Optimization

Edited by
R. S. Anderssen,
L. S. Jennings,
and D. M. Ryan

University of Queensland Press

TABLE OF CONTENTS

PROBLEMS OF OPTIMAL CONTROL THEORY:AN INTRODUCTION

B.D.O. Anderson

Electrical Engineering Department, University of Newcastle

1. INTRODUCTION

The aim of this paper is to introduce the reader to optimal control. We attempt first to answer the following questions:

(1) What is the context in which optimal control problems arise?

(2) What constitutes a formal statement of an optimal control problem?

Having posed the problem, we are then able to highlight some of the differences between the optimal control problem - essentially one of *dynamic optimization* - with the basic problems of nonlinear programming - which essentially deals with *static optimization*.

The practical implementation of optimal controls, which are computed when the optimal control problem is solved, is then discussed; the key point here is to distinguish the ideas of open-loop and closed-loop control.

Next, we pose optimal control problems in a discrete time framework; this is often done in practice, and results in a problem more closely tied to the problems of mathematical programming than when the optimal control problem is posed in a continuous time framework. Frequently though, the normal solution procedures of mathematical programming will be inadequate to solve the discrete time optimal control problem, and special procedures need to be used.

Several major theoretical results are then alluded to, and the implications of using these in the solution of optimal control problems is covered. Finally, we offer some further remarks on the engineering significance of optimal control.

A short reference list is provided. The only references given are books, and several words are said about each reference to enable the reader from outside the field to choose the book that best suits his requirements.

2. QUALITATIVE STATEMENT OF THE CONTROL PROBLEM

Optimal control problems are associated with a *plant*, possessing *controls* or input, and *states*. A plant is a physical object, and almost always in the context of optimal control it is a dynamical system, and these will be the only plants of concern to us here. Examples of such plants include a spinning top, a double pendulum, and electric circuit, an electromechanical servo system, and an inertial load in an environment where there is no damping. The behaviour of all these plants may be described by differential equations, thus

$$m\,\ddot{y} = F \tag{2.1}$$

for the last example above. The input for the plant in this case is F, the force applied to the inertial load, and the state is $[y, \dot{y}]$, or the position and velocity of the load.

For convenience in the differential equation description of a plant, it is common to adopt a *normal form* set of differential equations, i.e. one of the form

$$\dot{\underset{\sim}{x}} = \underset{\sim}{f}(\underset{\sim}{x},\underset{\sim}{u},t) \tag{2.2}$$

For example, (2.1) can be rewritten as

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u \tag{2.3}$$

where $x_1 \equiv y$, $x_2 \equiv \dot{y}$ and $u \equiv F/m$. The variable $\underset{\sim}{u}(.)$ in (2.2) is termed *the control* vector, and in the physical situation, one normally postulates that the user of the plant can adjust the value of $\underset{\sim}{u}(t)$ for each t, subject perhaps to some constraints. The variable $\underset{\sim}{x}(.)$ is normally termed the state vector, and can also be thought of as a response to the control $\underset{\sim}{u}(.)$, see Figure 1; at any particular instant of time t_1 knowledge of the state vector $\underset{\sim}{x}(t_1)$ and the control $\underset{\sim}{u}(\tau)$ for $\tau \geq t_1$, serves to determine plant behaviour for $\tau \geq t_1$, so that $\underset{\sim}{x}(t_1)$ is a summary of past behaviour of the plant, i.e. behaviour for $\tau < t_1$, that is sufficient for the purposes of obtaining future behaviour, i.e. behaviour for $\tau > t_1$.

In qualitative terms, the control problem is one of selecting $\underset{\sim}{u}(.)$ so that the state $\underset{\sim}{x}(.)$ varies in a manner that accomplishes some desirable objective. For the case of an inertial load, the objective might be to reduce the velocity of the load from some nonzero value down to zero. An additional objective might be to achieve this reduction within a certain elapsed time, or to cause the velocity to become zero at a particular distance from the origin, etc. Generally speaking, the objective derives from some aspect associated with the practical use and application of the plant.

At the same time as objectives are formulated, it is frequently the case that inequality constraints must be formulated too. For example, in the problem of reducing the velocity of an inertial load to zero, one might have the constraint that the force (here, the control) be bounded via $|F(t)| \leq F_{max}$, with F_{max} a prescribed constant. This bound might arise because of a need to limit g-forces in the load, or because the prime mover generating the force is practically limited in the force that it can apply.

The bound on $F(t)$ represents a control constraint. It is equally possible to have state constraints; for example, one can envisage a need to brake a vehicle from a nonzero velocity to zero velocity within a certain distance. For the plant represented by (2.3), this implies a constraint on $|x_1(t_f)-x_1(t_0)|$ where t_0 and t_f are respectively the initial time (the time at which the control commences) and the final time, (the time at which the control objective is attained).

Let us summarize what we have said to this point.

1. Optimal control ideas are applied to plants which are dynamical systems; they are usually represented by normal form differential equations of the form of (2.2), with $\underset{\sim}{u}(\cdot)$ denoting a control, and $\underset{\sim}{x}(\cdot)$ the state. The control can generally be freely chosen.
2. There is a desirable objective, or mode of behaviour of the system. This is normally determined from a consideration of the plant application.
3. In achieving the desirable objective, there may be constraints imposed on the control or the state.

Let us now attempt a further quantification of these ideas.

3. QUANTITATIVE DESCRIPTION OF OPTIMAL CONTROL PROBLEMS

In this section, our aim is to indicate the form of a mathematical statement of the optimal control problem.

Plant Description

We have already explained that we require a certain form of differential equation description of the plant, as in (2.2). To pose an optimal control problem, we need to also assume we know $x(t_0)$, an initial condition for (2.2).

[An alternative is to assume $x(t_0)$ is a random variable, but this leads into the difficult subject of stochastic optimal control].

Performance Index

The second thing that we need specified in mathematical terms is a *performance index*; this is a formula which allows us to assign a number to any control-state trajectory pair, and with a properly chosen performance index, *the smaller this number is, the more closely does the plant behaviour approach the desired plant behaviour.*

Let us return to the inertial load of (2.3). Suppose the aim is to apply a control over the time interval [0, 1] which will reduce a nonzero initial velocity to zero at the end of the interval. Then we might take as a performance index

$$V[x_1(0),\ x_2(0),\ u(\cdot)] = x_2^2(1) \tag{3.1}$$

Observe that if the final velocity is zero, the performance index is zero, and the smaller is the performance index, the closer is the final velocity to zero. Observe also that different initial states and different controls will lead to different values for $x_2(1)$ and thus for the performance index. This is the reason for the functional dependence on $\underset{\sim}{x}(0)$ and $u(\cdot)$ on the left side of (3.1).

Another aim might be to transfer from some initial state $[x_1(0),\ x_2(0)]$ to the zero state in as short a time as possible. In this case, we could take

$$V[x_1(0),\ x_2(0),\ u(\cdot)] = \int_{t_0}^{t_f} dt \tag{3.2}$$

where t_f is the first time for which $\underset{\sim}{x}(t_f) = 0$. If at the same time, we were concerned at the fuel used, as measured say by $|u|$, we might adopt for the index

$$V[x_1(0),\ x_2(0),\ u(\cdot)] = \int_{t_0}^{t_f} [k_1|u| + 1]dt \qquad (3.3)$$

and adjustment of k_1 would allow trade-off between the aims of minimizing fuel and minimizing time.

The above examples illustrate the following points:

1. A performance index may consist of final value terms, e.g. (3.1) or integral terms, e.g. (3.2) and (3.3), or a combination.
2. A performance index may involve the state and/or the control [e.g. (3.1) and (3.3)].
3. The final time may be fixed, e.g. (3.1), or indefinite, e.g. (3.2) and (3.3).
4. Associated with specification of the performance index, there may be an additional state equality constraint, e.g. $\underset{\sim}{x}(t_f) = 0$, associated with (3.2).

Just as we can write down the general form of equations describing a system, so we can write down the general form of performance index, as

$$V(\underset{\sim}{x}(t_0),\ \underset{\sim}{u}(\cdot)) = \int_{t_0}^{t_f} \ell[\underset{\sim}{x}(t),\ \underset{\sim}{u}(t),\ t]dt + \phi[\underset{\sim}{x}(t_f),\ t_f] \qquad (3.4)$$

with the possibility of a final value constraint

$$\underset{\sim}{\psi}[\underset{\sim}{x}(t_f)] = 0 \qquad (3.5)$$

and the possibility of t_f being free or fixed.

This form of index certainly excludes some which might seem logical candidates. However, more complex indices generally lead to unsolvable optimal control problems.

Constraints

As noted earlier, we can conceive of control and state variable inequality constraints. The simplest form of constraint with which optimal control deals are control constraints, of the form

$$u_{min}(t) \leq u(t) \leq u_{max}(t)$$

for a scalar control u(.). Here $u_{min}(.)$ and $u_{max}(.)$ are known functions. Practically, the constraint $|u(t)| \leq \alpha$ is very common, where α is a fixed constant. More generally, we can write the constraints as

$$\underset{\sim}{u}(t) \in U(t) \qquad t \in [t_0, t_f]$$

where U(t) is the constraint set at time t.

State variable constraints prove particularly awkward to handle, and penalty function methods are far more popular. Nevertheless, one can write a constraint as

$$\underset{\sim}{x}(t) \in X(t) \qquad t \in [t_0, t_f].$$

Functional constraints of the type

$$\int_{t_0}^{t_f} u^2(t)dt \leq \alpha$$

do not appear frequently in practice, are awkward to deal with, and are normally avoided simply by including a term $k_2 \int u^2 dt$ in the performance index. The constant k_2 is adjusted if necessary to ensure satisfaction of the constraint.

Let us now put the above ideas together to state the optimal control problem.

Optimal Control Problem

The problem is simply one of finding the control $\underline{u}(\cdot)$ which minimizes the performance index, provided that the various equality and inequality constraints are satisfied:

Given a system

$$\dot{\underline{x}} = \underline{f}(\underline{x},\underline{u},t)$$

with prescribed inital state $\underline{x}(t_0)$, find a control $\underline{u}(\cdot)$ to minimize

$$V[\underline{x}(t_0),\underline{u}(\cdot)] = \int_{t_0}^{t_f} \ell[\underline{x}(t),\underline{u}(t),t]dt + \phi[\underline{x}(t_f),t_f]$$

where t_f may be free or fixed, and $\ell(\cdot,\cdot,\cdot)$ and $\phi(\cdot,\cdot)$ are known functions, subject to satisfaction of the constraints

$$\underline{\psi}[\underline{x}(t_f)] = 0, \qquad \underline{u}(t) \in U(t), \qquad \underline{x}(t) \in X(t).$$

4. EXISTENCE AND UNIQUENESS PROBLEMS

There is no guarantee that there exists even one control which will achieve the minimization called for in the optimal control problem. There may also be a control which causes V to be $-\infty$, although most commonly, the functions $\ell(\cdot,\cdot,\cdot)$ and $\phi(.,.)$ are always nonnegative, so that V is also always nonnegative. Further, there may be many controls achieving the minimization. All these difficulties are of course reflections of standard difficulties in static optimization.

5. CONTRAST WITH STATIC OPTIMIZATION

We have already referred to some parallels between the dynamic optimization problems of optimal control and static optimization problems. Let us note further comparisons.

1. Both static and dynamic optimization are concerned with minimizing a performance index, subject to equality and inequality constraints.
2. The performance index in dynamic problems is a *functional*, rather than a function, of the controlling quantity $\underset{\sim}{u}(\cdot)$. Functions, rather than functionals, appear in static optimization.
3. In dynamic problems, one of the equality constraints, $\dot{\underset{\sim}{x}} - \underset{\sim}{f}(\underset{\sim}{x},\underset{\sim}{u},t) = 0$, has a dynamic character, and even a *causal* character, in the sense that knowing $\underset{\sim}{x}(t_0)$ and $\underset{\sim}{u}(\cdot)$, there is a corresponding trajectory $\underset{\sim}{x}(\cdot)$ computable forwards in time from t_0. Occurrence of such constraints in static optimization is rare, though an important quasi-exception is discrete time optimal control.

6. OPEN LOOP AND CLOSED LOOP CONTROL

There is one further peculiarity of optimal control that is apparently not generally shared with static optimization. This is the possibility of *practically implementing* the optimal control in one of two ways. The first way, *open-loop* control, is easy to understand. One solves the optimal control problem by some means, and thereby obtains a certain function of time, $\underset{\sim}{u}^*(\cdot)$, which achieves the minimization. Then one simply applies a control $\underset{\sim}{u}(\cdot)$ to the plant, ensuring that $\underset{\sim}{u}(t) = \underset{\sim}{u}^*(t)$.

Almost always, $\underset{\sim}{u}^*(\cdot)$ will be a different function for each inital state $\underset{\sim}{x}(t_0)$, so the hardware required to generate all the optimal controls, assuming they can be calculated, will in general be complex.

An alternative is to use what is known as *closed-loop* control. A deceptively simple theoretical result due to Bellman, but none the less one of major importance, states that, if an optimal control exists, the value it takes at time t, viz. $\underset{\sim}{u}^*(t)$, depends only on $\underset{\sim}{x}(t)$ and t, and in particular, is independent of $\underset{\sim}{x}(t_0)$. In other words, there exists a function $\underset{\sim}{k}(.,.)$, depending on

$\underset{\sim}{f}(.,.,.)$, $\ell(.,.,.,.)$, $\phi(.,.,.)$, $U(\cdot)$, $X(\cdot)$, and $\psi(\cdot)$, *but not on* $\underset{\sim}{x}(t_0)$ such that

$$\underset{\sim}{u}^*(t) = \underset{\sim}{k}(\underset{\sim}{x}(t), t)$$

for all $\underset{\sim}{x}(t_0)$.

The practical significance of this fact is that one can build a feedback controller, that generates $\underset{\sim}{u}^*(t)$ as a memoryless function of the system state, see Figure 2. This allows a major hardware simplification in implementing an optimal control, particularly since the same $\underset{\sim}{k}(.,.)$ works for all $\underset{\sim}{x}(t_0)$. Parenthetically, we note that $\underset{\sim}{k}(.,.)$ is frequently time-invariant, i.e.

$$\underset{\sim}{k}[\underset{\sim}{x}(t), t] \equiv \underset{\sim}{k}[\underset{\sim}{x}(t)]$$

and implementation is even simpler.

7. DISCRETE TIME OPTIMAL CONTROL PROBLEM

It may be that the plant which one desires to control is describable by a *difference equation* of the form

$$\underset{\sim}{x}(k+1) = \underset{\sim}{f}(\underset{\sim}{x}(k), \underset{\sim}{u}(k), k) \qquad k = k_0, k_0 + 1, \ldots \qquad (7.1)$$

Such a description is particularly appropriate when computers will play a role in implementing, as distinct from computing in advance, the optimal control. One then introduces a discrete time performance index, of the form

$$V[\underset{\sim}{x}(k_0), \underset{\sim}{u}(\cdot)] = \sum_{k=k_0}^{k_f-1} \ell[\underset{\sim}{x}(k), \underset{\sim}{u}(k), k] + \phi[\underset{\sim}{x}(k_f), k_f] \qquad (7.2)$$

with $\ell(.,.,.)$ and $\phi(.,.)$ chosen to reflect the aims of control. As for the continuous time problem, k_f may be free or fixed, and there may be an additional constraint $\underset{\sim}{\psi}[\underset{\sim}{x}(k_f)] = 0$. Further, $\underset{\sim}{u}(k)$ and $\underset{\sim}{x}(k)$ for each k may have to lie in a constraint set.

For finite k_f, V becomes a function of a finite number of variables, and V has to be minimized with respect to $\underset{\sim}{u}(k_0)$, $\underset{\sim}{u}(k_0+1)$,... $\underset{\sim}{u}(k_f-1)$. Therefore, one is faced with a problem of the sort tackled by static optimization theory: minimize a function, subject to a set of equality and inequality constraints.

The connection between discrete optimal control and nonlinear programming has been well exploited. Optimal control problems do however possess a structure not possessed by all nonlinear programming problems. This structure leads, for example, to the fact that the optimal control $\underset{\sim}{u}^*(k)$ at time k, assuming it exists, depends on $\underset{\sim}{x}(k)$ and k, but not on the particular $\underset{\sim}{x}(k_0)$. Further, it is this structure which may permit solution of the optimal control problem even if, when considered as a problem of nonlinear programming, it appears to have impossibly large dimension.

8. MAJOR THEORETICAL RESULTS - CONTINUOUS TIME PROBLEMS

There are two major theoretical results of optimal control (as well as an extraordinary number of lesser but still important results). These theoretical results are often used when actual solution of optimal control problems is attempted. They are generally titled: the Hamilton-Jacobi-Bellman (HJB) equation, and the Pontryagin Minimum (or Maximum) Principle (PMP).

The HJB equation is the less useful. As the names Hamilton and Jacobi may suggest, the HJB theory is an outgrowth of classical calculus of variations theory. It is therefore not surprising that its application demands absence of inequality constraints, and smoothness of the plant differential equation description and the functions appearing in the performance index. Even then, the equation may not be usable. The equation itself is a *nonlinear partial differential equation* whose solution provides the optimal performance index. The optimal control is obtained from the optimal performance index via a differentiation and minimization operation, and is automatically given in *closed-loop* form. The equation is rarely solvable. However, there do exist two major classes of problems, important from the engineering point of view, where solution is possible.

In contrast to the difficulties that constraints cause in applying the HJB equation, the bread and butter applications of the PMP are those involving control constraints. (State constraints remains awkward to handle). The end result of application of the PMP is a two point boundary value problem (TPBVP), most frequently of the form

$$\dot{\underset{\sim}{x}}^* = \underset{\sim}{f}_1(\underset{\sim}{x}^*,\underset{\sim}{p}^*,t) \qquad \underset{\sim}{x}^*(t_0) = \underset{\sim}{x}(t_0)$$

$$\dot{\underset{\sim}{p}}^* = \underset{\sim}{f}_2(\underset{\sim}{x}^*,\underset{\sim}{p}^*,t) \qquad \underset{\sim}{p}^*(t_f) = 0 \ . \qquad (8.1)$$

where $\underset{\sim}{f}_1(\cdot,\cdot,\cdot)$ and $\underset{\sim}{f}_2(\cdot,\cdot,\cdot)$ are known functions. The trajectory $\underset{\sim}{x}^*(\cdot)$ is the optimal trajectory, and the optimal control is computable, in *open-loop* form, via a minimization operation involving $\underset{\sim}{x}^*$ and $\underset{\sim}{p}^*$.

9. MAJOR THEORETICAL RESULTS - DISCRETE TIME PROBLEMS

It is a pleasant surprise to discover that the discrete time equivalent of the HJB equation, the Bellman Principle of Optimality, is free of almost all the problems associated with the HJB equation, and constitutes the basis for the majority of numerical techniques for solving discrete time optimal control problems.

The formal statement of the Principle proceeds roughly as follows.

Suppose we are dealing with the system (7.1), with initial state $\underset{\sim}{x}(k_0)$, and with performance index (7.2). Suppose $\underset{\sim}{u}^*(k_0)$, $\underset{\sim}{u}^*(k_0+1)$ ... is the optimal control. Let k_1 be an arbitrary time instant in the control interval, and let $\underset{\sim}{x}(k_1)$ be the state resulting from initial state $\underset{\sim}{x}(k_0)$, and control sequence $\underset{\sim}{u}^*(k_0),\ldots \underset{\sim}{u}^*(k_1-1)$.

Now consider a second optimal control problem, with the same system (7.1), with control interval (k_1, k_f), with initial state $\underline{x}(k_1)$, and with performance index

$$V[\underline{x}(k_1),\underline{u}(\cdot)] = \sum_{k=k_1}^{k_f-1} \ell[\underline{x}(k),\underline{u}(k),k] + \phi[\underline{x}(k_f),k_f] \qquad (9.1)$$

This is the same index as (7.2), except that the contributions to the performance index from controls and states over the interval $[k_0,k_1-1]$ have been lopped off.

The Principle states that the optimal control sequence for this second index is $\underline{u}^*(k_1),\ldots\ \underline{u}^*(k_f-1)$, i.e. the remaining part of the sequence for the first index. Thus, one can loosely say that subsequences of optimal control sequences are themselves optimal.

The nontrivial consequences of this result, particularly for tackling the problem of computing optimal controls, are enormous. To secure these advantages, a continuous time problem will frequently be artificially converted to a discrete time problem.

Just as there is a discrete time parallel to the HJB equation, there is a discrete time parallel to the PMP. Because of the power of the Principle of Optimality, it is not of such great importance as for continuous time systems.

10. SOLUTION OF OPTIMAL CONTROL PROBLEMS

Generally speaking, the solution of optimal control problems is a difficult computational task. Many methods have been developed, but here, we can do little more than list them.

It proves convenient to break down the methods into four categories:

(1) Dynamic Programming. This method applies to discrete time systems, and is based on the Principle of Optimality. It is flexible and deservedly popular. There are highly developed computational algorithms, some involving polynomial approximation, and others demanding theoretical insights, e.g. state increment dynamic programming, and differential dynamic programming.

(2) HJB Equation. This is almost useless, except for special classes of problems.

(3) "Direct" Methods. The characteristic of these methods is that they avoid the two point boundary value problem. The *gradient methods* (First order gradient, second order gradient, conjugate gradient, etc.) are those which compute a variation in the performance index value due to a variation in the control, and then vary the control to reduce the performance index. Expansion of control and trajectory in terms of basis functions, thus yielding a finite-dimensional minimization problem, is the basis of the *Ritz method*.

(4) Methods based on Solving Two Point Boundary Value Problem. *Quasilinearization* replaces the nonlinear TPBVP by a sequence of linear problems. *Invariant imbedding* attempts to imbed the nonlinear TPBVP within a set of easier solved problems. *Newton-Raphson* type procedures can be used to adjust boundary conditions until solution is obtained.

The different methods all have varying convergence characteristics. Naturally, convergence properties vary from problem to problem, and with the initial guess in an iterative procedure. Thus, a first order gradient method converges from most initial conditions, but is slow to converge near the optimum. Quasilinearization converges only if one is near the optimum, and then convergence is fast.

11. ENGINEERING SIGNIFICANCE OF OPTIMAL CONTROL

We have earlier noted that a performance index is normally chosen with some physical, or engineering, aim in mind. We make the point here that if the optimal control is realized in closed-loop form, the resulting closed-loop arrangement may have a significant number of valuable properties, over and above those directly associated with minimization of the performance index. Particularly is this the case with linear plants, which are frequently encountered in practice. The closed-loop system resulting from optimization with a performance index of the type

$$V = \int_{t_0}^{\infty} (\underline{u}'R\underline{u} + \underline{x}'Q\underline{x})dt$$

with Q and R positive definite symmetric matrices has good stability properties, and is tolerant of nonlinearities, noise, and plant parameter variation. In general terms, one could call it robust. Consequently, linear system design may often be done via optimal control methods, more to achieve these bonus properties than to ensure minimization of a particular performance index.

FIGURE 1

The control $u(\cdot)$ and its response $x(\cdot)$; $x(t_1)$ summarises plant behaviour up to t_1.

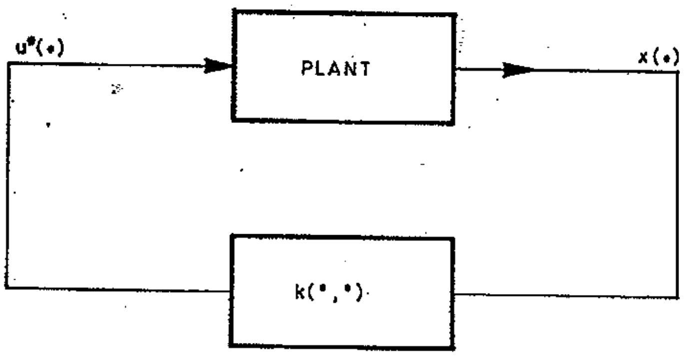

FIGURE 2

Feedback implementation of an optimal control via a memoryless function of the state.

References

For an introduction to optimal control problems and their numerical solution, the following books are recommended. They are listed in order of difficulty.

[1] Kirk, D.E. *Optimal Control Theory - An Introduction*, Prentice Hall Inc., New Jersey, 1970.

[2] Bryson, A.E. and Ho, Y.C. *Applied Optimal Control*, Blaisdell Pub. Co., Mass., 1969.

[3] Sage, A.P. *Optimum Systems Control*, Prentice Hall Inc., New Jersey, 1968.

For a rigorous introduction, with emphasis on the PMP, and with no material on computational procedures, see

[4] Athans, M. and Falb, P.L. *Optimal Control*, McGraw Hill Book Co., New York, 1966.

For a treatment of linear system design via optimal control, see

[5] Anderson, B.D.O. and Moore, J.B. *Linear Optimal Control*, Prentice Hall Inc., New Jersey, 1971.

For connections between optimal control and mathematical programming, see

[6] Tabak, D. and Kuo, B.C. *Optimal Control by Mathematical Programming*, Prentice Hall Inc., New Jersey, 1971.

[7] Canon, M.D., Cullum, C.D. and Polak, E. *Theory of Optimal Control and Mathematical Programming*, McGraw Hill Book Co., New York, 1970.