

A PRIMAL-DUAL ALGORITHM
FOR THE COMPUTATION OF OPTIMAL CONTROL*

David G. Luenberger
Stanford Electronics Labs, Stanford University,
Palo Alto, California

ABSTRACT

This paper describes a computational method, based on the finite-dimensional duality theorem of convex programming, for the solution of a wide class of optimal control problems. The method converts an optimal control problem having terminal constraints to a series of unconstrained problems that can be rapidly solved by standard gradient techniques. The method is applicable to problems having either fixed or free terminal time.

I. INTRODUCTION

In this paper a computational procedure is proposed for solving the following class of optimal control problems: given the dynamic system

$$\dot{x}(t) = f(x(t), u(t), t) \quad x(t_0) = x_0 \text{ fixed}; \quad (1)$$

find the control input function $u(t)$ that satisfies the constraints

$$u(t) \in U(x, t), \quad (2)$$

$$G(x(t_f), t_f) \leq \theta, \quad (3)$$

and minimizes the real-valued objective functional

$$\psi(x(t_f), t_f). \quad (4)$$

*This research was supported in part by Joint Services Contract Nonr 225(83) and in part by Contract DA-01-021-AMC-9000(Y) awarded to Stanford Research Institute.

COMPUTING METHODS IN OPTIMIZATION PROBLEMS

The function f takes its values in Euclidean n -space E_n and is assumed to be continuously differentiable with respect to x and u . The subset $U(x,t) \subseteq E_p$, which determines the admissible control values, is usually a connected set depending continuously on x and t . In addition to Eq. (2), the control function is required to be piecewise continuous.

The terminal time t_f may be either fixed or free. The only constraints on the state vector x occur at the terminal time t_f and are described by the inequality Eq.(3), where G is an m -dimensional vector-valued function of x and t . Problems having integral constraints or an integral objective functional of the form

$$\int_{t_0}^{t_f} l(x(t), u(t), t) dt$$

can, as is well known, be converted to problems of the type described above. As presented here, however, the procedure is not applicable to problems having state vector constraints at each $t \in [t_0, t_f]$, but the method can be modified to handle these problems as well.

The selection of a suitable computational procedure for such a problem depends essentially on the nature of the constraints (i.e., the set U and the function G). In the special case of no constraints and fixed final time, a variety of effective computational methods are available including the gradient method (1), the method of conjugate gradients (2,3), parallel tangents (4), and Newton's method (5). Most of these techniques are easily modified to account for the constraint $u(t) \in U[x,t]$, but often cannot be modified to handle the terminal constraints in any generally effective manner. In this paper it is shown that the efficiency of these routines for solving unconstrained control problems can, in conjunction with duality theory, be exploited to solve constrained problems.

In Sec.II the essential elements of duality theory for mathematical programming problems upon which the proposed method is based are presented. Then in Sec.III the proposed method is described for constrained optimal control problems having fixed terminal time. The method is extended to problems with free terminal time in Secs. IV, V, and VI. Examples of the method are presented in Sec.VII.

II. BASIC PRIMAL-DUAL ALGORITHM IN FINITE DIMENSIONS

Let Γ be a convex subset of an n -dimensional vector space X , and let ψ be a real-valued convex functional on Γ . Let F be a convex mapping from Γ into Euclidean m -space E_m .

COMPUTING METHODS IN OPTIMIZATION PROBLEMS

A general convex programming problem has the form

$$\left. \begin{array}{l} \text{minimize } \psi(x) \\ \text{subject to } x \in \Gamma, G(x) \leq \theta \end{array} \right\} \quad (5)$$

In order to avoid certain singularity difficulties, it is assumed that the origin $\theta \in E_m$ is an interior point of the set

$$\{z: G(x) \leq z \text{ for some } x \in \Gamma\}$$

This is referred to as the regularity condition.

Definition: The dual function ϕ , corresponding to problem (5), is defined on the set

$$\Gamma^* = \{\lambda \in E_m : \inf_{x \in \Gamma} [\psi(x) + \lambda'G(x)] > -\infty\}$$

by the equation

$$\phi(\lambda) = \inf_{x \in \Gamma} [\psi(x) + \lambda'G(x)]. \quad (6)$$

The following duality theorem, which is a direct consequence of the Lagrange saddle-value theorem (6), is applicable under the assumptions stated above[†] and forms the basis for the primal-dual method:

Theorem: Let

$$\mu_0 = \inf \{\psi(x) : x \in \Gamma, G(x) \leq \theta\} \quad (7)$$

be finite, then

$$\mu_0 = \max_{\lambda \in \Gamma^*} \phi(\lambda) \quad (8)$$

Furthermore, if λ_0 is the maximizer of Eq.(8) and the infimum of Eq. (7) is achieved by x_0 , then x_0 minimizes $\psi(x) + \lambda_0'G(x)$.

The primal-dual method for solving Eq.(7) consists of solving, instead, the dual problem Eq.(8), which has simpler constraints. Of course, evaluation of $\phi(\lambda)$ for any given λ itself requires the solution of a minimization problem, but one that is unconstrained.

[†] Define the set $A \subset R \times Z$ as $A = \{(r,z): \text{For some } x \in \Gamma, G(x) \leq z, f(x) \leq r\}$. Then the duality theorem holds if the regularity condition is satisfied and A is convex. Convexity of A is implied by convexity of Γ, ψ and G .

COMPUTING METHODS IN OPTIMIZATION PROBLEMS

In the course of the primal-dual procedure, the dual function is maximized either by a simple search technique if m is small (say $m = 1$ or 2), or more often by a gradient-based technique such as steepest ascent or conjugate gradients. Such gradient techniques for solving the dual problem are particularly attractive in view of the following observation:

Lemma: Given λ_0 let x_0 be a minimizer of $\psi(x) + \lambda_0'G(x)$, subject to $x \in \Gamma$. Then $\lambda_0'G(x_0) \leq 0$ implies $\varphi(\lambda_0 + \lambda) \leq \varphi(\lambda_0)$.

Proof: $\varphi(\lambda_0 + \lambda) = \inf \{f(x) + \lambda_0'G(x) + \lambda'G(x)\}$

$$\leq f(x_0) + \lambda_0'G(x_0) + \lambda'G(x_0)$$

$$\leq f(x_0) + \lambda_0'G(x_0)$$

$$= \varphi(\lambda_0).$$

The above statement says that the vector $G(x_0)$ defines a half-space of nonascent for the function φ and is therefore a so-called subgradient of φ at λ_0 . In particular, if φ has a gradient at λ_0 , it is equal to $G(x_0)$.

The primal-dual method for solving the constrained problem Eq(5) proceeds by first selecting $\lambda \geq \theta$, solving the unconstrained primal problem Eq.(6) to evaluate $\varphi(\lambda)$ and its gradient, and then changing λ , according to a gradient-based technique, to increase φ ; and so forth. A single constrained problem is therefore converted to two interwoven simpler problems.

III. FIXED TERMINAL-TIME CONTROL PROBLEMS

Consider the fixed terminal-time optimal control problem

$$\left. \begin{array}{l} \text{minimize } \psi(x(t_f)) \\ \text{subject to } G(x(t_f)) \leq \theta \\ u(t) \in U(x,t) \\ \dot{x}(t) = f(x,u,t) \end{array} \right\} \quad (9)$$

with $x(t_0)$ specified. Here the explicit dependence of ψ and G on t_f is suppressed since t_f is fixed. This problem can be solved by the primal-dual method by formulating the control problem in the finite-dimensional space of terminal state vectors $x(t_f)$ and considering the finite-dimensional problem of minimizing $\psi(x(t_f))$ subject to $G(x(t_f)) \leq \theta$ and $x(t_f) \in \Gamma$, where Γ is the set of end points attainable from $x(t_0)$ by application of admissible controls.

COMPUTING METHODS IN OPTIMIZATION PROBLEMS

Assuming that ψ and G are convex functions and that Γ is a convex set (this assumption is examined below), application of the dual method converts the constrained control problem to the problem:

$$\max_{\substack{\lambda \geq 0 \\ \lambda \in E_m}} \left\{ \min_{x(t_f) \in \Gamma} [\psi(x(t_f)) + \lambda'G(x(t_f))] \right\}$$

The interior minimization reduces to a control problem of type Eq. (9) but without terminal constraints. This problem is therefore susceptible to any of the standard minimization routines. The exterior problem can be solved by a finite-dimensional gradient-based technique as discussed in Sec.II.

In order to guarantee the success of this approach, some assumption must be made concerning the nature of the set Γ . Obviously, it is sufficient that Γ be convex, but in most practical problems involving nonlinear differential equations this assumption is probably not justified. This assumption can, however, be replaced by a weaker, more realistic, assumption.

Suppose that the terminal state vector $x(t_f)$ is partitioned as

$$x(t_f) = \begin{bmatrix} y \\ w \end{bmatrix},$$

where $\dim y = p$ and $\dim w = n - p$. Suppose further that in the original optimal control problem both the terminal objective functional ψ and terminal constraints G can be written as explicit functions of the vector y . Thus

$$\psi(x(t_f)) = \psi(y)$$

$$G(x(t_f)) = G(y).$$

Often in large complex control problems p is considerably less than n . In rocket guidance, for instance, objectives and constraints are often expressed in terms of position variables only, without explicit reference to velocities.

Definition: The optimal control problem Eq. (9) with $\psi(x(t_f)) = \psi(y)$, $G(x(t_f)) = G(y)$ is said to satisfy the weakened convexity hypothesis if the set

$\gamma = \{y: x(t_f) = \begin{bmatrix} y \\ w \end{bmatrix} \text{ for some } w \text{ and some admissible control}\}$ is convex.

The set γ consists of all reachable y points in the p -dimensional space. An equivalent characterization of γ is that it is the projection of the set Γ in E_n onto the space generated by the first p components of $\overset{n}{x}(t_f)$.

Convexity of γ is a weaker assumption than convexity of Γ since, for example, in the extreme case where $p = 1$, Γ need only be connected for γ to be convex.

As a general rule, for a given system, the weakened convexity assumption is more likely to be satisfied as p is reduced. Therefore, it appears that in considering various problem formulations the analyst might well strive to select ψ and G so as to minimize p .

Satisfaction of the weakened convexity assumption is, of course, sufficient to guarantee success of the primal-dual algorithm since the optimal control problem is then equivalent to

$$\begin{aligned} & \text{minimize} && \psi(y), \\ & \text{subject to} && G(y) \leq \theta \\ & && y \in \gamma \end{aligned}$$

IV. WEAKENED CONVEXITY FOR VARIABLE-TIME PROBLEMS

Now consider problems with unspecified terminal time of the form

$$\left. \begin{aligned} & \text{minimize} && \psi(x(t_f), t_f) \\ & \text{subject to} && \dot{x} = f(x, u, t), \quad x(t_0) \text{ given} \\ & && u(t) \in U(x, t) \\ & && G(x(t_f), t_f) \leq \theta. \end{aligned} \right\} \quad (10)$$

A problem of this form having particular interest is that of hitting a target, specified by $G[x(t_f)] \leq \theta$, in minimum time; i.e., $\psi(x(t_f), t_f) = t_f$.

Since the additional variable t_f is introduced in these problems, one must consider the geometry in the space $X \times T$. In $X \times T$, let Γ be the set of all points attainable from the dynamical system by use of admissible control functions, and, in a manner similar to that of the previous section, let γ be the projection of Γ onto the subspace $Y \times T$ where Y represents the space of variables appearing explicitly in ψ and G .

Let γ_t be the set in the space Y , which is the cross-section of γ at time t . From the discussion in the previous section, γ_t can be argued to be convex in many cases, although γ itself may not be. The sets γ_t , however, are for many systems, increasing as t increases, i.e. $\gamma_{t_1} \subset \gamma_{t_2}$ for $t_1 \leq t_2$. In these cases the t -axis can be rescaled monotonically so that γ , when drawn with respect to this new variable, say t' , is convex. Such a transformation from the variable t to the variable t' changes the form of the underlying differential equations

and the cost functional, but it is assumed that with respect to the new variable the usual convexity requirements are satisfied.

As an example, in the problem of hitting a target in minimum time, if γ is not convex but can be made convex by a monotone increasing transformation $t' = m(t)$, then the problem can be regarded as minimization of t' and all the convexity requirements will now be satisfied. Finally, observe that it is not actually necessary to perform the transformation to t' as indicated since minimization of t' is itself equivalent to minimization of t . It is only necessary to hypothesize that an appropriate transformation exists so that the theory is justified.

For the general variable terminal-time problem, it is sufficient to require satisfaction of the following:
 Weakened Convexity Assumption:[‡] Let Y be the subspace of X , which explicitly appears in $G(x(t_f), t_f)$ and $\psi(t_f, t_f)$. The optimal problem satisfies the weakened convexity hypothesis if there is a monotone increasing transformation of the t -axis such that the projection of the set Γ onto $Y \times T$ is convex, and the functions $\psi(y, t)$ and $G(y, t)$ are convex on $Y \times T$.

V. COMPUTATIONAL PROCEDURE FOR VARIABLE TERMINAL TIME

In this section the primal-dual algorithm is applied to variable terminal-time optimal control problems. The method is based on the following special duality theorem.

Theorem: Under the weakened convexity assumption, problem (10) is equivalent to

$$\max_{\lambda \geq 0} \min_{u \in U} \min_{t_f} \{ \psi(x(t_f), t_f) + \lambda' G(x(t_f), t_f) \}$$

Proof: The two minimizations are equivalent to minimization, over the convex set γ , of an objective function that is itself a convex function. Hence, the equation can be conceived as

$$\max_{\lambda \geq 0} \min_{(y, t) \in \gamma} \{ \psi(y, t) + \lambda' G(y, t) \}$$

[‡] The weakened convexity assumption implies convexity of the set A in footnote [†].

COMPUTING METHODS IN OPTIMIZATION PROBLEMS

Now by ordinary duality theory this is equivalent to

$$\begin{aligned} & \text{minimize } \psi(y,t) \\ & (y,t) \in \gamma \\ & \text{subject to } G(y,t) \leq \theta. \end{aligned}$$

This is the desired conclusion.

The algorithm for solving such problems is now read from Eq.(11) as follows:

- (1) For a given $\lambda \geq \theta$, use any standard algorithm to find $u \in U$ minimizing the objective functional

$$\min_{t_f} \{ \psi(x(t_f), t_f) + \lambda' G(x(t_f), t_f) \}$$

(This is discussed in greater detail below.)

- (2) Maximize with respect to λ by using any standard gradient-based routine by using the fact that the gradient is $G(x(t_f), t_f)$.

The weakened convexity assumption is sufficient to guarantee the success of this method. In practice the method has been successful for a large number of complex rocket problems where the weakened convexity assumption could not be verified. In general one would expect the method to succeed in many well-posed problems.

VI EVALUATION OF GRADIENT FOR CONTROL SUBPROBLEM

As a fundamental step in the procedure outlined in the preceding section, one must solve an optimal control problem of the form*

$$\begin{aligned} & \text{Minimize } \{ \min_{t_f} \psi(x(t_f), t_f) \} \\ & \text{subject to } \dot{x}(t) = f(x(t), u(t), t), \quad x(t_0) \text{ given} \\ & u(t) \in U(x, t). \end{aligned}$$

A number of effective algorithms are available for problems of this kind that require evaluation of the gradient of

$$J = \min_{t_f} \psi(x(t_f), t_f)$$

with respect to u . The problem can be written alternatively as

* Here $\psi + \lambda'G$ has, for convenience, been replaced by ψ .

COMPUTING METHODS IN OPTIMIZATION PROBLEMS

$$\begin{aligned} & \min_u \psi(x(t_f), t_f) \\ & \text{subject to } \dot{x} = f(x, u, t), \quad x(t_0) \text{ given} \\ & \quad u(t) \in U(x, t). \\ & t_f = \arg \min \psi(x(t), t) \end{aligned} \quad (12)$$

Therefore, the gradient can be computed for the problem of minimizing ψ subject to the usual differential equation constraint and the additional constraint Eq.(12), which serves as a stopping condition, since it determines t_f once a control function $u(t)$ has been selected.

To first-order, the incremental change in the cost ψ due to a change δu in control and a change δt_f in the final time, is

$$\delta\psi = \int_{t_0}^{t_f} \eta'(t) D(t) \delta u(t) dt + \dot{\psi} \Big|_{t_f} \cdot \delta t_f \quad (13)$$

where

$$\dot{\eta}(t) = -F'(t) \eta(t), \quad \eta(t_f) = \psi_x \Big|_{t_f}$$

$$F(t) = \begin{bmatrix} \frac{\partial f_i}{\partial x_j} \\ \frac{\partial f_i}{\partial u_j} \end{bmatrix}$$

$$D(t) = \begin{bmatrix} \frac{\partial f_i}{\partial x_j} \\ \frac{\partial f_i}{\partial u_j} \end{bmatrix}$$

The change δt_f is determined by the stopping condition Eq. (12) but is of no direct importance at present since the total derivative of ψ with respect to t , $\dot{\psi}$, is zero at t_f , and, hence, δt_f does not enter Eq. (13). Thus, the gradient of ψ with respect to u is simply the function $\eta'(t) D(t)$.

In a typical application, then, one changes u by some constant multiple of this gradient and integrates the equations of motion forward until the optimal t_f is found satisfying the stopping condition Eq.(12). In case there are constraints $u \in U(x, t)$, the standard methods for selecting u on the basis of the gradient can be employed.

If it turns out that δt_f is positive during one of these iteration steps, the new step will require selecting $u(t)$ on an interval that is larger than that on which it was previously defined. It probably makes little difference how u is selected on this interval, but one possible method, having certain theoretical justification, is to select $u(t)$ so as to minimize $\psi_x f(x, u, t) + \psi_t$ subject to $u \in U(x, t)$.

This tends to reduce $\psi_f + \psi_t$ to zero as soon as possible and hence minimizes the x_{new}^t . In the problems solved so far, however, the control function has been extended to a larger interval when necessary by using the same value as at the last available time instant.

VII. EXAMPLES

Two examples illustrating the type of results obtained by the primal-dual method are presented in this section. Both examples consider Van der Pol's equation, which in state variable form is

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= x_1 + (1 - x_1^2) x_2 + u, \end{aligned}$$

with initial conditions $x_1(0) = 1, x_2(0) = 0$

1. Fixed Time, Minimum Energy to a Line

In this problem one seeks to minimize

$$J = \frac{1}{2} \int_0^5 (x_1^2 + x_2^2 + u^2) dt$$

while satisfying the terminal constraint

$$-x_1(5) + x_2(5) = 1.$$

In dual form the problem becomes

$$\max_{\lambda} \min_u \frac{1}{2} \int_0^5 (x_1^2 + x_2^2 + u^2) dt + \lambda[-x_1(5) + x_2(5) - 1] \tag{14}$$

The derivative of the dual function is simply $H(\lambda) = -x_1(5) + x_2(5) - 1$, where $x_1(5)$ and $x_2(5)$ are taken as the terminal state variables resulting from the indicated minimization in Eq. (14). The problem was solved by finding λ to make $H(\lambda) = 0$, which corresponds to a stationary point of the dual function. The equation $H(\lambda) = 0$ was solved using regula-falsi (7) starting with $\lambda_1 = 0.0, \lambda_2 = 1.0$. In each case the unconstrained optimal control subproblem was solved using the method of parallel tangents. The resulting succession of trajectories is shown in Fig. 1.

2. Bounded Control, Minimum Time to a Disk

In this problem the control u is constrained to satisfy $|u(t)| \leq 1$, and the objective is to minimize time to the disk

COMPUTING METHODS IN OPTIMIZATION PROBLEMS

$$x_1^2 + (x_2 - 1)^2 \leq .04.$$

In solving the problem by the primal-dual method, the objective function was taken to be $t_f^2/10$ so that at each stage of the iterative process the modified objective function was

$$t_f^2/10 + \lambda[x_1^2 + (x_2 - 1)^2 - .04]$$

for some positive scalar λ . Each of the optimal control subproblems was without terminal constraints but was subject to $|u(t)| \leq 1$, which rendered solution by a gradient-based technique somewhat slow but still feasible.

Starting with $\lambda_1 = 0.5$ and $\lambda_2 = 1.5$, a stationary point of the dual function was found by regula-falsi. Representative trajectories, optimal for various λ , are shown in Fig. 2.

REFERENCES

1. A.E. Bryson and W.F. Denham, "A Steepest-Ascent Method for Solving Optimum Programming Problems," ASME, J. of Appl.Mech., June 1962, pp. 247-257.
2. L.S. Lasdon, S.K. Mitter, and A.D. Waren, "The Conjugate Gradient Method for Optimal Control Problems," IEEE Trans. on Automatic Control, Vol.AC-12, No.2, pp.132-138 (April 1967).
3. J.F. Sinnott and D.G. Luenberger, "Solution of Optimal Control Problems by the Method of Conjugate Gradients," Preprints of 1967 Joint Automatic Control Conference, Philadelphia, pp. 566-573.
4. B.V. Shah, R.J. Buehler, and O. Kempthorne, "Some Algorithms for Minimizing a Function of Several Variables," J. Soc. Indust. Appl. Math., Vol.12, No.1, p.p. 74-91 (March 1964).
5. H.J. Kelley, R.E. Kopp, and H.G. Moyer, "A Trajectory Optimization Technique Based upon the Theory of the Second Variation," Progress in Astronautics and Aeronautics, 14, pp.559-582 (Academic Press, New York, 1964).
6. Samuel Karlin, Mathematical Methods and Theory in Games, Programming, and Economics, Vol.1 (Addison-Wesley Publishing Company, Inc., Reading, Massachusetts, 1959).
7. A.Ralston, A First Course in Numerical Analysis, p.323 (McGraw-Hill, New York, 1965).




Fig. 1 Trajectories for Problem 1

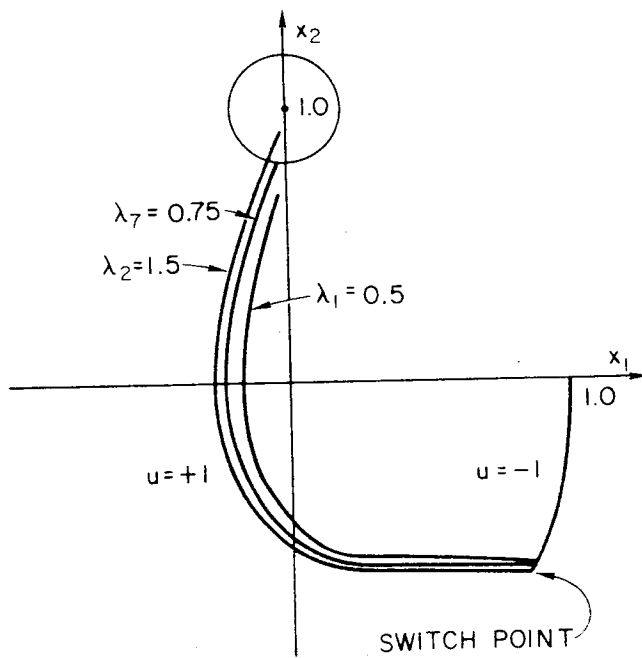


Fig. 2 Trajectories for Problem 2