

INITIAL GUESS SENSITIVITY IN COMPUTATIONAL OPTIMAL CONTROL PROBLEMS

JOHN T. BETTS

Applied Mathematical Analysis
24748 SE Mirromont Pl.
Issaquah, WA 98027, USA

STEPHEN L. CAMPBELL* AND CLAIRE DIGIROLAMO

Department of Mathematics
North Carolina State University
Raleigh, NC, 27695-8205, USA

(Communicated by Kok Lay Teo)

ABSTRACT. An optimal control problem is presented that exhibited unexpected initial guess dependence when being solved with direct transcription methods. This note presents that example and the cautionary tale it provides.

1. **Introduction.** A number of successful approaches for numerically solving optimal control problems convert the problem in some manner to a nonlinear programming problem (NLP) and then apply an optimization package to this NLP, and possibly iterate this process several times. Recently while investigating the computation of nonlinear operator bounds using direct transcription software we came across some unusual behavior. The purpose of this note is to give that example and discuss what cautionary tale it has to tell. At the beginning we shall be using the optimal control software GPOPS II which is a pseudospectral direct transcription code [2, 4, 5]. We will use the popular interior point solver IPOPT [6] to solve the NLP problem formulated in GPOPS II. Both of these pieces of software have been used by a number of people to solve a large variety of problems. GPOPS II comes with IPOPT already interfaced.

The specific problem we will discuss here is to

$$\max_w \int_0^T |z|^2 - \gamma^2 |w|^2 dt \quad (1a)$$

$$\dot{x} = -0.1x + w, \quad x(0) = 0 \quad (1b)$$

$$z = x^2 + 0.9w + 1 \quad (1c)$$

$$M^2 = \int_0^T |w|^2 dt. \quad (1d)$$

γ is a parameter and we wanted to solve this problem for a variety of γ . The maximization is over the function w defined on $[0, T]$ with L^2 norm equal to a

2010 *Mathematics Subject Classification.* Primary: 49M37; Secondary: 49K40, 49M25.

Key words and phrases. Optimal control, initial guesses, interior point methods, IPOPT, GPOPS II.

* Corresponding author: Stephen L Campbell.

constant $M > 0$. We changed the max to the min of the negative and got the following problem which was given to GPOPS II.

$$\min \quad -\eta(T) \tag{2a}$$

$$\dot{x} = -.1x + w \tag{2b}$$

$$\dot{\eta} = x^4 + x^2(2 + 1.8w) + w^2(.81 - \gamma^2) + 1.8w + 1 \tag{2c}$$

$$\dot{m} = w^2 \tag{2d}$$

$$x(0) = 0 \tag{2e}$$

$$m(0) = 0 \tag{2f}$$

$$\eta(0) = 0 \tag{2g}$$

$$m(T) = M^2. \tag{2h}$$

However, in solving this problem we noticed unusual behavior concerning the given initial guesses. This problem turns out to have two minima, one local and one global. But the strange thing was that running the exact same code a second time would sometimes produce a different result. Notice that given the constraint on w changing γ does not change mathematically the quantity being optimized. Our initial guess was linear interpolation between endpoint values, a default strategy that is often successful for simpler problems.

Let the state vector be $y = [x, \eta, m]$. The end point conditions for the initial guess are $y_0 = [0, 0, 0]$ and $y_f = [4, 4, M^2]$ for the state and $[w_0, w_f]$ for the control. Table 1 gives the end point conditions for the four initial guesses w_i used for w . The initial guesses for the state were the same on every solution.

i	w_0	w_f
1	1	0
2	-1	-1
3	1	-1
4	0	1

TABLE 1. Endpoints for Initial w guesses w_i for Example 2.

Using values of $\gamma^2 = 4.3, 4.4, 4.5, 4.6, 4.7$, and the four w_i from Table 1 we wrote a MATLAB script that computed the optimal cost and subtracted off the effect of the $\gamma^2 M^2$ term. The initial grid was 10 points and we used 4th order polynomials on each grid interval. The resulting $\eta(T)$ values are given in Table 2.

γ^2	w_1	w_2	w_3	w_4
4.3	18.0613	5.3457	18.0613	18.0613
4.4	18.0613	5.3457	18.0613	5.3457
4.5	18.0613	18.0613	5.3457	5.3457
4.6	18.0613	18.0613	18.0613	18.0613
4.7	5.3457	5.3457	18.0613	5.3457

TABLE 2. Adjusted Values of $\eta(T)$ for Example (2) using DT and $M = 2$, $T = 1$ and 3 iterations for different values of γ^2 and w_i .

Table 2 shows that the code is converging to one of the two extrema each time. However, there is no pattern as to which initial guesses w_i will converge to which

extrema. If we reran this script during the same MATLAB session we would get a similar but not necessarily identical table as shown in Table 3.

γ^2	w_1	w_2	w_3	w_4
4.3	18.0613	5.3457	18.0613	18.0613
4.4	18.0613	5.3457	18.0613	5.3457
4.5	18.0613	18.0613	5.3457	5.3457
4.6	5.3457	18.0613	5.3457	18.0613
4.7	5.3457	5.3457	18.0613	5.3457

TABLE 3. Rerun of the results reported in Table Table 2. Bold entries have changed from Table 2.

What is happening here? With optimal control software it can be very difficult to get good initial guesses. With industrial grade software there is often a whole library of regularizations and perturbations that are designed to hopefully get the optimization started.

However, the situation is even more challenging with interior point methods like IPOPT. The initial guess must be strictly feasible. Thus even a good initial guess is unlikely to be strictly feasible for the discretization [3]. In particular, a discretization of w which is feasible on one grid is not likely to be strictly feasible on the next grid. Thus the software has a number of tweaks to make the first iterate feasible. Sometimes this help with initialization can seem to have a random component as seen with this example. The exact cause of this behavior is not clear.

Problem formulation can play a role. Suppose that instead of (1a) we used the constraint on w and solved

$$\max_w \int_0^T \|z\|^2 dt - \gamma^2 M^2. \quad (3)$$

Then we get the results in Table 4.

γ^2	w_1	w_2	w_3	w_4
4.3	18.0613	18.0613	5.3457	18.0613
4.4	18.0613	18.0613	5.3457	5.3457
4.5	18.0613	18.0613	5.3457	18.0613
4.6	18.0613	18.0613	5.3457	18.0613
4.7	18.0613	18.0613	5.3457	5.3457

TABLE 4. Adjusted Values of $\eta(T)$ for Example (2) using DT and $M = 2$, $T = 1$ and 3 iterations for cost (3).

We see that changing the model formulation now results in two initial guesses converging to the global max independent of γ^2 , one converging to the local max independent of γ^2 , and the third initial guess still giving the different answers depending on γ^2 in a random sort of way. Figure 1 shows the controls that give the two maximums.

As a further examination of this phenomenon, we solved this same problem using the direct transcription software in the Sparse Optimization Suite (SOS). SOS has different discretizations than GPOPS II and a variety of NLP solvers.

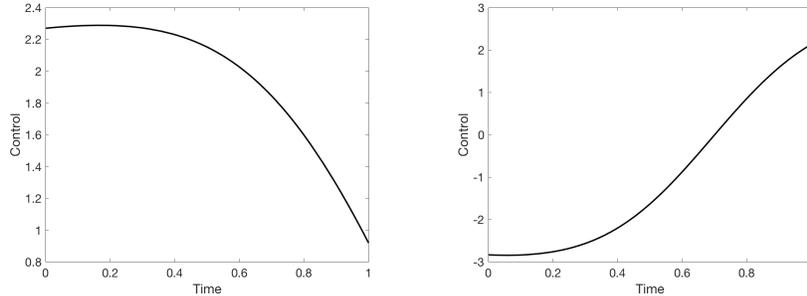


FIGURE 1. Graphs of computed optimal control w for Example (2) with DT approach and $M = 2, T = 1$.

We started with a uniform grid of 10 points and an initial guess of linear interpolation between the guess endpoints. The state guess endpoints were again taken as $[0 \ 0 \ 0]$ and $[4 \ 4 \ M^2]$. The initial discretization was Trapezoid and then switching after a couple of iterations to Hermite Simpson which is the default strategy in SOS. 50 initial guess pairs $[w_0 \ w_f]$ for the control were generated with w_0 and w_f each chosen randomly between -2 and 2. The problem was solved for each initial guess by both an interior point method we shall call IM and a Sequential Quadratic Programming (SQP) method. For every case the answer was one of the adjusted values $\overline{M} = 18.0613$ or $\underline{M} = 5.3457$. For both SQP and IM all but 4 values were \underline{M} . However, IM had \overline{M} at guesses 20, 25, 36, 40, and SQP at 20, 27, 36, 40. Thus we see that there are two local extrema to this problem. Usually the convergence was to \underline{M} but there were two initial guesses where SQP and IM converged to different values and the limits differed on those two guesses.

In contrast with GPOPS II, when running the computations with SOS, the results were repeatable, that is, the same guess always produced the same solution.

This SOS calculation suggests that the \underline{M} is in some sense dominant.

As a final experiment we took the solution for state, control, and time grids found by GPOPS II for a case that gave us the 18.0163 value. We then sampled the state and control on 10 even points using the `interp1` function from MATLAB. We then reran the GPOPS II code a number of times with a number of different γ^2 values. Every time it converged to 18.0163. This shows that even if the other value is more dominant that if we start close enough to the local extrema in this example we again get convergence to that extrema.

2. Conclusion. This example serves as a cautionary tale about the need to carefully examine the answers one computes when solving nonlinear optimal control problems. That different initial choices can lead to different extrema is well known and not surprising. But that the effect of regularization can seem to have a random component in it is not often mentioned. Run the same code with the same guess more than once and the answer can possibly change. Starting from an initial guess, which extrema the software converges to can depend on the solver and discretization and that difference can vary from one initial guess to another.

REFERENCES

- [1] J. T. Betts, *Methods for Optimal Control and Estimation using Nonlinear Programming*, SIAM, Philadelphia, 2010.
- [2] C. L. Darby, W. W. Hager and A. V. Rao, [An hp-adaptive pseudospectral method for solving optimal control problems](#), *Optimal Control Applications and Methods*, **32** (2011), 476–502.
- [3] A. Forsgren, [On warm starts for interior methods](#), in *System Modeling and Optimization*, IFIP International Federation for Information Processing (eds. F. Ceragioli, A. Dontchev, H. Furuta, K. Marti, and L. P. Pandolfi), Springer, Boston, **199** (2006), 51–66.
- [4] M. A. Patterson and A. V. Rao, [GPOPS II: A MATLAB software for solving multiple-phase optimal control problems using hp-adaptive Gaussian quadrature collocation methods and sparse nonlinear programming](#), *ACM Transactions Mathematical Software*, **41** (2014), 1–37.
- [5] A. V. Rao, D. A. Benson, C. Darby, M. A. Patterson, C. Franconin, I. Sanders and G. T. Huntington, [Algorithm 902: Gpops, a MATLAB software for solving multiple-phase optimal control problems using the Gauss pseudospectral method](#), *ACM Transactions Mathematical Software*, **37** (2010), 22:1–22:39.
- [6] A. Wächter and L. T. Biegler, [On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming](#), *Mathematical Programming*, **106** (2006), 25–57.
- [7] A. Wächter and L. T. Biegler, [Failure of global convergence for a class of interior point methods for nonlinear programming](#), *Mathematical Programming*, **88** (2000), 565–574.
- [8] E. A. Yildirim and S. J. Wright, [Warm start strategies in interior-point methods for linear programming](#), *SIAM J. Optimization*, **12** (2002), 782–810.
- [9] E. A. Yildirim, [Implementation of warm-start strategies in interior-point methods for linear programming in fixed dimensions](#), *Computational Optimization and Applications*, **41** (2008), 151–183.

Received July 2018; 1st revision October 2018; Final revision April 2019.

E-mail address: john.t.betts@comcast.net

E-mail address: slc@ncsu.edu

E-mail address: cmdigiro@ncsu.edu