Annals of Economic and Social Measurement, 1/4, 1972

THE DISCRETE TIME LINEAR-QUADRATIC-GAUSSIAN STOCHASTIC CONTROL PROBLEM*

BY MICHAEL ATHANS†

The purpose of this paper is to review in a tutorial fashion the role of the linear-quadratic stochastic control problem in discrete time system design. The design approach is motivated by considering the control of a nonlinear uncertain plant about a desired input–output response sequence. It is demonstrated how a design philosophy based on (a) deterministic optimal control, (b) deterministic perturbation control, (c) stochastic state estimation, and (d) linearized stochastic control, leads to an overall closed loop control system. The emphasis of the paper is on the philosophy of the design process, the modelling issue, and the formulation of the problem; the results are given for the sake of completeness, but no proofs are included. The systematic, off-line, nature of the design process is stressed throughout.

1. INTRODUCTION

This paper was motivated by the fact that most stochastic optimization problems in economics are most naturally described by difference equations. For this reason, it appeared appropriate that a summary paper describing a unified design philosophy based on advances in modern control theory, would contribute to the interchange of ideas between economists, management scientists, and control theorists. This paper then is a discrete-time version of the continuous time results (see reference 1) presented at the Princeton workshop. This paper focuses on the non-engineering aspects and interpretations of the theory.

It should be stressed that trends in stochastic control research by engineers has been greatly influenced by two factors

(*a*) a need to minimize on-line computations, and

(*b*) the requirements in many aerospace applications that the control system be realized by analog hardware.

In economic applications these requirements are not present, since the time period between decisions does allow for extensive digital computer calculations. Thus, one does have the luxury of examining more sophisticated decision and control algorithms, which however have increased computational requirements. Nonetheless, it is important to know what are the "bread-and-butter" tools in control practice, which requires a very modest amount of on-line digital computation. It is the purpose of this paper to state, in a summary form, this simple approach to the control of nonlinear stochastic systems.

The basic problem in engineering control system design almost invariably involves the *on-line* (i.e., real-time) feedback control of an uncertain, usually nonlinear, physical process. The engineer, usually likes to work with, and benefits from, a systematic approach to the design problem; such systematic approaches are often the outcome of past design experience.

* This research was carried out at the Decision and Control Sciences Group of the M.I.T. Electronic Systems Laboratory with support extended by NSF under grant GK-25781.

** The author wishes to express his sincere appreciation to Professor Gregory Chow and Mr. Kenneth Garbade of Princeton University, and to Professor S. Phillip Cooper and Dov Pekelman of the University of Chicago Business School whose discussions and comments greatly contributed to the motivation for and preparation of this paper.

Clearly, a "universal" system design approach must take into account

- the desired specifications
- actuator and sensor constraints
- measurement errors
- actuator errors
- design sensitivity due to plant parameter variations
- effects of unpredictable disturbances
- on-line vs. off-line computational requirements
- design simplicity

The purpose of this paper is to indicate how the available theory of optimal control and estimation for the so-called Linear-Quadratic-Gaussian problem provides such a unified design procedure. In particular, we wish to stress the advantages of this design process from the viewpoint of ease of computation since the theory provides us with equations that can be readily solved by modern digital computers. Thus, the success of the design process hinges on the capability of the designer to understand the physics of the problem and his ability to translate physical requirements and constraints into mathematical language. Once this crucial "modelling" has been done, the digital computer algorithms will readily generate the quantitative details of the design.

Towards this goal, this paper is structured in the following manner. In Section 2 we discuss the problem under consideration in the most general terms and we outline the design philosophy that we shall adopt. In Section 3 we discuss the deterministic aspects of the design problem, introduce the notions of the ideal input-state nominal time functions, dynamic linearization, provide the motivation for using quadratic criteria, and state the solution of the deterministic linear-quadratic problem. In Section 4, we analyze the deterministic design from the viewpoint of uncertainty and sensor constraints. This leads us to the problem of estimating the state variables of the uncertain physical process, on the basis of past measurements via the Kalman–Bucy theory. In Section 5 we "hook-up" together the stochastic estimator of Section 4 with the deterministic controller of Section 3 to obtain the desired compensator that translates actual sensor measurements to commanded control inputs. Section 6 contains a discussion of the results. Section 7 presents a brief overview of the adaptive control problem.

Of course, all the results outlined in this paper are available in one form or the other in the control literature. Hence, the hoped for contribution of this paper is that of unification, so that one can see how seemingly diverse topics in control theory can be combined to yield a systematic computer-aided design tool. Thus, the emphasis will be on philosophy, interpretation, and critical discussion of the interplay between physical processes and the mathematical models required to apply the powerful tools of modern control theory. We also hope to convey the fact that this approach to design involves both "art and science", so that creativity and know-how are (as always) the key ingredients of success.

2. The Physical Problem and the Design Philosophy

We commence our specific discussion with a brief description of the problem of controlling a physical dynamic process and a definition of the control problem.

2.1. *Physical Plant, Actuators, Sensors*

We shall deal with the interconnected entity of a *physical plant* driven by *actuators*; measurements can be made by *sensors*.

2.1.1. *Actuators*: The actuators are actual "devices" that translate *commanded inputs*, time sequences that can be specified by the designer, (e.g., commanded government expenditures into *actual plant inputs* (e.g., actual government expenditures). This translation is not exact; this is modelled by the *actuator uncertainties*. It is assumed that the actual plant inputs cannot be measured, at least at the times that the decisions have to be made.

2.1.2. *Plant*: The plant is a physical device that translates the *actual plant inputs* as well as other *plant disturbances* (e.g., probabilistic exogeneous variables) into a set of time sequences which we shall call the *physical state variables* of the plant (e.g., consumption, unemployment rate, interest rates, etc.). For our purposes, the plant state variables are the key physical variables that govern and specify completely the current behavior of the system. In economic applications they are the current and lagged values of the appropriate endogeneous variables. In the language of urban dynamics the state variables correspond to the level variables.

2.1.3. *Sensors*: We assume that it is either impossible or not desirable for physical or economic reasons, to measure all the plant state variables. The physical sensors are devices that indicate which physical variables (state variables and/or combinations thereof) can be indeed measured. However, the actual *sensor measurement signals* are different than the signals that are sensed; these errors are modelled by the inclusion of the *sensor error and uncertainty* signals, which take into account the measurement accuracy of any given sensor.

2.2. *Control System Objectives*

The desired behavior of the physical process as a whole is often judged on the basis of the actual time-evolution of all or some of the physical plant state variables. Of course, due to the sensor constraints one may not be able to deduce exactly what the plant is doing at each and every instant of time. Nonetheless, the plant state variables rather than the measurements are the key quantities that enter in the control problem formulation.

In many cases, the time evolution of the plant state variables may possess certain undesirable characteristics. These may be due to the effect of the plant disturbances and/or actuator errors, due to inherent plant instability or sluggish response. In such cases, one must be able to control the time evolution of the plant state variables by the adjustment of the time evolution of the commanded inputs, (which are the only variables that can be externally adjusted).

It is appropriate to remark at this point that the issue of unreliable measurements has been somewhat ignored in the mathematical economics literature. There is evidence to suggest that the measurements can be unreliable (see for example references [2] and [3]). On the engineering side, during the past decade, noisy measurements have been treated as the rule rather than the exception; this has contributed to certain distinct evolutionary developments in the research contributions of these two areas. Hence, the control system objective can be loosely stated as follows:

Find the time-evolution of the commanded inputs (decision variables, policy variables) such that the time evolution of the physical plant state variables is satisfactory for the task at hand.

2.3. *Control System Structure*

Since the control objective hinges on the time-evolution of the physical state variables, and since we cannot sense them directly in view of the sensor constraints, it is intuitively obvious that the actual values of the commanded inputs at the present time must somehow (at least, partially) depend upon the current, and perhaps past, values of the sensor measurements.

Thus, we are admitting right at the start that some *feedback* is necessary. This leads us to visualizing that we must construct a physical device which we shall call the *compensator*, whose task will be to translate the actual sensor measurement signals into the actual commanded inputs to the physical process.

We can now reformulate the control objective of Section 2.2, as follows:

Find the compensator, driven by the sensor measurement signals and generating the commanded inputs to the physical process, such that the time evolution of the physical plant state variables is satisfactory for the task at hand.

2.4. *The Design Philosophy*

It should be clear that the design of the compensator must hinge on

- Natural dynamics of the physical process both in the absence of uncertainty (deterministic) and in the presence of uncertainty (stochastic)
- The level of the uncertainty in the physical process (How big are the probable actuator errors? How large are the plant disturbances? How accurate are the sensors?)
- The precise notion of what characterizes, for any given application, a satisfactory time-evolution of the physical plant state variables.

In point of fact, our ability to construct such a compensator, must depend upon our ability to predict (exactly or approximately) what the physical plant state variables will be doing for any given

- commanded input time-sequences
- actuator errors, viewed as time-sequences
- plant disturbances viewed as time-sequences

Clearly the design issue is clouded because it involves the interplay between the natural dynamics of the physical process, the stochastic nature of the uncertainties, and the effects of the deterministic commanded inputs. Nonetheless, one can adopt a design philosophy that involves the following three basic steps:

Step 1. Deterministic Ideal Response Analysis and Design
Step 2. Stochastic Estimation Analysis and Design
Step 3. Stochastic Feedback Control System Design

In the rest of the paper, we shall elaborate on precisely what we mean by this three-part approach.

2.4.1. *Deterministic Ideal Response Analysis and Design (Step 1).* In this step, we *pretend* that there is absolutely *no uncertainty*. That is, we suppose that

- actuator errors *do not* exist
- there are *no* plant disturbances
- we can measure *exactly all* the physical plant state variables and output variables
- the actuator and plant dynamics are known *exactly*
- all parameter values are known *exactly*

Under these assumptions, we can predict *exactly* what the plant state and output variables will do for any given commanded inputs. If this is the case, then somehow (and this will be treated in detail in Section 3) we should be able to determine:

> *the ideal commanded inputs viewed as time-sequences which will give rise to an ideal set of plant state variable time-sequences for the application at hand.*

In short, the basic end product of this first step of the design process, is the specification of an *ideal deterministic commanded-input state-variable pair that incorporates the specifications of the application and the natural constraints and dynamics of the physical process.*

2.4.2. *Stochastic Estimation Analysis and Design (Step 2).* In this part of the design process, we reintroduce the uncertainty into our problem. In particular, we take into account that we cannot measure all of the plant state variables and that any measurement is subject to sensor errors.

The basic question that we answer at this step of the design process is the following:

> *Construct a device (state estimator, filter) that generates on the basis of the past sensor measurements a set of time-sequences which are as close as possible to the true values of the physical plant state and output variables at any instant of time.*

The way that this "state estimator" is constructed is the subject of Section 4. The reason that this step is essential to the design process becomes apparent in the next step.

2.4.3. *Stochastic Feedback Control System Design (Step 3).* Let us recapitulate for a moment on what we have constructed up to now. From Step 1 we have:

(1a) an ideal deterministic set of commanded input time sequences

(1b) an ideal deterministic set of desired plant state variable and output time sequences.

From Step 2 we have:

(2a) a set of estimated plant state and output variable time sequences (which are hopefully close to the true plant state and output variables in the uncertain stochastic environment.)

We now have the capability to compare the estimated state variables (from 2a) to the desired state variables (from 2b), at each and every instant of time. *Their differences constitute a set of estimated deviations of the actual plant state variables from their ideal desired values at each instant of time.* Thus, we have an approximate idea on how close is the response of the physical process to its desired one.

In general, due to the presence of uncertainties and plant disturbances, one would expect to observe such an estimated deviation. One can now reformulate the control objective as follows:

Design the compensator such that all estimated deviations of the plant state variables from their ideal desired values are close to zero for all instants of time.

It should be clear that if we keep applying the ideal deterministic commanded input time functions (from Step 1a) that the above objective will not be met, since the deterministic input was found under assumptions (no uncertainty!) that are violated. Hence, one would expect that the actual commanded inputs to the physical process must be somewhat different than the ideal deterministic inputs found in Step 1.

One can imagine that this is accomplished by constructing a set of control correction signals (generated on the basis of the estimated deviations of the state variables from their desired values) such that the actual commanded input to the physical process is the sum, at any instant of time, of the ideal inputs obtained in Step 1 and of the control corrections.

2.5. *Why Not Dynamic Programming?*

Since we are obviously dealing with a stochastic optimal feedback control problem, and since the only theoretical tool which is available to analyze this class of problems is *dynamic programming* (see reference 4), one may wonder why one does not attack the problem directly using the dynamic programming algorithm. The reason that dynamic programming is not used, is simply due to the *curse of dimensionality* (which is far more severe for stochastic problems as compared to deterministic problems). Thus, to obtain numerical solutions for realistic problems, we simply do not have, now and in the foreseeable future, digital computers with sufficient fast-access memories to solve this class of problems. It is for this very reason that the "suboptimal" three part approach to stochastic system design has gained popularity among engineering practitioners; at the very least, the computational requirements of this design approach are perfectly within the capabilities of modest digital computers.

3. Deterministic Ideal Response Analysis and Design (Step 1)

3.1. *Introduction*

As indicated in Section 2.4.1 the first step in the proposed design process *assumes* (only temporarily!) that the physical process operates in the absence of uncertainty. In this section we elaborate on the steps and techniques which culminate in the deterministic ideal pair of inputs and associated state variable response sequences.

Our objective here is to indicate that additional "nonphysical" uncertainties can be introduced even if the physical process is assumed to operate in an otherwise deterministic physical environment. These "nonphysical" uncertainties are the results of modelling approximations. They lead to a structure similar to that of the overall stochastic problem even if all physical plant state variables can be measured exactly.

3.2. *Deterministic Modelling*

It is essential for the overall design process that the physical process be modelled in a quantitative manner.* This of course requires a blending of natural or man made laws, experimentation, econometrics, etc., so as to determine the nominal parameter values of the physical process.

Quite often assumptions that are made at this point are:

1. Actuator dynamics are neglected
2. Sensor dynamics are neglected
3. The plant is modelled as a lumped system.

3.2.1. *Actuator-Plant Model.* Under these assumptions the actuator and plant are modelled by a nonlinear time varying vector difference equation:

$$\mathbf{x}(t+1) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t); \mathbf{x}(0) = \mathbf{x}_0 \qquad t = 0, 1, 2, \ldots \tag{3.1}$$

where:

$\mathbf{x}(t)$ is the *plant state vector*, an n-dimensional vector with components $x_1(t), x_2(t), \ldots, x_n(t)$, for each value of the discrete time t

$\mathbf{u}(t)$ is the *plant control vector*, an m-dimensional vector with components $u_1(t), u_2(t), \ldots, u_m(t)$, for each value of the discrete time t

$\mathbf{x}_0 = \mathbf{x}(0)$ is the initial state vector at the initial time $t = 0$

$\mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t)$ is the plant nonlinearity, a vector-valued nonlinear function with components $f_1(\mathbf{x}(t), \mathbf{u}(t), t), f_2(\mathbf{x}(t), \mathbf{u}(t), t), \ldots, f_n(\mathbf{x}(t), \mathbf{u}(t), t)$.

Remarks:

1. In general, actuator dynamics, if significant, can be absorbed together with the plant dynamics, thus increasing the dimensionality of the state vector $\mathbf{x}(t)$.
2. The function $\mathbf{f}(\cdot, \cdot, \cdot)$ contains parameters whose values (nominal) are assumed known.
3. The function $\mathbf{f}(\cdot, \cdot, \cdot)$ is assumed (for technical reasons) continuous and at least twice differentiable with respect to its arguments $\mathbf{x}(t)$, $\mathbf{u}(t)$, and t.
4. The effects of known exogeneous variables, known time-varying parameters, etc., is absorbed in the time dependence of $\mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t)$.

3.2.2. *Sensor model.* We let the *output vector* $\mathbf{y}(t)$ denote the r-dimensional vector that represents the variables that can be measured. Thus the components $y_1(t)$, $y_2(t), \ldots, y_r(t)$ of $\mathbf{y}(t)$ denote the scalar variables that can be measured by the sensors, at each instant of time t.

We assume that each output variable is at most a nonlinear combination of the state variables. This is modelled by

$$\mathbf{y}(t) = \mathbf{g}(\mathbf{x}(t), t) \tag{3.2}$$

where $\mathbf{g}(\mathbf{x}(t), t)$ is called the output nonlinearity, a vector-valued function with components $g_1(\mathbf{x}(t), t), g_2(\mathbf{x}(t), t), \ldots, g_r(\mathbf{x}(t), t)$.

* We do not mean to minimize the difficulties associated with this step. Modelling is an extremely difficult process. The dynamics of aerospace systems are very simple and well understood as compared say to the dynamics of chemical process control systems or socio-economic systems. This is one of the reasons that most of the successful applications of the L-Q-G techniques have been to aerospace systems up to this time.

Remarks:

1. Sensor dynamics, if significant, can be incorporated in the plant equation (3.1).
2. The vector $\mathbf{g}(\mathbf{x}(t), t)$ is assumed to be continuous and at least twice differentiable (once more for technical reasons).

3.3. *Ideal Input-State-Output Responses*

Under our assumptions the following is true

Given

- the current value state vector $\mathbf{x}(t)$
- the control input sequence $\mathbf{u}(\tau)$, $\tau = t, t+1, t+2, \ldots$

Then

one can compute exactly the unique future state sequence $\{\mathbf{x}(\tau)\}$, $\tau = t+1, t+2, \ldots$

one can compute exactly the unique future output sequence $\{\mathbf{y}(t)\}$, $\tau = t$, $t+1, t+2, \ldots$

This capability allows us to determine the ideal deterministic input-state pair for any given initial state $\mathbf{x}_0$. In general, one is interested in the operation of the system over a finite time interval $t = 0, 1, 2, \ldots T$. On the basis of the deterministic model one then defines:

Ideal Deterministic Input Time Sequence: $\{\mathbf{u}_0(t)\}, t = 0, 1, 2, \ldots, T-1$

Ideal Deterministic State Sequence: $\{\mathbf{x}_0(t)\}, t = 0, 1, 2, \ldots, T$

Ideal Deterministic Output Time Sequence: $\{\mathbf{y}_0(t)\}, t = 0, 1, 2, \ldots, T$

$$\mathbf{x}_0(t+1) = \mathbf{f}(\mathbf{x}_0(t), \mathbf{u}_0(t), t); \mathbf{x}(0) = \mathbf{x}_0 \tag{3.3}$$

$$\mathbf{y}_0(t) = \mathbf{g}(\mathbf{x}_0(t), t)$$

3.3.1. *Computation of ideal input-state response.* The design procedure requires that to each initial state $\mathbf{x}_0$ we associate an input-state pair of time sequences $\{\mathbf{u}_0(t)\}$ and $\{\mathbf{x}_0(t)\}, t = 0, 1, 2, \ldots, T$. The interpretation of $\{\mathbf{x}_0(t)\}$ is that it represents the desired state evolution of the system, provided that the system initial state is $\mathbf{x}_0$.

In principle, $\{\mathbf{u}_0(t)\}$ and $\{\mathbf{x}_0(t)\}$ can be obtained by "experience" coupled with digital computer simulation. However, there is a systematic approach to the determination of $\{\mathbf{u}_0(t)\}$ and $\{\mathbf{x}_0(t)\}$ via the solution of a nonlinear deterministic optimal control problem. This involves the definition by the designer of a (nonquadratic in general!) scalar valued cost functional (objective function).

$$I = \phi(\mathbf{x}(T)) + \sum_{t=0}^{T-1} L(\mathbf{x}(t), \mathbf{u}(t), t) \tag{3.4}$$

which incorporates any requirements on the terminal state $\mathbf{x}(T)$ by means of the penalty function $\phi(\mathbf{x}(T))$, and any state variable constraints, control variable constraints, and optimality criteria in the function $L(\mathbf{x}(t), \mathbf{u}(t), t)$. In this case then one can formulate a discrete-time deterministic optimal control problem of the form.

Given the system (3.1) *and the initial state* $\mathbf{x}_0$. *Find* $\{\mathbf{u}_0(t)\}$ *and the resultant* $\{\mathbf{x}_0(t)\}$ *such that the cost functional* (3.4) *is minimized.*

Remark

It is possible to formulate and solve deterministic optimal control problems with "hard or soft" constraints on $\mathbf{x}(T)$, $\{\mathbf{u}(t)\}$, and $\{\mathbf{x}(t)\}$. However, these constraints will be violated in the stochastic version of the problem, because of the unpredictable disturbances. It is for this reason that in this step, the deterministic optimal control problem is formulated in an unconstrained manner.

3.3.2. *The discrete matrix minimum principle.* The deterministic discrete optimal control problem can be in principle solved by dynamic programming; in practice this is, however, not possible due to the curse of dimensionality.

The appropriate theoretical tool is the so-called *discrete minimum principle*, which is the extension of Pontryagin's maximum principle for continuous time systems (see reference 5), to discrete-time systems. Appropriate references for the discrete minimum principle are 6 to 8. We remark that the discrete maximum principle is essentially equivalent to the Kuhn–Tucker theorem.

In the control literature, the dynamics of the system (3.1) are written in a somewhat different form so as to make the discrete version of the minimum principle bear strong resemblance to the continuous time version. The statement of the problem and of the discrete version of the minimum principle are as follows:

Problem

Given a system described by the vector difference equation*

$$\mathbf{x}(t+1) - \mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t); \qquad \mathbf{x}(0) = \mathbf{x}_0 \tag{3.4}$$

and the cost functional, with T fixed

$$I = \phi(\mathbf{x}(T)) + \sum_{t=0}^{T-1} L(\mathbf{x}(t), \mathbf{u}(t), t) \tag{3.5}$$

Find the optimal control sequence, denoted by $\{\mathbf{u}^*(t)\}$, $t = 0, 1, 2, \ldots, T-1$ such that I is minimized.

The Hamiltonian

It is convenient to define the scalar valued function H, called the Hamiltonian, as follows:

$$H = H(\mathbf{x}(t), \mathbf{u}(t), \mathbf{p}(t+1), t) \triangleq L(\mathbf{x}(t), \mathbf{u}(t), t) + \mathbf{p}'(t+1)\mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t) \tag{3.6}$$

In (3.6) the sequence $\{\mathbf{p}(t)\}$, $t = 0, 1, 2, \ldots, T$ is called the *costate sequence*; $\mathbf{p}(t)$, the *costate vector* at time t, is an n-dimensional column vector (the same dimension as the state vector).

The Discrete Minimum Principle

Assume that an optimal control sequence exists. Let $\{\mathbf{u}^*(t)\}$ denote the optimal control sequence, let $\{\mathbf{x}^*(t)\}$ denote the optimal state sequence. Then there exists a corresponding costate sequence $\{\mathbf{p}^*(t)\}$ such that the following conditions hold.

* Note that $\mathbf{f}(\cdot)$ in (3.4) is not the same as the $\mathbf{f}(\cdot)$ in (3.1).

A. Canonical Difference Equations

State Dynamics

$$\mathbf{x}^*(t+1) - \mathbf{x}^*(t) = \left.\frac{\partial H}{\partial \mathbf{p}(t+1)}\right|_* = \mathbf{f}(\mathbf{x}^*(t), \mathbf{u}^*(t), t) \tag{3.7}$$

Costate Dynamics

$$\mathbf{p}^*(t+1) - \mathbf{p}^*(t) = -\left.\frac{\partial H}{\partial \mathbf{x}(t)}\right|_* = -\left.\frac{\partial L}{\partial \mathbf{x}(t)}\right|_* - \left(\left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}(t)}\right|_*\right)' \mathbf{p}(t+1) \tag{3.8}$$

B. Boundary Conditions

At the initial time $t = 0$:

$$\mathbf{x}^*(0) = \mathbf{x}_0 \tag{3.9}$$

At the terminal time T:

$$\mathbf{p}^*(T) = \left.\frac{\partial \phi(\mathbf{x}(T))}{\partial \mathbf{x}(T)}\right|_* \tag{3.10}$$

C. Minimization of Hamiltonian

$$H(\mathbf{x}^*(t), \mathbf{u}^*(t), \mathbf{p}^*(t+1), t) \leq H(\mathbf{x}^*(t), \mathbf{u}(t), \mathbf{p}^*(t+1), t) \tag{3.11}$$

for each $t = 0, 1, 2, \ldots, T-1$ and all $\mathbf{u}(t) \varepsilon R_m$. Under appropriate smoothness conditions, (3.11) implies

$$\left.\frac{\partial H}{\partial \mathbf{u}(t)}\right|_* = \mathbf{0} \tag{3.12}$$

$$\left.\frac{\partial^2 H}{\partial \mathbf{u}(t)^2}\right|_* = \text{positive semidefinite } m \times m \text{ matrix} \tag{3.13}$$

Remarks

1. The precise conditions under which the discrete minimum principle is derived will not be given here; see reference [6].
2. The discrete minimum principle yields, in general, a set of necessary conditions for optimality. Sufficiency requires additional (convexity type) assumptions.

3.3.3. *Computational algorithms.* The control literature abounds with a host of computational algorithms, that use the necessary conditions of the minimum principle, which can be used to obtain in an iterative manner numerical solutions to the optimal control problem.

The detailed description of these algorithms falls beyond the scope of this paper. The interested reader should consult references [9], Chapter 7, [10], [11] and [12]. Such methods as steepest descent, conjugate directions, conjugate gradient, quasilinearization, Newton's method, etc., are extremely popular. We note that many of these algorithms are presented in the cited literature for the continuous-time optimal control problem; however, their transliteration to the discrete time case is trivial.

3.3.4. *Recapitulation.* The solution of the deterministic optimal control problem, using the discrete minimum principle and the associated computational algorithms, can be used to compute the ideal input $\{\mathbf{u}_0(t)\}$, state $\{\mathbf{x}_0(t)\}$, and output $\{\mathbf{y}_0(t)\}$ sequences. We remark that this requires *off-line* computation.

3.4. *Control Under the Deterministic Assumption*

Let us now examine the interrelationship between our deterministic mathematical model and the physical process which appears in a deterministic environment.

Let $\{\mathbf{u}(t)\}, \{\mathbf{x}(t)\}, \{\mathbf{y}(t)\}$ denote the *true* input, state, and output sequences of the physical process. By assumption, all can be measured exactly. Let us imagine that we conduct the following experiment. We let

$$\mathbf{u}(t) = \mathbf{u}_0(t), \qquad t = 0, 1, 2, \ldots, T-1 \tag{3.14}$$

that is, we excite the physical system with the ideal input found in Section 3.3. Let us then measure the true state $\mathbf{x}(t)$ and output $\mathbf{y}(t)$ of the physical system.

The natural question that arises is:

$$\text{Is } \mathbf{x}(t) = \mathbf{x}_0(t) \text{ for all } t = 0, 1, 2, \ldots, T$$

$$\text{is } \mathbf{y}(t) = \mathbf{y}_0(t) \text{ for all } t = 0, 1, 2, \ldots, T?$$

In general, the answer is: *no*. The reason is that $\mathbf{x}_0(t)$ and $\mathbf{y}_0(t)$ were computed using a mathematical model of the physical process. However, the designer has to make some approximations (often intentionally) to arrive at the mathematical model, often neglecting to include second-order effects. Even if the equations were exact "structurally", the values of the parameters used in the mathematical model are *nominal* ones and the true values may be slightly different. In addition, the actual initial state of the system $\mathbf{x}(0)$ may differ slightly from the ideally assumed one, $\mathbf{x}_0(0)$.

It then follows that *errors in the deterministic model* may by themselves contribute to deviations of the true physical plant state $\mathbf{x}(t)$ from its ideal deterministic one $\mathbf{x}_0(t)$. In fact, small initial deviations, caused by the difference $\mathbf{x}(0) - \mathbf{x}_0(0)$, may get worse and worse as time goes on.

3.5. *Deterministic Perturbation Control Problem*

If we agree that our design objective is to keep the actual plant state $\mathbf{x}(t)$ near its ideal desired value $\mathbf{x}_0(t)$, for all $t = 0, 1, \ldots, T$, then it is clear that the actual plant input sequence $\{\mathbf{u}(t)\}$ must be different from the precomputed ideal input sequence $\{\mathbf{u}_0(t)\}$.

This leads to define the following quantities

- *State perturbation vector*; $\delta\mathbf{x}(t)$:

$$\delta\mathbf{x}(t) \triangleq \mathbf{x}(t) - \mathbf{x}_0(t) \tag{3.15}$$

- *Output perturbation vector*; $\delta\mathbf{y}(t)$:

$$\delta\mathbf{y}(t) \triangleq \mathbf{y}(t) - \mathbf{y}_0(t) \tag{3.16}$$

- *Control Correction Vector*; $\delta\mathbf{u}(t)$:

$$\delta\mathbf{u}(t) \triangleq \mathbf{u}(t) - \mathbf{u}_0(t) \tag{3.17}$$

We can imagine that the control correction vector, $\delta\mathbf{u}(t)$, is generated by a *deterministic controller* which is possibly driven by

- the state perturbation sequence, $\{\delta\mathbf{x}(t)\}$
- the output perturbation sequence, $\{\delta\mathbf{y}(t)\}$

Thus, even in this deterministic case, one must use *feedback control* to take care of errors that are primarily associated with errors in modelling. We remark that this is one of the primary reasons for feedback, namely to make the system response be relatively insensitive to parameter variations.*

The control objective can then be stated as follows:

Given $\delta\mathbf{x}(t)$ *and* $\delta\mathbf{y}(t)$, *find* $\delta\mathbf{u}(\tau)$, $\tau = t, t+1, \ldots,$ *such that future state perturbation vectors,* $\delta\mathbf{x}(\tau)$, *are "small" for all* $\tau = t+1, \ldots, T$.

3.6. *The Linear-Quadratic Approach to the Deterministic Controller Design*

Since the compensator to be designed involves a relationship between $\delta\mathbf{x}(t)$, $\delta\mathbf{u}(t)$, and $\delta\mathbf{y}(t)$, it is natural to ask at this point how these quantities are related. The sought for relationship can be obtained by Taylor series expansions which lead to the use of dynamic linearization ideas.

3.6.1. *The linearized perturbation model.* The deterministic model for our system is still employed (since we have no other!). Thus, we assume that the true control $\mathbf{u}(t)$, true state $\mathbf{x}(t)$, and true output $\mathbf{y}(t)$ are related by

$$\mathbf{x}(t+1) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t) \tag{3.18}$$

$$\mathbf{y}(t) = \mathbf{g}(\mathbf{x}(t), t) \tag{3.19}$$

Similarly the ideal nominal control $\mathbf{u}_0(t)$, state $\mathbf{x}_0(t)$, and output $\mathbf{y}_0(t)$ are related by

$$\mathbf{x}_0(t+1) = \mathbf{f}(\mathbf{x}_0(t), \mathbf{u}_0(t), t) \tag{3.20}$$

$$\mathbf{y}_0(t) = \mathbf{g}(\mathbf{x}_0(t), t) \tag{3.21}$$

Expanding $f(\mathbf{x}(t), \mathbf{u}(t), t)$ and $\mathbf{g}(\mathbf{x}(t), t)$ about $\mathbf{x}_0(t)$, $\mathbf{u}_0(t)$ in a Taylor series expansion we obtain

$$\mathbf{f}(\mathbf{x}(t)), \mathbf{u}(t), t) = \mathbf{f}(\mathbf{x}_0(t), \mathbf{u}_0(t), t) + \left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right|_0 \delta\mathbf{x}(t) + \left.\frac{\partial \mathbf{f}}{\partial \mathbf{u}}\right|_0 \delta\mathbf{u}(t) + \alpha_0(\delta\mathbf{x}(t), \delta\mathbf{u}(t), t) \tag{3.22}$$

$$\mathbf{g}(\mathbf{x}(t), t) = \mathbf{g}(\mathbf{x}_0(t), t) + \left.\frac{\partial \mathbf{g}}{\partial \mathbf{x}}\right|_0 \delta\mathbf{x}(t) + \beta_0(\delta\mathbf{x}(t), t) \tag{3.23}$$

where $\alpha_0(\delta\mathbf{x}(t), \delta\mathbf{u}(t), t)$ and $\beta_0(\delta\mathbf{x}(t), t)$ denote the higher order terms in the Taylor series expansions.

From the above we readily deduce that

$$\delta\mathbf{x}(t+1) = \mathbf{A}_0(t)\delta\mathbf{x}(t) + \mathbf{B}_0(t)\delta\mathbf{u}(t) + \alpha_0(\delta\mathbf{x}(t), \delta\mathbf{u}(t), t) \tag{3.24}$$

$$\delta\mathbf{y}(t) = \mathbf{C}_0(t)\delta\mathbf{x}(t) + \beta_0(\delta\mathbf{x}(t), t) \tag{3.25}$$

* Parameter variations will also be discussed in Section 7 of this paper.

In the above we use the notations:

$$\mathbf{A}_0(t) \triangleq \left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right|_0 \triangleq \left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right|_{\substack{\mathbf{x}_0(t) \\ \mathbf{u}_0(t)}} \tag{3.26}$$

is an $n \times n$ time-varying matrix which is obtained by evaluating the elements of the Jacobian matrix $\partial \mathbf{f}/\partial \mathbf{x}$ along the *known* (precomputed) time sequences $\{\mathbf{x}_0(t)\}$ and $\{\mathbf{u}_0(t)\}$.

$$\mathbf{B}_0(t) \triangleq \left.\frac{\partial \mathbf{f}}{\partial \mathbf{u}}\right|_0 \triangleq \left.\frac{\partial \mathbf{f}}{\partial \mathbf{u}}\right|_{\substack{\mathbf{x}_0(t) \\ \mathbf{u}_0(t)}} \tag{3.27}$$

is an $n \times m$ time-varying matrix which is obtained by evaluating the elements of the Jacobian matrix $\partial \mathbf{f}/\partial \mathbf{u}$ along the *known* (precomputed) time sequences $\{\mathbf{x}_0(t)\}$ and $\{\mathbf{u}_0(t)\}$.

$$\mathbf{C}_0(t) \triangleq \left.\frac{\partial \mathbf{g}}{\partial \mathbf{x}}\right|_0 \triangleq \left.\frac{\partial \mathbf{g}}{\partial \mathbf{x}}\right|_{\mathbf{x}_0(t)} \tag{3.28}$$

is an $r \times n$ time-varying matrix which is obtained by evaluating the elements of the Jacobian matrix $\partial \mathbf{g}/\partial \mathbf{x}$ along the *known* desired state sequence $\{\mathbf{x}_0(t)\}$. The equations (3.24) and (3.25), *including the highest-order-terms* represent the *exact** relationship between $\delta\mathbf{x}(t)$, $\delta\mathbf{u}(t)$, and $\delta\mathbf{y}(t)$.

The *linearized perturbation model* is obtained by setting the higher order terms equal to zero in equations (3.24) and (3.25) to obtain

$$\delta\mathbf{x}(t+1) = \mathbf{A}_0(t)\delta\mathbf{x}(t) + \mathbf{B}_0(t)\delta\mathbf{u}(t) \tag{3.29}$$

$$\delta\mathbf{y}(t) = \mathbf{C}_0(t)\delta\mathbf{x}(t) \tag{3.30}$$

which is a standard state description of a *linear* discrete-time-varying system.

Remark:

The linear perturbation model (3.29) and (3.30) represents only an *approximate* relationship between $\delta\mathbf{x}(t)$, $\delta\mathbf{u}(t)$, and $\delta\mathbf{y}(t)$, while (3.24) and (3.25) represents an *exact* model.

3.6.2. *Justification of the quadratic criterion.* As we have indicated before, the modelling aspects of a problem represent an extremely important part of the design process. The type of model is up to the designer; its relative accuracy is not of primary importance as long as one knows what are the effects of the approximations to be made. Up to this point, the fact that the mathematical model $\mathbf{x}(t+1) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t)$ was only an *approximation* to reality, forced us to introduce *feedback* and to seek the feedback controller.

At this stage, we are also faced with a similar problem. *The designer may wish to use the approximate linear perturbation model* (3.29) *and* (3.30) *rather than the more accurate nonlinear model* (3.24) *and* (3.25). It is really up to him to do so, provided that he can anticipate the effects of this choice upon the overall design.

The fact that must be kept in mind is that one cannot simply ignore the higher order terms and hope that they are indeed going to be small.

In order to trust the validity of the linear model the designer must guarantee that: the higher order terms $\alpha_0(\delta\mathbf{x}(t), \delta\mathbf{u}(t), t)$ *and* $\beta_0(\delta\mathbf{x}(t), t)$ *are indeed "small".*

* Assuming that $\mathbf{x}(t+1) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t)$ and $\mathbf{v}(t) = \mathbf{g}(\mathbf{x}(t), t)$ are exact!

To see how this philosophy leads to the use of quadratic criteria and the linear-quadratic deterministic optimal control problem, it becomes necessary to examine in more detail the higher order terms.

If we use Taylor's theorem which allows us to truncate a Taylor series at an arbitrary point we can represent *exactly* the higher order terms as follows:

$$\alpha_0(\delta\mathbf{x}(t), \delta\mathbf{u}(t), t) = \sum_{i=1}^{n} \boldsymbol{\phi}_i \delta\mathbf{x}'(t) \left[\frac{\partial^2 f_i(\cdot)}{\partial \mathbf{x}^2(t)} \bigg|_{\sim} \delta\mathbf{x}(t) + \delta\mathbf{u}'(t) \frac{\partial^2 f_i(\cdot)}{\partial \mathbf{u}^2(t)} \bigg|_{\sim} \delta\mathbf{u}(t) + 2\delta\mathbf{x}'(t) \frac{\partial^2 f_i(\cdot)}{\partial \mathbf{x}(t) \partial \mathbf{u}(t)} \bigg|_{\sim} \delta\mathbf{u}(t) \right] \tag{3.31}$$

$$\beta_0(\delta\mathbf{x}(t), t) = \sum_{i=1}^{n} \boldsymbol{\phi}_i \left[\delta\mathbf{x}'(t) \frac{\partial^2 g_i(\cdot)}{\partial \mathbf{x}^2(t)} \bigg|_{\sim} \delta\mathbf{x}(t) \right] \tag{3.32}$$

where

(a) $\boldsymbol{\phi}_i$ are the natural basis vectors in R_n (i.e., $\boldsymbol{\phi}_1' = [1 \quad 0 \ldots 0]$).

(b) The several second derivative (Hessian) matrices are evaluated at values $\tilde{\mathbf{x}}(t)$, $\tilde{\mathbf{u}}(t)$ which are in general *different* than $\mathbf{x}_0(t)$ and $\mathbf{u}_0(t)$; the values of $\tilde{\mathbf{x}}(t)$ and $\tilde{\mathbf{u}}(t)$ are, of course, *not* provided by Taylor's theorem.

The advantage of viewing the higher order terms in this context is that one can readily see that they are *quadratic* in $\delta\mathbf{x}(t)$ and $\delta\mathbf{u}(t)$. It is also clear that they involve certain unknown parameters since we do not know what $\tilde{\mathbf{x}}(t)$ and $\tilde{\mathbf{u}}(t)$ are!

This approach now leads to the following philosophy:

To trust the validity of the linear model, one should select the $\{\delta\mathbf{u}(t)\}$ *such that*

$$\sum_{t=0}^{T-1} \|\alpha_0(\delta\mathbf{x}(t), \delta\mathbf{u}(t), t)\| = \text{minimum} \tag{3.33}$$

$$\sum_{t=0}^{T-1} \|\beta_0(\delta\mathbf{x}(t), t)\| = \text{minimum.} \tag{3.34}$$

Since $\alpha_0(\cdot)$ *and* $\beta_0(\cdot)$ *are quadratic in* $\delta\mathbf{x}(t)$ *and* $\delta\mathbf{u}(t)$, *one way of guaranteeing this is to select* $\delta\mathbf{u}(t)$ *so that the "standard" quadratic cost functional**

$$J_0 = \delta\mathbf{x}'(T)\mathbf{Q}_0(T)\,\delta\mathbf{x}(T) + \sum_{t=0}^{T-1} [\delta\mathbf{x}'(t)\mathbf{Q}_0(t)\,\delta\mathbf{x}(t) + \delta\mathbf{u}'(t)\mathbf{R}_0(t)\,\delta\mathbf{u}(t)] \tag{3.35}$$

is minimized, where $\mathbf{Q}_0(t)$, $t = 0, 1, \ldots, T$, *are symmetric, at least positive semi-definite, matrices and* $\mathbf{R}_0(t)$ *is a symmetric positive definite matrix.*

The weighting matrices $\{\mathbf{Q}_0(t)\}$ and $\{\mathbf{R}_0(t)\}$ are selected† by the designer as an upper bound to the effects of the second derivative matrices in equations (3.31) and (3.32); the matrix $\mathbf{Q}_0(T)$ and the terminal penalty cost $\delta\mathbf{x}'(T)\mathbf{Q}_0(T)\,\delta\mathbf{x}(T)$ are often included to insure that the $\delta\mathbf{x}(T)$ stay near zero at the terminal time, when the current actions of $\delta\mathbf{u}(t)$ are not felt (since they take at least one unit of time to excite the system).

We can see that the state dependent part $(\delta\mathbf{x}'(T)\mathbf{Q}_0(T)\,\delta\mathbf{x}(T)$ and $\delta\mathbf{x}'(t)\mathbf{Q}_0(t)\,\delta\mathbf{x}(t))$ of the quadratic cost functional are *consistent* with the control objective of Section 3.5 which was to keep $\delta\mathbf{x}(t)$ "small." The difference is that

* One could also include a cross term of the form $\delta\mathbf{x}'(t)\mathbf{M}_0(t)\delta\mathbf{u}(t)$ in (3.35). This causes no difficulty in its solution. (See for example, references [13] and [14].)

† The selection of the weighting matrices will be discussed in Section 3.9.

the vague "smallness" requirement has been translated into something very specific, namely, to a quadratic penalty on the state deviations $\delta\mathbf{x}(t)$ from their desired zero values.

The above arguments have hopefully communicated to the reader the notion that *quadratic criteria can be used to keep a linear model as honest as possible.* If the designer loved to work with nonlinear difference equations that were *quadratic*, then the Taylor series should have been terminated at the cubic terms and a *cubic* criterion should have been used to validate the quadratic model. Since the linear-quadratic problem has a "nice" solution, it may not be necessary to increase the complexity of the perturbation differential equation model further than the linear one.

3.7. *Formal Statement and Solution of the Deterministic Linear-Quadratic Problem*

Using the above philosophy (i.e., keeping our linearized model honest) we have arrived at the following precise mathematical optimization problem.

3.7.1. *The deterministic linear-quadratic problem.* Given the linear, deterministic, time-varying system

$$\delta\mathbf{x}(t+1) = \mathbf{A}_0(t)\,\delta\mathbf{x}(t) + \mathbf{B}_0(t)\,\delta\mathbf{u}(t). \tag{3.36}$$

Given a fixed time interval of interest $t = 0, 1, 2, \ldots, T$.

Find the control perturbation vector sequence $\{\delta\mathbf{u}(t)\}$, such that the following deterministic quadratic cost functional is minimized:

$$J_0 = \delta\mathbf{x}'(T)\mathbf{Q}_0(T)\,\delta\mathbf{x}(T) + \sum_{t=0}^{T-1} [\delta\mathbf{x}'(t)\mathbf{Q}_0(t)\,\delta\mathbf{x}(t) + \delta\mathbf{u}'(t)\mathbf{R}_0(t)\,\delta\mathbf{u}(t)] \tag{3.37}$$

where*

$$\mathbf{Q}_0(t) = \mathbf{Q}_0'(t) \geq \mathbf{0} \quad \text{for all } t = 0, 1, \ldots, T \qquad (n \times n \text{ matrix}) \tag{3.38}$$

$$\mathbf{R}_0(t) = \mathbf{R}_0'(t) > \mathbf{0} \quad \text{for all } t = 0, 1, \ldots, T-1 \qquad (m \times m \text{ matrix}). \tag{3.39}$$

3.7.2. *Solution of the linear-quadratic problem.* The optimal control perturbation vector, $\delta\mathbf{u}(t)$, is related to the state perturbation vector, $\delta\mathbf{x}(t)$, by means of the *linear time-varying feedback* relationship

$$\delta\mathbf{u}(t) = -\mathbf{G}_0(t)\,\delta\mathbf{x}(t); \qquad t = 0, 1, 2, \ldots, T-1 \tag{3.40}$$

where $\mathbf{G}_0(t)$ is a sequence of $m \times n$ control gain matrices, $t = 0, 1, \ldots, T-1$; the value of $\mathbf{G}_0(t)$ is given by

$$\mathbf{G}_0(t) = [\mathbf{B}_0'(t)\mathbf{K}_0(t+1)\mathbf{B}_0(t) + \mathbf{R}_0(t)]^{-1}\mathbf{B}_0'(t)\mathbf{K}_0(t+1)\mathbf{A}_0(t) \tag{3.41}$$

where the $n \times n$ matrix $\mathbf{K}_0(t)$ is the solution of the matrix difference equation

$$\begin{aligned}\mathbf{K}_0(t) = {} & \mathbf{Q}_0(t) + \mathbf{A}_0'(t)\mathbf{K}_0(t+1)\mathbf{A}_0(t) \\ & - \mathbf{A}_0'(t)\mathbf{K}_0(t+1)\mathbf{B}_0(t)[\mathbf{B}_0'(t)\mathbf{K}_0(t+1)\mathbf{B}_0(t) + \mathbf{R}_0(t)]^{-1} \\ & \times \mathbf{B}_0'(t)\mathbf{K}_0(t+1)\mathbf{A}_0(t)\end{aligned} \tag{3.42}$$

* The notation $\mathbf{A} \geq \mathbf{B}$ means that $\mathbf{A} - \mathbf{B}$ is positive semidefinite; $\mathbf{A} > \mathbf{B}$ means that $\mathbf{A} - \mathbf{B}$ is positive definite.

subject to the boundary condition at the terminal time T

$$\mathbf{K}_0(T) = \mathbf{Q}_0(T). \tag{3.43}$$

3.7.3. *Methods of proof*. There are several ways of proving the above result. One way is using the discrete minimum principle and by subsequent manipulation of the necessary conditions (see, for example, [15]). Another way is through the use of the dynamic programming (see, for example, references [16], [9], [17] to [20]).

3.8. *Discussion*

The solution of the deterministic linear-quadratic problem provides us with a deterministic feedback design that attempts to null out deviations of the true state $\mathbf{x}(t)$ from its ideal response $\mathbf{x}_0(t)$. From a practical viewpoint this deterministic design is appealing because the sequence of the control gain matrices $\mathbf{G}_0(t)$ can be completely *precomputed*.

The only *practical* deficiency of this scheme is associated with the fact that we *cannot* measure the true state vector $\mathbf{x}(t)$ so as to construct $\delta\mathbf{x}(t)$. This fact alone provides us with sufficient motivation to examine the stochastic aspects of the problem in Section 4.

3.9. *Selection of the Weighting Matrices* $\mathbf{Q}_0(t)$, $\mathbf{R}_0(t)$

The selection of the weighting matrices in the quadratic criterion (3.37) is not a simple matter. Usually, they are selected by the designer on the basis of experience coupled with alternate simulation runs for different trial values. There is no universal agreement on precisely how these are to be selected for any given application; in the design of classes of aerospace systems several workers have developed rules of thumb on the relative values of the elements of these weighting matrices.

In most practical applications, $\{\mathbf{Q}_0(t)\}$, and $\{\mathbf{R}_0(t)\}$ are selected to be diagonal. In this manner, specific components of the state perturbation vector $\delta\mathbf{x}(t)$ and of the control perturbation vector $\delta\mathbf{u}(t)$ can be penalized individually; it helps to have a "physical" set of state variables and control variables so that relative weightings can be rationally assigned.

For economic applications, and the effects of changing the weights the work of Pindyck (in this volume and in References [21] and [22]) has shed valuable insight. Needless to say, the book of Holt *et al.*, (reference 23) contains specific suggestions for problems in management science.

From a pragmatic viewpoint one can develop certain qualitative properties which can help the designer in the choice of these important design parameters (these properties are decided from the dependence of equation (3.42) upon $\mathbf{Q}_0(t)$, and $\mathbf{R}_0(t)$).

1. The larger $\|\mathbf{Q}_0(T)\|$, the "larger" the control gain matrix $\mathbf{G}_0(t)$ for values of time near the terminal time.
2. The larger $\|\mathbf{Q}_0(t)\|$, the "larger" the gain matrix $\mathbf{G}_0(t)$ and the "faster" the time during which state perturbations are reduced to small values.
3. The larger $\|\mathbf{R}_0(t)\|$, the "smaller" the gain matrix $\mathbf{G}_0(t)$ and the "slower" the system.

From the point of view of the justification of quadratic criteria and honesty of linearization the size of the state weighting matrix $\mathbf{Q}_0(t)$ should somehow be proportional to estimates of the second derivative matrices $\partial^2 f_i/\partial \mathbf{x}^2(t)$—see eq. (3.31)—while the control weighting matrix $\mathbf{R}_0(t)$ should be related in a proportional manner to the second derivative matrices $\partial^2 f_i/\partial \mathbf{u}^2(t)$. Estimates of these second derivative matrices can often be obtained by evaluating them at the "nominal" values $\mathbf{u}_0(t)$ and $\mathbf{x}_0(t)$.

An alternate procedure has been suggested in the context of *perturbation guidance* or *neighboring optimal control* ([9], pp. 177–197). This approach is motivated by the fact that one can use the solution of an optimal control problem to determine the optimal control $\mathbf{u}_0(t)$ and the optimal state $\mathbf{x}_0(t)$ as outlined in Section 3.3.1 of this paper. The basic idea uses the Hamiltonian function given by eq. (3.6).

Let us suppose then that we use the discrete minimum principle to deduce the necessary conditions for optimality, and then apply an iterative algorithm to solve the nonlinear two point boundary value problem. As we remarked in Section 3.3.1, this procedure will yield the optimal control sequence $\{\mathbf{u}_0(t)\}$ and optimal state sequence $\{\mathbf{x}_0(t)\}$. However, as a by-product, we also obtain the associated costate sequence $\{\mathbf{p}_0(t)\}$.

The key idea behind the neighboring optimal control is to assume that the actual controls and states are somewhat different than the optimal ones. One then can substitute eqs. (3.6) and (3.7) into (3.4) and obtain the increase in the cost which is *approximately* measured by the *second variation*, $\delta^2 J$, and given by

$$\delta^2 J = \delta \mathbf{x}'(T) \left.\frac{\partial^2 \phi}{\partial \mathbf{x}(T)^2}\right|_0 \delta \mathbf{x}(T) + \sum_{t=0}^{T-1} [\delta \mathbf{x}'(t) \mid \delta \mathbf{u}'(t)] \left[\begin{array}{c|c} \left.\dfrac{\partial^2 H}{\partial \mathbf{x}(t)^2}\right|_0 & \left.\dfrac{\partial^2 H}{\partial \mathbf{x}(t)\partial \mathbf{u}(t)}\right|_0 \\ \hline \left.\dfrac{\partial^2 H}{\partial \mathbf{u}(t)\partial \mathbf{x}(t)}\right|_0 & \left.\dfrac{\partial^2 H}{\partial \mathbf{u}(t)^2}\right|_0 \end{array}\right] \left[\begin{array}{c} \delta \mathbf{x}(t) \\ \hline \delta \mathbf{u}(t) \end{array}\right]. \tag{3.44}$$

One then seeks the control $\delta \mathbf{u}(t)$ which minimizes the second variation $\delta^2 J$ subject to the linear difference equation constraints relating $\delta \mathbf{x}(t)$ to $\delta \mathbf{u}(t)$. This leads to a deterministic linear-quadratic optimal control problem and can be viewed as another justification for quadratic criteria ([9], p. 193).

In the above the second-derivative matrices of the Hamiltonian H

$$\left.\frac{\partial^2 H}{\partial \mathbf{x}(t)^2}\right|_0, \quad \left.\frac{\partial^2 H}{\partial \mathbf{x}(t)\partial \mathbf{u}(t)}\right|_0, \quad \left.\frac{\partial^2 H}{\partial \mathbf{u}(t)\partial \mathbf{x}(t)}\right|_0, \quad \left.\frac{\partial^2 H}{\partial \mathbf{u}(t)^2}\right|_0$$

and

$$\left.\frac{\partial^2 \phi}{\partial \mathbf{x}(t)^2}\right|_0$$

are all evaluated along the sequences $\{\mathbf{x}_0(t)\}$, $\{\mathbf{p}_0(t)\}$, $\{\mathbf{u}_0(t)\}$.

Intuitively speaking, this approach attempts to minimize (to second order only!) increases in the cost functional. If we neglect the cross coupling terms one could then make the association

$$\mathbf{Q}_0(T) \sim \left.\frac{\partial^2 \phi}{\partial \mathbf{x}(T)^2}\right|_0 \tag{3.45}$$

$$\mathbf{Q}_0(t) \sim \left.\frac{\partial^2 H}{\partial \mathbf{x}(t)^2}\right|_0 \tag{3.46}$$

$$\mathbf{R}_0(t) \sim \left.\frac{\partial^2 H}{\partial \mathbf{u}(t)^2}\right|_0. \tag{3.47}$$

However, there is *no* guarantee that these matrices enjoy any of the definiteness properties required for global existence and uniqueness of solutions to the linear quadratic problem; if these definiteness assumptions are violated, then one may have to deal with singular problems. Also, note that this philosophy neglects the contribution of the third, fourth, etc. variations in the cost; if these were going to be taken into account, then the partial derivatives of the Hamiltonian would have to be evaluated *not* at $\mathbf{x}_0(t)$, $\mathbf{p}_0(t)$, $\mathbf{u}_0(t)$, but at some other time sequences that are *not* known (as it was the case with the approach of keeping the linearizations honest). Nonetheless, this approach can often give the designer some clue as to the way these weighting matrices should be selected.

4. Stochastic Estimation Analysis and Design (Step 2)

4.1. *Introduction*

We have seen that even under the deterministic assumption we require a feedback controller to take care of errors in modelling. The main practical disadvantage of the deterministic design step was that exact measurement of all state variables was necessary. This is seldom the case in practical applications.

Even if one could measure all of the state variables, one has to use physical devices (sensors) to carry out these measurements. Physical sensors yield (more or less) inaccurate measurements. Thus, this uncertainty in measurement must somehow be taken into account.

In addition, although the deterministic approach admitted errors in modelling (necessitating feedback) it did not explicitly take into account errors introduced by the actuators; furthermore, it did not take into account that in many practical applications there are disturbance inputs acting upon the physical process, which are not generated by the control system (e.g., exogeneous stochastic variables).

In this section we shall present the common means by which such "input" and "sensor" errors are introduced in the design process, and how they affect the generation of an estimate $\hat{\mathbf{x}}(t)$ of the true state vector $\mathbf{x}(t)$, through the use of the Kalman–Bucy filter. Towards this goal we present in Section 4.2. some philosophical remarks pertaining to the use of *white noise* to model uncertainties in the design process. In Section 4.3 we formulate the combined modelling problem of using linearized dynamics and white Gaussian noise. In Section 4.4 we discuss

the linear-gaussian estimation problem and state its solution via the Kalman filter.

4.2. *The Use of White Noise*

It is common engineering practice to use a probabilistic approach to the modelling and implications of physical uncertainty. The reason is that a probabilistic approach is characterized by the existence of an extensive mathematical theory which has been already developed. Alternate approaches to uncertainty (e.g., via fuzzy sets, bounded but unknown uncertainty) have not as yet reached, from a mathematical viewpoint, the theoretical sophistication of the probabilistic approach.

In the design of dynamical systems the existence in time of plant disturbance sequences and sensor errors is modelled by representing these uncertain time functions by means of *random sequences*. For example, suppose that $n(t)$ is a random sequence which represents the "noise" that is introduced by a sensor at any time t. Hence, we can model sensor uncertainty by

$$z(t) = s(t) + n(t) \tag{4.1}$$

where at time t

- $z(t)$ is the actual sensor measurement
- $s(t)$ is the actual variable to be measured
- $n(t)$ is additive measurement noise.

The statistical properties of $n(t)$ in essence define the accuracy of the sensor at time t. At any time t_1, the scalar $n(t_1)$ is viewed as a random variable. Its *probability density function* $p(n(t_1))$ summarizes the statistical knowledge at time t_1. However, since $n(t)$ is associated with a particular sensor, one must also specify any statistical properties of the random variables $n(t_1)$ and $n(t_2)$ at any two distinct instants of time t_1 and t_2. Such statistical information is specified by the *joint probability density function* $p(n(t_1), n(t_2))$ of the random variables $n(t_1)$ and $n(t_2)$.

If $n(t_1)$ and $n(t_2)$ are dependent, then from Bayes rule we have

$$p(n(t_2)/n(t_1)) = \frac{p(n(t_1), n(t_2))}{p(n(t_1))} \tag{4.2}$$

which loosely implies that *if we have observed* $n(t_1)$ *then we can say something about* $n(t_2)$*—e.g., estimate its average value—before we actually measure* $n(t_2)$.

If on the other hand $n(t_1)$ and $n(t_2)$ are *independent* (uncorrelated in the Gaussian case), then

$$p(n(t_1), n(t_2)) = p(n(t_1))p(n(t_2)) \tag{4.3}$$

and Bayes rule yields

$$p(n(t_2)/n(t_1)) = p(n(t_2)) \tag{4.4}$$

which means that *the fact that we have already observed* $n(t_1)$ *does not help us at all to improve our knowledge about* $n(t_2)$.

These properties have significant implications from the point of view of the software that we have to utilize in our control system. If *a sensor noise is modelled as a correlated random process*, then we must expect some sort of estimation algorithm (based on eq. (4.2)) which *attempts to guess properties of future values of sensor noise based upon past measurements*. If this can be done (at the expense of, perhaps excessive, on-line computation) one can expect an improved "noise removing filter."

If on the other hand, we model the noise $n(t)$ as "uncorrelated," then past measurements do not help us at all in future guessing. In this case, the noise is *unpredictable* and no estimation algorithm that attempts to guess future values of the noise is required (and no on-line computations are required in this respect).

The above discussion dealt with time-structure of the noise $n(t)$. Let us return to the statistical properties of the noise at any instant of time t_1. As we mentioned before, this statistical information is contained in the probability density function $p(n(t_1))$. It is well known that two important statistical parameters (from an applications viewpoint) are the mean

$$E\{n(t_1)\} \triangleq \bar{n}(t_1) \tag{4.5}$$

and the variance

$$\text{var}\,[n(t_1)] = E\{n(t_1) - \bar{n}(t_1))^2\}. \tag{4.6}$$

The mean $\bar{n}(t_1)$ is what we would expect to see *on the average*. The variance helps us understand *how much this average* $\bar{n}(t_1)$ *is to be believed*. A large variance means that the actual value $n(t_1)$ (in any given experiment) may be way-off (with a large probability) from its mean value. A small variance means that the mean is a pretty good guess.

It is the opinion of the author that *the use of white noise in control system design is primarily a modelling issue*. The designer has to make a judgement on how to model uncertainties via white noise. There are no available cook-book procedures for doing this; the success of the design depends on the ability of the designer to know the physics of his problem and to *subjectively* translate this into mathematical probabilistic models. We shall comment on these problems in Section 4.7 in some more detail.

4.2.1. *Mathematical description of white noise*. The mathematical specification of white noise is as follows.

Let $\mathbf{n}(t)$ be a vector valued Gaussian white noise process with mean

$$E\{\mathbf{n}(t)\} = \mathbf{0} \qquad \text{for all } t = 0, 1, 2, \ldots, T \tag{4.7}$$

and covariance matrix

$$\text{cov}\,[\mathbf{n}(t); \mathbf{n}(\tau)] = E\{\mathbf{n}(t)\mathbf{n}'(\tau)\} = \mathbf{N}(t)\,\delta_{t\tau} \tag{4.8}$$

where

$$\mathbf{N}(t) = \mathbf{N}'(t) \geq \mathbf{0} \tag{4.9}$$

and $\delta_{t\tau}$ is the Kroenecker delta. If $\mathbf{N}(t) = \mathbf{N} =$ constant for all t, then we deal with *stationary* white noise.

4.3. *Stochastic Modelling for Control System Design*

Let us now return to the modelling issues associated with the control system design problem.

4.3.1. *Actuator-plant-input disturbance models.* Recall that in the deterministic version the relation of the true commanded input to the actuator, $\mathbf{u}(t)$, and that the true plant state, $\mathbf{x}(t)$, were related by the deterministic model (see eq. (3.1))

$$\mathbf{x}(t+1) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t). \tag{4.10}$$

In the stochastic case we can model* the actuator-plant-input disturbance part of our physical process by the stochastic difference equation

$$\mathbf{x}(t+1) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t) + \boldsymbol{\xi}(t) \tag{4.11}$$

where $\boldsymbol{\xi}(t)$ is a white noise process. The addition of $\boldsymbol{\xi}(t)$ to the otherwise deterministic model (4.10) implies that the designer is communicating to the mathematics one or more of the following "facts of life":

(1) That there are additional stochastic disturbances that drive the system
(2) That the deterministic equations may be in error due to over-simplification
(3) That some of the parameters in $\mathbf{f}(\cdot, \cdot, \cdot)$ may not be exact (true parameters may vary slightly from their nominal values)
(4) That the actuators introduce errors.

If we examine eq. (4.11) we can see that since $\boldsymbol{\xi}(t)$ is white, then $\mathbf{x}(t+1)$ can change, in part, in an unpredictable way. The deterministic part of the equation, $\mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t)$, represents the contribution of our completely predictable model; the stochastic part of the equation, $\boldsymbol{\xi}(t)$, stresses the unpredictable element of the real world. Loosely speaking, the use of eq. (4.11) is a way of saying to the mathematics "watch out! The deterministic equation is in error, but I will not tell you the structure of the error, so that you will not try to second-guess it in the future."

4.3.2. *Sensor and measurement error modelling.* Recall that in the deterministic case the type of sensors that could be used led to the definition of the output vector $\mathbf{y}(t)$ whose components were the variables that could be measured by the available sensors. The deterministic model was—see eq. (3.2)—

$$\mathbf{y}(t) = \mathbf{g}(\mathbf{x}(t), t). \tag{4.12}$$

The simplest way of modelling sensor errors is to assume that the sensor that measures the output variable $y_i(t)$ yields the measurement (data) signal $z_i(t)$ which equals $y_i(t)$ and additive white noise $\theta_i(t)$

$$z_i(t) = y_i(t) + \theta_i(t) \tag{4.13}$$

or, in vector notation,

$$\mathbf{z}(t) = \mathbf{y}(t) + \boldsymbol{\theta}(t) = \mathbf{g}(\mathbf{x}(t), t) + \boldsymbol{\theta}(t) \tag{4.14}$$

where $\boldsymbol{\theta}(t)$ is vector-valued white noise.

* This is the simplest possible model; more complex models can be considered, e.g.,

$$\mathbf{x}(t+1) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), \boldsymbol{\xi}(t), t)$$

4.3.3. *Initial uncertainty.* In the deterministic context we assume that the initial state of the plant $\mathbf{x}(0) \triangleq \mathbf{x}_0$ was known. Since the state variables cannot be measured, we can no longer make this assumption. The simplest way of modelling this is to view the initial state $\mathbf{x}_0$ as a vector-valued Gaussian random variable. Its mean and covariance matrix represent our *a priori* information about the initial conditions of our system.

4.3.4. *Statistical description.* We can see that the uncertainty in the overall physical process has been modelled in three separate parts

(1) *Initial uncertainty*: The initial state $\mathbf{x}_0$ is viewed as a random variable
(2) *Plant uncertainty*: The system is driven by the white noise $\boldsymbol{\xi}(t)$ which implies that the next state $\mathbf{x}(t+1)$ has an unpredictable component
(3) *Measurement uncertainty*: The output vector is corrupted by the additive white noise $\boldsymbol{\theta}(t)$, so that the measurement vector $\mathbf{z}(t)$ has an unpredictable component.

The quantitative description of this uncertainty is as follows.

The initial state vector is Gaussian with known mean and covariance matrix, i.e.,

$$E\{\mathbf{x}_0\} \triangleq \bar{\mathbf{x}}_0 \text{ (assumed known)} \tag{4.15}$$

$$\operatorname{cov}[\mathbf{x}_0; \mathbf{x}_0] \triangleq E\{(\mathbf{x}_0 - \bar{\mathbf{x}}_0)(\mathbf{x}_0 - \bar{\mathbf{x}}_0)'\} = \boldsymbol{\Sigma}_0 \text{ (assumed known)}; \tag{4.16}$$

$$\boldsymbol{\Sigma}_0 = \boldsymbol{\Sigma}_0' \geq \mathbf{0}.$$

The plant driving noise $\boldsymbol{\xi}(t)$ is discrete-time white, Gaussian, with zero mean and known covariance matrix for all t, i.e.,

$$E\{\boldsymbol{\xi}(t)\} = \mathbf{0} \qquad \text{for all } t \tag{4.17}$$

$$\begin{aligned} \operatorname{cov}[\boldsymbol{\xi}(t); \boldsymbol{\xi}(\tau)] &\triangleq E\{\boldsymbol{\xi}(t)\boldsymbol{\xi}'(\tau)\} \\ &\triangleq \boldsymbol{\Xi}(t)\,\delta_{t\tau};\ \boldsymbol{\Xi}(t) = \boldsymbol{\Xi}'(t) \geq \mathbf{0} \qquad \text{for all } t \text{ (assumed known).} \end{aligned} \tag{4.18}$$

The measurement noise $\boldsymbol{\theta}(t)$ is white, Gaussian, with zero mean and known covariance matrix for all $t \geq t_0$, i.e.,

$$E\{\boldsymbol{\theta}(t)\} = \mathbf{0} \qquad \text{for all } t \tag{4.19}$$

$$\begin{aligned} \operatorname{cov}[\boldsymbol{\theta}(t); \boldsymbol{\theta}(\tau)] &\triangleq E\{\boldsymbol{\theta}(t)\boldsymbol{\theta}'(\tau)\} \\ &= \boldsymbol{\Theta}(t)\,\delta_{t\tau};\ \boldsymbol{\Theta}(t) = \boldsymbol{\Theta}'(t) \geq \mathbf{0} \qquad \text{for all } t \text{ (known).} \end{aligned} \tag{4.20}$$

Furthermore, one usually assumes that $\mathbf{x}_0$, $\boldsymbol{\xi}(t)$, and $\boldsymbol{\theta}(\tau)$ are mutually independent, i.e.,

$$\operatorname{cov}[\mathbf{x}_0; \boldsymbol{\xi}(t)] = \mathbf{0} \qquad \text{for all } t \tag{4.21}$$

$$\operatorname{cov}[\mathbf{x}_0; \boldsymbol{\theta}(t)] = \mathbf{0} \qquad \text{for all } t \tag{4.22}$$

$$\operatorname{cov}[\boldsymbol{\xi}(t); \boldsymbol{\theta}(\tau)] = \mathbf{0} \qquad \text{for all } t, \tau. \tag{4.23}$$

This assumption is reasonable in most physical applications.

We shall discuss later the selection of the covariance matrices $\boldsymbol{\Xi}(t)$ and $\boldsymbol{\Theta}(t)$, which govern the "strength" of the white noise sequences $\boldsymbol{\xi}(t)$ and $\boldsymbol{\theta}(t)$, respectively, as well as of the initial covariance matrix $\boldsymbol{\Sigma}_0$.

4.3.5. *Linearized dynamics.* Let us recall that one of the byproducts of the deterministic analysis was to specify an ideal deterministic time sequence, $\{\mathbf{u}_0(t)\}$ an ideal deterministic state sequence, $\{\mathbf{x}_0(t)\}$, and an ideal output sequence $\{\mathbf{y}_0(t)\}$. Our control system objective was to augment $\mathbf{u}_0(t)$ by the control correction vector $\delta\mathbf{u}(t)$ so that the commanded control $\mathbf{u}(t) = \mathbf{u}_0(t) + \delta\mathbf{u}(t)$ had the property that the state deviation vector $\delta\mathbf{x}(t) = \mathbf{x}(t) - \mathbf{x}_0(t)$ was "small" for all future values of t.

Our control objective has not changed except that $\mathbf{x}(t)$, $\mathbf{u}(t)$, $\delta\mathbf{x}(t)$, and $\delta\mathbf{u}(t)$ are now random sequences (rather than deterministic). We still would expect to generate $\delta\mathbf{u}(t)$ by means of some *feedback* arrangement which is based on the actual sensor measurements $\mathbf{z}(t)$.

Let us also recall that associated with the ideal state response $\mathbf{x}_0(t)$ we had an ideal measurement vector $\mathbf{y}_0(t)$ (see eq. (3.4)) and an output perturbation vector $\delta\mathbf{y}(t) = \mathbf{y}(t) - \mathbf{y}_0(t)$ (see eq. (3.7)).

Since our measurement vector is given by $\mathbf{z}(t) = \mathbf{y}(t) + \boldsymbol{\theta}(t)$, then

$$\mathbf{z}(t) = \mathbf{y}_0(t) + \delta\mathbf{y}(t) + \boldsymbol{\theta}(t). \tag{4.24}$$

Arbitrarily we define

$$\mathbf{z}(t) \triangleq \mathbf{z}_0(t) + \delta\mathbf{z}(t) \tag{4.25}$$

where

$$\mathbf{z}_0(t) \triangleq \mathbf{y}_0(t) = \mathbf{g}(\mathbf{x}_0(t), t) \tag{4.26}$$

and

$$\delta\mathbf{z}(t) = \delta\mathbf{y}(t) + \boldsymbol{\theta}(t). \tag{4.27}$$

Note that $\mathbf{z}_0(t)$ is a deterministic precomputable quantity. Hence $\delta\mathbf{z}(t)$ can be evaluated.

A repeat of the development of the Taylor series expansions (see Section 3.6.1) about $\mathbf{x}_0(t)$, $\mathbf{u}_0(t)$, $\mathbf{y}_0(t)$, using the stochastic models leads to the following set of equations:

$$\delta\mathbf{x}(t+1) = \mathbf{A}_0(t)\,\delta\mathbf{x}(t) + \mathbf{B}_0(t)\,\delta\mathbf{u}(t) + \boldsymbol{\xi}(t) + \boldsymbol{\alpha}_0(\delta\mathbf{x}(t), \delta\mathbf{u}(t) \tag{4.28}$$

$$\delta\mathbf{y}(t) = \mathbf{C}_0(t)\,\delta\mathbf{x}(t) + \boldsymbol{\beta}_0(\delta\mathbf{x}(t)) \tag{4.29}$$

$$\delta\mathbf{z}(t) = \mathbf{C}_0(t)\,\delta\mathbf{x}(t) + \boldsymbol{\theta}(t) + \boldsymbol{\beta}_0(\delta\mathbf{x}(t)). \tag{4.30}$$

In the above

(a) The matrices $\mathbf{A}_0(t)$, $\mathbf{B}_0(t)$, $\mathbf{C}_0(t)$ are still given by eqs. (3.26), (3.27), and (3.28) respectively. They are *deterministic and precomputable.*

(b) The vectors $\boldsymbol{\alpha}_0(\cdot,\cdot)$ and $\boldsymbol{\beta}_0(\cdot)$ represent the effects of the quadratic and higher order terms; they are stochastic sequences, since at least $\{\delta\mathbf{x}(t)\}$ is a stochastic sequence.

Once more we *define* the linearized stochastic (approximate) model by:

$$\mathbf{x}(t+1) = \mathbf{A}_0(t)\,\delta\mathbf{x}(t) + \mathbf{B}_0(t)\,\delta\mathbf{u}(t) + \boldsymbol{\xi}(t) \tag{4.31}$$

$$\delta\mathbf{z}(t) = \mathbf{C}_0(t)\,\delta\mathbf{x}(t) + \boldsymbol{\theta}(t) \tag{4.32}$$

simply by ignoring $\boldsymbol{\alpha}_0(\cdot)$ and $\boldsymbol{\beta}_0(\cdot)$ in eqs. (4.28) and (4.30) respectively.

To be sure eqs. (4.31) and (4.32) represent only approximations to eqs. (4.28) and (4.30). However, both equations contain the white noise driving term which at least is a "flag" to the mathematics that the linearized equations are "in error."

We can now see even more clearly the role of the white noises in modelling. Up to this point, the noise $\boldsymbol{\xi}(t)$ could be used to model input uncertainties and deterministic modelling errors. Now we see that it can also be used to model the fact that the higher order terms have been neglected in the use of eq. (4.31) instead of (4.28). Thus, the choice of the covariance matrix for $\boldsymbol{\xi}(t)$

$$\text{cov}\,[\boldsymbol{\xi}(t);\boldsymbol{\xi}(\tau)] = \boldsymbol{\Xi}(t)\,\delta_{t\tau} \tag{4.33}$$

i.e., the value of $\boldsymbol{\Xi}(t)$ selected by the designer, should incorporate his judgment on the importance of the higher order terms in the validity of the linearized model. Thus, the "more nonlinear" the system dynamics, the "larger" $\boldsymbol{\Xi}(t)$ should be used.

The white noise $\boldsymbol{\theta}(t)$ (assumed independent of $\boldsymbol{\xi}(t)$) in the observation equation (4.32) plays a similar role. Not only should it reflect the inherent uncertainty of the measurements due to sensor inaccuracies, but it should also be used to model the implications of neglecting $\boldsymbol{\beta}_0(\cdot)$ in eq. (4.30) to obtain the linear equation (4.32). Since

$$\text{cov}\,[\boldsymbol{\theta}(t);\boldsymbol{\theta}(\tau)] = \boldsymbol{\Theta}(t)\,\delta_{t\tau} \tag{4.34}$$

then the "more nonlinear" the output nonlinearity $\mathbf{g}(\mathbf{x}(t), t)$ is, the "larger" $\boldsymbol{\Theta}(t)$ should be selected.

Admittedly, we are cheating in our quest for linear models. However, the use of white noises allows us to communicate to the mathematics our estimate of the "degree of cheating." This is extremely important because as we shall see in the next section we shall ask some very precise questions of the mathematics. If we ask precise but stupid questions, we shall get precise but stupid answers!

4.4. *The Estimation (Filtering) Problem*

We have seen in eq. (4.25) that we can construct the signal $\delta\mathbf{x}(t)$ from the actual sensor measurement $\mathbf{z}(t)$. The state perturbation $\delta\mathbf{x}(t)$ is still the deviation of the actual state $\mathbf{x}(t)$ from the desired ideal state response $\mathbf{x}_0(t)$. However, $\delta\mathbf{x}(t)$ *cannot be measured directly*; it is related, however, to the available signal $\delta\mathbf{z}(t)$ by eq. (4.31). The future evolution of $\delta\mathbf{x}(t)$ can be influenced by the control correction vector $\delta\mathbf{u}(t)$ according to eq. (4.31).

We still want to keep $\delta\mathbf{x}(t)$ small by selecting $\delta\mathbf{u}(t)$. We have seen how this can be done in the deterministic case if $\delta\mathbf{x}(t)$ were known. Since in this case $\delta\mathbf{x}(t)$ is not directly available, then we can ask the following question:

> *Is it possible to generate a "good" estimate* $\delta\hat{\mathbf{x}}(t)$ *of* $\delta\mathbf{x}(t)$, *based on the measurements made up to time* t, *for any given time function* $\delta\mathbf{u}(t)$?

The Kalman filter presents a precise way of obtaining such an estimate.

4.4.1. *Statement of the filtering problem.* Given the linear dynamic stochastic system

$$\delta\mathbf{x}(t+1) = \mathbf{A}_0(t)\,\delta\mathbf{x}(t) + \mathbf{B}_0(t)\,\delta\mathbf{u}(t) + \boldsymbol{\xi}(t) \tag{4.35}$$

and the linear stochastic measurement equation

$$\delta \mathbf{z}(t) = \mathbf{C}_0(t)\,\delta \mathbf{x}(t) + \boldsymbol{\theta}(t). \tag{4.36}$$

It is assumed that $\mathbf{A}_0(t)$, $\mathbf{B}_0(t)$, $\mathbf{C}_0(t)$, $\delta \mathbf{u}(t)$ are deterministic and known. It is assumed that the white noise $\boldsymbol{\xi}(t)$ has the statistics specified by eqs. (4.17) and (4.18). It is also assumed that the white noise $\boldsymbol{\theta}(t)$ has the statistics specified by eqs. (4.19) and (4.20). It can also be shown that $\delta \mathbf{x}(t_0)$ is a Gaussian random vector with mean (see eq. (4.15)).

$$E\{\delta \mathbf{x}(t_0)\} \triangleq \delta \bar{\mathbf{x}}_0 = \bar{\mathbf{x}}_0 - \mathbf{x}_0(0) \qquad \text{(known)} \tag{4.37}$$

and covariance matrix $\boldsymbol{\Sigma}_0$ (see eq. (4.16)).

Given the measured signal* $\delta \mathbf{z}(\tau)$, for all $\tau = 1, 2, \ldots, t$.

Find a vector $\delta \hat{\mathbf{x}}(t)$, an estimate of the true $\delta \mathbf{x}(t)$ which is "optimal" in a well defined statistical sense.

We remark that the linear-gaussian nature of the hypotheses allows us to define a variety of optimization criteria (least-squares, minimum variance, maximum likelihood, etc.). They all lead to the same answer. For example, one can show that the above assumptions imply that the *a posteriori* density function of $\delta \mathbf{x}(t)$

$$p(\delta \mathbf{x}(t) | \delta \mathbf{z}(\tau) : \tau = 1, 2, \ldots, t) \tag{4.38}$$

is Gaussian and $\delta \hat{\mathbf{x}}(t)$, as generated by the Kalman filter, is the conditional mean (see [9] and [18]).

4.4.2. *The discrete Kalman filter.* The easiest way of writing the equations of the discrete-time Kalman filter is to divide the calculations into two cycles

(a) a predict cycle, and

(b) an update cycle.

This subdivision motivates us to a somewhat different notation.

We let t denote the present value of time. We assume that we have available

(a) the past measurements up to and including the current measurement

$$\delta \mathbf{z}(1), \delta \mathbf{z}(2), \ldots, \delta \mathbf{z}(t)$$

(b) The past control corrections

$$\delta \mathbf{u}(0), \delta \mathbf{u}(1), \ldots, \delta \mathbf{u}(t-1).$$

It is convenient to summarize this *information set* at time t by $S(t)$; thus $S(t)$ is the set

$$S(t) = \{\delta \mathbf{z}(1), \delta \mathbf{z}(2), \ldots, \delta \mathbf{z}(t), \delta \mathbf{u}(0), \ldots, \delta \mathbf{u}(t-1)\}. \tag{4.39}$$

Now we define the following:

(a) $\delta \hat{\mathbf{x}}(t|t)$ is called the *updated estimate* of $\delta \mathbf{x}(t)$ given the information set $S(t)$; under the linearity and Gaussian assumptions $\delta \hat{\mathbf{x}}(t|t)$ is the conditional mean of $\delta \mathbf{x}(t)$, i.e.,

$$\delta \hat{\mathbf{x}}(t|t) = E\{\delta \mathbf{x}(t) | S(t)\}. \tag{4.40}$$

* Here we assume that if the initial time is $t = 0$, the first measurement is taken at time $t = 1$. This assumption is by no means crucial and it can be replaced by assuming that the first measurement occurs at $t = 0$.

(b) $\Sigma_0(t|t)$ is the *updated covariance matrix* of $\delta\mathbf{x}(t)$ given the information set $S(t)$; under the linearity and Gaussian assumptions $\Sigma(t|t)$ is the conditional covariance of $\delta\mathbf{x}(t)$, i.e.,

$$\Sigma_0(t|t) = \operatorname{cov}[\delta\mathbf{x}(t); \delta\mathbf{x}(t)|S(t)] \tag{4.41}$$
$$= E\{(\delta\mathbf{x}(t) - \delta\hat{\mathbf{x}}(t|t))(\delta\mathbf{x}(t) - \delta\hat{\mathbf{x}}(t|t))'|S(t)\}.$$

(c) $\delta\hat{\mathbf{x}}(t+1|t)$ is called the one-step predicted estimate of $\delta\mathbf{x}(t+1)$ given the information set $S(t)$, i.e., before the measurement $\delta\mathbf{z}(t+1)$ at time $t+1$ is obtained; under the linearity and Gaussian assumptions

$$\delta\hat{\mathbf{x}}(t+1|t) = E\{\delta\mathbf{x}(t+1)|S(t), \delta\mathbf{u}(t)\}. \tag{4.42}$$

(d) $\Sigma_0(t+1|t)$ is the *one-step predicted covariance matrix* of $\delta\mathbf{x}(t+1)$ given the information set $S(t)$; under the linearity and Gaussian assumptions

$$\Sigma_0(t+1|t) = \operatorname{cov}[\delta\mathbf{x}(t+1); \delta\mathbf{x}(t+1)|S(t), \delta\mathbf{u}(t)] \tag{4.43}$$
$$= E\{(\delta\mathbf{x}(t+1) - \widehat{\delta\mathbf{x}}(t+1|t))(\delta\mathbf{x}(t+1)$$
$$- \widehat{\delta\mathbf{x}}(t+1|t)'|S(t), \delta\mathbf{u}(t)\}.$$

Thus in the expressions for $\delta\mathbf{x}(\cdot|\cdot)$ and $\Sigma_0(\cdot|\cdot)$ the first symbol denotes the actual value of time while the second time denotes the last value of time at which the information was utilized.

Under the linear-Gaussian assumptions the Kalman filter generates the above conditional means. Its detailed description is as follows:

Initialization: At $t = 0$

$$\delta\mathbf{x}(0|0) = \delta\bar{\mathbf{x}}_0 = E\{\delta\mathbf{x}(0)\} = \text{prior mean} \tag{4.44}$$

$$\Sigma_0(0|0) = \Sigma_0 = \operatorname{cov}[\delta\mathbf{x}(0); \delta\mathbf{x}(0)] = \text{prior covariance.} \tag{4.45}$$

One now proceeds in a recursive manner: for any $t = 0, 1, 2, \ldots$ assume that $\delta\mathbf{x}(t|t)$ and $\Sigma_0(t|t)$ are available. One then needs an algorithm that is based upon

(1) the value of $\delta\mathbf{u}(t)$

(2) the measurement $\delta\mathbf{z}(t+1)$

one generates

(1) $\widehat{\delta\mathbf{x}}(t+1|t+1) = E\{\delta\mathbf{x}(t+1)|S(t+1)\}$

(2) $\Sigma_0(t+1|t+1) = \operatorname{cov}[\delta\mathbf{x}(t+1); \delta\mathbf{x}(t+1)|S(t+1)]$.

The predict cycle

First one generates the one-step predicted estimate $\widehat{\delta\mathbf{x}}(t+1|t)$ by

$$\widehat{\delta\mathbf{x}}(t+1|t) = \mathbf{A}_0(t)\,\widehat{\delta\mathbf{x}}(t|t) + \mathbf{B}_0(t)\,\delta\mathbf{u}(t). \tag{4.46}$$

Next, one generates the one step predicted covariance matrix $\Sigma(t+1|t)$ by

$$\Sigma_0(t+1|t) = \mathbf{A}_0(t)\Sigma_0(t|t)\mathbf{A}_0'(t) + \Xi(t). \tag{4.47}$$

The update cycle

At time $t+1$ one makes the measurement $\delta\mathbf{z}(t+1)$. This can be used to improve the estimates obtained in the predict cycle. First one computes the

updated covariance $\Sigma_0(t+1|t+1)$ from the matrix difference equation

$$\begin{aligned}(4.48)\quad \Sigma_0(t+1|t+1) &= \Sigma_0(t+1|t)\\ &\quad - \Sigma_0(t+1|t)C_0'(t+1)[C_0(t+1)\Sigma_0(t+1|t)C_0'(t+1)\\ &\quad + \Theta(t+1)]^{-1}C_0(t+1)\Sigma_0(t+1|t).\end{aligned}$$

Next one obtains the updated estimate $\widehat{\delta x}(t+1|t+1)$ of $\widehat{\delta x}(t+1)$ from the vector difference equation

$$\begin{aligned}(4.49)\quad \widehat{\delta x}(t+1|t+1) &= \widehat{\delta x}(t+1|t) + \Sigma_0(t+1|t+1)C_0'(t+1)\Theta^{-1}(t+1)\\ &\quad \times [\delta z(t+1) - C_0(t+1)\,\widehat{\delta x}(t+1|t)].\end{aligned}$$

Remarks

1. Note that the equations that propagate the covariance matrices, (4.47) and (4.48), are independent of
   (a) the actual applied $\delta u(t)$
   (b) the actual measurement $\delta z(t+1)$
   i.e., independent of the information set. Hence the sequence of $\Sigma_0(t|t)$, $t = 0, 1, 2, \ldots$ can be computed *off-line*. The only on-line computations are those specified by (4.46) and (4.49), i.e., the propagation of the estimates.
2. If the Gaussian assumption is removed, the Kalman filter does no longer generate the conditional mean. However, *among the class of linear estimators*, the estimate $\widehat{\delta x}(t|t)$ is optimal in a least squares sense.
3. The quantity

$$(4.50)\qquad \delta r(t+1) \triangleq \delta z(t+1) - C_0(t+1)\,\widehat{\delta x}(t+1|t)$$

   is often called the *residual* or *innovations sequence*. This is the difference between the actual measurement, $\delta z(t+1)$, and what we expected the measurement to be, $C_0(t+k)\,\widehat{\delta x}(t+1|t)$. If the linearity assumptions are valid, then the residual sequence is *white* with zero mean and covariance matrix

$$(4.51)\qquad \operatorname{cov}[\delta r(t); \delta r(\tau)] = [C_0(t)\Sigma_0(t|t)C_0'(t) + \Theta(t)]\,\delta_{t\tau}.$$

   We remark that *on-line* test for the whiteness of the residual sequence are often used to calculate the degree of modelling error (see [30]).
4. By defining the so-called Kalman filter gain matrix $H_0(t+1)$ by

$$(4.52)\qquad H_0(t+1) = \Sigma_0(t+1|t+1)C_0'(t+1)\Theta^{-1}(t+1).$$

   which can be computed *off-line*, and by substituting eqs. (4.46), (4.50), and (4.52) into eq. (4.49) we obtain

$$(4.53)\quad \widehat{\delta x}(t+1|t+1) = \underbrace{A_0(t)\,\widehat{\delta x}(t|t) + B_0\,\delta u(t)}_{\widehat{\delta x}(t+1|t)} + H_0(t+1)\,\delta r(t+1).$$

This shows that the "larger" the filter gain matrix $H_0(t+1)$ the more the residual is used (and hence the actual measurement) to correct the predicted

estimate $\hat{\delta x}(t+1|t)$. The factors that contribute to a "large" gain matrix $\mathbf{H}_0(t+1)$ are (see eq. (4.52)).

(a) large current uncertainty, i.e., large $\Sigma_0(t+1|t+1)$

(b) large signal to noise ratio, i.e., large $\mathbf{C}_0(t+1)$

(c) small measurement noise, i.e., small $\Theta(t+1)$ and hence large $\Theta^{-1}(t+1)$.

Thus, the Kalman filter combines in a systematic way the state of knowledge about the system uncertainty, so as to decide each time that a noisy measurement is made, its relative value in correcting the available estimates.

4.4.3. *Derivations of the Kalman–Bucy filter.* Since the original publication of Kalman [24] there have been many different derivations of these results each contributing to enhanced understanding to the advantages and shortcomings of these techniques (see references 9, 12, 18, 25 to 29) as well as extensions to the nonlinear case (see references 9, 25).

4.5. *Discussion*

Most of the difficulties that are encountered with the Kalman–Bucy filter are primarily related to

(a) model mismatching (i.e., the model used in the implementation of the Kalman–Bucy filter is different than the physical process), and

(b) correct selection of Σ_0, and of the white noise covariance matrices $\Xi(t)$ and $\Theta(t)$.

In pure filtering situations these contribute to the so called *divergence* of the Kalman filter. There are several analyses that have been carried out that considered the effects and implications of selecting the wrong covariance matrices (see, for example, [33] pp. 376–419).

The existence of unknown biases in the noises $\xi(t)$ and $\theta(t)$ are not as troublesome since they can be estimated by an augmented Kalman filter, at the expense of introducing additional state variables. Some research efforts have been directed toward simultaneous estimation of the state variables and the covariance matrices (see [30]).

The sensitivity, and possible divergence, of the Kalman–Bucy filter is then intimately related to the modelling issues. If we view the (wrong) linearized model as a constraint, then the designer can attempt to minimize the filter sensitivity by judicious choice of the covariance matrices $\Xi(t)$ and $\Theta(t)$. Considerable success has been obtained in certain classes of application problems (re-entry vehicle tracking, orbit determination) by increasing the covariance matrix $\Xi(t)$ to compensate for modelling errors, which arise primarily in the dynamical equations. However, these techniques were developed only after excessive Monte-Carlo simulations and trial-and-error approaches. There is need for systematic approaches to this most important problem of selecting $\Xi(t)$ and $\Theta(t)$, and this represents an exciting research area.

Loosely speaking, the effect of increasing the magnitude of the covariance matrix $\Xi(t)$ (fake plant noise) results in larger values of the covariance matrix $\Sigma_0(t|t)$—see eqs. (4.47) and (4.48)—and this leads to an increase in the filter gain matrix $\mathbf{H}_0(t+1)$—see eq. (4.52). Qualitatively speaking, the residuals are then weighted more (the filter is paying more attention to the *actual* measurements to compensate for errors in the *a priori* values of $\mathbf{A}_0(t)$, $\mathbf{B}_0(t)$, $\Xi(t)$, $\Theta(t)$ and Σ_0) and

one obtains a high-gain filter. Thus, an increase in $\Xi(t)$ causes the filter to have a wider "bandwidth." This bandwidth interpretation is useful since an increased $\Xi(t)$ means that the plant white noise $\xi(t)$ has more power and, hence, causes more "wiggles" in the actual state $\mathbf{x}(t)$; the filter must estimate these "wiggles" in $\mathbf{x}(t)$ and this requires higher "bandwidth." Of course, a higher "bandwidth" passes more of the measurement noise $\theta(t)$ and this is the price that one must pay. Hence, the choice of distinct pairs of $\Xi(t)$ and $\Theta(t)$ by the designer can be interpreted as one way of controlling the filter bandwidth. In fact, it appears that the class of applications in which increased values of $\Xi(t)$ "cured" the sensitivity of the Kalman–Bucy filter were characterized by relatively accurate measurements (low values of $\Theta(t)$).

The above discussions point out the relative effects of using *white* vs. *colored noise* in the modelling stage. If we model the plant uncertainties as *colored noise* (which may be more realistic since modelling errors are certainly not white), then we may get a better filter but at the expense of adding extra state variables in the dynamics. The issue of using colored *measurement* noise has been investigated (see, for example, [9] and [27]); its accurate modelling will certainly yield better results than its replacement with white noise. However, in the majority of applications, measurement noise is relatively white. Hence, in such applications, one would not expect too much improvement by the more accurate modelling of the measurement noise.

In short, there are no general techniques currently available that can be applied with confidence by the designer when he has to select the noise covariance matrices $\Xi(t)$ and $\Theta(t)$. Nonetheless, physical intuition, common sense, and off-line simulations represent effective tools that have been used to obtain excellent designs.

This brings us to a final *word of caution*. The *ad-hoc* techniques that have been developed for decreasing the sensitivity of Kalman filters do *not* necessarily carry over when the problem is one of stochastic control (in which the Kalman filter is a subsystem in the compensator). Many of the sensitivity problems that arise in filtering can be traced to the lack of a valid trajectory for linearization purposes. In the control problem, one does have a much more valid trajectory —$\mathbf{u}_0(t)$, $\mathbf{x}_0(t)$, $\mathbf{y}_0(t)$—on which to base the linearizations. The reason is that one would select the control to keep the system near its desired precomputed trajectory. Hence, *even if a Kalman filter is "by itself" relatively sensitive, this does not necessarily imply that, when it is used in the control problem (as part of the compensator), the closed-loop control system will be as sensitive.* Intuitively speaking, in the latter problem there are many more feedback loops that help to reduce sensitivity. Thus, the selection of the matrices Σ_0, $\Xi(t)$ and $\Theta(t)$ by the designer, should depend on whether or not the problem is that of state estimation or stochastic control. Additional discussion on this point will be presented in the sequel.

5. Stochastic Control System Design (Step 3)

5.1. *Introduction*

We have seen how the linearized Kalman–Bucy filter can be designed so as to generate in real time the estimated deviation $\delta\hat{\mathbf{x}}(t|t)$ of the actual plant state $\mathbf{x}(t)$

from its ideal deterministic response $\mathbf{x}_0(t)$. Of course $\delta\mathbf{x}(t)$ also depends on the control correction vector $\delta\mathbf{u}(t)$. Hence, one can now think of the final step of the design process as the techniques necessary for generating *on-line* the control correction vector $\delta\mathbf{u}(t)$ as a function of the measurements so as to keep $\delta\mathbf{x}(t)$ small.

The remarkable property of the "linear-quadratic-gaussian" control problem is that the optimal control correction $\delta\mathbf{u}(t)$ is generated from the estimated state deviation $\delta\hat{\mathbf{x}}(t|t)$ generated by the Kalman filter by means of the relationship

$$\delta\mathbf{u}(t) = -\mathbf{G}_0(t)\,\delta\hat{\mathbf{x}}(t|t) \tag{5.1}$$

where the gain matrix $\mathbf{G}_0(t)$ is precisely the one determined in the solution of the deterministic linear-quadratic problem (see Section 3.7 and eq. (3.40)). Recall that the deterministic solution was

$$\delta\mathbf{u}(t) = -\mathbf{G}_0(t)\,\delta\mathbf{x}(t) \tag{5.2}$$

under the assumption that the complete state perturbation vector $\delta\mathbf{x}(t)$ is measured exactly. Furthermore, recall that in the statement and solution of the filtering problem (see Section 4.4) the control correction vector $\delta\mathbf{u}(t)$ was *assumed deterministic*. Clearly, from eq. (5.1), $\delta\mathbf{u}(t)$ is *not* deterministic (since $\delta\hat{\mathbf{x}}(t)$ is a random process). Thus, it is neither apparent nor intuitively obvious why the generation of the control correction vector according to eq. (5.1) should be "optimal" since in the true stochastic problem

(a) The deterministic assumptions on $\delta\mathbf{x}(t)$ that led to the generation of $\mathbf{u}_0(t)$ are violated, and

(b) The deterministic assumptions on $\delta\mathbf{u}(t)$ that led to the generation of $\delta\hat{\mathbf{x}}(t|t)$ are also violated.

Thus, the purpose of this section is to precisely state how the overall "linear-quadratic-gaussian" problem solution *separates* into the solution of a "linear-quadratic" deterministic problem and the solution of a "linear-gaussian" estimation problem. The key theorem that shows this property is often called the *separation theorem* (see references [9], [26], and [31] to [36]).

We remark that what is referred to as the "separation theorem" in the control literature and the "certainty equivalence principle" in the economic literature (see [23], [37], [38]) are essentially the same thing; there are structural differences because in the "certainty equivalence" principle one needs the conditional mean; in the "separation theorem" the conditional mean is explicitly generated by the Kalman filter; because of this fact one can explicitly obtain many interesting properties of the overall stochastic control system.

5.2. *The Linear-Quadratic-Gaussian Problem*

We have seen in Section 4.3.5 that the (approximate) linearized relation between the actual state deviation vector $\delta\mathbf{x}(t)$ and the control correction vector $\delta\mathbf{u}(t)$ is

$$\delta\mathbf{x}(t+1) = \mathbf{A}_0(t)\,\delta\mathbf{x}(t) + \mathbf{B}_0(t)\,\delta\mathbf{u}(t) + \boldsymbol{\xi}(t) \tag{5.3}$$

while the true relation was that of eq. (4.28) which includes the effects of the higher order terms in the function $\boldsymbol{\alpha}_0(\delta\mathbf{x}(t), \delta\mathbf{u}(t), t)$.

Similarly, we have seen that the (approximate) linearized measurement relation between $\delta\mathbf{z}(t)$ and $\delta\mathbf{x}(t)$ is

$$\delta\mathbf{z}(t) = \mathbf{C}_0(t)\,\delta\mathbf{x}(t) + \boldsymbol{\theta}(t) \tag{5.4}$$

while the true relation was that of eq. (4.30) which includes the effects of the higher order terms in the function $\boldsymbol{\beta}_0(\delta\mathbf{x}(t), t)$.

We can motivate the use of quadratic criteria by mimicking the development of 3.6.2 in the deterministic case; there we remarked that use of Taylor's theorem can be used to justify the fact that the control correction vector $\delta\mathbf{u}(t)$ could be selected so as to "maximize the validity of the linearized model" by minimizing the quadratic cost given by eq. (3.37), i.e.,

$$J_0 = \delta\mathbf{x}'(T)\mathbf{Q}_0(T)\mathbf{x}(T) + \sum_{\tau=0}^{T-1} [\delta\mathbf{x}'(\tau)\mathbf{Q}_0(\tau)\,\delta\mathbf{x}(\tau) + \delta\mathbf{u}'(\tau)\mathbf{R}_0(\tau)\,\delta\mathbf{u}(\tau)]. \tag{5.5}$$

However, in our case J_0 is a scalar-valued random variable, because both $\delta\mathbf{x}(t)$ and $\delta\mathbf{u}(t)$ are random sequences.

Great care must be exercised in order to arrive at a well-formulated stochastic optimal control problem. There are two issues that demand precision

(a) Precisely what type of an expectation should be used in the cost functional?

(b) Precisely what is the admissible class of control that will be allowed in the optimization?

Such issues have often been slurred over in the literature; hence, there have been many derivations of the right result using erroneous formulations (see [34] and [39] for a critical discussion).

For the correct formulation of the cost functional to be minimized, consider the situation at any time (the present time). In addition to any *a priori* information, the following information set $S(\tau)$ is available

$$S(\tau) = \{\delta\mathbf{z}(1), \ldots, \delta\mathbf{z}(\tau), \delta\mathbf{u}(0), \ldots, \delta\mathbf{u}(\tau - 1)\}. \tag{5.6}$$

Then it makes sense to minimize the conditional expectation of the cost-to-go, denoted by

$$\mathbf{J}_0(\tau) = E\{\delta\mathbf{x}'(T)\mathbf{Q}_0(T)\,\delta\mathbf{x}(T) + \sum_{t=\tau}^{T-1} [\delta\mathbf{x}'(t)\mathbf{Q}_0(t)\,\delta\mathbf{x}(t) + \delta\mathbf{u}'(t)\mathbf{R}_0(t)\,\delta\mathbf{u}(t)] \| S(\tau)\}. \tag{5.7}$$

The way that the minimization is to be carried out is by the judicious choice of the control corrections from now on

$$\delta\mathbf{u}(\tau), \delta\mathbf{u}(\tau + 1), \ldots, \delta\mathbf{u}(T - 1).$$

However, we must be careful in communicating to the mathematics what we are allowing the control corrections to depend upon.

To obtain realizable controls, that utilize the maximum information, we can specify that the $\delta\mathbf{u}(t)$ at any time t, now and in the future, should depend on all information available up to time t, namely $S(t)$. Mathematically then we demand

$$\delta\mathbf{u}(t) = \boldsymbol{\phi}(S(t), t) \tag{5.8}$$

where $\boldsymbol{\phi}(\cdot, \cdot)$ is a deterministic map of all past measurements and controls, and perhaps of the time t. It should be noted that the structure (5.8) communicates

to the mathematics that we expect to make future observations (from time τ on) and that future controls will be functions of these measurements.

The use of the cost functional (5.6) implies also that we wish to "maximize on the average, the validity of our linearized stochastic models." Since $\boldsymbol{\xi}(t)$ and $\boldsymbol{\theta}(t)$ are white, and hence unpredictable, in individual experiments they may cause the system to deviate significantly from the region in which the linearization is more-or-less valid. Since we have no control over the specific outcome of the white noise processes, we cannot guarantee the validity of the linearization for any specific experiment. However, we can attempt to design the control system so as to optimize its average behavior.

5.2.1. *Formal statement of the linear-quadratic-gaussian stochastic control problem.* Given the linearized dynamical system (5.3) and the linearized observation equation (eq. (5.4). Given the information set $S(t)$. Find a system that generates the control correction vector $\boldsymbol{\delta u}(\tau)$, according to (5.8) such that the "average cost to go" given by (5.7) is minimum. The weighting matrices $\mathbf{Q}_0(\tau)$, $\mathbf{R}_0(\tau)$ are those defined in Section 3.7.1, while the statistical properties of the noises are those given in Section 4.3.4.

5.2.2. *The separation theorem: Solution of the linear-quadratic-gaussian stochastic control problem.* The optimal control correction vector $\boldsymbol{\delta u}(t)$ is generated by

$$\boldsymbol{\delta u}(t) = -\mathbf{G}_0(t)\,\boldsymbol{\delta}\hat{\mathbf{x}}(t|t). \tag{5.9}$$

Specification of $\mathbf{G}_0(t)$:

The control gain matrix $\mathbf{G}_0(t)$ is obtained by the solution of the deterministic linear quadratic problem (see Section 3.7.2) *forgetting completely the stochastic aspects*. Thus $\mathbf{G}_0(t)$ is given by eqs. (3.41), (3.42), and (3.43).

Specification of $\widehat{\boldsymbol{\delta x}}(t|t)$:

The vector $\widehat{\boldsymbol{\delta x}}(t|t)$ is generated by the Kalman–Bucy filter (see Section 4.4) under the assumption that $\boldsymbol{\delta u}(t)$ is deterministic and *forgetting* completely the control problem. Thus, $\widehat{\boldsymbol{\delta x}}(t|t)$ is specified by eqs. (4.44) to (4.49).

5.2.3. *The minimum value of the cost to go.* It is also possible to evaluate the minimum value, $\bar{J}_0^*(\tau)$, of the cost-to-go; the formula is

$$\begin{aligned}\bar{J}_0^*(\tau) = {} & \widehat{\boldsymbol{\delta x}}'(\tau|\tau)\mathbf{K}_0(\tau)\,\widehat{\boldsymbol{\delta x}}(\tau|\tau) \\ & + \operatorname{tr}\left[\mathbf{K}_0(\tau)\boldsymbol{\Sigma}_0(\tau|\tau)\right] + \sum_{t=\tau}^{T-1} \operatorname{tr}\left[\mathbf{K}_0(t+1)\boldsymbol{\Xi}(t)\right] \\ & + \sum_{t=\tau}^{T-1} \operatorname{tr}\left[\mathbf{K}_0(t+1)\mathbf{B}_0(t)\mathbf{G}_0(t)\boldsymbol{\Sigma}_0(t|t)\mathbf{A}_0'(t)\right]\end{aligned} \tag{5.10}$$

where (see eq. (3.41))

$$\mathbf{G}_0(t) = \left[\mathbf{B}_0'(t)\mathbf{K}_0(t+1)\mathbf{B}_0(t) + \mathbf{R}_0(t)\right]^{-1}\mathbf{B}_0'(t)\mathbf{K}_0(t+1)\mathbf{A}_0(t). \tag{5.11}$$

Equation (5.10) has important interpretations and is extremely valuable in assessing the effects of uncertainties coupled with the control doctrine. Each of

the four terms in eq. (5.10) has a special significance, and for this reason we shall discuss them separately.

1. The first term

$$\hat{\delta x}'(\tau|\tau)\mathbf{K}_0(\tau)\hat{\delta x}(\tau|\tau) \tag{5.12}$$

in (5.10) represents the effect of the current estimate $\hat{\delta x}(\tau|\tau)$ of the deviation $\delta x(\tau)$ of the actual state $x(\tau)$ from the desired one $x_0(\tau)$. Note that this term cannot be evaluated *off-line*, because $\hat{\delta x}(\tau|\tau)$ depends on the actual measurements, although the matrix $\mathbf{K}_0(\tau)$ can be computed *off-line* by eqs. (3.42) and (3.43).

2. The second term

$$\text{tr}\,[\mathbf{K}_0(\tau)\Sigma_0(\tau|\tau)] \tag{5.13}$$

in eq. (5.10) represents the increase in the cost due to the *current* uncertainty in $\delta x(\tau)$; recall that under the linearity and Gaussian assumptions $\Sigma_0(\tau|\tau)$ is the conditional covariance matrix of $\delta x(\tau)$. This term can be computed *off-line* since both $\mathbf{K}_0(\tau)$ and $\Sigma_0(\tau|\tau)$ can be calculated before the actual system actually is placed in operation. Note that this term couples the effects of the control cost functional—via $\mathbf{K}_0(\tau)$—and the current accuracy of estimation—via $\Sigma_0(\tau|\tau)$.

3. The third term

$$\sum_{t=\tau}^{T-1} \text{tr}\,[\mathbf{K}_0(t+1)\Xi(t)] \tag{5.14}$$

summarizes the contributions of the future plant white noise sequence; the more noisy the system, the larger the covariance matrix $\Xi(t)$ and the larger the stochastic cost. Once more this term can be calculated *off-line* since it does not depend upon the actual measurements.

4. The last term

$$\sum_{t=\tau}^{T-1} \text{tr}\,[\mathbf{K}_0(t+1)\mathbf{B}_0(t)\mathbf{G}_0(t)\Sigma_0(t|t)\mathbf{A}_0'(t)] \tag{5.15}$$

in eqs. (5.10) summarizes the contributions of future uncertainties in the estimate of $\delta x(t)$, which are reflected by the values of the covariance matrix $\Sigma_0(t|t)$. Once more this quantity can be computed off-line. Note that the effects of future measurement accuracy are reflected in this term since $\Sigma_0(t|t)$, $t = \tau, \tau+1, \ldots, T-1$ depends on the measurement noise covariance matrix $\Theta(t)$ (see eqs. (4.47) and (4.48)). Once more the future accuracy of estimation, dictated by the Kalman filter and quantified by $\Sigma_0(t|t)$, couples to the control objectives, which are quantified by the values of $\mathbf{K}_0(t+1)$ and $\mathbf{G}_0(t)$.

5.3. *The Special Case of Noiseless Measurements*

In many economic problems the assumption is made that all state variables can be measured exactly; in our terminology this means that $\delta x(t)$ is known exactly. Under this assumption the "conditional mean" is the measurement itself, i.e.,

$$\hat{\delta x}(t|t) = \delta x(t) \tag{5.16}$$

and the conditional covariance is zero, i.e.,

$$\Sigma_0(t|t) = 0. \tag{5.17}$$

For such problems, one of course does not need the Kalman filter, even if the system is still stochastic ($\Xi(t) \neq 0$). The optimal control is still generated by (5.9) with the constraint (5.16).

The noiseless measurements naturally decrease the minimum value of the cost. In view of (5.16) we see that the second and fourth terms in eq. (5.10) vanish; the stochastic aspects of the problem are reflected in the third term (5.14) of the cost function. For additional remarks on this problem see ref [40].

5.4. *Methods of Proof*

The most fail-safe method of deriving the separation theorem is via dynamic programming; references [9], [18] to [20], [28], [32] to [40] contain such derivations although the methods of proof can be quite different, and the level of rigour quite variable. It is worthwhile to note that [32] employs a clever transformation of the noisy measurement problem to an equivalent noiseless measurement problem, using the whiteness of the residual or innovations sequence. The appropriate way of using deterministic optimization techniques, i.e., the discrete minimum principle, to solve this stochastic optimal control problem can be found in [39] and [41].

5.5. *Discussion*

We shall now make some brief remarks regarding the interpretation that should be attached to the formal solution to the linear-quadratic-gaussian problem. First, we shall discuss how trade-off studies regarding the accuracy as well as the type of sensors and actuators to be used affect the solution to the control problem as a whole. Let us suppose that the weighting matrices $\mathbf{Q}_0(t)$, and $\mathbf{R}_0(t)$ have been somehow selected. In this context, the solution $\mathbf{K}_0(t)$ of the control equation (3.42) is available.

5.5.1. *Sensor selection.* Let us suppose that we are faced with the problem of selecting between two types of measurement devices which, except for accuracy, perform otherwise the same tasks. Suppose that the more accurate sensor(s) is characterized by a measurement noise covariance matrix $\Theta_1(t)$ while the less accurate by $\Theta_2(t)$, such that $\Theta_1(t) \leq \Theta_2(t)$. On the other hand, the more accurate sensor(s) cost more money. For each sensor, we can solve the filter problem equations (4.47)–(4.48) and obtain the corresponding error covariance matrices, say $\Sigma_1(t|t)$ and $\Sigma_2(t|t)$; it turns out that $\Sigma_1(t) \leq \Sigma_2(t)$, i.e., use of the more accurate sensor improves the estimation accuracy. In fact, from the filtering viewpoint the increase in state estimation accuracy may justify the increase in monetary cost. However, it does *not* necessarily follow that the expected improvement in the control system performance will necessarily justify the monetary cost. The reason is that only the last term (5.15) in the minimum cost functional (5.10) will decrease (since $\Sigma_1(t|t) \leq \Sigma_2(t|t)$) from the use of the more accurate sensor. However, the relative percentage in increased performance is also governed by the first three terms in (5.10). It may turn out that for a doubling of invested money we may

double estimation accuracy, but only buy a few percent in improving the control system performance as measured by (5.10). In particular, if we assume that at the initial time the state deviations are small, then we can carry out off-line (non-Monte-Carlo) studies by examining essentially percentage changes in the last three terms of (5.10).

Similar remarks can be made regarding the selection of the number of sensors. In our context, this would change the $\mathbf{C}_0(t)$ matrix (both dimensionwise and numerically) in eq. (4.48); this in turn will change the value of the $\mathbf{\Sigma}_0(t|t)$ matrix that affects only the last term in eq. (5.10).

5.5.2. *Actuator and model accuracy tradeoffs.* In a similar vein we can carry out tradeoff studies which involve the selection of the plant noise covariance matrix $\mathbf{\Xi}(t)$. As we have remarked before, this models the total uncertainty in the dynamics (due to actuator errors as well as modelling errors). Let us suppose that we can "buy" two sets of actuators characterized by $\mathbf{\Xi}_1(t) \leq \mathbf{\Xi}_2(t)$ so that the first are more accurate (and more costly) than the second. Once more, from eqs. (4.47)–(4.48) we can deduce that $\mathbf{\Sigma}_1(t|t) \leq \mathbf{\Sigma}_2(t|t)$; i.e., more money busy increased state estimation accuracy (one can make a similar argument that more accurate modelling requires more engineering and experimentation time). As far as the effects of less plant uncertainty on the control system performance is concerned, different values of $\mathbf{\Xi}(t)$ affect the last *two* terms in the cost (5.10)—directly in the third term (5.14) and indirectly via $\mathbf{\Sigma}_0(t|t)$ in the fourth term, (5.15).

Even more interesting (off-line and non-Monte-Carlo!) tradeoffs can be carried out in the wisest allocation of funds partly to buy some better sensors, partly to buy some better actuators, and partly to invest in additional engineering time for better modelling.

5.6. *On The Selection of Weighting and Covariance Matrices*

We conclude this section with some remarks pertaining to the selection of the control weighting matrices and the noise covariance matrices.

The fact that from *a mathematical point of view* the separation theorem allows us to solve the control and filter problems separately, does not imply that these two problems should be solved separately by two distinct design groups and "hooked together" by the supervisor. Unfortunately, this is how the theorems have been used in many engineering designs leading to unsuccessful results. For this reason, we shall briefly elaborate on the proper usage of this theorem.

In general, if we could solve the overall nonlinear nonquadratic stochastic control problem, the optimal design would *not* obey this nice separation property. Since we cannot solve the true problem, we employ the linear-quadratic-gaussian approach to arrive at a set of problems that we can solve. The key question is then: *What is available to the designer to control the goodness of the overall design so that he can obtain a satisfactory system?* The answer to this question is the selection of the four matrix sequences $\{\mathbf{Q}_0(t)\}$, $\{\mathbf{R}_0(t)\}$, $\{\mathbf{\Xi}(t)\}$ and $\{\mathbf{\Theta}(t)\}$. For any arbitrary selection, the mathematical problem separates. However, this *does not mean that* $\mathbf{Q}_0(t)$, $\mathbf{R}_0(t)$, and $\mathbf{\Xi}(t)$, $\mathbf{\Theta}(t)$ should be selected independently of each other. The discussion of section 5.3 points this out to a certain degree since the changes in

uncertainty by changes in $\Xi(t)$ and $\Theta(t)$ are modulated by the values of $\mathbf{R}_0(t)$ and $\mathbf{Q}_0(t)$ (via $\mathbf{K}_0(t)$) in the cost (5.10).

Unfortunately, there seems to be no published literature on the above point. We have already commented that $\mathbf{Q}_0(t)$ and $\mathbf{R}_0(t)$ can be used in the quadratic criterion to maximize linearization validity (see section 3.6). We have also commented that $\Xi(t)$ and $\Theta(t)$ can be used to communicate to the mathematics the existence of modelling errors due to linearization (see section 4.3.5). Clearly both sets of matrices are partly used for the same purpose: hence, their selection should not be done independently. Unfortunately, as we have stressed throughout this paper, there are no systematic procedures available for the specification of $\mathbf{Q}_0(t)$, $\mathbf{R}_0(t)$, $\Xi(t)$, and $\Theta(t)$. Additional theoretical research and applications studies are necessary.

6. Summary of the L-Q-G Approach to Design

We have outlined the philosophy, assumptions, formulation, and mathematical characterization of a design process for controlling a nonlinear uncertain system about a desired trajectory, through the use of the so-called linear-quadratic-gaussian problem. This design process represents a relatively well understood by-product of modern control theory. Of course, successful control systems have been designed using alternate approaches. However, this design process is characterized by a clear cut division of responsibilities between the modelling and the calculation aspects of the problem.

We outline below, for the sake of completeness, the thirteen key steps in the design process. All of these are carried out off-line; the on-line computational requirements are minimal.

PART A: DETERMINISTIC MODELLING

Step 1: Determine a deterministic model of the plant; this yields

$$\mathbf{x}(t+1) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t).$$

Step 2: Determine a deterministic model of the sensors; this yields

$$\mathbf{y}(t) = \mathbf{g}(\mathbf{x}(t), t).$$

Step 3: Determine ideal input-state-output sequences (perhaps using the discrete minimum principle)

$$\{\mathbf{u}_0(t)\}: \text{ideal input sequence}$$

$$\{\mathbf{x}_0(t)\}: \text{ideal state sequence}$$

$$\{\mathbf{y}_0(t)\}: \text{ideal output sequence}$$

for all $t = 0, 1, 2, \ldots, T$.

PART B: STOCHASTIC MODELLING

Step 4: Model uncertainty in initial plant state

$$\text{Select mean: } \mathbf{x}_0 = E(\mathbf{x}(0)\}.$$

$$\text{Select covariance: } \Sigma_0 = \text{cov}\,[\mathbf{x}(0); \mathbf{x}(0)].$$

Step 5: Model input uncertainty

Select covariance: $\boldsymbol{\Xi}(t)\ \ \delta_{tt} = \text{cov}\,[\boldsymbol{\xi}(t); \boldsymbol{\xi}(\tau)]$.

Step 6: Model sensor uncertainty

Select covariance $\boldsymbol{\Theta}(t)\,\delta_{tt} = \text{cov}\,[\boldsymbol{\theta}(t); \boldsymbol{\theta}(\tau)]$.

PART C: LINEARIZATION

Step 7: Establish matrices $\mathbf{A}_0(t)$, $\mathbf{B}_0(t)$, and $\mathbf{C}_0(t)$ from information in Steps 1, 2, 3

$$\mathbf{A}_0(t) = \left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}(t)}\right|_0, \qquad \mathbf{B}_0(t) = \left.\frac{\partial \mathbf{f}}{\partial \mathbf{u}(t)}\right|_0, \qquad \mathbf{C}_0(t) = \left.\frac{\partial \mathbf{g}}{\partial \mathbf{x}(t)}\right|_0.$$

Step 8: Depending on "degree of nonlinearity" select weighting matrices $\mathbf{Q}_0(t)$, $\mathbf{R}_0(t)$ with due consideration of the values of $\boldsymbol{\Sigma}_0$, $\boldsymbol{\Xi}(t)$, and $\boldsymbol{\Theta}(t)$.

PART D: CONTROL PROBLEM CALCULATIONS (OFF-LINE)

Step 9: Using the weighting matrices $\mathbf{Q}_0(t)$, $\mathbf{R}_0(t)$ of Step 8 and the matrices $\mathbf{A}_0(t)$, $\mathbf{B}_0(t)$ of Step 7 solve backward in time the matrix difference equation

$$\mathbf{K}_0(t) = \mathbf{Q}_0(t) + \mathbf{A}_0'(t)\mathbf{K}_0(t+1)\mathbf{A}_0(t) - \mathbf{A}_0'(t)\mathbf{K}_0(t+1)\mathbf{B}_0(t)$$
$$\times [\mathbf{B}_0'(t)\mathbf{K}_0(t+1)\mathbf{B}_0(t) + \mathbf{R}_0(t)]^{-1}\mathbf{B}_0'(t)\mathbf{K}_0(t+1)\mathbf{A}_0(t)$$

with $\mathbf{K}_0(T) = \mathbf{Q}_0(T)$.

Step 10: Compute the control gain matrix $\mathbf{G}_0(t)$

$$\mathbf{G}_0(t) = [\mathbf{B}_0'(t)\mathbf{K}_0(t+1)\mathbf{B}_0(t) + \mathbf{R}_0(t)]^{-1}\mathbf{B}_0'(t)\mathbf{K}_0(t+1)\mathbf{A}_0(t).$$

PART E: FILTER PROBLEM CALCULATIONS (OFF-LINE)

Step 11: Using the covariance matrices $\boldsymbol{\Sigma}_0$, $\boldsymbol{\Xi}(t)$, and $\boldsymbol{\Theta}(t)$, established in Steps 4–6, and the matrices $\mathbf{A}_0(t)$, $\mathbf{C}_0(t)$ of Step 7 solve forward in time the matrix difference equations

$$\boldsymbol{\Sigma}_0(t+1|t) = \mathbf{A}_0(t)\boldsymbol{\Sigma}_0(t|t)\mathbf{A}_0'(t) + \boldsymbol{\Xi}(t)$$

$$\boldsymbol{\Sigma}_0(t+1|t+1) = \boldsymbol{\Sigma}_0(t+1|t) + \boldsymbol{\Sigma}_0(t+1|t+1)\mathbf{C}_0'(t+1)$$
$$\times [\mathbf{C}_0(t+1)\boldsymbol{\Sigma}_0(t+1|t)\mathbf{C}_0'(t+1) + \boldsymbol{\Theta}(t+1)]^{-1}\mathbf{C}_0(t+1)\boldsymbol{\Sigma}_0(t+1|t)$$

with $\boldsymbol{\Sigma}_0(0|0) = \boldsymbol{\Sigma}_0$.

Step 12: Compute the filter gain matrix

$$\mathbf{H}_0(t+1) = \boldsymbol{\Sigma}_0(t+1|t+1)\mathbf{C}_0'(t+1)\boldsymbol{\Theta}^{-1}(t+1).$$

PART F: ON-LINE CALCULATIONS

From the actual measurements $\mathbf{z}(1), \mathbf{z}(2), \ldots$

(a) Compute $\delta\mathbf{z}(1), \delta\mathbf{z}(2), \ldots,$ by:

$$\delta\mathbf{z}(t) = \mathbf{z}(t) - \mathbf{g}(\mathbf{x}_0(t), t).$$

(b) Compute estimated deviations $\delta\hat{\mathbf{x}}(t|t)$ and control correction $\delta\mathbf{u}(t)$ by:

$$\delta\hat{\mathbf{x}}(t+1|t) = \mathbf{A}_0(t)\,\delta\hat{\mathbf{x}}(t|t) + \mathbf{B}_0(t)\,\delta\mathbf{u}(t)$$

$$\delta\mathbf{r}(t+1) = \delta\mathbf{z}(t+1) - \mathbf{C}_0(t+1)\,\delta\hat{\mathbf{x}}(t+1|t)$$

$$\delta\hat{\mathbf{x}}(t+1|t+1) = \delta\hat{\mathbf{x}}(t+1|t) + \mathbf{H}_0(t+1)\,\delta\mathbf{r}(t+1)$$

$$\delta\mathbf{u}(t) = -\mathbf{G}_0(t)\,\delta\hat{\mathbf{x}}(t|t)$$

$$\delta\hat{\mathbf{x}}(0|0) = E\{\mathbf{x}(0)\} - \mathbf{x}_0(0).$$

(c) Compute actual control $\mathbf{u}(t)$ by:

$$\mathbf{u}(t) = \mathbf{u}_0(t) + \delta\mathbf{u}(t).$$

The step-by-step development should convince the reader of the crucial importance of the modelling issue in this design process. The ability of the designer to translate physical quantities into their mathematical counterparts in Steps 1, 2, 3, 4, 5, 6 and 8 is absolutely essential. Once the modelling has been carried out, the remaining steps 7, 9, 10, 11, 12, and 13 are purely mechanical.

7. Trends in Adaptive Control

As remarked before, the use of the linear-quadratic-Gaussian problem for feedback control system design for nonlinear stochastic systems, represents the *simplest* possible approach. It hinges on the assumption that the feedback control system can do a good job of returning the state of the system to its desired nominal trajectory. From an engineering viewpoint, the major advantage of this approach is that most of the complex calculations can be carried out prior to the actual system operation, and the on-line calculations are minimal.

One of the most important problems in engineering, economics, management science, and urban systems arises when certain of the parameters of the dynamical system are not known exactly. These parameters may be constant or time-varying (in a deterministic or stochastic manner). Thus, such systems may be described by the difference equation

$$\mathbf{x}(t+1) = \mathbf{f}(\mathbf{x}(t), \mathbf{p}(t), \mathbf{u}(t), t) + \xi(t) \tag{7.1}$$

where $\mathbf{x}(t)$ is the state, $\mathbf{u}(t)$ is the control, $\mathbf{p}(t)$ is the parameter vector and $\xi(t)$ is the additive plant white noise.

If the parameters are constant, then

$$\mathbf{p}(t+1) = \mathbf{p}(t)$$

and $\mathbf{p}(t)$ can be viewed as a random vector. At the opposite extreme, the parameters may change in a stochastic manner

$$\mathbf{p}(t+1) = \mathbf{h}(\mathbf{p}(t), t) + \gamma(t) \tag{7.2}$$

where $\mathbf{h}(\cdot,\cdot)$ summarizes the deterministic variation, and $\gamma(t)$ is another white noise.

As we have remarked before, the use of feedback does tend to compensate to a certain extent the effects of parameter variations. Thus, if we have a nominal

sequence of parameters, denoted by $\mathbf{p}_0(t)$, then we can use the suggested approach with some degree of confidence *provided that we are sure from the start that the parameter variation vector*

$$\delta\mathbf{p}(t) \triangleq \mathbf{p}(t) - \mathbf{p}_0(t) \tag{7.3}$$

is "small."

We emphasize that there is a distinct difference between having $\delta\mathbf{x}(t)$ "small," and $\delta\mathbf{p}(t)$ "small." In the case of the state, the control can influence $\mathbf{x}(t)$, and hence $\delta\mathbf{x}(t)$; it was this fact that was crucial in the use of linearization. On the other hand, as it can be evidenced by eqs. (7.2) and (7.3), *the control* $\mathbf{u}(t)$ *cannot influence the dynamic evolution of the parameter vector* $\mathbf{p}(t)$; *thus, there is no way of using the control to maximize the honesty of any linearizations about the nominal parameter sequence* $\mathbf{p}_0(t)$.

For this reason, if one suspects that $\delta\mathbf{p}(t)$ can become large, then the approach described in this paper may not lead to a satisfactory design.

There are techniques available that can incorporate the effects of unknown parameters. It should be stressed right at the start that these techniques require much more on-line computation than the procedure described. For many engineering applications, the extra on-line computational requirements makes such techniques impractical. However, there are some engineering applications, and certainly economic applications, which are characterized by sufficient time between measurements and decisions so that the increased on-line computations become feasible.

In the remainder of this section we shall present a very brief review of the engineering literature on this problem. We remark that this is called the *adaptive control problem*; the references [42]–[58] contain a (non-exhaustive) sample of pertinent works.

It should be self-evident that since the actual parameter sequence $\mathbf{p}(t)$ may differ from the *a priori* nominal parameter sequence $\mathbf{p}_0(t)$, then on the basis of measurements (noisy or not) one needs to construct a parameter estimate sequence

$$\hat{\mathbf{p}}(t|t) \tag{7.4}$$

in addition to a state estimate sequence

$$\hat{\mathbf{x}}(t|t). \tag{7.5}$$

Depending on the structure of the equations *this estimation problem* (even for linear systems!) *may be a nonlinear one.* In this case, one employs an extended Kalman filter or more complex estimation algorithm (see reference [25]) to generate the estimates $\hat{\mathbf{p}}(t|t)$ and $\hat{\mathbf{x}}(t|t)$. It is beyond the scope of this paper to delve in detail on the detailed structure of these nonlinear estimation algorithms. However, it is worthwhile to remark that the propagation of the covariance matrices essentially the equivalent to (4.47) and (4.48) cannot be done off-line, because at each step one has to do the linearizations about the current estimate.

The interesting aspects of the adaptive control problem pertain to the role of the control. In essence, there is a *dual* nature (see reference [42]) to the control; one for ordinary control, and another for helping us estimate more accurately the parameters. Thus, *although the control cannot be used to change the time-evolution*

of the parameters, it can control the conditional covariance matrix of $\mathbf{p}(t)$, *which reflects the accuracy of parameter estimation.*

Progress in this area, in the engineering literature, has been almost totally devoted to the linear case, i.e., where the state dynamics have the form

$$\mathbf{x}(t+1) = \mathbf{A}(\mathbf{p}(t), t)\mathbf{x}(t) + \mathbf{B}(\mathbf{p}(t), t)\mathbf{u}(t) + \boldsymbol{\xi}(t) \tag{7.6}$$

the parameter dynamics have the form

$$\mathbf{p}(t+1) = \mathbf{H}(t)\mathbf{p}(t) + \boldsymbol{\gamma}(t) \tag{7.7}$$

and the measurement equations have the form

$$\mathbf{z}(t) = \mathbf{C}(\mathbf{p}(t), t)\mathbf{x}(t) + \boldsymbol{\theta}(t) \tag{7.8}$$

where the matrices $\mathbf{A}(\cdot)$, $\mathbf{B}(\cdot)$, $\mathbf{C}(\cdot)$ depend upon the parameter vector, and $\boldsymbol{\xi}(t)$, $\boldsymbol{\gamma}(t)$, and $\boldsymbol{\theta}(t)$ are white noise sequences.

The criterion used is quadratic. Thus, if $\mathbf{p}(t)$ were known exactly for all t the separation theorem would yield the optimal stochastic control.

One technique that can be used given $\hat{\mathbf{x}}(t|t)$ and $\hat{\mathbf{p}}(t|t)$ is that of *enforced separation.* In this case one predicts the future values of the parameter vector from (7.7) to obtain

$$\hat{\mathbf{p}}(\tau+1|t) = \mathbf{H}(\tau)\hat{\mathbf{p}}(\tau|t); \qquad \tau = t, t+1, \ldots, T. \tag{7.9}$$

This predicted sequence of parameter vectors fixes the future values of $\mathbf{A}(\cdot)$ and $\mathbf{B}(\cdot)$ and one can now calculate the control gain (by an equation like (3.41)), by solving *on-line* a difference equation like (3.42)–(3.43). Note that this process must be repeated when the next observation is made, since, in general,

$$\hat{\mathbf{p}}(t+1|t+1) \neq \hat{\mathbf{p}}(t|t). \tag{7.10}$$

Thus, in the *enforced separation adaptive control scheme*, for linear-quadratic-Gaussian problems, the structure of the separation theorem is preserved; however, the corresponding covariance matrices and control gains must be computed on-line each and every time a new measurement has been made.

Another technique (that does *not* lead to a separation type result) is that of open-loop feedback optimal (OLFO) control (see [45]). This technique once more requires that $\hat{\mathbf{x}}(t|t)$ and $\hat{\mathbf{p}}(t|t)$ be generated. Next, the assumption is made that no more measurements will be made; hence the deterministic open-loop optimal control can be found that minimizes the conditional expectation of the quadratic cost functional. This open-loop control is applied *only* at the current time period; when a new measurement is made, one repeats the entire set of computation (see references [52] to [56] for details). The OLFO approach yields a design that has the properties

(a) the control gain matrix not only depends upon the predicted average values of the unknon parameters, but also upon the predicted parameter covariance matrix

(b) the computations are far more complex than in the *enforced separation* scheme.

In the economic literature it appears that reference [58] is essentially an OLFO approach.

A somewhat different approach to adaptive control, yet consistent with an OLFO viewpoint, is to assume that the actual parameters of the system are one out of N possible sets. Each one of the N possibilities forms a hypothesis; N Kalman filters in parallel generate the state estimate under each hypothesis. One can then construct, from the residuals and state estimates of the Kalman filters, in real time the conditional probability that each hypothesis is correct, and these can be used to construct the adaptive control (see reference [57] for details).

There is little doubt that in the next five years the adaptive control problem will receive more and more attention in the stochastic control literature. At the present time there are several approaches, but with little cross-evaluation. A unified treatment of this important class of problems is still lacking.

Department of Electrical Engineering
Massachusetts Institute of Technology

References

[1] M. Athans, "The Role and Use of the Linear-Quadratic-Gaussian Problem in Control System Design," *IEEE Transactions on Automatic Control*, Vol. AC-16, No. 6, December, 1971, pp. 529–552.
[2] A. Zellner, "A Statistical Analysis of Provisional Estimates of Gross National Product and Its Components, of Selected National Income Components, and of Personal Saving," *Journal of American Statistical Association*, LIII, No. 281, March, 1958, pp. 54–65.
[3] O. Morgenstern, "On the Accuracy of Economic Observations," Princeton University Press, Princeton, N.J., 1963.
[4] R. Bellman, "Dynamic Programming," Princeton University Press, N.J., 1956.
[5] L. S. Pontryagin et al., "The Mathematical Theory of Optimal Processes," J. Wiley and Sons, Interscience, N.Y., 1962.
[6] T. M. Holtzman and H. Halkin, "Directional Convexity and the Maximum Principle for Discrete Systems," *Journal SIAM on Control*, Vol. 4, 1966, pp. 263–275.
[7] H. Halkin, "Optimal Control for Systems Described by Difference Equations," in *Advances in Control Systems* (C. T. Leondes, ed.), Academic Press, N.Y., 1964.
[8] M. D. Cannon, C. D. Cullum, and E. Polak, "Theory of Optimal Control and Mathematical Programming," McGraw-Hill Book Co., N.Y., 1970.
[9] A. E. Bryson, Jr. and Y-C Ho, "Applied Optimal Control," Blaisdell Publishing Company, Waltham, Mass., 1969.
[10] D. H. Jacobson and D. Q. Mayne, "Differential Dynamic Programming," Elsevier Publishing Co., N.Y., 1970.
[11] A. V. Balakrishnan and L. W. Neustadt (eds.), "Computing Methods in Optimization Problems," Academic Press, Inc., N.Y., 1964.
[12] A. P. Sage, "Optimum Systems Control," Prentice-Hall, Inc., Englewood Hills, N.J., 1968.
[13] A. H. Levis, R. A. Schlueter, and M. Athans, "On the Behavior of Optimal Sampled Data Regulators," *International Journal on Control*, Vol. 13, No. 2, February, 1971, pp. 343–361.
[14] P. Dorato and A. H. Levis, "Optimal Linear Regulators: the Discrete-Time Case," *IEEE Trans. Auto. Control*, Vol. AC-16, No. 6, December, 1971, pp. 613–620.
[15] D. L. Kleinman and M. Athans, "The Discrete Minimum Principle with Application to the Linear Regulator Problem," Report ESL-R-260, Electronics Systems Laboratory, MIT, Cambridge, Mass., 1966.
[16] R. E. Kalman and R. W. Koepcke, "Optimal Synthesis of Linear Damping Control Systems Using Generalized Performance Indices," *Trans. ASME*, Vol. 80, 1958, pp. 1800–1826.
[17] R. E. Larson, "State Increment Dynamic Programming," Elsevier, N.Y., 1968.
[18] R. C. K. Lee, "Optimal Estimation, Identification, and Control, MIT Press, Cambridge, Mass., 1964.
[19] J. S. Tou, "Optimum Design of Digital Systems," Academic Press, N.Y., 1963.
[20] ———, "Modern Control Theory," McGraw-Hill Book Co., N.Y., 1964.

[21] R. S. Pindyck, "An Application of the Linear Quadratic Tracking Problem to Economic Stabilization Policy," *IEEE Trans. Auto. Control*, Vol. AC-17, No. 3, June, 1972, pp. 287–300.
[22] ———, "Optimal Stabilization Policies via Deterministic Control," this issue, pp. 385–390.
[23] C. C. Holt, F. Modigliani, S. F. Muth, and H. A. Simon, "Planning Production, Inventories, and Work Force," Prentice-Hall, Inc., Englewood Hills, N.J., 1960.
[24] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," Trans. ASME, *Journal of Basic Engineering*, Vol. 82, March, 1960, pp. 34–35.
[25] A. H. Jazwinski, "Stochastic Processes and Filtering Theory," Academic Press, N.Y., 1970.
[26] K. J. Astrom, "Introduction to Stochastic Control Theory", Academic Press, N.Y., 1970.
[27] A. P. Sage and J. L. Melsa, "Estimation Theory with Applications to Communications and Control," McGraw-Hill Book Co., 1971.
[28] J. S. Meditch, "Stochastic Optimal Linear Estimation and Control," McGraw-Hill Book Co., 1969.
[29] I. B. Rhodes, "A Tutorial Introduction to Estimation and Filtering," *IEEE Trans. Auto. Control*, Vol. AC-16, No. 6, December, 1971, pp. 688–706.
[30] R. K. Mehra, "On the Identification of Variances and Adaptive Kalman Filtering," *IEEE Trans. Auto. Control*, Vol. AC-15, No. 2, April, 1970, pp. 175–184.
[31] H. J. Kushner, "Introduction to Stochastic Control," Holt, Rinehart, and Winston, N.Y., 1971.
[32] R. E. Meier, R. E. Larson, and A. J. Tether, "Dynamic Programming for Stochastic Control of Discrete Systems," *IEEE Trans. Auto. Control*, Vol. AC-16, No. 6, December, 1971, pp. 767–776.
[33] C. Striebel, "Sufficient Statistics in the Optimum Control of Stochastic Systems," *Journal Mathematical Analysis and Applications*, Vol. 12, 1965, pp. 576–592.
[34] H. S. Witsenhausen, "Separation of Estimation and Control for Discrete Time Systems," *Proc. IEEE*, December, 1971.
[35] T. L. Gunckel and G. F. Franklin, "A General Solution for Linear Sampled Data Control," Trans. ASME, *Journal of Basic Engineering*, Vol. 85 D, June, 1963, pp. 197–201.
[36] P. D. Joseph and J. T. Tou, "On Linear Control Theory," *AIEE Trans.* (Applications and Industry), Vol. 80, September, 1961, pp. 193–196.
[37] H. Theil, "A Note on Certainty Equivalence in Dynamic Planning," *Econometrica*, Vol. 25, 1957, pp. 346–349.
[38] H. A. Simon, "Dynamic Programming under Uncertainty with a Quadratic Criterion Function," *Econometrica*, Vol. 24, 1956, pp. 74–81.
[39] L. C. Kramer, "On Stochastic Control and Optimal Measurement Strategies," Report ESL-R-462, Electronic Systems Laboratory, MIT, Cambridge, Mass., November, 1971.
[40] G. C. Chow, "How Much Could Be Gained by Optimal Stochastic Control Policies," this issue, pp. 391–406.
[41] L. C. Kramer and M. Athans, "On the Application of Deterministic Optimization Methods to Stochastic Control Problems," *Proc. 1972 Joint Automatic Control Conference*, Stanford, Calif., August, 1972.
[42] A. A. Feldbaum, "Dual Control Theory," *Automation and Remote Control*, Part I, Vol. 21, pp. 874–880, 1961; Part II, Vol. 21, pp. 1033–1046, 1961; Part III, Vol. 22, pp. 1–12, 1962; Part IV, Vol. 22, pp. 109–121, 1962.
[43] H. A. Spang, "Optimum Control of an Unknown Linear Plant Using Bayesian Estimation of the Error, *IEEE Trans. Auto. Control*, Vol. AC-10, 1965, pp. 80–93.
[44] Y. Bar-Shalom and R. Sivan, "The Optimal Control of Discrete Time Systems with Random Parameters," *IEEE Trans. Auto. Control*, Vol. AC-14, 1969, pp. 3–8.
[45] S. E. Dreyfus, "Some Types of Optimal Control of Stochastic Systems," *SIAM Journal of Control*, Vol. 2, 1964, pp. 120–134.
[46] J. J. Florentin, "Optimal Probing, Adaptive Control of a Simple Bayesian System," *Journal of Electronics and Control*, Vol. 2, 1962.
[47] S. Farrison, R. E. Graham, and R. C. Skelton, "Identification and Control of Linear Discrete Systems," *IEEE Trans. Auto. Control*, Vol. AC-12, 1967, pp. 438–442.
[48] W. S. Murphy, "Optimal Stochastic Control of Discrete Linear Systems with Unknown Gain," *IEEE Trans. Auto. Control*, Vol. AC-13, 1968, pp. 338–344.
[49] G. Stein and G. N. Saridis, "A Parameter Adaptive Control Technique," *Automatica*, Vol. 5, 1969.
[50] G. N. Saridis and R. N. Lobbia, "Parameter Identification and Control of Linear Discrete Time Systems," *Proceedings 1971 Joint Automatic Control Conference*, St. Louis, Mo., pp. 724–730.
[51] G. T. Schmidt, "A New Technique for Identification and Control of Systems with Unknown Parameters," Ph.D. Thesis, MIT, Cambridge, Mass., October, 1970.
[52] E. Tse, "On the Optimal Control of Linear Systems with Incomplete Information," Report ESL-R-412, MIT Electronic Systems Laboratory, Cambridge, Mass., January, 1970.

[53] E. Tse and M. Athans, "Adaptive Stochastic Control for Linear Systems," Parts I and II, *Proceedings IEEE Decision and Control Conference*, Austin, Texas, December, 1970.
[54] ———, "Adaptive Stochastic Control for a Class of Linear Systems," *IEEE Trans. Auto. Control*, Vol. AC-17, February, 1972, pp. 38–52.
[55] R. Ku, "Adaptive Control of Stochastic Linear Systems with Unknown Parameters," Report ESL-R-477, MIT Electronic Systems Laboratory, Cambridge, Mass., May, 1972.
[56] R. Ku and M. Athans, "On the Adaptive Control of Linear Systems Using the Open Loop Feedback Optimal Approach," IEEE Decision and Control Conference, New Orleans, La., December, 1972.
[57] D. G. Lainiotis, J. G. Deshpande, and T. N. Upadhyay, "Optimal Adaptive Control: A Nonlinear Separation Theorem," *International Journal on Control*, Vol. 15, 1972, pp. 877–888.
[58] E. Chase-McRae, "Linear Decision with Experimentation," this issue, pp. 437–447.