Title

Consistent numerical methods for state and control constrained trajectory optimisation with parameter dependency

Journal

International Journal of Control, 94(9)

ISSN

0020-7179

Authors

Walton, Claire
Kaminer, Isaac
Gong, Qi

Publication Date

2021-09-02

DOI

10.1080/00207179.2020.1717633

Peer reviewed

Consistent Numerical Methods for State and Control Constrained Trajectory Optimization with Parameter Dependency

ARTICLE HISTORY
Compiled December 9, 2019

ABSTRACT
This paper describes and proves the consistency of a flexible numerical method for producing solutions to state and control constrained control problems with parameter dependencies. This method allows for the use of a variety of underlying discretization schemes, which can be catered to differing numerical challenges of specific problems, such as rapid convergence or large parameter spaces. The paper first provides a broad formulation for optimal control problems with parameter dependencies which includes multiple types of state, control, and end time constraints to enable a wide scope of application. For this formulation, the consistency of these methods for state and control constrained problems is then proved. Finally, a numerical example of an optimal search problem with constraints is demonstrated.

KEYWORDS
optimal control, nonlinear control, numerical methods, parameter dependency, parameter uncertainty, distributed-parameter systems, ensemble control, optimal search

1. Introduction

Instances of optimal control problems with parameter dependency have arisen in multiple recent control applications. In these applications, performance is optimized over a set of parameters which impact cost and possibly system dynamics given access to the system only through a parameter independent control. For instance, the ensemble control problem deals with the control of a family or continuum of systems $x(t, \theta)$ whose dynamics

$$\frac{\partial x}{\partial t} = f(x(t, \theta), u(t), \theta)$$

depend continuously on parameter $\theta \in \mathbb{R}^{n_\theta}$ and are driven by parameter independent control $u(t)$. Optimizing the behavior of the 'ensemble' of systems over all parameter values creates control problems which must track the behavior over a range of parameter values and incorporate this range of performances in its cost metric. Ensemble control problems arise in quantum applications such as medical imaging or quantum experiments, where an external application of electromagnetic pulses is applied to a system with environmental inhomogeneity [Brockett and Khaneja (2000); Li and Khaneja (2007); Ruths and Li (2012)]. Parameter dependency can also present itself in spatially distributed problems with no spatial feedback control. For example, the optimal search problem [Foraker (2011); Phelps, Gong, Royset, Walton, and Kaminer

(2014)], which aims to optimize the search for uncertain targets over a search region. When uncertain target location is modeled as a deterministic quantity conditionally dependent on a set of unknown parameter values, such as initial location, then expected probability of success becomes dependent on the parameter set. This creates a cost function in the form of an integral over parameter space:

$$J[x,u] = \int_\Theta \Big[F\left(x(T,\theta),\theta\right) + \int_0^T r(x(t,\theta),u(t),t,\theta)dt \Big] d\theta$$

where $\int_\Theta$ represents the possibly multidimensional integral over parameter space. Additional examples include robotics, where manufacturing error can create parameter uncertainty in regards to vehicle part sizes and subsequent dynamics [Becker (2012)] and in chemical engineering, where the model of parameter dependency has been utilized for the optimization of batch processes under uncertainty [Ruppen, Benthack, and Bonvin (1995); Terwiesch, Ravemark, Schenker, and Rippin (1998)].

A key issue in addressing optimal control problems with parameter dependencies, which tend to be analytically intractable, has been the development of numerical algorithms. Recently, there has been much progress in this area. For instance, for nonlinear finite-dimensional optimization problems with parameter uncertainty, Robust Optimization (RO) frameworks have been developed to address the minimization of mean performance given constraints on variance or other risk metrics, such as in [Darlington, Pantelides, Rustem, and Tanyi (2000)]. In continuous time optimal control, the method of polynomial chaos has been applied to a variety of problems with amenable problem structures, such as quadratic costs or linear dynamics [Fisher and Bhattacharya (2011); Hover and Triantafyllou (2006)].

As an approach to general nonlinear control problems, there have been many variations on forms of direct discretization put forth. This can be summarized informally as first choosing a set of nodes $\{\theta_i^M\}_{i=1}^M$ from the parameter domain Θ and an associated set of integration weights $\{\alpha_i^M\}_{i=1}^M$. States, if parameter dependent, are then propagated at these nodes:

$$\frac{dx_i^M}{dt}(t) = f(x_i^M(t), u(t), \theta_i^M), \quad i = 1, \ldots, M$$

and the cost function is estimated as a weighted sum:

$$\int_\Theta \Big[F\left(x(T,\theta),\theta\right) + \int_0^T r(x(t,\theta),u(t),t,\theta)dt \Big] d\theta \approx$$

$$\sum_{i=1}^M \alpha_i^M \Big[F\left(x_i^M(T),\theta_i^M\right) + \int_0^T r(x_i^M(t),u(t),t,\theta_i^M)dt \Big].$$

This creates an approximate problem—a standard control problem which can be solved using other established methods. There are many approaches to node selection, with distinct tradeoffs. Monte Carlo sampling, for instance, provides slow converge but is a way to circumvent the dimensional growth problem of a high dimension parameter space [Phelps, Royset, and Gong (2016)]. Conversely, the multi-dimensional pseudospectral approach can provide rapid convergence but node selection grows exponen-

tially with dimension [Ruths and Li (2012)]. Other options include Riemann-Stieltjes integration [Ross, Proulx, Karpenko, and Gong (2015)] and using the sigma points from the unscented transform [Ross, Karpenko, and Proulx (2016)]. A more generic approach is that of [Phelps et al. (2014)] and [Walton, Phelps, Gong, and Kaminer (2016)], which considers all node and weight schemes which provide convergent quadrature for continuous functions.

Approach	multi-dimensional pseudospectral	Monte Carlo sampling	Riemann-Stieltjes	unscented transform	collocation with convergent quadrature		
Reference	(i)	(ii)	(iii)	(iv)	(v)	(vi)	this paper
Constraints:							
control	Yes	Yes	No	No	Yes	Yes	Yes
state-control	Yes	No	No	Yes	No	No	Yes
end time state	Yes	No	No	Yes	No	No	Yes
distributional state-control	No	No	No	No	No	No	Yes
distributional end time state	No	No	Yes	No	No	No	Yes
uncertain dynamics	Yes	Yes	Yes	Yes	No	Yes	Yes
consistency proof	Yes	Yes	No	No	Yes	Yes	Yes

Reference	
(i)	Ruths and Li (2012)
(ii)	Phelps et al. (2016)
(iii)	Ross et al. (2015)
(iv)	Ross et al. (2016)
(v)	Phelps et al. (2014)
(vi)	Walton et al. (2016)

Table 1. Overview for Direct Discretization

The variety of node generation approaches in the literature have been applied to diverse formulations of this problem. Differences include whether state dynamics are influenced by the uncertain parameter, whether control and/or state constraints are incorporated, and whether any convergence results are proven. Table 1 provides an overview of results for direct discretization methods applied to finite horizon problems with uncertain parameters. The goal of this paper is to provide a method which addresses the full variety of constraint needs and to provide the subsequent proof of consistency (that convergent numerical solutions of approximate problems converge to optimal solutions of the original problem) of the numerical method. An additional goal is to provide this for a broad family of numerical methods, by building on the convergent quadrature approach. This method is inclusive of both multi-dimensional pseudospectral and Riemann-Stieltjes, as well methods that can mitigate the dimensional growth problems of quadrature, such as sparse grid methods.

We address the variety of parameter dependency forms and constraint needs through the following class of optimal control problems:

Problem P. *Determine the state and control pair,* (x, u)*, that minimizes the cost*

function:

$$J[x,u] = \int_\Theta \Big[F\left(x(T,\theta),\theta\right) + \int_0^T r(x(t,\theta),u(t),t,\theta)dt \Big] d\theta \tag{1}$$

subject to:

$$\frac{\partial x}{\partial t}(t,\theta) \;=\; f(x(t,\theta),u(t),\theta), \quad t \in [0,T], \quad \theta \in \Theta, \tag{2}$$

$$x(0,\theta) \;= x_0(\theta), \qquad \theta \in \Theta, \tag{3}$$

$$g(u(t)) \;\leq\; 0, \qquad t \in [0,T], \tag{4}$$

$$e(x(T,\theta),\theta) \;\leq\; 0, \qquad \theta \in \Theta, \tag{5}$$

$$h(x(t,\theta),\theta) \;\leq\; 0, \quad t \in [0,T], \quad \theta \in \Theta, \tag{6}$$

$$\Psi_e \left(\int_\Theta e_I\left(x(T,\theta),\theta\right) d\theta \right) \leq 0, \tag{7}$$

$$\Psi_h \left(\int_\Theta h_I\left(x(t,\theta),\theta\right) d\theta \right) \leq 0, \quad t \in [0,T]. \tag{8}$$

Component functions have dimensions: $x : \mathbb{R} \times \mathbb{R}^{n_\theta} \mapsto \mathbb{R}^{n_x}$, $u : \mathbb{R} \mapsto \mathbb{R}^{n_u}$, $F : \mathbb{R}^{n_x} \times \mathbb{R}^{n_\theta} \mapsto \mathbb{R}$, $r : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R} \times \mathbb{R}^{n_\theta} \mapsto \mathbb{R}$, $x_0 : \mathbb{R}^{n_\theta} \mapsto \mathbb{R}^{n_x}$, $e : \mathbb{R}^{n_x} \times \mathbb{R}^{n_\theta} \mapsto \mathbb{R}^{n_e}$, $g : \mathbb{R}^{n_u} \mapsto \mathbb{R}^{n_g}$, $h : \mathbb{R}^{n_x} \times \mathbb{R}^{n_\theta} \mapsto \mathbb{R}^{n_h}$, $e_I : \mathbb{R}^{n_x} \times \mathbb{R}^{n_\theta} \mapsto \mathbb{R}$, $\Psi_e : \mathbb{R} \mapsto \mathbb{R}$, $h_I : \mathbb{R}^{n_x} \times \mathbb{R}^{n_\theta} \mapsto \mathbb{R}$, $\Psi_h : \mathbb{R} \mapsto \mathbb{R}$. Additional conditions imposed on the state and control space and component functions for the numerical method are specified in Section 2.

In Problem **P**, the set Θ is the domain of the parameters $\theta \in \mathbb{R}^{n_\theta}$. The format of the cost function J is that of the integral over Θ of a Mayer-Bolza type cost dependent on the uncertain parameters. In regards to state constraints, the nature of the problem, in which a single control input is applied over multiple values of parameters θ, creates a variety of possible boundary conditions and constraints which can arise in applications. In this formulation, constraints over all parameter values, in the form of [Li and Khaneja (2006, 2007, 2009); Ross, Proulx, and Karpenko (2014b)] are included with conditions (5) and (6), as are the control constraints found in [Phelps et al. (2014)] and [Walton et al. (2016)] with condition (4). State constraints bounded over all parameter values, however, are rather strict, and guaranteed feasibility of a problems with such constraints has not been established with great generality for control systems with parameter uncertainty. Due to these limitations and also the strictness of conditions which constrain all parameter values in cases where low probability parameter values may not be as critical to restrict, applications may also consider constraints on aggregate performance, as suggested in [Ross, Proulx, and Karpenko (2014a)]. The

constraints on the aggregate performance found in [Ross et al. (2014a)] are generalized here to conditions (7) and (8). These conditions allow for limits to be set on quantities such as the expectation or variance of states over a distribution of parameter values.

As with all control formulations, the feasibility of individual problems is a difficult issue and not necessarily guaranteed. However, there are plenty of available examples which demonstrate the feasibility of parameter uncertainty problems under diverse constraints. For example, the constraint $e(x(T,\theta),\theta) \leq 0$ with

$$e(x(T,\theta),\theta) = \begin{bmatrix} x(T,\theta) - \gamma(\theta) - \epsilon \\ \gamma(\theta) - x(T,\theta) - \epsilon \end{bmatrix}, \tag{9}$$

limits the end time states $x(T,\theta)$ to be within $[-\epsilon,\epsilon]$ of an end time goal curve $\gamma(\theta)$. In [Li and Khaneja (2006, 2007, 2009)], feasibility results for a class of linear parameter uncertainty systems are derived, in the sense that for all ϵ there exists a final time T such that the systems can be driven to within an ϵ-ball of a goal curve $\gamma(\theta)$. (The case of nonlinear systems with parameter uncertainty, however, is an open question, and in fact it is easy to construct systems which will be unable to satisfy such a constraint. For instance, [Becker (2012)] provides a scaled nonholonomic unicycle system which is provably unable to satisfy an end constraint on orientation over all parameters.) A path constraint example can be found in [Ross et al. (2016)], which uses such a constraint of type (6) to avoid singular gimbal trajectories for an agile spacecraft. The spatially distributed application of optimal search [Pursiheimo (1976)] also provides many intuitive instances of feasible constrained problems—attainable search performance constraints, for instance. We provide such an example in Section 4.

For this general Problem **P**, we provide a discretization-based numerical algorithm, and proofs of its feasibility (that it generates solutions which are feasible for Problem **P**) and consistency (that convergent numerical solutions converge to optimal solutions of Problem **P**). These results build on the work of [Phelps et al. (2014)] and [Walton et al. (2016)] to extend to the broader problem formulation presented here. The structure of this paper is as follows. Section 2 provides the additional regularity conditions imposed on Problem **P**. Section 3 presents a numerical method for generating solutions and a proof of the consistency of this method. Finally, Section 4 provides an example with numerical solution.

2. Regularity Assumptions

In order to address Problem **P** numerically and analytically in the next sections, we impose the following regularity assumptions:

Assumption 2.1. *The functions f, r, F, e, h, e_I, h_I, Ψ_e, and Ψ_h are C^1, and $x_0 : \Theta \mapsto \mathbb{R}^{n_x}$ is continuous.*

Due to their compact domains, the C^1 functions thus satisfy Lipschitz conditions.

Assumption 2.2. *Control solutions u are C^0 on $[0,T]$.*

From Assumption 2.1, this implies state solutions are C^1 on $\Theta \times [0,T]$.

Assumption 2.3. *The function g is continuous and the set $U = \{\nu \in \mathbb{R}^{n_u} | g(\nu) \leq 0\}$ is compact.*

Assumption 2.4. *The set Θ is compact, and there exists a compact set $X \subset \mathbb{R}^{n_x}$ such that for all $u(t) \in U$ and $\theta \in \Theta$, $x(t,\theta) \in X$ for all $t \in [0,T]$, where $x(t,\theta) = x_0 + \int_0^t f(x(s,\theta),u(s),\theta)ds$.*

This assumption essentially requires for all bounded controls that there is no $\theta \in \Theta$ for which the state has a finite escape time. A large class of nonlinear systems satisfy this assumption, for example, input-to-state stable systems and systems for which f is globally Lipschitz or satisfies a linear growth condition.

3. Numerical Approximation Scheme

We introduce an approximation of Problem **P**, referred to as Problem $\mathbf{P^M}$. Problem $\mathbf{P^M}$ is created by approximating the parameter space, Θ, with a numerical integration scheme which is defined in terms of a finite set of M nodes $\{\theta_i^M\}_{i=1}^M$ and an associated set of M weights $\{\alpha_i^M\}_{i=1}^M \subset \mathbb{R}$. Throughout the paper, M is used to denote the number of nodes used in this approximation of parameter space. The only requirement on the numerical integration scheme is that it satisfies the following assumptions:

Assumption 3.1. *For each $M \in \mathbb{N}$, the integration scheme is defined by a set of nodes $\{\theta_i^M\}_{i=1}^M \subset \Theta$ and an associated set of weights $\{\alpha_i^M\}_{i=1}^M \subset \mathbb{R}$ such that for any continuous function $f : \Theta \to \mathbb{R}$,*

$$\int_\Theta f(\theta)d\theta = \lim_{M\to\infty} \sum_{i=1}^M f(\theta_i^M)\alpha_i^M.$$

Remark 1. For a series of functions $\{f_M\}$, note that if $f_M : \Theta \to \mathbb{R}$ is continuous for all $M \in \mathbb{N}$ and $\{f_M\}$ converges uniformly to f on Θ with respect to the Euclidean norm, then the following also holds as a result of Assumption 3.1:

$$\int_\Theta f(\theta)d\theta = \lim_{M\to\infty} \sum_{i=1}^M f_M(\theta_i^M)\alpha_i^M.$$

This property is used later.

Remark 2. For a given function $h : \Theta \to \mathbb{R}$, $h \in C^1$ and given integration scheme satisfying Assumption 3.1, there exists a constant $K \in \mathbb{R}$ and function $\epsilon(M) : \mathbb{N} \to R$ such that $\lim_{M\to\infty} \epsilon(M) = 0$ and the numerical integration error is bounded by:

$$\Big| \int_\Theta f(\theta)d\theta - \sum_{i=1}^M f(\theta_i^M)\alpha_i^M \Big| \leq K\epsilon(M). \tag{10}$$

We note that Θ has been defined as a compact domain. Many numerical integration methods satisfy the convergence requirement of Assumption 3.1 on compact domain, for instance Gaussian quadrature, composite-Simpson, and Clenshaw-Curtis. Remark 2 will be used to quantify slack bounds for the integral constraints of Problem **P** given by equations (7) and (8), whose integrated functions h_I and e_I have been defined as satisfying the C^1 requirement of Remark 2. Although numerical integration errors are often given in terms of higher order derivatives, bounds satisfying Remark 2

can be provided given the compact domain and subsequently bounded first derivative of the function. In [Trefethen (2008)], for instance, such bounds are given for Gaussian quadrature and Clenshaw-Curtis in terms of the bounded first derivative of the function. Reference [Davis and Rabinowitz (2007)] provides such bounds for quadrature schemes which are exact for polynomials of some degree. Alternate discretization methods with lower dimensional growth rates or the computation of problems with high parameter space dimension, in which the full tensor product required by multi-dimensional collocation may be intractable, may also be chosen, for instance Smolyak sparse grid methods satisfy both assumptions [Gerstner and Griebel (1998)].

For a given set of nodes $\{\theta_i^M\}_{i=1}^M$ and quadrature weights $\{\alpha_i^M\}_{i=1}^M \subset \mathbb{R}$, the approximate problem is defined as follows:

Problem $\mathbf{P^M}$. *Determine the control function u that minimizes*

$$J^M[X^M, u] = \sum_{i=1}^{M} \left[F\left(x_i^M(T), \theta_i^M\right) + \int_0^T r(x_i^M(t), u(t), t, \theta_i^M)dt \right] \alpha_i^M \tag{11}$$

subject to:

$$\frac{dx_i^M}{dt}(t) = f(x_i^M(t), u(t), \theta_i^M), t \in [0,T], i = 1, \ldots, M, \tag{12}$$

$$x_i^M(0) = x_0(\theta_i^M), \quad i = 1, \ldots, M, \tag{13}$$

$$g(u(t)) \leq 0, \quad t \in [0,T], \tag{14}$$

$$e(x_i^M(T), \theta_i^M) \leq 0, \quad i = 1, \ldots, M, \tag{15}$$

$$h(x_i^M(t), \theta_i^M) \leq 0, \quad t \in [0,T], i = 1, \ldots, M, \tag{16}$$

$$\Psi_e \left(\sum_{i=1}^{M} e_I\left(x_i^M(T), \theta_i^M\right) \alpha_i^M \right) \leq K_e \epsilon_e(M), \tag{17}$$

$$\Psi_h \left(\sum_{i=1}^{M} h_I\left(x_i^M(t), \theta_i^M\right) \alpha_i^M \right) \leq K_h \epsilon_h(M), \quad t \in [0,T]. \tag{18}$$

This approximate problem is a result of the enforcement of the state dynamics and conditions (3)-(8) at the collocation nodes $\{\theta_i^M\}_{i=1}^M$ and the enforcement at these nodes of an approximation of the integral constraints (7) and (8) through quadrature. Due to the discretization of parameter space, the state space for Problem $\mathbf{P^M}$ is of a different dimension than that of Problem $\mathbf{P}$. The state variables x_i^M, $i = 1, \ldots, M$ specified by the ODEs in (12) are functions of time rather than time and parameter space, and the dimension of the entire state space, $X^M(t) = [x_1^M(t), \ldots, x_M^M(t)]$, has dimension $n_x \times M$, where n_x is the dimension of the original state space. The dimension of the control space, however, remains the same.

We take a moment to discuss the slack constraints of equations (17) and (18) versus the exact constraints which precede them. In contrast to direct discretization methods in the time domain which must introduce slack constraints into all state constraints, such as the pseudospectral method for standard control, Problem $\mathbf{P^M}$ is only approximating the additional parameter space. Constraints which are defined pointwise with respect to this parameter space, such as the dynamics, can be enforced exactly. This does not mean that the original problem, Problem $\mathbf{P}$, is necessarily feasible without its own slack constraints. Recall that equation (9) in Section 1 for instance, provides an example of an end time constraint which is only provably feasible within an ϵ-ball. This 'slack' constraint of the original problem is first defined in terms of its own ϵ-slacks, which are in turn enforced strictly via equation (15). Equations (17) and (18) in contrast are made approximate from the approximation of integration. The constants K_e and K_h come from Remark 2 combined with the Lipschitz constants of Ψ_e and Ψ_h respectively. For equation (18), the constant from Remark 2 must be determined using the max over the time interval $[0, T]$.

The resulting Problem $\mathbf{P^M}$ is a standard optimal control problem, no longer distinguished by an additional parameter domain, for which there are many available methods of solution. The following theorem establishes the property of consistency for solutions of Problem $\mathbf{P^M}$ in regards to Problem $\mathbf{P}$. Consistency is not a proof of the feasibility of the original Problem $\mathbf{P}$—that is a design task. Rather, this proof addresses the reliability of the numerical solutions for a feasible problem. The proof thus assumes Problem $\mathbf{P}$ is feasible. From that it follows that Problem $\mathbf{P}^M$ is feasible, since any discretization of Problem $\mathbf{P}$ at the nodes will generate a feasible solution to Problem $\mathbf{P}^M$. The question the proof then addresses, consistency, is the property that if optimal controls to Problem $\mathbf{P^M}$ converge as the number of nodes $M \to \infty$, they converge to feasible, optimal controls of Problem $\mathbf{P}$. Convergence is referred to in the following sense:

Definition 3.2. Uniform Accumulation Point - A function f is called a uniform accumulation point of the sequence of functions $\{f_n\}_{n=0}^{\infty}$ if $\exists$ a subsequence of $\{f_n\}_{n=0}^{\infty}$ that uniformly converges to f. Similarly, a vector $v \in \mathbb{R}^M$ is called a uniform accumulation point of the sequence of vectors $\{v_n\}_{n=0}^{\infty}$ if $\exists$ a subsequence of $\{v_n\}_{n=0}^{\infty}$ that converges to v.

Note that Definition 3.2 applies to limits of subsequences, not the limits of the sequence in entirety. To express this we adopt the following notation. Let V be an infinite subset of the index set $\{0, 1, 2, \dots\}$. If for a sequence $\{x_n\}_{N=0}^{\infty}$, the subsequence $\{x_n | n \in V\}$ has a limit point x, we will refer to this with the notation $\lim_{n \in V} x_n = x$. For uniform accumulation points of the controls of Problem $\mathbf{P^M}$ as $M \to \infty$, the following holds:

Theorem 3.3. Control Consistency: *Let $\{u^{M*}\}_{M \in V}$ be a sequence of optimal controls for Problem $\mathbf{P^M}$ with an accumulation point u^{∞}. Given Assumptions 2.2 - 2.4 , then u^{∞} is an optimal control for Problem $\mathbf{P}$.*

3.1. Theorem 3.3 Proof

In the following proof, and all subsequent proofs in this paper, $\| \cdot \|$ refers to the Euclidean norm. Let $\{u^{M*}\}_{M \in V}$ be a set of optimal controls for the Problem P^M such

that $\lim_{M\in V}\{u^{M*}\} = u^{\infty}$. Let $x^{\infty}(t,\theta)$ be the solution to the dynamical system:

$$\frac{\partial x^{\infty}}{\partial t}(t,\theta) = f(x^{\infty}(t,\theta), u^{\infty}(t), \theta),\ x^{\infty}(0,\theta) = x_0(\theta) \tag{19}$$

and let $\{x^{M*}(t,\theta)\}$ be the sequence of solutions to the dynamical systems:

$$\frac{\partial x^{M*}}{\partial t}(t,\theta) = f(x^{M*}(t,\theta), u^{M*}(t), \theta), x^{M*}(0,\theta) = x_0(\theta) \tag{20}$$

for $M \in V$. We note, first of all, that since Problem P and Problem P^M share the same feasible control set U, the solutions to equations (19) and (20) exist, due to Assumption 2.4. We note, second of all, that these solutions are *not* necessarily feasible solutions to Problem P, as they may violate the state constraints. We will revisit this issue after the following lemma.

Lemma 3.4. *The sequence $\{x^{M*}(t,\theta)\}$ converges pointwise to $x^{\infty}(t,\theta)$ and this convergence is uniform in θ.*

Proof: From their definitions, we have:

$$\|x^{M*}(t,\theta) - x^{\infty}(t,\theta)\| \leq \int_0^t \|f(x^{M*}(\tau,\theta), u^{M*}(\tau), \theta) - f(x^{\infty}(\tau,\theta), u^{\infty}(\tau), \theta)\| d\tau$$

Since f is C^1 on a compact domain by Assumptions 2.3, 2.4, and 2.1, the Lipschitz condition applies, yielding:

$$\|x^{M*}(t,\theta) - x^{\infty}(t,\theta)\| \leq \int_0^t c\Big(\|x^{M*}(\tau,\theta) - x^{\infty}(\tau,\theta)\| + \|u^{M*}(\tau) - u^{\infty}(\tau)\|\Big) d\tau$$

for Lipschitz constant c. Since u^{M*} and u^{∞} are in the compact set U, they are bounded. Thus the Dominated Convergence Theorem applies and we have:

$$\lim_{M\in V} \int_0^t \|u^{M*}(\tau) - u^{\infty}(\tau)\| d\tau = 0.$$

For any t and δ_u we can therefore pick an N such that for all $M > N$, $M \in V$:

$$\int_0^t \|u^{M*}(\tau) - u^{\infty}(\tau)\| d\tau < \delta_u.$$

This provides us with the inequality

$$\|x^{M*}(t,\theta) - x^{\infty}(t,\theta)\| \leq cT\delta_u + c\int_0^t \|x^{M*}(\tau,\theta) - x^{\infty}(\tau,\theta)\| d\tau.$$

By Gronwall's Inequality, $\|x^{M*}(t,\theta) - x^{\infty}(t,\theta)\| \leq cT\delta_u e^{cT}$. Since for each value of t and θ, this quantity can be made arbitrarily small, $\{x^{M*}(t,\theta)\}$ converges pointwise to $x^{\infty}(t,\theta)$. Furthermore, since δ_u, though it may depend on t, does not depend on the value of θ, this convergence is uniform in θ.

We now revisit feasibility. State feasibility in this context is defined as a state generated through the state dynamics by an admissible control which satisfies all

boundary conditions and state constraints.

Lemma 3.5. *State $x^\infty(t,\theta)$ is a feasible state for Problem P.*

Proof: By definition, $x^\infty(t,\theta)$ satisfies Problem P's conditions (2) and (3). It furthermore satisfies control constraints (4), as since U is compact, the set of feasible controls is trivially closed in the topology of pointwise convergence. We thus consider the satisfaction of conditions (5), (6), (7), and (8).

Condition (5): When $\theta = \theta_i^M$, then $x^{M*}(t,\theta_i^M) = x_i^{M*}(t)$, where x_i^{M*} is the optimal state for Problem P^M generated by the optimal control $u^{M*}(t)$. Thus from constraint (15) $x^{M*}(t,\theta)$ satisfies constraint (5) at the collocated constraints:

$$e(x^{M*}(T,\theta_i^M),\theta_i^M) \leq 0, i = 1,\ldots,M.$$

For an arbitrary $\theta \in \Theta$ and an arbitrary node θ_i^M, we have

$$\|e(x^\infty(T,\theta),\theta) - e(x^{M*}(T,\theta_i^M),\theta_i^M)\|$$

$$= \|e(x^\infty(T,\theta),\theta) - e(x^{M*}(T,\theta),\theta) + e(x^{M*}(T,\theta),\theta) - e(x^{M*}(T,\theta_i^M),\theta_i^M)\|$$

$$\leq \|e(x^\infty(T,\theta),\theta) - e(x^{M*}(T,\theta),\theta)\| + \|e(x^{M*}(T,\theta),\theta) - e(x^{M*}(T,\theta_i^M),\theta_i^M)\|.$$

Due to Assumptions 2.1 and 2.4, e satisfies Lipschitz conditions with respect to each argument. Thus:

$$\|e(x^\infty(T,\theta),\theta) - e(x^{M*}(T,\theta),\theta)\| \leq c_x\|x^\infty(T,\theta) - x^{M*}(T,\theta)\|$$

for Lipschitz constant c_x with respect to the first argument of e. By the results of Lemma 3.4, for any ϵ_x there exists an M' such that $\forall M > M'$ this quantity is less than ϵ_x and this M' is independent of θ. Similarly,

$$\|e(x^{M*}(T,\theta),\theta) - e(x^{M*}(T,\theta_i^M),\theta_i^M)\| \leq c_x\|x^{M*}(T,\theta) - x^{M*}(T,\theta_i^M)\| + c_\theta\|\theta - \theta_i^M\|$$

where c_θ is the Lipschitz constant with respect to the second argument of e. Furthermore,

$$\|x^{M*}(T,\theta) - x^{M*}(T,\theta_i^M)\| \leq \int_0^T \|f(x^{M*}(\tau,\theta),u^{M*}(\tau),\theta) - f(x^{M*}(\tau,\theta_i^M),u^{M*}(\tau),\theta_i^M)\|d\tau$$

$$\leq \int_0^T c\left(\|x^{M*}(\tau,\theta) - x^{M*}(\tau,\theta_i^M)\| + \|\theta - \theta_i^M\|\right)d\tau$$

Thus, choosing now θ_i^M for each $M > M'$ to be a node within ϵ_θ of θ, as was done in Lemma 3.4 we can apply the Dominated Convergence Theorem to the integral of the second quantity, yielding:

$$\|x^{M*}(T,\theta) - x^{M*}(T,\theta_i^M)\| \leq c\int_0^T \|x^{M*}(\tau,\theta) - x^{M*}(\tau,\theta_i^M)\|d\tau + cT\epsilon_\theta.$$

By Gronwall's Inequality, $\|x^{M*}(T,\theta) - x^{M*}(T,\theta_i^M)\| \leq cT\epsilon_\theta e^{cT}$*, which through choosing* ϵ_θ *can be made smaller than any* $\bar{\epsilon}_\theta$*. Together these inequalities yield that for arbitrary* ϵ_x *and* $\bar{\epsilon}_\theta$*, there exists an* M'*, independent of* θ *such that for all* $M > M'$ *there exists a node* θ_i^M *within* ϵ_θ *of* θ *and:*

$$\|e(x^\infty(T,\theta),\theta) - e(x^{M*}(T,\theta_i^M),\theta_i^M)\| \leq \epsilon_x + \bar{\epsilon}_\theta.$$

If we assume by contradiction that $e(x^\infty(T,\theta),\theta) > 0$ *for some* θ *then since* $e(x^{M*}(T,\theta_i^M),\theta_i^M) \leq 0$*, the difference in quantities is some constant number greater than zero. Choice of* $\epsilon_x + \bar{\epsilon}_\theta$ *less than this number creates a contradiction.*

Condition (6) follows from similar arguments.

Condition (7): The satisfaction of the integral constraints will be shown through similar arguments as above with the addition of the convergence property of Remark 1. By the uniform convergence of $\{x^{M*}(t,\theta)\}$ *to* $x^\infty(t,\theta)$ *with regards to* θ *given by Lemma 3.4, and by the Lipschitz condition on* e_I *provided by Assumptions 2.1 and 2.4, it can be seen that the sequence of functions* $\{e_I\left(x^{M*}(T,\theta_i^M),\theta_i^M\right)\}$ *converges to* $e_I\left(x^\infty(T,\theta),\theta\right)$ *and that this convergence is uniform in* θ*. By Remark 1 then:*

$$\lim_{M\in V}\sum_{i=1}^{M} e_I\left(x^{M*}(T,\theta_i^M),\theta_i^M\right)\alpha_i^M = \int_\Theta e_I\left(x^\infty(T,\theta),\theta\right)d\theta$$

and by the continuity of Ψ_e *from Assumption 2.1 we have*

$$\Psi_e\left(\lim_{M\in V}\sum_{i=1}^{M} e_I\left(x^{M*}(T,\theta_i^M),\theta_i^M\right)\alpha_i^M\right) = \lim_{M\in V}\Psi_e\left(\sum_{i=1}^{M} e_I\left(x^{M*}(T,\theta_i^M),\theta_i^M\right)\alpha_i^M\right)$$

$$= \Psi_e\left(\int_\Theta e_I\left(x^\infty(T,\theta),\theta\right)d\theta\right).$$

Since the collocated constraints are satisfied:

$$\Psi_e\left(\sum_{i=1}^{M} e_I\left(x_i^M(T),\theta_i^M\right)\alpha_i^M\right) \leq K_e\epsilon_e(M), \quad i = 1,\ldots,M$$

with $\lim_{M\to\infty}\epsilon_e(M) = 0$*, and since Assumptions 2.1 and 2.4 also provide a Lipschitz condition for* Ψ_e*, the arguments of the proof of Condition (5) follow from this point on. One finds that this contradicts the possibility that* $\Psi_e\left(\int_\Theta e_I\left(x^\infty(T,\theta),\theta\right)d\theta\right) > 0$ *thus proving satisfaction of the constraint.*

Condition (8) follows from similar arguments.

Using Lemmas 3.4 and 3.5, we now establish Theorem 3.3. By Lemma 3.5, the pair (x^∞, u^∞) is feasible for Problem P. The limit of the difference in cost values satisfies:

$$\lim_{M\in V}\left\|J^M[X^{M*},u^{M*}] - J[x^\infty,u^\infty]\right\|$$

$$= \lim_{M\in V}\Bigg\|\sum_{i=1}^{M}\Big[F(x^{M*}(T,\theta_i^M))+\int_0^T r(x^{M*}(t,\theta_i^M),u^{M*}(t),t,\theta_i^M)dt\Big]\alpha_i^M$$

$$-\int_\Theta\Big[F\left(x^\infty(T,\theta)\right)+\int_0^T r(x^\infty(t,\theta),u^\infty(t),t,\theta)dt\Big]d\theta\Bigg\|$$

$$\leq \lim_{M\in V}\left\|\sum_{i=1}^{M}F\left(x^{M*}(T,\theta_i^M)\right)\alpha_i^M-\int_\Theta F\left(x^\infty(T,\theta)\right)d\theta\right\|$$

$$+\lim_{M\in V}\left\|\sum_{i=1}^{M}\int_0^T r(x^{M*}(t,\theta_i^M),u^{M*}(t),t,\theta_i^M)dt-\int_\Theta\int_0^T r(x^\infty(t,\theta),u^\infty(t),t,\theta)dtd\theta\right\|.$$

We examine the quantity:

$$\lim_{M\in V}\left\|\int_0^T r(x^{M*}(t,\theta),u^{M*}(t),t,\theta)dt-\int_0^T r(x^\infty(t,\theta),u^\infty(t),t,\theta)dt\right\|$$

$$\leq\lim_{M\in V}\int_0^T\|r(x^{M*}(t,\theta),u^{M*}(t),t,\theta)-r(x^\infty(t,\theta),u^\infty(t),t,\theta)\|dt.$$

From the continuity of r on a compact domain, we apply the Lipschitz condition to get:

$$\int_0^T\|r(x^{M*}(t,\theta),u^{M*}(t),t,\theta)-r(x^\infty(t,\theta),u^\infty(t),t,\theta)\|dt$$

$$\leq\int_0^T c\big(\|x^{M*}(t,\theta)-x^\infty(t)\|+\|u^{M*}(t)-u^\infty(t)\|\big)dt.$$

The results of Lemma 3.4, the compactness of X and U, and the assumption that u^∞ is an accumulation point of $\{u^{M*}\}_{M\in V}$, enable us to apply the Dominated Convergence Theorem. Thus:

$$\lim_{M\in V}\int_0^T c\big(\|x^{M*}(t,\theta)-x^\infty(t,\theta)\|+\|u^{M*}(t)-u^\infty(t)\|\big)dt=0$$

and this convergence must be uniform in θ due to the uniform convergence of $x^{M*}(t,\theta)$. We can therefore conclude that:

$$\lim_{M\in V}\int_0^T r(x^{M*}(t,\theta),u^{M*}(t),t,\theta)dt=\int_0^T r(x^\infty(t,\theta),u^\infty(t),t,\theta)dt$$

and that the convergence is uniform. This enables the use of Remark 1, which provides:

$$\left\| \lim_{M \in V} \sum_{i=1}^{M} \int_0^T r(x^{M*}(t, \theta_i^M), u^{M*}(t), t, \theta_i^M) dt \; - \int_\Theta \int_0^T r(x^\infty(t, \theta), u^\infty(t), t, \theta) dt d\theta \right\| = 0.$$

Similar arguments show:

$$\left\| \lim_{M \in V} \sum_{i=1}^{M} F(x^{M*}(T, \theta_i^M)) \alpha_i^M \; - \int_\Theta F\left(x^\infty(T, \theta)\right) d\theta \right\| = 0.$$

Thus $\lim_{M \in V} J^M[X^{M*}, u^{M*}] = J[x^\infty, u^\infty]$.

Assume that u^∞ is not an optimal control for Problem P. Then there exists some admissible control u such $J[x, u] < J[x^\infty, u^\infty]$, where x is a feasible state for Problem P defined by:

$$\frac{\partial x}{\partial t}(t, \theta) = f(x(t, \theta), u(t), \theta), \quad x(0, \theta) = x_0(\theta).$$

As can be seen by the definition of Problem P^M, the set of states given by $X^M = \{x(t, \theta_i^M)\}_{i=1}^M$ is a feasible solution to Problem P^M (note that this assertion requires the slackness in constraints (17) and (18)). Furthermore, $\lim_{M \in V} J^M[X^M, u] = J[x, u]$ through identical arguments as the convergence of the optimal cost. The optimality of X^{M*} for Problem P^M creates the following inequalities:

$$J[x^\infty, u^\infty] = \lim_{M \in V} J^M[X^{M*}, u^{M*}] < \lim_{M \in V} J^M[X^M, u],$$

$$\lim_{M \in V} J^M[X^M, u] = J[x, u] < J[x^\infty, u^\infty]$$

which are in contradiction. Thus u^∞ must in fact be an optimal solution to Problem P.

4. Example: Optimal Search

Here we present an example application of optimal motion planning with parameter dependency: the task of path planning for sonar search over an area. The search problem can be considered an example of a coverage path planning problem. This challenge of covering a region can be generically defined as the task of producing a path for a vehicle which reaches all regions of interest while avoiding obstacles. The problem has been studied in robotics and reviews of approaches can be found in [Choset (2001)] and [Galceran and Carreras (2013)]. Applications of covering algorithms include vacuuming, snow cleanup, lawn mowing, window cleaning, painting, and topographical mapping. Commercial products, such as the Roomba, incorporate a mix of local and global planning in their design to attain satisfactory confidence in their performance [Hess, Beinhofer, and Burgard (2014)]. However, though global coverage plans have been developed which provide assurances of eventual complete or near-complete coverage, the question of optimality in performance is still largely unaddressed. One reason for this is that performance is ultimately not a feature solely of the geometry of the

coverage path, the production of which has been the main focus of research [Choset (2001); Galceran and Carreras (2013)]. Rather, performance is additionally dependent on specific equipment capabilities as well as dynamic features, such as the duration of coverage and the impact of kinematics on equipment. Spatial approaches to coverage path planning which minimize path overlap for example, such as [Galceran and Carreras (2012)], will not necessarily produce maximal performance if a slow-acting device hasn't allotted enough time to a region.

In this example, we model the SeaFox [MarineLink (2004)] unmanned surface vehicle (USV), pictured in Figure 1 searching for a target on a $100m^2$ surface region. Param-

Figure 1. The Naval Postgraduate School's SeaFox USV

eters $\theta \in \Theta \subset \mathbb{R}^2$ parameterize the two-dimensional surface region where the goal target could be located beneath, with uniform probability density function given by $\phi(\theta)$. The dynamics of the vehicle are represented by the horizontal location (x_1, x_2), heading angle ψ, and velocity v, and given by:

$$\frac{dx}{dt}(t) = \begin{bmatrix} \dot{x}_1(t) \\ \dot{x}_2(t) \\ \dot{\psi}(t) \\ \dot{v}(t) \end{bmatrix} = \begin{bmatrix} v \sin \psi(t) \\ v \cos \psi(t) \\ u_1(t) \\ u_2(t) \end{bmatrix}, \quad x(0) = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1.57 \end{bmatrix}.$$

The functions u_1 and u_2 are the control inputs guiding the vehicle's motion. They are limited by constraints of the form:

$$|u_1(t)| \leq K_1, \quad |u_2(t)| \leq K_2.$$

Velocity is also constrained:

$$0 \leq v \leq K_v.$$

Coefficient values are provided in Table 2.

A range-limited sonar is modeled by an indicator function, $I(x, \theta)$ which returns 1 if the point θ is within range of location $x \in \mathbb{R}^2$ and 0 if it doesn't. To comply with the regularity assumptions of the numerical algorithm, we further approximate this indicator function by a continuously differentiable approximation of the indicator

function, $I_{smooth}(x,\theta)$, which we define as:

$$I_{smooth}(x,\theta) = \frac{1}{2}\left(1 + erf\left(\frac{c_1 - c_2 x}{\sigma}\right)\right),$$

where

$$erf(x) = \frac{2}{\sqrt{\pi}}\int_0^x e^{-t^2} dt.$$

Coefficient values are given in Table 2. Figures 2 and 3 illustrate the smoothed indicator concept and its implementation with the values in Table 2. For an object within the

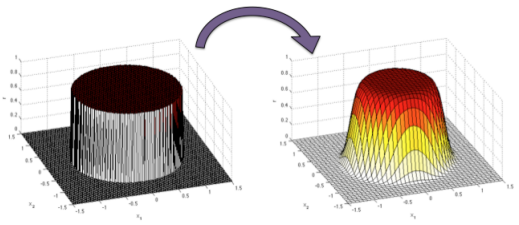

Figure 2. Smoothed sensor range indicator function

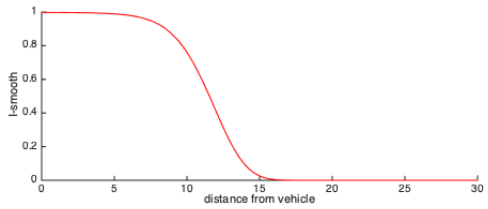

Figure 3. I_{smooth} for the values in Table 2

vehicle's range, we assume that the vehicle acts on it with a constant probabilistic detection rate, r_0. The probability that the vehicle detects the object in time interval $[t, t+\Delta t]$ is approximated by $r_0 I_{smooth}(x,\theta)\Delta t$. The probability of *not* detecting the target at location θ by time t, given by $P_{ND}(t,\theta)$, is thus given by:

$$P_{ND}(t+\Delta t,\theta) = P_{ND}(t,\theta)[1 - r_0 I_{smooth}(x,\theta)\Delta t]$$

which as $\Delta t \to 0$ creates the dynamics:

$$\frac{\partial P_{ND}}{\partial t}(t,\theta) = -r_0 I_{smooth}(x,\theta) P_{ND}(t,\theta), \quad P_{ND}(0,\theta) = \phi(\theta).$$

This is the exponential detection model, derived in [Koopman (1956)].

In the scenarios that follow, the initial probability at each location has been set to $\phi(\theta) = 1$. Together, x and P_{ND} create the state space of our control problem, referred to with the augmented state variable z:

$$z(t,\theta) = \begin{bmatrix} x(t) \\ P_{ND}(t,\theta) \end{bmatrix}.$$

We consider three path plans and contrast their performance. Performance is gauged by the expected final probability of missing the target after searching:

$$J = \frac{\int_\Theta P_{ND}(T,\theta)d\theta}{\int_\Theta d\theta}. \tag{21}$$

Table 2. Scenario Values for Paths 1 - 3

Aperture coefficients	c_1	20
	c_2	.15
	σ	10
max detection rate	r_0	.5
Control constraints	K_1	.125
	K_2	1
Velocity constraint	K_v	12.5 m/s
search region	Θ	$[0,100] \times [0,100]$ m^2
Initial probability	$\phi(\theta)$	$\phi(\theta) = 1$
Final time	T	360 s
Switchback interval	2ρ	2×12.5
Path 1 velocity	v	1.57 m/s

Path 1, a baseline, is a heuristic solution generated by following a lawnmower pattern at constant velocity. The vehicle region is over $\Theta = [0,100] \times [0,100]$ m^2. The width between paths has been chosen as 2ρ where ρ is an estimate of the effective radius of the aperture, based on I_{smooth} and allowing for a small amount of overlap; its value is set as $\rho = 12.5$. The time interval has been chosen to allow for this lawnmower path to finish: $T = 360$. Control constraints on $u_1(t)$ enforce a maximum turning rate, $\dot{\psi}_{max}$, for executing the lawnmower path's switchbacks, and the velocity of the vehicle is determined by ρ and this turning rate using the equation $\rho\dot{\psi}_{max} = v = 1.57$. Figure 4 shows the performance of the lawnmower pattern, Path 1. Path 1's average probability of failure, J_1, is 0.2982. The lowest failure probability $P_{ND}(T,\theta)$ at point θ it attains is 0.1195 and the highest is 1 (i.e. there are spots it completely misses, where if the target is located there, probability of failure to detect is 1.

Path 2 is generated by solving the optimal control problem to minimize the failure probability in equation (21). This solution has been generated by discretizing Θ with 25×25 nodes and time with 150 nodes, both domains using LGL pseudospectral nodes. We note that the choice of solution method used in the time domain is independent of the method chosen for the parameter domain. In this example, we have chosen to use matching methods. However, any method appropriate to the approximate problem established in Section 3, Problem $\mathbf{P^M}$, can be applied to the time domain. Discretized state and control constraints (14) - (18) may create an approximate problem which, due to the constraints, may benefit from a method from a time domain method which caters more specifically to constrained optimal control problems, such as the local pseudospectral method [Darby, Garg, and Rao (2011)], or the symplectic pseudospectral method [Wang, Peng, Zhang, Chen, and Zhong (2017)]. This can be ascertained from the structure of Problem $\mathbf{P}$. In this example, which has smooth state and control solutions, we find the global pseudospectral method utilized to be sufficient.

Figure 5 shows the performance of Path 2. Path 2's average failure probability, J_2, is 0.1795, the lowest failure probability at point θ it attains is 0.0710 and the highest is 0.9694. As one would expect, the optimal solution performs much better than

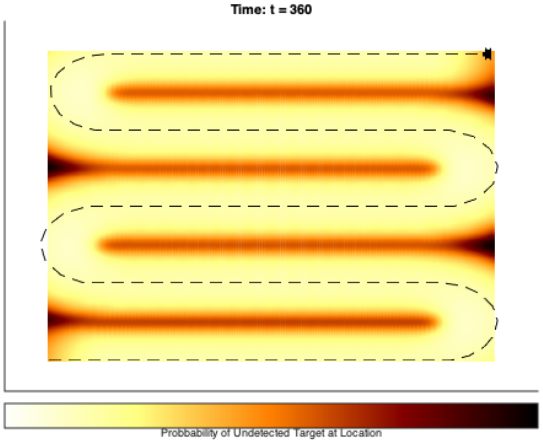

Figure 4. Path 1 at final time T

Path 1, providing around a 40% reduction in expected failure. However, depending on one's standards, Path 2 may still not be ideal. It is notable that the highest failure probability, 0.9694, reveals spots which have been mostly skipped in order to create the best average. An alternate goal for a vehicle's searching path could be to a thoroughly search the region to within a certain risk threshold.

Path 3 has been generated by minimizing the expected failure given by equation (21) subject to the risk constraint:

$$P_{ND}(T, \theta) \leq \gamma, \ \forall \theta \in \Theta \tag{22}$$

with $\gamma = 0.5$. This problem sets a floor for the vehicle's performance. As with Path 2, Path 3 is generated by discretizing Θ with 25×25 nodes and time with 150 nodes, both domains using LGL pseudospectral nodes. Figure 6 shows the performance of Path 3. Path 3's average final performance, J_3, is worse than Path 2's, at 0.2305. However, it's maximum non-detection probabilities, γ, is much improved at 0.6917.

Table 3. Performance Values

	Avg final P_{ND}	Low P_{ND}	High P_{ND}	Computation time
Path 1 (heuristic)	0.2982	0.1195	1	n/a
Path 2 (optimal avg)	0.1795	0.0710	0.9694	12.61 min
Path 3 (constrained worst case)	0.2305	0.0638	0.6917	26.67 min

Table 3 provides a summary of the performance of all three paths. Each path's control solution is assessed in performance by propagation through a refined 300×300 LGL-pseudospectral quadrature grid of Θ to estimate its performance in the non-approximated problem. We note that Path 3, calculated from a 25×25 discretization of Θ, does not yet strictly satisfy the constraint inequality $P_{ND}(T, \theta) \leq 0.5$ at all $\theta \in \Theta$ when the control solution is propagated through the refined 300×300 grid. The strict

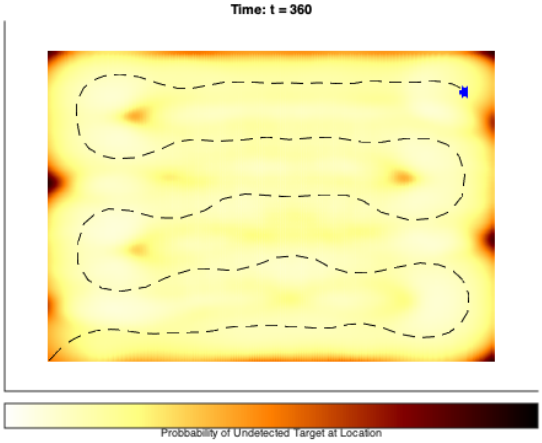

Figure 5. Path 2 at final time T

satisfaction of the constraint at the discretized nodes yields an asymptotic point-wise convergence. Figures 7 and 8 show the convergence of the constraint as discretization increases.

5. Conclusions

The tools presented in this paper enable the computation of numerical solutions to a wide variety of optimal control problems with parameter dependency in the cost and/or dynamics. These problems include those with state constraints over the parameter domain, constraints over the distribution (such as constrained mean and other statistics), and also minimum time problems. The computational framework presented allows for the flexible choice of many numerical schemes such as pseudospectral methods or sparse grid methods, as well as the ability to use separate methods in the parameter domain versus the time domain, allowing for the application in the time domain of the plentiful available computational methods developed for standard optimal control problems.

These tools enable multiple new avenues for application. As Section 4 mentions, the optimal control framework presented in this paper is applicable to many types of area coverage problems. These problems are increasingly relevant for autonomous vehicles as these vehicles gain in durability and longterm deployability. Problems such as aerial and marine mapping, search, networked communication coverage, and efficient routing can all be approached as control problems with (spatial) parameter dependencies. The use of optimal control for these problems provides the opportunity to generate path plans optimized over specific vehicle and sensor dynamics. This has the potential to provide substantial performance improvements, especially when planning for multi-vehicle teams with heterogenous dynamics.

The framework also provides a tool for many motion planning problems with uncertainty. The parameter domain can be used to consider uncertain locations, as in the optimal search problem in Section 4. It can also be used to incorporate model uncertainty through estimated parameter ranges or as the coefficients of series expansion

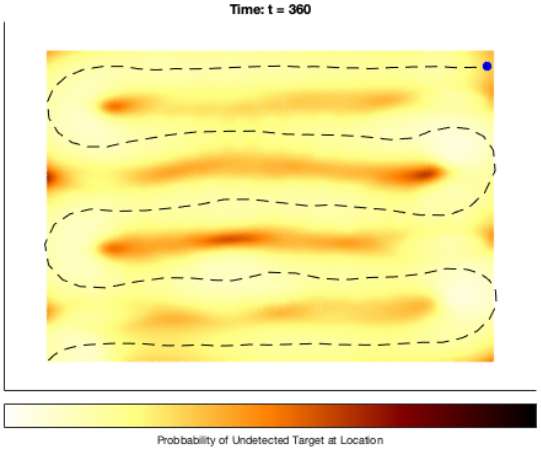

Figure 6. Path 3 at final time T

for function approximation. These provide approaches for optimizing plans against dynamic targets with unknown dynamics models. Another interesting future direction of application may be in the estimation of the unknown parameters themselves. With sensor-based observations driving updated parameter estimates (for instance through Bayesian updates), path planning for the mobile sensor platforms to optimize observation trajectories for estimating the unknown parameters becomes an optimal control problem dependent on those parameters.

The wealth of potential applications provide many future directions for this problem, as well as many technical challenges. A challenge in application will be the available methods for solving the resulting standard control problem, which may have to deal with a high number of states and state constraints. Due to the high number of state variables generated through parameter discretization, efficiency is also key. However, as the example in this paper demonstrates, the implementation of this approach is already feasible. Future work includes building problem models for the applications described above and investigating efficient numerical methods for both parameter and time domain implementations.

References

Becker, A. T. (2012). *Ensemble control of robotic systems* (Unpublished doctoral dissertation). University of Illinois at Urbana-Champaign.

Brockett, R., & Khaneja, N. (2000). On the stochastic control of quantum ensembles. In *System theory* (pp. 75–96). Springer.

Choset, H. (2001). Coverage for robotics - a survey of recent results. *Annals of Mathematics and Artificial Intelligence*, *31*, 113-126.

Darby, C. L., Garg, D., & Rao, A. V. (2011). Costate estimation using multiple-interval pseudospectral methods. *Journal of spacecraft and rockets*, *48*(5), 856–866.

Darlington, J., Pantelides, C., Rustem, B., & Tanyi, B. (2000). Decreasing the sensitivity of open-loop optimal solutions in decision making under uncertainty. *European Journal of Operational Research*, *121*, 343-362.

Davis, P. J., & Rabinowitz, P. (2007). *Methods of numerical integration*. Courier Corporation.

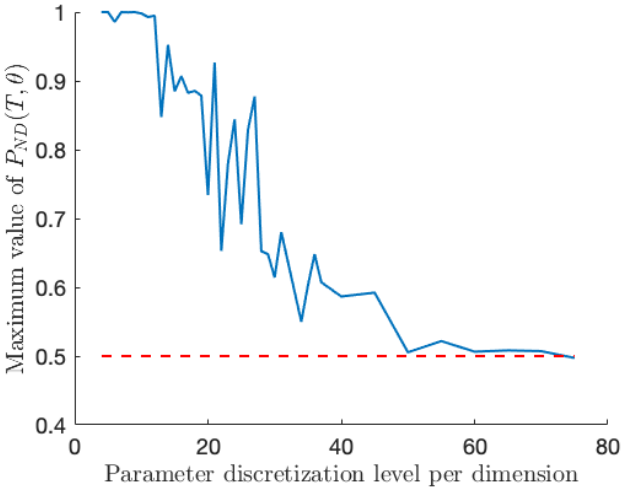

Figure 7. Convergence of maximum constraint violation as discretization increases

Fisher, J., & Bhattacharya, R. (2011, January). Optimal trajectory generation with probabilistic system uncertainty using polynomial chaos. *Journal of Dynamic Systems, Measurement, and Control*, *133*.

Foraker, J. (2011). *Optimal search for moving targets in continuous time and space using consistent approximations* (Unpublished doctoral dissertation). Naval Postgraduate School.

Galceran, E., & Carreras, M. (2012). Coverage path planning for marine habitat mapping. In *Oceans* (pp. 1–8).

Galceran, E., & Carreras, M. (2013). A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, *61*(12), 1258–1276.

Gerstner, T., & Griebel, M. (1998). Numerical integration using sparse grids. *Numerical Algorithms*, *18*(3-4), 209.

Hess, J., Beinhofer, M., & Burgard, W. (2014). A probabilistic approach to high-confidence cleaning guarantees for low-cost cleaning robots. In *2014 ieee international conference on robotics and automation* (pp. 5600–5605).

Hover, F. S., & Triantafyllou, M. S. (2006). Application of polynomial chaos in stability and control. *Automatica*, *42*, 789-795.

Koopman, B. (1956, October). The theory of search ii: Target detection. *Operations Research*, *4*(5), 503-531.

Li, J.-S., & Khaneja, N. (2006). Control of inhomogeneous quantum ensembles. *Physical Review A*, *73*(3), 030302.

Li, J.-S., & Khaneja, N. (2007, December). Ensemble control of linear systems. *Proceedings of the 46th IEEE Conference on Decision and Control*.

Li, J.-S., & Khaneja, N. (2009). Ensemble control of bloch equations. *IEEE Transactions on Automatic Control*, *54*(3), 528–536.

MarineLink. (2004, April). *Little usv, big applications.* Retrieved 2019-05-15, from `https://www.marinelink.com/news/applications-little-big323082`

Phelps, C., Gong, Q., Royset, J. O., Walton, C., & Kaminer, I. (2014). Consistent approximation of a nonlinear optimal control problem with uncertain parameters. *Automatica*, *50*(12), 2987–2997.

Phelps, C., Royset, J. O., & Gong, Q. (2016). Optimal control of uncertain systems using sample average approximations. *SIAM Journal on Control and Optimization*, *54*(1), 1–29.

Pursiheimo, U. (1976). A control theory approach in the theory of search when the motion of the target is conditionally deterministic with stochastic parameters. *Applied Mathematics and Optimization*, *2*, 259-264.

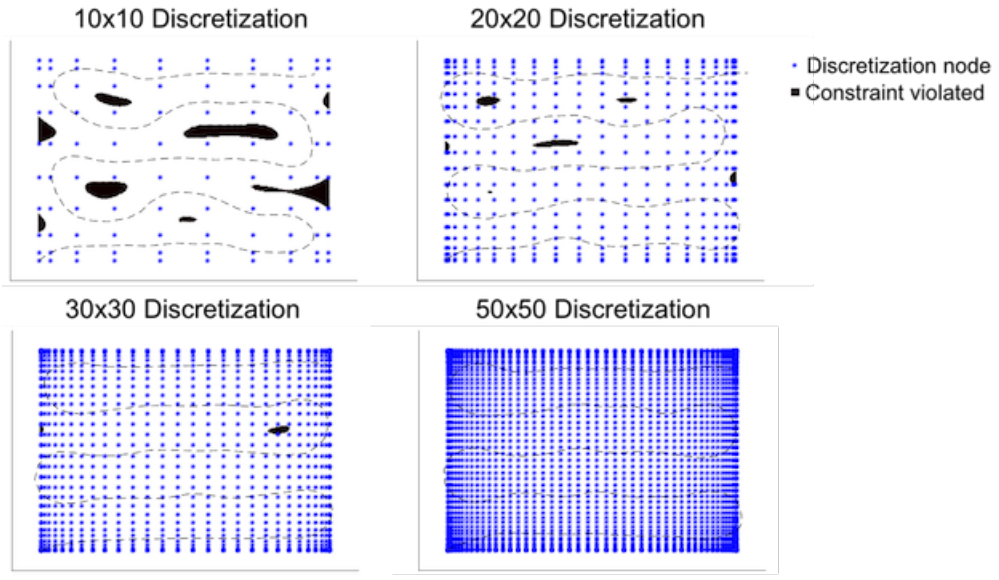

Figure 8. Convergence of constraint satisfaction as discretization increases

Ross, I. M., Karpenko, M., & Proulx, R. J. (2016). Path constraints in tychastic and unscented optimal control: Theory, application and experimental results. In *American control conference (acc), 2016* (pp. 2918–2923).

Ross, I. M., Proulx, R. J., & Karpenko, M. (2014a). Unscented optimal control for orbital and proximity operations in an uncertain environment: A new zermelo problem. In *Proceedings of aiaa space and astronautics forum and exposition: Aiaa/aas astrodynamics specialist conference. san diego, ca.*

Ross, I. M., Proulx, R. J., & Karpenko, M. (2014b). Unscented optimal control for space flight. In *Proceedings of the 24th international symposium on space flight dynamics. laurel, md.*

Ross, I. M., Proulx, R. J., Karpenko, M., & Gong, Q. (2015). Riemann-Stieltjes optimal control problems for uncertain dynamic systems. *Journal of Guidance, Control, and Dynamics*, *38*(7), 1251–1263.

Ruppen, D., Benthack, C., & Bonvin, D. (1995). Optimization of batch reactor operation under parametric uncertainty- computational aspects. *Journal of Process Control*, *5*(4), 235-240.

Ruths, J., & Li, J.-S. (2012). Optimal control of inhomogeneous ensembles. *Transactions on Automatic Control*, *57*(8).

Terwiesch, P., Ravemark, D., Schenker, B., & Rippin, D. W. T. (1998). Semi-batch process optimization under uncertainty: Theory and experiments. *Computers and Chemical Engineering*, *22.1*, 201-213.

Trefethen, L. N. (2008). Is gauss quadrature better than clenchaw–curtis. *SIAM Review*, *50*(1), 67–87.

Walton, C., Phelps, C., Gong, Q., & Kaminer, I. (2016). A numerical algorithm for optimal control of systems with parameter uncertainty. *10th IFAC Symposium on Nonlinear Control Systems NOLCOS 2016, IFAC-PapersOnLine*, *49*(18), 468–475.

Wang, X., Peng, H., Zhang, S., Chen, B., & Zhong, W. (2017). A symplectic pseudospectral method for nonlinear optimal control problems with inequality constraints. *ISA transactions*, *68*, 335–352.