An External Active-Set Strategy for Solving Optimal Control Problems

Elijah Polak
Hoam Chung
S. Shankar Sastry

Electrical Engineering and Computer Sciences
University of California at Berkeley

July 16, 2007

Acknowledgement

An External Active-Set Strategy for Solving Optimal Control Problems

Elijah Polak, Hoam Chung, and Shankar Sastry [*]

Abstract

We present a new, *external*, active constraints set strategy for solving nonlinear programming problems with a large number of inequality constraints that arise in the process of discretizing continuous-time optimal control problems with state-space constraints. This strategy constructs a sequence of inequality constrained nonlinear programming problems, containing a progressively larger subset of the constraints in the original problem, and submits them to a nonlinear programming solver for a fixed number of iterations. We prove that this scheme computes a solution of the original problem and show by means of numerical experiments that this strategy results in reductions in computing time ranging from a factor of 6 to a factor of over 100.

1   Introduction

Optimal control problems with state space constraints are usually solved by discretizing the dynamics, which results in the conversion of the continuous-time optimal control problem into a discrete-time optimal control problem. A discrete-time optimal control problem is, in fact, a nonlinear programming problem and hence can be solved by nonlinear programming algorithms. The distinguishing features of these nonlinear programming problems are (i) the required gradients can be computed using adjoint equations, and (ii) although

---

[*]Authors are with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA, 94720-1770 USA e-mail: polak/hachung/sastry at eecs.berkeley.edu.

they have very large numbers of inequality constraints, relatively few of these inequalities are active.

An important example of optimal control problems with state space constraints arises in the control of unmanned aerial vehicles (UAV's) using receding horizon control (RHC). RHC is a form of sample-data control that determines the control to be applied over the next sampling interval by solving an optimal control problem during the current sampling interval. The optimal control problems for RHC control of UAV's are characterized by large numbers of collision avoidance inequalities and by expensive evaluations of the gradients of these inequality defining functions. Since potential collisions are confined to relatively short segments of the UAV trajectories, most of the collision avoidance inequalities are inactive. In the case of UAV's, the sampling intervals are short and hence the viability of the RHC scheme is largely determined by the speed of the nonlinear programming solvers.

Unfortunately, standard nonlinear programming packages, including the excellent set found in TOMLAB [2], including SNOPT [5], NPSOL [1], Schittkowski SQP [8], and KNITRO[1] [3], are not designed to exploit the fact that a problem with a large number of nonlinear inequality constraints may have few active constraints.

In this paper, we present a new, *external* active constraint set strategy for solving nonlinear programming problems with a large number of inequality constraints. This strategy constructs a sequence of inequality constrained nonlinear programming problems, containing a progressively larger subset of the constraints in the original problem, and submits them to a nonlinear programming solver for a fixed number of iterations. We prove that this scheme computes a solution of the original problem and show by means of numerical experiments that when applied to UAV control, this strategy results in reductions in computing time ranging from a factor of 6 to a factor of over 100. Our new strategy is particularly effective when used with nonlinear programming solvers that allow warm starts. It may be useful to observe that a related strategy [7] for solving semi-infinite minimax problems using log-sum-exponential smoothing has proved to be equally effective.

In Section II we present a motivational optimal control example, in Section III we state our strategy in the form of an algorithm and provide theoretical justification for it, in Section IV we present numerical results, and

---

[1]KNITRO is a collection of optimization algorithms, and we use the algorithm option 'Interior/Direct' with quasi-Newton symmetric rank one updates in this paper.

our concluding remarks are in Section V.

Notation *We will denote elements of a vector by superscripts (e.g., x^i) and elements of a sequence or a set by subscripts (e.g., η_k).* □

2 Optimal Control Example

In order to motivate our approach, consider the following simple, fixed-time optimal control problem, which consists of minimizing the sum of the energy used by a UAV and the square of the distance to a desired destination point. The UAV trajectory is required to stay out of a circle. As we will see, while the number of state space constrains is as large as the number of points used to discretize the dynamics, only those that touch the forbidden circle are active. A geometric representation of the constraints is shown in Fig. 1, which presents the given initial trajectory and the computed optimal trajectory for the example below.

For the sake of simplicity, we assume that the UAV flies at a constant speed v and that the scalar control u determines the yaw rate of the UAV. In order to state the optimal control problem as an end-point problem defined on $[0, 1]$, we rescale the state dynamics using the actual terminal time T and augment the 3-dimensional physical state with a fourth component, x^4, so that

$$x^4(t) = \int_0^t \frac{T}{2} u^2(\tau) d\tau \tag{1}$$

represents the energy used. The resulting dynamics have the form

$$\frac{dx}{dt} = \begin{bmatrix} Tv \cos x^3 \\ Tv \sin x^3 \\ Tu \\ \frac{T}{2} u^2 \end{bmatrix} \triangleq h(x(t), u(t)) \tag{2}$$

with the initial state $x(0)$ given. We will denote the solution of the dynamic equation (2) by $x(t, u)$, with $t \in [0, 1]$. The optimal control problem now assumes the form

$$\min_{u \in \mathbb{R}} f^0(u) \triangleq x^4(1, u) + (x^1(1, u) - 10)^2 + (x^2(1, u) - 10)^2, \tag{3}$$

3
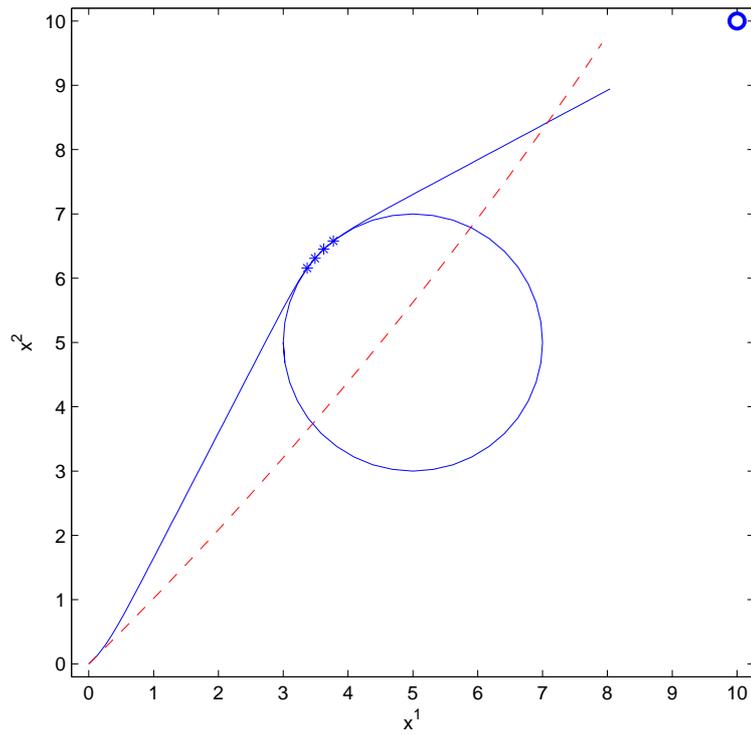
Figure 1: Initial trajectory (dashed red) and optimal trajectory (solid blue).
Active constraints (constraints within feasibility tolerance) are marked as '*'.

subject to:

$$f^t(u) \triangleq (x^1(t,u) - 5)^2 + (x^2(t,u) - 5)^2 \geq 2^2, \ t \in [0,1]. \tag{4}$$

In order to solve this problem, we must discretize the dynamics. We use Euler's method to obtain

$$\bar{x}(t_{k+1}) - \bar{x}(t_k) = \Delta h(\bar{x}(t_k), \bar{u}(t_k)), \ \bar{x}(0) = x(0), \tag{5}$$

with $\Delta \triangleq 1/N$, $N \in \mathbb{N}$, $t_k \triangleq k\Delta$ and $k \in \{0, 1, \ldots, N\}$. We use an overbar to distinguish between the exact variables and the discretized variables. We will denote the solution of the discretized dynamics by $\bar{x}(t_k, \bar{u})$, $k = 0, 1, \ldots, N$. Letting

$$\bar{u} \triangleq (\bar{u}(t_0), \bar{u}(t_1), \ldots, \bar{u}(t_{N-1})), \tag{6}$$

we obtain the discretized optimal control problem

$$\min_{\bar{u} \in \mathbb{R}^N} \bar{f}^0(\bar{u}) \triangleq \bar{x}^4(1, \bar{u}) + (\bar{x}^1(1, \bar{u}) - 10)^2 + (\bar{x}^2(1, \bar{u}) - 10)^2, \tag{7}$$

subject to the constraints

$$\bar{f}^k(\bar{u}) \triangleq -(\bar{x}^1(t_k, \bar{u}) - 5)^2 - (\bar{x}^2(t_k, \bar{u}) - 5)^2 + 2^2 \leq 0, \ k \in \{1, \ldots, N\}. \tag{8}$$

Clearly, (7), (8) is a mathematical programming problem which is distinguished from ordinary mathematical programming problems only by the fact that adjoint equations can be used in the computation of the gradients of the functions $f^k(\cdot)$, $k = 0, 1, \ldots, N$.

3   The Algorithm

We now proceed in a more abstract setting. Consider the inequality constrained minimization problem:
P$_\mathbf{q}$ $\qquad\qquad\qquad \min\{f^0(\eta) \mid f^j(\eta) \leq 0, j \in \mathbf{q}\}, \tag{9}$

where $\eta \in \mathbb{R}^n$, and $\mathbf{q} \triangleq \{1, \ldots, q\}$. We assume that functions $f^j : \mathbb{R}^n \to \mathbb{R}$ are at least once continuously differentiable. In the context of discrete optimal control problems, η can be either a discretized control sequence or an initial state - discretized control sequence pair[2].

---

[2]see [6] Chapter 4 for a detailed treatment of discrete time optimal control problems

Next we define the index set $\mathbf{q}_\epsilon(\eta)$ with $\epsilon > 0$ by

$$\mathbf{q}_\epsilon(\eta) \triangleq \{j \in \mathbf{q} \mid f^j(\eta) \geq \psi_+(\eta) - \epsilon\}, \tag{10}$$

where

$$\psi(\eta) \triangleq \max_{j \in \mathbf{q}} f^j(\eta), \tag{11}$$

and

$$\psi_+(\eta) \triangleq \max\{0, \psi(\eta)\}. \tag{12}$$

Definition. *We say that an algorithm defined by a recursion of the form*

$$\eta_{k+1} = A(\eta_k), \tag{13}$$

for solving inequality constrained problems of the form (9), is convergent *if any accumulation point of a sequence $\{\eta_i\}_{i=0}^\infty$, constructed according to the recursion (13), is a feasible stationary point.* $\qquad\square$

Finally, we assume that we have a convergent algorithm for solving inequality constrained problems of the form (9), represented by the recursion function $A(\cdot)$, i.e., given a point η_k the algorithm constructs its successor η_{k+1} according to the rule (13).

Algorithm 1: Active-Set Algorithm for Inequality Constrained Minimization Problems.

Data: η_0, $\epsilon > 0$, $N_{iter} \in \mathbb{N}$
Step 0: Set $i = 0$, $\mathbf{Q}_0 = \mathbf{q}_\epsilon(\eta_0)$.
Step 1: Set $\zeta_0 = \eta_i$ and perform N_{iter} iterations of the form $\zeta_{k+1} = A(\zeta_k)$ on the problem

$$\mathbf{P}_{\mathbf{Q}_i} \qquad\qquad \min\{f^0(\zeta) \mid f^j(\zeta) \leq 0,\ j \in \mathbf{Q}_i\} \tag{14}$$

to obtain $\zeta_{N_{iter}}$ and set $\eta_{i+1} = \zeta_{N_{iter}}$.
Step 2: Compute $\psi(\eta_{i+1})$.
if $\zeta_{N_{iter}}$ is returned as a global, local, or stationary solution of $\mathbf{P}_{\mathbf{Q}_i}$ and $\psi(\eta_{i+1}) \leq 0$, **then**
   STOP,
else
   Compute

$$\mathbf{Q}_{i+1} = \mathbf{Q}_i \cup \mathbf{q}_\epsilon(\eta_{i+1}), \tag{15}$$

   and set $i = i + 1$, and go to Step 1.
end if

6

Lemma 2. *Suppose that $\epsilon > 0$ and that the sequence $\{\eta_i\}_{i=0}^{\infty}$, in $\mathbb{R}^n$, is such that $\eta_i \to \hat{\eta}$ as $i \to \infty$. Then there exists an i_0 such that for all $i \geq i_0$, $\mathbf{q}_0(\hat{\eta}) \subset \mathbf{q}_\epsilon(\hat{\eta})$.*

Proof. By definition (10), for any $j \in \mathbf{q}_0(\hat{\eta})$,

$$f^j(\hat{\eta}) - \psi_+(\hat{\eta}) \leq 0. \tag{16}$$

First suppose that $\psi_+(\hat{\eta}) = \psi(\hat{\eta}) \geq 0$. Then the set $\mathbf{q}_0(\hat{\eta})$ is nonempty. Since all the functions $f^j(\cdot)$ and $\psi(\cdot)$ are continuous, $[f^j(\eta_i) - \psi_+(\eta_i)] \to [f^j(\hat{\eta}) - \psi_+(\hat{\eta})]$ as $i \to \infty$. Hence there must exist an i_0 such that for all $i \geq i_0$ and $j \in \mathbf{q}_0(\hat{\eta})$,

$$f^j(\hat{\eta}) - \psi_+(\hat{\eta}) \geq -\epsilon, \tag{17}$$

which proves that for all $i \geq i_0$, $\mathbf{q}_0(\hat{\eta}) \subset \mathbf{q}_\epsilon(\hat{\eta})$.

Next suppose that $\psi(\hat{\eta}) < 0$. Then $\psi_+(\hat{\eta}) = 0$ and the set $\mathbf{q}_0(\hat{\eta})$ is empty. Since the empty set is a subset of any set, the desired result follows. $\qquad\square$

Lemma 3. *Suppose that $\epsilon > 0$ and that the sequence $\{\eta_i\}_{i=0}^{\infty}$, in $\mathbb{R}^n$, is such that $\eta_i \to \hat{\eta}$ as $i \to \infty$ and that $\mathbf{Q} = \cup_{i=0}^{\infty} \mathbf{q}_\epsilon(\eta_i) \subset \mathbf{q}$. For any $\eta \in \mathbb{R}^n$, let*

$$\psi_{\mathbf{Q}}(\hat{\eta}) = \max_{j \in \mathbf{Q}} f^j(\eta). \tag{18}$$

If $\psi_{\mathbf{Q}}(\hat{\eta}) \leq 0$, then $\psi(\hat{\eta}) \leq 0$.

Proof. By Lemma 2, $\mathbf{q}_0(\hat{\eta}) \subset \mathbf{Q}$. Since by assumption, $f^j(\hat{\eta}) \leq 0$ for all $j \in \mathbf{Q}$, it follows that $\psi(\hat{\eta}) \leq 0$. $\qquad\square$

Lemma 4. *Suppose that $\mathbf{Q} \subset \mathbf{q}$ and consider the problem*
$\mathbf{P_Q} \qquad\qquad\qquad \min\{f^0(\eta) | f^j(\eta) \leq 0, j \in \mathbf{Q}\}. \tag{19}$

Suppose that $\hat{\eta} \in \mathbb{R}^n$ is feasible for $\mathbf{P_q}$, i.e, $f^j(\eta) \leq 0$ for all $j \in \mathbf{q}$.

(a) *If $\hat{\eta}$ is a global minimizer for $\mathbf{P_Q}$, then it is also a global minimizer for $\mathbf{P_q}$.*

(b) *If $\hat{\eta}$ is a local minimizer for $\mathbf{P_Q}$, then it is also a local minimizer for $\mathbf{P_q}$.*

(c) *If $\hat{\eta}$ is a stationary point for $\mathbf{P_Q}$, i.e., it satisfies the F. John conditions [4] (or Theorem 2.2.4, p. 188 in [6]), then it is also a stationary point for $\mathbf{P_q}$.*

Proof. Clearly, since $\hat{\eta}$ is feasible for $\mathbf{P_q}$ it is also feasible for $\mathbf{P_Q}$.

(a) Suppose that $\hat{\eta}$ is not a global minimizer for $\mathbf{P_q}$. Then there exists an η^* such that $f^j(\eta^*) \le 0$ for all $j \in \mathbf{q}$ and $f^0(\eta^*) < f^0(\hat{\eta})$. Now, η^* is also feasible for $\mathbf{P_Q}$ and hence $\hat{\eta}$ cannot be a global minimizer for $\mathbf{P_Q}$, a contradiction.

(b) Suppose that $\hat{\eta}$ is not a local minimizer for $\mathbf{P_q}$. Then there exists a sequence $\{\eta_i\}_{i=0}^{\infty}$ such that $\eta_i \to \hat{\eta}$, $f^0(\eta_i) < f^0(\hat{\eta})$ and $f^j(\eta_i) \le 0$ for all i and $j \in \mathbf{q}$. But this contradicts the assumption that $\hat{\eta}$ is a local minimizer for $\mathbf{P_Q}$.

(c) Since $\hat{\eta}$ satisfies the F. John conditions for $\mathbf{P_Q}$, there exist multipliers $\mu^0 \ge 0$, $\mu^j \ge 0$, $j \in \mathbf{Q}$, such that $\mu^0 + \sum_{j \in Q} \mu^j = 1$,

$$\mu^0 \nabla f^0(\hat{\eta}) + \sum_{j \in Q} \mu^j \nabla f^j(\hat{\eta}) = 0 \tag{20}$$

and

$$\sum_{j \in Q} \mu^j f^j(\hat{\eta}) = 0. \tag{21}$$

Clearly, $\hat{\eta}$ also satisfies the F. John conditions for $\mathbf{P_q}$ with multipliers $\mu^j = 0$ for all $j \notin \mathbf{Q}$ and otherwise as for $\mathbf{P_Q}$. $\square$

Combining the above lemmas, we get the following convergence result.

Theorem 5. *Suppose that the problem $\mathbf{P_q}$ has feasible solutions, i.e., there exist vectors η^* such that $f^j(\eta^*) \le 0$ for all $j \in \mathbf{q}$.*

(a) *If Algorithm 1 constructs a finite sequence $\{\eta_i\}_{i=0}^k$, exiting in Step 2, with $i + 1 = k$, then η_k is a global, local, or stationary solution for $\mathbf{P_q}$, depending on the exit message from the solver defined by $A(\cdot)$.*

(b) *If $\{\eta_i\}_{i=0}^{\infty}$ is an infinite sequence constructed by Algorithm 1 in solving $\mathbf{P_q}$. Then any accumulation point[3] $\hat{\eta}$ of this sequence is feasible and stationary for $\mathbf{P_q}$.*

Proof. **(a)** If sequence $\{\eta_i\}_{i=0}^k$ is finite, then, by the exit rule, it is feasible for $\mathbf{P_q}$ and it is a global, local, or stationary solution for $\mathbf{P_{Q_i}}$. It now follows from Lemma 4, that it is also a global, local, or stationary solution for $\mathbf{P_q}$.

(b) Since the sets $\mathbf{Q}_i$ grow monotonically, and since $\mathbf{q}$ is finite, there must

---

[3]A point $\hat{\eta}$ is said to be an accumulation point of the sequence $\{\eta_i\}_{i=0}^{\infty}$, if there exists an infinite subsequence, indexed by $K \subset \mathbb{N}$, $\{\eta_i\}_{i \in K}$, such that $\eta_i \xrightarrow{K} \hat{\eta}$. as $i \xrightarrow{K} \infty$

exist an i_0 and a set $\mathbf{Q} \subset \mathbf{q}$, such that $\mathbf{Q}_i = \mathbf{Q}$ for all $i \geq i_0$. Next, it follows from the fact that $A(\cdot)$ is convergent, that for any accumulation point $\hat{\eta}$, $\psi_{\mathbf{Q}}(\hat{\eta}) \leq 0$ and hence, from Lemma 3 that $\psi(\hat{\eta}) \leq 0$, i.e., that $\hat{\eta}$ is a feasible point for $\mathbf{P_q}$. Since for any accumulation point $\hat{\eta}$, by Lemma 2, $\mathbf{q}_0(\hat{\eta}) \subset \mathbf{Q}$, it now follows from the fact that $A(\cdot)$ is convergent and Lemma 4 that any accumulation point $\hat{\eta}$ is stationary for $\mathbf{P_q}$. $\qquad\square$

4   Numerical Results

All numerical experiments were performed using MATLAB V7.2 and TOM-LAB V5.5 [2] in Windows XP, on a desktop computer with an Intel Xeon 3.2GHz processor with 3GB RAM. Optimization solvers tested in this paper were the Schittkowski SQP algorithm with cubic line search [8], NPSOL 5.02 [1], SNOPT 6.2 [5], and KNITRO [3]. It should clear from the form of Algorithm 1, that our strategy benefits considerably from warm starts of the nonlinear programming solvers, to be used after constructing the active set $\mathbf{Q}_i$. Hence it is desirable to use solvers with as extensive a warm start capability as possible, so that one can transmit the last value of important information from the last iteration of a solver on the problem $\mathbf{P_{Q_i}}$ as initial conditions for solving the problem $\mathbf{P_{Q_{i+1}}}$. SNOPT allows the user to provide initial variables and their states and slack variables. NPSOL allows the user to provide initial variables and their states, Lagrange multipliers, as well as an initial Hessian approximation matrix for quasi-Newton updates. conSolve, the TOMLAB implementation of the Schittkowski SQP algorithm, allows the user to provide an initial solution vector and initial Hessian matrix approximation. KNITRO allows the user to provide only the initial solution vector. For maximum efficiency, this data must be saved at the end of the $i-$th run and transmitted as initial data for the $i+1$-th run of the solver.

As we will see from our numerical results, the performance of Algorithm 1 is much more sensitive to the parameters N_{iter} and ϵ when using a solver, such as KNITRO, that has minimal warm start features, than when it is using a solver with extensive warm start features.

4.1   Single UAV Example

Our first set of results are for the optimal control problem (7), (8) described in Section II, with initial state $x_0 = (0, 0, \pi/4, 0)$. All computations were

initialized using the control $\bar{u}_0 = 8 \times 10^{-3}\mathbf{1}_{1 \times N}$. The Parameters used in the numerical experiments were $T = 25$, $v = 0.5$, and $N = 64$, thus there were 64 inequality constraints (8) in the discretized optimal control problem.

The numerical results are summarized in the Tables 1 - 4. In these tables, N_{grad}, the total number of gradient evaluations, and t_{CPU}, the total CPU time for achieving an optimal solution using Algorithm 1, are defined as follows:

$$N_{grad} = \sum_{i=0}^{i_T} |Q_i| \times \text{number of gradient function calls during } i\text{-th inner iteration}$$

$$t_{CPU} = \sum_{i=0}^{i_T} \Big[\text{CPU time spent for } i\text{-th inner iteration}$$

$$+ \text{CPU time spent for setting up } i\text{-th inner iteration} \Big].$$
$$(22)$$

In the above, and in the tables, i_T is the value of the iteration index i at which Algorithm 1 is terminated by the termination tests incorporated in the optimization solver used, and i_{stab} is the value of index i at which $|\mathbf{Q}|$ is stabilized. Also, $\%_{RAW}$, the percentage of t_{CPU} with respect to the computation time with the *raw* algorithm, i.e. using the solver with the full set of constraints (shown in the last row of each table), is used in tables.

Observations.

- For all solvers, there exist parameter pairs of ϵ and N_{iter} which result in more than an 80% reduction in total CPU time, as compared to using these solvers without our active-set strategy.

- When using an optimization solver without an extensive warm start functionality, the performance of Algorithm 1 is much more sensitive to the value of N_{iter} than when using a solver with an extensive warm start functionality. Thus with SNOPT and NPSOL, Algorithm 1 is not very sensitive to the values of N_{iter}. With conSolve, Algorithm 1 yields better performance when $N_{iter} \geq 10$ is given. With KNITRO, Algorithm 1 does not work at all well with $N_{iter} \leq 40$, but it was possible to achieve an 80% reduction in computing time (similar to the other solvers), as compared to the direct use of KNITRO, with proper values of N_{iter} and ϵ. (Table 4).

Table 1: External Active-Set Strategy with Schittkowski SQP Algorithm, Single UAV Example.

Data #	ϵ	N_{iter}	i_T	f^0	N_{grad}	$|\mathbf{Q}|$	i_{stab}	t_{CPU}	$\%_{RAW}$
01	1	10	13	5.0367	7996	27	6	46.432	109.3
02	1	20	5	5.0367	5953	21	3	32.847	77.3
03	1	30	3	5.0367	3112	16	3	17.772	41.8
04	0.1	10	10	5.0367	1478	8	6	11.265	26.5
05	0.1	20	6	5.0367	1757	8	4	12.908	30.4
06	0.1	30	3	5.0367	986	8	3	7.5872	17.9
07	0.01	10	12	5.0367	1114	5	5	10.383	24.4
08	0.01	20	5	5.0367	752	6	4	7.0621	16.6
09	0.01	30	4	5.0367	676	5	4	7.5543	17.8
Raw				5.0367	9600	64		42.500	100

Table 2: External Active-Set Strategy with NPSOL, Single UAV Example.

Data #	ϵ	N_{iter}	i_T	f^0	N_{grad}	$|\mathbf{Q}|$	i_{stab}	t_{CPU}	$\%_{RAW}$
01	1	10	11	5.0367	3214	16	6	18.196	38.0
02	1	20	5	5.0367	2625	16	3	14.211	29.7
03	1	30	3	5.0367	2113	16	2	11.278	23.5
04	0.1	10	12	5.0367	1578	8	7	10.544	22.0
05	0.1	20	6	5.0367	1411	8	4	9.0533	18.9
06	0.1	30	4	5.0367	1465	8	3	9.1705	19.1
07	0.01	10	12	5.0367	915	5	7	7.3035	15.2
08	0.01	20	7	5.0367	920	5	5	7.3169	15.3
09	0.01	30	5	5.0367	808	5	4	6.7290	14.0
Raw				5.0367	11136	64		47.922	100

Table 3: External Active-Set Strategy with SNOPT, Single UAV Example.

Data #	ϵ	N_{iter}	i_T	f^0	N_{grad}	$|\mathbf{Q}|$	i_{stab}	t_{CPU}	$\%_{RAW}$
01	1	10	9	5.0367	2196	16	3	12.671	50.5
02	1	20	6	5.0367	1871	16	2	10.471	41.7
03	1	30	2	5.0367	893	16	2	4.866	19.4
04	0.1	10	7	5.0367	732	8	3	4.9606	19.8
05	0.1	20	6	5.0367	808	8	3	5.4314	21.6
06	0.1	30	2	5.0364	794	7	2	5.1596	20.6
07	0.01	10	9	5.0367	575	5	4	4.6431	18.5
08	0.01	20	6	5.0367	395	5	4	3.4343	13.7
09	0.01	30	5	5.0367	500	5	4	4.2915	17.1
Raw				5.0367	5760	64		25.094	100

Table 4: External Active-Set Strategy with KNITRO, Single UAV Example.

Data #	ϵ	N_{iter}	i_T	f^0	N_{grad}	$|\mathbf{Q}|$	i_{stab}	t_{CPU}	$\%_{RAW}$
01	1	50	60	5.0369	48689	16	2	301.07	1187.2
02	1	75	1	5.0367	770	11	1	5.2344	20.6
03	1	100	1	5.0367	770	11	1	5.1875	20.5
04	0.1	50	2	5.0367	492	7	2	4.7225	18.6
05	0.1	75	2	5.0367	492	7	2	4.7224	18.6
06	0.1	100	2	5.0367	492	7	2	4.7851	18.9
07	0.01	50	4	5.0367	376	5	4	5.6167	22.1
08	0.01	75	4	5.0367	376	5	4	5.5700	22.0
09	0.01	100	4	5.0367	376	5	4	5.5855	22.0
Raw				5.0367	5568	64		25.359	100

4.2 Multi-UAV Example

The next example consists of controlling multiple UAVs. Suppose that we have N_a UAVs, each with the same dynamics as (2), and we want them to stay in a circular region centered at the origin, without incurring any collisions. The stay-in-a-circle constraints are described by the following equations:

$$f_{\text{bnd}}^{t,i}(u^i) \triangleq x^{1,i}(t, u^i)^2 + x^{2,i}(t, u^i)^2 \leq r_{\text{bnd}}^2, \ t \in [0, 1], \tag{23}$$

where $i \in \{1, 2, \ldots, N_a\}$ denotes the UAV index. The collision avoidance constraints are given by

$$f_{\text{ca}}^{t,(i,j)}(u^i, u^j) \triangleq (x^{1,i}(t, u^i) - x^{1,j}(t, u^j))^2 + (x^{2,i}(t, u^i) - x^{2,j}(t, u^j))^2 \geq r_{\text{ca}}^2, \ t \in [0, 1], \tag{24}$$

where (i, j) is an element of the set of all 2-combinations of the index set $\{1, 2, \ldots, N_a\}$. After discretization, as in the previous example, the constraints become

$$\bar{f}_{\text{bnd}}^{k,i}(\bar{u}^i) \triangleq \bar{x}^{1,i}(t, \bar{u}^i)^2 + \bar{x}^{2,i}(t, \bar{u}^i)^2 \leq r_{\text{bnd}}^2, \ k \in \{1, \ldots, N\}, \tag{25}$$

and

$$\bar{f}_{\text{ca}}^{k,(i,j)}(\bar{u}^i, \bar{u}^j) \triangleq (\bar{x}^{1,i}(k, \bar{u}^i) - \bar{x}^{1,j}(k, \bar{u}^j))^2 + (\bar{x}^{2,i}(k, \bar{u}^i) - \bar{x}^{2,j}(k, \bar{u}^j))^2 \\ \geq r_{\text{ca}}^2, \ k \in \{1, \ldots, N\}. \tag{26}$$

Finally, we obtain the following discretized optimal control problem for the multi-UAV example:

$$\min_{\bar{u}^i \in \mathbb{R}^N, \ i \in \{1, \ldots, N_a\}} \bar{f}^0(\bar{u}) \triangleq \sum_{i=1}^{N_a} \bar{x}^{4,i}(1, \bar{u}^i) \tag{27}$$

subject to the dynamics of each UAV and the constraints (25) and (26). The total number of inequality constraints in this problem is $NN_a(N_a - 1)/2 + NN_a$.

The dynamics of the UAV's are nonlinear and the non-collision constraints are non-convex. Hence this problem has many local minima. Consequently, the solution trajectory may depend on the initial control and on the evolution of the active sets during computation.

For numerical experiments, we set $r_{\mathrm{bnd}} = 4$, $r_{\mathrm{ca}} = 1$, $N = 64$ and $N_a = 4$, resulting in 640 nonlinear inequality constraints. The initial conditions and initial controls for each agent are set as

$$
\begin{aligned}
x_0^1 &= (2.5, 2.5, \pi, 0), & \bar{u}_0^1 &= -1.25 \times 10^{-1} \mathbf{1}_{1 \times N} \\
x_0^2 &= (-2.5, 2, -\pi/2, 0), & \bar{u}_0^2 &= 1.25 \times 10^{-1} \mathbf{1}_{1 \times N} \\
x_0^3 &= (-2.5, -2.5, -\pi/4, 0), & \bar{u}_0^3 &= 1.25 \times 10^{-1} \mathbf{1}_{1 \times N} \\
x_0^4 &= (2, -2.5, \pi/2, 0), & \bar{u}_0^4 &= 2.50 \times 10^{-1} \mathbf{1}_{1 \times N}.
\end{aligned}
\tag{28}
$$

Fig. 2 shows a locally optimal solution for this problem. There are only 8 active constraints at the end. These are all associated with staying in the circle; there are no active non-collision constraints. When properly adjusted, Algorithm 1 accumulates fewer than 16 constraints. Consequently, the reduction in the number of gradient computations is huge, and the savings in CPU time is more dramatic than in the single UAV case. There exist parameter pairs with conSolve and KNITRO that achieve more than 95% savings in computation time. As can be seen from our tables, in several cases using NPSOL and SNOPT, our algorithm used less than 1/100 of the CPU time required by NPSOL or SNOPT to solve the example problem directly, i.e., using the full set of constraints.

5   Conclusion

We have presented an external active-set strategy for solving discrete-time optimal control problems with state-space constraints, using nonlinear programming solvers. Our numerical results show that this strategy results in considerable savings in computer time. In our examples, the savings ranged from a factor ranging from 6 to 9, on a problem with 64 constraints, to a factor ranging from around 20 to a factor around of 135 on a problem with 640 constraints. The results depend on the nonlinear programming solver. There is reason to believe that the larger the number of inequalities in the discrete-time optimal control problem, the larger the computational savings will be. This observation is consistent with the two examples presented in this paper. Finally, it should be obvious that one can add refinements to our algorithm, such as restarting it after a certain number of iterations, so as to avoid accumulating constraints that are not close to being active at the solution.
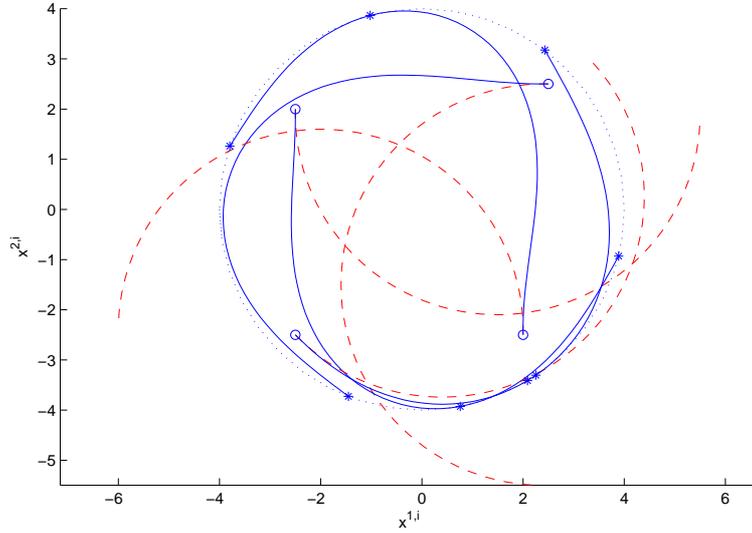
Figure 2: Initial trajectory (dashed red) and an optimal trajectory (solid blue). Bounding circular region is represented by the dotted blue circle. Active constraints (constraints within feasibility tolerance) are marked as '*' and initial positions are marked as 'o'.

Table 5: External Active-Set Strategy with Schittkowski SQP Algorithm, Four-UAV example

Data #	ϵ	N_{iter}	i_T	f^0	N_{grad}	$|\mathbf{Q}|$	i_{stab}	t_{CPU}	$\%_{RAW}$
01	1	10	17	1.3790	27623	122	12	292.40	15.2
02	1	20	16	1.0708	41556	154	16	422.21	22.0
03	1	30	7	1.1585	16626	77	7	180.81	9.4
04	0.1	10	14	1.4203	6444	44	14	90.890	4.7
05	0.1	20	11	0.80255	9463	30	10	129.09	6.7
06	0.1	30	12	1.0970	16393	32	12	228.97	11.9
07	0.01	10	19	1.0708	4477	28	19	83.182	4.3
08	0.01	20	16	1.0708	6318	26	16	115.36	6.0
09	0.01	30	11	0.80255	5810	16	11	109.42	5.7
Raw				3.3564	274560	640		1917.8	100

15

Table 6: External Active-Set Strategy with NPSOL, Four-UAV example

Data #	ϵ	N_{iter}	i_T	f^0	N_{grad}	$|\mathbf{Q}|$	i_{stab}	t_{CPU}	$\%_{RAW}$
01	1	10	12	1.1077	12384	109	12	133.10	1.9
02	1	20	10	1.1585	15089	99	9	158.90	2.3
03	1	30	8	1.1585	17717	93	8	182.61	2.6
04	0.1	10	13	0.80255	4540	34	12	56.871	0.8
05	0.1	20	10	0.80255	6184	34	10	75.456	1.1
06	0.1	30	9	0.80255	7543	30	9	93.732	1.3
07	0.01	10	14	0.80255	1997	13	12	33.211	0.5
08	0.01	20	11	0.80255	2844	15	11	44.847	0.6
09	0.01	30	11	0.80255	3872	16	11	58.878	0.8
Raw				3.3564	1011840	640		6976.1	100

Table 7: External Active-Set Strategy with SNOPT, Four-UAV Example

Data #	ϵ	N_{iter}	i_T	f^0	N_{grad}	$|\mathbf{Q}|$	i_{stab}	t_{CPU}	$\%_{RAW}$
01	1	10	7	1.1585	3212	77	7	36.927	1.6
02	1	20	7	1.1585	3505	77	7	39.880	1.7
03	1	30	7	1.1585	3646	77	7	41.797	1.8
04	0.1	10	9	1.1579	2089	31	9	29.475	1.3
05	0.1	20	9	1.1579	2197	31	9	31.326	1.4
06	0.1	30	9	1.1579	2250	31	9	31.813	1.4
07	0.01	10	11	0.80255	924	16	11	16.806	0.7
08	0.01	20	11	0.80255	965	16	11	17.608	0.8
09	0.01	30	11	0.80255	965	16	11	17.466	0.8
Raw				3.3564	313600	640		2318.4	100

Table 8: External Active-Set Strategy with KNITRO, Four-UAV Example

Data #	ϵ	N_{iter}	i_T	f^0	N_{grad}	$|\mathbf{Q}|$	i_{stab}	t_{CPU}	$\%_{RAW}$
01	1	100	100	*[1]	*	*	*	*	*
02	1	200	100	*	*	*	*	*	*
03	1	300	9	1.1479	80614	90	9	892.13	22.6
04	0.1	100	100	*	*	*	*	*	*
05	0.1	200	10	0.80255	21826	35	10	307.89	7.8
06	0.1	300	10	0.80255	21826	35	10	312.59	7.9
07	0.01	100	17	0.80255	16254	16	12	265.26	6.7
08	0.01	200	11	0.80255	7733	16	11	141.71	3.6
09	0.01	300	11	0.80255	7733	16	11	141.47	3.6
Raw				3.3485	479360	640		3943.5	100

[1] '*' indicates that no meaningful data is available since the algorithm returns without an optimum after 100 iterations.

Acknowledgment

References

[1] Philip E. Gill, Walter Murray, Michael A. Saunders, and Margaret H. Wright. User's guide for NPSOL 5.0: A fortran package for nonlinear programming. Technical Report SOL 86-2, Systems Optimization Laboratory, Department of Operations Research, Stanford University, 1998.

[2] Kenneth Holmström, Anders O. Göran, and Marcus M. Edvall. *User's Guide for TOMLAB*. Tomlab Optimization Inc., December 2006.

[3] Kenneth Holmström, Anders O. Göran, and Marcus M. Edvall. *User's Guide for TOMLAB/KNITRO v5.1.* Tomlab Optimization Inc., April 2007.

[4] F. John. Extremum problems with inequlaities as side condtions. In K. O. Friedrichs, O. W. Neugebauer, and J. J. Stoker, editors, *Studies and Essays: Courant Anniversary Volume*, pages 187–204. Interscience Publishers, Inc., New York, 1948.

[5] W. Murray, P. E. Gill, and M. A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Journal on Optimization*, 12:979–1006, 2002.

[6] E. Polak. *Optimization: Algorithms and Consistent Approximations*, volume 124 of *Applied Mathematical Sciences*. Springer, 1997.

[7] E. Polak, R. S. Womersley, and H. X. Yin. An algorithm based on active sets and smoothing for discretized semi-infinite minimax problems. *Journal of Optimization Theory and Applications*, 2007. in press.

[8] K. Schittkowski. On the convergence of a sequential quadratic programming method with an augmented lagrangian line search function. Technical report, Systems Optimization laboratory, Stanford University, 1982.