

TOWARD GLOBAL OPTIMIZATION OF A DIFFICULT OPTIMAL CONTROL PROBLEM

R. LUUS

(Department of Chemical Engineering, University of Toronto, Toronto, ON M5S 3E5, CANADA)

Received: July 5, 2001

The optimal control problem introduced in 1980 by Yeo [1] exhibits numerous local optima which have not been reported in the literature. When 50 time stages of equal length are used, there are only two local optima. When 80 time stages are taken, there are over 50 local optima. It was found that there are at least 15 local optima which are better than the one reported in the literature as the global optimum. With the use of 100 time stages of equal length, the number of local optima is increased substantially to over 150. With 100 time stages of equal length, the global optimum solution found with iterative dynamic programming (IDP) gives $I = 0.119044$. This value of the performance index is 0.2% better than the value that had been accepted in the literature as the global optimum. IDP provides a good means of getting the global optimum solution for this problem, but numerous runs are necessary due to the presence of a large number of local optima.

Keywords: optimal singular control, iterative dynamic programming, IDP, global optimization, sensitivity

Introduction

Establishing the optimal control policy for nonlinear control problems is usually quite difficult, especially when the system exhibits low sensitivity of the performance index to a change in the control policy. When the number of time stages is small, then the optimal control policy can usually be readily obtained either by iterative dynamic programming [2,3] or with the LJ optimization procedure [4]. For many problems a small number of flexible stage lengths, where the stage lengths are incorporated into optimization, can be used to give very accurate results by enabling the optimal switching times to be accurately determined. However, there are some problems where the structure of the optimal control policy changes quite substantially when the number of time stages is increased. This is especially true with the system introduced by Yeo [1]. The goal of this paper is to illustrate the interesting features presented by this system and to show how the optimal control policy can be established in a reliable way.

Yeo's optimal control problem

One of the earliest examples tried with iterative dynamic programming [5,6], is the optimal control problem used by Yeo [1] to illustrate the use of quasilinearization to solve nonlinear singular control problems. The system is described by

$$\frac{dx_1}{dt} = x_2 \quad (1)$$

$$\frac{dx_2}{dt} = -x_3 u + 16x_5 - 8 \quad (2)$$

$$\frac{dx_3}{dt} = u \quad (3)$$

$$\frac{dx_4}{dt} = x_1^2 + x_2^2 + 0.0005(x_2 + 16x_5 - 8 - 0.1x_3 u^2)^2 \quad (4)$$

$$\frac{dx_5}{dt} = 1 \quad (5)$$

with $x(0) = [0 \ -1 \ -\sqrt{5} \ 0 \ 0]^T$, and final time $t_f = 1$. The constraints on the control are

$$-4 \leq u \leq 10 \tag{6}$$

It is required to find the control policy that minimizes the performance index

$$I = x_4(t_f) \tag{7}$$

This system was used for global optimization by Rosen and Luus [7] with the introduction of a line search into sequential quadratic programming (SQP). With the use of $P = 80$ time stages of equal length they obtained $I = 0.119258$, and concluded by extrapolation that the minimum value for the performance index for the continuous case (infinite number of stages) should be $I = 0.11923$.

This system was recently investigated by Esposito and Floudas [8], with $P = 10$ time stages, who found two local optima as was found by Luus [3]. The differential equations are well behaved, but to get six figure accuracy for the performance index, it was necessary to use double precision and to use a tight local error tolerance of 10^{-8} in the subroutine DVERK [9].

The optimal value of the performance index depends very much on the number of time stages that are used. The goal here is to illustrate the nature of the optimal solution when the number of time stages of equal length is increased.

Table 1 Affect of the number of time stages P on the value of the performance index after 50 passes, each consisting of 20 iterations

Number of time stages P	Performance index I	CPU time seconds
10	0.120114	5.9
20	0.119438	9.0
40	0.119292	28.9
50	0.119277	47.1
60	0.119272	61.1
80	0.119134	110.5
100	0.119085	180.7

Results and Discussion

In using iterative dynamic programming (IDP), we used the multi-pass approach as outlined by Luus [3, p. 106], using 20 iterations per pass. The region contraction factor was chosen as $\gamma = 0.95$ and the region restoration factor $\eta = 0.85$. A single grid point was used at each time stage, and 3 candidates for control, chosen as $0, \pm r$ were used at each grid point. The initial value for control was chosen as 7.0 and the initial region size $r = 4$ was used. The computational results carried out on a PentiumII/400 personal computer in double precision

after 50 passes for different number of time stages P are given in Table 1.

It is noted that IDP yields the optimum substantially faster than the global optimization procedure of Esposito and Floudas [8]. Even with $P = 100$ the computation time is almost negligible.

The use of $P = 60$ provides a very small refinement to the control policy obtained with $P = 40$ and $P = 50$ for which the optimal control policy is shown in Fig. 1.

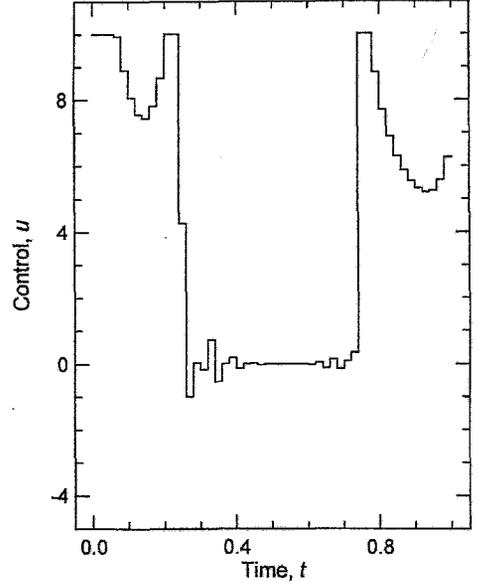


Fig.1 Optimal control policy for $P = 50$ giving $I = 0.119277$

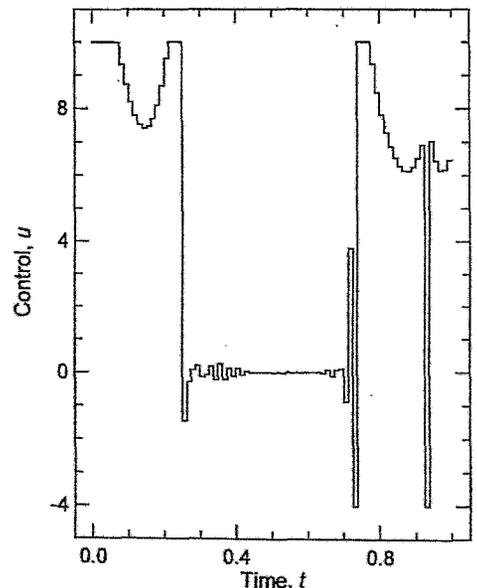


Fig.2 Optimal control policy for $P = 80$, giving $I = 0.1191325$

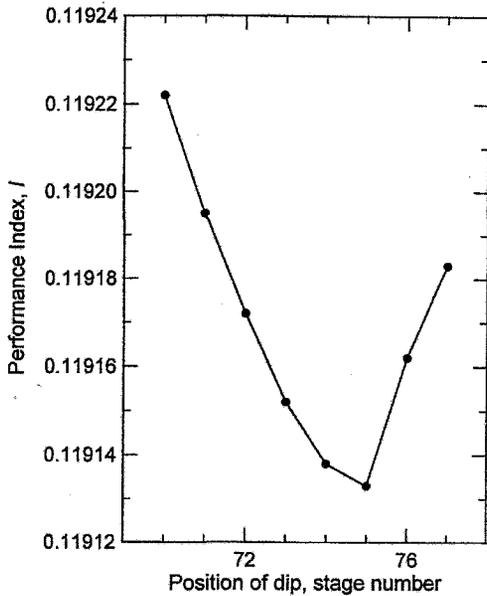


Fig.3 Local optima as a function of the location of the dip in the control policy for $P = 80$

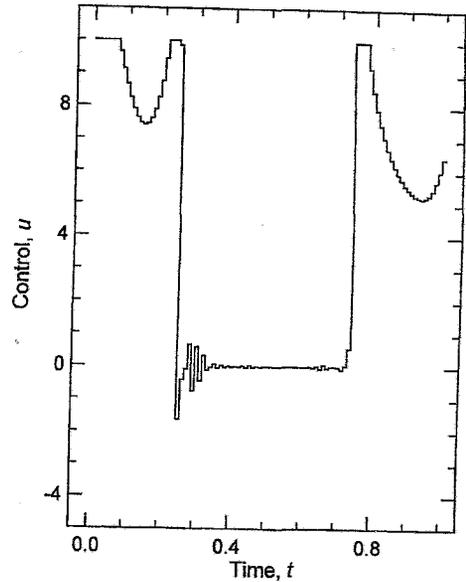


Fig.5 Control policy for local optimum for $P = 100$ giving $I = 0.119250$

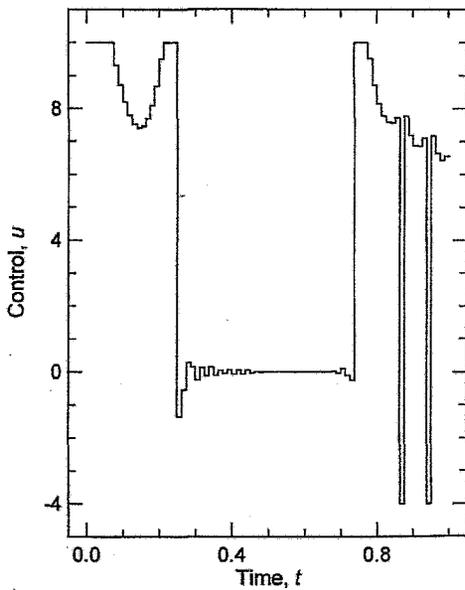


Fig.4 Control policy for a local optimum for $P = 80$, giving $I = 0.1191603$

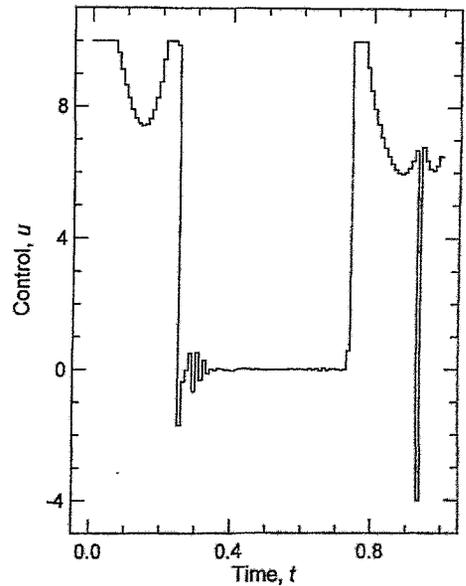


Fig.6 Control policy for local optimum for $P = 100$, giving $I = 0.119092$

Since there is a significant improvement in the performance index for $P = 80$ stages, the run was repeated, using 40 iterations per pass and allowing 100 passes. The performance index was marginally decreased to $I = 0.119133$. As is shown in Fig.2, the use of $P = 80$ yields a totally different optimal control policy. By using a greater number of time stages, the shorter stage lengths allow sudden changes to be made in the control policy. These abrupt changes in the control policy here improve the performance index quite significantly.

There is a dip in the control policy to -4 at stage 59, and in the time interval $(0.8, 1.0)$ there is a dip to -4 at stage 75, with a significant improvement in the resulting performance index over the value of 0.119258 , where there are no dips to the minimum value of control. As is shown in Fig.3, the position of the latter dip affects the performance index. With $P = 80$ we can also get local optima with two dips in the time interval $(0.8, 1.0)$, as is shown in Fig.4, where the dips at stages 70 and 76 give a performance index $I = 0.119160$.

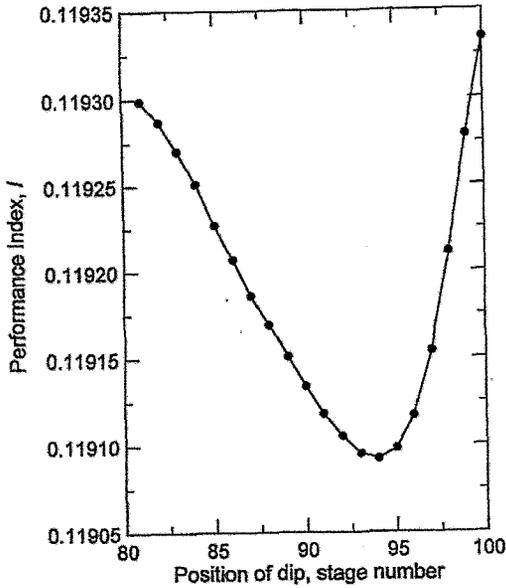


Fig. 7 Local optima as a function of the location of a single dip in the control profile with $P = 100$

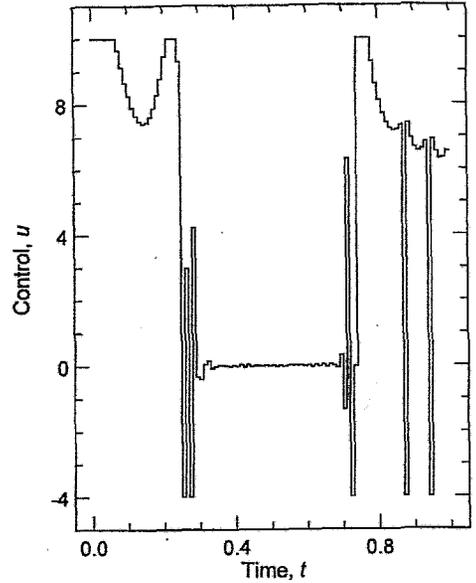


Fig. 9 Optimal control policy for $P = 100$ giving $I = 0.1190437$

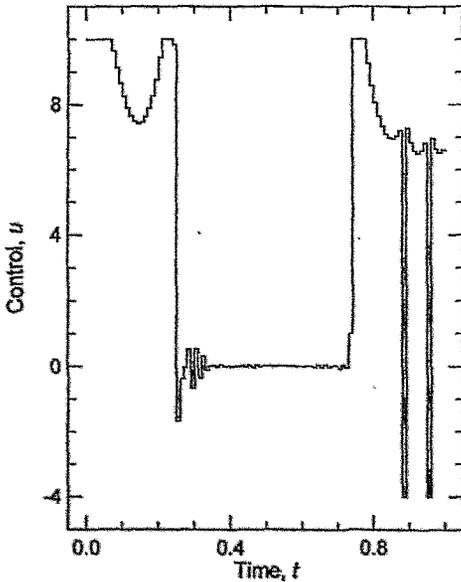


Fig. 8 Control profile for a local optimum with $P = 100$, giving $I = 0.119045$

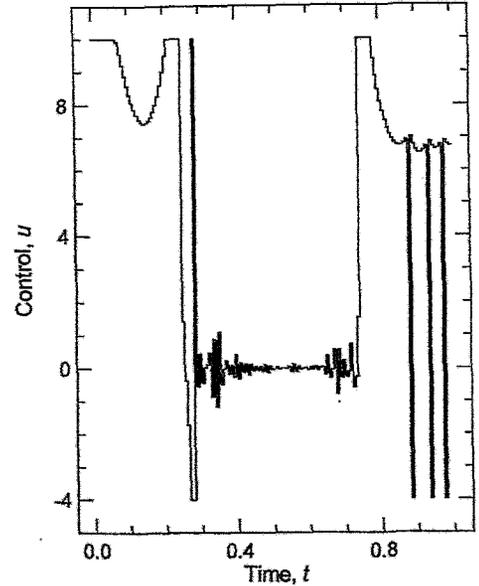


Fig. 10 Control policy for $P = 150$, giving $I = 0.118977$

The global optimum for $P = 80$ with $I = 0.1191325$ occurs with a single dip at stage 75.

With the use of $P = 100$ time stages of equal length the number of local optima is quite large. The local optimum control policy in which there are no dipoles, yielding $I = 0.119250$ is shown in Fig. 5. The best local optimum with a single dip at stage 94, giving $I = 0.119092$ is shown in Fig. 6. As is shown in Fig. 7, the position of the single dip affects the value of the performance index. As can be seen 14 local optima with a single dip give a better value for the performance index than the local optimum without any dipoles.

What makes this problem challenging is the existence of double dipoles in the time interval (0.8, 1.0), many of which give better values for the performance index. In fact, the best local optimum with dipoles at stages 88 and 95, giving $I = 0.119045$, is shown in Fig. 8. Further improvement to $I = 0.1190437$ is obtained with the control policy shown in Fig. 9, where there are additional 3 dipoles in the time interval (0.2, 0.8). There are over 150 local optima for this system with $P = 100$ stages of equal length. Here, with IDP only the most likely candidates for global optimum were scanned to yield the global optimum of $I = 0.119044$.

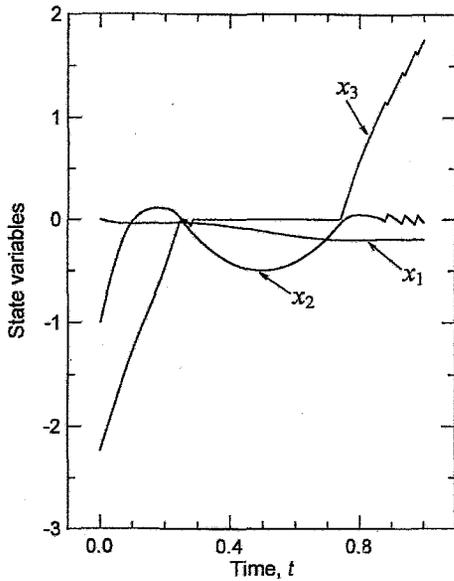


Fig.11 State trajectories corresponding to the control policy in x Figure 10

The question of what happens when the number of time stages is further increased arises. It is obvious that the number of local optima will increase. In fact, a run was performed with $P = 150$, yielding $I = 0.118977$, with the control policy shown in Fig.10. This result was obtained with the conditions used for Table 1, so this result is expected to be close to the global optimum. It is interesting to note that now there are three dips in the time interval (0.8, 1.0). The corresponding trajectories for the first three state variables is given in Fig.11. There is no doubt that further noticeable improvement of the performance index can be obtained by using even a larger number of time stages, such as 200. Alternatively, one could use time stages of varying length to refine the results obtained here. Research in this area is continuing.

Conclusions

By increasing the number of time stages to solve the optimal control problem presented by Yeo [1], some very interesting features about this system were noted. With the use of 80 time stages numerous local optima, that had been observed for the first time, were found. A large number of these local optima were better than the best local optimum that was obtained by SQP. The use of 100 time stages yielded even a larger number of local optima. By scanning numerous local optima, the global optimum with 100 stages of equal length was determined as $I = 0.119044$. This value is 0.2% better than the best value reported in the literature. With the use of 150 stages, the performance index was further improved to 0.118977, showing that further increases in the number of time stages, or the use of stages of varying length should yield even better values for the performance index.

This example illustrates the difficulty in determining with absolute certainty the global optimum of some nonlinear systems. IDP proved to be an effective tool in the establishment of local optima. Through the interpretation of the results obtained for the local optima, the global optimum for 100 time stages was readily obtained. However, the complexity of the problem is realized, so the results obtained here provide only a step toward obtaining the global optimum to the problem for which we could conceptually use an infinite number of time stages.

Acknowledgement

Financial support from the Natural Sciences and Engineering Research Council of Canada is gratefully appreciated.

SYMBOLS

I	performance index to be minimized
P	number of time stages of equal length
r	region size
t	time
t_f	final time
u	control variable
x_i	i^{th} state variable
x	state vector

Greek letters

γ	region contraction factor by which the region size is reduced after every iteration
η	restoration factor by which the region size is restored after every pass

Abbreviations

IDP	iterative dynamic programming
SQP	sequential quadratic programming

REFERENCES

1. YEO B.P.: *Int. J. Control*, 1980, 32, 723-730.
2. LUUS R.: *Control and Intelligent Systems*, 1998, 26, 1-8.
3. LUUS R.: *Iterative Dynamic Programming*, Chapman & Hall CRC, London, UK, 2000.
4. LUUS R. and HENNESSY D.: *Ind. Eng. Chem. Res.*, 1999, 38, 1948-1955.
5. Luus R.: *Hung. J. Ind. Chem.*, 1989, 17, 523-543.
6. Luus R.: *Int. J. Control* 1990, 19, 995-1013.
7. ROSEN O. and LUUS R.: *J. Optimization Theory & Applic.*, 1992, 73, 547-562.
8. ESPOSITO W.R. and FLOUDAS C.A.: *J. Global Optimization*, 2000, 17, 97-126.
9. HULL T.E., ENRIGHT W.D. and JACKSON K.R.: *User Guide to DVERK -- a Subroutine for Solving Nonstiff ODE's*, Report 100, 1976, Department of Computer Science, University of Toronto, Canada