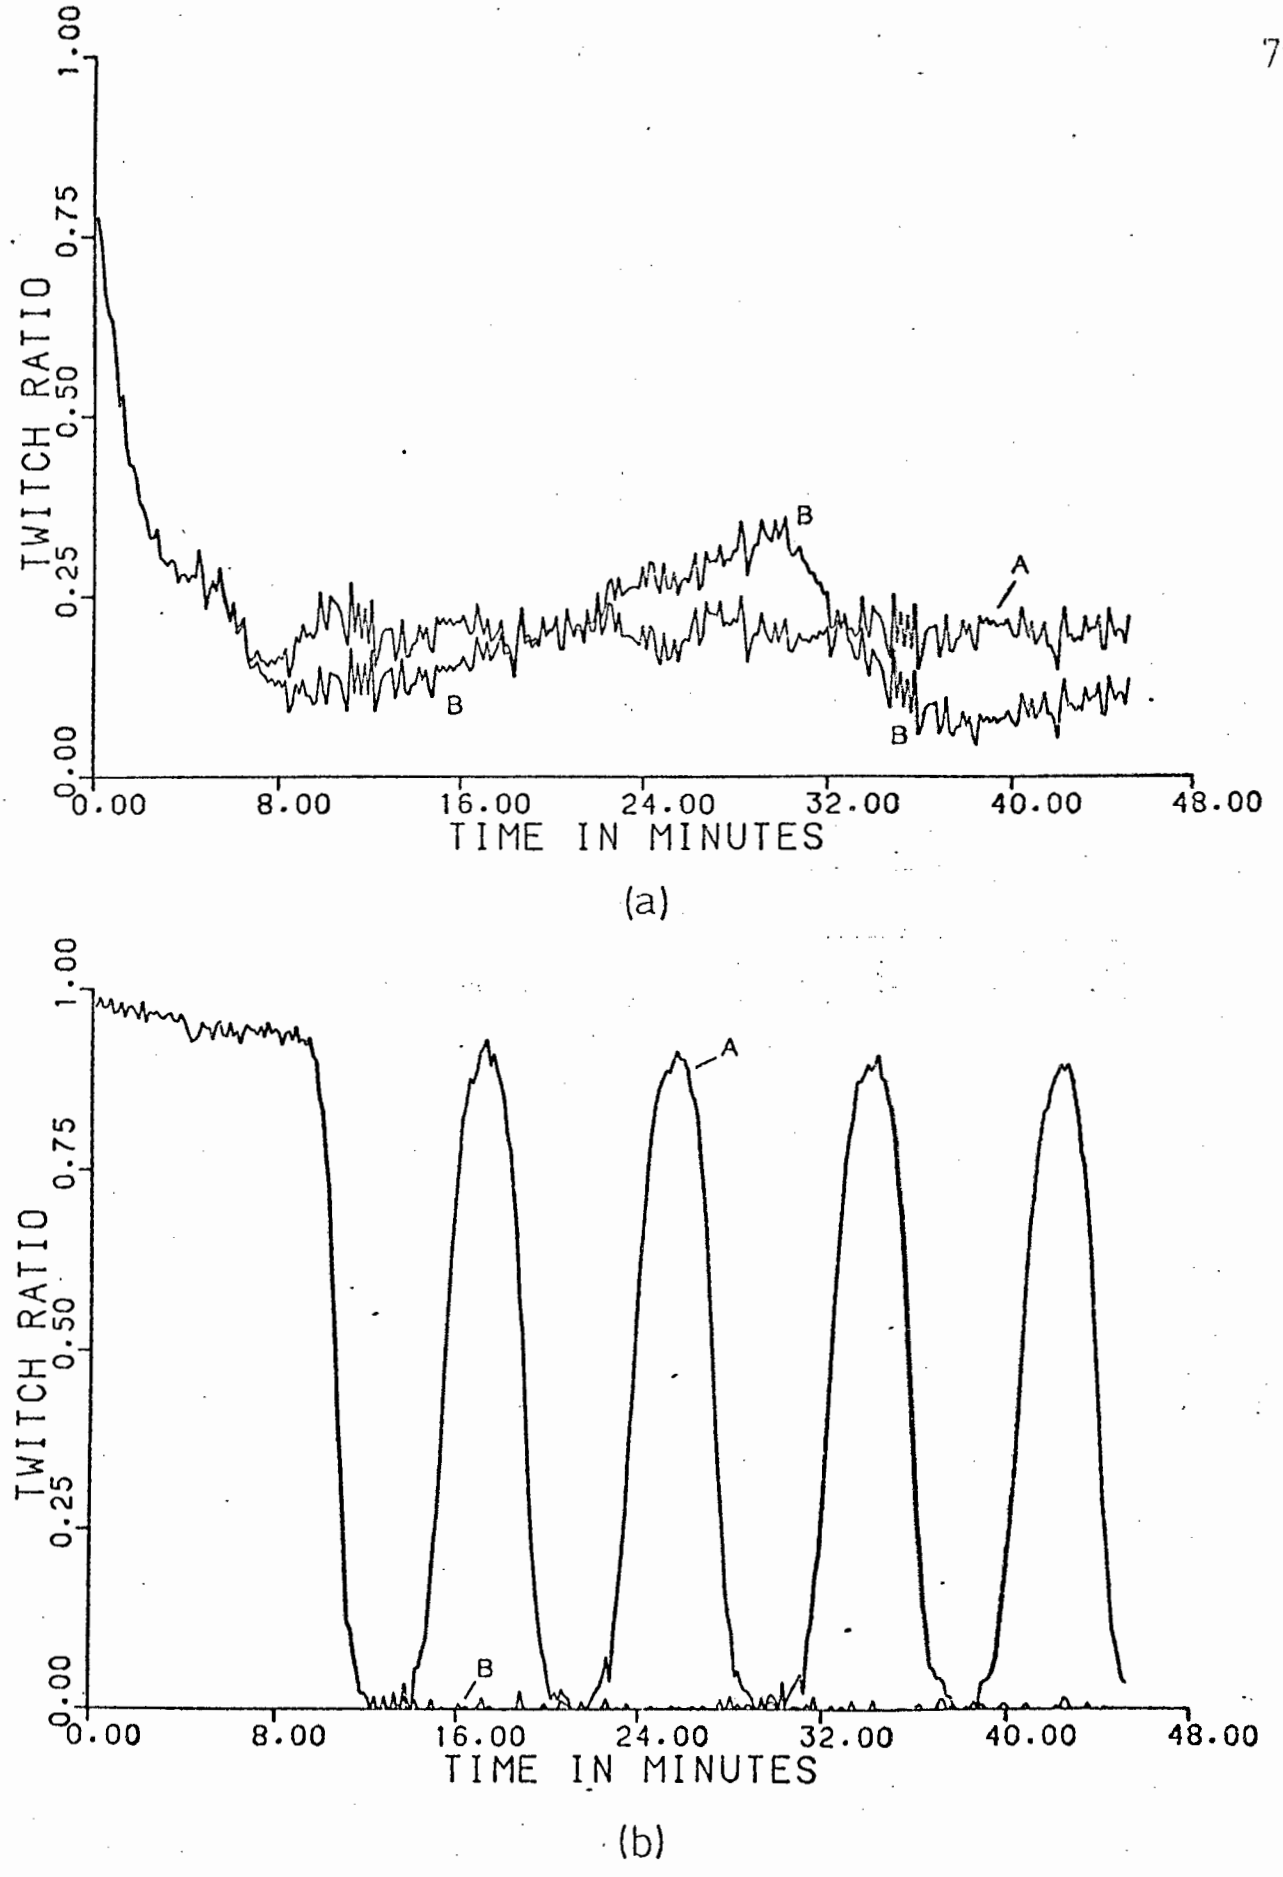

Figure 4.3 : PID controller used on three simulated patients. Curve A is the unmodified version; in curve B the controller is constrained to disallow negative drug commands. Relaxation is induced initially by administering a 6 mg bolus injection of dTC and the controllers are switched on once this has taken maximum effect.
(a) sensitive patient (no. 4), (b) insensitive patient (no. 3), (c) average patient (no. 9).

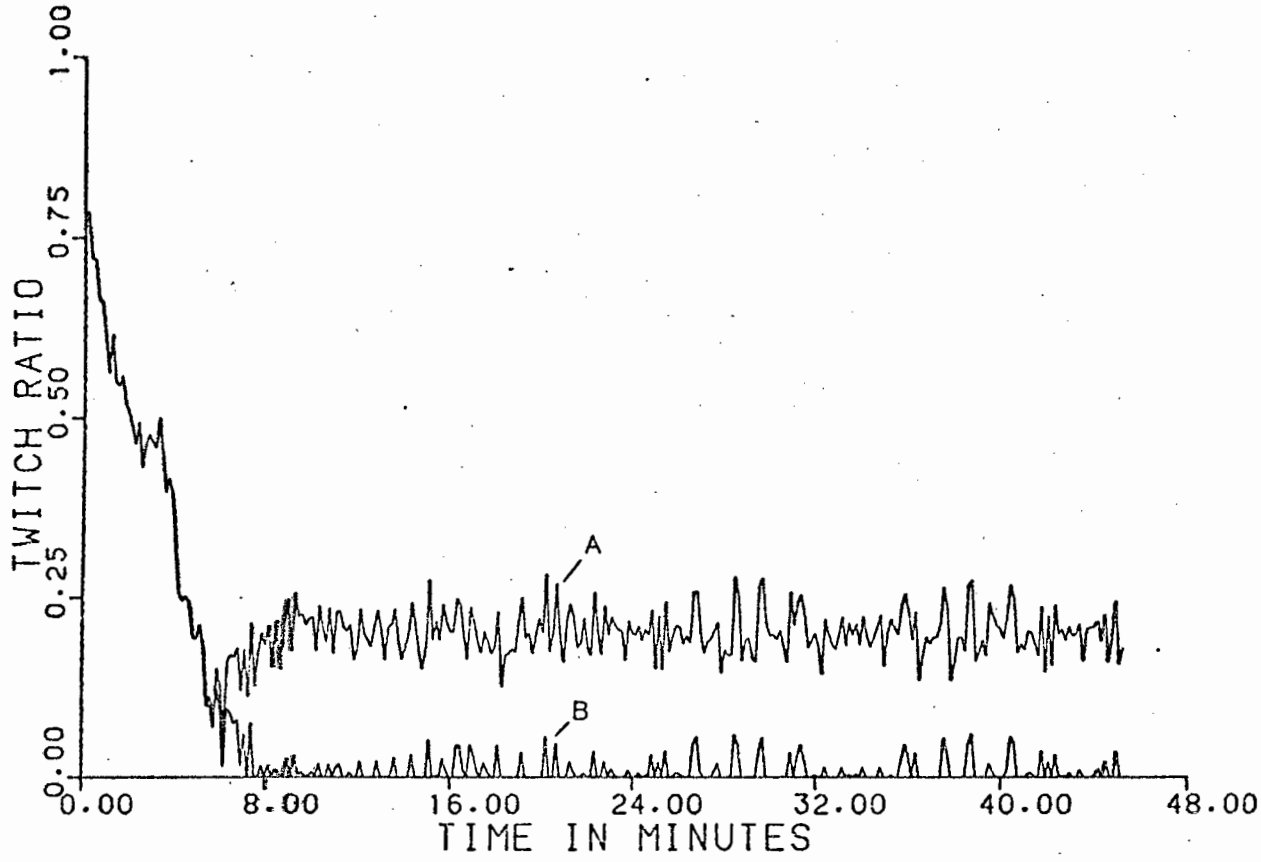

(c)

figure 4.3 (contd.)

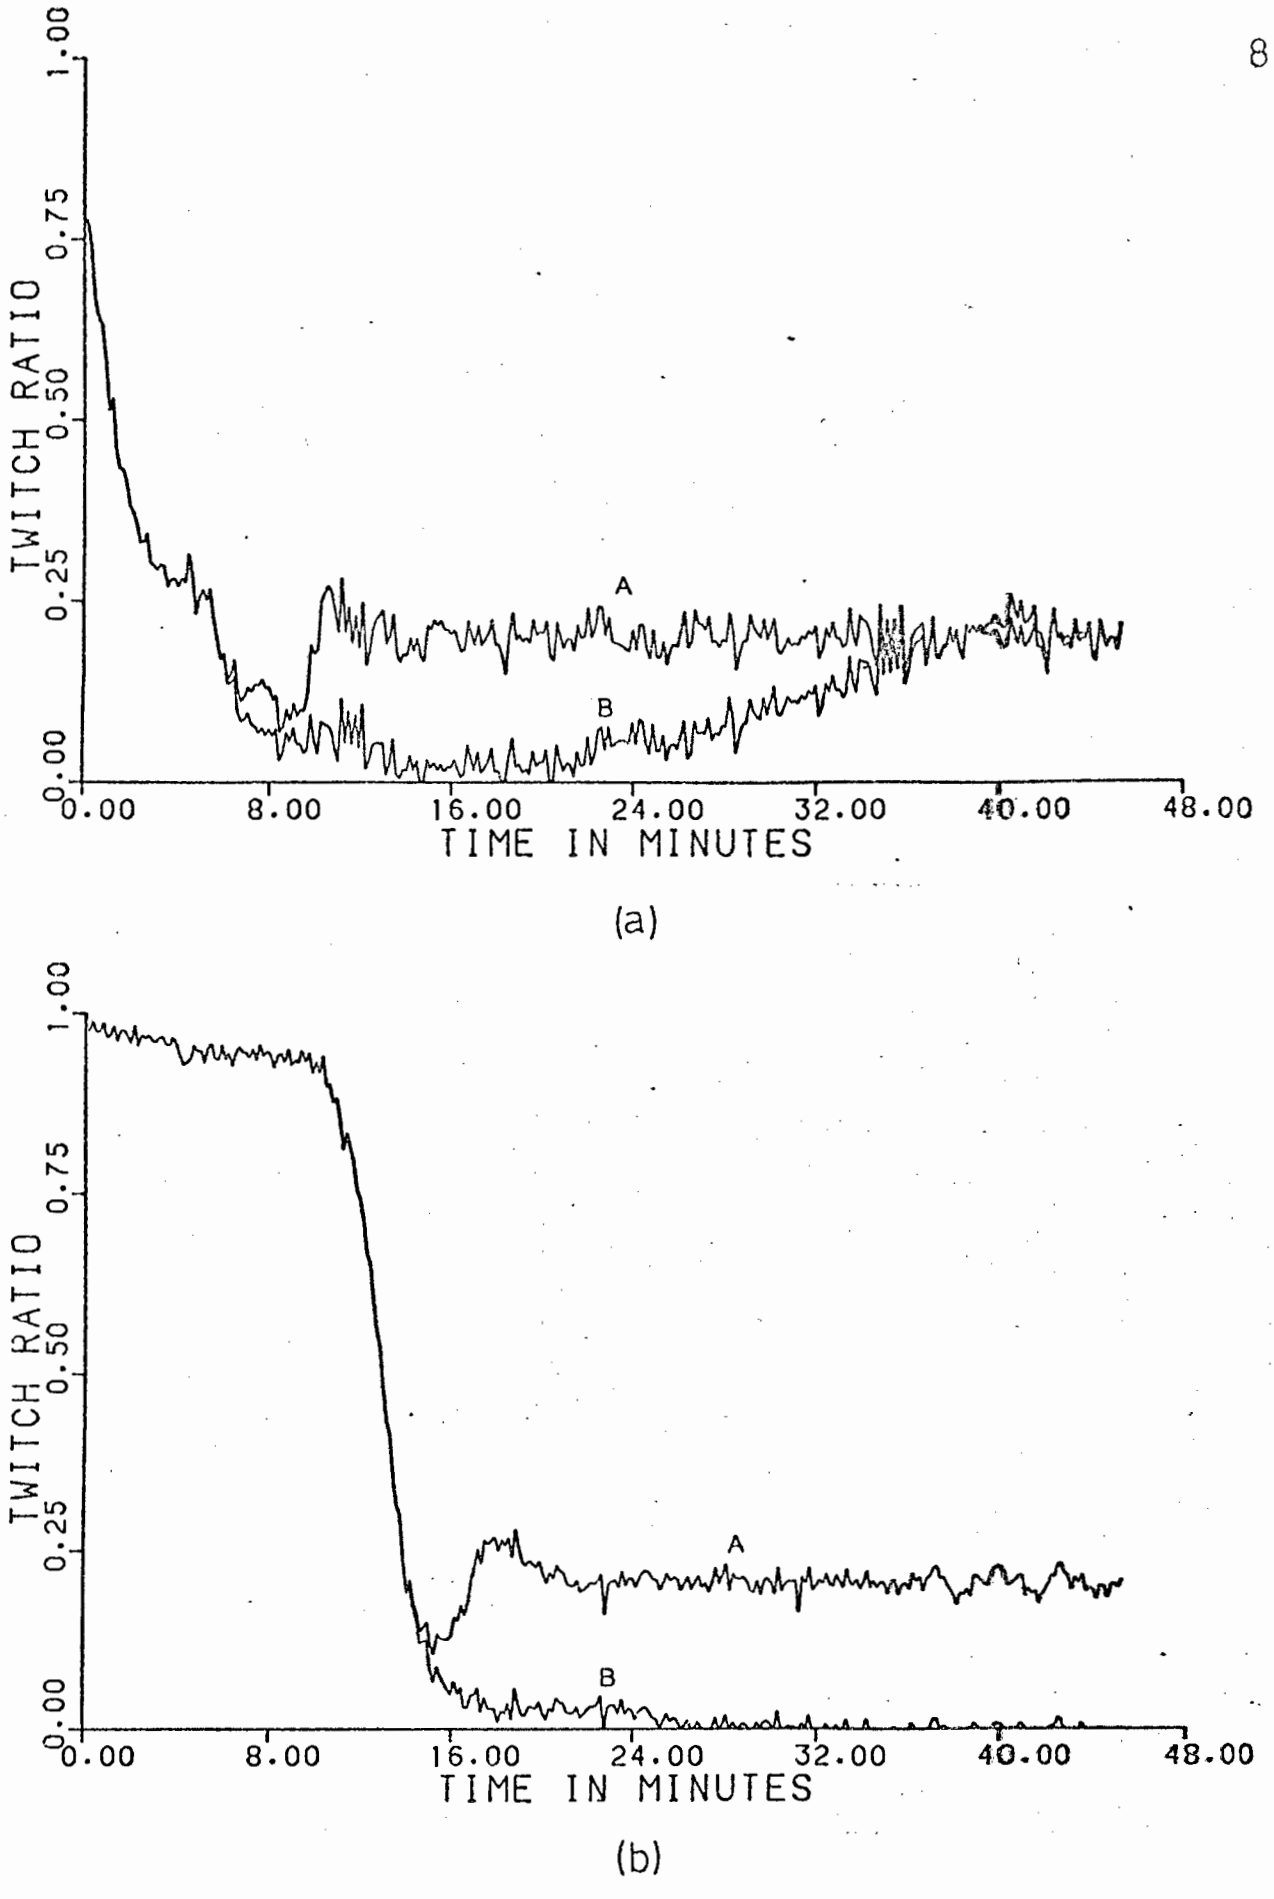

Figure 4.4 : Clarke and Gawthrop self-tuning controller used on three simulated patients. Curve A is the unmodified version; in curve B the controller is constrained to disallow negative drug commands. Relaxation is induced initially by administering a 6 mg bolus injection of dTC and the controllers are switched on once this has taken maximum effect. (a) sensitive patient (no. 4), (b) insensitive patient (no. 3), (c) average patient (no. 9).

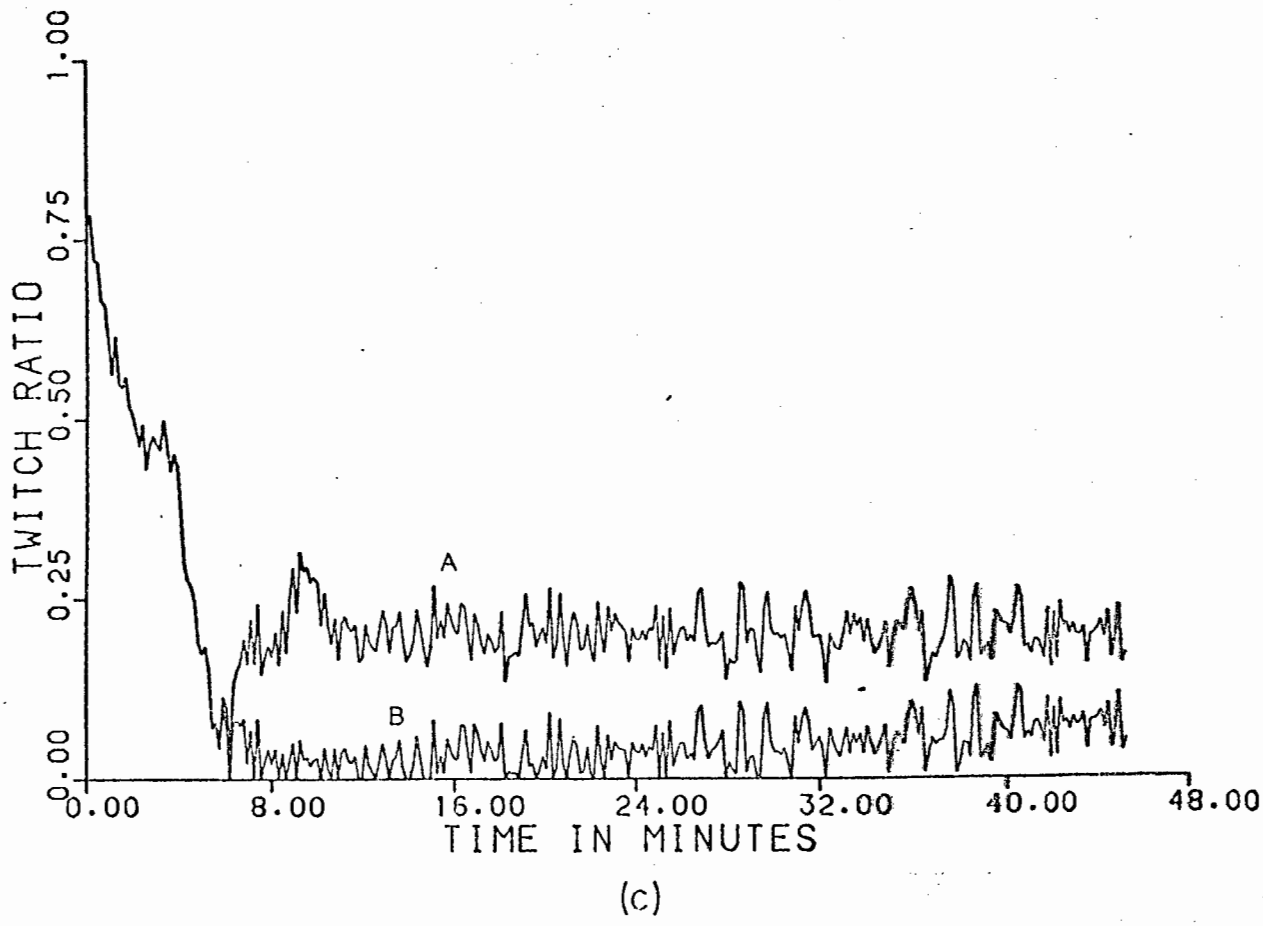

(c)

figure 4.4 (contd.)

4.4 Patient variability

It was noted in section 4.3 that the self-tuning algorithm of Clarke and Gawthrop was not influenced by variability in the patient response to the extent that the fixed parameter PID controller was. Linkens et al.,(1982) have advocated self-tuning control for the non-depolarizing relaxant pancuronium whose duration of action is somewhat shorter than dTC, (Bevan,1983). Bradlow et al.,(1983) investigated the suitability of a fixed parameter Dahlin's controller and concluded that the variability in patient response was large enough to render it useless for control of dTC induced muscle relaxation.

It is convenient to interpret the variability in the patient models as physically determinable parameters but without the need for applying closed loop control. This is done by exciting the patient model with a single bolus dose which is calculated to depress the twitch ratio by 80% in each case. The three measures of interest to the induction phase are:

1. the sensitivity of the patient (dose per unit body mass)

2. the initial dead time (time for twitch ratio to drop to Ø,85)

3. the time to reach maximum induced blocking effect (dose independant).

These measures were calculated by a programme described in Appendix J for each of the 44 patients whose data was available and are listed in Table 4.1. In the table the T4 and T1 ratio models are treated separately. The variation in patient sensitivity is 5,2:1 and that in time to maximum effect is 7,3:1. The greatest initial dead time is 4,5 minutes. (The dead time quoted in the table is that in response to the optimum loading dose. A more conservative loading dose, as would be administered in closed loop control, would lengthen this period. It will be noticed that several of the patients have zero initial dead time. This is a model inaccuracy as models that have parameter BT4 less than -1,73 cannot attain a twitch ratio of above Ø,85 as an inspection of the logistic relationship equation 3.7 shows.)

Admittedly the entries of this table pertain to patients who were relaxed under different clinical conditions, but, as this is likely to be the rule in practice, this fact should be regarded as adding to the patient variability.

PATIENT	SENSITIVITY (μgdTC/Kg to TWITCH RATIO OF 20%)	MASS (Kg)	TIME TO MAX. EFFECT (MINS)	INITIAL DEAD-TIME (MINS)
1	123,4	74	4,1	0,5
2	130,3	65	7,8	0,5
3	315,9	84	9,1	1,2
4	102,3	70	5,0	0,0
5	203,5	49	17,3	2,0
6	179,4	83	6,3	0,2
7	166,2	48	19,8	4,5
8	123,9	60	19,3	0,0
9	157,2	73	3,5	0,0
10	112,9	70	3,5	0,2
11	94,2	60	4,7	0,0
12	111,8	60	10,8	0,0
13	136,5	50	11,3	0,5
14	91,3	60	4,3	0,0
15	168,6	50	16,2	0,0
16	120,6	60	4,5	0,2
17	94,6	55	4,7	0,2
18	133,0	60	12,3	0,0
19	103,0	60	9,3	0,0
20	107,6	60	16,7	0,0
21	125,5	60	11,3	1,2
22	110,9	70	7,8	0,0
23	139,8	60	3,2	0,0
24	104,7	60	7,2	0,0
25	109,2	55	6,0	0,0
26	92,4	50	11,0	1,0
MEAN	133,0	62	9,1	
STD. DEVN.	47,1	10	5,2	
27	155,2	120	2,3	0,0
28	121,4	67	7,7	0,0
29	194,6	65	3,7	0,2
30	208,5	60	5,8	0,2
31	232,6	60	15,3	0,8
32	274,3	55	12,0	0,0
33	427,7	75	7,0	1,0
34	362,7	43	10,0	0,7
35	176,9	76	6,2	0,0
36	208,1	80	16,7	0,0
37	201,1	73	3,2	0,0
38	182,2	65	15,8	0,0
39	233,8	54	13,7	0,0
40	504,5	50	9,2	0,0
41	559,0	46	12,7	1,3
42	106,9	60	5,0	0,0
43	299,1	65	9,7	0,0
44	165,2	75	11,3	0,0
MEAN	256,3	66	9,3	
STD. DEVN.	128,2	17	4,5	

Table 4.1 : Variability in the pharmacological response of patients estimated from their clinical trials. For patients 1 - 26 the T4 ratio was modelled. For the remainder the T1 ratio was modelled.

4.5 Controller design

Both the PID and Clarke/Gawthrop controllers failed in the relaxation phase when constrained to give non-negative drug commands. It was therefore considered more reasonable to approach the muscle relaxation control problem afresh than attempt to modify existing algorithms to work with control commands of a single polarity. The development of the resulting algorithm is described here. It is a dual-mode controller as different strategies are used in the initial phase and in the relaxation and regulation phases that follow. The structure of the controller is shown in figure 4.5.

The Initial Phase Controller is introduced to deal with the non-linearity which the pharamacodynamic relationship imposes and the consequent considerable period of initial dead-time exhibited by certain patients as was noted in Table 4.1. It is designed to administer a sequence of bolus doses up until onset of change in blocking effect. These are calculated to minimize the possibility of overshoot of the set-point. It is programmed to operate on a conservative estimate of the parameter vector which defines the patient model, $\underline{\theta}_E$.

The controller used in the relaxation and regulation phases predicts the future levels of patient paralysis from the doses already administered and locates the maximum. Then it calculates by a numerical method the present dose which will take that maximum to the set-point. The estimates of the true patient state $\underline{X}_T$ and model parameter vector $\underline{\theta}_D$ are required for this. These estimates are designated $\hat{\underline{X}}_T$ and $\hat{\underline{\theta}}_D$.

$$\underline{X}_T(k) = \left[x_T(k) \ , \ x_T(k-1)\right]^T$$

where $x_T(k)$ is the true state and can be thought of as the "smoothed" value of $T_{1in}(k)$ - BT4.

$$\underline{\theta}_D = \left[p_1 \ , p_2 \ , \ g \ , \ BT4\right]^T$$

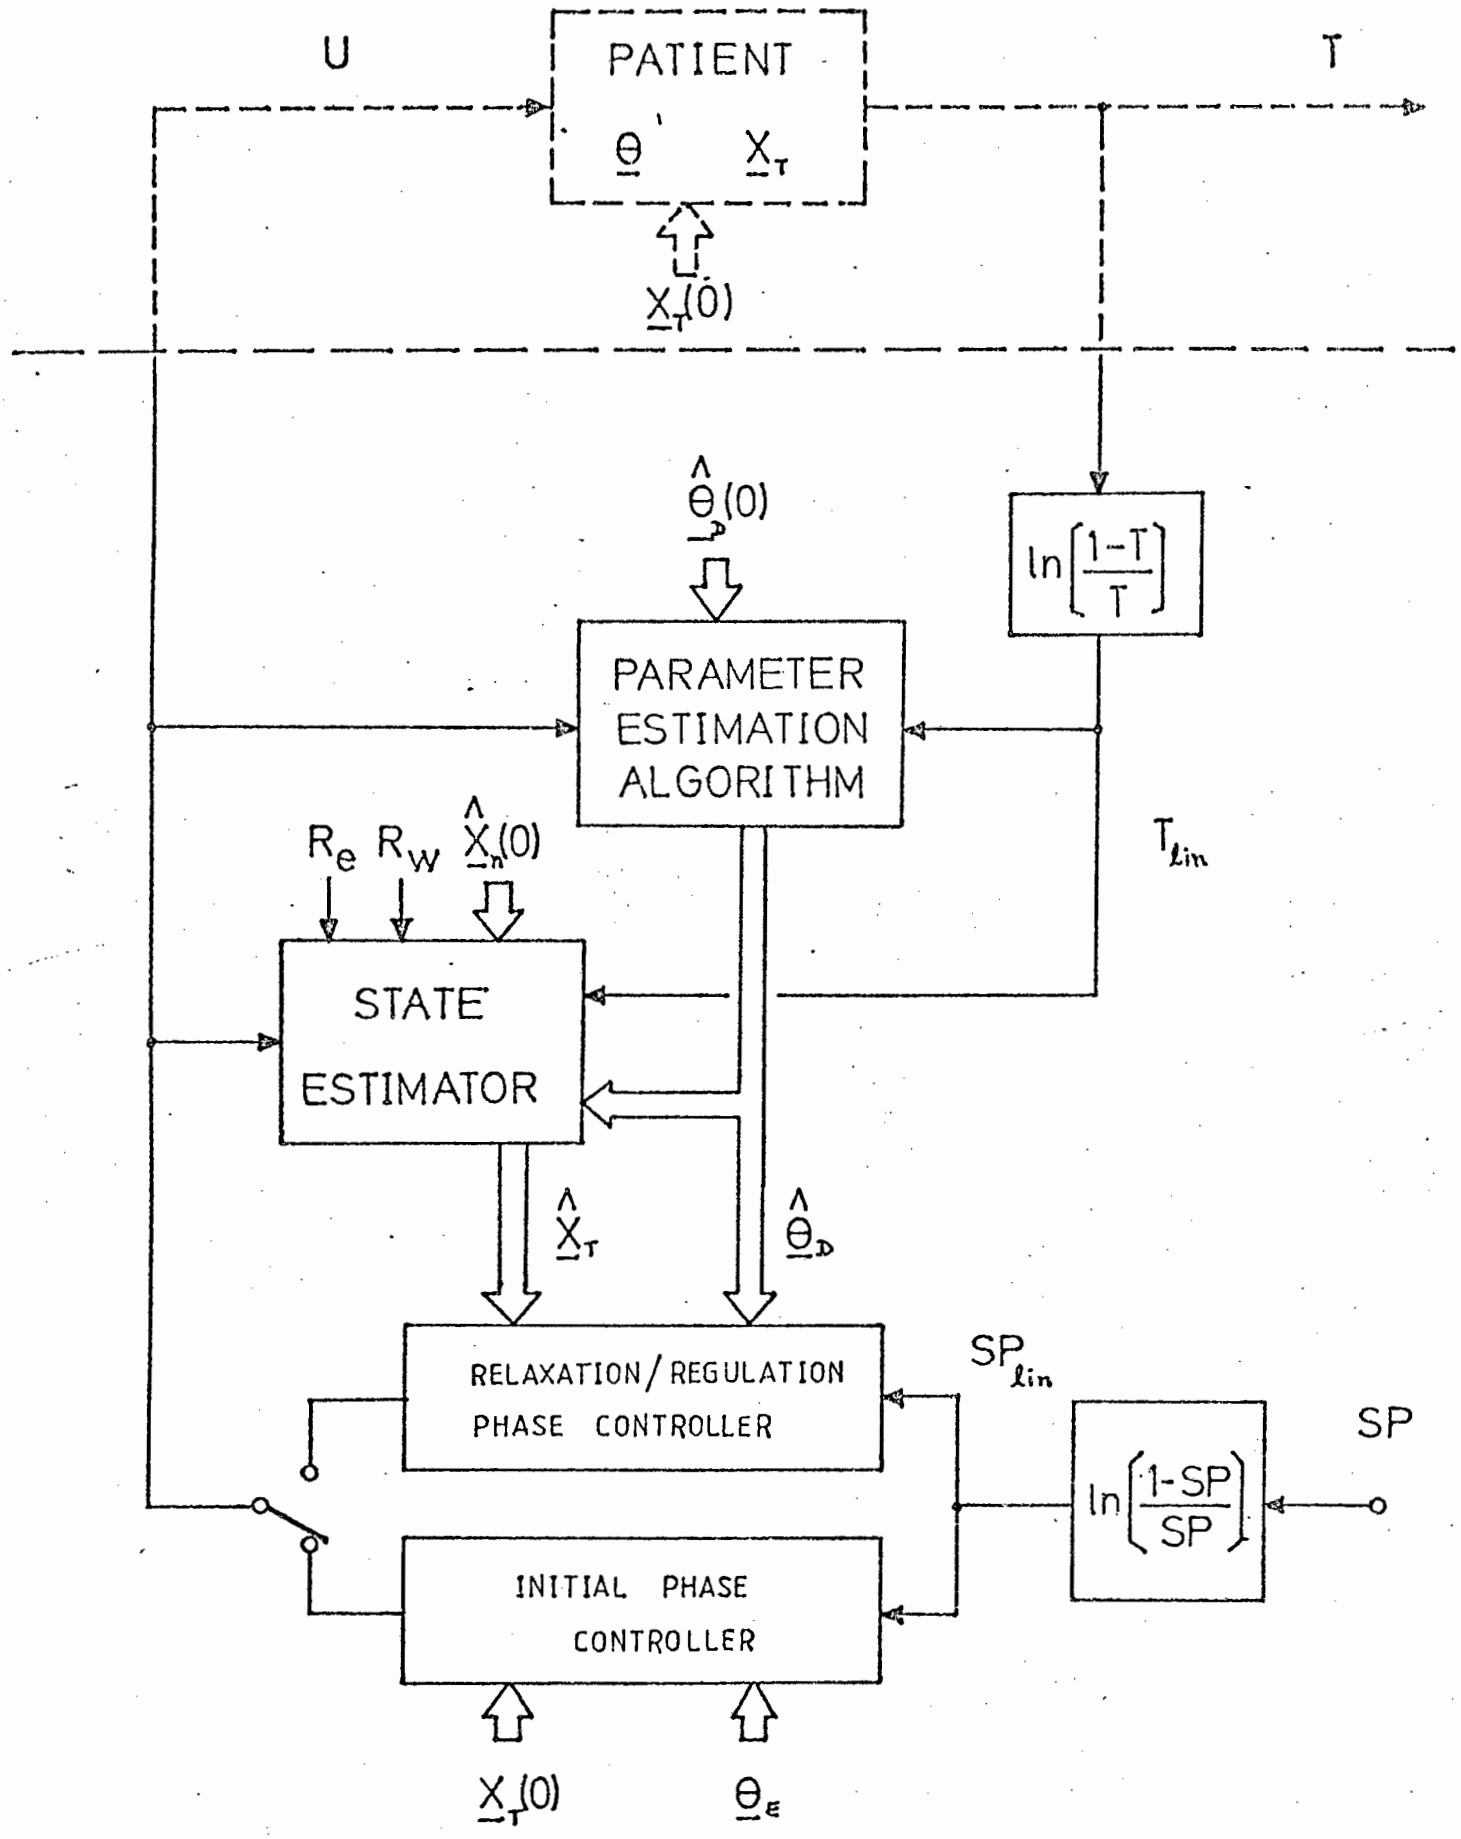

Figure 4.5 : Structure of the on-line controller. The transition from the initial phase to the relaxation phase is made when the twitch ratio being controlled drops below Ø,85.

Three alternative implementations of the controller of figure 4.5 have been investigated: (The name in brackets refers to the FORTRAN programme which implements the algorithm and is given in Appendix A).

1. (CONTRL.BOLUS4) - simple linear State Estimator; Least Squares parameter estimation using Steepest Descent with numerical approximation of the gradient vector. This is the simplist of the three controllers and is the one finally implemented on the microcomputer.

2. (CONTRL.BOLUS) - non-linear state estimation by a Kalman filter; Least Squares parameter estimation by Steepest Descent with exact calculation of the gradient vector. This version is compared with the controller in 1. to justify the simplifications made there.

3. (CONTRL.BOLUS5) - simple linear State Estimator; Least Squares parameter estimation by subroutine VAØ5A of the Harwell library. This is used to compare the sophisticated parameter estimation algorithm with the simpler method of Steepest Descent.

4.5.1 The Initial Phase Controller

In the case of patients who exhibit a considerable initial dead-time, it is not obvious how to proceed administering the relaxant with safety during this period. The Initial Phase Controller is designed to induce muscle relaxation as rapidly as possible even in the cases where the initial dead-time is large. At the same time it minimizes the possibility of overshoot of the set-point due to an overly large dose.

The algorithm uses a patient model parameter vector $\underline{\theta}_E$ specially designed to cater for the extremes in response anticipated. From a given group of patient models, the poles are chosen to be those of the patient with the longest time to reach maximum effect in response to a bolus dose (this time is dependent upon the poles in the model only). Parameter BT4 is chosen as the largest in the group for the reason to be explained below. With these three parameters fixed, the model sensitivity is determined by the remaining parameter, 'g'. This is selected so that the resulting model has a sensitivity equal to that of the most sensitive patient model in the group.

The first dose commanded by the controller is calculated to take the conservative patient model just to the chosen set-point. It is clear that this should not overdose any real patient. The calculated response of the conservative model is curve A in figure 4.6. It can be seen from the figure

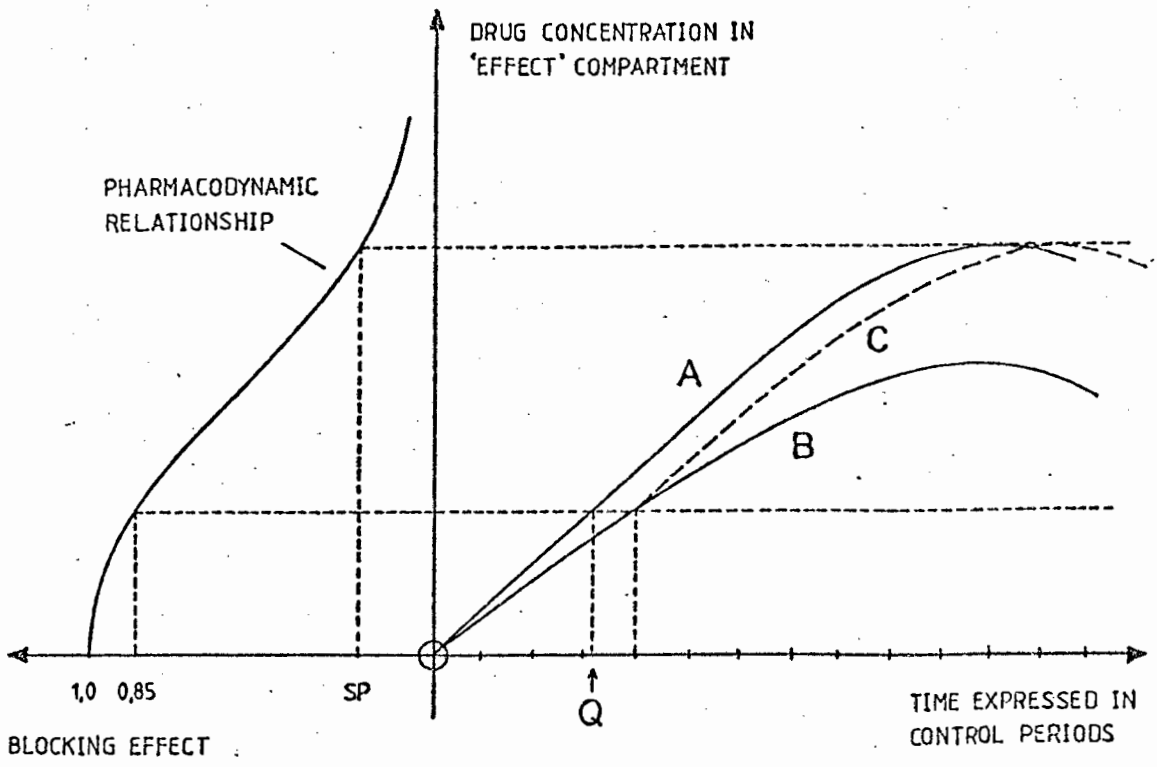

Figure 4.6 : Drug concentration in the effect compartment in response to the initial injection sequence which is calculated as described in the text.

that the response takes a time Q to reach the threshold which defines the initial dead-time. The threshold was set at a twitch ratio of Ø,85 because this was found to be a reliable indication of the onset of block from the clinical trials undertaken in the course of the research. The Initial Phase Controller issues no further drug command during this time. If the twitch ratio of the patient being controlled has not dropped below the threshold at the control period following Q, one can reasonably conclude that the patient is less sensitive than the model defined by $\underline{\theta}_E$. This is so because the poles in the model were set to the slowest patient response while the largest value of BT4 was selected. (This has the effect of lifting the threshold in figure 4.6 to the greatest level possible hence increasing the crossing time as much as possible). In fact at this time instant one can conclude that the patient being controlled has a maximum sensitivity defined by curve B and so redefine the gain parameter in $\underline{\theta}_E$ accordingly. It is now possible to calculate a second dose to be administered at this instant, based on the new $\underline{\theta}_E$, which will just take the model to the set-point (along curve C). This process is repeated at subsequent control periods until the patient response crosses the threshold. The calculation of the initial injection sequence is then terminated and the Relaxation/Regulation Phase Controller switched on. The implementation of this algorithm is described in Chapter 5.

It is worth noting that whilst the initial dead-time of the patient model is attributable to the non-linearity in the pharmacodynamic relationship, the dead-time exhibited by real patients is known to be caused by the approximately 75% surplus receptor sites on the muscle surface which need to be occupied by molecules of the relaxant before relaxation can commence (Ward, 1975).

4.5.2 Control algorithm for relaxation and regulation phases

In these phases of the control it is necessary to bring the patient to a desired level of paralysis as quickly as possible and without significant overshoot then to maintain that level within a sufficiently narrow band about the set-point.

The control algorithm is best described using the state formulation of the patient. The deterministic patient model equation 3.11 can be written as

$$\underline{X}(k+\ell) = \Phi\, \underline{X}(k+\ell-1) + \Gamma \cdot u(k) \quad \ldots \quad 4.1$$

$$x(k) = H\, \underline{X}(k)$$

where

$$\Phi = \begin{bmatrix} (p_1 + p_2) & -p_1 p_2 \\ 1 & 0 \end{bmatrix}$$

$$\Gamma = \begin{bmatrix} g \, , & 0 \end{bmatrix}^T$$

$$H = \begin{bmatrix} 1 \, , & 0 \end{bmatrix}$$

and

$$\underline{X}(k) = \begin{bmatrix} x(k) \, , & x(k-1) \end{bmatrix}^T$$

If $\underline{X}$ equals the true patient state vector $\underline{X}_T$ and if the model parameters are derived from the true parameter vector $\underline{\theta}_D$, the true patient state n steps into the future and with no subsequent doses administered can be calculated on entry to the relaxation phase by

$$\underline{X}(k+n) = \Phi^n\, \underline{X}(k)$$

The point at which this predicted level of response to a given dose is a maximum can be found by examining $\underline{X}(k+n)$ for each successive 'n' until the sampling period is found for which

$$x(k+n) \leqslant x(k+n-1)$$

At the termination of initial phase control the resulting maximum level of relaxation $x_{max} = x(k+n-1)$ is expected to be less than or equal to the desired set-point SP_{des} calculated by equation 3.9

$$SP_{des} = \ln\left[\frac{1 - SP}{SP}\right] - BT4$$

as conservative doses are administered during the initial phase. It follows that there will be a unique bolus dose u(k) which will just take the induced level of relaxation to the set-point. It can easily be found by performing a numerical search. The proceedure is shown in the flowchart of figure 4.7. The problem in practice is that neither the true patient state $\underline{X}_T(k)$ nor the true vector of model parameters $\underline{\theta}_D$ are known. They are consequently replaced by their estimates $\hat{\underline{X}}_T(k)$ and $\hat{\underline{\theta}}_D(k)$ and these are supplied by a State Estimator and parameter estimation routine respectively, which are described in the following two sections.

Use of the estimated vectors $\hat{\underline{X}}_T(k)$ and $\hat{\underline{\theta}}(k)$ in lieu of the actual values has a danger associated with it when the process is a positive input system because the errors passed on to the control algorithm may result in drug commands which are too large and these cannot be counteracted by subsequent "negative" doses. This will inevitably lead to overshoot of the set-point. This is particularily annoying for systems which have long time constants.

To minimize the possibility of overshoot in the relaxation phase, a trial set-point, SP_t is introduced and brought to the desired set-point SP_{des} in a number of steps from the predicted maximum level of relaxation induced by the initial injection sequence. Thus the controller attempts to make the transition to the desired set-point in a reasonable time period and some overshoot of the intermediate set-points can be tolerated while the parameter estimates converge. Specifically the trial set-point is calculated by equation 4.2

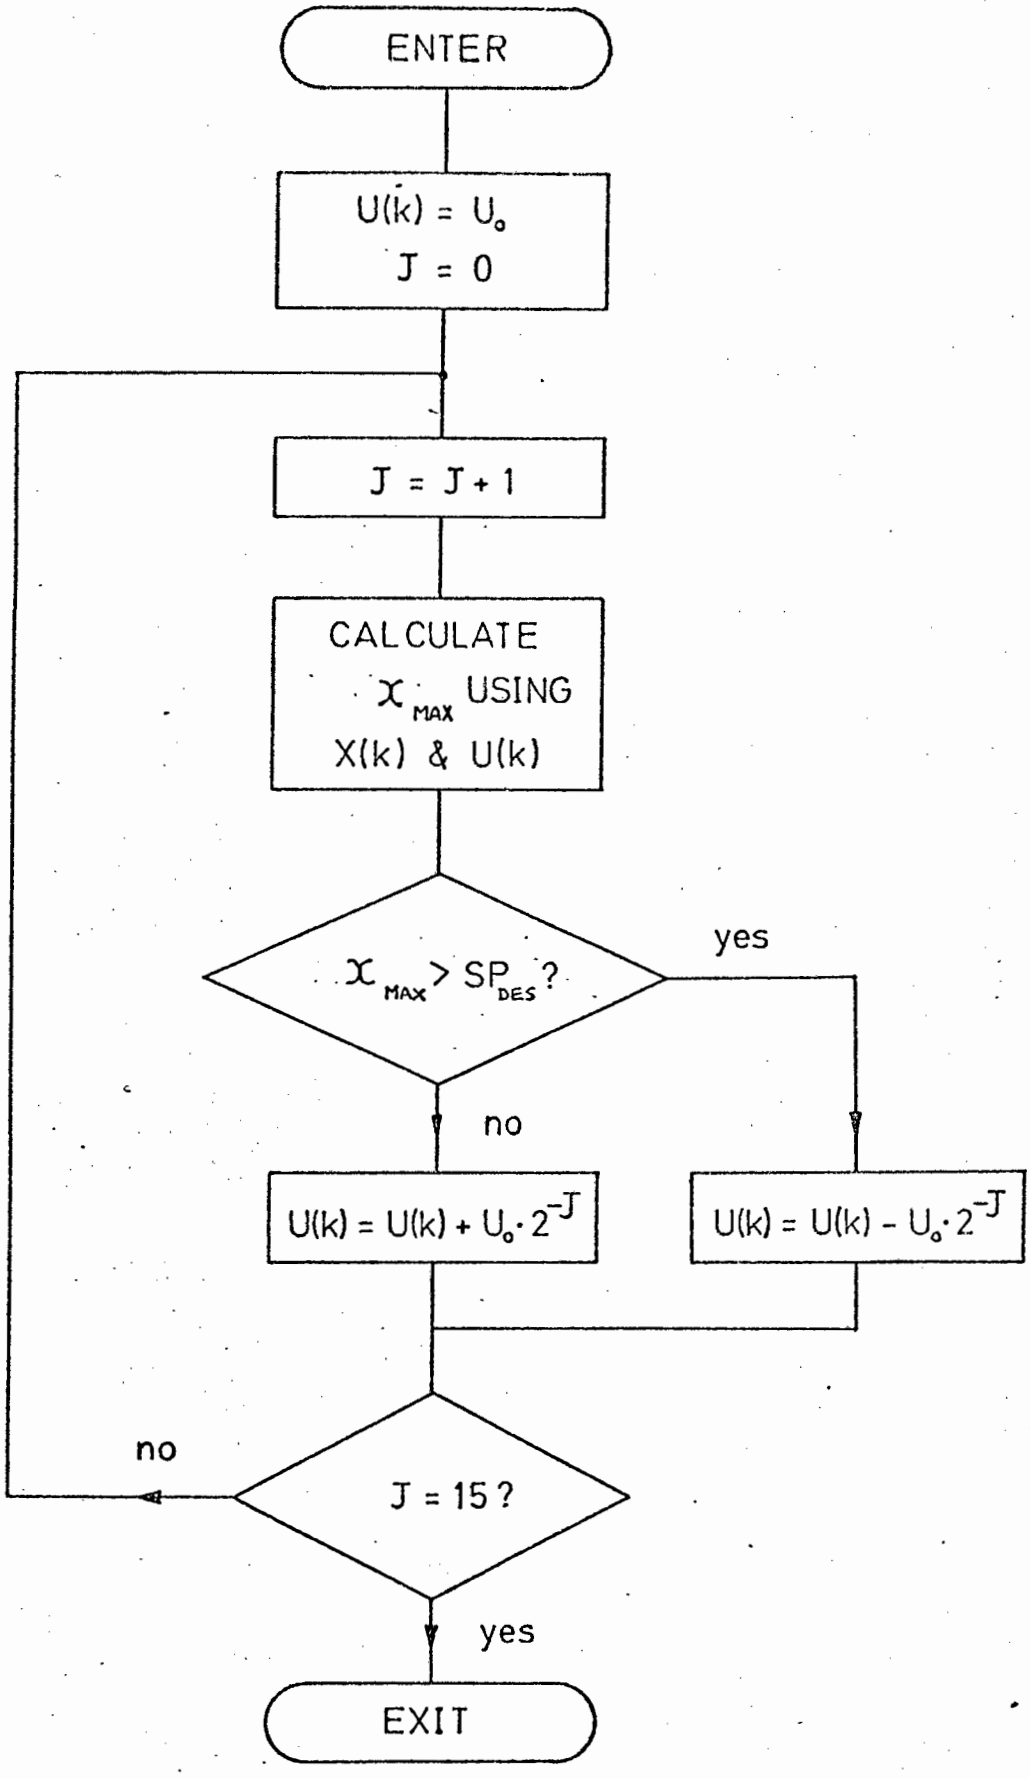

Figure 4.7 : Flowchart of the controller algorithm which describes the numerical search for the present time drug command u(k) which will just take the predicted level of relaxation to the chosen set-point.

$$SP_t(k) = \hat{x}_{max}(k) + \frac{SP_{des} - \hat{x}_{max}(k)}{\frac{L}{T_c} - i} \quad \ldots \quad 4.2$$

$\hat{x}_{max}(k)$ is the predicted maximum level of relaxation induced on the basis of doses already administered to the patient. It is determined at each control period before calculating the present time drug command. 'L' is the intended length of the relaxation phase (typically 60 sampling periods). 'T_c' is the spacing between successive drug commands and referred to as the control period (typically 6 sampling periods). The counter 'i' runs from 1 to (L/T_c-1) and is incremented at each control period. Thus the final trial set-point is the desired set-point.

4.5.3 The State Estimator

The State Estimator calculates the present time estimate of the true state vector $\underline{X}_T$ required in the control algorithm. It should be derived from a Kalman filter designed for composite deterministic and autoregressive stochastic patient models equations 3.11 and 3.17.

$$Y(z^{-1}) = \frac{z^{-\ell} g}{(1 - p_1 z^{-1})(1 - p_2 z^{-1})} U(z^{-1}) + \frac{E(z^{-1})}{D(z^{-1})} \quad \ldots \quad 4.3$$

where the white noise source E has variance R_e. If

$$a_1 = -(p_1 + p_2)$$

$$a_2 = p_1 p_2$$

$$A(z^{-1}) = 1 + a_1 z^{-1} + a_2 z^{-2}$$

$$D(z^{-1}) = 1 + d_1 z^{-1} + \ldots + d_n z^{-n} \quad ; \quad n = 5$$

equation 4.3 can be expressed as

$$D(z^{-1})\, Y(z^{-1}) = D(z^{-1})\, X(z^{-1}) + E(z^{-1}) \quad \ldots \quad 4.4$$

and

$$A(z^{-1})\ X(z^{-1}) = z^{-\ell}.g.U(z^{-1}) \qquad \ldots\ 4.5$$

where $X(z^{-1})$ is the output of the deterministic patient model. Equation 4.5 may be expressed as

$$\begin{bmatrix} x_1(k+1) \\ x_2(k+1) \\ x_3(k+1) \\ x_4(k+1) \\ x_5(k+1) \\ x_6(k+1) \end{bmatrix} = \begin{bmatrix} -a_1 & -a_2 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x_1(k) \\ x_2(k) \\ x_3(k) \\ x_4(k) \\ x_5(k) \\ x_6(k) \end{bmatrix} + \begin{bmatrix} g \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} u(k-\ell+1)$$

$$\underline{X}_n(k+1) = \phi_n\ \underline{X}_n(k) + \Gamma_n\ u(k-\ell+1) \qquad \ldots\ 4.6$$

where the state vector is defined by

$$\underline{X}_n(k) = \left[x_1(k)\ ,\ x_2(k)\ \ldots\ x_6(k)\right]^T$$
$$= \left[x(k)\ ,\ x(k-1)\ ,\ \ldots\ ,\ x(k-5)\right]^T$$

Equation 4.6 is now generalised to include a process driving noise source $w(k)$ of variance R_w to get it into the desired form. $w(k)$ nominally accounts for model inaccuracies. Thus

$$\underline{X}_n(k+1) = \phi_n\ \underline{X}_n(k) + \Gamma_n u(k-\ell+1) + \Gamma_n w(k) \qquad \ldots\ 4.7$$

Equation 4.4 may now be expressed as

$$y_F(k) = H_n \underline{X}_n(k) + e(k) \qquad \ldots\ 4.8$$

where

$$H_n = \left[1\ ,\ d_1\ ,\ d_2\ ,\ \ldots\ ,\ d_5\right] \qquad \ldots\ 4.9$$

and

$$Y_F(z^{-1}) = D(z^{-1})\ Y(z^{-1})$$

The usual Kalman state prediction/correction, covariance prediction/correction and gain generation equations, see for example Franklin and Powell, (1980), can be written immediately for the process described by equations 4.7 and 4.8 :

$$\left.\begin{aligned}
\bar{\underline{X}}_n(k+1) &= \underline{\Phi}_n\,\hat{\underline{X}}_n(k) + \Gamma_n u(k-\ell+1) \\
\hat{\underline{X}}_n(k+1) &= \bar{\underline{X}}_n(k+1) + G(k+1)\left[y_F(k+1) - H_n\bar{\underline{X}}_n(k+1)\right] \\
P_p(k+1) &= \underline{\Phi}_n\,P_c(k)\,\underline{\Phi}_n^T + \Gamma_n R_w \Gamma_n^T \\
G(k+1) &= P_p(k+1)\,H_n^T\left[H_n\,P_p(k+1)\,H_n^T + R_e\right]^{-1} \\
P_c(k+1) &= P_p(k+1) - G(k+1)\,H_n\,P_p(k+1)
\end{aligned}\right\} \quad \ldots \quad 4.10$$

where $\bar{\underline{X}}_n$ and P_p are the one step ahead predictions of the state vector and convariance matrix respectively while $\hat{\underline{X}}_n$, P_c are their corrections. The validity of equations 4.10 depends on 'v' and 'w' being zero mean and white noise sources.

Whether the conditions imposed on the noise sources always hold in pratice is doubtful. 'w' is little more than a mathematical artifice. The choice of R_w is really one of filter bandwidth and may be used to tune the filter to the process. In fact it is the ratio R_w/R_e which sets the gains in the steady state.

The filter is initialized by

$$\hat{\underline{X}}_n(0) = \left[0\,,\,0\,,\,0\,,\,0\,,\,0\,,\,0\right]^T$$

$$P_p(1) = \begin{bmatrix}
g^2R_w & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix}$$

The first of these follows because the initial state is known to be zero and the second follows from equations 4.10 and the fact that $P_c(0) = 0$ for there is no

uncertainty in the initial states.

R_w is needed to initialize the predicted covariance matrix. This is somewhat difficult to determine as explained above. It is chosen as that value which gives the best average simulated control on the original batch of 1Ø patients.

This completes the design of the Kalman filter. The elements x_1 and x_2 of the corrected state vector $\hat{\underline{X}}_n$ are used as the estimates of the true states in $\underline{X}_T$ required by the control algorithm.

In order to observe their transient behavior, the Kalman gains were calculated for patient 9 who is of average sensitivity from time zero until the steady state was reached. This was done off-line by programme CONTRL.RICATI which is described in Appendix B. The patient model parameters and measurement noise variance were taken from Table 3.3. The process noise variance was set to the typical value of $R_w = Ø,Ø42$. (This value is derived in section 4.6.2). The gains are plotted in figure 4.8. There are six elements in accordance with the definition of the State Vector in equation 4.6.

The gains reach their steady state in about 1Ø minutes. They are smallest initially because the initial patient state is known exactly.

The full Kalman filter requires:

1. the parameters of the patient noise model

2. R_e and R_w separately

3. online solution of the Riccati equation which forms the major computational burden of the Kalman filter.

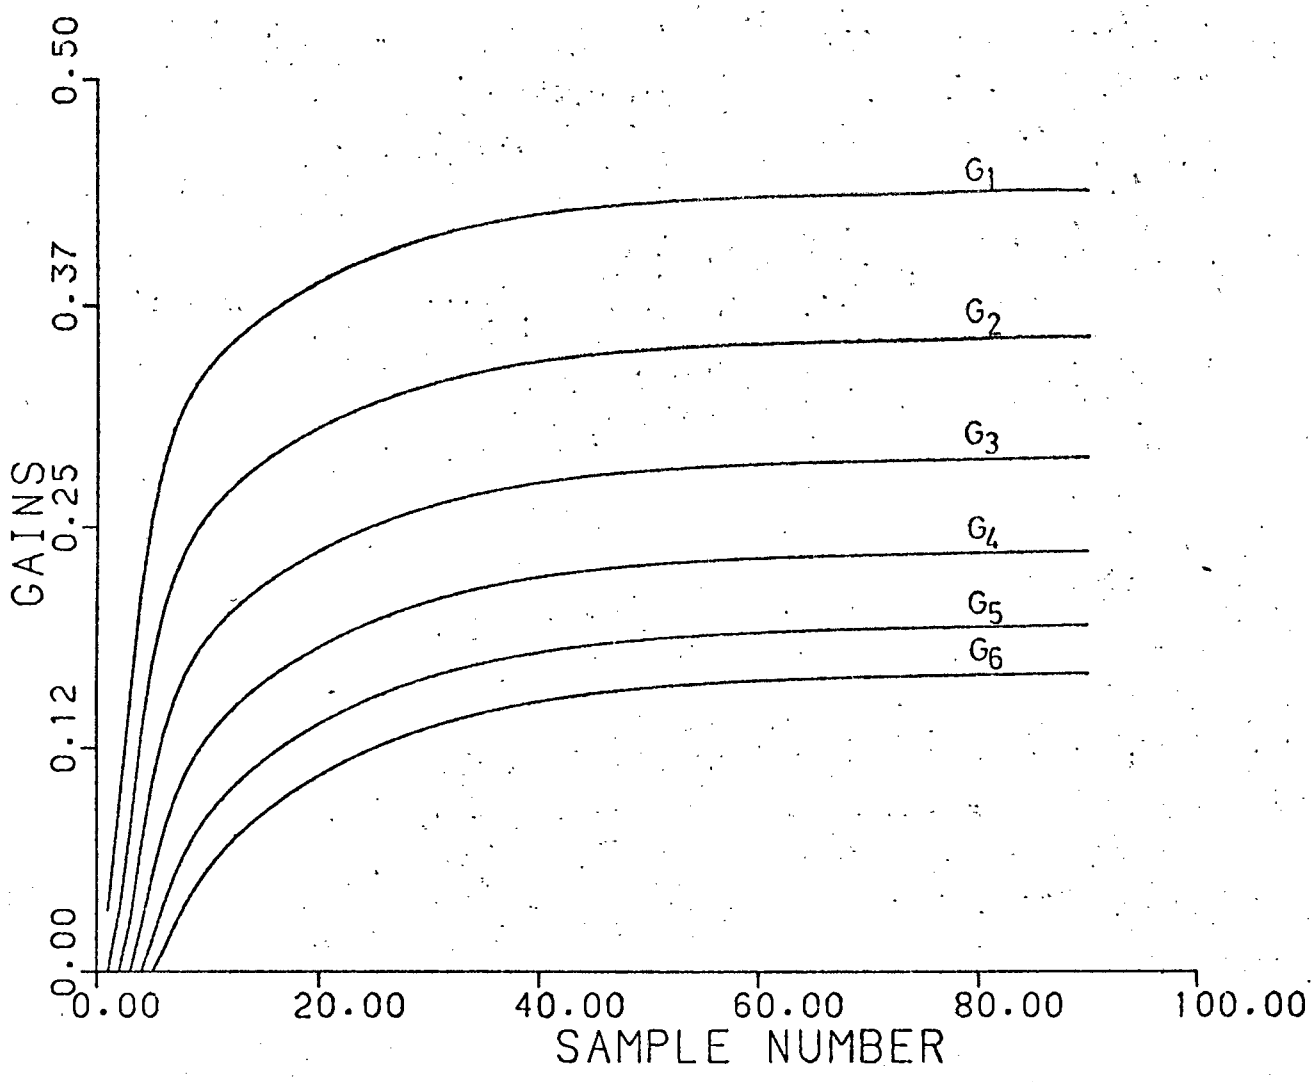

Figure 4.8 : Gain elements of the 6 state Kalman filter. These pertain to patient 9 who is of average sensitivity.

It is of interest to know whether a simpler linear Estimator formed using the steady state limit of the Kalman gains and disregarding the noise correlation will perform satisfactorily. In place of the three requirements above, such an Estimator only needs R_w/R_e to tune it. The gains can be calculated off-line. As the noise correlations are disregarded, the form of $\underline{\Phi}, \Gamma, H, \underline{X}$ would be the same as those of the simpler deterministic system defined by equation 4.1. Accordingly the controllers CONTRL.BOLUS and CONTRL.BOLUS4 were designed, the former using the full Kalman filter to estimate the states and the latter the simplified linear Estimator.

The two controllers are evaluated in chapter 6. In the event it turns out that there is no significant difference between them and so attention is focussed on the linear State Estimator henceforth.

One method of checking the performance of the State Estimator is to plot the estimated level of relaxation calculated by it together with the actual patient response. The trade-off between the tracking and noise smoothing abilities of the Estimator can then be assessed visually as its bandwidth is changed. Accordingly a provision was written into the FORTRAN programme that is the equivalent of the controller implemented on the microcomputer which enables it to operate on the patient data captured during clinical trials rather than on that of an simulated patient. The estimated patient response may then be plotted together with the actual patient data. Details are given in Appendix A.

4.5.4 Parameter Estimation Algorithm

The purpose of the parameter estimation algorithm is to provide an estimate of the patient parameter vector $\underline{\theta}_D$ for use by the control algorithm in the relaxation and regulation phases. In this section three parameter estimators are considered:

1. A sophisticated Least Squares routine written in FORTRAN and intended for use only in the control of simulated patients on the UNIVAC computer. It is used as a standard against which to compare the performance of simpler methods. It incorporates the optimization package VAØ5A used to model the patient response off-line as described in chapter 3. It allows any combination of the parameters to be estimated but is complex in construction. The effort required to translate it into assembly language for use on the microcomputer was

not considered to be justified.

2. A simple Least Squares routine using Steepest Descent and written in FORTRAN. It calculates the gradient vector exactly.

3. A simplification of the Steepest Descent method of 2. above and written in FORTRAN for direct translation into assembly language and implementation on the microcomputer. It uses a simple numerical approximation to the gradient vector. At the time of the design, simplicity of the algorithm was an important consideration because the departmental laboratories were unable to support a high level language for a micro-processor at the time that this research was undertaken.

In all three parameter estimation algorithms, the Output error method is used to compare patient and model using the structure of figure 3.2(b) developed in chapter 3. The noise parameters are not estimated for they are only required by the controller which uses the Kalman filter. In this case they are assumed to be known to the controller a priori in order to attain the best performance of which this controller is capable.

Equation Error was considered less suitable than Output Error because the patient disturbance is known to be correlated. Equation Error requires that a noise model be formulated and its parameters estimated along with those of the deterministic model otherwize convergence of the latter to the correct limit cannot occur. A noise model is unnecessary for Output Error (Moore,1982) nor is knowledge of its parameters required subsequently in calculating the controller commands.

The loss function for N collected data is

$$V(\underline{\theta}_D) = \sum_{k=1}^{N} r_c^2(k).W(k) \quad \ldots \quad 4.11$$

In the optimization process it was decided to minimize V numerically. While a recursive Output Error algorithm does exist in the form of the Automatic Compensator version (ROAC) due to Landau, see for example Dugard and Landau, (1980), it needs to be modified to operate on a discontinuous data record and to estimate the parameter BT4 which shows itself as an off-set in the linearized twitch data as stated in section 4.2. This has been done for recursive algorithms which use Equation Error. Franklin and Powell,(1980) show how a sequence of weights can be incorporated into Recursive Least Squares. Isermann,(1982) shows how to estimate an off-set for three recursive parameter estimation algorithms. It may be possible to extend these to Output Error.

The Steepest Descent method uses a simple linear search in the direction of the gradient vector (Adby and Dempster, 1974). Although slow to converge near its minimum, Steepest Descent requires only the first partial derivative of V with respect to the parameters being estimated. If the parameter vector is

$$\underline{\theta} = \left[\theta_1 \ , \ \theta_2 \ , \ \ldots \ , \ \theta_n\right]^T$$

the unit vector in the direction of steepest descent in parameter space is

$$\underline{\xi} = \frac{\left[\frac{\partial V}{\partial \theta_1} , \frac{\partial V}{\partial \theta_2} , \ldots , \frac{\partial V}{\partial \theta_n}\right]^T}{\sqrt{\left(\frac{\partial V}{\partial \theta_1}\right)^2+\left(\frac{\partial V}{\partial \theta_2}\right)^2+ \ldots +\left(\frac{\partial V}{\partial \theta_n}\right)^2}}$$

The partial derivatives in question can be found by either:

1. differentiating V in equation 4.11 with respect to $\underline{\theta}_D$ analytically, or

2. approximating the partial derivatives numerically.

1. Analytical calculation of the partial derivatives

From equation 4.11

$$\frac{\partial V}{\partial \theta_i} = 2\sum_{k=1}^{N} r_c(k).W(k).\frac{\partial r_c(k)}{\partial \theta_i} \quad \ldots \quad 4.12$$

From the formation of the Output Error in figure 3.3(b), the residual is

$$r_c(k) = T_{lin}(k) - \left[\frac{z^{-\ell} g}{(1 - p_1 z^{-1})(1 - p_2 z^{-1})} u(k) + BT4\right]$$

or

$$r_c(k) - (p_1 + p_2)r_c(k-1) + p_1p_2r_c(k-2) = T_{lin}(k) - (p_1 + p_2)T_{lin}(k-1)$$
$$+ p_1p_2T_{lin}(k-2) - g.u(k-\ell) - BT4\left[1 - (p_1 + p_2) + p_1p_2\right]$$

Thus

$$\left.\begin{aligned}
\frac{\partial r_c(k)}{\partial p_1} &= (p_1 + p_2)\frac{\partial r_c(k-1)}{\partial p_1} - p_1p_2\frac{\partial r_c(k-2)}{\partial p_1} + r_c(k-1) - \\
&\quad p_2r_c(k-2) - T_{lin}(k-1) + p_2T_{lin}(k-2) + BT4(1 - p_2) \\
\frac{\partial r_c(k)}{\partial p_2} &= (p_1 + p_2)\frac{\partial r_c(k-1)}{\partial p_2} - p_1p_2\frac{\partial r_c(k-2)}{\partial p_2} + r_c(k-1) - \\
&\quad p_1r_c(k-2) - T_{lin}(k-1) + p_1T_{lin}(k-2) + BT4(1 - p_1) \\
\frac{\partial r_c(k)}{\partial g} &= (p_1 + p_2)\frac{\partial r_c(k-1)}{\partial g} - p_1p_2\frac{\partial r_c(k-2)}{\partial g} - u(k-\ell) \\
\frac{\partial r_c(k)}{\partial BT4} &= -1
\end{aligned}\right\} \quad \ldots \quad 4.13$$

As patient and model are initially at rest,

$$r_c(k) = \text{constant} - BT4 \quad , -\infty < k \leq 0$$

Thus

$$\left.\frac{\partial r_c(k)}{\partial \theta_i}\right|_{i = p_1, p_2, g} = 0 \quad , -\infty < k \leq 0$$

and equation 4.13 can be used to calculate $\partial r_c(k)/\partial \theta$, $k = 1,N$. These are substituted into equation 4.12 to calculate $\partial V/\partial \theta$ as required.

2. Numerical approximation of the partial derivatives

A somewhat simpler method of finding $\partial V/\partial\theta$ is to use

$$\frac{\partial V}{\partial \theta_i} \simeq \frac{V(\theta_1,\theta_2,\ldots,\theta_{i-1},\theta_i+\text{DSTEP},\theta_{i+1},\ldots,\theta_n) - V(\underline{\theta})}{\text{DSTEP}} \quad \ldots \quad 4.14$$

for small enough DSTEP.

The alternative methods have approximately the same computational load for both involve a summation over all the collected data. The numerical approximation however is somewhat simpler to programme. Both methods were used. DSTEP was set to 2^{-13}, this value being found to give comparable estimation accuracy to the more sophisticated routine VA05A.

With $\underline{\xi}$ the unit vector in the direction of steepest descent calculated, the parameter vector is updated according to

$$\underline{\theta}_{i+1} = \underline{\theta}_i - \lambda\underline{\xi} \quad \ldots \quad 4.15$$

where λ is the constant found by the linear search to give greatest reduction of the loss function in the chosen direction. The range of λ was selected to be $[1,2^{-13}]$ again after relative comparison with VA05A. This is an increase in precision of 2^3 over the range used in the original controller evaluation (Rametti and Bradlow,1983) and eliminates the ripple on the estimated parameters noted there.

The number of iterations of equation 4.15 per sampling period was set at 1 in the original control simulations reported in the above paper but was subsequently increased to take advantage of the processing power of the 8088/8087 combination which exceeded earlier expectations. The scheme adopted is given in Table 4.2.

SAMPLING PERIOD	NO. OF ITERATIONS OF EQUATION 4.15
ℓ	Parameter Estimation inhibited (algorithm stability)
ℓ - 50	4
51 - 70	3
71 - 100	2
101 - 200	1
> 200	Parameter Estimation inhibited (computational load)

Table 4.2 : Number of iterations of the Steepest Descent algorithm per sampling period throughout the interval that parameters are estimated.

It should also be noted that the twitch data is not discounted by the use of some forgetting factor as is commonly done with recursive estimators so allowing them to adapt to changing parameters. The advantage of the infinite memory scheme used is that the parameter estimator cannot windup even if the input signal is not persistently exciting (Astrom, 1983). When the control loop is closed around the patient the drug commands are determined solely by the on-line controller and generally persistent excitation cannot be guaranteed.

The parameters of the patient model should be sensibly initialized with knowledge gained from previous modelling studies performed on a batch of patients. These studies should use

1. the inverse logistic relationship

2. a fixed value of the patient circulatory delay 'ℓ' taken to be the average value of 5 sampling periods. This is not determined on-line as it is a discrete parameter.

The patient model parameters estimated off-line and presented in Table 3.3 conform to these requirements.

The question arises as to which parameters in the vector $\underline{\theta}_D$

$$\underline{\theta}_D = \left[p_1 \; ; \; p_2 \; ; \; g \; ; \; BT4 \right]^T$$

should be estimated on-line. If the sophisticated algorithm VA05A is used, all the parameters may be estimated provided care is exercised to delay the estimation of the poles for model stability (to keep them within the unit circle). Experience has shown that approximately 50 and 100 data need to be collected before the smaller and larger poles respectively can be estimated reliably. Thus estimation of the poles will not benefit the control greatly in the relaxation phase. In any event the controller performance is least affected by the choice of the larger pole (Rametti and Bradlow, 1983). It was therefore decided that the poles would not be estimated in the version of the controller implemented on the microcomputer but set to "averaged" values. These are discussed in chapter 6.

The decision not to estimate the poles on-line is reinforced by considering the means of model parameters obtained from groups of patients in different clinical trials. If the larger pole of the patient model is fixed at an average value, physical significance can be associated with the remaining three parameters as follows:

1. the shape of the pharmacodynamic relationship is determined by BT4,

2. the time to maximum effect is determined by the minor pole,
   and with the minor pole and BT4 fixed,

3. the sensitivity of the patient model is determined by the gain

parameter.

It is known that there is interpatient variability in these three quantities but there is also a variation between groups of patients in which the clinical methods were different. Patients 1 - 10 and 11 - 26 are two such groups. In both cases the T4 ratio was modelled. The t-test described by Spiegel,(1961) was used to compare these groups and revealed that both the average patient sensitivity and BT4 were significantly different at a 5% confidence level whereas the time to maximum effect was not. (The data for this test was taken from Table 3.3).

It so happens that the Steepest Descent algorithm described above was found to be inefficient when attempting to tune more than two parameters simultaneously so the decision to tune g and BT4 alone was fortunate.

In the course of the clinical trials, it was discovered that numerical difficulties were encountered in using the Steepest Descent algorithm as proposed because the sensitivity of the loss function to the parameters g and BT4 . is very different. It was consequently found necessary to estimate the parameter BT4/100 as opposed to BT4 and to apply the reverse scaling factor to this parameter when it was required in the controller calculations.

The sequence of weights W(k) in equation 4.11 is set to 1 for twitch ratio data in the range 80% to 10% and zero outside it. Ham et al.,(1979) proposed limits 90% and 10% but is was found that a twitch ratio of 90% is not always a positive indication of the onset of block. Consequently the relaxation phase is only entered at 85% and a further band of 5 % is then allowed to ensure that the data is reliable before commencing parameter estimation at 80%.

Despite this restriction placed on the data used in the parameter estimation, all twitch ratios between 0% and 100% were used by the controller in

calculating the mass of drug to be administered at each control period. If data is excluded from the control calculation, there is a potential source of error or instability.

4.6 Comparison between alternative controllers in which the patient states are estimated by Kalman filter and by simple linear Estimator

In this section a comparison is made between the muscle relaxation controller which uses a simple linear State Estimator and a numerical approximation to the gradient vector in the Steepest Descent algorithm and the more elaborate controller which estimates the states with a Kalman filter and calculates the gradient vector analytically. All controller settings other than for the State Estimator are given in section 6.2.

Controller performance is rated in terms of four criteria which are defined below.

4.6.1 Definition of Performance Criteria

The disadvantage of the present method of administration of dTC viz. a large loading dose and subsequent "top-up" doses is that on average the peak plasma concentration of the drug is far in excess of the level required to induce relaxation during surgical proceedures. This could conceivably be avoided if the anaesthetist were to "experiment" on each patient by administering a number of smaller doses but this would lengthen the induction time.

Two criteria by which to judge the controller's performance are then the time to reach the chosen set-point and the resulting overshoot. There is usually some trade-off between these.

In the regulation phase, the clinical requirements of the controller are less demanding. It might be sufficient to keep the twitch ratio above the 10% T1 ratio level defined by Ham et al.,(1979). The criteria that were chosen for this phase of control are more stringent than this however. The quality of control is judged by the off-set from the set-point and the standard deviation about this level.

Formal definitions of the four performance criteria for use in this thesis are therefore:

1. Time to set-point - the time taken to bring an initially unrelaxed patient to the set-point. This is the time in minutes between the instant that the loading dose is administered and a smoothed estimate of the twitch response reaches the set-point.

2. The initial overshoot of the set-point - this is expressed as a percentage of the range of possible twitch ratios [0,1].

3. Off-set in the regulation phase - this is the difference between the patient's twitch ratio and the set-point averaged throughout the regulation phase and expressed as a percentage of the range of possible twitch ratios.

4. Standard Deviation in the regulation phase - this is the standard deviation about the mean off-set in the regulation phase and expressed as a percentage of the range of possible twitch ratios.

In comparing the average performance of different versions of the controller a statistical test is necessary to decide whether or not there is a significant difference between them. For this purpose the test of hypotheses based on the Z-score for large samples (greater than 30) and t-score for small samples as described in Spiegel,(1961) are used. The confidence interval is chosen as 5% throughout the thesis.

4.6.2 Selection of the measurement and process noise variances for the Kalman filter.

The Kalman filter requires the parameters of the patient noise model and the noise variance R_e. These are however unknown to the controller as they are not estimated on-line. For the purposes of the present comparison of controllers however, this information derived from Table 3.3 is supplied to the controller. A real controller would not have access to this information however and presumably could not perform as well.

The best average values of the process noise variance R_w to be used for control between sampling periods 1 and 200 and then beyond sampling period 200 for regulation were found by means of a simple manual search. For each of the trial values of R_w given in Table 4.3, the average controller performance expressed in terms of the four defined criteria was calculated for the group of simulated patients 1 - 10. (Patient 7 was excluded because the controller became unstable in this case). It is clear from the table that R_w = 0,042 is close to the best choice which can be made for control between sampling periods 1 and 200. Control beyond this point is best with a value of R_w around 0,0042 which is 10% of the former.

R_w		Time to Set-point (Mins)	Initial overshoot (%)	Off-set from Set-point (%)	Std. devn. from mean (%)
Sampling period 1 - 200	beyond Sampling period 200				
0,4200	0,0420	11,8 ± 4,0	8,7 ± 5,5	4,0 ± 2,0	4,4 ± 1,7
0,1330	0,0130	11,9 ± 4,4	6,3 ± 5,6	2,4 ± 2,0	3,9 ± 1,8
0,0750	0,0075	11,9 ± 4,7	5,0 ± 5,6	2,0 ± 2,2	3,7 ± 1,9
0,0420	0,0042	11,7 ± 4,7	3,6 ± 6,6	1,7 ± 2,2	3,7 ± 2,1
0,0240	0,0024	12,0 ± 5,6	5,5 ± 5,8	1,5 ± 2,3	3,8 ± 2,2
0,0130	0,0013	12,8 ± 6,3	6,2 ± 6,4	1,6 ± 2,7	3,9 ± 2,4
0,0042	0,00042	13,5 ± 7,3	7,8 ± 7,0	2,1 ± 4,3	4,3 ± 3,1

(a)

R_w		Time to Set-point (Mins)	Initial overshoot (%)	Off-set from Set-point (%)	Std. devn from mean (%)
Sampling period 1 - 200	beyond Sampling period 200				
0,2000	0,0200	11,7 ± 4,1	8,8 ± 6,0	4,8 ± 3,2	4,2 ± 1,8
0,0630	0,0063	12,4 ± 3,8	5,5 ± 5,2	3,2 ± 1,8	3,8 ± 1,8
0,0360	0,0036	12,3 ± 3,7	5,6 ± 5,1	2,9 ± 1,8	3,7 ± 1,9
0,0200	0,0020	12,2 ± 3,9	3,7 ± 5,0	2,5 ± 1,6	3,7 ± 1,8
0,0110	0,0011	12,7 ± 3,7	4,1 ± 5,5	2,2 ± 1,7	3,7 ± 1,9
0,0063	0,00063	12,7 ± 3,6	4,8 ± 5,5	1,8 ± 1,7	3,6 ± 1,9
0,0020	0,00020	13,1 ± 3,7	4,0 ± 5,2	1,5 ± 1,8	3,6 ± 2,0

(b)

Table 4.3 : Average performance of two controllers on the batch of simulated patients 1 - 10.
(a). Controller which estimates patient states with Kalman filter and calculates alalytically the gradient vector for Steepest Descent minimization.
(b). Controller which uses the linear State Estimator and a numerical approximation to the gradient vector.

The linear State Estimator is the steady state limit of a simple Kalman filter designed for an average patient model. It disregards noise correlations. In this case it is the ratio of the measurement and process noise variances which set the gains. In keeping with Rametti and Bradlow, (1983), the measurement noise variance R_e will be set at Ø,Ø1 for convenience and R_w considered exclusively hence forth. The parameters of the patient model used for this State Estimator are given in section 6.2.

Again a search was made for the best values of R_w for use in the linear State Estimator. For control between sampling periods 1 and 2ØØ, R_w = Ø,Ø2 was close to best as Table 4.3 shows. Beyond sampling period 2ØØ however, regulation improved as R_w was decreased - even down to Ø,ØØØ2. Such small values are dangerous because the tracking ability of the Estimator is seriously impaired. (The choice R_w = Ø,ØØØ2 during a particular clinical trial caused controller malfunction. This case is discussed separately in chapter 6). The regulation phase value of R_w was consequently set at Ø,ØØ2 which is 1Ø% of the best value found for the relaxation phase. This was the finding with the controller which used the full Kalman filter, as stated above.

4.6.3 Performance of the two controllers

The performance of the alternative controllers for each of the simulated patients in the batch 1 - 1Ø and their average performance is given in Table 4.4.

The average performance of the controller which uses the full Kalman filter and calculates the Steepest Descent gradient vector analytically is slightly better than the other though not significantly so according to the Student's t-test. Furthermore it gave unacceptable control of patient 7. This justifies

PATIENT	Kalman filter and analytical calculation of gradient vector				linear State Estimator and numerical approximation to gradient vector			
	TIME TO SET-POINT (MINS)	INITIAL OVERSHOOT (%)	OFF-SET FROM SET-POINT (%)	STD. DEVN FROM MEAN (%)	TIME TO SET-POINT (MINS)	INITIAL OVERSHOOT (%)	OFF-SET FROM SET-POINT (%)	STD. DEVN FROM MEAN (%)
1	8,3	0,0	1,8	4,3	10,3	6,8	3,9	4,5
2	11,3	0,0	0,7	2,2	11,3	0,0	0,7	2,2
3	21,2	1,9	0,1	1,1	20,2	2,1	0,2	1,1
4	8,6	0,0	-0,5	2,5	8,3	0,0	2,3	3,0
5	15,4	13,4	5,2	4,5	16,0	10,9	3,6	3,9
6	13,0	0,0	0,8	2,3	13,9	1,1	0,9	2,3
7	UNSTABLE				15,7	16,7	6,4	6,5
8	9,4	16,8	5,4	7,0	10,0	12,4	4,6	6,0
9	12,8	0,0	1,9	3,1	11,6	0,0	3,0	3,9
10	5,1	0,0	-0,2	6,7	8,8	0,0	3,1	6,7
MEAN	11,68	3,57	1,69	3,74	12,23	3,70	2,48	3,73
STD. DEVN	4,70	6,62	2,20	2,05	3,87	5,02	1,55	1,82

Table 4.4 : Comparison of controllers which :
(a). use the full Kalman filter and analytical calculation of the gradient vector,
(b). use a linear State Estimator and numerical approximation to the gradient vector.
The average figures exclude patient 7.

the use of the simple linear State Estimator and the numerical approximation to the gradient vector.

4.7 The controller settings

The control algorithm implemented on the microcomputer has a total of 8 settings which are entered at the keyboard prior to initiation of control. They are:

- The Extreme Parameter Vector of the Initial Phase Controller (4 parameters)
- The Initial Parameter Vector of the Relaxation/Regulation Phase Controller (4 parameters).
- The State Estimator gains (2 parameters for relaxation phase)
  (2 parameters for regulation phase)
- Parameter L : desired length of relaxation phase
- Parameter ℓ : patient model circulatory delay in sampling periods
- Parameter T_c : the control period
- Parameter SP : chosen set-point on range (Ø,1)
- The patient body mass.

The patient mass is either known or can be estimated. The desired length of the relaxation phase was taken as 6Ø sampling periods in every case. This corresponds to 9 minutes at a sampling period of 9 seconds or 1Ø minutes at 1Ø seconds. This was chosen because the best average induction time from a single bolus of dTC for set-points in the range considered in this work is just over 8

minutes (see Table 6.5).

The choice of the other 6 settings is discussed in section 6.2.

4.8 Computational implications of the control algorithm

The parameter estimation method proposed is not a recursive technique but a batch method repeated at each sampling period. The only information that links estimations at successive sampling periods is the parameter vector which is being updated. There are 2 weaknesses of this method:

1. The need for an expanding memory. Input/Output data must be stored from the beginning of the parameter estimation process i.e. from the start of control. (This is also true for the State Estimator which re-estimates from start to present time at each sampling period so as to make best use of the improving patient parameter estimates). This however is a minor disadvantage because an operation of 10 hours duration would require 58K Bytes of RAM for the arrays. This is inexpensive and simple to install.

2. The computational load increases linearly with the length of the data record. It follows that parameter estimation must be terminated at some point. This is done at 200 samples. However it is not a serious disadvantage in the muscle relaxation problem because the patient is then well into the regulation phase and the estimated parameters have largely converged as explained in section 6.5.

The estimated computational load placed on the 8087 NDP by each of the parameter estimation routine, State Estimator and Control algorithm at sampling period 200 is given in Table 4.5. This is just before parameter estimation is inhibited and thus shows the greatest execution time of the parameter estimation algorithm. It embraces the mathematical functions used and the data transfers to and from memory. These are easily the most time consuming instructions executed. Hence the total execution time of approximately 4,4 seconds per sampling period shown in the table is a good approximation to the load imposed by the online controller.

After sampling period 200 parameter estimation is inhibited and thus it is only the execution time of the control algorithm and State Estimator which need to be considered. The former is fixed at 0,82 seconds and the latter increases linearly with time. It follows that total execution time will reach the limit of 8 seconds between successive patient stimulations in approximately 30 hours. Execution time therefore imposes no real restriction on the usefulness of the controller.

High internal precision of the microcomputer is required if the gradient vector of the Steepest Descent method is calculated as proposed in equation 4.14. With DSTEP = 2^{-13} one might expect a difference of as little as 2^{-19} between the loss functions in the numerator. Thus the internal precision required of the computer greatly exceeds that of the Input/Output data (twitch ratio is 8 bit). The 8087 represents its floating point numbers in either the Short-real (32 bit) or Long-real (64 bit) formats. These have mantissas of 23 and 52 bits respectively. The Long-real format was used in calculating the gradient vector, the Short-real representation was adequate for all else.

Routine \ Instruction type	Number of instructions executed								MAX. TIME PER SAMPLING PERIOD (SECS)
	Memory Load to 8087 transfer	8087 to memory transfer	+	-	X	÷	$\sqrt{}$	$\log_2$	
Parameter Estimation Routine	34260	28655	17036	11410	28486	15	1	0	3,43
State Estimator	1177	789	781	391	1368	0	0	0	0,13
Control algorithm	11373	8169	3217	3236	4864	24	0	2	0,82
8087 NDP Instruction execution speed in µSec @ 3.33MHz CLOCK	15,6	29,1	34,2	34,2	38,1	68,7	54,9	285,0	

Table 4.5 : Estimate of the computational load imposed on the 8087 NDP. The calculation is done at sampling period 200. The total execution time which is the sum of the entries in the last column, is 4,4 seconds per sampling period.

CHAPTER 5

Controller Implementation

NOTE

The symbols used in this chapter which relate directly to the controller software are explained in Appendix C and not in the Symbol Table.

5.1 Macro flowchart of the controller implemented on the microcomputer

It was explained in section 4.5 that three versions of the muscle relaxation controller were designed each having the structure shown in figure 4.5. One of these used a sophisticated optimization method in the parameter estimation section, the other two used Steepest Descent. Of these two, one used a Kalman filter to estimate the necessary states, the other, and simplest algorithm, a linear State Estimator. It was this simplest algorithm (referred to as BOLUS4) which was implemented on the microcomputer. This was done in accordance with the macro-flowchart shown in figure 5.1. The assosiated software to which reference is made in this chapter is contained in Appendix C.

The five units within the controller of figure 4.5 correspond to the five numbered sections shown on the flowchart. They are

1. The data linearization section

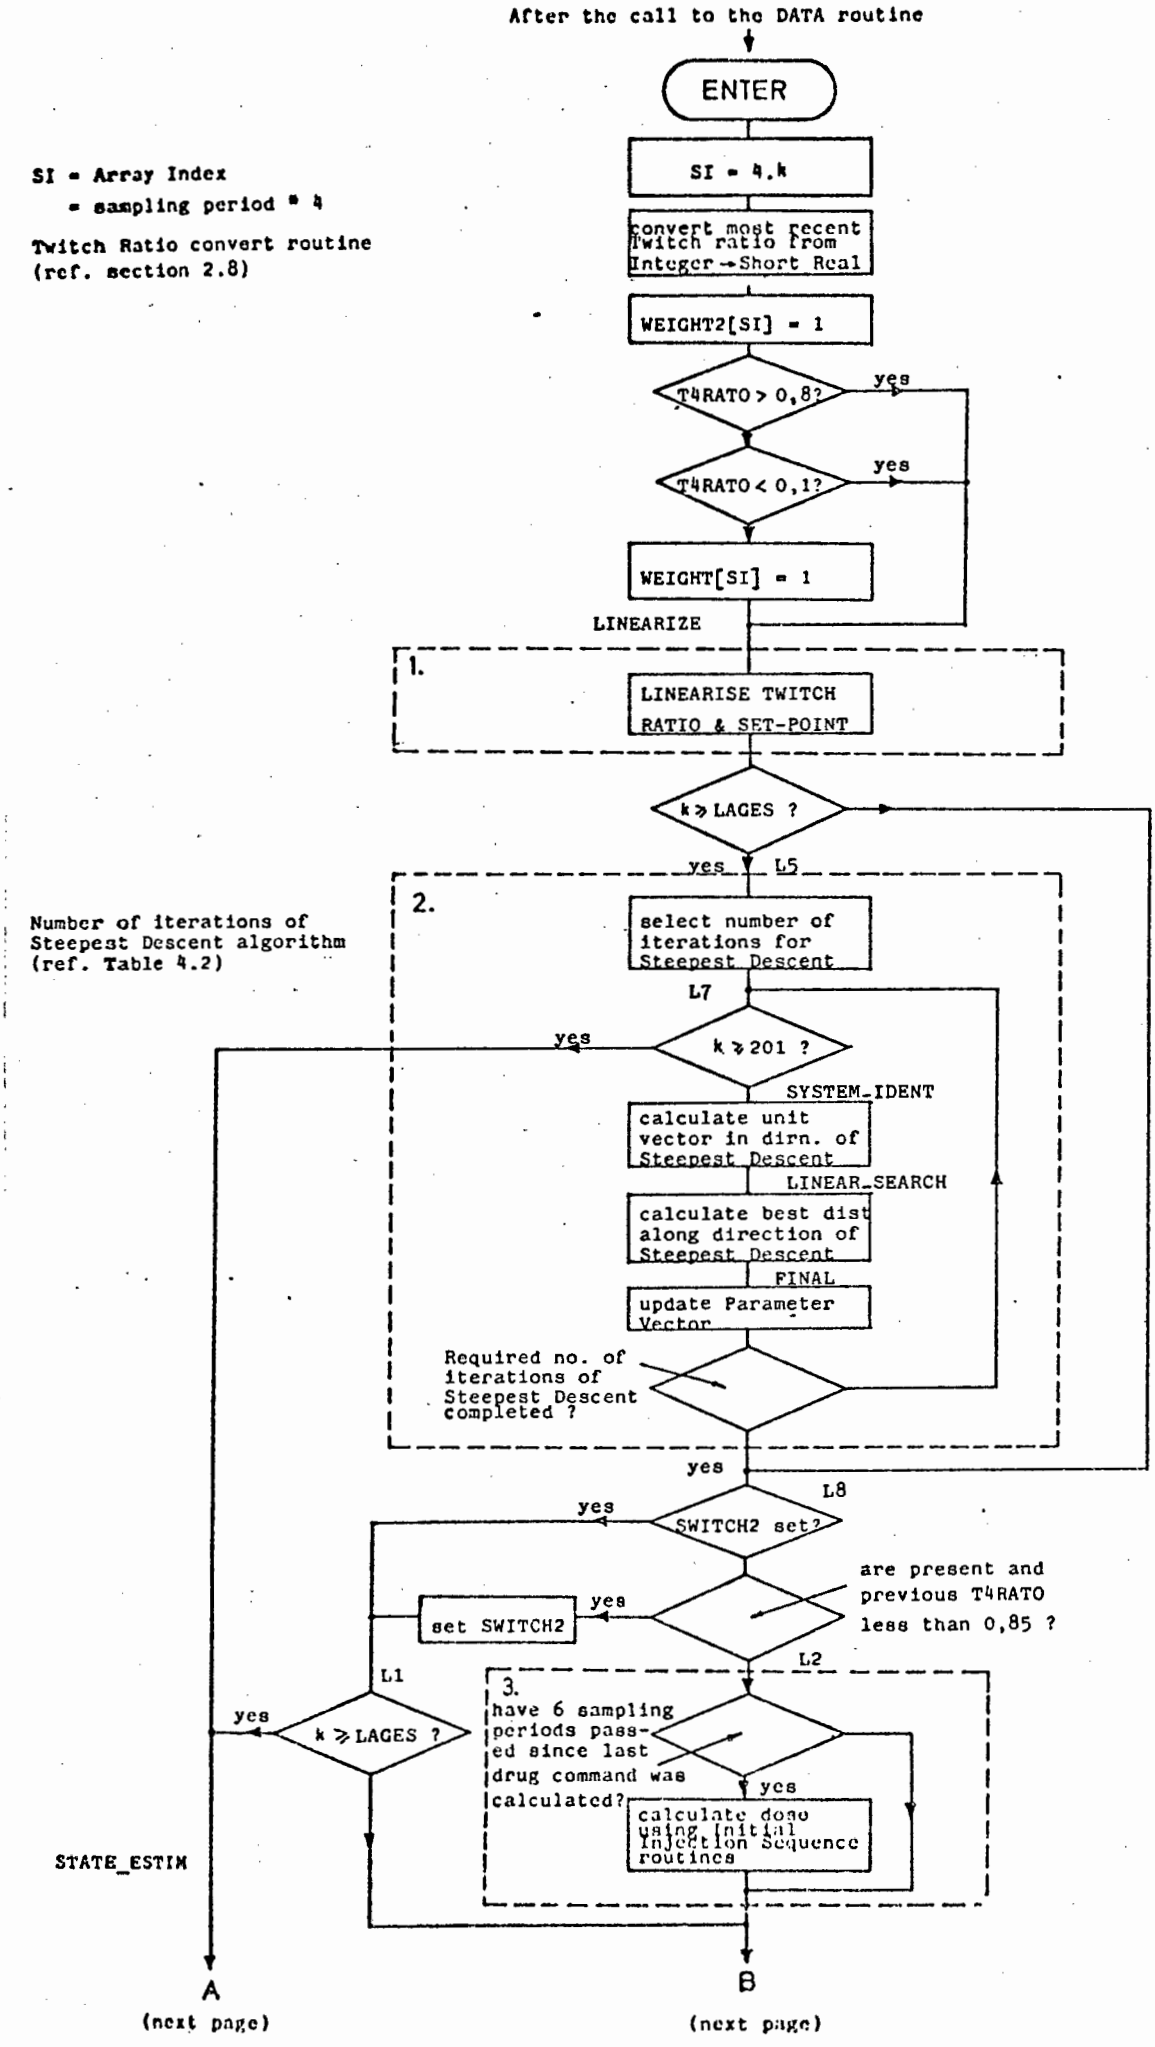
After the call to the DATA routine
ENTER
SI = 4.k
convert most recent Twitch ratio from Integer→Short Real
WEIGHT2[SI] = 1
T4RATO > 0,8?
yes
T4RATO < 0,1?
yes
WEIGHT[SI] = 1
LINEARIZE
1.
LINEARISE TWITCH RATIO & SET-POINT
k ⩾ LAGES ?
yes
L5
2.
select number of iterations for Steepest Descent
L7
yes
k ⩾ 201 ?
SYSTEM_IDENT
calculate unit vector in dirn. of Steepest Descent
LINEAR_SEARCH
calculate best dist along direction of Steepest Descent
FINAL
update Parameter Vector
Required no. of iterations of Steepest Descent completed ?
yes
L8
yes
SWITCH2 set?
are present and previous T4RATO less than 0,85 ?
yes
set SWITCH2
L2
L1
3.
have 6 sampling periods passed since last drug command was calculated?
yes
calculate dose using Initial Injection Sequence routines
yes
k ⩾ LAGES ?
STATE_ESTIM
A
(next page)
B
(next page)
SI = Array Index
= sampling period * 4
Twitch Ratio convert routine
(ref. section 2.8)
Number of iterations of Steepest Descent algorithm
(ref. Table 4.2)

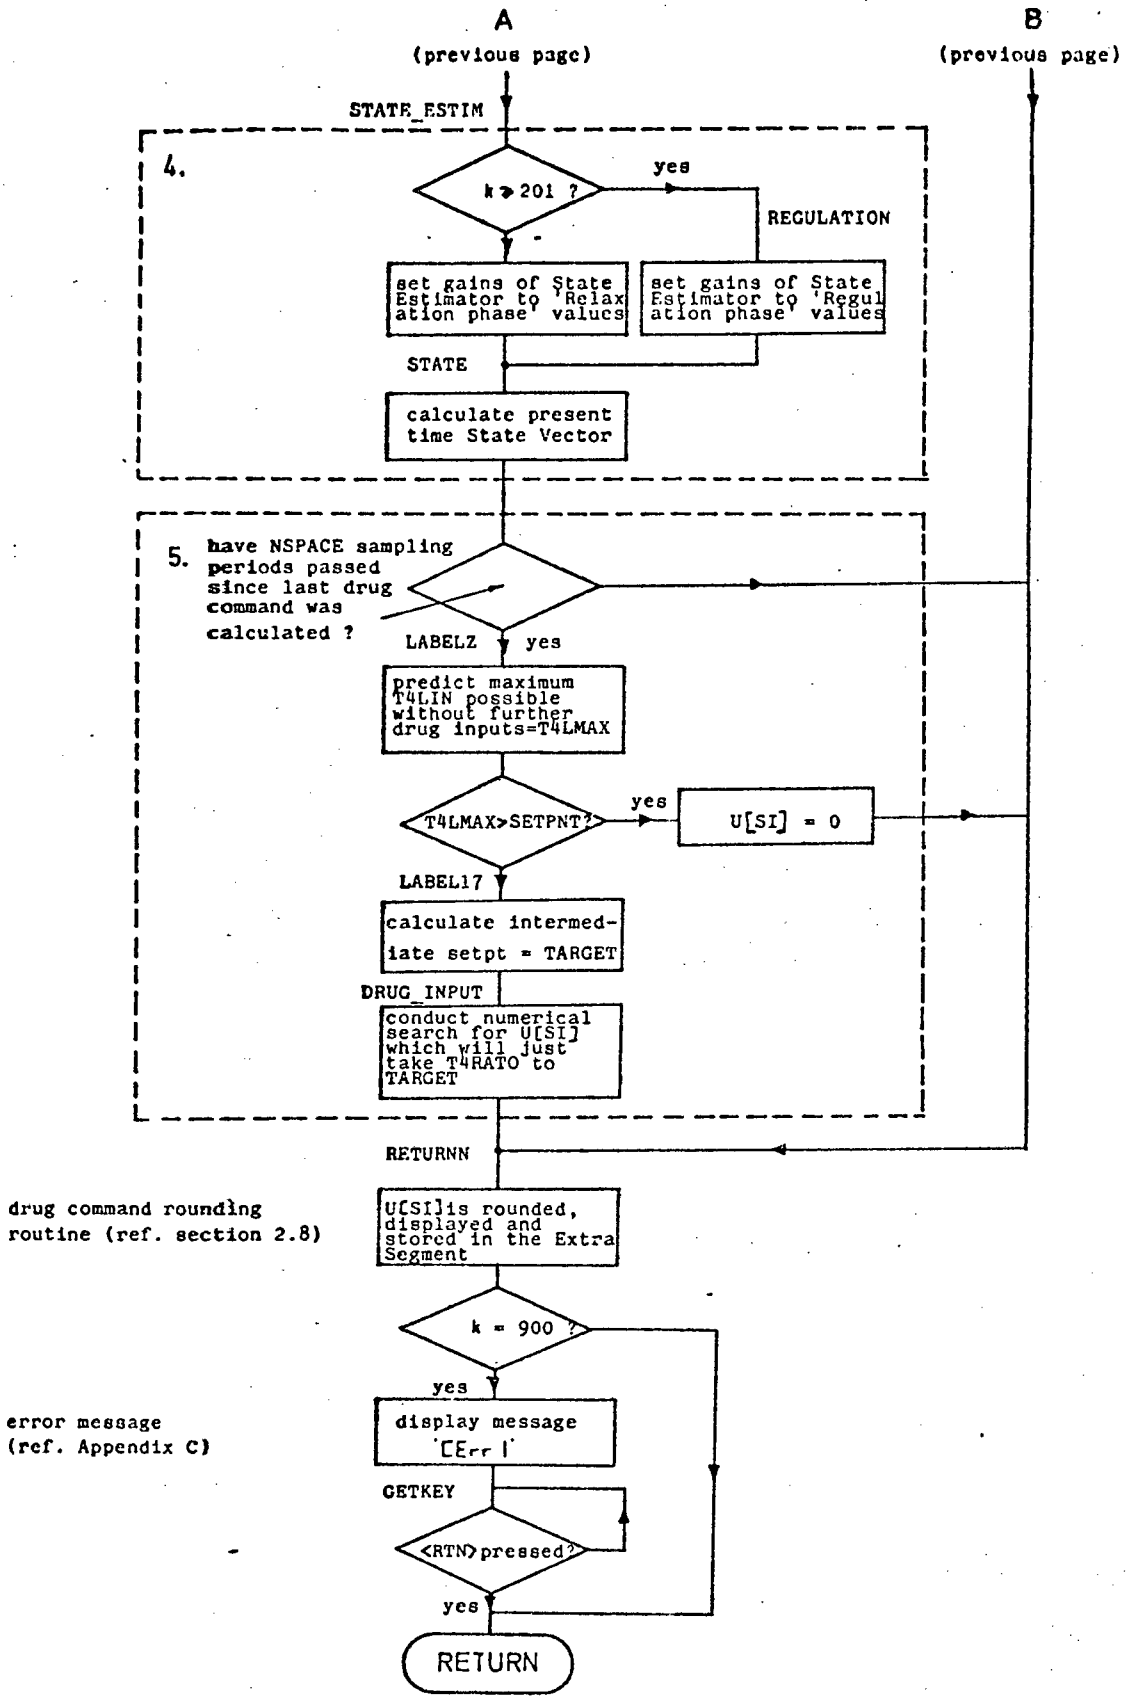

Figure 5.1 : Macro-flowchart of the BOLUS4 controller as implemented on the microcomputer showing the major functions and associated branching.

2. The parameter estimation algorithm

3. The Initial Phase Controller

4. The State Estimator

5. The Relaxation/Regulation Phase Controller.

and are discussed separately in sections 5.2 to 5.6.

There are in addition a number of minor functions performed by the controller as indicated on the flowchart. Marginal notes are included there which refer to the relevant section of the thesis where the operation is explained.

The arrays WEIGHT and WEIGHT2 are cleared before first entry to the control algorithm. Each sampling period that the routine is entered, the appropriate element of the WEIGHT2 array is set to 1. This occurs whenever the controller is not in the "pause" state. This array is used by the State Estimator. The WEIGHT array determines which twitch data are used by the parameter estimation algorithm. Elements are set to 1 for twitch ratios between 10% and 80% and 0 otherwise.

5.2 Data linearization

It was explained in section 4.2 that the twitch ratio data is immediately linearized upon entry to the controller using

$$T_{lin}(k) = \ln\left[\frac{1 - T(k)}{T(k)}\right]$$

The same is true of the set-point. Thus

$$SP_{lin}(k) = \ln\left[\frac{1 - SP(k)}{SP(k)}\right]$$

The LOG_E subroutine of the BOLUS4 software module calculates

$$y = \ln\left[\frac{1 - x}{x}\right]$$

for this purpose. It calculates natural logarithms using the theorem

$$\log_e a = \frac{\log_2 a}{\log_2 e}$$

for the 8087 operates exclusively with logarithms of base 2.

5.3 Parameter Estimation algorithm

The unit vector in the direction of steepest descent is found by calculating the loss function between patient data and model thrice:

1. with the nominal parameter vector $\underline{\theta}_D = \left[p_1 ; p_2 ; g ; BT4\right]^T$ and resultant loss function VLOSS0.

2. with perturbation on BT4 $\underline{\theta}_D = \left[p_1 ; p_2 ; g ; BT4 + TUNE*SCALE\right]^T$ resulting in VLOSS1.

3. with perturbation on g $\quad \underline{\theta}_D = \left[p_1 \; ; \; p_2; g + \text{TUNE} \; ; \; BT4\right]^T$

resulting in VLOSS2.

TUNE is 2^{-13}. SCALE is the scaling factor of 1ØØ for BT4 as discussed in section 4.5.4. The unit vector for the two parameters then follows from equation 4.14 and is

$$\underline{\xi} = \frac{\left[DG \; ; \; DBT4\right]^T}{\sqrt{DG^2 + DBT4^2}}$$

where

$$DBT4 = VLOSS1 - VLOSSØ$$

$$DG = VLOSS2 - VLOSSØ$$

The linear search for the optimum distance in the direction of steepest descent proceeds according to the flowchart of figure 5.2 which is self-explanatory.

Subroutine ARITHØ of the BOLUS4 module calculates the loss function whenever it is required by the parameter estimation algorithm. It uses the subroutines UNIT_R and SUBSCRIPT. The operation of these two routines is described first.

Subroutine UNIT_R

Given equation 3.11 which defines the patient model

$$X(z^{-1}) = \frac{z^{-\ell} g}{(1 - p_1 z^{-1})(1 - p_2 z^{-1})} U(z^{-1}) \quad \ldots \quad 3.11$$

and equation 3.14 which defines the residual when using the inverse logistic

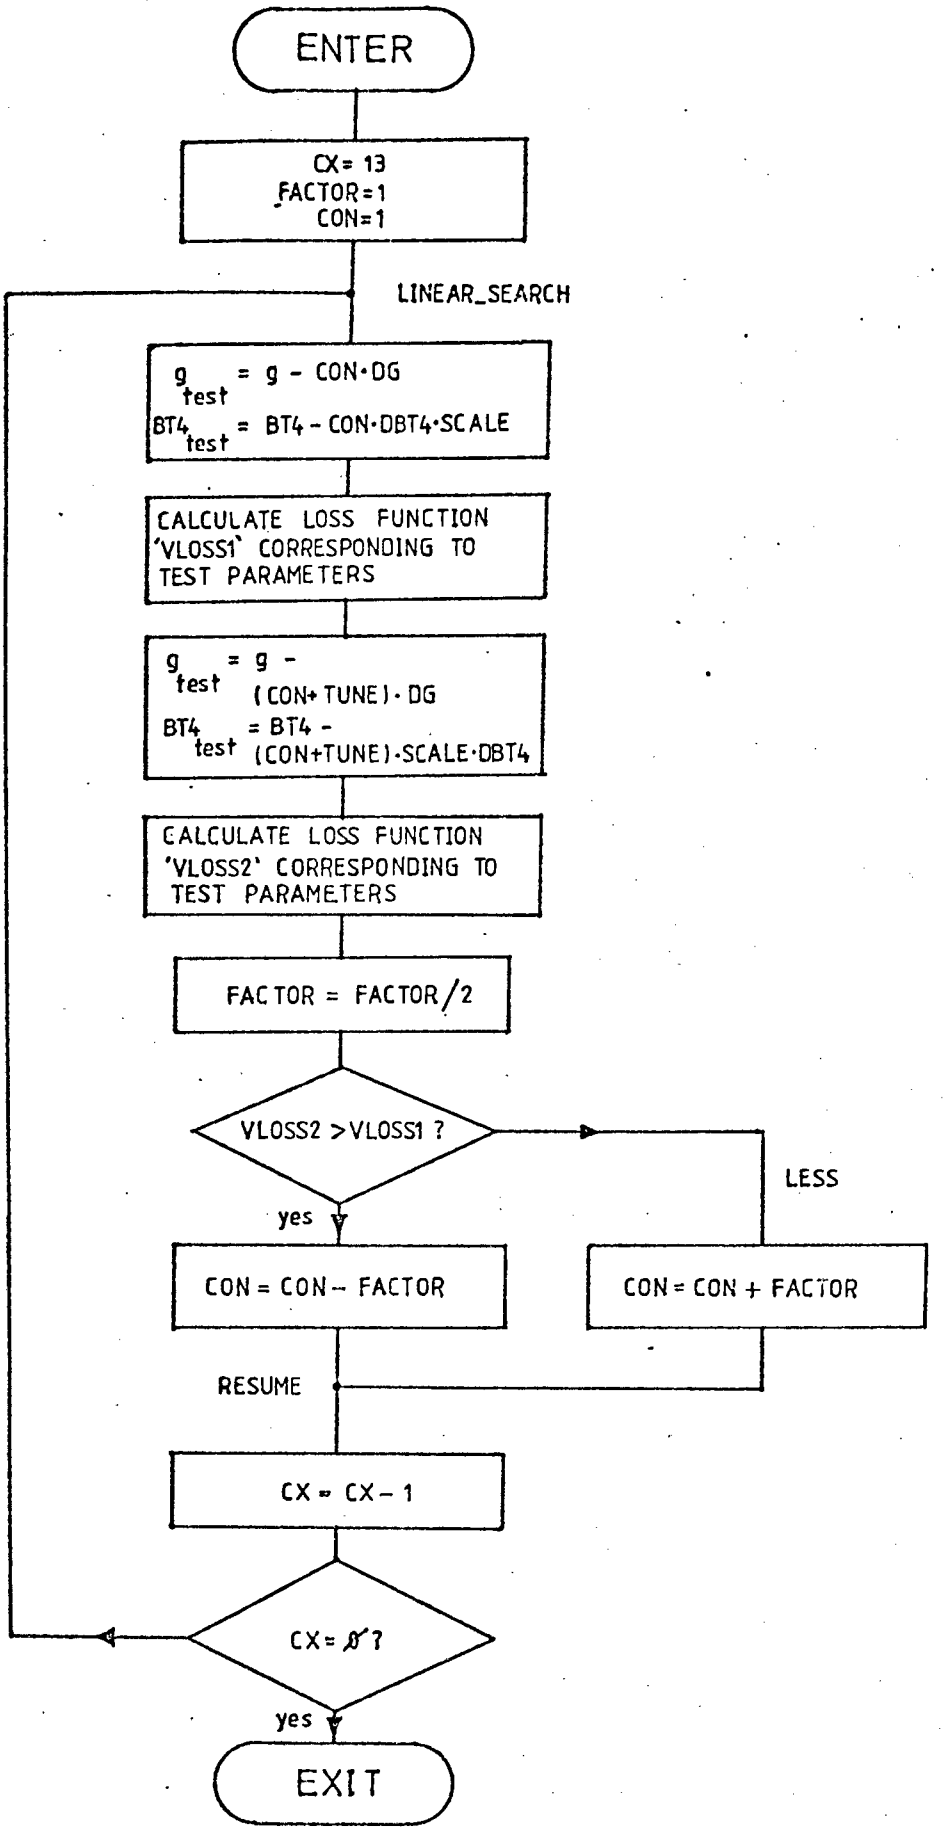

Figure 5.2 : Flowchart of the linear search used in the Steepest Descent parameter estimation algorithm.

relationship,

$$r_c(k) = \left[\ln\left[\frac{1 - T(k)}{T(k)}\right] - X(k) - BT4\right] \cdot W(k) \quad \ldots \quad 3.14$$

equation 3.11 can be written as a Difference Equation with changed nomenclature thus

$$\left.\begin{aligned} UNIT &= A1*XX[\emptyset] + A2*XX[4] + P[8]*U(k - LAGES) \\ x(k) &= UNIT \end{aligned}\right\} \quad \ldots \quad 5.1$$

where

$$\begin{aligned} A1 &= P[\emptyset] + P[4] \\ A2 &= -P[\emptyset]*P[4] \\ XX[\emptyset] &= x(k-1) \\ XX[4] &= x(k-2) \end{aligned}$$

and

$$\begin{aligned} P[\emptyset] &= p_1 \\ P[4] &= p_2 \\ P[8] &= g \\ LAGES &= \ell \end{aligned}$$

Equation 3.14 can be expressed as

$$R = T4LIN(k) - UNIT - P[12] \quad \ldots \quad 5.2$$

where

$$T4LIN(k) = \ln\left[\frac{1 - T(k)}{T(k)}\right]$$

$$P[12] = BT4$$

and the weight W(k) will be incorporated into equation 5.3 to follow. Subroutine UNIT_R calculates the variables UNIT defined in equation 5.1 and R defined in equation 5.2. The subscripts 'k' and 'k-LAGES' are DI and SI respectively in this routine. As this subroutine is used extensively, it was coded to hold the results of intermediate calculations on the 8087 stack thereby minimizing execution time.

Subroutine SUBSCRIPT

Subroutine SUBSCRIPT initializes LOOPCOUNT, SI and DI as follows

LOOPCOUNT = NFIRST

SI = 4.(NFIRST-LAGES)

DI = 4.NFIRST

The summation limits NFIRST and NLAST are set by the calling routine.

Subroutine ARITHØ calculates the loss function defined in equation 5.3

$$VLOSS = \sum_{\substack{LOOPCOUNT \\ = NFIRST}}^{NLAST} R^2 * WEIGHT[DI] \quad \ldots \quad 5.3$$

according to the simplified flowchart of figure 5.3. The summation limits NFIRST and NLAST are defined in subroutine BOLUS4 and are equal to LAGES (typically 5) and k (the present sampling period) respectively when used with ARITHØ.

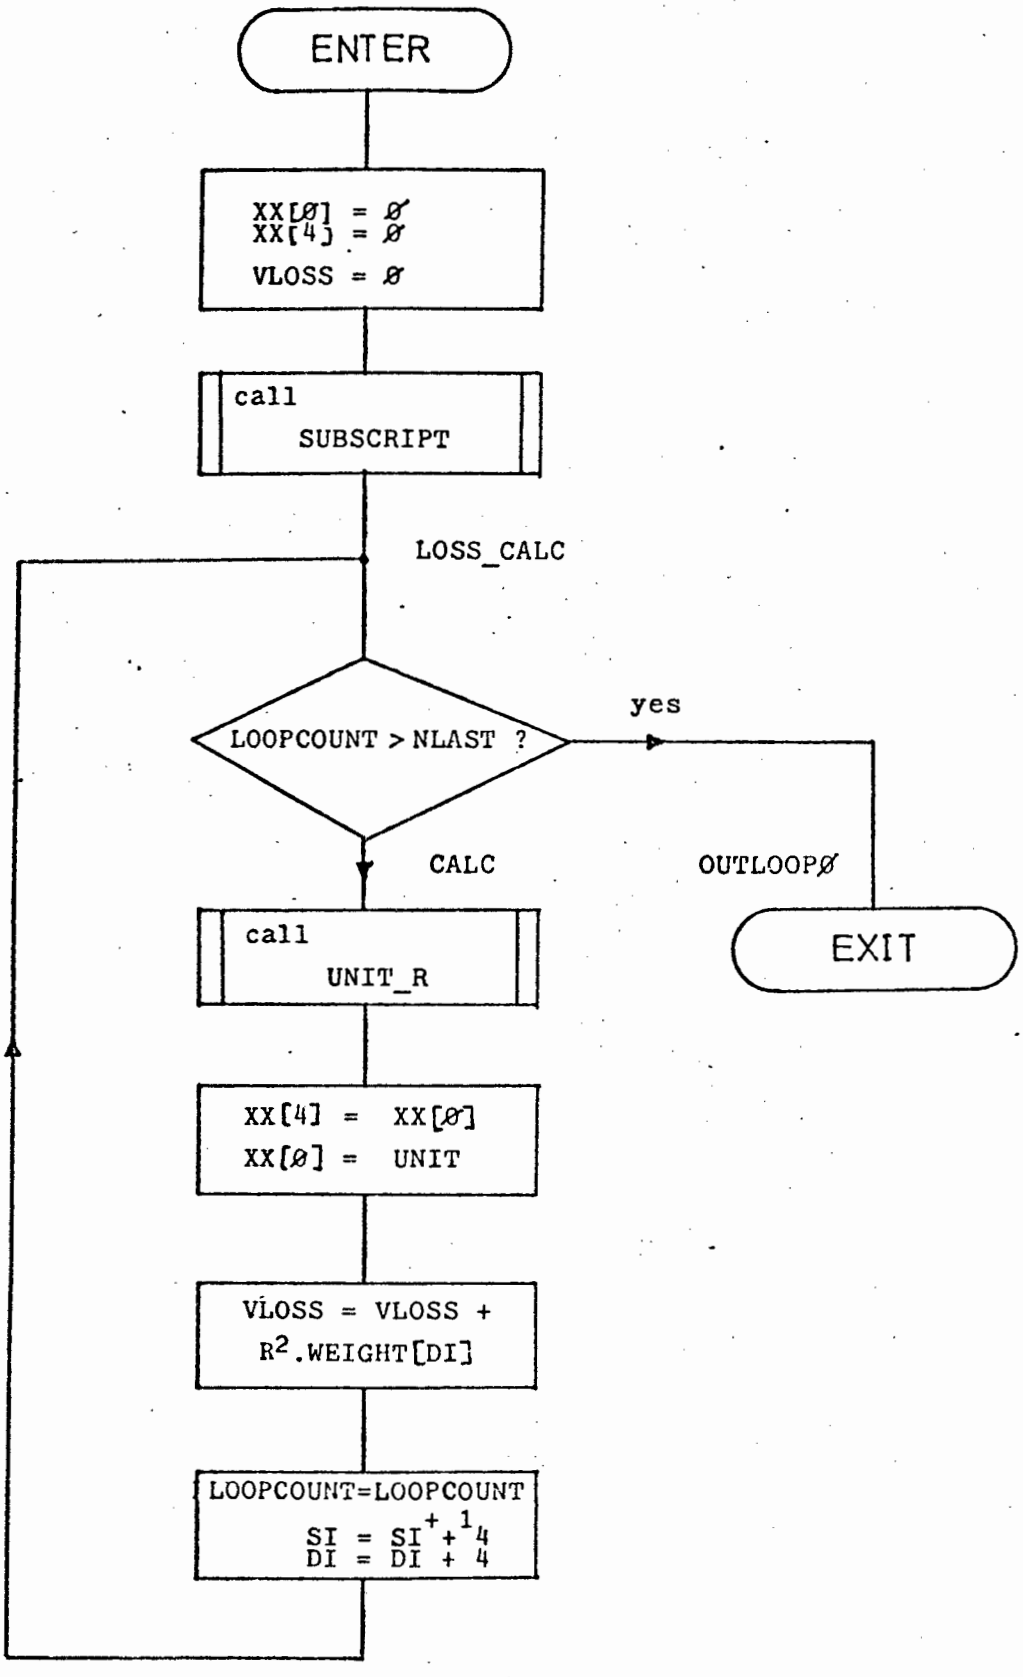

Figure 5.3 : Simplified flowchart of subroutine ARITHØ.

5.4 State Estimation

Subroutine ARITH1 performs the state estimation. It is similar in construction to subroutine ARITHØ. Its operation is shown by the simplified flowchart of figure 5.4. The limits NFIRST and NLAST are LAGES and k respectively. The output of the routine is the present time estimated state vector

$$\begin{bmatrix} X[\emptyset] \\ X[4] \end{bmatrix} = \underline{\hat{X}}_T(k)$$

Two choices of the gains G1 and G2 are required. The first is in use from the start until parameter estimation is inhibited at sampling period 2ØØ - this includes the relaxation phase and a portion of the regulation phase. The other choice of gains begins at sampling period 2Ø1 and continues throughout the regulation phase as shown in the macro flowchart of figure 5.1. In both cases the gains are calculated off-line as the steady state limit of the Kalman gains for a given noise variance ratio R_w/R_e and given patient model parameters.

5.5 The Relaxation/Regulation Phase Controller

Subroutine ARITH2 is the heart of the Relaxation/Regulation Phase Controller. It calculates the greatest depth of relaxation which will be achieved for given patient model parameters and a known sequence of drug commands. Its construction is similar to subroutine ARITHØ. The limits NFIRST and NLAST are equal to 'k+1' and 'k+1ØØ' respectively. The principle of operation can be deduced from the flowchart of figure 5.5. On entry the initial states are set equal to the smoothed present time values calculated by the State Estimator. The UNIT_R subroutine is then called within a loop which at each

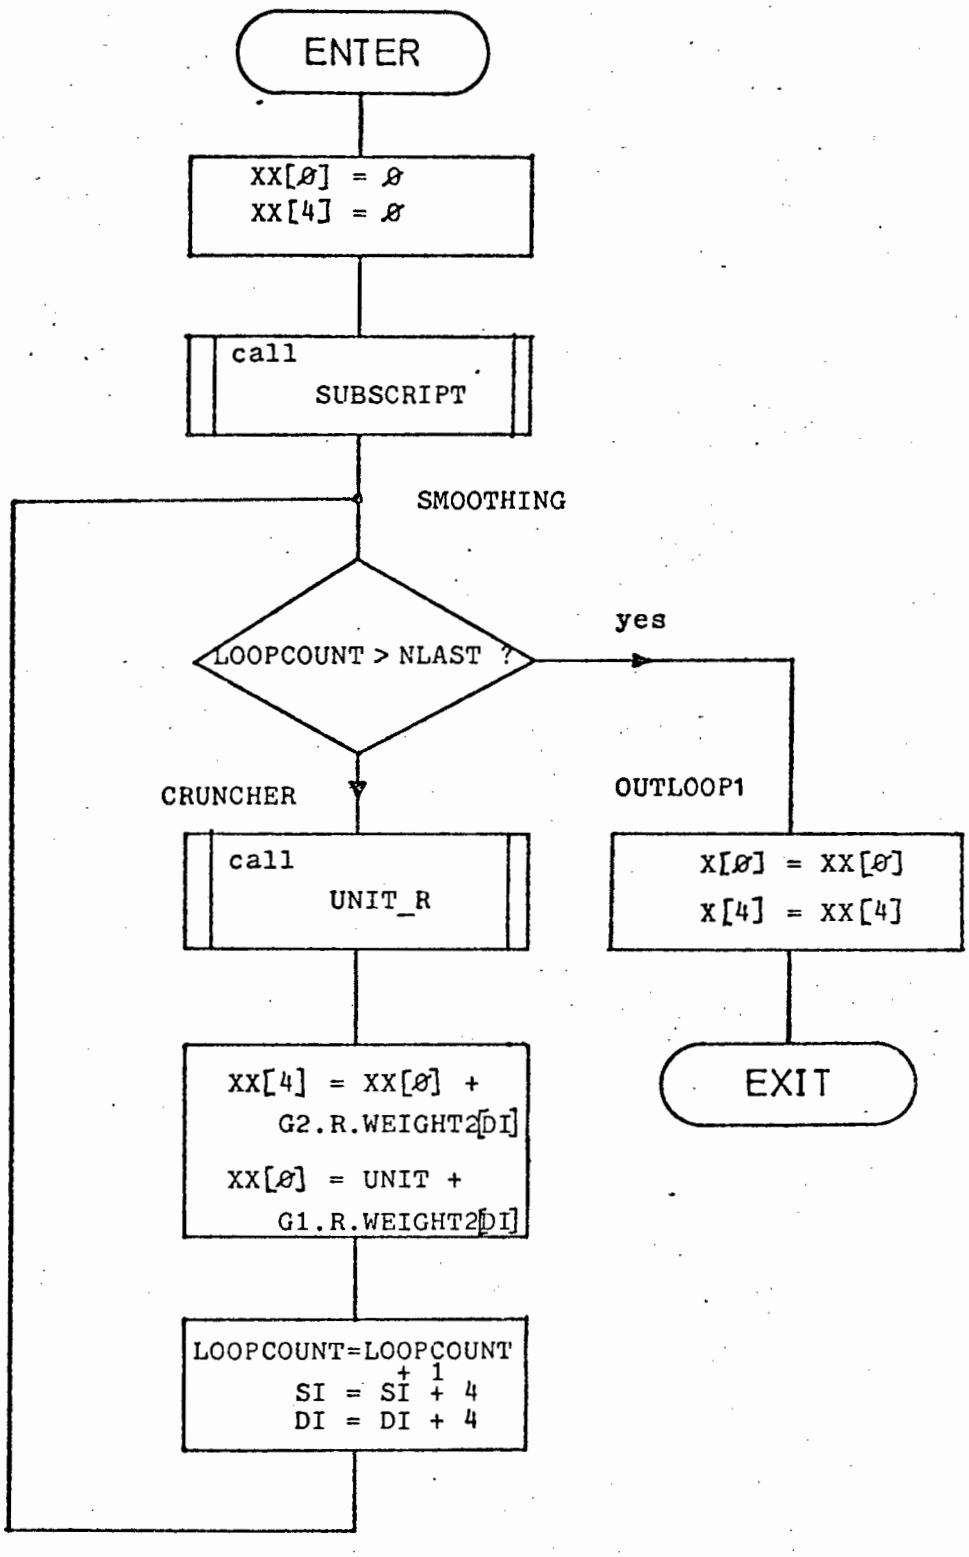

Figure 5.4 : Simplified flowchart of subroutine ARITH1.

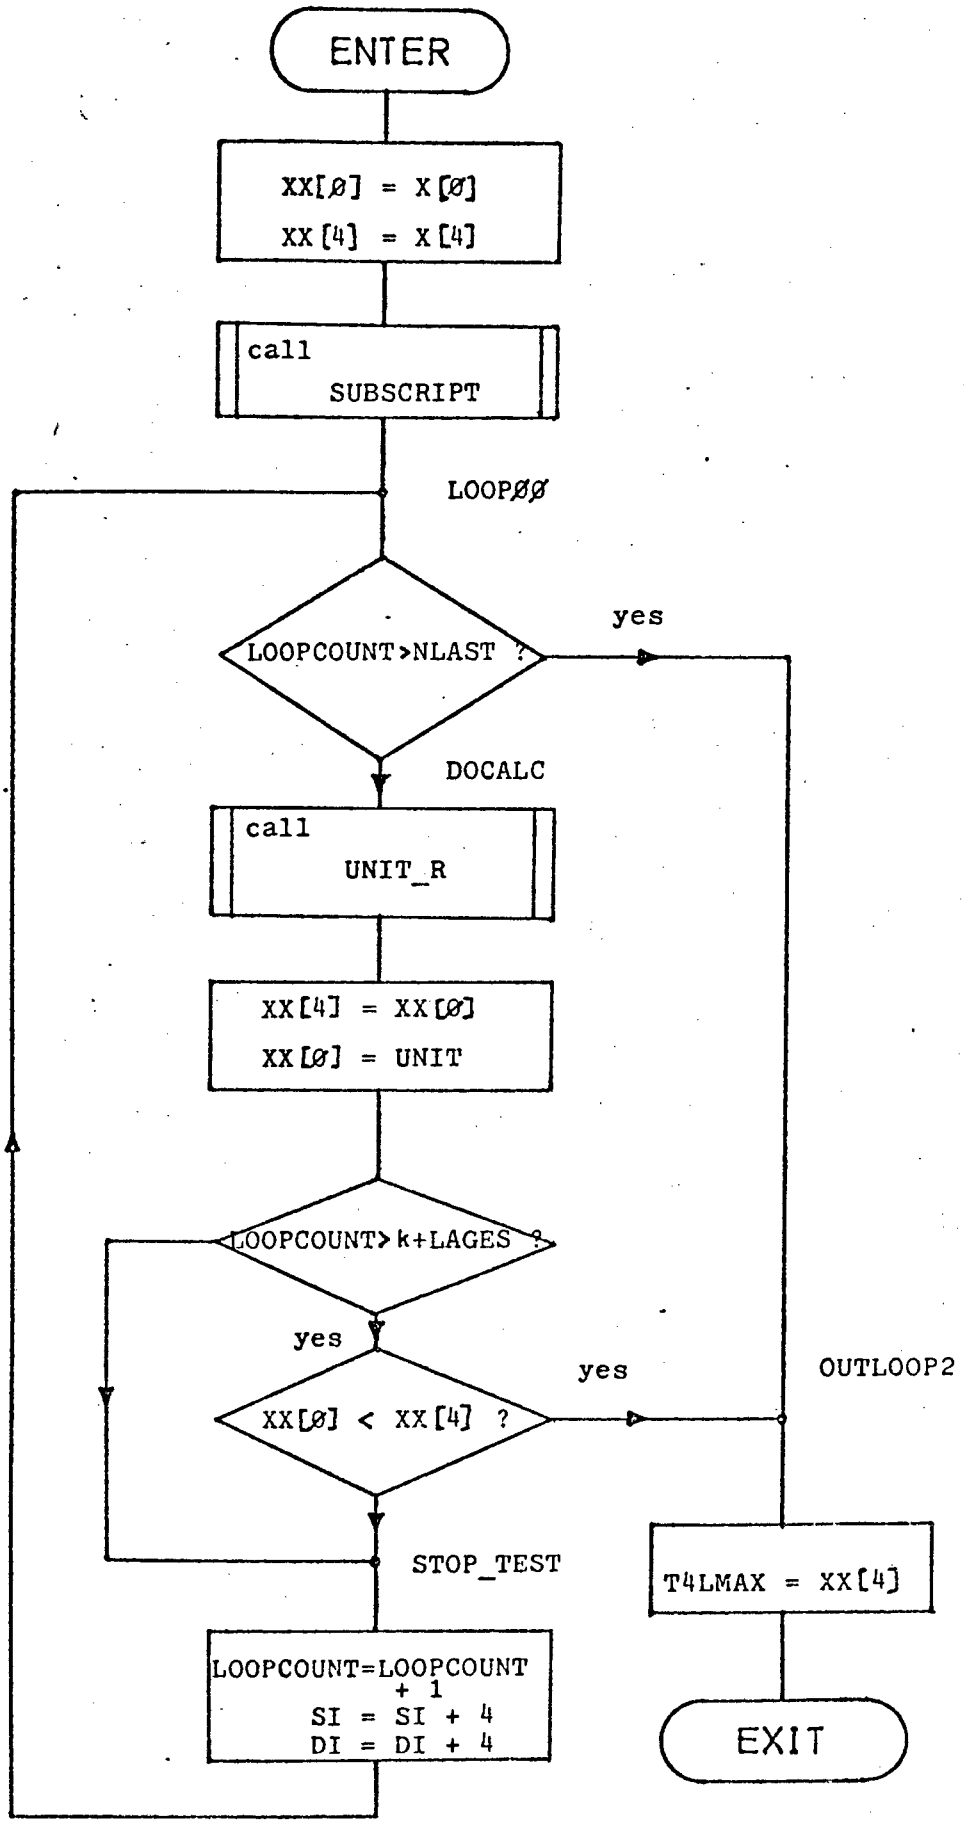

Figure 5.5 : Simplified flowchart of subroutine ARITH2.

iteration predicts the states at the following sampling period beginning at NFIRST (which is set to one sampling period ahead of present time) and terminating when the maximum has been located (i.e. when XX[Ø] < XX[4]) or when LOOPCOUNT passes NLAST. (This latter exit is a safety precaution. NLAST is made large enough - 1ØØ sampling periods ahead of the present time - that the maximum level of relaxation will always be found first unless an absurd choice of patient poles is made). Finally T4LMAX is set equal to the maximum level of relaxation that was found.

The Relaxation/Regulation Phase Controller calculates a drug command once every NSPACE sampling periods (typically 6). In these cases subroutine ARITH2 is first used to determine whether the level of relaxation induced by the drug already administered will reach or pass the set-point. If this is found to be so, the present drug command is set to zero and the calculation is thus completed.

If it is estimated that the previously administered drug is insufficient to take the patient to the set-point, the necessary supplementary dose is calculated. This is found numerically by the algorithm which is flowcharted in figure 4.7 and the details are not repeated here except to note that the set-point used in this calculation is not necessarily that chosen by the user. Rather it is a trial value (referred to as TARGET in the software) which is brought from zero to the chosen set-point in a number of steps. The flowchart of figure 5.6 shows how TARGET is calculated. It is based on equation 4.2. The counter INDEX2 (set to zero before first entry to subroutine BOLUS4) is incremented each time a drug command is calculated i.e. once every NSPACE sampling periods.

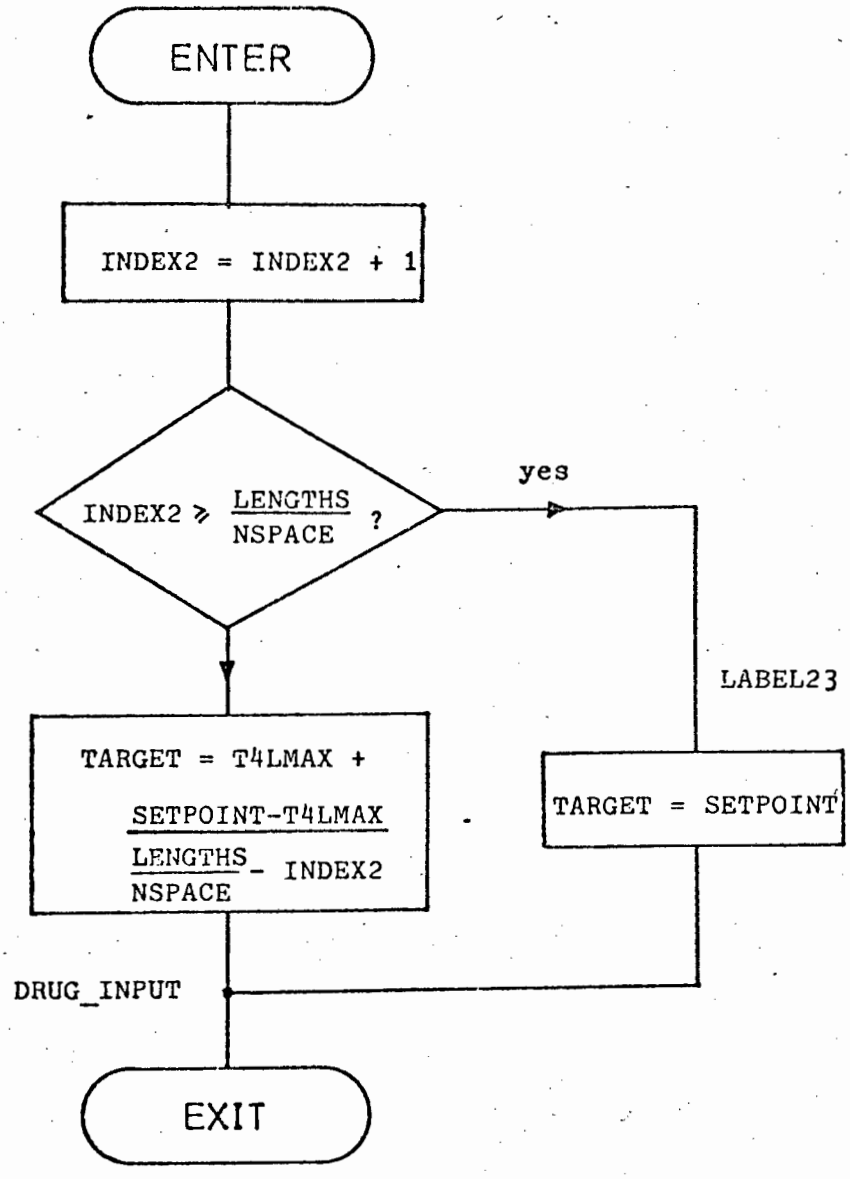

Figure 5.6 : Flowchart which shows how the trial set-point TARGET is found.

5.6 The Initial Phase Controller

The purpose of the Initial Phase Controller is to administer drug from the start of the operation until the twitch ratio being controlled drops below Ø,85. The principle of operation was explained in section 4.5.1. and its implementation is now discussed.

The controller calculates drug commands once every 6 sampling periods starting at Ø. The software consists of 4 subroutines:

1. INITIAL_INPUT. This is called by subroutine BOLUS4 and calculates the initial sequence of injections. It calls subroutines GAIN_CALC and MASS_CALC. (Note that this subroutine would not be called if the measured twitch ratio were below Ø,85).

2. GAIN_CALC. This recalculates the gain element of the Extreme Parameter Vector such that the patient model so defined and excited by the sequence of doses already administered, just reaches the 85% threshold at the time the subroutine is called.

3. MASS_CALC. This calculates the present time drug command which will just take the patient model defined by the Extreme Parameter Vector to the required set-point.

4. T4LIN_CALC. This is called by subroutines GAIN_CALC and MASS_CALC. It calculates the depth of relaxation reached by a patient model defined by the Extreme Parameter Vector given the sequences of doses already administered. This can be found either at a specified time or it can calculate the greatest depth of relaxation reached.

In the flowcharts of figures 5.7 to 5.1Ø discussed below, the elements of the Extreme Parameter Vector are defined as follows:

$$PE[Ø] = p_1$$
$$PE[4] = p_2$$
$$PE[8] = g$$
$$PE[12] = BT4$$

'SI' is an array index designating "present time".

The INITIAL_INPUT subroutine

The operation of this subroutine is shown by the flowchart of figure 5.7. When called the first time it passes the Extreme Parameter Vector chosen by the user to the MASS_CALC subroutine. This calculates the initial dose as that which will just take the patient model to the set-point.

When INITIAL_INPUT is called subsequently a somewhat different method is used. GAIN_CALC is first called to re-calculate the gain parameter.

Until the output of the patient model calculated on the basis of the Extreme Parameter Vector reaches Ø,85, the re-calculated gain parameter GTEMP will be greater than that of the Extreme Parameter Vector (for only a more sensitive patient could hope to reach the Ø,85 threshold before the extreme patient defined by the vector in which case the INITIAL_INPUT routine would have been terminated). Thus until such time as the patient model response reaches Ø,85 no drug is administered and the Extreme Parameter Vector left unchanged.

The first time INITIAL_INPUT is called after the Extreme Parameter Vector patient model reaches Ø,85, the re-calculated gain element is smaller than that in the Extreme Parameter Vector and is used to replace it. The dose to be administered at this time instant is found by calling MASS_CALC. It is rounded to the nearest Ø,1 mg by subroutine U_CONVERT.

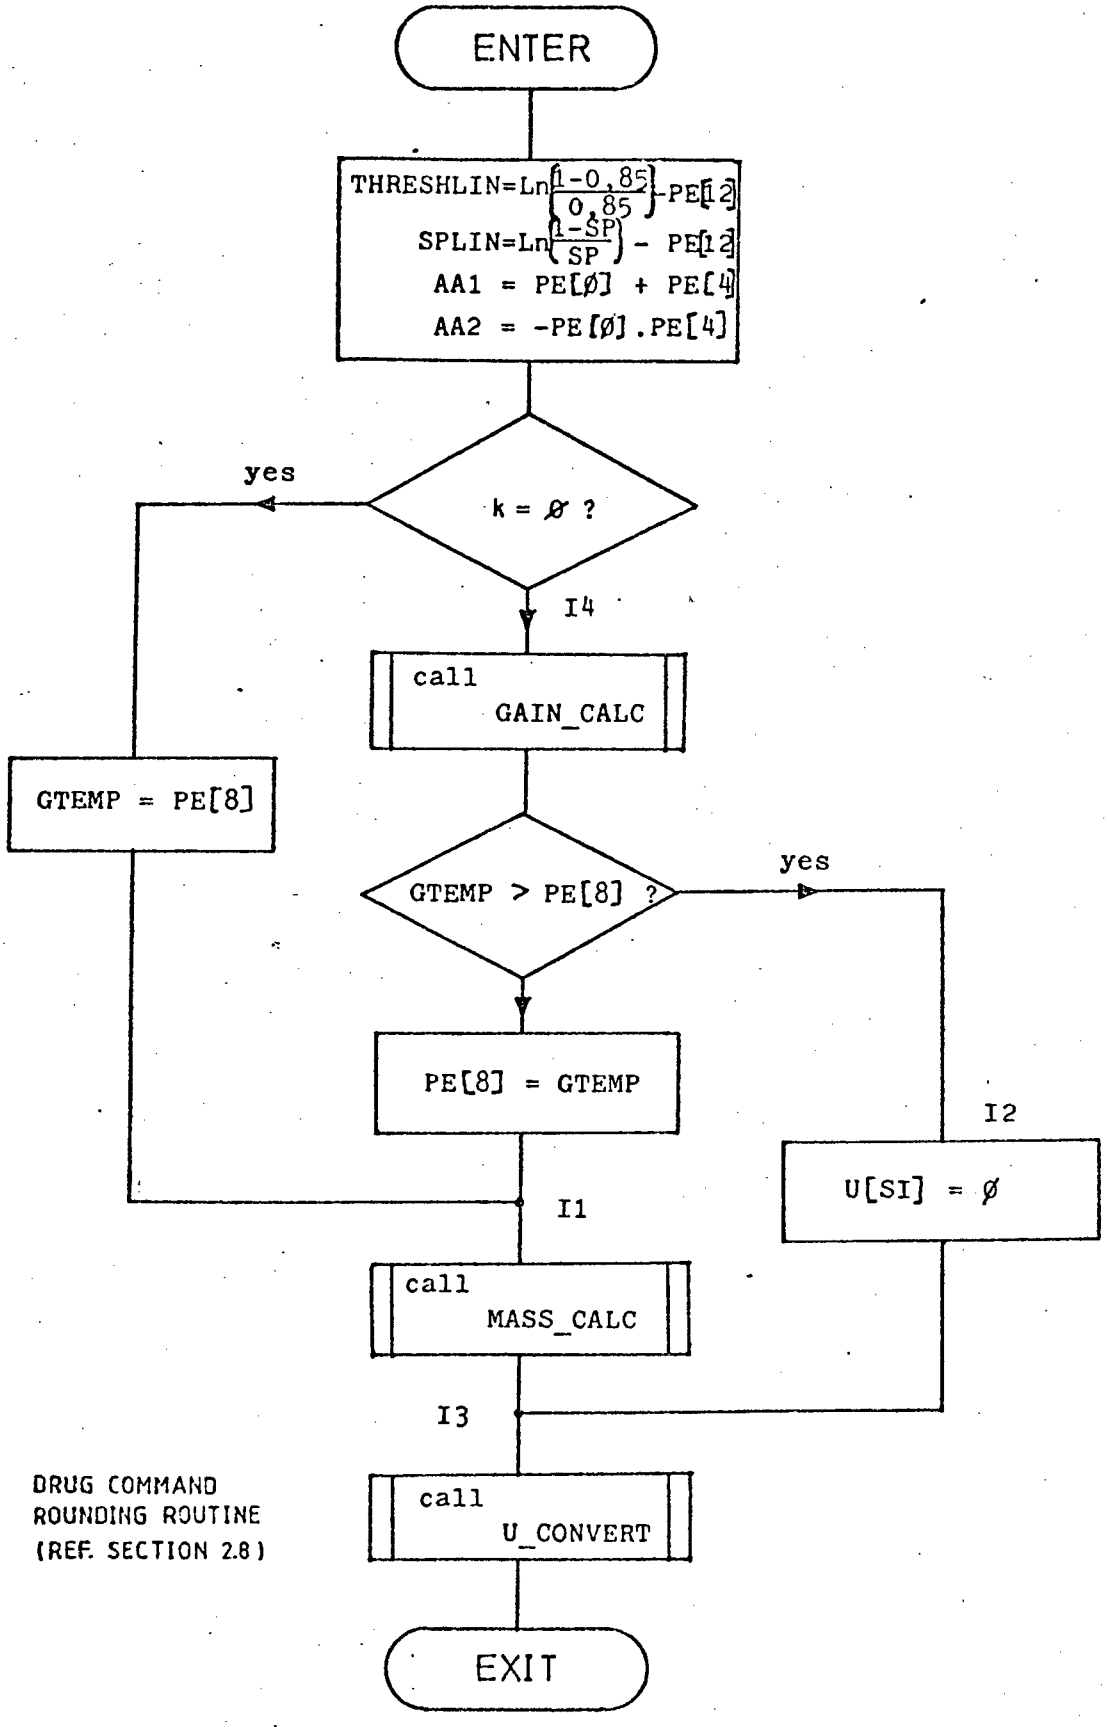

Figure 5.7 : Simplified flowchart of subroutine INITIAL_INPUT.

The GAIN CALC subroutine

The operation of this subroutine is shown by the flowchart of figure 5.8. Subroutine T4LIN_CALC (described below) is used to find the depth of relaxation 'T4LIN_END' at the present time 'k' from the known initial state of zero for the fixed sequence of doses already administered. This is done 1Ø times while the gain parameter is varied in a search for that gain 'GTEMP' which takes the patient model output to the Ø,85 ratio threshold at sampling period 'k'.

The MASS CALC subroutine

The operation of this subroutine is shown in the flowchart of figure 5.9. Subroutine T4LIN_CALC is here used to find the greatest depth of relaxation 'T4LIN_END' (which will occur some time after the present sampling period 'k') starting at the known initial state of zero and with a fixed Extreme Parameter Vector. This is done 1Ø times while a search is made for that drug bolus U[SI] to be administered at the present time which will just take the patient model to the set-point.

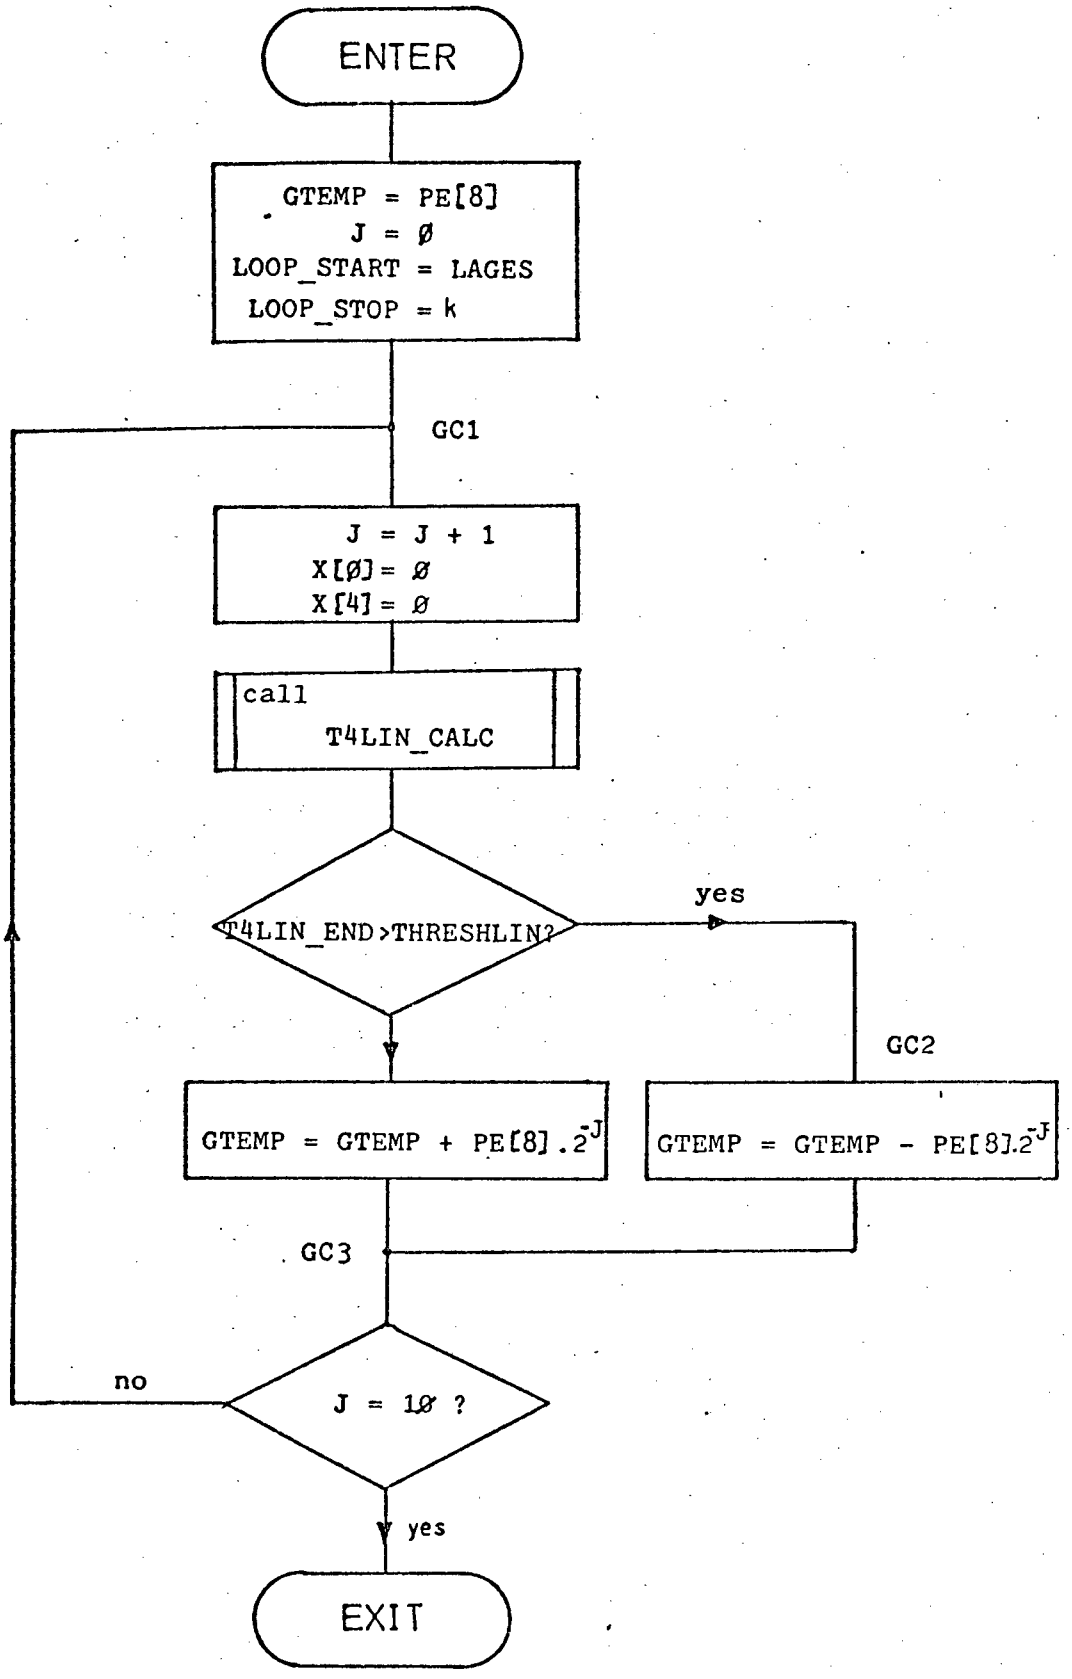

Figure 5.8 : Simplified flowchart of subroutine GAIN_CALC.

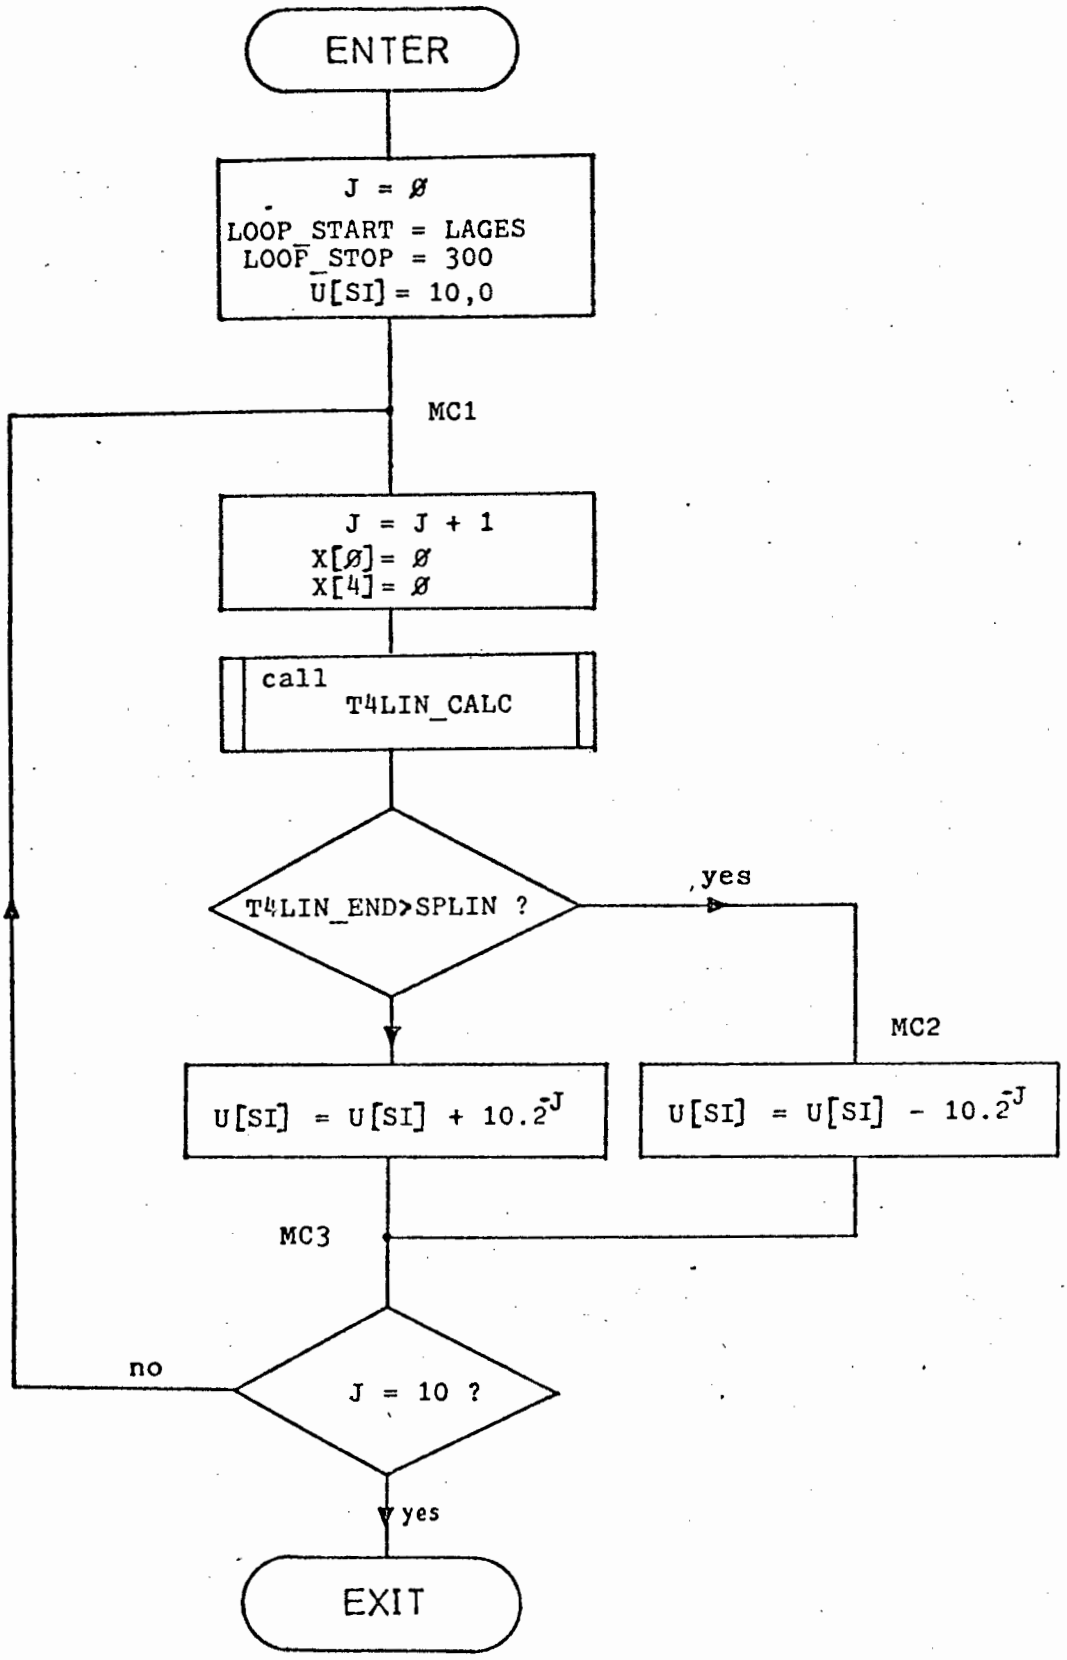

Figure 5.9 : Simplified flowchart of subroutine MASS_CALC.

The T4LIN CALC subroutine

The operation of this subroutine is shown in the flowchart of figure 5.10. The equation

XX[Ø] = AA1.X[Ø] + AA2.X[4] + GTEMP.U(LOOPCOUNT - LAGES)

used there predicts the level of relaxation at sampling period 'LOOPCOUNT+1' based on information up to and including sampling period 'LOOPCOUNT'. It is based on equation 5.1 and is executed in a loop. LOOPCOUNT is incremented at each iteration. Thus the predicted level of relaxation is calculated beginning at sampling period 'LOOP_START+1' and terminating at whichever of the conditions

1. the maximum level of relaxation

2. LOOPCOUNT = LOOP_STOP

is first encountered. This level of relaxation is retained as 'T4LIN_END' on exit from the subroutine. Note that LOOPCOUNT is shown in the flowchart to be incremented by 4. This is a consequence of the assembly language used. Successive elements of the Short-real arrays are stored 4 bytes apart.

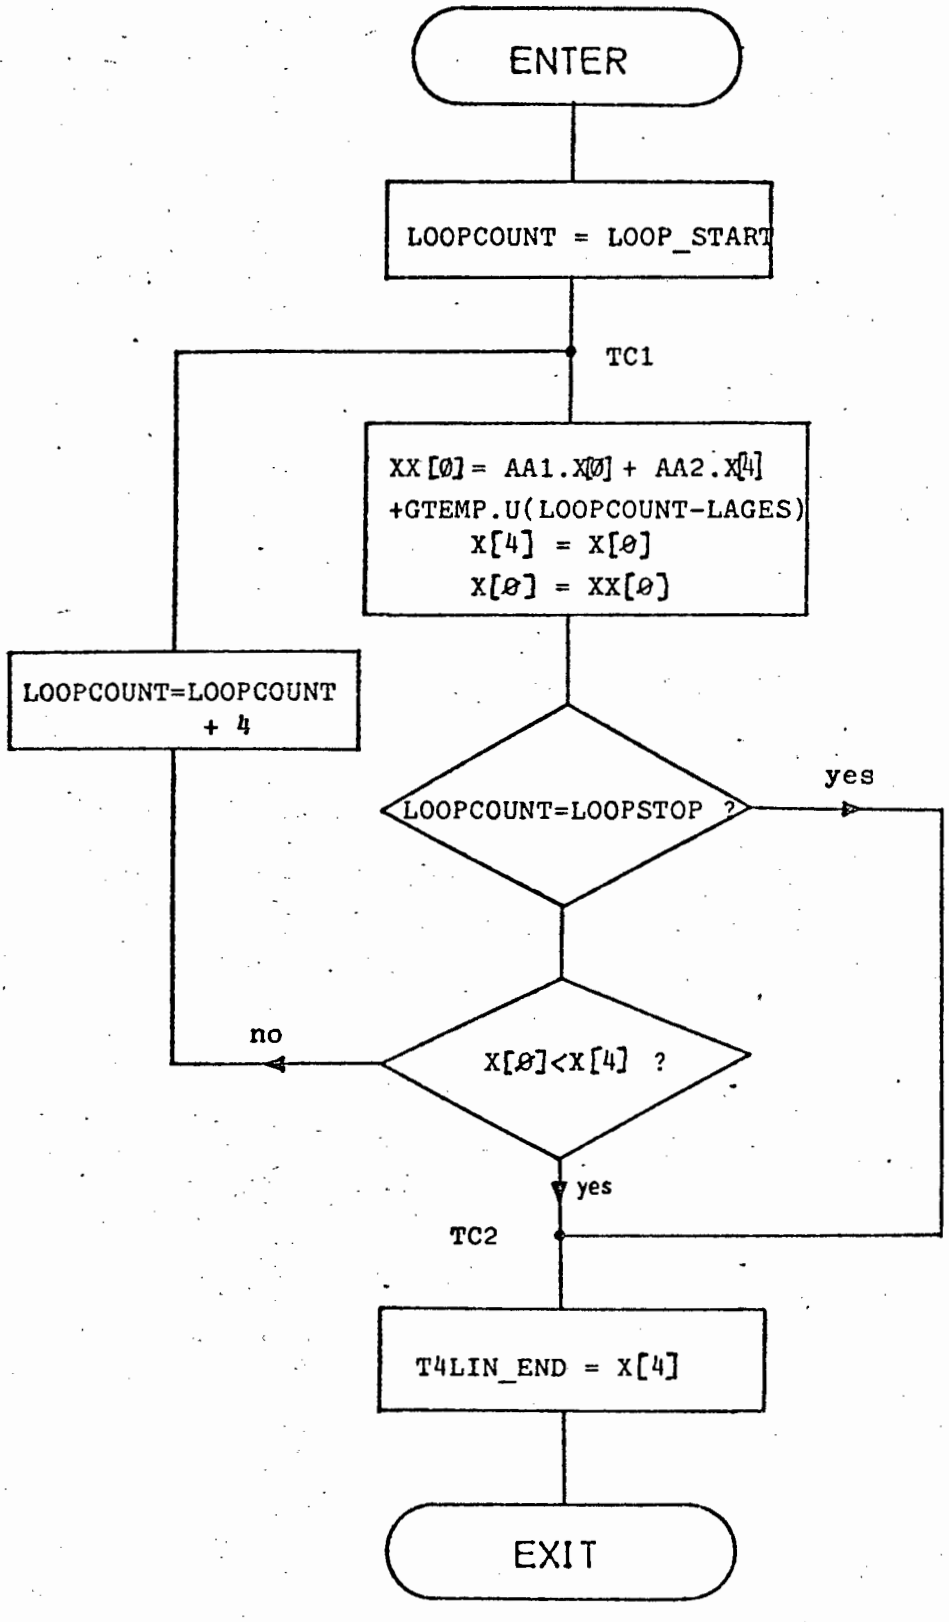

Figure 5.1Ø : Simplified flowchart of subroutine T4LIN_CALC.

CHAPTER 6

Controller Evaluation

6.1 Introduction

The purpose of this chapter is three-fold:

1. to decide upon controller settings for use in the clinical trials

2. to report on the clinical trials undertaken which are to show that the controller is reliable and useful. Here it is shown that the parameter estimation algorithm within the controller functions correctly, also that parameter estimation was useful in the trials undertaken. In addition an in-depth investigation is made of the controller for different choices of bandwidth of the State Estimator.

3. An improved method of initializing the controller is proposed. The actual controller performance from clinical trials and trials on simulations of these patients which use the improved method of initialization are compared to the ideal performance possible for dTC with any controller.

There are 3 batches of patients involved in the analysis of this chapter:

1. patients 1 - 1Ø, the original ten in which the T4 ratio was modelled but no automatic control done.

2. patients 11 - 26. These are taken from clinical trials on a group of 19 patients in which the T4 ratio was controlled

3. patients 27 - 44. These are taken from clinical trials on a group of 19 patients in which the T1 ratio was controlled.

The clinical conditions under which these patients were prepared have been stated in section 3.8.1.

6.2 Controller parameter settings

This section is aimed at finding suitable values with which to initialize the control algorithm and at finding the sensitivity of the controller to the other settings.

There are 6 settings which need to be decided upon:

- The Extreme Parameter Vector (for initial phase control)
- The Initial Parameter Vector (for relaxation and regulation phase control)
- The gains of the State Estimator
- The time spacing between drug commands
- The estimated patient circulatory dead-time
- The set-point.

1. The Extreme Parameter Vector

Two methods were used to select the Extreme Parameter Vector for the research. The resulting vectors are listed in Table 6.1. The gain parameter used in individual cases is found by dividing the third element of the parameter vector given in the table by the patient mass. This incorporates the reasonable assumption that patient sensitivity should be inversely proportional to body mass.

METHOD OF SELECTION	TYPE OF CONTROL	$POLE_1$	$POLE_2$	Average patient mass * gain	BT4
1	T4	0,9964	0,989	0,994	-3,52
	T1	0,9964	0,989	0,497	-3,52
2	T4	0,99574	0,973	1,541	-2,33
	T1	0,99709	0,972	1,172	-2,73

Table 6.1 : The choices of Extreme Parameter Vector used in the course of the research. The average patient mass is taken as 70 Kg in the table.

Method 1

This is based on the models of the original batch of patients 1 - 10.

In deciding upon the Extreme Parameter Vector for T4 control, the poles and gain were set to those of patient 7 who had the longest initial dead-time (slowest response). BT4 is the average value of the

10 patients. (In modelling these 10 patients the logistic relationship was used and the circulatory dead-time was fixed by visual examination of each patient's response. The range of validity of the patient data was 10% - 90% and not 10% - 80% as used subsequently. For this reason the parameter values differ from those given in Table 3.2).

The models of patients 1 - 10 pertain to the T4 response. In deciding upon the Extreme Parameter Vector for the T1 trials, the poles of patient 7 were again used for these ought to be fixed by the pharmacokinetics of the patient and thus should not be affected by which measurement is chosen to represent muscle relaxation. In the absence of further information, BT4 was left at the average value found for the 10 patients. The only data available with which to compare the T4 and T1 ratios pertained to patients 1 - 10. It appeared that approximately twice as much drug was needed to depress the T1 ratio by 80% compared to the T4 ratio and therefore the gain element of the Extreme Parameter Vector was halved for T1 control.

Method 2

This selection was made on completion of the clinical trials and was used only to control simulations of the patients entered into these trials.

The vector was calculated by programme CONTRL.VECTOR described in Appendix J and is based on the model parameters of a batch of patients. It is believed to be close to the best choice which can be made from available patient data.

Two batches of patients were used for the calculation - both taken from the clinical trials. The model parameters of patients 11-26 were

used to calculate a vector for T4 control and patients 27 - 44 were used to determine the vector for T1 control.

The calculation is done as follows:

The larger pole is taken to be the average of those in the batch. The largest * value of BT4 is selected in accordance with the description of the Initial Phase Controller given in section 4.5.1. The sensitivity and time to maximum effect is calculated for each patient model in the batch. The most sensitive* response and greatest* time to maximum effect are noted. Given this latter time and the larger pole already fixed, the other pole may be calculated for the time to maximum effect is determined by the poles alone. The gain parameter of the vector is then set to that value which will give the resulting model a sensitivity equal to the most sensitive of the batch.

* The extremes referred to above are defined as one standard deviation away from the mean value. This is done so that the values used are based on the statistical properties of the entire batch rather than on an isolated model value. If a Gaussian distribution is assumed, only 16% of the patients will exceed the extremes used.

2. The Initial Parameter Vector

Four methods were used to select the Initial Parameter Vector for the research. The corresponding vectors are listed in Table 6.2.

Method 1

The T4 response of the original batch of patients 1 - 10 was modelled. The logistic relationship was used to represent the pharmacodynamics and the circulatory dead-time was fixed by visual

METHOD OF SELECTION	TYPE OF CONTROL	$POLE_1$	$POLE_2$	GAIN	BT4
1	T4	0,9983	0,949	0,042	-3,52
2	T4	0,9970	0,980	0,020	-3,52
3	T4	0,9958	0,952	0,039	-1,25
	T1	0,9958	0,952	0,028	-1,25
4	T4	0,9957	0,956	0,024	-1,48
	T1	0,9971	0,955	0,0094	-1,50

Table 6.2 : The choices of Initial Parameter Vector used in the course of the research.

inspection of the patient response.

The Initial Parameter Vector was set to the mean of the model parameters so obtained except that the gain element was made 50% larger in order to allow for sensitive patients.

Method 2

The Initial Parameter Vector of method 1 was modified in the following way. The smaller pole was purposely increased from the mean of Ø,949 to Ø,98Ø to allow for a slower patient response to the drug and hence to avoid over-dosing such patients because the poles are not estimated on-line. The larger pole was set to Ø,997Ø. (There is negligible difference in controller performance for the cases that this pole is set to Ø,997Ø and to the mean of Ø,9983). BT4 was set to the mean identified value as in method 1. Finally the gain parameter was

set by visual inspection of the controller's performance on the 1Ø simulated patients. It was chosen as the smallest value (hence fastest response) for which there was no overshoot of the set-point in the relaxation phase.

This choice of parameters was used in the controller evaluation on the batch of simulated patients 1 - 1Ø.

Method 3

The patient responses in the first 13 clinical trials were examined visually. The T4 ratio was controlled in each of these. The single patient who had an average response (no. 19) was selected. The poles and BT4 in the fitted model of this patient were used in the Initial Parameter Vector. As this vector was to be used in further T4 clinical trials as well as T1 trials, the gain parameter was selected as follows:

- for the T4 trials it was increased by 4Ø% over that of the patient model as a safety precaution in controlling future sensitive patients.

- for the T1 trials it was left at the model value of Ø,Ø28. (As this is an average value for the T4 response it should be a conservative one for the T1 response as required).

The choice of the gain and BT4 parameters in the Initial Parameter Vector is less critical than in the Extreme Parameter Vector because these parameters are estimated on-line.

Method 4

In this case the vector is calculated by the programme CONTRL.VECTOR as described for the Extreme Parameter Vector. The calculation differs from the former case in that it is based on the average patient sensitivity and time to maximum effect response rather than the patient extremes.

The objection to method 1 is that it does not average the sensitivity of the patients to the drug as sensitivity is a non-linear function of all 4 model parameters. Consequently only 2 of the T4 clinical trials were performed using it.

Method 2 was used on 11 of the T4 clinical trials but in order to investigate what effect the Initial Parameter Vector has on controller performance it was changed to method 3 which was used for the remaining clinical trials.

Method 4 was not used in the clinical trials but in the control of simulated patients in the controller evaluation which followed.

3. The gains of the State Estimator

The gains of the linear State Estimator were chosen to be the steady state limit of the Kalman gains for a filter designed for an average patient model. The model was choice 2 of Table 6.2 except that the gain parameter was decreased by 20% from the conservative value of 0,02.

The steady state gains are calculated off-line by the programme CONTRL.RICATI described in Appendix B.

If the patient were absolutely noise-free, a large value of R_w ought to be best in every control situation, for this corresponds to a large bandwidth of the Estimator which compensates for modelling errors in the patient. Noise-free simulations were consequently made of patients 1 - 10 and each was controlled in two ways:

1. Parameter estimation and $R_w = 2$

2. Parameter estimation and $R_w = 0$

For all patients, except no. 10, method 1 was best as might be expected. Patient 3, the most insensitive demonstrates this behaviour most clearly and is shown in figure 6.1. It must however be noted that noise-free simulations promise over-optimistic results from the controller.

When noise is present, it is necessary to reduce R_w because a large bandwidth renders the estimated states prone to noise. The controller is then liable to overshoot during the relaxation phase and give poor regulation in the steady state.

In section 4.6.2 the sensitivity of the controller to R_w for noisy simulations of patients 1 - 10 was investigated. A value of 0,02 was found best for the relaxation phase when the controller used a linear State Estimator. The regulation phase value was set there to 10% of this.

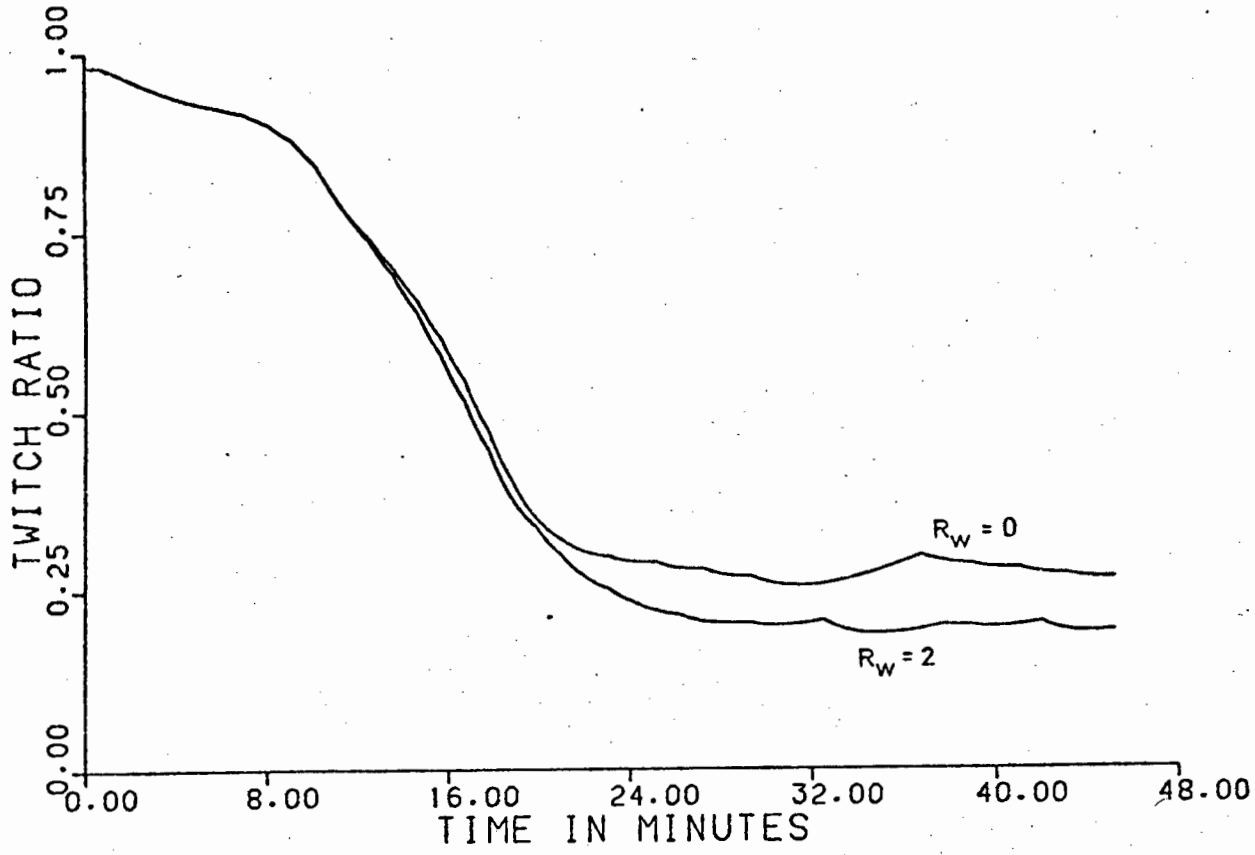

Figure 6.1 : The effect which the variance Rw has on closed loop control in a noise-free simulation of patient 3.

During the clinical trials, the nominal value of R_W chosen for the relaxation and regulation phases of control was Ø,Ø2 and Ø,ØØ2 respectively. However other choices were made in certain cases to establish the sensitivity of the controller to this parameter. These are discussed in section 6.6.

4. The time spacing between the control drug injections

In approximately the first half of the clinical trials the relaxant was administered manually at the prompting of the controller. (During this period the drug pump was not yet completed). It was therefore of interest to see how far the drug commands could be separated before the control deteriorated significantly. Simulated

control of patients 1 - 10 revealed that it could be extended up to about 2 minutes before this was so. See figure 6.2 which shows the control of the most sensitive patient no. 4 of this batch for spacings of 1,2 and 3 minutes. This patient showed the most rapid deterioration in control with increasing time between injections. In the light of this investigation a spacing of 1 minute was decided upon.

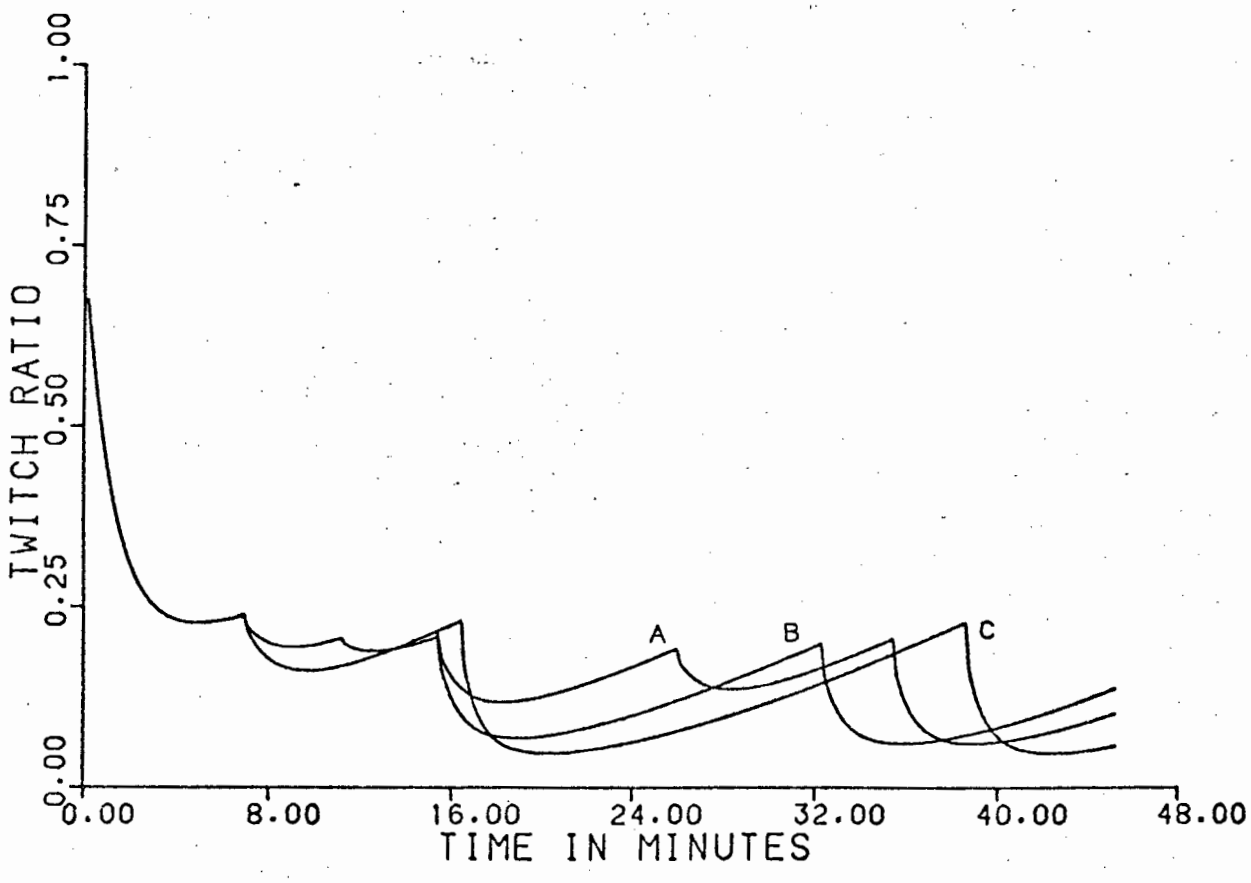

Figure 6.2 : The effect which different spacings between the drug injections has on controlling the most sensitive simulated patient (no. 4) of the original batch of 10 patients. The spacings are 7, 14 and 21 sampling periods (1, 2 and 3 mins.) for curves A, B and C respectively. A noise-free simulation is shown for clarity.

In the clinical trials the spacing between injections was set at 7 sampling periods in the first five trials and 6 (i.e. one minute) thereafter.

5. The estimated patient circulatory dead-time

The patient circulatory dead-time is not estimated on-line and consequently a fixed value must be assumed for it in the controller. The off-line analysis of chapter 3 showed that the mean fitted value of this delay for the 44 patients is 5 sampling periods (at sampling periods of both 9,Ø5 and 1Ø.Ø seconds).

It is of interest to know how sensitive the controller is to mis-matches in this parameter over the range likely to be encountered in practice. The batch of simulated patients 1 - 1Ø was consequently controlled with the FORTRAN version of the controller used in clinical trials. The circulatory delay of each patient was set in turn to 1,5 and 9 sampling periods whereas the controller assumed a value of 5 throughout. The average performance of the controller under these 3 conditions is shown in Table 6.3. The t-test reveals there is no

PERFORMANCE CRITERION	PATIENT CIRCULATORY DELAY (SAMPLING PERIODS)		
	1	5	9
Time to set-point (mins)	15,6±6,3	15,0±6,8	14,7±5,7
Initial overshoot of set-point (%)	7,0±5,6	4,6±2,7	7,9±5,1
Off-set from set-point in regulation phase (%)	4,5±3,1	4,1±3,0	5,2±3,1
Standard deviation from mean in regulation phase(%)	3,9±1,8	3,8±1,7	4,6±1,8
COLUMN	1	2	3

Table 6.3 : Average performance of the controller (version implemented on the microcomputer) on simulated patients 1 - 1Ø. The controller assumes a fixed circulatory delay of 5 sampling periods.

significant difference between the cases where it is mis-estimated (columns 1 and 3) and where it is correctly chosen (column 2). Sensitivity to this parameter is therefore quite weak.

The estimated circulatory dead-time was set at 5 sampling periods (50 secs) for the controller evaluation on the batch of simulated patients 1 - 10. During the the clinical trials it was set at 5 sampling periods in twenty two cases and 6 sampling periods in twenty trials.

6. The Set-point

The set-point was chosen as a twitch ratio of 0,2 for the controller evaluation on the batch of simulated patients 1 - 10. For the clinical trials it was chosen as 0,2 in twenty nine cases and 0,15 in ten. In two cases it was set initially at 0,2 and changed to 0,15 in the course of the operation. During one trial it was changed from 0,2 to 0,3.

6.3 The methods used to evaluate the controller

There are 3 issues which must be considered in judging the performance of the control algorithm implemented on the microcomputer:

1. The accuracy of the parameter estimation algorithm.

This is determined by verifying that the on-line parameter estimation routine locates the minimum of the loss function consistently in every clinical trial.

2. The efficacy of parameter estimation within the controller.

This is measured in terms of the four performance criteria defined in section 4.6.1 which are used to judge the controller's performance on the real patients as well as that of a number of alternative controllers on realistic simulations of of these same patients when modelled as described in chapter 3. The alternative controllers show the relative value of estimating different numbers of parameters or estimating them from better initial conditions. Finally the criteria are used to determine the ideal controller response i.e. that which the controller would have achieved had all the parameters of the patient model been known to it a priori.

The alternative controllers used in evaluating the efficacy of parameter estimation have the following features:

1. Only parameters g and BT4 are estimated. All conditions used in the clinical trials are duplicated. By comparing the performance of each simulated patient to the actual clinical control results for the same patients, this method is used to judge how realistically the patients have been simulated for control purposes.

2. The smaller pole is estimated in addition to g and BT4. This is to test whether there is an advantage to be gained in implementing a more complex algorithm for this application.

3. All 4 parameters are estimated. This is to determine the further value in estimating the larger pole.

4. No parameters are estimated. This is to show any advantages which were gained by parameter estimation in the trials undertaken.

The controller which estimates g and BT4 is also used with improved estimates of the Initial and Extreme Parameter Vectors obtained in the light of experience gained from the clinical trials. The purpose is to discover what improved performance this offers and to investigate whether the choice of the Initial Parameter Vector affects the controller performance when parameters are estimated.

3. The choice of bandwidth of the State Estimator

This investigation was approached in two ways:

1. Some of the clinical trials were conducted with bandwidths which differed from that given in section 6.2 which was derived from a control study on simulated patients 1 - 10.

2. In an off-line analysis, the controller was made to operate on the actual Input/Output data recorded during each clinical trial. The process noise variance R_w was then changed as desired and plots of the estimated level of relaxation were made for each choice. Then a visual examination of the smoothed twitch response superimposed on the raw data gives an idea of the sensitivity of the smoothed estimate to R_w.

In addition to the visual inspection, one can note the sequence of drug commands issued by the controller for different bandwidths of the

State Estimator set by R_w. Successive boli which fluctuate excessively in magnitude indicate that the controller is too noise prone (R_w too large) while boli which are consistently too large or too small indicate that the tracking ability of the Estimator is inadequate (R_w too small). (These boli are not used in the calculation of subsequent drug commands however for if they differ from those actually administered in the operating theatre they would cause a patient response other than that logged during the clinical trial).

6.4 Accuracy of the on-line parameter estimation algorithm

In 31 of the clinical trials the parameters g and BT4 that were estimated on-line were logged onto cassette tape together with the twitch data and controller commands. Thus it is possible to compute the residual sum of squares or loss function being minimized at any point in the control operation.

It is also possible to repeat the parameter estimation off-line on the Input/Output data record captured on tape using other methods of parameter estimation. This can be used to check the accuracy of the on-line parameter estimation method.

Accordingly, the loss function as minimized on-line was calculated at the end of the parameter estimation period (sample 200) from the data logged onto tape and this was compared to that of an off-line Least Squares estimation on the same data with the same parameters estimated but using subroutine VA05A of the Harwell library in place of Steepest Descent. The parameters not estimated were fixed at the nominal values assumed by the controller. This analysis was done on the university's UNIVAC 11/08 computer.

The comparison is made in Table 6.4. Also shown there are the loss functions obtained when no parameters are estimated (column 2), 3 parameters estimated (column 4) and all 4 parameters estimated (column 5).

The simpler Steepest Descent method is on average only 3% inferior to the other. In only 1 case is it more than 10% worse.

The F-test was applied to determine the significance of estimating increasingly many patient parameters. When g and BT4 were estimated, a significantly better fit over the no-estimation case was always obtained at a 95% confidence interval. When g, BT4 and the smaller pole were estimated, the loss function was also significantly reduced for every patient over the case where only g and BT4 were estimated. In 5 of the 23 cases however no additional significant improvement was obtained in estimating all four parameters compared with estimating g, BT4 and the smaller pole. These patients are indicated by asterisks in column 5.

The purpose of investigating the behavior of the loss-function as increasingly many parameters are estimated, is to determine whether there is a correlation with the improvement in control as the corresponding parameters are estimated on-line for use by the controller.

6.5 Controller performance and efficacy of parameter estimation

The results of the clinical trials and the performance of each of the controllers described in section 6.3 on simulations of these same patients were rated in terms of the 4 performance criteria discussed in section 4.6.1. These are tabulated in full in Appendix I for both T4 and T1 control. An abbreviated table was derived from the figures quoted in this appendix by averaging the

PATIENT	LOSS FUNCTIONS				
	ON-LINE ESTIMATION OF g , BT4	NO PARAM. ESTIMATION	OFF-LINE PARAMETER ESTIMATION		
			g , BT4	g,BT4 & smaller pole	g,BT4 & both poles
22	15,03	617,7	14,97	9,51	5,44
23	7,92	493,1	7,89	6,75	6,67 *
24	5,65	595,0	5,65	5,55	5,07
25	5,64	452,1	5,62	5,06	5,00 *
26	4,27	30,7	4,25	2,83	2,41
27	4,30	6729,6	3,57	3,13	3,12 *
28	5,13	602,9	5,09	4,77	3,23
29	4,30	1862,4	4,12	3,90	1,78
30	6,04	1980,7	5,93	5,18	5,17 *
31	29,84	2237,7	29,76	6,11	2,87
32	18,23	2627,6	17,94	3,79	2,32
33	27,78	17361,3	26,56	14,10	1,08
34	2,91	2042,2	2,77	2,49	2,36
35	4,34	1975,6	4,27	4,07	1,42
36	45,27	5232,6	44,91	3,00	2,38
37	30,42	1152,0	30,28	10,31	1,71
38	23,29	2066,4	23,28	5,99	1,41
39	27,47	1639,1	27,46	14,62	6,71
40	11,35	9151,7	10,45	2,19	1,70
41	11,70	8460,1	10,87	2,02	1,48
42	13,05	141,0	13,01	3,13	2,57
43	3,32	25,3	3,05	2,15	2,14 *
44	8,68	15,8	8,58	7,58	7,50
COLUMN	1	2	3	4	5

Table 6.4 : Comparison of loss functions as calculated on-line within the controller and off-line by a Least Squares method using subroutine VAØ5A. The patients for whom estimation of the larger pole did not result in a significant reduction of the loss function compared to the case where the other 3 parameters are estimated are indicated by an asterisk in column 5.

patients' responses for each of the criteria. These average values are contained in Table 6.5.

Column 1 gives the results for the actual clinical trials performed, while columns 2 to 10 give the results of various simulation studies using the patient models obtained from these trials.

The results of the clinical trials

Column 1 shows that the average time taken to induce relaxation to the set-point for the real patients is about 12 minutes and the average overshoot about 6%. Thereafter relaxation is maintained within an average of 2% of the set-point during the regulation phase. Therefore a lower bound of 16% can be placed on the set-point in order to avoid the twitch ratio dropping below 10% as required by Ham et al.,(1979).

The effectiveness of the patient simulations

The simulated control which duplicates the conditions of the clinical trials, i.e. only g and BT4 are estimated, was included to establish whether the patients were simulated realistically enough for control purposes. These results, shown in column 3, were compared with those of the clinical trials in column 1 using the Z-score test. By statistical test there is no significant difference between corresponding entries thus justifying comparisons between the various controllers.

The efficacy of parameter estimation

Column 2 is intended to show how the average performance of the controller deteriorates when no parameters are estimated on-line.It is therefore compared

TYPE OF PATIENT	REAL	SIMULATED							
EXTREME PARAMETER VECTOR	AS USED IN CLINICAL TRIALS						CALCULATED BEST	CLINICAL TRIALS VALUE	KNOWN a priori
INITIAL PARAMETER VECTOR	AS USED IN CLINICAL TRIALS					CALCULATED BEST			KNOWN a priori
PARAMETERS ESTIMATED	g,BT4	NIL	g,BT4	g,BT4, smaller pole	g,BT4, both poles	NIL	g,BT4	g,BT4	NIL	NIL
TIME TO SET-POINT (MINS.)	12,1±6,2	16,0±9,0	11,2±5,3	10,9±5,1	10,9±5,1	11,4±5,6	10,4±4,9	11,1±3,5	11,5±5,2	8,1±3,1
INITIAL OVERSHOOT OF SET-POINT (%)	5,9±3,4	3,2±4,1	4,9±4,6	4,9±4,3	5,0±4,5	3,9±4,2	4,4±4,3	3,8±3,7	2,9±3,9	0,9±1,6
OFF-SET FROM SET-POINT IN REGULATION PHASE (%)	2,0±1,6	0,3±2,1	2,7±1,7	2,2±1,6	2,0±2,1	2,0±1,6	2,1±1,2	2,0±1,3	1,1±1,2	0,8±0,9
STANDARD DEVIATION FROM MEAN IN REGULATION PHASE (%)	3,4±1,6	3,4±1,2	3,5±1,2	3,4±1,1	3,5±1,3	3,4±1,1	3,4±1,1	3,3±1,2	3,0±1,2	2,8±1,1
COLUMN	1	2	3	4	5	6	7	8	9	10

Table 6.5 : Average performance of the muscle relaxation controller in 34 sucessful clinical trials and a comparison of various other controllers on realistic simulations of the same patients.

with the simulation of column 3 in which g and BT4 are estimated. Now the induction period is lengthened significantly by 5 minutes on average when no parameters are estimated and the mean off-set in the regulation phase is significantly reduced. The other changes are insignificant.

In clinical terms the length of the induction period is of considerable significance whereas the off-set in the regulation phase is not.

The value of estimating the poles

Table 6.4 has shown that estimation of the smaller pole always resulted in a significant reduction in the loss function between patient and model over the case where just g and BT4 were estimated. The same was true in 18 of the 23 cases when the larger pole was estimated compared to the case where only the other three parameters were estimated. This significance is not reflected in the resulting performance of the controller however as columns 3,4 and 5 of Table 6.5 are not statistically different although the performance does improve marginally as first the smaller, then the larger pole is estimated. Further, this improvement is mainly confined to the regulation phase where is it of no clinical value. The reason for this is that the poles cannot be estimated from the beginning of the relaxation phase because there is a risk of controller instability if the estimation of these additional parameters is based on too few data. Simulations showed that 50 and 100 data need to be collected for reliable estimation of the smaller and larger poles respectively. Thus estimation of the smaller pole can begin only towards the end of the relaxation phase.

Choice of the Initial Parameter Vector in the clinical trials

Four methods were given in section 6.2 for choosing the Initial Parameter Vector of which three (methods 2, 3 and 4) warrented serious consideration.

The controller's performance in the clinical trials was investigated separately for the two batches of patients in which the Initial Parameter Vector had been chosen according to methods 2 and 3. The average performance for these two batches of patients is given in Table 6.6 from which it is seen that the change in Initial Parameter Vector shortened the induction period by 3,6 minutes on average (although this is not significant as judged by the Student's t-test). Thus a slight preference was given to initialization by method 3.

SELECTION OF INITIAL PARAMETER VECTOR IN SECTION 6.2	METHOD 2	METHOD 3
TIME TO SET-POINT(MINS)	14,3±8,1	10,7±5,5
INITIAL OVERSHOOT OF SET-POINT (%)	4,3±3,7	6,9±3,5
OFF-SET FROM SET-POINT IN REGULATION PHASE (%)	0,9±1,2	2,8±1,4
STANDARD DEVIATION FROM MEAN IN REGULATION PHASE (%)	4,1±1,7	3,2±1,5

Table 6.6 : Average performance of the muscle relaxation controller taken from clinical trials in which 2 different Initial Parameter Vectors were used.

Best choice of the Initial Parameter Vector

Selection of the Initial Parameter Vector by method 4 of section 6.2 was used in the control of simulations of the patients who underwent clinical trials.

Columns 6 and 7 of Table 6.5 show the average performance of two controllers so initialized. The controller of column 6 does not estimate

parameters on-line; that of column 7 estimates g and BT4. There is no significant difference between them. Neither is there a significant difference between controllers which estimate only g and BT4 when initialized differently (columns 3 and 7) except for the off-set (the poles are set differently in these cases and are not estimated on-line).

Choice of the Extreme Parameter Vector

This investigation is confined to simulated patients only. Two methods are given in section 6.2 for selecting the Extreme Parameter Vector. The average performance of the controller which uses method 1 is given in column 7 of Table 6.5 and that which uses method 2 is given in column 8. Both controllers estimate g and BT4 only and use the "Best" Initial Parameter Vector (method 4).

There is no significant difference between the average performances of these controllers.

Theoretical limits of controller performance

The performance of the controller for various numbers of parameters estimated was examined above. The greatest benefit which can be derived from parameter estimation in individual simulation trials is easily found by replacing the Initial Parameter Vector with the patient model parameters (assumed to be known a priori) in each trial and then inhibiting parameter estimation. This performance is the theoretical limit of which any parameter estimation algorithm is capable when used in conjunction with the controller.

If both the Initial and Extreme Parameter Vectors are replaced by the model parameters of the simulated patient being controlled and parameter estimation then inhibited, the resulting controller performance is not only the best which

can be obtained from this controller with any setting of the Parameter Vectors, it is in practical terms very close to the best of which any controller is capable. This is so because with the Extreme Parameter Vector so chosen, the transition to the set-point is made with a single bolus dose and without overshoot while the regulation performance is near ideal as the patient model parameters are exactly known throughout that phase of control.

A comparison between the average results of the clinical trials (column 1 of Table 6.5) and those of the best self-tuning controller (column 9) and best controller (column 10) is instructive for it shows that the actual performance achieved was very close to that ideally possible. (In this comparison all controller settings other than the Parameter Vectors were identical to those used in the clinical trials).

This section is concluded with 4 sets of graphs. The first of these, figure 6.3, shows the value of parameter estimation in the simulation trials on the batch of patients 1 - 10. The patient is the most insensitive (no. 3). The patient reaches the set-point in about 28 minutes when the parameters are estimated but has not reached it by 45 minutes when they are not. The convergence of the parameters g and BT4 which are estimated on-line is shown in figure 6.3 (b). The Initial Parameter Vector used was choice 2 of Table 6.2. In the light of the subsequent clinical trials it is known to be sub-optimal.

Figures 6.4, 6.5 and 6.6 show the control achieved clinically for three patients representing the extremes of response encountered during the T1 clinical trials. They were patients 44, 34 and 33 whose dTC consumption during the first hour of operation was least (figure 6.4), average (figure 6.5) and greatest (figure 6.6). The convergence of the parameters g and BT4 which are estimated on-line is also shown. It is seen there that parameter convergence is virtually complete when estimation is stopped at 33 minutes. This was true of all the patients. The dashed line in each of the figures is the desired

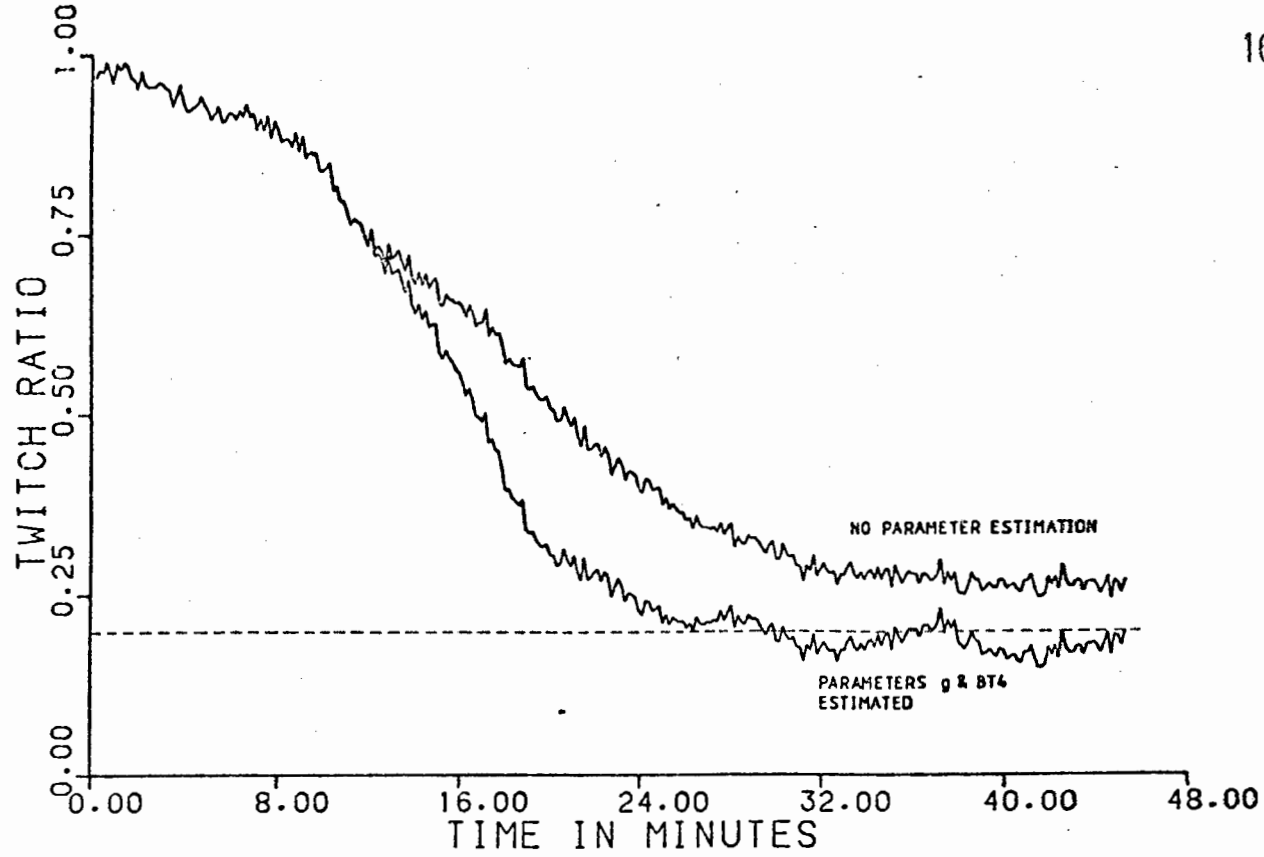

(a)

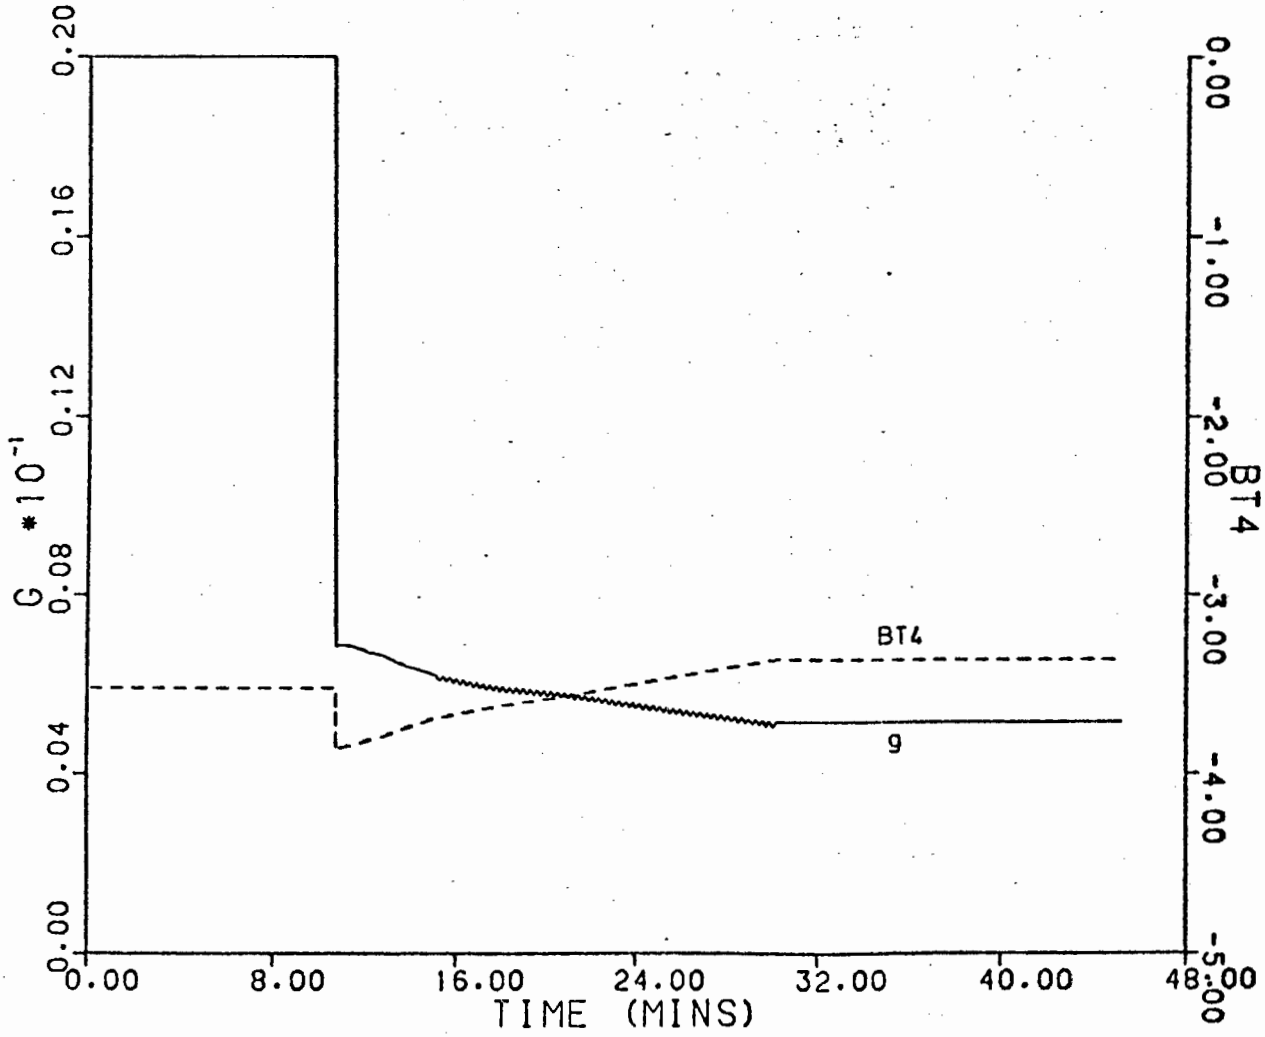

(b)

Figure 6.3 : The value of parameter estimation in the control of simulated patient 3 who is the most insensitive of the original batch of 10. (a) Control with and without parameter estimation. The dashed line is the desired set-point. (b) Parameters 'g' and BT4 as they are estimated on-line.

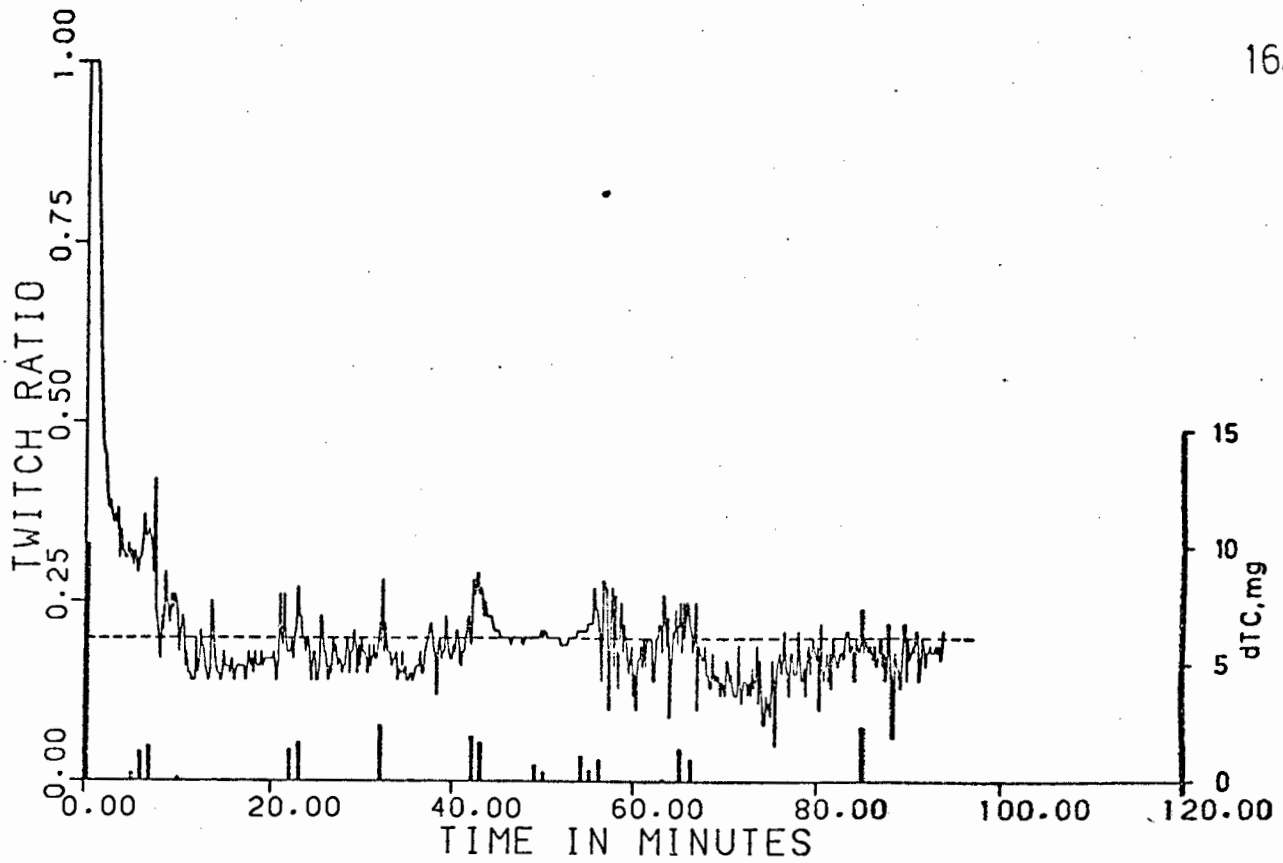

(a)

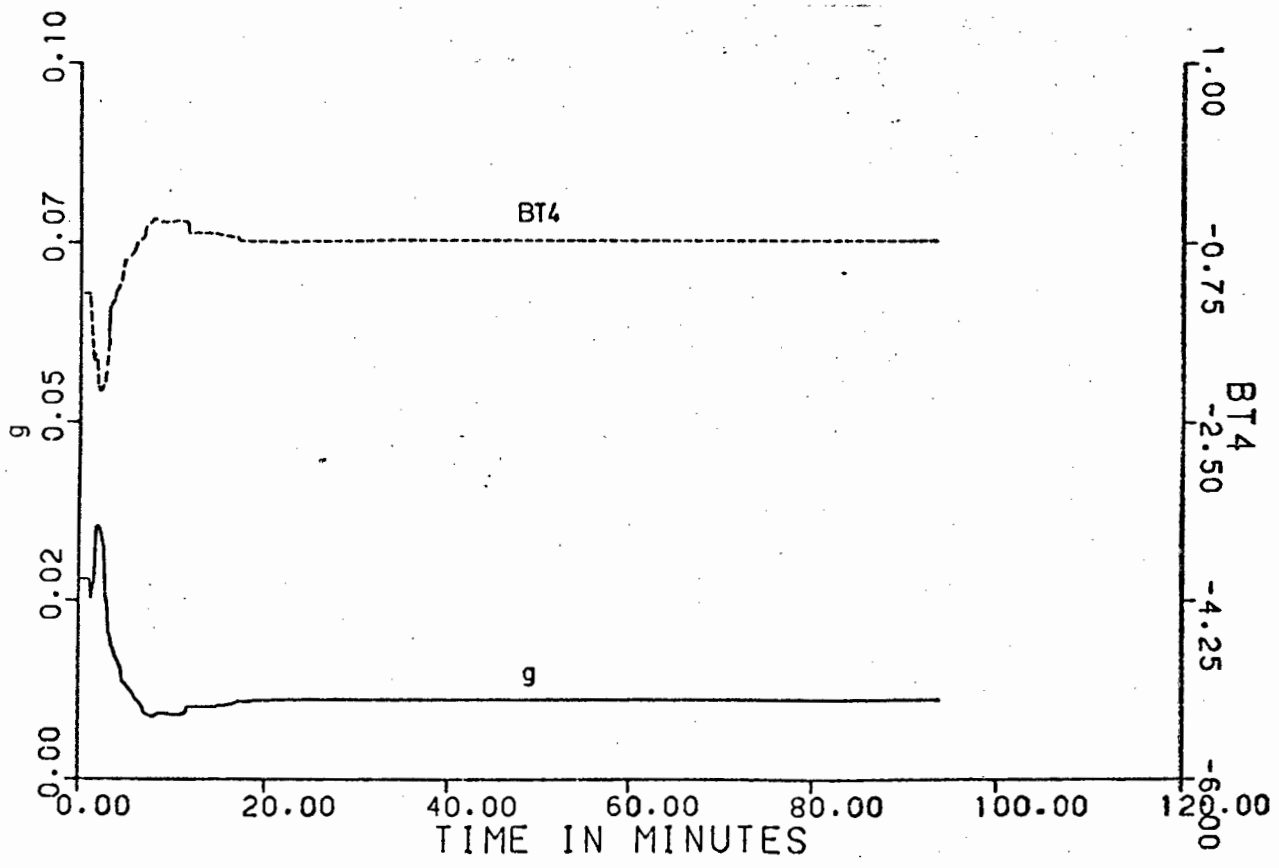

(b)

Figure 6.4 : On-line control of the Tl ratio for patient 44 who required the least amount of dTC to induce and maintain relaxation during the first hour of operation. Fig 6.4(a) - Actual muscle relaxation response measured clinically (solid line). The dashed line is the desired set-point. Fig 6.4(b) - Model parameters 'g' and BT4 as they were estimated on-line during the clinical trial.

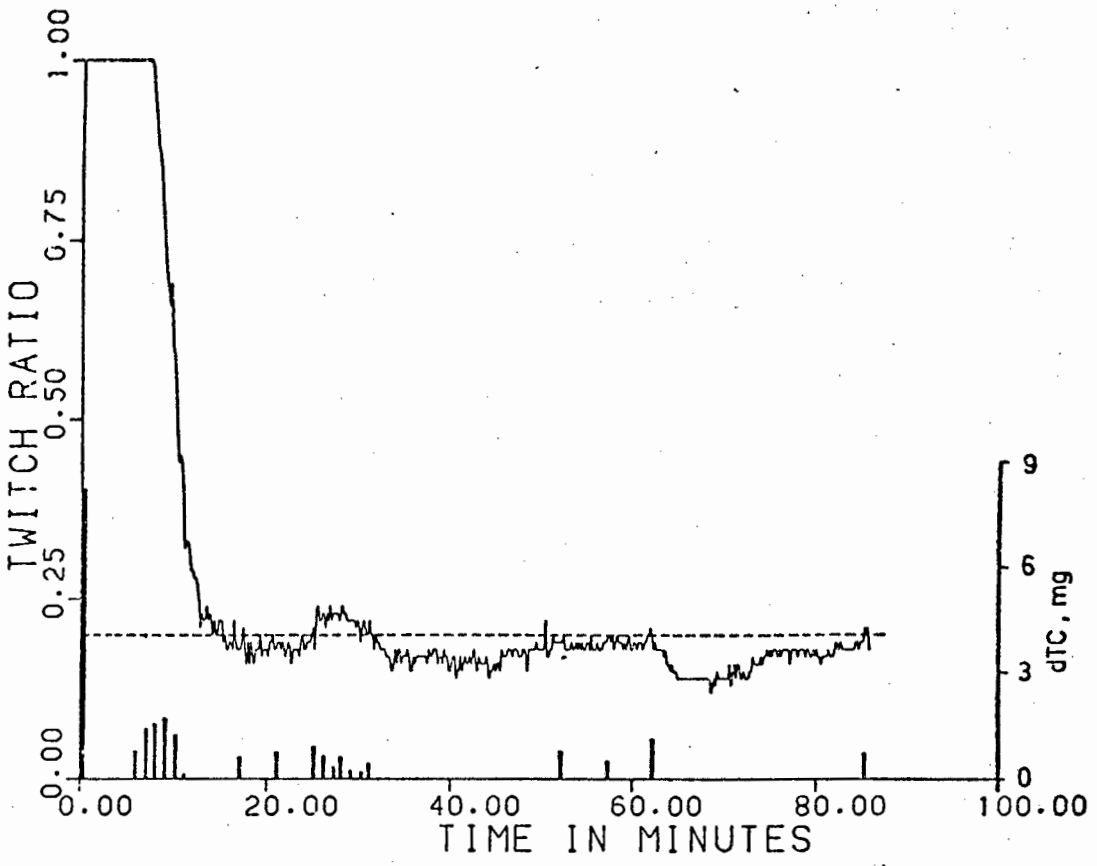

(a)

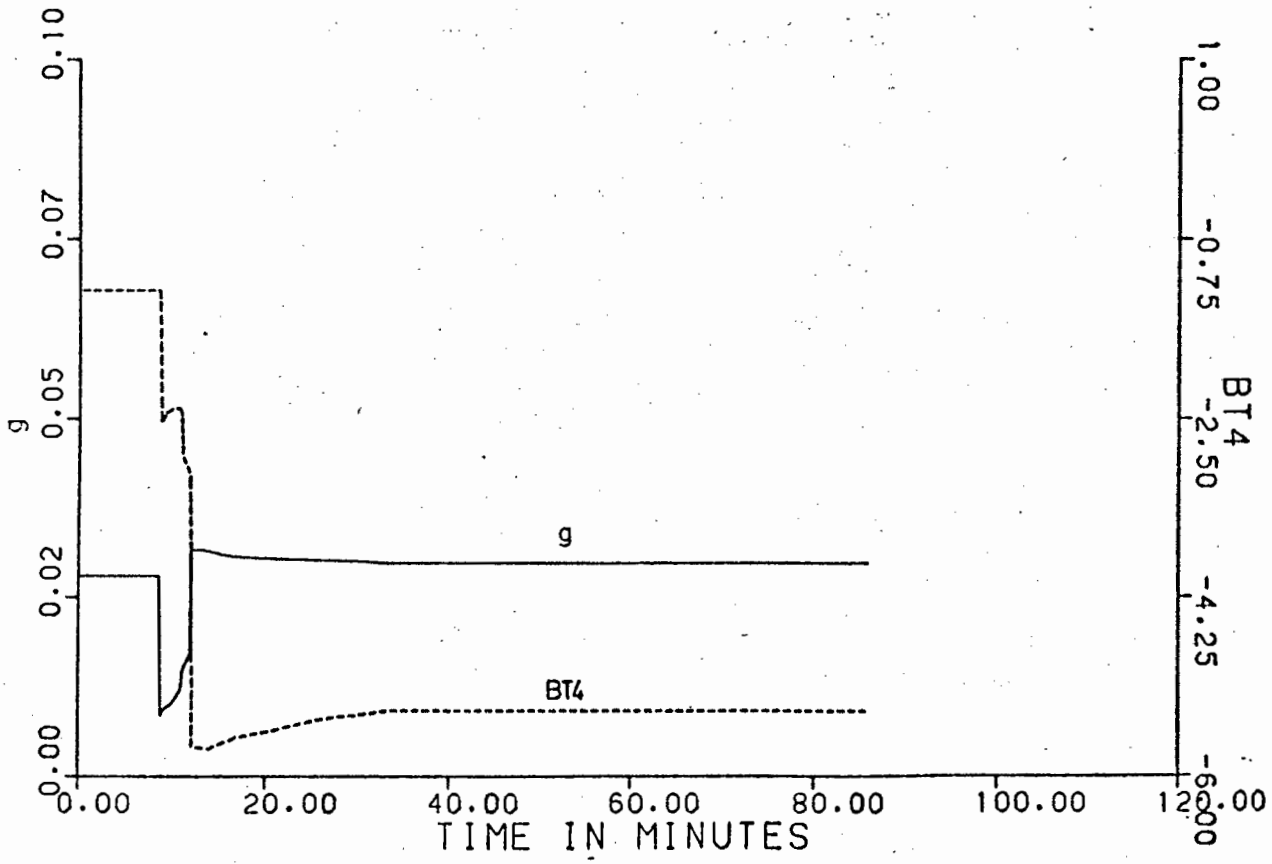

(b)

Figure 6.5 : On-line control of the T1 ratio for patient 34 who required an average amount of dTC to induce and maintain relaxation during the first hour of the operation. Fig 6.5(a) - Actual response measured clinically (solid line) and the desired set-point (dashed line). Fig 6.5(b) - Model parameters 'g' and BT4 as they were estimated on-line during the clinical trial.

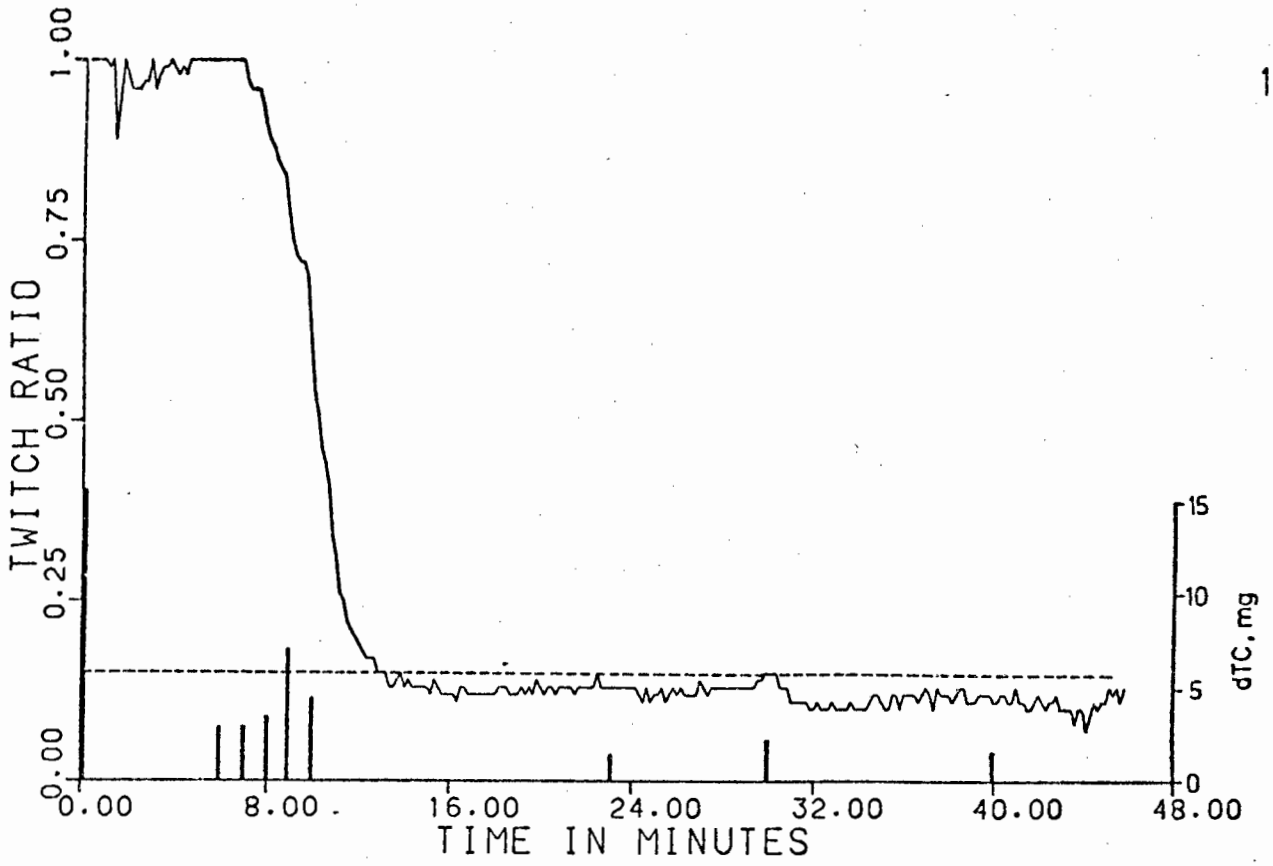

(a)

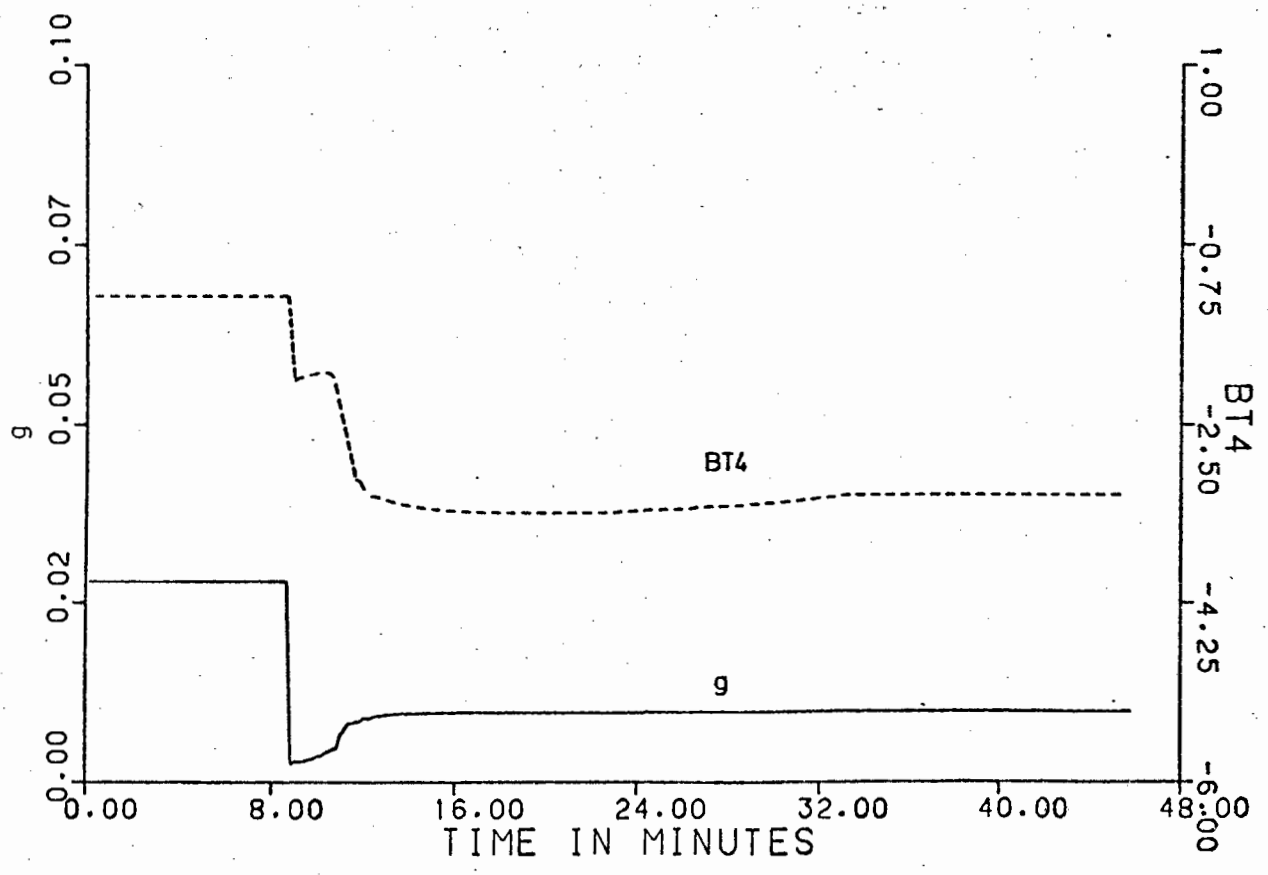

(b)

Figure 6.6 : On-line control of the T1 ratio for patient 33 who required the greatest amount of dTC to induce and maintain relaxation during the first hour of the operation. Fig 6.6(a) - Actual response measured clinically (solid line) and the desired set-point (dashed line). Fig 6.6(b) - Model parameters 'g' and BT4 as they were estimated on-line during the clinical trial.
The bandwidth of the State Estimator in the regulation phase is inadequate (Rw = Ø,ØØØ2) and doses are therefore administered while the patient is over-relaxed.

set-point. The drug commands issued by the controller have been superimposed on these three figures at the times they were administered. They are seen to be quite randomly spaced.

6.6 Analysis of the State Estimator

The settings of the State Estimator noise variance R_w given in section 6.2 were based on an investigation of controller performance on simulated patients only. In view of this further refinement is needed when the controller operates on real patients.

Accordingly various choices of R_w were made in the T1 clinical trials so that the sensitivity of the controller to this parameter could be tested. The relaxation and regulation phases were treated separately - two choices of R_w being made in the former (Ø,Ø2 and Ø,ØØ2) and three in the latter (Ø,Ø2; Ø,ØØ2 and Ø,ØØØ2). The investigation brought three points to light:

1. The choice of R_w = Ø,ØØØ2 for the regulation phase was too small, for in the single case it was chosen the controller showed a tendancy to administer drug while the patient was still over-relaxed - a sign that the estimated response was not tracking the measured data well. This was patient 33 shown in figure 6.6. This choice of R_w was thus not used again.

2. As long as R_w is not too small, the average controller performance is not particularily sensitive to the choice made (as can be seen in Table 6.7) for the Student's t-test reveals that there is no significant

difference between performance with the alternative settings for either phase of control.

3. The choice R_w = Ø,Ø2 for the regulation phase may be too large because patient 37 in that group had poor regulation. In this case one of the drug commands given was obviously too large. Analysis of this patient is continued below.

	RELAXATION PHASE		REGULATION PHASE		
R_w	0,02	0,002	0,02	0,002	0,0002
NUMBER OF CLINICAL TRIALS	11	7	6	11	1
TIME TO SET-POINT (MIN)	12,1±5,7	9,5±5,0	-	-	-
INITIAL OVERSHOOT OF SET-POINT (%)	7,2±2,9	7,7±3,2	-	-	-
OFF-SET FROM SET-POINT IN REGULATION PHASE (%)	-	-	3,5±1,4	2,8±1,3	3,0
STD. DEVN. FROM MEAN IN REGULATION PHASE (%)	-	-	4,0±1,9	3,0±1,3	1,3

Table 6.7 : The sensitivity of the controller to the Parameter Rw of the State Estimator. The figures quoted are averages taken from the 18 sucessful Tl clinical trials.

In an off-line analysis performed on completion of the clinical trials, the estimated state of relaxation was plotted for each successful clinical trial for three choices of R_w - 2,Ø; Ø,Ø2 and Ø,ØØØ2. These values were used in both the relaxation and regulation phases of control.

The graphs obtained show the trade-off between tracking ability and noise rejection as can be seen from figure 6.7. The choice of R_w = Ø,Ø2 in this figure seems reasonable except that it does not reject the false datum received at 9 minutes. The estimate is consequently disturbed for the next 1,5 minutes. In fact this choice of bandwidth for this patient proved to be slightly too large in the clinical trial. This follows from the odd behaviour exhibited at 37 and 52 minutes where the doses given were obviously too large. This behaviour was investigated more fully by examining the print-out of drug commands obtained off-line for various choices of R_w using the technique described before. The findings for this patient are listed in Table 6.8. (The entries for R_w = Ø,Ø2 are identical to those issued by the controller during the clinical trial). For this patient a choice of R_w = Ø,ØØ2 for the regulation phase would seem to have been appropriate and so taken together with results of the bandwidth investigation done on 2 groups of patients during the clinical trials, the choice of R_w = Ø,ØØ2 appears preferable for all in the regulation phase.

TIME (MINS.)	R_w			
	0,0002	0,002	0,02	0,2
37	-	1,2	2,5	4,1
38	-	-	0,6	-
51	-	-	-	1,8
52	-	2,6	4,0	4,8

Table 6.8 : Reconstructed control experiment showing the doses of dTC in milligrams which would have been given to patient 37 during a portion of the regulation phase for different choices of Rw.

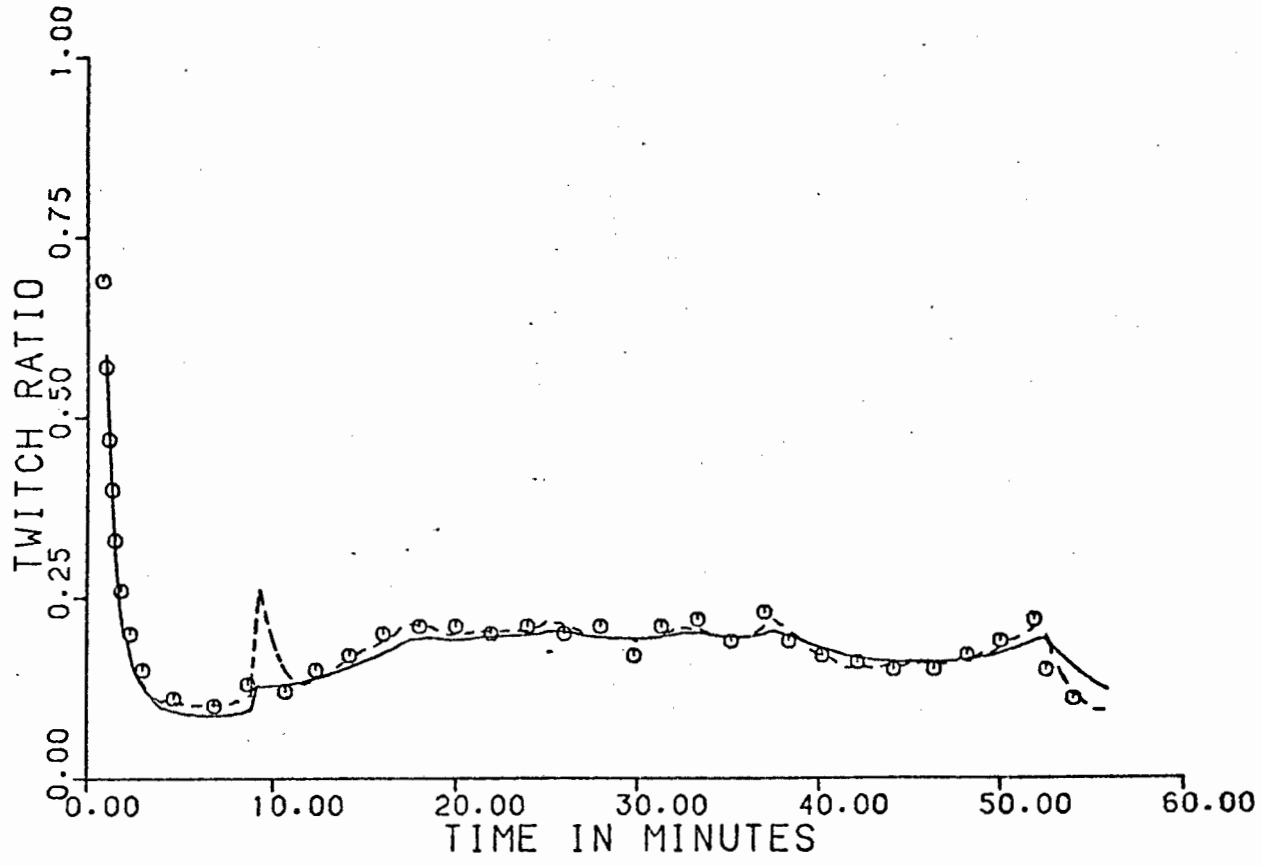

(a)

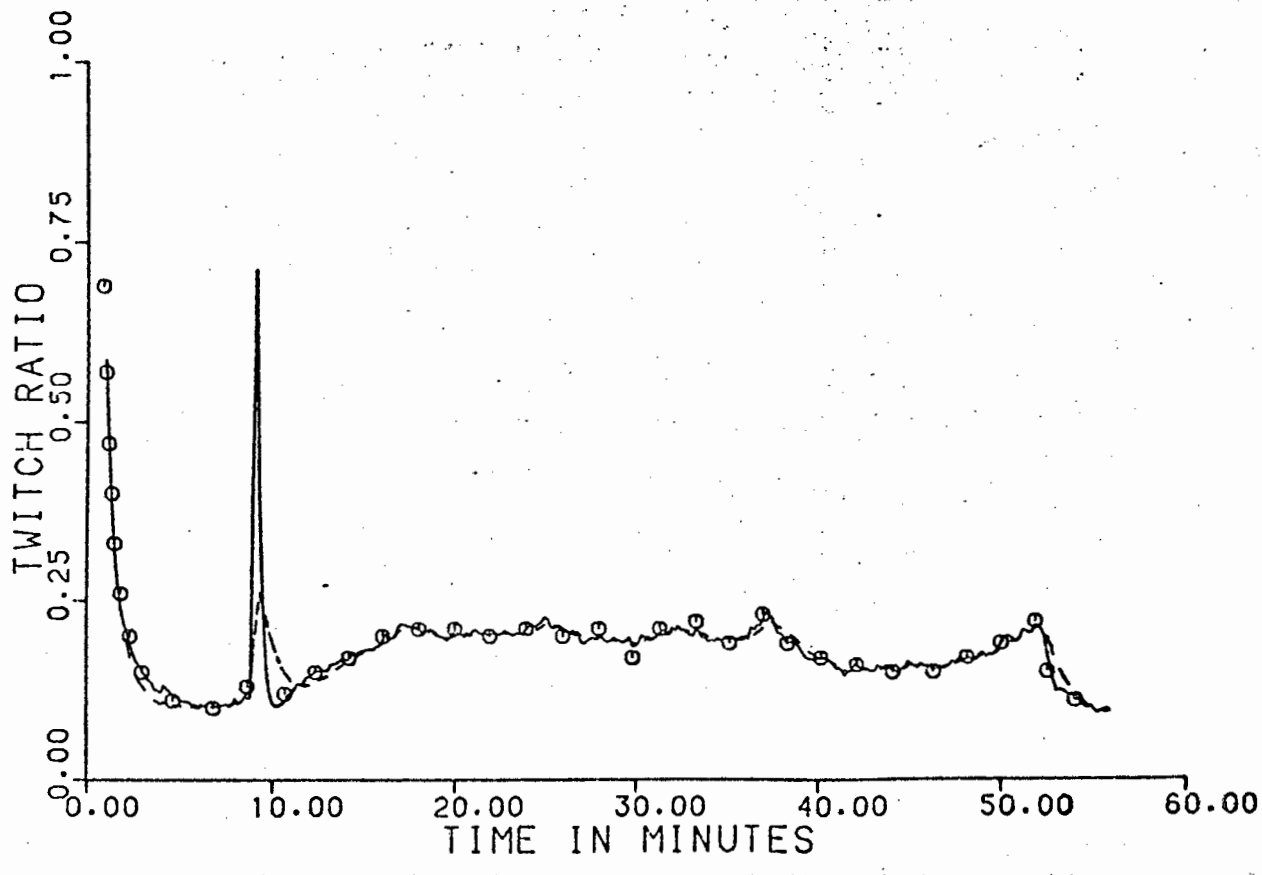

(b)

Figure 6.7 : Estimated levels of relaxation for patient 37 taken from the data of the clinical trial. In both figures (a) and (b), the dashed line represents the choice Rw = Ø,Ø2. In Fig 6.7(a) the solid line represents Rw = Ø,ØØØ2 and in Fig 6.7(b) it represents Rw = 2,Ø.

6.7 Drug Consumption

The variation in the sensitivity of the patients to dTC noted in Table 4.1 is reflected in the total mass of drug required during the operations performed on the patients. Table 6.9 shows this drug consumption for the first hour of operation and the mean infusion rate during the regulation phase for the patients of the successful clinical trials.

In all but 3 cases the controller used a smaller total mass of drug during the first hour than would have been the case had a bolus dose of Ø,5 mg/Kg been administered at the start of the operation. More important however is the fact that, as Table 6.9 shows, the total drug consumption is in all cases very close to the ideal which could have been achieved if all the patient model parameters were known a priori.

6.8 Trials which were unsuccessful for reasons other than equipment failure

Of the 42 patients entered into the clinical trials, control was unsuccessful in 6 cases due to the signal-to-noise ratio of the measurement system being inadequate. These are discussed in chapter 2.

In 1 case control had to be aborted due to an inadequacy in the software of the controller. This occurred amongst the T1 trials. In the course of controlling this patient, the T1 ratio showed a sudden and inexplicable jump which the controller interpreted as a rapid recovery and which it attempted to counter by giving an inappropriately large bolus dose of 7,9 mg. This kept the patient in an over-relaxed state until the end of the operation. The results of this case are therefore not discussed further here.

TOTAL DRUG CONSUMPTION PER HOUR OF OPERATION (mg/Kg)						DRUG INFUSION RATE IN REGULATION PHASE (µg/Kg/min)					
T1 CONTROL			T4 CONTROL			T1 CONTROL			T4 CONTROL		
PATIENT	ACTUAL	IDEAL	PATIENT	ACTUAL	IDEAL	PATIENT	ACTUAL	IDEAL	PATIENT	ACTUAL	IDEAL
27	0,26	0,28	11	0,18	0,17	27	2,1	2,4	11	0,8	1,1
28	0,27	0,24	12	0,27	0,27	28	1,9	2,5	12	1,5	1,2
29	0,35	0,33	13	0,19	0,19	29	3,3	3,0	13	1,7	1,7
30	0,47	0,46	14	0,27	0,26	30	4,0	4,4	14	1,1	0,2
31	0,34	0,35	15	0,31	0,34	31	3,1	3,4	15	3,4	3,4
32	0,35	0,36	16	0,27	0,28	32	2,9	3,1	16	1,5	1,6
33	0,71	0,70	17	0,11	0,11	33	2,1	2,1	17	0,8	0,8
34	0,37	0,37	18	0,27	0,32	34	2,5	2,1	18	1,8	2,7
35	0,25	0,24	19	0,19	0,20	35	1,8	1,7	19	2,1	2,3
36	0,29	0,27	20	0,14	0,14	36	1,3	2,2	20	1,4	1,1
37	0,42	0,34	21	0,37	0,40	37	3,7	3,3	21	4,2	4,3
38	0,25	0,28	22	0,26	0,28	38	2,0	2,9	22	3,1	3,2
39	0,25	0,24	23	0,29	0,28	39	1,7	1,6	23	1,8	1,6
40	0,47	0,49	24	0,23	0,24	40	3,1	3,5	24	2,2	2,5
41	0,57	0,56	25	0,37	0,34	41	3,7	3,6	25	2,7	2,5
42	0,30	0,28	26	0,25	0,24	42	1,6	2,7	26	1,5	2,2
43	0,61	0,67				43	5,0	7,3			
44	0,25	0,27				44	2,6	2,8			
MEAN	0,38	0,37	MEAN	0,25	0,25	MEAN	2,69	3,03	MEAN	1,98	2,03
SD. DEV	0,14	0,14	SD. DEV	0,07	0,08	SD. DEV	0,99	1,28	SD. DEV	0,95	1,07

Table 6.9 : Actual drug consumption per hour of relaxation and mean infusion rate in the regulation phase. The 'ACTUAL' performance is taken from the clinical trials; 'IDEAL' performance is that which would have been possible had the parameters of the patient model been known to the controller a priori.

As control of the T2 ratio failed in 3 of the 4 cases where it was attempted, the single successful case has not been included in the analysis of this chapter.

6.9 Discussion

In this chapter the controller has been evaluated on both simulated and real patients. The controller was tested on simulated patients in order to select settings for use in the clinical trials.

The accuracy of the parameter estimation algorithm was verified and its efficacy was investigated. It was found that there is no advantage to be gained in using a minimization technique more sophisticated than Steepest Descent in the controller. This follows because there was no significant difference in controller performance between the clinical trials and simulated trials which used identical controller settings and estimated the same parameters but which used subroutine VA05A in place of Steepest Descent.

When the controller has been correctly initialized, the interpatient variability in response to dTC is insufficient to justify on-line parameter estimation with the control algorithm used. This is bourne out by figure 6.8 which shows control of simulated patients 44 and 33 who required respectively the least and the greatest mass of relaxant for the first hour of operation during the clinical trials. The best Initial Parameter Vector, choice 4 of Table 6.2 is used. Estimation of g and BT4 is seen to result in minimal improvement over the case where no parameters are estimated.

The fact that parameter estimation does not improve controller performance significantly when correct initialization of the controller is possible can be

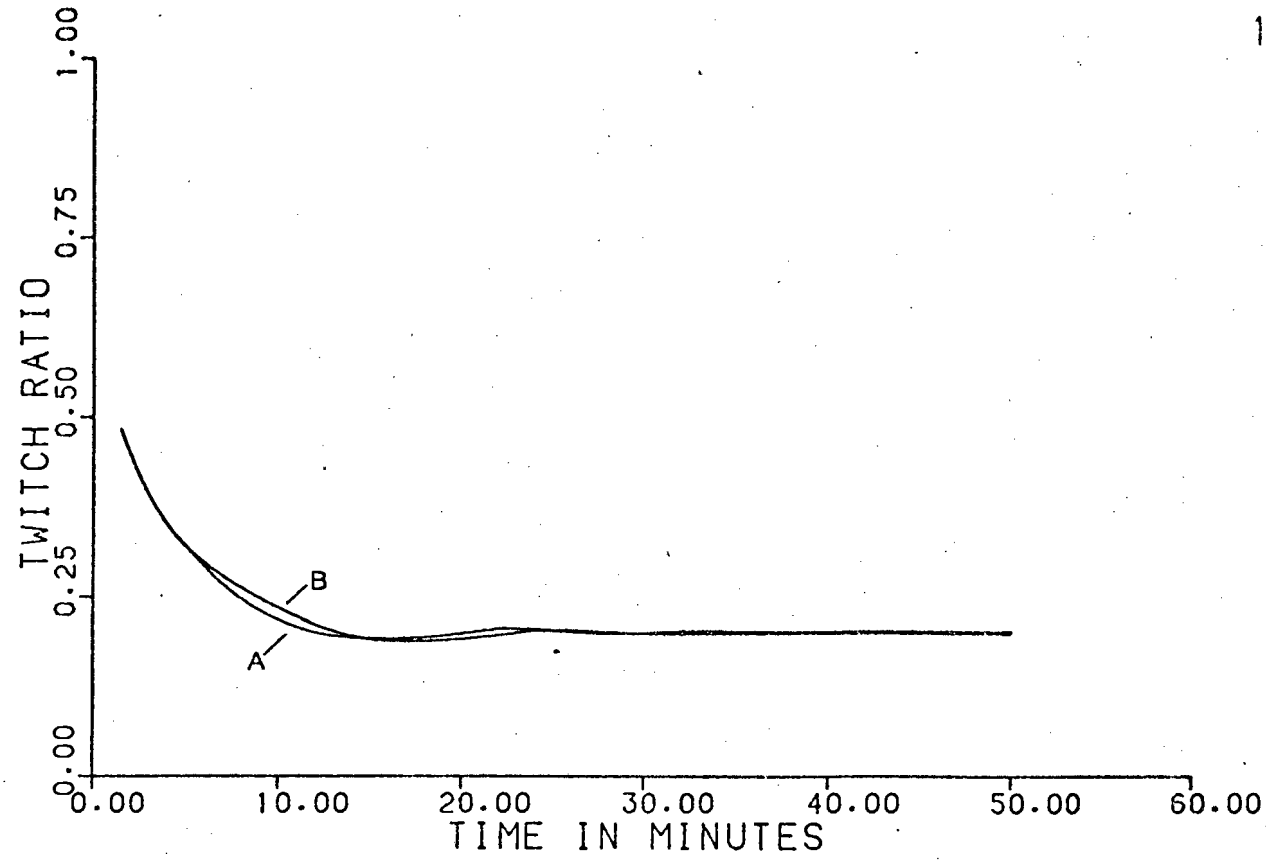

(a)

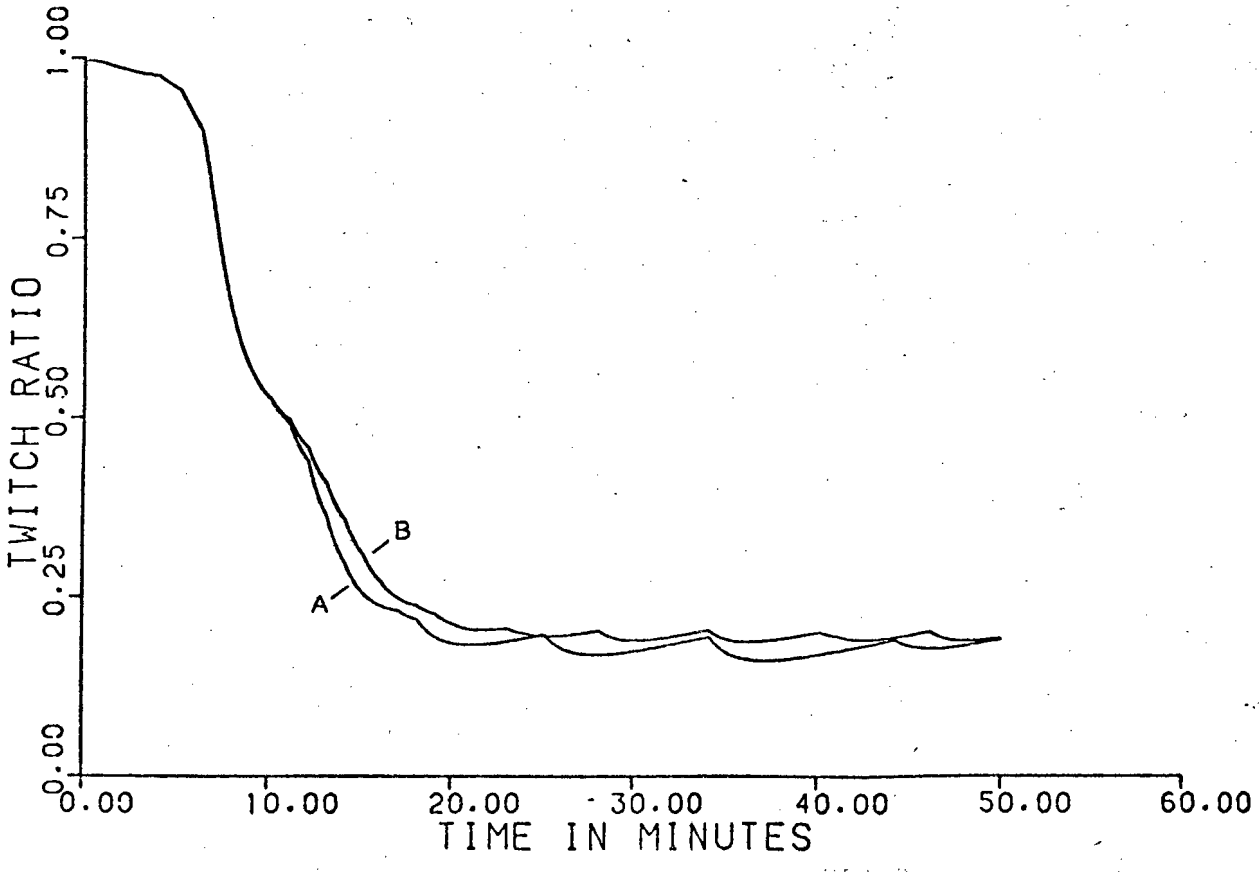

(b)

Figure 6.8 : Control of simulated patients 44 (figure (a)), and 33 (figure (b)) who required respectively the least and the greatest mass of relaxant during the first hour of operation. Curve A - 'g' and BT4 estimated, curve B - no parameters estimated. In both cases the best Extreme and Initial Parameter Vectors were used.

attributed to there being insufficient variability in the patients' sensitivity to dTC. In section 4.4 this variation was shown to be 5,2:1.

When the a priori information on which to initialize the controller is inadequate, as was the case in this research, on-line parameter estimation can benefit the control significantly. This was shown by the significant reduction in induction time gained by estimating g and BT4, also by the fact that the controllers which estimate g and BT4 from different initial conditions do not have significantly different performance. This demonstrates that sensitivity to initialization is significantly less when parameters are estimated on-line which is a major advantage.

It was noted in Table 6.5 that on-line estimation of g and BT4 shortened the induction time significantly but that this was accompanied by a significant increase in the off-set from set-point. However clinically a reduction in the induction time is of far greater value than reducing deviations from the set-point. The value of on-line parameter estimation in the clinical trials is thus verified by this shortening of the induction period.

The patient response in the regulation phase will be asymmetrical about the set-point because a drug command is calculated immediately that the response recovers to the set-point. The patient thus spends most of the time in a slightly over-relaxed state. In this sense the controller has been designed to operate with an off-set.

It is a simple matter to remove much of this off-set by writing a provision in the controller algorithm.

The selection of the Extreme Parameter Vector should be a clinical decision. The "best" choice made in this chapter gave no better average performance than that used in clinical trials. This is so because the best

choice is that which is safe for extreme patients who are in the minority. The average performance could be improved at the expense of these patients. Use of a truely extreme vector might be more realistic on a shorter acting drug where some lengthening of the induction time can be tolerated and the risk of initial overshoot correspondingly reduced.

The investigation into the State Estimator bandwidth revealed that controller performance is not very sensitive to this choice provided the bandwidth is not too small. Further, the settings chosen following the control study on the batch of simulated patients 1 - 10 proved suitable for use in the clinical trials.

Finally the clinical trials in which the T1 ratio was controlled have shown that it is necessary to pre-filter the twitch measurements before passing them to the control algorithm so as to exclude data which are obviously false.

CHAPTER 7

Conclusion

Individual chapters of this thesis were devoted to descriptions of the hardware/software of the control system, parameter estimation theory, controller design and the evaluation of the control system.

This chapter presents the conclusions which can be drawn from the results of the research and from the experience gained in undertaking it. Finally recommendations are made for future work.

7.1 Hardware/Software of the control system

All the software of the control system was written in assembly language because the departmental laboratories were unable to support a high-level language at the time that the research was undertaken. A great portion could however have been written in a high-level language and the use of assembly language confined to the few critical sections of code. The Microcomputer Development System that was used makes provision for such a mixture.

The microcomputer system could have been implemented on an IBM Personal Computer. This can accomadate the 8087 NDP. The data would be logged on Floppy Disk rather than on cassette tape. This product was however not on the market when the research was initiated. In any event the hardware interface to the Twitch Measurement System and Drug Pump would still need to be designed.

7.2 System modelling theory

The three purposes of system modelling in this research were :

1. to obtain a model structure on which to base the design of the controller
2. to obtain patient parameters with which to initialize the control algorithm
3. to enable realistic simulations of patients to be made for controller evaluation

The two stage process in which the patient deterministic and noise models were fitted independently simplified the practical aspects of parameter estimation. Thus the circulatory dead-time could be found by examining the behavior of the loss function of correlated residuals without the need of simultaneously finding the appropriate number of parameters in the noise model.

The autoregressive form of the patient noise model that was used did result in an inconvieniently large number of parameters (more than five) in 7 of the 44 patients. This was of no eventual consequence because the noise model was not used in the controller design.

7.3 Controller design

Whilst simulations have shown that a conventional PID controller may well be suitable for the regulation phase of control, it appears that the non-linearity in the patient pharmacodynamics together with the restriction that

the patient is a 'positive input system' render such a controller unsuitable for the relaxation phase. However the dual-mode controller designed for this research functioned well in all phases of control.

In the control of muscle relaxation, on-line parameter estimation is of value if :

1. there is inadequate a priori information on which to initialize the controller

2. there is a large inter-patient variability in the response to the relaxant

In the clinical trials undertaken, on-line parameter estimation was used for both of these reasons and did result in a significantly shortend relaxation phase. However it was also shown by simulation that this improvement would have been achieved without on-line parameter estimation had a better choice of the Initial Parameter Vector been made. It therefore appears that the variability in the patient response is insufficient to justify on-line parameter estimation with this type of controller.

The variability in the patient sensitivity (i.e. the single bolus dose which will depress the twitch ratio of an unrelaxed patient by 80%) to dTC encountered in the clinical trials is 5,2:1. This figure is probably conservative as the patients were all adults and of grades ASA I or II. Further there were only three sets of clinical conditions used in the experiments. It was shown in section 4.5.4 that the average patient sensitivity was significantly different in two of these groups. One might therefore expect the value of on-line parameter estimation to be greater in general than that found

in the clinical trials.

The controller can be used with other relaxants which have a similar pharmacokinetic model structure. The variation in patient sensitivity to these relaxants is unknown at present and needs to be investigated before the value of on-line parameter estimation can be established with greater certainty.

The on-line parameter estimation algorithm uses Steepest Descent to minimize the loss function. The technique is simple and although somewhat inelegant (it is a batch processing method which is repeated at each sampling period), it is adequately suitable for the application. The accuracy is close to that of the sophisticated optimization package used to estimate the parameters off-line for comparison purposes. Simulated control trials have shown that use of this package in the controller does not result in significantly improved performance over the Steepest Descent method when estimating the same parameters. Further, had the poles of the patient model been estimated on-line in addition to parameters 'g' and 'BT4', the efficacy of the controller would only have improved in the regulation phase where the off-set would have been smaller. However this gain is of no clinical significance.

The fact that the control algorithm needs an expanding memory (section 4.8) does not detract from its performance. It is a simple matter, and inexpensive, to include as much memory as might be required for an operation of any practical duration.

7.4 Clinical considerations

The research has shown closed loop control of all phases of dTC induced muscle relaxation to be possible on humans within the safety bounds derived from Ham et al.,(1979). The controller is able to attain a degree of paralysis suitable for both lower and upper abdominal surgery. See Uys et al.,(1984).

Whilst control of both the T4 and T1 ratios was successful, the depth of relaxation in the former case was not really adequate - even for lower abdominal surgery - and it is concluded that use of the T4 ratio should be confined to the recovery phase of the operation to assess when the patient has recovered from the effects of the relaxant.

In clinical terms, the performance of the controller is close to the optimum which could be achieved with curare.

The important advantage of on-line control is that the patient is always maintained at a state of relaxation at which there is a measurable response - this applies even to patients who showed an initial overshoot of the set-point - so that the danger of re-curarization (Viby-Mogensen et al.,1979) is minimized at all times.

Relaxation was controlled at a peripheral muscle and not at the site of interest to the surgeon. However the degree of paralysis was always found to be adequate for his purposes in the clinical trials where the T1 ratio was controlled. One may conclude that the present method of stimulus and effect measurement is acceptable for the types of surgical proceedure stated.

7.5 Recommendations for future work

The Nerve Stimulator used in this research was of the constant voltage type. It performed satisfactorily and it was possible to obtain a supramaximal response from every patient. It is however the current density at the electrodes which determines the magnitude of the stimulus and in the case of the stimulator actually used this is dependant on the impedance between the electrodes, which could conceivably change in the course of the operation. Constant current stimulus is therefore preferable because it is insensitive to this impedance.

It is useful to have a visual indication of the twitch response for this assists in setting the supramaximal level and assessing the degree of background noise. A chart recorder was used for this purpose during the trials but a CRT based unit of the type used to display the ECG response in the theatre is more convienient.

When the effect measurement is chosen to be the T2 or T4 ratio, there are two checks which are made on the validity of the measured twitch response namely that the time spacing between the twitches in the group is that which is expected and that the received pulse train decreases monotonically. A visual examination of the measured patient response has shown that these measures are effective in rejecting false data. However neither of them can be used when controlling the T1 ratio. There is consequently no guard to prevent data which are obviously false from being passed to the control algorithm. This resulted in instability in the case of the patient discussed in section 6.8. The use of a data pre-filtering algorithm is therefore suggested.

LIST OF REFERENCES

- Adby, P.R. and Dempster, M.A.H. (1974) Introduction to optimisation methods. Chapman and Hall, London.

- Albisser, A.M. (1979) Devices for the control of diabetes mellitus. Proc. IEEE, 67, 1308 - 1320.

- Ali, H.H. and Saverese, J.J. (1976) Monitoring of Neuromuscular Function. Anesthesiology, 45, 216 - 249.

- Ali, H.H., Utting, J.E. and Gray, T.C. (1971) Quantitative assessment of Residual antidepolarizing Block. Brit. J. Anaesth., 43, 473 - 483.

- Arnsparger, J.M., McInnis, B.C., Glover, J.R. et al. (1983) Adaptive Control of Blood Pressure, IEEE Trans. on Biomedical Engineering, BME-30, 168 - 176.

- Astrom, K.J. (1983) Theory and Applications of Adaptive Control - A Survey, Automatica, 19, 471 - 486.

- Astrom, K.J. and Eykhoff, P. (1971) System Identification - Survey, Automatica, 7, 123 - 162.

- Astrom, K.J. and Soderstrom, T. (1974) Uniqueness of the Maximum Likelihood Estimates of the Parameters of an ARMA Model. IEEE Trans. on Automatic Control, AC19, 769 - 773.

- Astrom, K.J. and Wittenmark, B. (1973) On Self Tuning Regulators. Automatica, 9, 185 - 199.

- Astrom, K.J. and Wittenmark, B. (1984) Computer controlled systems; theory and design. Prentice-Hall, Englewood Cliffs, N.J. Ø7632.

- Auer, L.M. and Rodler, H. (1981) Microprocessor-control of drug infusion for automatic blood-pressure control, Med. & Biol. Eng. & Comput., 19, 171 - 174.

- Balakrishnan, A.V. and Peterka, V. (1969) Identification in Automatic Control Systems. Automatica, 5, 817 - 829.

- Berger, C.S. and Brown, W.A. (1984) Adaptive control of muscle relaxation. Proc. IFAC/IFIP Symposium on Real Time Digital Control Applications, Guadalajara, Mexico, 1Ø9 - 119.

- Bevan, D.R. (1983) Neuromuscular Blocking Drugs. Can. Anaesth. Soc. J., 3Ø, 556 - 561.

- Bradlow, H.S. and Rametti, L.B. (1982) A model for use in automated control of muscle relaxation during surgical proceedures. Sixth IFAC Symposium on Identification and System Parameter Estimation, Washington, 1Ø98 - 11Ø3.

- Bradlow, H.S., Rametti, L.B., Uys, P.C. and Coetzee, W.P. (1984) A micro-computer based muscle relaxation monitor and controller, for clinical use. To be published in Med. & Biol. Eng. & Comput.

- Bradlow, H.S., Sherlock, B.G, Thornington, R.E. and Solomon, E.S. (1983) Model determination for use in automated control of drug dosages : application to d-tubocurarine. Med. & Biol. Eng. & Comput., 21, 119 - 127.

- Brammer, R.F. (1972) Controllability in linear autonomous systems with positive controllers, SIAM J. Control and Optimization, 10, 339 - 353.

- Brown, B.H., Asbury, J., Linkens, D.A., Perks, R. and Anthony, M. (1980) Closed-loop control of muscle relaxation during surgery. Clin. Phys. Physiol. Meas., 1, 203 - 210.

- Brown, W.A., Lampard, D.G., Ng, K.C. et al. (1981) Microprocessor control of muscle relaxation, Proc. IFAC 8th Triennial World Congress on Control Science and Technology for the progress of Society, Kyoto, Japan.

- Cass, N.M., Lampard, D.G., Brown, W.A. and Coles, J.R. (1976) Computer controlled muscle relaxation : a comparison of four muscle relaxants in the sheep, Anaesthesia and Intensive Care, IV (#1), 16 - 22.

- Clarke, D.W. and Gawthrop, P.J. (1975) Self-tuning controller, Proc. IEE, 122, 929 - 934.

- Collins, S.M. and Arzbaecher, R.C. (1979) Computer control of cardiac arrhythmia, Frontiers of Engineering in Health Care : Proc. 1st Annual IEEE Conference on Engineering in Medicine and Biology, Denver, 276 - 279.

- Day, A. (1981) SABUS - Triple RS232 Card, Internal Report, Department of Electrical and Electronic Engineering, University of Cape Town.

- Dugard, L. and Landau, I.D. (1980) Recursive Output Error Identification Algorithms Theory and Evaluation, Automatica, 16, 443 - 462.

- Eva, J.A. (1980) A Digital Cassette Tape Recorder Controller. B.Sc (Eng.) Thesis 539, Department of Electrical and Electronic Engineering, University of Cape Town.

- Eykhoff, P. (1982) On the coherence among the multitude of system identification methods. Sixth IFAC Symposium on Identification and System Parameter Estimation, Washington, 737 - 748.

- Franklin, G.F. and Powell, J.D. (1980) Digital control of dynamic systems. Addison-Wesley, World Student Series.

- Gibaldi, M., Levy, G. and Hayton, W. (1972) Kinetics of the elimination and neuromuscular blocking effect of d-tubocurarine in man. Anesthesiology, 36, 213 - 218.

- Ham, J., Miller, R.D., Sheiner, L.B. and Matteo, R.S. (1979) Dosage-schedule independence of d-tubocurarine pharmacokinetics and pharmacodynamics, and recovery of neuromuscular function. Anesthesiology, 50, 528 - 533.

- Hopper, M.J. (Editor) (1974) Subroutine VA05A of the Harwell subroutine library, Harwell. Atomic Energy Authority, United Kingdom.

- Hull, C.J., English, M.J.M. and Sibbald, A. (1980) Fazadinium and Pancuronium : a pharmacodynamic study. Br. J. Anaesth., 52, 1209 - 1220.

- International Electrical Commission (IEC). (1977) Safety of medical electrical equipment, Part 1 : General Requirements. IEC Standard Publication 601 - 1, Geneva.

- ISCOR. (1981) SABUS 32k Static Memory Module - 140/3. Pulse (Nov. 81), 61 - 63.

- Isermann, R. (1982) Parameter Adaptive Control Algorithms - A Tutorial. Automatica, 18, 513 - 528.

- Jacobs, O.L.R. (1982) Case-studies in direct digital control of stochastic systems. O.U.E.L. Report no. 1430/82, Department of Engineering Science, University of Oxford.

- Jacobs, O.L.R. and Reasbeck, M.P. (1982) On-line estimation in the control of post-operative pain. 6th IFAC Symposium on Identification and System Parameter Estimation, Washington, 228 - 233.

- Jelliffe, R.W. (1983) Computer-Controlled Administration of Cardiovascular Drugs. Progress in Cardiovascular Diseases, 26, 1 - 14.

- Kaufman, H., Roy, R. and Xu, X. (1982) Model reference adaptive control of drug infusion rate. 6th IFAC Symposium on Identification and System Parameter Estimation, Washington, 216 - 221.

- Koivo, A.J. (1980) Automatic Continuous-Time Blood Pressure Control in Dogs by means of Hypotensive Drug Injection. IEEE Trans. on Biomedical Engineering, BME-27, 574 - 581.

- Koivo, A.J. (1981) Microprocessor-Based Controller for Pharmacodynamic Applications. IEEE Trans. on Automatic Control, AC-26, 1208 - 1212.

- Linkens, D.A., Asbury, A.J., Rimmer, S.J. and Menad, M. (1982) Self-tuning control of muscle relaxation during anaesthesia, IEEE Conference on applications of adaptive and multivariable control. Hull, 96 - 102.

- Luenberger, D.G. (1978) Introduction to Dynamic Systems - Theory, Models and Applications, Wiley.

- Mandel, J. (1964) The statistical analysis of experimental data., Interscience, New York.

- Meijer, D.K.F., Weitering, J.G., Vermeer, G.A. et al. (1979) Comparative pharmacokinetics of d-Tubocurarine and metocurine in man. Anesthesiology, 51, 402 - 407.

- Moore, J.B. (1982) Global Convergence of Output Error Recursions in Coloured Noise. IEEE Trans. on Auto. Control, AC-27, 1189 - 1199.

- Morris, P., Tatnall, M.L. and Montgomery, F.J. (1983) Controlled anaesthesia : A clinical evaluation of an approach using patient characteristics identified during uptake. Br. J. Anaesth., 55, 1065 - 1075.

- Olson, W.H. (1978) Electrical Safety. In Webster, J.G. (Ed.), Medical Instrumentation : Application and design. Houghton Mifflin Co., Boston. 667 - 707.

- Page, B.S. (1980) Real-time Microprocomputer System for Medical Data. B.Sc (Eng.) Thesis 554, Department of Electrical and Electronic Engineering, University of Cape Town.

- Putman, A.E. (1980) SABUS - EPROM, RAM and Keyboard Display Controller. Internal Report, Department of Electrical and Electronic Engineering, University of Cape Town.

- Rametti, L.B. and Bradlow, H.S. (1983) On-line control of d-tubocurarine induced muscle relaxation : a simulation study. Med. & Biol. Eng. & Comput., 21, 710 - 717.

- Rametti, L.B., Bradlow, H.S. and Uys, P.C. (1984) Self-tuning feedback control of d-Tubocurarine induced muscle relaxation. To be published in Med. & Biol. Eng. & Comput.

- SABUS Committee. (1981) SABUS - The South African standard bus system for microprocessors. Pulse (Nov., 1981), 49 - 50.

- Sheiner, L.B., Stanski, D.R., Vozeh, S. et al. (1979) Simultaneous modelling of pharmacokinetics and pharmacodynamics : application to d-tubocrarine. Clin. Pharmacol. Ther., 25, 358 - 371.

- Sheppard, L.C., Jannett, T., Walker, F. et al. (1982) Computer controlled muscle relaxation system for Succinylcholine administration. Private communication, University of Alabama, Birmingham.

- Sherlock, D.A. (1981). S64 Bus 8086/7/8/9 Multiprocessing CPU Module 013. Internal Report, Department of Electrical and Electronic Engineering, University of Cape Town.

- Slate, J.B. and Sheppard, L.C. (1982) Automatic control of blood pressure by drug infusion. IEE Proc., 129, Pt. A, 639 - 645.

- Smith, J.A. (1979). A High Resolution Digitizer. B.Sc (Eng.) Thesis 485, Department of Electrical and Electronic Engineering, University of Cape Town.

- Soderstrom, T., Ljung, L. and Gustavsson, I. (1978) A Theoretical Analysis of Recursive Identification Methods. Automatica, 14, 231 - 244.

- Spiegel, M.R. (1961) Theory and Problems of Statistics. Schaum Publishing Co., New York.

- Stanski, D.R., Ham, J., Miller, R.D. et al. (1979) Pharmacokinetics and pharmacodynamics of d-tubocurarine during nitrous oxide - narcotic and halothane anesthesia in man. Anesthesiology, 5, 235 - 241.

- Stanski, D.R. and Watkins, W.D. (1982) Drug disposition in anesthesia. Grune and Stratton, New York.

- Strumpher, J. (1982). A real-time microcomputer system for medical control. B.Sc (Eng.) Thesis 663, Department of Electrical and Electronic Engineering, University of Cape Town.

- Talmon, J.L. and van den Boom, A.J.W. (1973) On the estimation of the transfer function parameters of process - and noise dynamics using a single-stage estimator. Proc. of the 3rd IFAC Symposium on Identification and System Parameter Estimation, The Hague/Delft, 711 - 720.

- Uys, P.C., Bradlow, H.S. and Rametti, L.B. (1984) Controlled muscle relaxation : An evaluation of a computer controlled system to achieve and maintain a predetermined degree of muscle relaxation during anaesthesia. To be submitted for publication.

- Viby-Mogensen, J. (1982). Clinical assessment of Neuromuscular Transmission. Br. J. Anaesth., 54, 209 - 233.

- Viby-Mogensen, J., Jogensen, B.C. and Ording, H. (1979) Residual curarization in the recovery room. Anesthesiology, 50, 539 - 541.

- Walker, B.K., Chia, T., Stern, K.S. et al. (1982) Parameter identification and adaptive control for blood pressure. 6th IFAC Symposium on Identification and System Parameter Estimation, Washington, 1255 - 1260.

- Waud, R.E. (1975) Serum d-tubocurarine concentration and twitch height. Anesthesiology, 43, 381 - 382.

- Wellstead, P.E., Prager, D. and Zanker, P. (1979) Pole assignment self-tuning regulator. Proc. IEE, 126, 781 - 787.

- Wingard, L.B. and Cook, D.R. (1976) Pharmacodynamics of tubocurarine in humans. Br. J. Anaesth., 48, 839 - 844.

- Young, P.C. (1970) An instrumental variable method for real-time identification of a noisy process. Automatica, 6, 271 - 287.