

**The Design and Implementation of Motion Planning Problems Given
Parameter Uncertainty**

by

Claire Walton

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Applied Mathematics and Statistics

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, SANTA CRUZ

Committee in charge:
, Chair
Professor Isaac Kaminer
Professor Hongyun Wang

Winter 2015

**The Design and Implementation of Motion Planning Problems Given
Parameter Uncertainty**

Copyright 2015

by

Claire Walton

Contents

List of Figures	iii
List of Tables	v
1 Introduction	1
1.1 Motivating Problem: The Optimal Search Problem	1
1.1.1 Searcher and Detection Models	3
1.1.2 Target Models	5
1.1.3 Optimal Control Formulation	7
1.2 Expanded Modeling Frameworks for Multi-Agent Motion Planning . . .	10
1.2.1 Mathematical Properties	17
1.3 Thesis Contributions	19
1.3.1 Multi-Agent Models	19
1.3.2 Necessary Conditions for Optimality	19
1.3.3 Computational Algorithm	20
1.4 Thesis Outline	21
2 Kinematic and Performance Models	23
2.1 Kinematic Models	23
2.1.1 Heading and Velocity Driven Herding	30
2.1.2 Force Driven Herding	36
2.2 Performance Models	38
2.2.1 Strategic Search	40
2.2.2 Shooting Problems	43
2.2.2.1 The Kamikaze Swarm Scenario	48
2.2.2.2 Mutual Attrition	52
2.2.3 Path Coverage	55
3 Analytical Foundations	59
3.1 General Problem Statement	59
3.2 Necessary Conditions for Optimality	62
3.2.1 Necessary Conditions for Open Control Regions	63
3.2.2 Example Analytic Solution Via Necessary Conditions	73

4	Computational Approach	77
4.1	Numerical Algorithm	78
4.1.1	Approximating Parameter Space	79
4.1.2	Convergence of Primal Variables	80
4.1.3	Approximating the Time Domain	88
4.1.4	Convergence of Dual Problems	89
4.1.5	Linear Quadratic System	98
4.2	Implementations	102
4.2.1	Ensemble Control of Nonholonomic Unicycles	103
4.2.2	Ensemble Control of Harmonic Oscillators	108
4.2.3	Channel Search	116
4.2.4	Kamikaze Swarm Scenario	121
5	Conclusions and Future Work	123
5.1	Increasing Computational Speed	124
5.1.0.1	Monte Carlo Methods	125
5.1.0.2	Quasi-Monte Carlo Methods	125
5.1.0.3	Sparse Grid Methods	127
5.1.0.4	Parallelization	129
5.2	Expanding Mathematical Framework	131
5.2.0.5	Feedback Control	131
5.2.0.6	End Time State Constraints	132
5.3	Furthering Applications	133
5.3.1	Protective Herding	133
5.3.2	Longterm Autonomous Flight with Solar-Powered Thermal- ing Gliders	134
	Bibliography	137

List of Figures

1.1	Diagram of Components of the Classic Optimal Search Problem	2
1.2	Example Detection Rate Function: Poisson Scan Model	4
1.3	Example of a distribution of conditionally deterministic target trajectories with two parameters made from the optimization of the optimal control problem described in section II.C of [1]. The target trajectories are designed to minimize the time between interception of a high valued unit (HVV) given velocity, starting time, and curvature constraints. The position of the HVV is given by the diamond icon.	7
1.4	Illustration from Chung et al. (2011) [2] of a general taxonomy of search problems as outlined by Benkowski et al. (1991) [3]	12
1.5	Diagram of the Types of Quantities Involved in Optimal Search	14
1.6	Diagram of Example Interactions Found in an Oil Cleanup	15
1.7	Multiple Influences Wanted	16
2.1	Geometry of Radial Evasion	24
2.2	Geometry of the Derivative of a Trajectory in Relation to Heading Angle	25
2.3	Non-uniqueness of Tangent	27
2.4	Relative Position Vector and Normalized Relative Position Vector	28
2.5	Relative Position Vectors Between Multiple Agents	30
2.6	Diagram of Heading Vector	31
2.7	Example of Average Heading Policy with Two Herders	33
2.8	Beta functions can be used to simulate a maximal firing effectiveness at some distance from the firing agent	45
2.9	“Half Beta” functions can be used to simulate decreasing firing effectiveness over distance capped by a finite range	45
2.10	Effect of an angularly decaying multiplier on an example attrition rate function	47
2.11	Resulting angularly decaying function reflecting POV limitations	47
2.12	Structure of Kamikaze Scenario	48
2.13	Structure of Mutual Attrition Scenario	52
2.14	Example: Smoothed Aperture-Based Rate Function	56

4.1	Diagram of Goals for Primal and Dual Relations for Parameter Uncertainty Control. Red lines designate the remaining needed results.	92
4.2	Diagram of Primal and Dual Relations for Standard Control Combined with Parameter Uncertainty Control	96
4.3	Convergence of analytical solutions for $J(x^*, u^*)$ and $J^M(x_M^*, u_M^*)$ using Euler's method vs Legendre pseudospectral. The dimension of the problem has been set at $K = 2$, and a Beta(10,10) distribution over the domain $[0,1]$ has been given to each parameter.	102
4.4	Optimal Controls for Unicycle Ensemble Problem	105
4.5	Optimal Trajectories for Unicycle Ensemble Problem	106
4.6	Values of the mapped duals corresponding to $\lambda_1(t, \omega)$ and $\lambda_2(t, \omega)$ returned by the NLP solver, plotted over all values of parameter nodes . .	108
4.7	Optimal Control Values, $u^*(t)$, for Nominal Oscillator System	112
4.8	Optimal State Values, $x_1^*(t)$, $x_1^*(t)$, for Nominal Oscillator System	112
4.9	Perturbed state values, $x_1^*(t, \omega)$, $x_1^*(t, \omega)$, created by the optimal control to the nominal system. Color is determined by $ \omega - \hat{\omega} $	114
4.10	State values created by the optimal control to the general parameter uncertainty problem. Color is determined by $ \omega - \hat{\omega} $	114
4.11	Values for u_1 and $\int_{\Omega} \lambda_1(t, \omega) d\omega$	115
4.12	Values for u_2 and $\int_{\Omega} \lambda_2(t, \omega) d\omega$	115
4.13	Hamiltonian values for general parameter uncertainty problem.	115
4.14	Snapshots of numerical solution to 'Channel Search Problem'. Colors represent the log probability density value of an undetected target at a point at time t , i.e. the evolution over time of the probability density of target location given parameter value compounded with the probability of target non-detection for given parameter value . 'Low' = $4.14 * 10^{-6}$, 'High' = $2.9317 * 10^{-4}$	119
4.15	Searching trajectories. Image (a) illustrates optimal control solutions, image (b) straight line patrols, and image (c) looping patrols.	120
4.16	Comparison of the performance of the optimal control solution vs heuristic methods.	120
4.17	Snapshots of numerical solution to 'Kamikaze Shooting Problem'. The magenta icon indicates the position of the HVU and the green icon is defender. Colors represent the log probability density value of a surviving attacker at a point at time t . 'Low' = $4.14 * 10^{-6}$, 'High' = $2.9317 * 10^{-4}$	122
5.1	High Dimensional Kamikaze Scenario Implemented with Monte Carlo; $d = 20$, $M = 200$	126
5.2	Cost and Error for Sparse Grid vs Full Grid System of [4] with $d = 4$. .	128
5.3	Snapshots of Protective Herding Scenario with 2 Defenders and 2 Dependent Attacker Swarms with 50 Members Each	135

List of Tables

4.1	Parameter Values for Channel Search Problem	118
4.2	Run Times Vs Nodes for Channel Search Problem	121

Abstract

The Design and Implementation of Motion Planning Problems Given Parameter
Uncertainty

by

Claire Walton

Doctor of Philosophy in Applied Mathematics and Statistics

University of California, Santa Cruz

, Chair

This dissertation explores the potential for utilizing direct methods in optimal control to solve trajectory optimization problems with uncertain parameters. Parameter uncertainty extends traditional optimal control problems by inserting constant but unknown uncertainty into problem components such as the cost function or the state dynamics. The objective in these problems becomes to minimize the cost function, subject to all available information, such as a range of values or prior distribution for the uncertain parameter. Research into this topic has historically been motivated by applications in optimal search theory. However, the development of more general numerical methods and optimality conditions creates the potential to address a greater variety of problems.

The goal of this dissertation is to facilitate the maturation of optimal control problems with parameter uncertainty into a tool of wider applicability. This is ap-

proached by addressing three aspects of progressing prior research: the development of more realistic and interactive kinematic and performance models for application in problems with parameter uncertainty, the development of a general mathematical framework for parameter uncertainty problems, and a numerical algorithm for generating solutions.

Acknowledgments

I would like to thank, first and foremost, my advisor Qi Gong for his wealth of ideas and infinite patience, and my colleagues Isaac Kaminer and Johannes Royset for their many contributions. I would also like to thank my parents for their support over all the years, and my siblings for their support during many late night phone calls. Finally, I would like to thank Nancy MacAllister for her generosity and hospitality, and Michael Rienstra for his amazing willingness to go on grocery runs.

Dedicated to my family.

Chapter 1

Introduction

1.1 Motivating Problem: The Optimal Search Problem

An area of interest in robotics and autonomous vehicle research is the development of trajectory optimization algorithms given uncertainty. In such situations, autonomous agents are faced with the task of optimizing their behavior under a given performance criterion while taking into consideration environmental or infrastructural features which may have an amount of uncertainty. This problem can arise in many situations, including search and rescue operations, robotic guidance, missile defense, and combat situations.

A much-studied problem of trajectory optimization under uncertainty is The Optimal Search Problem, developed in the 1940's. The optimal search problem considers the question of how to optimize the probability of detection of a non-evasive target with certain uncertain features, given detection equipment capabilities and limitations on the allocation of search effort. The problem has been studied extensively in the fields of

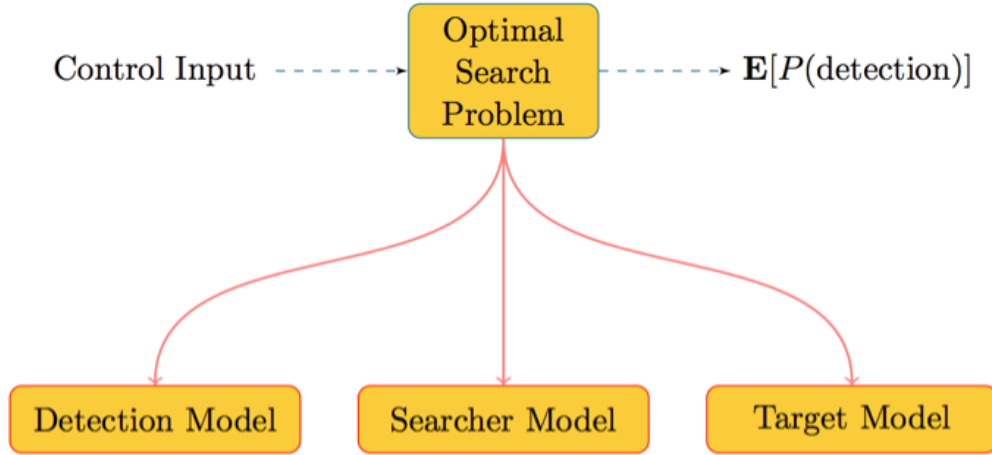


Figure 1.1: Diagram of Components of the Classic Optimal Search Problem

applied mathematics and operations research and has applications in search and rescue operations, robotic missions, missile defense, and tactical situations; a review is given by [5].

In order to construct The Optimal Search Problem, three major components must be modeled: a criteria for evaluating detection probabilities given some allocated search effort, a model for allocating that search effort (this could be through the movement of some ship, or through a more discrete process of allocating power to different point in a sensor grid), and a model to capture what knowledge we have of target motion and location (Figure 1.1).

The modeling decisions for each of these components determine the resulting mathematical problem, and as with most scenarios, the resulting mathematical problems can be diverse. One of the common frameworks, studied by for instance [6], [1], and [7], configures the search problem as an optimal control problem influenced by parameter uncertainty. This particular framework has served as the starting point for the models

in this thesis, so the following will briefly present its details.

1.1.1 Searcher and Detection Models

In this framework, first developed by B.O. Koopman in [8], search effort is applied by autonomous searching agents characterized over time by a vector of states $x(t) \in \mathbb{R}^{n_x}$. These states are determined by a dynamical system:

$$\dot{x}(t) = f(x(t), u(t)), \quad x(0) = x_0$$

with a control input $u(t) \in \mathbb{R}^{n_u}$ and known initial condition $x_0 \in \mathbb{R}^{n_x}$. Each of these agents is equipped with detection equipment, for instance radar or sonar. Establishing a model to quantify the effectiveness of such equipment was one of the earliest contributions to The Optimal Search Problem, and the model derived during WWII in [9]—the so-called ‘exponential detection model’—remains ubiquitous in the literature today. In general, detection through visual input or through sensor input such as sonar and radar is probabilistic in nature. There is a chance that an object’s presence can be missed, even when it’s within sensor range, however intuitively this probability decreases as attention is focused on it for a longer period of time.

The exponential detection model follows from the assumption that a feature termed the *instantaneous rate of detection* can be effectively modeled. Given the position of a searcher at $x(t)$ and a target at $y(t) \in \mathbb{R}^{n_y}$, this instantaneous rate of detection is a function $r(x(t), y(t)) : \mathbb{R}^{n_x} \times \mathbb{R}^{n_y} \rightarrow \mathbb{R}$ such that the probability of detection in a sufficiently small interval $[t, t + \Delta t]$ is independent from previous time intervals and




Figure 1.2: Example Detection Rate Function: Poisson Scan Model

given by the quantity $r(x(t), y(t))\Delta t$. The rate function $r(x(t), y(t))$ is chosen to model the qualities of sensor equipment such as acoustic and sonar sensors, which have rapid enough sweep rates to be modeled as continuous processes (Figure 1.2). Proceeding with these assumptions, a formula for the probability of target detection by a single searcher can be derived. If we denote the probability of *non-detection* of a target with the function $P(t)$, the independence of the time intervals creates the following difference equation:

$$P(t + \Delta t) = P(t) [1 - r(x(t), y(t))\Delta t].$$

Rearranging gives:

$$\frac{P(t + \Delta t) - P(t)}{\Delta t} = -P(t)r(x(t), y(t))$$

which as $\Delta t \rightarrow 0$ becomes:

$$\dot{P}(t) = -P(t)r(x(t), y(t)).$$

This differential equation has the exponential solution:

$$P(t) = e^{-\int_0^t r(x(\tau), y(\tau)) d\tau}.$$

1.1.2 Target Models

The values of $y(t)$ in the above probability are considered to be uncertain (otherwise the task of searching wouldn't have much purpose!). This uncertainty can be modeled in a variety of ways, including as a diffusion process with stochastic parameters, [10], and Markovian motion, [11]. The model of conditionally deterministic motion assumes that the motion of the targets is given entirely by a function of time and a parameter ω . In other words, the target motion becomes conditional on ω :

$$y(t|\omega) = h(t, \omega)$$

in a fashion which can be specified by some function $h(t, \omega)$. This parameter ω is an element of a bounded parameter space $\Omega \subset \mathbb{R}^n$ and furthermore has a known probability density function $p : \Omega \rightarrow \mathbb{R}$.

As an illustration of the kind of probabilistic parameters that can be chosen in a search model we can consider the classic channel patrol problem, created by [8] in 1946, and subsequently studied analytically in 1982 by [12] and computationally in 2011 by [7]. In this problem a searcher is given the task of optimally patrolling a channel of water of given width. Intruders into this channel are known to be moving in a straight line through the channel at a constant velocity. Some of the specific features of their

motion, however, may be known only probabilistically, such as starting position, starting time, or velocity. This creates a distribution of possible positions for each intruder at each moment in time which is then incorporated into optimization decisions.

The model of conditionally deterministic motion can also be used to describe more complicated processes. A particularly interesting example of this is that conditionally deterministic motion can encompass numerical solutions to previously computed optimization problems. This fact is utilized by [1] to create optimal search trajectories against a swarm of attackers headed towards a high-valued unit (HVV) moving along a pre-determined trajectory. Prior to the implementation of the optimal search algorithm, the probability density function of parameter space is approximated through discretization and a set of each attacker's conditional trajectories over time is created. These trajectories are computed as a solution to an optimal control problem which seeks to minimize the time between attacker interception of the HVV given velocity, starting time, and curvature constraints.

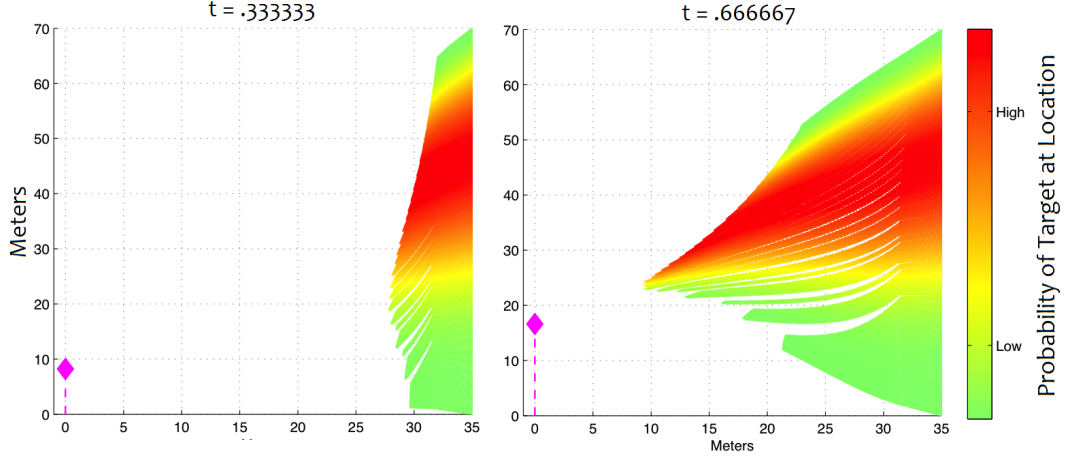


Figure 1.3: Example of a distribution of conditionally deterministic target trajectories with two parameters made from the optimization of the optimal control problem described in section II.C of [1]. The target trajectories are designed to minimize the time between interception of a high valued unit (HVV) given velocity, starting time, and curvature constraints. The position of the HVU is given by the diamond icon.

1.1.3 Optimal Control Formulation

The conditioning of y on ω leads to a probability of target detection which is itself a conditional probability, i.e. $P(t|\omega)$, defined by:

$$P(t|\omega) = e^{-\int_0^t r(x(\tau), y(\tau|\omega)) d\tau}.$$

A natural performance measure is to minimize the expectation of the probability of non-detection over a time interval $[0, T]$. This creates the objective:

$$\min J = 1 - \int_{\Omega} e^{\int_0^T r(x(t), y(t|\omega)) dt} p(\omega) d\omega$$

in which the existence of uncertain parameters in the problem has presented itself through a (potentially high dimensional) integration over a parameter space. If we

assume independence in the detection of each target by each searcher, we can also use these principles to derive detection probabilities in the case of multiple searchers and targets. For K searchers and L targets, let the searcher states be given by:

$$\dot{x}_k(t) = f_k(x_k(t), u_k(t)), \quad x_k(0) = x_{k,0}, \quad k = 1, \dots, K$$

the targets states by:

$$y_l(t|\omega) = h_l(t, \omega)$$

and each searcher's detection rates by:

$$r_k(x_k(t), y_l(t|\omega)).$$

The probability of the k -th searcher not detecting the l -th target by time $t + \Delta t$ is:

$$P_{k,l}(t + \Delta t|\omega) = P_{k,l}(t) [1 - r_k(x_k(t), y_l(t|\omega))\Delta t].$$

The ‘worst-case’ probability of none of the searchers detecting any of the targets by time $t + \Delta t$ is determined by the relation $P(t|\omega) = \prod_{l=1}^L \prod_{k=1}^K P_{k,l}(t + \Delta t|\omega)$ which creates the difference equation:

$$\begin{aligned} P(t + \Delta t|\omega) &= P(t) \prod_{l=1}^L \prod_{k=1}^K [1 - r_k(x_k(t), y_l(t|\omega))\Delta t] \\ &= P(t) [1 - \sum_{l=1}^L \sum_{k=1}^K r_k(x_k(t), y_l(t|\omega))\Delta t + (\Delta t)^2(\dots) + \dots] \end{aligned}$$

As $\Delta t \rightarrow 0$ the higher order terms with respect to Δt disappear and so we get:

$$\dot{P}(t|\omega) = P(t|\omega)[1 - \sum_{l=1}^L \sum_{k=1}^K r_k(x_k(t), y_l(t|\omega))]$$

and

$$P(t|\omega) = e^{-\int_0^t \sum_{l=1}^L \sum_{k=1}^K r(x_k(\tau), y_l(\tau|\omega)) d\tau}.$$

This creates the objective:

$$\min J = 1 - \int_{\Omega} e^{\int_0^T \sum_{l=1}^L \sum_{k=1}^K r(x_k(t), y_l(t|\omega)) dt} p(\omega) d\omega$$

The added integration over parameter space as well as the exponential function encompassing the time integral distinguish this cost function from the standard types encountered in optimal control. For a variety of derived cost functions, this framework for The Optimal Search Problem creates the following class of optimal control problems:

The Optimal Search Problem: Given probability density function $p : \Omega \rightarrow \mathbb{R}$, determine the control $u : [0, T] \rightarrow U \in \mathbb{R}^{n_u}$ that minimizes the cost:

$$J = \int_{\Omega} \left[F(x(T), \omega) + G \left(\int_0^T r(x(t), u(t), t, \omega) dt \right) \right] p(\omega) d\omega$$

subject to the searcher dynamics:

$$\dot{x}(t) = f(x(t), u(t)), \quad x(0) = x_0$$

and control constraint

$$g(u(t)) \leq 0, \forall t \in [0, T].$$

This class of control problems allows for the influence of the uncertain parameter on a broad range of possible cost functions while retaining deterministic state equations. As a class of control problems, it has been studied by multiple researchers, including [13], [6], [1], and [14].

1.2 Expanded Modeling Frameworks for Multi-Agent Motion Planning

In his 1989 review of the state of the subject, [5], Stone divides the study of The Optimal Search Problem, referred to in generality as ‘Search Theory’, into four historical periods:

·Classical	1942-1965	·Algorithmic	1975-1985
·Mathematical	1965-1975	·Dynamic	1985-1989

The ‘Classical Period’ provided several of the constructions which still define many problems today, including the concept of search density and the exponential detection model. The subsequent ‘Mathematical Period’ focused on providing analytical tools, often in the form of necessary conditions on optimality derived using the calculus of variations. In order to create analytical tractability, targets were usually assumed to be either stationary or moving along a linear track at constant velocity. The ‘Algorithmic Period’ saw the further development of algorithms for finding optimal or near optimal

numerical solutions for simplified systems. Heuristic algorithmic methods were explored during this period, as well as the initial application of optimal control. The modern ‘Dynamic Period’ shifted the focus to dynamic feedback control solutions.

It is notable that, though written in 1989, until just the last few years Stone’s survey has remained an accurate description of the state of much of modern research into search theory, especially in the continuous time-and-space case. Though research on numerical methods has progressed significantly since 1989, many older questions have not been revisited. In particular, until 2011 [1], these concluding remarks from Stone’s review were still the case:

“As I look over the list of accomplishments and developments arising from search theory, I am struck by some missing elements. None of the standard optimization methods from nonlinear programming, control theory, or dynamic programming have made much of a contribution to solving search problems...These problems appear to occupy a small but special niche.” [5]

Thanks to the developments of [1] [7], [14], and [15], however, these remarks no longer hold. Due to these developments, and the progress made in numerical optimal control in general, we are now able to use more efficient numerical methods, such as pseudospectral collocation, than have been previously implemented. We are furthermore able to flexibly incorporate complicated objectives, state and control constraints, and nonlinear searcher and target dynamics.

In fact, the flexibility of modern numerical methods to solve optimal control problems with parameter uncertainty for a large class of objective functions and nonlinear state dynamics is a new development both in search theory and in optimal control. Many of these methods were originally developed for use on The Optimal Search Problem, however they have now created a foundation which can be leveraged to address a

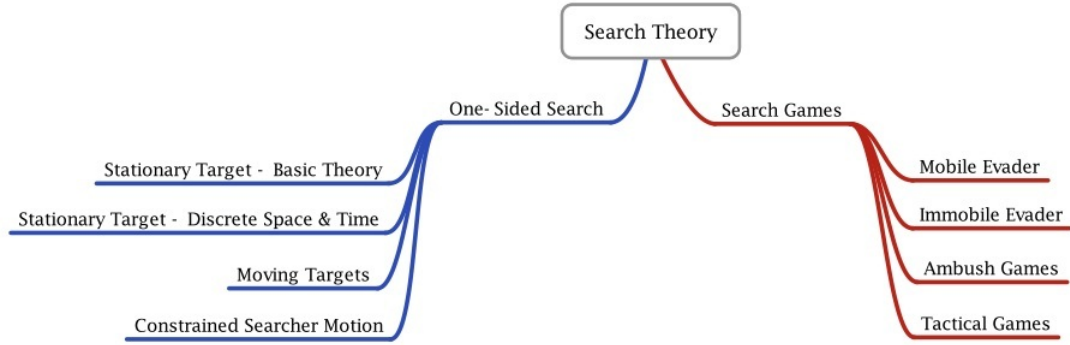


Figure 1.4: Illustration from Chung et al. (2011) [2] of a general taxonomy of search problems as outlined by Benkowski et al. (1991) [3]

greater variety of problems. One initial step in utilizing this potential is to expand the framework provided by The Optimal Search Problem to address more general multi-agent problems, with the goal of connecting the particular topic of optimal search with the broader field of pursuit-evasion research.

The modeling of optimal search problems and of pursuit-evasion scenarios share a common substrate: mobile autonomous agents with potentially incomplete knowledge of each other are cooperating or competing with the aim of optimizing some objective. Because of this similarity, the family of problems created by optimal search and pursuit-evasion scenarios can be interpreted as a single area of study. This is the approach taken by Benkowski et al. [3] and Chung et al. [2] who categorize pursuit-evasion problems under the umbrella of ‘search theory,’ using the taxonomy replicated here in Figure 1.4. Within this rubric, ‘one-sided search’ refers to the optimization of motion planning under the assumption of non-evasive target motion. ‘Search games’, also denoted as ‘adversarial search’, refers to the optimization of pursuit-evasion scenarios created by granting the adversary a higher degree of intelligence or responsiveness.

Despite these commonalities, however, the differing demands of problems with evasive targets versus those with non-evasive targets have so far divided the mathematical techniques applied into disparate approaches. Some of these differing demands can be seen to fundamentally require distinct approaches—for instance, in a pursuit-evasion differential game where the goal is to simultaneously find optimal strategies for both the pursuer and for the evader, the multi-objective optimization problem created will require a different approach than a single-objective optimization problem. However, there are many valuable problems with evasive targets which do result in single-objective optimization problems.

In the case of these problems, we would like to identify the design aspects of The Optimal Search Problem as derived above which may be open to elaboration or translation to these types of problems, as well as identify any complications that may arise numerically or theoretically by implementing expansions of the optimal search framework.

A More General Structure

From a modeling perspective, The Optimal Search Problem deals with two types of quantities. The first type is what Koopman refers to as the ‘kinematic bases’ of the problem.[8] These are the physically moving agents—the searchers and the uncertain targets. Attached to these bases is a set of performance measures—the probabilities of detection between the agents. These performance measures are subsequently used to create an objective function, J .

We now consider the task of modeling more “realistic scenarios,” using this

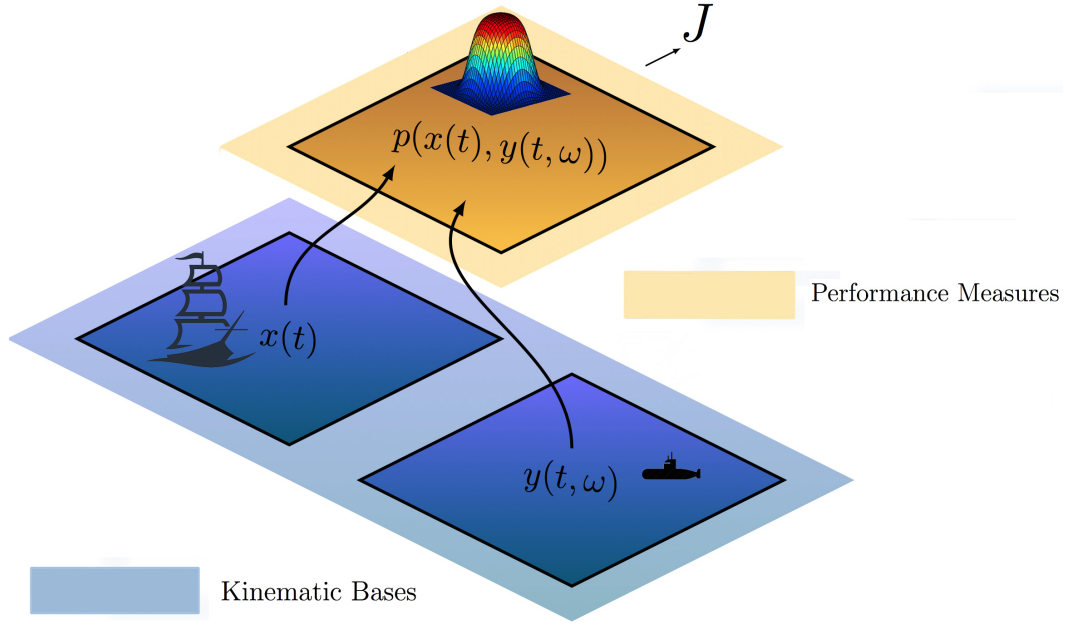


Figure 1.5: Diagram of the Types of Quantities Involved in Optimal Search

framework of kinematic bases and performance measures. As can be seen in the diagram, though The Optimal Search Problem has been challenging historically, as a model of interacting agents it is quite simple. The performance measures in the classic optimal search problem are created by merely two directions of influence—the influence of a searcher’s state on performance, and the influence of the conditional attacker states on performance. In reality, many scenarios require the modeling of more directions of influence than this simple model incorporates.

For example, consider the challenge of optimizing the trajectory of a ship with the mission of cleaning up an oil spill. The possible positions of the oil over time given unknown starting location but known ocean currents could be pre-computed and subsequently modeled as a conditionally deterministic function. If the goal of a clean-up was merely to “find” all the oil, then this problem would be a good candidate for

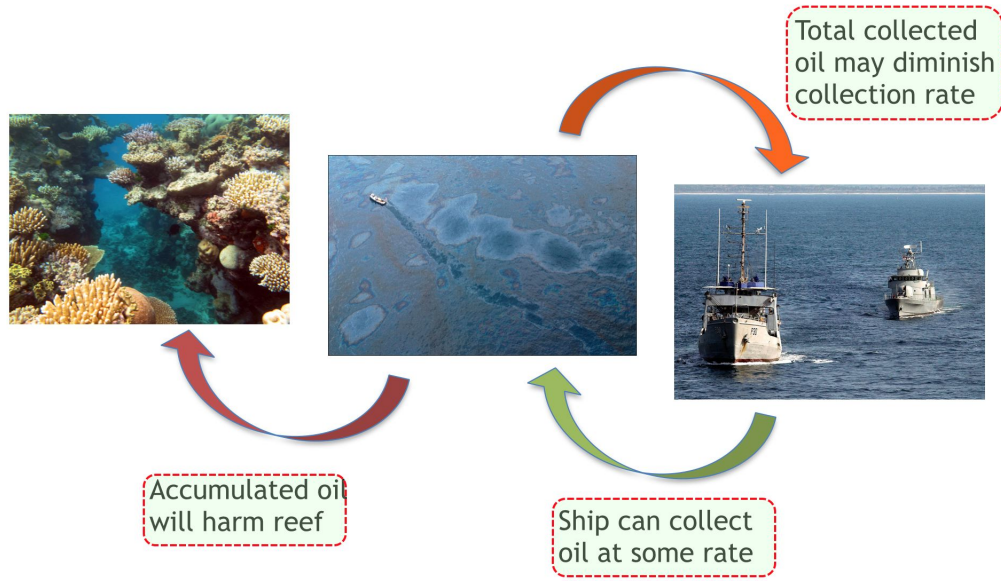


Figure 1.6: Diagram of Example Interactions Found in an Oil Cleanup

modeling as an optimal search problem. However, in many situations the objective of maximizing oil discovery would often not be a sufficiently nuanced strategy. For instance, the protection of a sensitive but valuable at-risk ecosystem may need to be prioritized during the clean up. In order to do this another layer of interaction needs to be modeled—that of the oil accumulating throughout the ecosystem, as influenced both by searcher states and discovered oil. We could furthermore imagine another layer of interaction—one in which the accumulated oil slows the ship, impacting searcher dynamics.

Another example can be found in the optimization of tactical trajectories. An area of interest in current Naval research is the development of guidance laws and optimal trajectories in situations with multiple defenders, multiple attackers, and a high-valued unit (HVV) that requires protection. This situation arises in tactical operations and in missile defense and guidance. Progress has been made both in creating heuristic




Figure 1.7: Multiple Influences Wanted

algorithms to tackle this scenario and in extending the results of optimal control theory to applicable formats. In [1], the problem was addressed using optimal search to maximize the probability of detecting attackers. However applying optimal search only indirectly optimizes HVU survival. A different approach would be to directly optimize the probability of survival by modeling the influence of the attackers on the HVU as well as the effectiveness of the defenders against the attackers.

The common theme is the need to expand the framework to include the possibility of interaction between relevant kinematic bases or performance measures. A basic challenge on this front is that, for the sort of multi-agent scenarios described above, models of agent interaction and objective functions for assessing their behavior simply haven't been developed. The exponential detection model provides us with a performance criterion which can be used as the objective function in the evaluation of optimal

search problems, however the derivation of objectives for more interactive situations is a more open question. Section 2 will develop several new models for interactive performance measures and kinematic bases and derive possible objective functions for such scenarios. Several new applications of optimal control with parameter uncertainty will then be put forth.

1.2.1 Mathematical Properties

The individual features of the kinematic bases in turn affect the features of the performance measures. For instance, conditioning the target motion y on an uncertain parameter ω makes the performance measures attached to y conditional on ω as well. The uncertainty is a transitive property which travels from the kinematic base to the performance measure. Another property which can be viewed as having a transitive effect—at least informally—is the differential constraint on $x(t)$. Cumulative performance measures by their nature depend on the entirety of the behavior of our state x from $t_0 = 0$ to t . This creates only two possibilities: either a performance measure, $P(t, \omega)$, is defined in terms of a function of an integral over time, i.e.:

$$P(t, \omega) = G \left(\int_0^t r(x(t), u(t), y(t, \omega)) dt \right) p(\omega)$$

or it is defined as a differential constraint as well:

$$\dot{P}(t, \omega) = g(P(t, \omega), x(t), u(t), t, \omega).$$

The third possibility, that the integral $\int_0^t r(x(t), u(t), t, \omega) dt$ could be solved for explicitly as a function of t and ω , is precluded by the fact that $\dot{x}(t)$ is subject to our control.

The classic optimal search problem, as derived above, in fact does end up with its performance measure defined as a differential equation. For instance, as we saw in the single search, single target case, the conditional probability $P(t, \omega)$ of not detecting the target by time t is given by the differential equation:

$$\dot{P}(t, \omega) = -P(t, \omega)r(x(t), y(t, \omega))$$

The problem then benefits from its structure, which allows for the exponential solution of this differential equation, i.e. a translation to the form:

$$G\left(\int_0^t r(x(t), u(t), y(t, \omega))dt\right)p(\omega).$$

The multi-agent ‘worst-case’ problem also takes advantage of this. The key effect of this is that the influence of uncertainty is never transmitted to the “state variables”—which in optimal control consists of all the variables which are defined in terms of a differential constraint. This is what allows The Optimal Search Problem to take the form it has been studied in—an optimal control problem with a nonstandard cost function but deterministic state variables. The reduction to this form, however, is a precarious feature since in general the differential equations involved in the performance measures may not yield a closed-form answer.

1.3 Thesis Contributions

The goal of this thesis is to facilitate the maturation of optimal control problems with parameter uncertainty into a tool of wider applicability. This is approached by assailing the three aspects identified as necessary for progressing the research devoted to The Optimal Search Problem: the design of useful interactive performance models and kinematic models, the development of a general framework for dealing with parameter uncertainty in both the cost and dynamics of optimal control problems, and a numerical algorithm for generating solutions.

1.3.1 Multi-Agent Models

Section 2 provides several new models for both kinematic and performance interaction in multi-agent systems. Several systems for herding—where the spatial trajectory of a target is influenced, with the goal of guiding the target towards desirable location—are developed. For performance interaction, a framework for modeling tactical attrition between multiple targets and agents is developed, leading to the publication of [15]. Additionally, models for a strategic search problem and a path coverage problem are provided.

1.3.2 Necessary Conditions for Optimality

In standard control problems, the necessary conditions for optimality provided by Pontryagin’s Minimum Principle, [16], provide tools for both theoretical and numerical analysis of solutions. Necessary conditions for sub-cases of the mathematical problem studied in this thesis were derived by [17] and [13]. Section 3 adds to this

collection of conditions by deriving necessary conditions optimality for a more general class of optimal control problems with parameter uncertainty in the case of open control regions. Stationarity conditions on a Hamiltonian function distinct from that found in standard control are provided both with respect to the control and with respect to time. These conditions complement the conditions which arise later, in Section 4.1.4, in that they apply even in the case of when the numerical algorithm fails to converge.

1.3.3 Computational Algorithm

Section 4 provides contributions to both the theoretical and practical sides of the issue of computation. A computational algorithm is given along with consistency results proving the optimality of convergent answers. As a joint work with Chris Phelps and Qi Gong at UCSC and Johannes Royset and Isaac Kaminer at the Naval Postgraduate School, the consistency of this algorithm when applied to optimal control problems with parameter uncertainty influencing the cost function has been analyzed in [14]. The consistency of these methods for problems with parameter uncertainty influencing both cost and dynamics is briefly considered in [18] by interpolating states and costates over the parameter domain and assuming the convergence of these interpolations. This thesis and the pending paper [19] further consider the consistency of this computational algorithm by analyzing consistency in terms of the states and costates propagated by control solutions rather than in terms of interpolations. The goal of the consistency results derived in this thesis is to be applicable to engineering applications, in which control solutions will be inputted and their effect propagated. Additional effort is also made to clarify the relationship of these results with the Covector Mapping Theory,

described in [20], which provides tools for numerical verification of solutions.

In addition to the theoretical results supporting the computational approach, a series of implementations is provided, to demonstrate both the computation efficacy of the algorithm and the behavior of problems of this type. These examples are implemented using a software package designed by the author which uses the methods described in Section 4. This software allows for a variety of options for discretization methods, and also automates the creation of features which are crucial to the efficient implementation such as gradient information and sparsity pattern. This software has been used to provide the numerical examples published in [14], [21], [18], and [15] and is used to provide further examples in Section 4.

1.4 Thesis Outline

The development of a family of interactive kinematic dynamics is the focus of Section 2.1 and the development of a family of interactive performance measures is the focus of Section 2.2.

Section 3 then puts forth a general mathematical framework for dealing with all of these scenarios. This framework places the scenarios in a class of nonstandard optimal control problems where dependence on an unknown parameter is a feature which can be found in both the cost function and the dynamics. Necessary conditions for optimality are derived using the method of the calculus of variations, for the case when controls are in an open set. An additional condition guaranteeing constant Hamiltonian value for optimal solutions is also derived. An example problem is then solved analytically

using these conditions.

Section 4 then introduces a numerical algorithm for finding solutions to this class of optimal control problems. Convergence properties of this method are proved for both the primal and dual problems. Subsequently, in Section 4.2, a variety of examples scenarios are implemented.

Lastly, Section 5 will discuss future directions of research.

Chapter 2

Kinematic and Performance Models

2.1 Kinematic Models

In addition to additional performance measures, to move towards modeling more ‘realistic’ scenarios, we would like to incorporate models of kinematic influence, where the agents under our control are able to physically influence the dynamics of the uncertain targets. Of particular interest is a type of interaction which can be referred to as ‘herding’—where the spatial trajectory of a target is influenced, with the goal of guiding the target towards desirable locations. Herding can be thought of as one-sided pursuit-evasion. The control inputs of the pursuer—i.e. the herder—are available for optimization whereas the evaders—i.e. the targets—respond in a purely reactive manner. This section will present several frameworks for modeling herding, with a focus on providing simple but flexible schemes which enable the use of calibratable nonlinear, multi-agent inputs.

Motivating Example: Radial Evasion System

A system for multi-agent herding is developed in the dissertation of [22]. In this model, the framework of [23], which presents herding dynamics for *fixed distances* between herder and target and *constant velocities* for both agents, is extended to encompass variable distances between agents and a distance-based velocity tactic. For a single target, single herder, the heading of the target in this framework is set to be along the path of sight to the herder, moving away from the herding. The heading angle, θ , then satisfies the following:

$$\tan \theta = \frac{x_2 - y_2}{x_1 - y_1}$$

and θ is determined by:

$$\theta = \tan^{-1} \left(\frac{x_2 - y_2}{x_1 - y_1} \right).$$

The velocity of the target speeds up as the herder approaches and is set as:




Figure 2.1: Geometry of Radial Evasion

$$v(r) = \frac{1}{r}$$




Figure 2.2: Geometry of the Derivative of a Trajectory in Relation to Heading Angle

where

$$r = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2}$$

Using the facts that $\tan \theta$ is additionally constrained by $\tan \theta = \frac{dy_2}{dy_1} = \frac{\dot{y}_2}{\dot{y}_1}$ (Figure 2.2) and that $\dot{y}_1^2 + \dot{y}_2^2 = v^2(r)$, a solution to the differential equations for \dot{y}_1 and \dot{y}_2 can be found by substitution:

$$\dot{y}_1^2 + \dot{y}_2^2 = \dot{y}_1^2 + (\tan^2 \theta) \dot{y}_1^2 = \dot{y}_1^2 \left(1 + \frac{\sin^2 \theta}{\cos^2 \theta} \right) = \frac{1}{r^2}$$

$$\implies \dot{y}_1^2 (\cos^2 \theta + \sin^2 \theta) = \dot{y}_1^2 = (\cos^2 \theta) \frac{1}{r^2}$$

and similarly

$$\dot{y}_2^2 = (\sin^2 \theta) \frac{1}{r^2}.$$

For target motion to be directed away from the herder, the solution is:

$$\begin{cases} \dot{y}_1 = -\frac{1}{r} \cos \theta \\ \dot{y}_2 = -\frac{1}{r} \sin \theta \end{cases} \quad (2.1)$$

Several methods are proposed in [22] for extending this to the case of multiple herders and targets. In the first method, herders and targets are identical in number and paired. In the second method, one herder may effect multiple targets, all driven by the same dynamics as described above. Finally, in the third method, the quantities $-\frac{1}{r} \cos \theta$ and $-\frac{1}{r} \sin \theta$ are thought of not as specifying a velocity and a heading, but as specifying the components of an individual herder's ‘influence.’ These influences are then added together, yielding the system:

$$\begin{cases} \dot{y}_{l,1} = -\sum_{k=1}^K \frac{1}{r_{lk}} \cos \theta_{lk} \\ \dot{y}_{l,2} = -\sum_{k=1}^K \frac{1}{r_{lk}} \sin \theta_{lk} \end{cases} \quad (2.2)$$

where r_{lk} and θ_{lk} are the relative radii and angles between the l -th target and the k -th herder.

This previously established model provides several key structural elements one might like to retain in subsequent models, such as the dependency on relative distance and position and considering herding influence in terms of inputs into the target's heading and velocity. However, it also contains several features which might be considered undesirable. One such feature is its use of $\tan \theta$ as the quantity which receives heading input. Since $\tan \theta$ is only unique in value within an interval of size π , the use of




Figure 2.3: Non-uniqueness of Tangent

$\tan^{-1} \left(\frac{x_2 - y_2}{x_1 - y_1} \right)$ to specify θ is only valid when θ is restricted to such an interval. The simulations in [22] did in fact remain in such a region, with targets always in front of herders and maintaining relative angles of less than $\frac{\pi}{2}$. However in large scale herding simulations, the possibility that the herder and target may switch orientations at some point cannot be ignored. Another feature which may benefit from modification is the method in which this herding model extends to multiple agents. Additive herding components have a possible physical analogue in the application of force, however their direct input into the velocity field as in this model has no straightforward interpretation. This makes limiting the modeled physical capabilities of the targets (as well as the herders) more difficult. For instance, in the multi-agent model above, the original relation to heading and velocity is complicated. The magnitudes of the ‘influence’ components, $\frac{1}{r_{lk}}$, no longer cancel out to yield the heading, $\frac{dy_{l,2}}{dy_{l,1}}$, as they did in the single target case, making the calculation of the heading laborious and constraints on its values difficult. Similarly, the trigonometric terms no longer cancel out in the velocity expression, $\dot{y}_{l,1}^2 + \dot{y}_{l,2}^2$, complicating velocity constraints.




Figure 2.4: Relative Position Vector and Normalized Relative Position Vector

The next sections will address some of these opportunities for development. First, we will reformulate this example model to avoid the restrictions tangent imposes and in doing so introduce new notation to simplify discussion. Then we will provide two flexible frameworks for modeling different systems of herding—one which allows for multi-agent interactions while retaining a hold on velocity and heading, and another which conceives of ‘influences’ as cumulative forces.

Relative Position Vectors

As mentioned above, the use of $\tan^{-1}\left(\frac{x_2-y_2}{x_1-y_1}\right)$ to specify θ is problematic due to the restricted interval in which $\tan \theta$ is unique. This can of course be addressed in a piecewise fashion, by choosing the appropriate solution of $\tan^{-1}\left(\frac{x_2-y_2}{x_1-y_1}\right)$ depending on the signs of $(x_1 - y_1)$ and $(x_2 - y_2)$. However, it would be preferable to have system equations which don’t require this sort of piecemeal approach, both for the sake of analytic evaluations and numerical efficiency. The obvious alternative to using the polar system (r, θ) to describe the relative position between agents is to describe the relative position in terms of vectors in Cartesian coordinates. Let the relative position

vector, R , be defined by:

$$R = \begin{pmatrix} R_1 \\ R_2 \end{pmatrix} = \begin{pmatrix} y_1 - x_1 \\ y_2 - x_2 \end{pmatrix}.$$

When R is normalized to a unit vector by dividing by $\|R\| = \sqrt{R_1^2 + R_2^2}$, we have:

$$\frac{R_1}{\|R\|} = -\cos \theta, \quad \frac{R_2}{\|R\|} = -\sin \theta.$$

We can thus translate the equations from the model above into a new set of equations, now in terms of the relative position vector.

$$\begin{cases} \dot{y}_1 = -\frac{1}{r} \cos \theta \\ \dot{y}_2 = -\frac{1}{r} \sin \theta \end{cases} \rightarrow \begin{cases} \dot{y}_1 = \frac{1}{\|R\|} \cdot \frac{R_1}{\|R\|} \\ \dot{y}_2 = \frac{1}{\|R\|} \cdot \frac{R_2}{\|R\|} \end{cases} \quad (2.3)$$

In the multi-agent case, relative position vectors will be considered between each target and herder. For targets y_l , $l = 1, \dots, L$ and herders x_k , $k = 1, \dots, K$ we can designate these as R_{lk} where

$$R_{lk} = \begin{pmatrix} R_{lk,1} \\ R_{lk,2} \end{pmatrix} = \begin{pmatrix} y_{l,1} - x_{k,1} \\ y_{l,2} - x_{k,2} \end{pmatrix}.$$




Figure 2.5: Relative Position Vectors Between Multiple Agents

2.1.1 Heading and Velocity Driven Herding

In the system of Equation 2.3, the dynamics of the system have two key components. The term $\frac{1}{\|R\|}$ gives the total velocity of the system and the relative position vector R determines the heading. We can generalize systems which specify these features—heading and velocity—with dynamics of the form:

$$\begin{cases} \dot{y}_1 = \frac{H_1(R)}{\|H(R)\|} v(R) \\ \dot{y}_2 = \frac{H_2(R)}{\|H(R)\|} v(R) \end{cases}$$

where

$$H(R) = \begin{pmatrix} H_1(R) \\ H_2(R) \end{pmatrix}, \quad v(R) > 0$$




Figure 2.6: Diagram of Heading Vector

and R is the relative position vector. The instantaneous total velocity of such a system is given by:

$$\begin{aligned}\sqrt{\dot{y}_1^2 + \dot{y}_2^2} &= \sqrt{\left(\frac{H_1(R)}{\|H(R)\|}\right)^2 v^2(R) + \left(\frac{H_2(R)}{\|H(R)\|}\right)^2 v^2(R)} \\ &= \sqrt{\frac{H_1^2(R) + H_2^2(R)}{\|H(R)\|^2} v^2(R)} = \sqrt{v^2(R)} = v(R)\end{aligned}$$

The slope of the tangent line, $\frac{dy_2}{dy_1}$, is given by $\frac{H_2}{H_1}$ with, as Figure 2.6 shows, the vector H then specifying the heading. The components H and v thus specify the heading and velocity of the target's trajectory. The radial evasion policy shown above, for example, is one instance of this form, with:

$$H = R, \quad v = \frac{1}{\|R\|}$$

For a K herders and L targets, this can be generalized as:

$$\begin{cases} \dot{y}_{l,1} = \frac{H_{l,1}(R_{l1}, \dots, R_{lK})}{\|H_l(R_{l1}, \dots, R_{lK})\|} v_l(R_{l1}, \dots, R_{lK}) \\ \dot{y}_{l,2} = \frac{H_{l,2}(R_{l1}, \dots, R_{lK})}{\|H_l(R_{l1}, \dots, R_{lK})\|} v_l(R_{l1}, \dots, R_{lK}) \end{cases}$$

$$H_l(R_{l1}, \dots, R_{lK}) = \begin{pmatrix} H_{l,1}(R_{l1}, \dots, R_{lK}) \\ H_{l,2}(R_{l1}, \dots, R_{lK}) \end{pmatrix}, \quad v_l(R_{l1}, \dots, R_{lK}) > 0, \quad l = 1, \dots, L$$

The next few sections will provide several examples of herding policies of this form.

Weighted Average Headings

One way to extend the radial evasion policy to multiple agents is to sum in some fashion the heading vectors contributed by each herder. This is comparable to the approach of Equation 2.2, which sums the contributions from each herder as “influences,” however it differs in that the effect can be limited to heading—the evasive velocity can be modeled separately.

As an example of heading summation, consider the situation diagrammed in Figure 2.7. In this scenario, two herders are located symmetrically on either side of a target at $(y_{l,1}, y_{l,2})$. Assuming both herders are considered of equal importance, then the direction of travel which maximizes the average distance from the herders is given by the sum of the two radial evasion heading vectors. In the scenario in Figure 2.7, the two herders are equidistant from the target, so a direct sum of each radial evasion heading vectors yields the optimal heading. However when the herders vary in distance, the relative position vectors need to be normalized before summation—or else herders *farther*




Figure 2.7: Example of Average Heading Policy with Two Herders

away would have more influence than closer herders due to their larger relative position vectors, a feature which is clearly unrealistic. Additionally, in most circumstances, greater emphasis might be placed on evading nearby herders. These two features can be incorporated in the following summation format:

$$H_l = \alpha_{l1} \frac{R_{l1}}{\|R_{l1}\|} + \alpha_{l2} \frac{R_{l2}}{\|R_{l2}\|} + \cdots + \alpha_{lK} \frac{R_{lK}}{\|R_{lK}\|}$$

where α_{lk} is some function of herder distance $\|R_{lK}\|$.

An alternate statement of this weighted average form can be made using potential functions. Let $V_{lk}(\|R_{lk}\|)$ be some function of radial distance. Then the gradient of V_{lk} with respect to the components of R_{lk} , i.e. $R_{lk,1}$ and $R_{lk,2}$, is:

$$\nabla V_{lk} = \left(\frac{\partial V_{lk}}{\partial R_{lk,1}}, \frac{\partial V_{lk}}{\partial R_{lk,2}} \right) = \left(\frac{dV_{lk}}{d\|R_{lk}\|} \cdot \frac{\partial \|R_{lk}\|}{\partial R_{lk,1}}, \frac{dV_{lk}}{d\|R_{lk}\|} \cdot \frac{\partial \|R_{lk}\|}{\partial R_{lk,2}} \right)$$

$$= \frac{dV_{lk}}{d\|R_{lk}\|} \left(\frac{\partial\|R_{lk}\|}{\partial R_{lk,1}}, \frac{\partial\|R_{lk}\|}{\partial R_{lk,2}} \right).$$

The partial derivative $\frac{\partial\|R_{lk}\|}{\partial R_{lk,1}}$ is given by:

$$\frac{\partial\|R_{lk}\|}{\partial R_{lk,1}} = \frac{\partial}{\partial R_{lk,1}} \left(\sqrt{R_{lk,1}^2 + R_{lk,2}^2} \right) = \frac{1}{2} \cdot \frac{1}{\sqrt{R_{lk,1}^2 + R_{lk,2}^2}} \cdot 2R_{lk,1} = \frac{R_{lk,1}}{\|R_{lk}\|}.$$

Similarly, $\frac{\partial\|R_{lk}\|}{\partial R_{lk,2}} = \frac{R_{lk,2}}{\|R_{lk}\|}$. Therefore

$$(\nabla V_{lk})^T = \frac{dV_{lk}}{d\|R_{lk}\|} \begin{pmatrix} R_{lk,1}/\|R_{lk}\| \\ R_{lk,2}/\|R_{lk}\| \end{pmatrix} = \frac{dV_{lk}}{d\|R_{lk}\|} \cdot \frac{R_{lk}}{\|R_{lk}\|}.$$

Thus we can also express a weighted average heading in terms of potential functions:

$$H_l = (\nabla V_{l1})^T + (\nabla V_{l2})^T + \cdots + (\nabla V_{lK})^T$$

where the potential functions are chosen such that $\frac{dV_{lk}}{d\|R_{lk}\|}$ provides a meaningful weight for a herder's relative effect at distance $\|R_{lk}\|$.

Example: Radially Decaying Heading Influence 1

This system provides an example of a herding system where the influence on a target's heading decays with distance from herder. For $k = 1, \dots, K$, $l = 1, \dots, L$, let

$$H_{lk} = \alpha_{lk} \frac{R_{lk}}{\|R_{lk}\|}, \quad \alpha_{lk} = \frac{c_{lk}}{(d_{lk}\|R_{lk}\| + 1)^2} \quad (2.4)$$

where c_{lk} and d_{lk} are calibrated constants. The maximum heading influence of each herder has a finite limit, c_{lk} , and monotonically decreases towards zero as radial distance increases. We may, as in the example above, also throw in an attractive herding influence. For instance, as a contrast, let

$$H_{l0} = -\frac{R_{l0}}{\|R_{l0}\|} \cdot c_{l0}$$

where c_{l0} is a calibrated constant. This provides a tracking influence that *doesn't* decay with distance—i.e the target always reorients towards its mission to track the HVU. The weighted average heading vector is given by:

$$H_l = \sum_{k=0}^K H_{lk}$$

This system can also be represented using potentials, with:

$$V_{l0}(\|R_{l0}\|) = -c_{l0}\|R_{l0}\|$$

and

$$V_{lk}(\|R_{lk}\|) = \frac{-c_{lk}/d_{lk}}{(d_{lk}\|R_{lk}\| + 1)}, \quad k = 1, \dots, K$$

Example: Radially Decaying Heading Influence 2

This system, developed by Panos Lambrianides, uses potential functions to model both the herding effects of K herders repelling L targets and, additionally, an attractive herding influence—the impulse of the targets to also ‘track’ towards an HVU,

with position given by $x_0 = (x_{0,1}, x_{0,2})$. For $k = 1, \dots, K$, let

$$V_{lk}(\|R_{lk}\|) = -\frac{1}{\sigma_D} e^{\left(\frac{-\|R_{lk}\|^2}{\sigma_D^2}\right)}$$

and for the case of the tracked HVU let:

$$V_{l0}(\|R_{l0}\|) = \frac{1}{\sigma_{HVU}} e^{\left(\frac{-\|R_{l0}\|^2}{\sigma_D^2}\right)}$$

where R_{l0} is the relative position vector between the l -th herder and the HVU at x_0 .

Then a weighted average heading vector incorporating both herding and tracking can be given by:

$$H_l = \sum_{k=0}^K (\nabla V_{lk})^T$$

2.1.2 Force Driven Herding

Another approach to modeling herding is through the application of physical force, directed from the herder at the target. For a single target moving in two-dimension space with relative position to a herder given by R , the creates a system of the form:

$$\left\{ \begin{array}{l} \dot{y}_1 = v_1 \\ \dot{y}_2 = v_2 \\ \dot{v}_1 = F_1(R) \\ \dot{v}_2 = F_2(R) \end{array} \right. \quad (2.5)$$

where the target is now modeled as a four dimensional system with both velocity and acceleration components and the herding input is now an influence on target acceleration. As mentioned in the previous section, though additive herding components are questionable when inputted into the velocity field, they have a physical foundation when it comes to the application of force. For K herders and L targets, additive herding components can be generalized as:

$$\left\{ \begin{array}{l} \dot{y}_{l,1} = v_{l,1} \\ \dot{y}_{l,2} = v_{l,2} \\ \dot{v}_{l,1} = \sum_{k=1}^K F_{lk}(R_{lk}) \\ \dot{v}_{l,2} = \sum_{k=1}^K F_{lk}(R_{lk}) \end{array} \right. \quad (2.6)$$

For radially directed force, the force functions can be decomposed in the same fashion as the velocity inputs in Section 2.1.1:

$$F_{lk}(R_{lk}) = m_{lk}(R_{lk}) \frac{R_{lk}}{\|R_{lk}\|}$$

where $m_{lk}(R_{lk})$ is the magnitude of force F_{lk} and $\frac{R_{lk}}{\|R_{lk}\|}$ gives the normalized direction vector. Section 5.3.1 discusses the implementation of herding models in a multi-agent tactical scenario.

2.2 Performance Models

To construct performance objectives for more ‘realistic’ multi-agent scenarios, the development of frameworks for multiple, interacting performance measures is needed. Given a set of kinematic bases with possible parameter uncertainty, $x(t, \omega)$ and $y(t, \omega)$, and a set of performance measures:

$$Q(t, \omega) = \begin{pmatrix} Q_1(t, \omega) \\ Q_2(t, \omega) \\ \vdots \\ Q_J(t, \omega) \end{pmatrix}$$

in full generality this concept could of course lead to any nonlinear performance dynamics:

$$\dot{Q}(t, \omega) = g(Q(t, \omega), x(t, \omega), y(t, \omega), u(t), t, \omega).$$

This section will narrow the focus down from that generality to a family of performance dynamics that arises from a broader application of the methods developed for The Optimal Search Problem. In regards to its performance metric, the method used by the The Optimal Search Problem is the exponential detection model, developed in [8]. In this model, the relatively small intervals between ‘glimpses of detection’ of sonar and radar sensors allow for the generalization to the use of a continuous rate of detection function, which yields a difference equation of the form:

$$Q_j(t + \Delta t, \omega) = Q_j(t, \omega) [1 - r(x(t, \omega), y_j(t, \omega))\Delta t]$$

where $Q_j(t, \omega)$ is the probability of non-detection for the j -th target. As $\Delta t \rightarrow 0$ becomes:

$$\dot{Q}_j(t, \omega) = -Q_j(t, \omega)r(x(t, \omega), y_j(t, \omega)).$$

The product of these states is then solved in closed form as an exponential expression. In The Optimal Search Problem, this last step, solving as an exponential, is possible because the probabilities Q_j are decoupled—the probability of one target remaining undetected doesn't depend on the success of the others. If the rate of detection were not decoupled, we would instead have performance dynamics of the following form:

$$\dot{Q}_j(t, \omega) = -Q_j(t, \omega)r(Q(t, \omega)x(t, \omega), y_j(t, \omega)).$$

Outside of the application of search and the act detection, we can identify a more general concept at play in this equation: attrition. Attrition is meant in the following sense:

attrition: *the action or process of gradually reducing the strength or effectiveness of someone or something through sustained attack or pressure*

With this terminology, for example, the rate of detection function described in The Optimal Search Problem would be interpreted as the rate of attrition of a target's probability of remaining undetected. When this concept is generalized, we find that the design effort put into The Optimal Search Problem can be leveraged to address a wide range of problems and that this method of modeling through attrition, calibrated rate functions, and optimal control is applicable to many relevant gauges. The next few sections will explore the application of this method to several areas: search, path coverage problems, and tactical 'shooting' problems.

2.2.1 Strategic Search

Target evasiveness is a reality of many real-life search problems. While search can apply to passive or willing targets, oftentimes the reason the target has not been found is because it doesn't want to be. In such cases the plan of detection involves a trade-off between proximity to the target and covertness of the act of detection.

An example of such a scenario can be found in the goal of detecting covert operation bases, whose presence may be revealed through voluntary emissions, such as heat signatures. For the sake of simplicity, let us imagine these emissions as binary decisions on the part of the base: either the operation continues, or it abruptly shuts down due to the detection of unwanted searching agents. Let us also assume that the act of detecting these unwanted agents can be modeled using the principles of the classic optimal search problem.

That is, given an uncertain but conditionally deterministic base location, $y(\omega)$, and searcher states (position, for instance, and perhaps velocity) $x(t)$, the base's detection capabilities can be modeled using an instantaneous rate function, $r_{base}(y(\omega), x(t))$ such that in a small time interval $[t, t + \Delta t]$ the probability of agent detection can be modeled by $r_{base}(y(\omega), x(t))\Delta t$. Then, the probability of base remaining unaware of the searching agents, represented by $Q(t, \omega)$, as in the previous search derivation obeys the following difference equation:

$$Q(t + \Delta t, \omega) = Q(t, \omega) [1 - r_{base}(y(\omega), x(t))\Delta t]$$

which as Δt goes to zero has again the exponential solution

$$Q(t, \omega) = e^{-\int_0^t r_{\text{base}}(y(\omega), x(\tau)) d\tau}$$

or alternately:

$$\dot{Q}(t, \omega) = -Q(t, \omega) r_{\text{base}}(y(\omega), x(\tau)).$$

The probability that the base is still in operation is equivalent to $Q(t, \omega)$.

The act of the searcher's detection of the base is now predicated on two things: its innate capability of detecting an active base and the probability that the base is still in operation. Thus if the innate rate of detection for the searching agents is given by $r_{\text{searcher}}(x(t), y(\omega))$, then the overall instantaneous rate of detection is given by $Q(t, \omega) r_{\text{base}}(x(t), y(\omega))$. Letting $P(t, \omega)$ represent the probability that the searcher has not detected the base by time t , we have:

$$P(t + \Delta t, \omega) = P(t, \omega) [1 - Q(t, \omega) r_{\text{searcher}}(x(t), y(\omega)) \Delta t]$$

which as Δt goes to zero becomes:

$$P(t, \omega) = e^{-\int_0^t Q(t, \omega) r_{\text{searcher}}(y(\omega), x(\tau)) d\tau}$$

or:

$$\dot{P}(t, \omega) = -P(t, \omega) Q(t, \omega) r_{\text{searcher}}(y(\omega), x(\tau)).$$

The final problem can be written in two forms:

Problem B1: Given $p : \Omega \rightarrow \mathbb{R}$, determine the control $u : [0, T] \rightarrow U \in \mathbb{R}^{n_u}$ that minimizes the cost:

$$J = \int_{\Omega} \left[1 - e^{-\int_0^t \left(e^{-\int_0^s r_{\text{base}}(y(\omega), x(s)) ds} \right) r_{\text{searcher}}(y(\omega), x(\tau)) d\tau} \right] p(\omega) d\omega$$

subject to:

$$\begin{aligned} \dot{x}(t) &= f(x(t), u(t)), & x(0) &= x_0 \\ g(u(t)) &\leq 0, \forall t \in [0, T] \end{aligned}$$

Problem B2: Given $p : \Omega \rightarrow \mathbb{R}$, determine the control $u : [0, T] \rightarrow U \in \mathbb{R}^{n_u}$ that minimizes the cost:

$$J = \int_{\Omega} [1 - P(x(T, \omega), \omega)] p(\omega) d\omega$$

subject to:

$$\dot{Q}(t, \omega) = -Q(t, \omega)r_{\text{base}}(y(\omega), x(\tau)), \quad Q(0, \omega) = 0 \quad \forall \omega \in \Omega$$

$$\dot{P}(t, \omega) = -P(t, \omega)Q(t, \omega)r_{\text{base}}(y(\omega), x(\tau)), \quad P(0, \omega) = 0 \quad \forall \omega \in \Omega$$

$$\dot{x}(t) = f(x(t), u(t)), \quad x(0) = x_0$$

$$g(u(t)) \leq 0, \quad \forall t \in [0, T]$$

2.2.2 Shooting Problems

A problem which arises in combat modeling is the optimization of strategies given antagonistic weapons fire. Objectives in such scenarios could include optimizing average agent survival numbers, for instance, or optimizing the survival of a specific strategic unit. In either case, the quantity in need of modeling is the probability of survival given weapon accuracy and damage magnitudes.

In [24], Washburn models this probabilistic survival using the concept of a ‘damage function’ $D(x, y)$. The damage function returns the probability that a target at y is destroyed by a shot fired from x . The value of the damage function is taken to depend on radial distance. The concept of the damage function in [24] is developed solely for the case of a single shot fired between two stationary agents, however its basic idea can also be applied to rapid-fire situations with moving agents.

Consider the case of fire rapid enough to be modeled as a continuous barrage of damage. The probability of destruction of given by the rate of damage and the time

spent in its presence. As such, the probability of destruction in a time interval $[t, t + \Delta t]$ can be modeled in the form:

$$r_D(x(t, \omega), y(t, \omega))\Delta t$$

where r_D is a *rate of damage* function, i.e. an *attrition rate*, calibrated to depend on factors like distance and relative location and orientation. For a target at location $y(t, \omega)$, let $P(t, \omega)$ be the probability that the target has survived at time t . If there are no other mitigating probabilities involved (such as the probability of shooter survival), then $P(t + \Delta t, \omega)$ is given by:

$$P(t + \Delta t, \omega) = P(t, \omega) [1 - r_D(x(t, \omega), y(t, \omega))\Delta t]$$

which as $\Delta t \rightarrow 0$ becomes:

$$\dot{P}(t, \omega) = -P(t, \omega)r_D(x(t, \omega), y(t, \omega)).$$

Compound probabilities, between for instance agents whose probability of survival is coupled with the survival probabilities of other agents, can be derived using these methods as well. Elaborations are provided in the scenarios in Sections 2.2.2.1 and 2.2.2.2.

Designing Damage Rate Functions

To model shooting problems, a variety of the specific properties of the damage rates between agents involved may be desirable to calibrate. For instance, if one is modeling a hit-or-miss situation—where an agent hitting a target results in probable

destruction but missing it causes little to no damage—then the range of effective fire would be limited to a small space around the respective attacker, to simulate that the attacker can only impart damage when it is very close to the target. On the other hand, if attackers are vehicles or ships with long-range weapons capabilities, then the distance from the HVU at which they become effective may be larger. Another relevant feature is possible point-of-view (POV) constraints, which limit the angles of effective fire. Finally, any model for shooting problems must be able to address the divided effort available per target in situations where firing is constrained to be allotted only towards single targets at any one moment.

Fortunately, large classes of continuous functions exist which can be calibrated to model these features. For example, an available option is to follow the Poisson Scan Model and use a cumulative normal distribution centered around agent position. Parameters of the distribution can be modified to change the steepness of this function to be simulate the decay of a firing rate function over distance. Benefits of this choice are that the distribution is normalized and existing literature on radar and sonar detection functions may be helpful in calibrating parameters. Another family of functions available




Figure 2.8: Beta functions can be used to simulate a maximal firing effectiveness at some distance from the firing agent




Figure 2.9: “Half Beta” functions can be used to simulate decreasing firing effectiveness over distance capped by a finite range

for use is the family of beta functions. A beta function with positive parameters $\alpha, \beta \geq 2$ can be connected smoothly (C^1) as a piecewise function:

$$f(x) = \begin{cases} x^\alpha(1-x)^\beta, & 0 \leq x \leq 1 \\ 0, & \text{else} \end{cases}$$

This can be modified to have arbitrary radius and magnitude and can then be centered around an appropriate point.

The finite radius of the beta function and the ease of algebraically manipulating aid modeling possibilities. For instance, the beta function extended through radial symmetry creates a donut shape which can be used to model situations where offense effectiveness is maximized at some distance from the agent. This distance can be chosen as a parameter of the model. Beta functions can also be used to mimic the cumulative normal distribution, but with the added advantage of having a finite radius of effectiveness. This is created taking the “half beta” function:

$$f(x) = \begin{cases} (x+c)^\alpha(1-(x+c))^\beta, & c \leq x \leq 1 \\ 0, & \text{else} \\ c = \frac{\alpha}{\alpha+\beta} \end{cases}$$

and extending it through radial symmetry. In addition to modeling different radii and peaks of effectiveness, smooth attrition rate functions can be created which model point-of-view (POV) constraints. This can be accomplished by including an angularly decaying multiplier which decreases either with arc length or arc angle.




Figure 2.10: Effect of an angularly decaying multiplier on an example attrition rate function




Figure 2.11: Resulting angularly decaying function reflecting POV limitations

The issue of divided fire effort is more subtle, as smooth functions are desired for optimal control implementation. This prevents the use of switching structures which move an agent's attention discretely from one target to another or which divide attention between the number of targets within a bounded range. However, we can employ methods to smoothly approximate divided attention. For example, using the notation that $x_k(t)$ is the location of the k -th defender and $y_l(t)$ is the location of the l -th attacker, we can divide a defender's attrition function by the following sum of Gaussian distributions:

$$\sum_{l=1}^L \Phi \left(\frac{\rho - \|x_k(t) - y_l(t)\|}{\sigma} \right)$$

where Φ is a cumulative normal distribution. With the standard deviation σ set to a small number, this will closely approximate the number of attackers within a radius ρ of defender x_k at time t .

2.2.2.1 The Kamikaze Swarm Scenario

We consider the scenario where a swarm of L attackers is headed towards a moving high-valued unit (HVU). K defenders are dispatched to neutralize the attackers before they destroy the HVU. The objective now is to minimize the probability of destruction of the HVU over the time interval $[0, T]$. Each defender's state, x_k , is governed by the dynamics:

$$\dot{x}_k = f_k(x_k, u_k), \quad x_k(0) = \xi_k \quad (2.7)$$

These are referenced as the single vectors $x = [x_1, \dots, x_K]^T \in \mathbb{R}^{N_x}$, $u = [u_1, \dots, u_K]^T \in \mathbb{R}^{N_u}$, and $f(x, u) = [f_1(x_1, u_1), \dots, f_K(x_K, u_K)]$. The trajectory of the HVU is not influenced by a control and is set by a function $x_0(t) \in \mathbb{R}^{N_x}$. Attacker trajectories are influenced by parameters $\omega \in \Omega$ with probability density function $\phi : \Omega \rightarrow \mathbb{R}$. Each attacker has a conditionally deterministic trajectory given by $y_l(t, \omega)$. The positions for all L attackers are also combined into the single vector $y = [y_1, \dots, y_L]^T \in \mathbb{R}^{N_y}$.

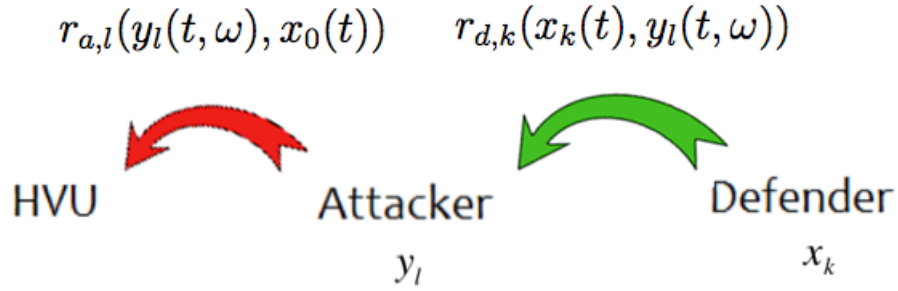


Figure 2.12: Structure of Kamikaze Scenario

For this ‘kamikaze’ situation, we will assume a single-minded focus of respective agents. The attackers have as their sole mission the destruction of the HVU, with no

firepower to spare on other agents. They focus all their fire on the HVU, hoping to evade the defenders long enough to succeed in its destruction. The defenders focus their fire on the attackers with the goal of protecting the HVU. The HVU relies on the protection of the others and is unable to fire on the attackers itself. We will assume that the firing rates of the agents are rapid and as such can be modeled as continuous quantities. Firing rates are interpreted such that if $r(t)$ is the instantaneous rate of fire directed against an agent at time t , then the probability of an agent's destruction in a sufficiently small time interval $[t, t + \Delta t]$ is given by the quantity $r(t)\Delta t$. Let $r_{d,k}(x_k(t), y_l(t, \omega))$ represent the rate of fire of the k -th defender against the l -th attacker for a parameter value $\omega \in \Omega$. The probability that the l -th attacker survives at time $t + \Delta t$ conditional on ω is then given by

$$q_l(t + \Delta t, \omega) = q_l(t, \omega) \prod_{k=1}^K (1 - r_{d,k}(x_k(t), y_l(t, \omega))\Delta t)$$

which becomes:

$$q_l(t + \Delta t, \omega) = q_l(t, \omega) \left(1 - \sum_{k=1}^K r_{d,k}(x_k(t), y_l(t, \omega))\Delta t + \text{h.o.t.} \right).$$

We take the limit and find that this results in a proportional attrition rate:

$$\dot{q}_l(t, \omega) = -q_l(t, \omega) \sum_{k=1}^K r_{d,k}(x_k(t), y_l(t, \omega))$$

and yields the expression:

$$q_l(t, \omega) = e^{-\int_0^t \sum_{k=1}^K r_{d,k}(x_k(\tau), y_l(\tau, \omega)) d\tau}.$$

We now let $r_{a,l}(y_l(t, \omega), x_0(t))$ be the rate of fire of the l -th attacker, if it has survived, against HVU. The probability of destruction of the HVU in a small time interval $[t, t+\Delta t]$ is determined by the rate of possible fire against it compounded with the probability that the attackers have survived to emit that firepower. Thus the probability that the HVU survives at time $t + \Delta t$ is given by:

$$p(t + \Delta t, \omega) = p(t, \omega) \prod_{l=1}^L (1 - q_l(t, \omega) r_{a,l}(y_l(t, \omega), x_0(t)) \Delta t).$$

After similar manipulations to those above this yields:

$$\begin{aligned} p(t, \omega) &= e^{-\int_0^t (\sum_{l=1}^L q_l(\tau, \omega) r_{a,l}(y_l(\tau, \omega), x_0(\tau))) d\tau} \\ &= e^{-\int_0^t \left(\sum_{l=1}^L r_{a,l}(y_l(\tau, \omega), x_0(\tau)) e^{-\int_0^\tau \sum_{k=1}^K r_{d,k}(x_k(s), y_l(s, \omega)) ds} \right) d\tau} \end{aligned}$$

Section 4.2.4 provides an implementation of this Kamikaze Swarm model.

Optimal Control Formulation

Optimizing the expectation of this probability over Ω creates a nonstandard optimal control problem of the same form as that of problem **P**:

The Kamikaze Shooting Problem: Given the probability density function $\phi : \Omega \rightarrow \mathbb{R}$ and conditionally deterministic attacker trajectories $y(t, \omega)$, determine the control

$u : [0, T] \rightarrow U \in \mathbb{R}^{n_u}$ that minimizes the expectation

$$J = \int_{\Omega} G \left(\int_0^T r(x(\tau), y(\tau), \omega) d\tau \right) \phi(\omega) d\omega \quad (2.8)$$

where

$$r(x(\tau), y(\tau), \omega) = \sum_{l=1}^L r_{a,l}(y_l(\tau, \omega), x_0(\tau)) e^{-\int_0^{\tau} \sum_{k=1}^K r_{d,k}(x_k(s), y_l(s, \omega)) ds}$$

and

$$G(z) = 1 - e^{-z}$$

subject to the searcher dynamics

$$\dot{x}(t) = f(x(t), u(t))$$

with initial condition $x(0) = x_0$ and control constraint $g(u(t)) \leq 0, \forall t \in [0, T]$.

A notable feature of this problem is that, as with The Optimal Search Problem, it is still possible to transport the influence of the uncertain parameter entirely to the cost function. This is convenient for numerical implementation, as uncertainty in the state dynamics increases the number of decision variables created by the numerical algorithm described in Section 4. The next section will look at mutual attrition, where that capability is no longer the case.

2.2.2.2 Mutual Attrition

We now consider again the case of L attackers, K defenders, and an HVU, however with the increased symmetry that attackers are able to fire on both the HVU or the defenders. As with Section 2.2.2.1, defender states $x_k(t)$ are deterministic and governed by the dynamics of Equation 2.7, attacker trajectories are conditionally deterministic and given by $y_l(t, \omega)$, and the objective is to minimize the probability of destruction of the HVU over the time interval $[0, T]$.




Figure 2.13: Structure of Mutual Attrition Scenario

Let $r_{d,k}(x_k(t), y_l(t, \omega))$ again represent the rate of fire of the k -th defender against the l -th attacker. In contrast to Section 2.2.2.1, the probability that the l -th attacker survives at time $t + \Delta t$ is now conditional on two quantities: the probabilities of destruction caused by fire *and* the probability that the defender is still alive to emit that fire. Let $q_l(t, \omega)$ be the probability that the l -th attacker has survived at time t and let $p_k(t, \omega)$ be the probability that the k -th defender has survived. Then the probability

that the l -th attacker has survived by time $t + \Delta t$ is given by:

$$q_l(t + \Delta t, \omega) = q_l(t, \omega) \prod_{k=1}^K (1 - p_k(t, \omega) r_{d,k}(x_k(t), y_l(t, \omega)) \Delta t)$$

which becomes:

$$q_l(t + \Delta t, \omega) = q_l(t, \omega) \left(1 - \sum_{l=1}^K p_k(t, \omega) r_{d,k}(x_k(t), y_l(t, \omega)) \Delta t + \text{h.o.t.} \right).$$

and has the limit as $\Delta t \rightarrow 0$ of:

$$\dot{q}_l(t, \omega) = -q_l(t, \omega) \sum_{k=1}^K p_k(t, \omega) r_{d,k}(x_k(t), y_l(t, \omega)) \quad (2.9)$$

Let $r_{a,l}(y_l(t, \omega), x_k(t))$ be the rate of fire of the l -th attacker against the k -th defender.

Just as with the attacker survival probability, the defender survival probability depends both on the chance of damage by fire and the probability that the attacker survives to emit damage. Thus the probability that the k -th defender has survived by time $t + \Delta t$ is given by:

$$p_k(t + \Delta t, \omega) = p_k(t, \omega) \prod_{l=1}^L (1 - q_l(t, \omega) r_{a,l}(y_l(t, \omega), x_k(t)) \Delta t).$$

which has the limit:

$$\dot{p}_k(t, \omega) = -p_k(t, \omega) \sum_{l=1}^L q_l(t, \omega) r_{a,l}(y_l(t, \omega), x_k(t)). \quad (2.10)$$

Finally, $r_{0,l}(y_l(t, \omega), x_0(t))$ be the rate of fire of the l -th attacker against the HVU and

let $p_0(t, \omega)$ be the probability that the HVU has survived at time t . Following the arguments above, $p_0(t, \omega)$ obeys the dynamics:

$$\dot{p}_0(t, \omega) = -p_0(t, \omega) \sum_{l=1}^L q_l(t, \omega) r_{0,l}(y_l(t, \omega), x_0(t)). \quad (2.11)$$

Because Equations 2.9 and 2.10 are coupled, the explicit exponential solutions of Section 2.2.2.1 are no longer an option. Thus, the dynamics presented in Equations 2.9 - 2.11 must remain as part of the resulting problem's state dynamics.

Mutual Attrition Scenario: Given the probability density function $\phi : \Omega \rightarrow \mathbb{R}$ and conditionally deterministic attacker trajectories $y(t, \omega)$, determine the controls $u_k : [0, T] \rightarrow U \in \mathbb{R}^{n_u}$, $k = 1, \dots, K$, that minimize:

$$J = \int_{\Omega} [1 - p_0(T, \omega)] p(\omega) d\omega$$

subject to:

$$\begin{aligned} \dot{x}_k(t) &= f(x_k(t), u_k(t)), & x_k(0) &= x_k^0 \\ \dot{p}_0(t, \omega) &= -p_0(t, \omega) \sum_{l=1}^L q_l(t, \omega) r_{0,l}(y_l(t, \omega), x_0(t)), & p_0(0, \omega) &= 0 \\ \dot{p}_k(t, \omega) &= -p_k(t, \omega) \sum_{l=1}^L q_l(t, \omega) r_{a,l}(y_l(t, \omega), x_k(t)), & p_k(0, \omega) &= 0 \\ \dot{q}_l(t, \omega) &= -q_l(t, \omega) \sum_{k=1}^K p_k(t, \omega) r_{d,k}(x_k(t), y_l(t, \omega)), & q_l(0, \omega) &= 0 \end{aligned}$$

for all $l = 1, \dots, L$, $k = 1, \dots, K$, and $\omega \in \Omega$ with possible control constraint $g(u(t)) \leq$

0, $\forall t \in [0, T]$. Section 5.3.1 discusses the implementation of this mutual attrition model in a multi-agent tactical scenario.

2.2.3 Path Coverage

Another type of problem amenable to modeling through attrition which has arisen repeatedly in modern robotics is that of ‘path coverage,’ where the goal of a robot is to ‘cover’ in some fashion all points in a region. 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 [25]. However, though global coverage plans have been developed which provide assurances of complete or near-complete coverage (see [26] for a survey of results in this area), the question of optimality is still largely unaddressed.

It is relevant to point out that the optimal search problem is, in essence, also a path coverage problem. Given a region to be searched, the goal of the optimal search problem is to ‘cover’ it with sensor readings in a fashion which minimizes un-missed spots (targets). As such, the methods of the optimal search problem can be naturally extended to cover other path coverage problems.

For instance, consider the example of an automated vacuum cleaner. To begin with, we can consider a simplified model of a vacuum: A vacuum moving over a two-dimensional surface Ω with a probabilistic rate of success determined solely by the vacuum’s innate capabilities (such as airflow). In this simple model, the vacuum can be characterized as having a certain chance for picking up any dust which lies below the

vacuum's active aperture, determined by the vacuum's innate success rate, say r_0 , such that if the vacuum's aperture is over an region of dust, then the probability in a small time interval $[t, t + \Delta t]$ of picking up a particular mote of dust is $r_0 \Delta t$. For a dust mote at location $\omega \in \Omega \subset \mathbb{R}^2$ and vacuum at position x , the rate of success is thus:

$$r(x, \omega) = \begin{cases} r_0 & \text{if } \omega \text{ is in aperture determined by } x \\ 0 & \text{else} \end{cases}$$

This rate of success function can be smoothed to provide a continuous rate function, as demonstrated in Figure 2.14. If there is initially dust at ω , then the probability of an




Figure 2.14: Example: Smoothed Aperture-Based Rate Function

unclean spot at location ω and time t is given by:

$$P(t + \Delta t | \text{initial dust at } \omega) = P(t | \text{initial dust at } \omega) [1 - r(x(t), \omega) \Delta t]$$

which as $\Delta t \rightarrow 0$ has the solution:

$$P(t|\text{initial dust at } \omega) = e^{-\int_0^t r(x(\tau), \omega) d\tau}.$$

The total probability of an unclean spot at location ω and time t is given by the joint probability of there having been dust at ω to begin with and of dust there having not been picked up. Letting $\phi(\omega)$ be a density distribution for dust in Ω , then the expectation for a clean room by time T is given by:

$$J = 1 - \int_{\Omega} e^{\int_0^T r(x(t), \omega) dt} \phi(\omega) d\omega$$

This provides a cost function for optimizing a global plan, given the dynamical constraints of the vacuum. Elaborations on this simplified model can be undertaken to increase its realism. For instance, suppose dust is not in fact constrained to a two-dimensional sheen but is present in varying levels throughout the room. A prior probability density distribution is constructed (using, for instance, the vacuum's past history in a room and the mapping methods of [25]) which provides a starting distribution $\phi(\omega_1, \omega_2, \omega_3)$ over three components: two location variables, ω_1 and ω_2 , and vertical density range ω_3 . The rate of vacuum success is then modeled with vertical density as a penalty. For example:

$$r(\omega_1, \omega_2, \omega_3) = \tilde{r}_0(x, \omega_1, \omega_2) e^{-c_0 \omega_3}$$

with

$$\tilde{r}_0(x, \omega_1, \omega_2) = \begin{cases} r_0 & \text{if } [\omega_1, \omega_2] \text{ is in aperture determined by } x \\ 0 & \text{else} \end{cases}$$

Chapter 3

Analytical Foundations

3.1 General Problem Statement

The Optimal Search Problem provides an initial template for optimal control problems with parameter uncertainty. A feature we can identify as fundamental to such problem is the inclusion of an extra layer of integration over parameter space—this is necessary for the evaluation of useful statistics of the problem, such as expectation and variance. On the other hand, as the previous examples of parameter uncertainty demonstrate, dependence on the unknown parameter is a feature which can be found in both the cost function and the dynamics. This thesis will focus on the following class of generalized optimal control problem for parameter uncertainty, which incorporates both of these features:

Problem P: Determine the control $u : [0, T] \rightarrow \mathbb{R}^{n_u}$ that minimizes the cost:

$$J = \int_{\Omega} \left[F(x(T, \omega), \omega) + \int_0^T r(x(t, \omega), u(t), t, \omega) dt \right] d\omega \quad (3.1)$$

subject to the dynamics:

$$\dot{x}(t, \omega) = f(x(t, \omega), u(t)), \quad x(0, \omega) = x_0(\omega) \quad (3.2)$$

and control constraints

$$g(u(t)) \leq 0, \quad \forall t \in [0, T] \quad (3.3)$$

with $x : [0, T] \times \Omega \rightarrow \mathbb{R}^{n_x}$, $F : \mathbb{R}^{n_x} \times \Omega \rightarrow \mathbb{R}$, $r : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times [0, T] \times \Omega \rightarrow \mathbb{R}$, $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$, and $g(u(t)) : \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{\mathbb{N}_g}$.

The format of the cost functional in this class of problems is that of the integral over Ω of a Mayer-Bolza type cost with parameter ω , i.e. the integral over Ω of a cost of the form:

$$F(x(1, \omega), \omega) + \int_0^1 r(x(t, \omega), u(t), t, \omega) dt.$$

This parameter can represent a range of values for a feature of the system or a time-invariant stochastic parameter with a known probability density function. In the case of a stochastic parameter, this format, of the integration over Ω of a Mayer-Bolza cost, is one which can encompass expressions of quantities such as the expectation or variance of costs which depend on a parameter θ . In such cases, the probability density function becomes a component of $F(\cdot)$ and $r(\cdot)$.

Instances of this class of problems have arisen in multiple recent applications. For example, the ensemble control problem [27, 28, 29], deals with the control of a family

or continuum of systems $x(t, \omega)$ whose dynamics depend continuously on a parameter ω . That is, $\dot{x}(t, \omega) = f(x(t, \omega), u(t), \omega)$. Optimizing the behavior of these systems over all parameter values creates control problems in the form of Problem **P**. An additional example can be found in chemical engineering, where instances of Problem **P** have been utilized for the optimization of batch processes under uncertainty [30, 31].

Relation to The Optimal Search Problem

Problem **P** and The Optimal Search Problem differ in the inclusion of uncertainty in the dynamics, but also, seemingly, in the form of their cost function. The Optimal Search Problem's cost function included the term:

$$G\left(\int_0^T r(x(t), u(t), t, \omega) dt\right)$$

with the time integral contained by an outer function, $G(\cdot)$, whereas Problem **P** contains no such outer function. It is worth noting, however, that this latter difference is superficial. When the additional flexibility of uncertain dynamics is allowed, and problem with such a term in the cost function can be transformed to fit into problem **P**. This can be accomplished by augmenting the state space with the following auxiliary variable:

$$\dot{z}(t, \omega) = r(x(t, \omega), u(t), t, \omega), \quad z(0, \omega) = 0, \quad \forall \omega \in \Omega.$$

When written in terms of this variable, the optimal search cost function becomes:

$$J = \int_{\Omega} \{F(x(T, \omega) + G(z(T, \omega)))\} d\omega$$

which is an end cost function that fits in the format of Problem **P**. From this transformation, we see that The Optimal Search Problem is a sub-case of this more general class of problems.

3.2 Necessary Conditions for Optimality

Variations of this problem were studied extensively in the 1970's, with a focus on deriving analytical conditions for optimality. These variations were almost exclusively studied in reference to the optimal search problem, in which controlled agents were deterministic and targets were affected in some fashion by uncertainty. In this pursuit, two relevant results were derived which apply to cases which overlap with Problem **P**.

The first is that of Pursiheimo, [17], who developed necessary and sufficient conditions for optimality for a search problem in which ‘search effort’ $r(\cdot, t)$ was available for allocation. Unlike the optimal search problem described in the Introduction, this search effort was not constrained by a searcher moving along a dynamical constraint. Rather, the search effort was constrained over some target space \mathbb{R}^n and time interval $[0, T]$ by:

$$r(x, t) \geq 0, \forall x \in \mathbb{R}^n, \quad \forall t \in [0, T]$$

and

$$\int_{\mathbb{R}^n} r(x, t) dx = 1, \quad \forall t \in [0, T]$$

For a target trajectory $y(t, \omega) \in \mathbb{R}^n$, the search effort yielded a probability of detection

given by:

$$J = \int_{\Omega} G \left(\int_0^T r(y(t, \omega), t) dt \right) d\omega$$

The second result of note is that of Lukka, [13], who did look at the problem of searchers moving along a dynamical constraint. Lukka developed necessary (but not sufficient) conditions for the problem of minimizing

$$J = \int_{\Omega} G \left(\int_0^T r(x(t), u(t), t, \omega) dt \right) d\omega$$

subject to:

$$\dot{x}(t) = u(t), \quad x(0) = x_0, \quad x(T) = x_T$$

with control constraint $|u(t)| \leq K$ and the additional requirements that $u(t) \in L_{\infty}([0, T])$ and $x(t)$ is absolutely continuous on $[0, T]$. The next section will add to this collection by proving necessary conditions for Problem **P** in the case when the control constraints define an open region for feasible controls.

3.2.1 Necessary Conditions for Open Control Regions

In most of the applications discussed in this thesis, a state $x(t, \omega)$ evolves with the independent variable t and is considered to be conditioned on a stochastic parameter ω . However, because ω has been construed as a constant unknown parameter, as opposed to a stochastic process, an alternative interpretation is available to us: We can equivalently interpret this problem as a problem with multiple independent variables: t , and (the possibly vector-valued) variable ω . Though the interpretation of ω as an

additional independent variable does not always have a physical analogue in recent engineering applications, mathematically it reveals a large relevant area of previous research which can be applied fruitfully to this problem.

Optimal control problems with multiple independent variables have been studied previously under the topic of *distributed control*, in the context of spatial problems with p.d.e. constraints. The problem considered here is neither a subset nor an extension of distributed control. In addition to the difference of p.d.e. constraints, in distributed control problems, the control itself is a function of the additional variables—e.g. in comparable notation a control of the form $u(t, \omega)$. However, there are commonalities, such as the extra layer of integration, and these commonalities make utilizable several of the techniques developed for distributed control.

In the 1960's, necessary conditions and a Pontryagin-like minimum principle were established for the Mayer-Bolza problem with multiple integrals for distributed controls and very general boundary conditions, control constraints, and systems of p.d.e constraints [32], [33], [34]. This section applies several of the techniques found in those proofs to this Problem **P**, yielding necessary conditions for optimality in the case of open control regions. To apply these methods, the following regularity conditions are imposed on the class of problems:

Assumption 1 *Controls $u(t)$ are arbitrary piecewise continuous functions whose values are contained in some set U . Furthermore, the control region U is an open set.¹ States $x(t, \omega)$ take values in some set $X \subset \mathbb{R}$.*

¹Note: Assumption 1 makes these results inapplicable to problems with an active control constraint $g(u) \leq 0$.

Assumption 2 *The functions $\phi(\omega)$, F , r , and f are continuous and twice continuously differentiable with respect to their arguments over their respective domains of Ω , X , $X \times U \times \Omega$, $X \times \Omega$.*

We now define a dual problem to Problem **P**, referred to as Problem **P^λ**, as follows:

Problem P^λ: For feasible solution (x, u) to Problem **P**, find $\lambda : [0, T] \times \Omega \rightarrow \mathbb{R}^{N_x}$ that satisfies the following conditions:

$$\dot{\lambda}(t, \omega) = -\frac{\partial H(x, \lambda, u, t, \omega)}{\partial x} \quad (3.4)$$

$$\lambda(T, \omega) = \frac{\partial F(x(T, \omega), \omega)}{\partial x} \Big|_{\Omega} \quad (3.5)$$

where $H : X \times X \times U \times [0, T] \times \Omega \rightarrow \mathbb{R}$ is defined as:

$$H(x, \lambda, u, t, \omega) = \lambda f(x(t, \omega), u(t), \omega) + r(x(t, \omega), u(t), t, \omega) \quad (3.6)$$

Note, Equations 3.4 and 3.5 create a linear system with a unique solution ([16]).

Furthermore, the continuous differentiability of f and r with respect to x guarantees the continuity of $\dot{\lambda}$ and thus the continuous differentiability of λ . As a final note, the partial derivatives $\frac{\partial H}{\partial x}$ and $\frac{\partial F}{\partial x}$ are taken to return row vectors, of dimension $1 \times N_x$. Thus the dimensions of λ are a transposition of the dimensions of x and the dynamics function f which are taken to be $N_x \times 1$.

Theorem 1 *Let $u^*(t)$ be an optimal control to Problem **P** under Assumptions 1 and 2, let $x^*(t, \omega)$ be the corresponding optimal trajectory and let $\lambda^*(t, \omega)$ be the solution to*

Problem \mathbf{P}^λ for (x^*, u^*) . Then the following holds for all $t \in [0, T]$:

$$\left(\frac{\partial}{\partial u} \int_{\Omega} H(x^*, \lambda^*, u, t, \omega) d\omega \right) \Big|_{u=u^*(t)} = 0 \quad (3.7)$$

Proof:

Inserting $H(x, \lambda, u, t, \omega)$ into $J[x, u]$ we find:

$$J[x, u] = \int_{\Omega} \left\{ F(x(T, \omega), \omega) + \int_0^T [H(x, \lambda, u, t, \omega) - \lambda(t, \omega) \dot{x}] dt \right\} d\omega.$$

We now add and subtract the quantity $\dot{\lambda}x$ to arrive at the following:

$$\begin{aligned} J[x, u] &= \int_{\Omega} F(x(T, \omega), \omega) d\omega - \int_{\Omega} \left\{ \int_0^T [\dot{\lambda}x + \lambda \dot{x}] dt \right\} d\omega \\ &\quad + \int_{\Omega} \left\{ \int_0^T [H(x, \lambda, u, t, \omega) + \dot{\lambda}x] dt \right\} d\omega \end{aligned}$$

which can be split in to two pieces:

$$J_1[x, u] = \int_{\Omega} F(x(T, \omega), \omega) d\omega - \int_{\Omega} \left\{ \int_0^T [\dot{\lambda}x + \lambda \dot{x}] dt \right\} d\omega \quad (3.8)$$

$$J_2[x, u] = \int_{\Omega} \left\{ \int_0^T [H(x, \lambda, u, t, \omega) + \dot{\lambda}x] dt \right\} d\omega. \quad (3.9)$$

We will proceed by examining the first variation of $J[x, u]$ with respect to these two pieces. We take the first variation in the sense used in the calculus of variations, and described for example in [35] and [36]. A necessary condition of extremal trajectories of Problem \mathbf{P} is that the first variation of the functional along these extremals is zero.

The approach of [33], which we follow here, is to demonstrate that the stationarity of this first variation enforces Theorem 1.

Let $u^*(t)$ be an extremal control of the functional $J[x, u]$. One can define a single parameter family of functions as:

$$u(t, \epsilon) = u^*(t) + \epsilon h(t)$$

for a continuously differentiable function h . Let h be further constrained to be zero-valued at $t = 0$ and $t = T$. With no other restrictions on h , the openness of U guarantees the existence of a sufficiently small ϵ such $u(t, \epsilon)$ is in the set of admissible controls. For sufficiently small ϵ , this creates a unique solution $x(t, \omega, \epsilon)$ to the system

$$\dot{x}(t, \omega, \epsilon) = f(x(t, \omega, \epsilon), u(t, \epsilon), \omega), \quad x(0, \omega, \epsilon) = x_0(\omega),$$

which is continuously differentiable with respect to ϵ . (The derivative of $u(t, \epsilon)$ with respect to ϵ is $h(t)$, and thus $u(t, \epsilon)$ is also continuously differentiable with respect to ϵ .) The functional $J[u, x]$ for the controls and states $u(t, \epsilon)$ and $x(t, \omega, \epsilon)$ is also a function of ϵ and the first variation δJ is defined as the differential²:

$$\delta J[x, u] = \left. \frac{d}{d\epsilon} J[x(t, \omega, \epsilon), u^*(t) + \epsilon h(t)] \right|_{\epsilon=0} \cdot \epsilon \quad (3.10)$$

²For a more elaborate treatment of taking the first variation of a functional with dynamical constraints, see [36], Chapter 3.

As this is a simple derivative, we have that:

$$\begin{aligned}
\delta J[x, u] &= \frac{d}{d\epsilon} (J_1[x(t, \omega, \epsilon), u^*(t) + \epsilon h(t)] + J_2[x(t, \omega, \epsilon), u^*(t) + \epsilon h(t)]) \Big|_{\epsilon=0} \cdot \epsilon \\
&= \frac{d}{d\epsilon} J_1[x(t, \omega, \epsilon), u^*(t) + \epsilon h(t)] \Big|_{\epsilon=0} \cdot \epsilon + \frac{d}{d\epsilon} J_2[x(t, \omega, \epsilon), u^*(t) + \epsilon h(t)] \Big|_{\epsilon=0} \cdot \epsilon \\
&\quad \delta J_1[x, u] + \delta J_2[x, u].
\end{aligned}$$

We can thus examine the first variation of J in terms of the variations of J_1 and J_2 .

We first observe that at $\epsilon = 0$, we have $u(t, 0) = u^*(t)$ and thus $x(t, \omega, 0) = x^*(t, \omega)$, the extremal state for $J[x, u]$ corresponding to $u^*(t)$. Let $\lambda^*(t, \omega)$ be the costate defined by Equation 3.4 with respect to the values x^* , u^* . The continuous differentiability of all relevant functions with respect to all relevant arguments allows us to move the differentiation inside the integration, yielding the following:

$$\begin{aligned}
\delta J_2[x, u] &= \int_{\Omega} \int_0^T \epsilon \left\{ \frac{d}{d\epsilon} H(x(t, \omega, \epsilon), \lambda^*(t, \omega), u^*(t) + \epsilon h(t), t, \omega) \right. \\
&\quad \left. + \dot{\lambda}^*(t, \omega) x(t, \omega, \epsilon) \right\} \Big|_{\epsilon=0} dt d\omega
\end{aligned}$$

We define

$$\delta x = \epsilon \cdot \frac{\partial x}{\partial \epsilon} \Big|_{\epsilon=0}, \quad \delta u = \epsilon \cdot \frac{\partial u}{\partial \epsilon} \Big|_{\epsilon=0} = \epsilon h(t)$$

and find that

$$\delta J_2[x, u] = \int_{\Omega} \int_0^T \left\{ \delta x \left(\frac{\partial H(x^*(t, \omega), \lambda^*(t, \omega), u, t, \omega)}{\partial x} \Big|_{u=u^*} + \dot{\lambda}^*(t, \omega) \right) \right.$$

$$+\delta u \left. \frac{\partial H(x^*(t, \omega), \lambda^*(t, \omega), u, t, \omega)}{\partial u} \right|_{u=u^*} \Big\} dt d\omega$$

We invoke equation 3.4 and arrive at:

$$\delta J_2[x, u] = \epsilon \int_{\Omega} \int_0^T \left. \frac{\partial H(x^*(t, \omega), \lambda^*(t, \omega), u, t, \omega)}{\partial u} \right|_{u=u^*} h(t) dt d\omega. \quad (3.11)$$

We next deal with the variation of J_1 around u^* and x^* . By observing that

$$\int_0^T (\dot{\lambda}x + \lambda \dot{x}) dt = \int_0^T \frac{\partial}{\partial t}(\lambda x) dt = \lambda(T, \omega)x(T, \omega) - \lambda(0, \omega)x(0, \omega)$$

due to integration by parts, we can rewrite J_1 as

$$J_1[x, u] = \int_{\Omega} \{F(x(T, \omega), \omega) - (\lambda(T, \omega)x(T, \omega) - \lambda(0, \omega)x(0, \omega))\} d\omega$$

The first variation is

$$\delta J_1[x^*, u^*] = \epsilon \int_{\Omega} \left. \frac{d}{d\epsilon} \{F(x(T, \omega, \epsilon), \omega) - (\lambda^*(T, \omega)x(T, \omega, \epsilon) - \lambda^*(0, \omega)x(0, \omega, \epsilon))\} \right|_{\epsilon=0} d\omega$$

We define

$$\begin{aligned} \delta x(T, \omega, \epsilon) &= \epsilon \cdot \left. \frac{\partial x(T, \omega, \epsilon)}{\partial \epsilon} \right|_{\epsilon=0}, \\ \delta x(0, \omega, \epsilon) &= \epsilon \cdot \left. \frac{\partial x(0, \omega, \epsilon)}{\partial \epsilon} \right|_{\epsilon=0} \end{aligned}$$

Due to our initial conditions, $x(0, \omega, \epsilon)$ was defined as a constant with respect to ϵ , so

$\delta x(0, \omega, \epsilon) = 0$. The variation is thus

$$\delta J_1[x^*, u^*] = \int_{\Omega} \delta x(T, \omega, \epsilon) \left\{ \frac{\partial F(x^*(T, \omega), \omega)}{\partial x} - \lambda^*(T, \omega) \right\} d\omega \quad (3.12)$$

We now invoke equation 3.5. This sets the variation δJ_1 equal to zero. Under these conditions then $\delta J = \delta J_2$ and so δJ_2 is zero-valued over extremal trajectories. Removing the non-zero quantity epsilon from equation 3.11, we have

$$\int_{\Omega} \int_0^T \frac{\partial H(x^*(t, \omega), \lambda^*(t, \omega), u, t, \omega)}{\partial u} h(t) dt d\omega \Big|_{u=u^*} = 0.$$

Ideally, one would like to apply the Fundamental Lemma of the Calculus of Variations to this quantity to yield the stationarity of a Hamiltonian with respect to u . However, the restriction of the controls to be functions of t but not ω , which constrained our choice of h , limits us in this step. To apply the lemma, h must be fixed along the boundary of integration—in this case, h is only fixed along the boundary of $[0, T]$. We address this with the following maneuver, allowed to us by our continuity assumptions and the dependence of h and u solely on t :

$$\begin{aligned} \int_{\Omega} \int_0^T \frac{\partial H(x^*(t, \omega), \lambda^*(t, \omega), u, t, \omega)}{\partial u} h(t) dt d\omega \Big|_{u=u^*} &= \\ \int_0^T h(t) \left(\frac{\partial}{\partial u} \int_{\Omega} H(x^*(t, \omega), \lambda^*(t, \omega), u, t, \omega) d\omega \right) dt \Big|_{u=u^*} &= 0 \end{aligned}$$

The fundamental lemma of the calculus of variations now applies yielding

$$\left(\frac{\partial}{\partial u} \int_{\Omega} H(x^*, \lambda^*, u, t, \omega) d\omega \right) \Big|_{u=u^*(t)} = 0$$

which proves the theorem. The new quantity

$$\mathbf{H}(x^*, \lambda^*, u, t) = \int_{\Omega} H(x^*, \lambda^*, u, t, \omega) d\omega \quad (3.13)$$

provides a stationary Hamiltonian for Problem **P**.

The result of Theorem 1 also facilitates a second result, which is useful for the numerical verification of optimal answers.

Theorem 2 *Let $u^*(t)$ be an optimal control to Problem **P** under Assumptions 1 and 2, let $x^*(t, \omega)$ be the corresponding optimal trajectory and let $\lambda^*(t, \omega)$ be the solution to Problem \mathbf{P}^λ for (x^*, u^*) . Let $\mathbf{H}(x^*, \lambda^*, u, t)$ be the Hamiltonian defined by Equation 3.13. If the component functions, f and r , of Problem **P** are time-invariant, then:*

$$\frac{d\mathbf{H}(x^*, \lambda^*, u^*, t)}{dt} = 0$$

Proof:

The total derivative of \mathbf{H} with respect to time is given by:

$$\frac{d\mathbf{H}(x^*, \lambda^*, u^*, t)}{dt} = \frac{d}{dt} \int_{\Omega} H(x^*, \lambda^*, u^*, t, \omega) d\omega$$

The continuous differentiability of all relevant arguments allows for the time derivative to be moved inside the integral, yielding:

$$\frac{d\mathbf{H}(x^*, \lambda^*, u^*, t)}{dt} = \int_{\Omega} \frac{dH(x^*, \lambda^*, u^*, t, \omega)}{dt} d\omega$$

Since the component functions of $H(x^*, \lambda^*, u^*, t, \omega)$ do not depend explicitly on time, the total derivative of $H(x^*, \lambda^*, u^*, t, \omega)$ with respect to time is given by:

$$\begin{aligned} \frac{dH(x^*, \lambda^*, u^*, t, \omega)}{dt} &= \frac{\partial H(x^*, \lambda^*, u^*, t, \omega)}{\partial x} \dot{x}^*(t, \omega) + \\ &\quad \frac{\partial H(x^*, \lambda^*, u^*, t, \omega)}{\partial \lambda} \dot{\lambda}^{*T}(t, \omega) + \frac{\partial H(x^*, \lambda^*, u^*, t, \omega)}{\partial u} \dot{u}^*(t) \end{aligned}$$

By Equation 3.4,

$$\frac{\partial H(x^*, \lambda^*, u^*, t, \omega)}{\partial x} = -\lambda^*(t, \omega)$$

and by Equation 3.6,

$$\frac{\partial H(x^*, \lambda^*, u^*, t, \omega)}{\partial \lambda} = f^T(x(t, \omega), u(t, \omega)) = \dot{x}^{*T}(t, \omega).$$

Thus:

$$\begin{aligned} &\frac{\partial H(x^*, \lambda^*, u^*, t, \omega)}{\partial x} \dot{x}^*(t, \omega) + \frac{\partial H(x^*, \lambda^*, u^*, t, \omega)}{\partial \lambda} \dot{\lambda}^{*T}(t, \omega) \\ &= -\lambda^*(t, \omega) \dot{x}^*(t, \omega) + \dot{x}^{*T}(t, \omega) \dot{\lambda}^{*T}(t, \omega) = 0. \end{aligned}$$

This leaves the total derivative of \mathbf{H} with the values:

$$\frac{d\mathbf{H}(x^*, \lambda^*, u^*, t)}{dt} = \int_{\Omega} \frac{dH(x^*, \lambda^*, u^*, t, \omega)}{dt} d\omega = \int_{\Omega} \frac{\partial H(x^*, \lambda^*, u^*, t, \omega)}{\partial u} \dot{u}^*(t) d\omega$$

Since $\dot{u}^*(t)$ does not depend on ω , this becomes:

$$\frac{d\mathbf{H}(x^*, \lambda^*, u^*, t)}{dt} = \dot{u}^*(t) \int_{\Omega} \frac{\partial H(x^*, \lambda^*, u^*, t, \omega)}{\partial u} d\omega$$

which by Theorem 1 is zero-valued. Thus $\frac{d\mathbf{H}(x^*, \lambda^*, u^*, t)}{dt} = 0$.

3.2.2 Example Analytic Solution Via Necessary Conditions

As is the case with many optimality conditions, the condition derived above is realistically only a tool for providing analytic solutions to problems in limited, simple cases. Its larger use is as a computational verification tool, as will be demonstrated in the examples in the next chapter. However, this section will provide an example of a solvable system, to demonstrate that, though the derived Hamiltonian differs in form from the Hamiltonian found in standard control problems (incorporating, as does Problem **P**, an extra layer of integration), the necessary conditions it provides are specific enough to yield solutions.

To demonstrate this, we examine the following linear quadratic problem: Given a continuous probability density function $\phi : \Omega \rightarrow \mathbb{R}$, determine the control $u : [0, T] \rightarrow \mathbb{R}$ that minimizes the expectation

$$J = \int_{\Omega} \int_0^T (x^2(t, \omega) + u^2(t)) dt \phi(\omega) d\omega \quad (3.14)$$

subject to the dynamics

$$\dot{x}(t, \omega) = u(t) + \omega, \quad x(0, \omega) = x_0$$

This problem could be interpreted as the minimization of a robot's distance from a target translated to the origin, with a penalty function $u^2(t)$ to keep the control within reasonable bounds and with an unknown additive noise influencing the robot's motion. Define the function $H(x, \lambda, u, t, \omega)$ by

$$H(x, \lambda, u, t, \omega) = \lambda(u + \omega) + (x^2 + u^2)\phi(\omega) \quad (3.15)$$

and let the costate variable $\lambda(t, \omega)$ be defined, as per Equation 3.4, as the solution to:

$$\dot{\lambda}(t, \omega) = -\frac{\partial H(x, \lambda, u, t, \omega)}{\partial x} = -2x\phi(\omega) \quad (3.16)$$

$$\lambda(T, \omega) = 0 \quad (3.17)$$

From Theorem 1 we have

$$\left(\frac{\partial}{\partial u} \int_{\Omega} H(x^*, \lambda^*, u, t, \omega) d\omega \right) \Big|_{u=u^*(t)} = 0 \quad (3.18)$$

The integral inside this expression becomes:

$$\int_{\Omega} H(x^*, \lambda^*, u, t, \omega) d\omega = \int_{\Omega} (\lambda^*(t, \omega)(u(t) + \omega) + [x^*(t, \omega)]^2 + [u^*(t, \omega)]^2) \phi(\omega) d\omega$$

$$\begin{aligned}
&= \int_{\Omega} \lambda^*(t, \omega) u^*(t) \phi(\omega) d\omega + \int_{\Omega} \lambda^*(t, \omega) \omega \phi(\omega) d\omega \\
&\quad + \int_{\Omega} [x^*(t, \omega)]^2 \phi(\omega) d\omega + \int_{\Omega} [u^*(t)]^2 \phi(\omega) d\omega \\
&= [u^*(t)]^2 + u^*(t) \int_{\Omega} \lambda^*(t, \omega) \phi(\omega) d\omega + \int_{\Omega} \lambda^*(t, \omega) \omega \phi(\omega) d\omega + \int_{\Omega} [x^*(t, \omega)]^2 \phi(\omega) d\omega.
\end{aligned}$$

Note that the last two terms do not depend on u . Thus:

$$\begin{aligned}
\frac{\partial}{\partial u} \int_{\Omega} H(x^*, \lambda^*, u, t, \omega) d\omega &= \frac{\partial}{\partial u} \left([u^*(t)]^2 + u^*(t) \int_{\Omega} \lambda^*(t, \omega) \phi(\omega) d\omega \right) \\
&= 2u + \int_{\Omega} \lambda(t, \omega) \phi(\omega) d\omega = 0
\end{aligned}$$

From this we find:

$$u = -\frac{1}{2} \int_{\Omega} \lambda(t, \omega) \phi(\omega) d\omega$$

which, since $\dot{x} = u + \omega$, can be rewritten as:

$$\dot{x} - \omega = -\frac{1}{2} \int_{\Omega} \lambda(t, \omega) \phi(\omega) d\omega.$$

This implies that:

$$\dot{x} = -\frac{1}{2} \int_{\Omega} \lambda(t, \omega) \phi(\omega) d\omega + \omega. \quad (3.19)$$

This equation along with Equation 3.16 creates a system of ODE's. To solve this system, we will utilize the following auxiliary variables:

$$\tilde{\lambda}(t) = \int_{\Omega} \lambda(t, \omega) \phi(\omega) d\omega$$

$$\tilde{x}(t) = \int_{\Omega} x(t, \omega) \phi(\omega) d\omega.$$

From these, we have:

$$\begin{aligned} \dot{\tilde{x}} &= \int_{\Omega} \dot{x}(t, \omega) \phi(\omega) d\omega = \int_{\Omega} \left[-\frac{1}{2} \int_{\Omega} \lambda(t, \omega) \phi(\omega) d\omega + \omega \right] \phi(\omega) d\omega \\ &= -\frac{1}{2} \int_{\Omega} \lambda(t, \omega) \phi(\omega) d\omega + E[\omega] = -\frac{1}{2} \tilde{\lambda} + E[\omega] \end{aligned}$$

where the redundant integral has integrated to one after the first layer of integrations yields quantities which are no longer dependent on ω . Furthermore:

$$\dot{\tilde{\lambda}} = \int_{\Omega} \dot{\lambda}(t, \omega) \phi(\omega) d\omega = -2 \int_{\Omega} x(t, \omega) \phi(\omega) d\omega = -2\tilde{x}.$$

Thus in terms of the auxiliary variables we have the linear system:

$$\begin{bmatrix} \dot{\tilde{x}}(t) \\ \dot{\tilde{\lambda}}(t) \end{bmatrix} = \begin{bmatrix} 0 & -\frac{1}{2} \\ -2 & 0 \end{bmatrix} \begin{bmatrix} \tilde{x}(t) \\ \tilde{\lambda} \end{bmatrix} + \begin{bmatrix} E[\omega] \\ 0 \end{bmatrix}, \quad \begin{bmatrix} \dot{\tilde{x}}(0) \\ \dot{\tilde{\lambda}}(T) \end{bmatrix} = \begin{bmatrix} x_0 \\ 0 \end{bmatrix}. \quad (3.20)$$

Upon solving this we find that:

$$\tilde{\lambda}(t) = -\frac{2e^{-t}(e - e^t - e^{2+t} + e^{1+2t})}{1 + e^2} E[\omega]$$

and utilizing that fact that $u = -\frac{1}{2}\tilde{\lambda}(t)$, we find the optimal control is:

$$u^*(t) = \frac{e^{-t}(e - e^t - e^{2+t} + e^{1+2t})}{1 + e^2} E[\omega].$$

Chapter 4

Computational Approach

Recently, there has been much progress in developing numerical methods for tackling optimization problems with parameter uncertainties. For instance, for nonlinear finite-dimensional optimization problems, 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 [37]. In 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, [38], [39]. More general nonlinear control problems were approached in [40, 14], where a consistent numerical method was provided for a class of problems with time-invariant parameter uncertainty effecting only the cost function; and in [29], where a multi-dimensional pseudospectral collocation scheme for discretizing both time and parameter space was developed for problems of the form of Problem **P**.

Of these approaches, only [29] and [40, 14] address the convergence properties of the proposed methods. This chapter will contribute to that list by presenting and

algorithm along with convergence results, developed in collaboration with Chris Phelps. These results extend the methods of [14] to a larger class of problems. The problems tackled by the algorithm described in this chapter are comparable in scope to those tackled in [29]. However, a key distinction is that, in this approach, the approximation of parameter space has been entirely separated from the approximation of the time domain. This separation enables a much wider variety of numerical methods than just pseudospectral collocation to be applied as means of approximation, in both the parameter domains and time domain.

4.1 Numerical Algorithm

The computational algorithm used to generate numerical solutions in this thesis is accomplished through two stages of approximation:

1. Approximate Problem \mathbf{P} through discretizing parameter space. This approximation is carried out using M parameter ‘nodes’, in a manner described below in Section 4.1.1. This yields an approximate problem, Problem $\mathbf{P}^{\mathbf{M}}$, which is a standard optimal control problem in the Mayer-Bolza form.
2. Approximate Problem $\mathbf{P}^{\mathbf{M}}$ through discretizing the time domain. This approximation is carried out using N time ‘nodes’, in a manner described below in Section 4.1.3. This yields another approximate problem, Problem $\mathbf{P}^{\mathbf{MN}}$. This final problem is a finite-dimensional nonlinear programming problem, which is solved using the SQP algorithm of the commercial solver SNOPT [41].

4.1.1 Approximating Parameter Space

We first introduce an approximation of Problem \mathbf{P} , referred to as Problem \mathbf{P}^M . Problem \mathbf{P}^M is created by approximating the parameter space, Ω , with a numerical integration scheme which satisfies the form and requirements of Assumption 3 below.

Assumption 3 *For each $M \in \mathbb{N}$, there is a set of nodes $\{\omega_i^M\}_{i=1}^M \subset \Omega$ and an associated set of weights $\{\alpha_i^M\}_{i=1}^M \subset \mathbb{R}$, such that for any continuous function $h : \Omega \rightarrow \mathbb{R}$,*

$$\int_{\Theta} h(\omega) d\omega = \lim_{M \rightarrow \infty} \sum_{i=1}^M h(\omega_i^M) \alpha_i^M.$$

Remark 1 *Note that if $h_M : \Theta \rightarrow \mathbb{R}$ is continuous for all $M \in \mathbb{N}$ and $\{h_M\}$ converges uniformly to h , then*

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

This property is used later in the proof of Theorem 3.

Many methods, for instance Gaussian quadrature and composite-Simpson, satisfy the above assumption. This numerical integration scheme is defined in terms of a finite set of M nodes $\{\omega_i^M\}_{i=1}^M$ and an associated set of M weights $\{\alpha_i^M\}_{i=1}^M \subset \mathbb{R}$. Throughout what follows, M is used to denote the number of nodes used in this approximation of parameter space.

After establishing the nodes of our chosen integration scheme, the state vector $x : [0, T] \times \Omega \mapsto \mathbb{R}^{n_x}$ is approximated by a set of state variables $[\bar{x}_1^M, \dots, \bar{x}_M^M]$, where

\bar{x}_i^M is the solution to the dynamical system:

$$\begin{cases} \dot{\bar{x}}_i^M(t) = f(\bar{x}_i^M(t), u(t), \omega_i^M) \\ \bar{x}_i^M(0) = x_0(\omega_i^M), \end{cases} \quad i = 1, \dots, M. \quad (4.1)$$

The notation $\bar{X}_M = [\bar{x}_1^M, \dots, \bar{x}_M^M]$ is used to denote the states of the approximate problem where the dependence on the parameter ω has been discretized.

The numerical integration scheme for parameter space creates an approximate objective functional for each $M \in \mathbb{N}$, defined by:

$$\bar{J}^M[\bar{X}_M, u] = \sum_{i=1}^M \left[F(\bar{x}_i^M(T), \omega_i^M) + \int_0^T r(\bar{x}_i^M(t), u(t), t, \omega_i^M) dt \right] \alpha_i^M. \quad (4.2)$$

This discretization of parameter space thus leads to the following approximate problem:

Problem \mathbf{P}^M : Determine the function pair (\bar{X}_M, u) that minimizes the cost functional (4.2) subject to the dynamics (4.1) and the control constraint of Problem **P**.

4.1.2 Convergence of Primal Variables

In this section the convergence properties of Problem \mathbf{P}^M as the number of nodes M tends to infinity is addressed. The property which will be particularly addressed is that of ‘consistency’—the property that if optimal solutions to Problem \mathbf{P}^M converge as the number of nodes $M \rightarrow \infty$, they converge to feasible, optimal solutions of Problem **P**. Before addressing this property, it is necessary to impose the following regularity assumptions on Problem **P**:

Assumption 4 *Feasible controls u are those which satisfy all problem constraints and*

are in $L_\infty([0, T]; \mathbb{R}^{n_u})$. The function $g : \mathbb{R}^{n_u} \mapsto \mathbb{R}^{n_g}$ is continuous and the set $U = \{\nu \in \mathbb{R}^{n_u} | g(\nu) \leq 0\}$ is compact.

Assumption 5 *There exists a compact set $X \subset \mathbb{R}^{n_x}$ such that for each feasible control $u \in U$ and $\omega \in \Omega, t \in [0, T]$, $x(t, \omega) \in X$ where $x(t, \omega) = x_0(\omega) + \int_0^t f(x(s, \omega), u(s), \omega) ds$ for all $t \in [0, T]$.*

Assumption 6 *The functions f and r are C^1 . The set Ω is compact and $x_0 : \Omega \mapsto \mathbb{R}^{n_x}$ is continuous. Moreover, for the compact sets X and U defined in Assumptions 4-5 and for each $t \in [0, T]$, $\omega \in \Omega$, the Jacobians $r_x(\cdot, \cdot, t, \theta)$ and $f_x(\cdot, \cdot, \theta)$ are Lipschitz on the set $X \times U$, and the corresponding Lipschitz constants L_r and L_f are uniformly bounded in θ and t . The function $F(\cdot, \theta)$ is C^1 on X for all $\omega \in \Omega$; in addition, F and F_x are continuous with respect to ω .*

The concept of ‘convergent’ solutions of Problem \mathbf{P}^M that the concept of consistency will be applied to is that of ‘uniform accumulation points.’ This concept is defined as follows:

Definition 1 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 1 applies to limits of subsequences and not just the limits of the sequence in entirety. The convergence assumptions placed on Problem \mathbf{P}^M will be of

this form, so it is useful to introduce a notation that enables the convenient mention of subsequence limits. 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 $\{x_i | i \in V\}$ has a limit point x , we will refer to this with the notation:

$$\lim_{n \in V} x_n = x$$

This is in contrast to the usual limit notation, $\lim_{n \rightarrow \infty} x_n$, which means that the sequence in entirety converges.

Because of the process used in the approximation, the state space for Problem \mathbf{P}^M is of a different dimension than that of Problem \mathbf{P} . States of Problem \mathbf{P}^M are functions of time, i.e. the functions $\bar{X}^M(t)$, whereas states of Problem \mathbf{P} are functions of time and ω . This aspect of the problem differs from the related work of [40, 14], in which the state space was unaffected by the discretization of parameter space. Although the states spaces of the two problem differ, however, their control space remains the same. The following theorem utilizes this latter fact to demonstrate that the accumulation points of optimal *controls* for Problem \mathbf{P}^M yield optimal *controls* for Problem \mathbf{P} . The theorem relies on the following two lemmas:

Lemma 1 *Feasible controls for Problem \mathbf{P}^M are also feasible controls for Problem B , and the set of feasible controls is closed in the topology of pointwise convergence.*

Proof:

As problems \mathbf{P} and \mathbf{P}^M share the same control constraint, the only concern in using a feasible control u_M from Problem \mathbf{P}^M would be that there exists a value of ω , not

equal to a discretized value ω_i^M , such that the solution $\dot{x} = f(x, u_M, \omega)$ does not exist for some t in $[0, T]$. However, since \mathbf{P} and \mathbf{P}^M share the same control, u_M is in the set U , and thus this possibility is prevented by Assumption 5. Furthermore, since U is compact, the set of feasible controls is trivially closed in the topology of pointwise convergence.

Lemma 2 *Let $\{u_M\}$ be a sequence of feasible controls for Problem \mathbf{P}^M with an accumulation point u^∞ for the infinite set $V \subset \mathbb{N}$. Let $x^\infty(t, \omega)$ be the solution to the dynamical system:*

$$\dot{x}^\infty(t, \omega) = f(x^\infty(t, \omega), u^\infty(t), \omega), \quad x^\infty(0, \omega) = x_0(\omega)$$

and let $\{x_M(t, \omega)\}$ be the sequence of solutions to the dynamical systems:

$$\dot{x}_M(t, \omega) = f(x_M(t, \omega), u_M(t), \omega), \quad x_M(0, \omega) = x_0(\omega), \quad M \in V$$

The pairs $(x^\infty(t, \omega), u^\infty(t))$ and $\{(x_M(t, \omega), u_M(t))\}$ are all feasible solutions to Problem \mathbf{P} . Furthermore, the sequence $\{x_M(t, \omega)\}$ converges pointwise to $x^\infty(t, \omega)$ and this convergence is uniform in ω .

Proof:

By Lemma 2, both u^∞ and $\{u_M\}$ are feasible controls for Problem **P**, which guarantees the feasibility of $x^\infty(t, \omega)$ and $\{x_M(t, \omega)\}$. From their definitions, we have:

$$\|x_M(t, \omega) - x^\infty(t, \omega)\| \leq \int_0^t \|f(x_M(\tau, \omega), u_M(\tau), \omega) - f(x^\infty(\tau, \omega), u(\tau), \omega)\| dt$$

Since f is C^1 on a compact domain, the Lipschitz condition applies, yielding:

$$\|x_M(t, \omega) - x^\infty(t, \omega)\| \leq \int_0^t L \left(\|x_M(\tau, \omega) - x^\infty(\tau, \omega)\| + \|u_M(\tau) - u(\tau)\| \right) dt$$

Since u_M and u^∞ are in the compact set U , they are bounded. Thus the Dominated Convergence Theorem applies and we have:

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

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

$$\int_0^t \|u_M(\tau) - u(\tau)\| dt < \delta_u.$$

This provides us with the inequality

$$\|x_M(t, \omega) - x(t, \omega)\| \leq LT\delta_u + L \int_0^t \|x_M(\tau, \omega) - x(\tau, \omega)\| dt.$$

By Gronwall's Inequality,

$$\|x_M(t, \omega) - x(t, \omega)\| \leq LT\delta_u e^{LT}.$$

Since for each value of t and ω , this quantity can be made arbitrarily small, $\{x_M(t, \omega)\}$ converges pointwise to $x^\infty(t, \omega)$. Furthermore, since δ_u , though it may depend on t , does not depend on the value of ω , this convergence is uniform in ω .

Theorem 3 *Let $\{u_M^*\}_{M \in V}$ be a sequence of optimal controls for Problem \mathbf{P}^M with an accumulation point u^∞ for the infinite set $V \subset \mathbb{N}$. Then u^∞ is an optimal control for Problem \mathbf{P} .*

Proof:

Let $\{u_M^*\}_{M \in V}$ be a set of optimal controls for the Problem \mathbf{P}^M such that $\lim_{M \in V} \{u_M^*\} = u^\infty$. Let $x^\infty(t, \omega)$ be the solution to the dynamical system:

$$\dot{x}^\infty(t, \omega) = f(x^\infty(t, \omega), u^\infty(t), \omega), \quad x^\infty(0, \omega) = x_0(\omega)$$

and let $\{x_M(t, \omega)\}$ be the sequence of solutions to the dynamical systems:

$$\dot{x}_M(t, \omega) = f(x_M(t, \omega), u_M^*(t), \omega), \quad x_M(0, \omega) = x_0(\omega), \quad M \in V$$

Notice that this latter state is *not* a feasible state for Problem \mathbf{P}^M ; it is the feasible state for Problem \mathbf{P} generated by the control $u_M^*(t)$. However, when $\omega = \omega_i^M$, then $x_M(t, \omega_i^M) = \bar{x}_i^{*M}(t)$, where \bar{x}_i^{*M} is the optimal state for Problem \mathbf{P}^M generated by the

optimal control u_M^* . Thus:

$$\begin{aligned}\bar{J}^M[\bar{X}_M^*, u_M^*] &= \sum_{i=1}^M \left[F(\bar{x}_i^{*M}(T), \omega_i^M) + \int_0^T r(\bar{x}_i^{*M}(t), u_M^*(t), t, \omega_i^M) dt \right] \alpha_i^M \\ &= \sum_{i=1}^M \left[F(x_M(T, \omega_i^M)) + \int_0^T r(x_M(t, \omega_i^M), u_M^*(t), t, \omega_i^M) dt \right] \alpha_i^M\end{aligned}$$

The limit of the difference in values between $\bar{J}^M[\bar{X}_M^*, u_M^*]$ and $J[x^\infty, u^\infty]$ is thus:

$$\begin{aligned}\lim_{M \in V} \left\| \bar{J}^M[\bar{X}_M^*, u_M^*] - J[x^\infty, u^\infty] \right\| &= \\ \lim_{M \in V} \left\| \sum_{i=1}^M \left[F(x_M(T, \omega_i^M)) + \int_0^T r(x_M(t, \omega_i^M), u_M^*(t), t, \omega_i^M) dt \right] \alpha_i^M - \right. \\ &\quad \left. \int_{\Omega} \left[F(x^\infty(T, \omega)) + \int_0^T r(x^\infty(t, \omega), u^\infty(t), t, \omega) dt \right] d\omega \right\| \\ &\leq \left\| \lim_{M \in V} \sum_{i=1}^M F(x_M(T, \omega_i^M)) \alpha_i^M - \int_{\Omega} F(x^\infty(T, \omega)) d\omega \right\| \\ &\quad + \left\| \lim_{M \in V} \sum_{i=1}^M \int_0^T r(x_M(t, \omega_i^M), u_M^*(t), t, \omega_i^M) dt - \int_{\Omega} \int_0^T r(x^\infty(t, \omega), u^\infty(t), t, \omega) dt d\omega \right\|\end{aligned}$$

On the other hand, we can examine the quantity:

$$\begin{aligned}\lim_{M \in V} \left\| \int_0^T r(x_M(t, \omega), u_M^*(t), \omega) - \int_0^T r(x^\infty(t, \omega), u^\infty(t), \omega) \right\| \\ \leq \lim_{M \in V} \int_0^T \|r(x_M(t, \omega), u_M^*(t), \omega) - r(x^\infty(t, \omega), u^\infty(t), \omega)\| dt\end{aligned}$$

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

$$\begin{aligned} & \int_0^T \|r(x_M(t, \omega), u_M^*(t), \omega) - r(x^\infty(t, \omega), u^\infty(t), \omega)\| dt \\ & \leq \int_0^T L(\|x_M(t, \omega) - x^\infty(t)\| + \|u_M^*(t) - u^\infty(t)\|) dt \end{aligned}$$

The results of Lemma 2 and the compactness of X and U enable us to apply the Dominated Convergence Theorem. Thus:

$$\lim_{M \in V} \int_0^T L(\|x_M(t, \omega) - x^\infty(t, \omega)\| + \|u_M^*(t) - u^\infty(t)\|) dt = 0$$

and this convergence must be uniform in ω due to the uniform convergence of $x_M(t, \omega)$.

We can therefore conclude that:

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

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

$$\left\| \lim_{M \in V} \sum_{i=1}^M \int_0^T r(x_M(t, \omega_i^M), u_M^*(t), \omega_i^M) dt - \int_\Omega \int_0^T r(x^\infty(t, \omega), u^\infty(t), \omega) dt d\omega \right\| = 0$$

Similar arguments show that:

$$\left\| \lim_{M \in V} \sum_{i=1}^M F(x_M(T, \omega_i^M)) \alpha_i^M - \int_\Omega F(x^\infty(T, \omega)) d\omega \right\| = 0$$

Thus we find that:

$$\lim_{M \in V} \bar{J}^M[\bar{X}_M^*, u_M^*] = J[(x^\infty, u^\infty)]$$

Following arguments identical to those in [40], we find that the convergence of this limit shows u^∞ to be an optimal control for Problem **P**.

4.1.3 Approximating the Time Domain

The next step after creating the standard control problem, Problem **P^M**, through discretization of parameter space, is to solve the resulting problem. To implement problem **P^M**, this thesis utilizes the method of ‘direct pseudospectral collocation.’ As described in [42] and [43], in direct collocation, both the state and control are approximated over a discretized time grid $\pi^N = [t_0, \dots, t_N]$ as:

$$x(t) \approx x^N(t) = \sum_{k=0}^N \bar{x}^{Nk} \phi_k(t), \quad u(t) \approx u^N(t) = \sum_{k=0}^N \bar{u}^{Nk} \psi_k(t)$$

where $\{\phi_k\}_{k=0}^N, \{\psi_k\}_{k=0}^N$ are sets of interpolating functions. In local collocation methods, the degree of the interpolating functions $\{\phi_k\}_{k=0}^N$ is constant with respect to N , whereas in global collocation methods, like the pseudospectral method, the degree increases with N . The choice of interpolating functions also creates a differentiation scheme such that if $\bar{x}^N = [\bar{x}^{N0}, \dots, \bar{x}^{NN}]$, differentiation can be approximated as the matrix multiplication $D^N \bar{x}^N$, where D^N is determined by the values of $\{\dot{\phi}_k\}$. This in coordination with a quadrature scheme with weights $\{b_k^N\}_{k=0}^N$ for the grid π^N creates the following problem:

Problem P^{MN}: Determine the decision variables $\bar{x}_i^N = [\bar{x}_i^{N0}, \dots, \bar{x}_i^{NN}]$, $i = 1, \dots, M$

and $\bar{u}^N = [\bar{u}^{N0}, \dots, \bar{u}^{NN}]$ that minimize the sum:

$$J^{MN} = \sum_{i=1}^M a_i^M \left[F(\bar{x}_i^{NN}, \omega_i^M) + \sum_{k=0}^N b_k^N r(\bar{x}_i^{Nk}, \bar{u}^{Nk}, t_k, \omega_i^M) \right]$$

Subject to

$$D^N \bar{x}_i^N - f(\bar{x}_i^N, \bar{u}^N) = 0, \quad \bar{x}_i^{N0} = x_0(\omega_i^M)$$

$$g(\bar{u}^{Nk}) \leq 0 \text{ for all } k = 0, \dots, N$$

This is a finite-dimensional nonlinear programming problem, which can be solved using the SQP algorithm of the commercial solver SNOPT [41].

4.1.4 Convergence of Dual Problems

As mentioned, Problem \mathbf{P}^M is a standard optimal control, in the Bolza form. It is therefore subject to the necessary conditions for optimality applicable to such problems. For Problem \mathbf{P}^M , the dual variables are proscribed by the following problem:

Problem $\mathbf{P}^{M\lambda}$: For feasible solution (\bar{X}_M, u) to Problem \mathbf{P}^M , find $\bar{\Lambda}(t) = [\bar{\lambda}_1^M(t), \dots, \bar{\lambda}_M^M(t)]$,

$\bar{\lambda}_i^M : [0, T] \rightarrow \mathbb{R}^{N_x}$, that satisfies the following conditions:

$$\dot{\bar{\lambda}}_i^M(t) = - \frac{\partial \bar{H}^M(\bar{x}_i^M, \bar{\lambda}_i^M, u, t)}{\partial x_i^M} \quad (4.3)$$

$$\bar{\lambda}_i^M(T) = \alpha_i^M \frac{\partial F(\bar{x}_i^M, \omega_i^M)}{\partial \bar{x}_i^M} \Big|_{\Omega} \quad (4.4)$$

where \bar{H}^M is defined as:

$$\bar{H}^M(\bar{X}_M, \bar{\Lambda}_M, u, t) = \sum_{i=1}^M [\bar{\lambda}_i^M f(\bar{x}_i^M(t), u(t), \omega_i^M) + \alpha_i^M r(\bar{x}_i^M(t), u(t), t, \omega_i^M)] \quad (4.5)$$

As with the necessary conditions of Section 3.2.1, the partial derivatives $\frac{\partial \bar{H}^M}{\partial x}$ and $\frac{\partial F}{\partial x}$ are taken to return row vectors, of dimension $1 \times N_x$. Thus the dimensions of $\bar{\lambda}_i^M$ are a transposition of the dimensions of \bar{x}_i^M which are taken to be $N_x \times 1$.

An alternate approach to solving Problem **P** overall is to approximate the necessary conditions of Section 3.2.1, i.e. Problem **P** ^{λ} , directly rather than to approximate Problem **P**. Discretizing Problem **P** ^{λ} amounts to enforcing the values of the dynamics in Equations 3.4 and 3.5, as well as the primary state dynamics, at the parameter collocation nodes $\{\omega_i^M\}$ (just as the state dynamics of Problem **P** are collocated in Equation 4.1). This creates the system of equations:

$$\begin{cases} \dot{\lambda}(t, \omega_i^M) = -\frac{\partial H(x, u, t, \omega_i^M)}{\partial x} \\ \lambda(T, \omega_i^M) = \frac{\partial F(x(T, \omega_i^M), \omega_i^M)}{\partial x} \end{cases} \quad i = 1, \dots, M. \quad (4.6)$$

where H is defined as per Equation 3.6 as:

$$H(x, \lambda, u, t, \omega) = \lambda f(x(t, \omega), u(t), \omega) + r(x(t, \omega), u(t), t, \omega)$$

This system of equations can actually be written compactly using the quadrature ap-

proximation of the stationary Hamiltonian defined in Equation 3.13:

$$\begin{aligned}\mathbf{H}(x, \lambda, u, t) &= \int_{\Omega} H(x(t, \omega), \lambda(t, \omega), u(t), t, \omega) d\omega \\ &\approx \sum_{i=1}^M \alpha_i^M H(x(t, \omega_i^M), \lambda(t, \omega_i^M), u(t), t, \omega_i^M) = \tilde{H}^M(x, \lambda, u, t).\end{aligned}$$

Letting $\tilde{\Lambda}(t) = [\tilde{\lambda}_1^M(t), \dots, \tilde{\lambda}_M^M(t)] = [\lambda(t, \omega_1^M), \dots, \lambda(t, \omega_M^M)]$, and letting $\tilde{X}_M = [\tilde{x}_1^M(t), \dots, \tilde{x}_M^M(t)]$

denote the discretized states from Equation 4.1, Equation 4.6 can be written as.

$$\begin{cases} \dot{\tilde{\lambda}}_i^M(t) = -\frac{1}{\alpha_i^M} \cdot \frac{\partial \tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}, u, t)}{\partial \tilde{x}_i^M} \\ \tilde{\lambda}_i^M(T) = \frac{\partial F(\tilde{x}_i^M(T), \omega_i^M)}{\partial \tilde{x}_i^M} \end{cases}, \quad i = 1, \dots, M. \quad (4.7)$$

Thus we reach the following discretized dual problem:

Problem $\mathbf{P}^{\lambda M}$: For feasible controls u and solutions \tilde{X}_M to Equation 4.1, find $\tilde{\Lambda}(t) =$

$[\tilde{\lambda}_1^M(t), \dots, \tilde{\lambda}_M^M(t)]$, $\tilde{\lambda}_i^M : [0, T] \rightarrow \mathbb{R}^{N_x}$, that satisfies the following conditions:

$$\dot{\tilde{\lambda}}_i^M(t) = -\frac{1}{\alpha_i^M} \cdot \frac{\partial \tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}, u, t)}{\partial \tilde{x}_i^M} \quad (4.8)$$

$$\tilde{\lambda}_i^M(T) = \frac{\partial F(\tilde{x}_i^M(T), \omega_i^M)}{\partial \tilde{x}_i^M} \quad (4.9)$$

where \tilde{H}^M is defined as:

$$\tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}_M, u, t) = \sum_{i=1}^M \left[\alpha_i^M \tilde{\lambda}_i^M f(\tilde{x}_i^M(t), u(t), \omega_i^M) + \alpha_i^M r(\tilde{x}_i^M(t), u(t), t, \omega_i^M) \right] \quad (4.10)$$

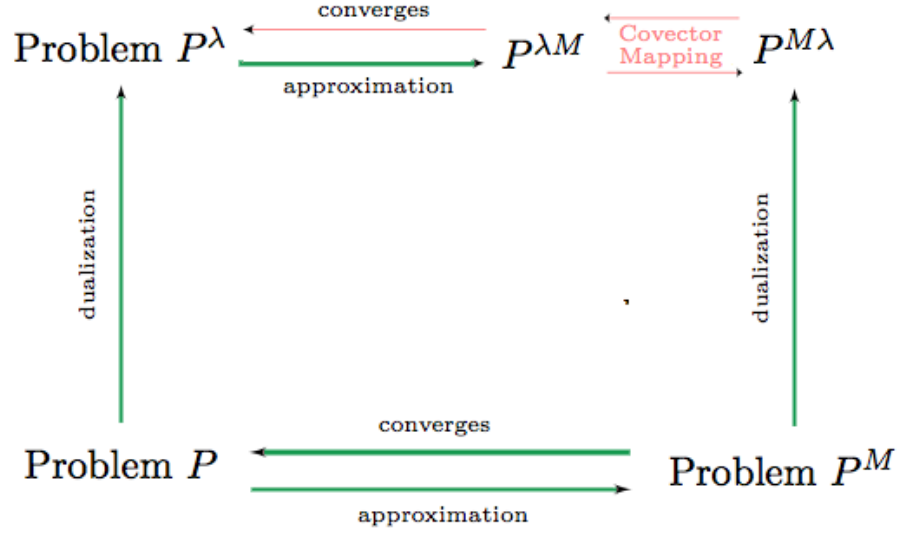


Figure 4.1: Diagram of Goals for Primal and Dual Relations for Parameter Uncertainty Control. Red lines designate the remaining needed results.

In the spirit of the Covector Mapping Theorem for direct methods applied to standard control problems, discussed for instance in references [20], [42], and [44], several results are desirable for the relationship between these dual problems. The desired results are diagrammed in Figure 4.1. In the case of this particular problem, unlike standard control, the collocation of the relevant dynamics involves no approximation of differentiation (since the discretization is in the parameter domain rather than the time domain) and thus the mapping of covectors between Problem $\mathbf{P}^{M\lambda}$ and Problem $\mathbf{P}^{\lambda M}$ is straightforward.

Lemma 3 *The mapping:*

$$(\bar{x}_i^M, \bar{u}) \mapsto (\tilde{x}_i^M, \tilde{u}), \quad \frac{\bar{\lambda}_i^M}{\alpha_i^M} \mapsto \tilde{\lambda}_i^M,$$

for $i = 1, \dots, M$ is a bijective mapping from solutions of Problem $\mathbf{P}^{M\lambda}$ to Problem

$\mathbf{P}^{\lambda \mathbf{M}}$.

Lemma 3 follows immediately through substitution. To pave the way for convergence, we also have the following:

Lemma 4 *Let $\{u_M\}$ be a sequence of optimal controls for Problem $\mathbf{P}^{\mathbf{M}}$ with an accumulation point u^∞ for the infinite set $V \subset \mathbb{N}$. Let $(x^\infty(t, \omega), \lambda^\infty(t, \omega))$ be the solution to the dynamical system:*

$$\begin{cases} \dot{x}^\infty(t, \omega) = f(x^\infty(t, \omega), u^\infty(t, \omega), \omega), & x^\infty(0, \omega) = x_0(\omega) \\ \dot{\lambda}^\infty(t, \omega) = -\frac{\partial H(x^\infty(t, \omega), \lambda^\infty(t, \omega), u^\infty(t, \omega), \omega)}{\partial x}, & \lambda^\infty(T, \omega) = \frac{\partial F(x^\infty(T, \omega), \omega)}{\partial x} \end{cases} \quad (4.11)$$

where H is defined as per Equation 3.6, and let $\{(x_M(t, \omega), \lambda_M(t, \omega))\}$ for $M \in V$ be the sequence of solutions to the dynamical systems:

$$\begin{cases} \dot{x}_M(t, \omega) = f(x_M(t, \omega), u_M(t, \omega), \omega), & x_M(0, \omega) = x_0(\omega) \\ \dot{\lambda}_M(t, \omega) = -\frac{\partial H(x_M(t, \omega), \lambda_M(t, \omega), u_M(t, \omega), \omega)}{\partial x}, & \lambda_M(T, \omega) = \frac{\partial F(x_M(T, \omega), \omega)}{\partial x} \end{cases} \quad (4.12)$$

Then, the sequence $\{(x_M(t, \omega), \lambda_M(t, \omega))\}$ converges pointwise to $(x^\infty(t, \omega), \lambda^\infty(t, \omega))$ and this convergence is uniform in ω .

The convergence of $\{x_M(t, \omega)\}$ is given by Lemma 2. The convergence of the sequence of solutions $\{\lambda_M(t, \omega)\}$ is guaranteed by the optimality of $\{u_M\}$. The convergence of $\{\lambda_M(t, \omega)\}$ then follows the same arguments given the convergence of $\{x_M(t, \omega)\}$, utilizing the regularity assumptions placed on the derivatives of F , r , and f with respect to x to enable the use of Lipschitz conditions on the costate dynamics and transversality

conditions.

Remark 2 Note that $\lambda_M(t, \omega)$ is not a costate of Problem $\mathbf{P}^{\lambda\mathbf{M}}$, since it is a function of ω . However, when $\omega = \omega_i^M$, then $\lambda_M(t, \omega_i^M) = \tilde{\lambda}_i^M(t)$, where $\tilde{\lambda}_i^M$ is the costate of Problem $\mathbf{P}^{\lambda\mathbf{M}}$ generated by the pair of solutions to Problem $\mathbf{P}^{\mathbf{M}}$, (\tilde{x}_i^M, u_M^*) . In other words, the function $\lambda_M(t, \omega)$ matches the costate values at all collocation nodes. Since these values satisfy the dynamics equations of Problem $\mathbf{P}^{\lambda\mathbf{M}}$, a further implication of this is that the values of $\lambda_M(t, \omega_i^M)$ produce feasible solutions to Problem $\mathbf{P}^{\lambda\mathbf{M}}$.

Remark 3 Since the functions $\{(x_M(t, \omega), \lambda_M(t, \omega))\}$ obey the respective identities $x_M(t, \omega_i^M) = \tilde{x}_i^M(t)$ and $\lambda_M(t, \omega_i^M) = \tilde{\lambda}_i^M(t)$, their convergence to $(x^\infty(t, \omega), \lambda^\infty(t, \omega))$ also implies the convergence of the sequence of discretized primals and duals, $\{\tilde{X}_M\}$ and $\{\tilde{\Lambda}_M\}$, to accumulation points given by the relations

$$\lim_{M \in V} \tilde{x}_i^M(t) = x^\infty(t, \omega_i^M), \quad \lim_{M \in V} \tilde{\lambda}_i^M(t) = \lambda^\infty(t, \omega_i^M)$$

Building on this lemma, we have the following theorem:

Theorem 4 Let $\{\tilde{X}_M, \tilde{\Lambda}_M, u_M\}$ be a sequence of solutions for Problem $\mathbf{P}^{\lambda\mathbf{M}}$ with an accumulation point $\{\tilde{X}^\infty, \tilde{\Lambda}^\infty, u^\infty\}$ for the infinite set $V \subset \mathbb{N}$. Let $(x^\infty, \lambda^\infty, u^\infty)$ be the solutions to Problem \mathbf{P}^λ for the control u^∞ . Then

$$\lim_{M \in V} \tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}_M, u_M, t) = \mathbf{H}(x^\infty, \lambda^\infty, u^\infty, t)$$

where \tilde{H}^M is the Hamiltonian of Problem $\mathbf{P}^{\lambda\mathbf{M}}$ as defined by Equation 4.10 and \mathbf{H} is the Hamiltonian of Problem \mathbf{P} as defined by Equation 3.13.

Proof:

Let $\{(x_M(t, \omega), \lambda_M(t, \omega))\}$ for $M \in V$ be the sequence of solutions defined by Equation 4.12 and let $(x^\infty(t, \omega), \lambda^\infty(t, \omega))$ be the accumulation functions defined by Equation 4.11. Incorporating Remarks 2 and 3, we have:

$$\begin{aligned} \lim_{M \in V} \tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}_M, u_M, t) &= \\ \lim_{M \in V} \sum_{i=1}^M \alpha_i^M \left[\tilde{\lambda}_i^M(t) f(\tilde{x}_i^M(t), u(t), \omega_i^M) + r(\tilde{x}_i^M(t), u(t), t, \omega_i^M) \right] &= \\ \lim_{M \in V} \sum_{i=1}^M \alpha_i^M \left[\lambda_M(t, \omega_i^M) f(x_M(t, \omega_i^M), u(t), \omega_i^M) + r(x_M(t, \omega_i^M), u(t), t, \omega_i^M) \right] \end{aligned}$$

Due to the results of Lemma 4, and applying Remark 1 on the convergence of the quadrature scheme for uniformly convergent sequences of continuous functions, we find that:

$$\begin{aligned} \lim_{M \in V} \tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}_M, u_M, t) &= \\ \int_{\Omega} [\lambda^\infty(t, \omega) f(x^\infty(t, \omega), u^\infty(t), \omega) + r(x^\infty(t, \omega), u^\infty(t), t, \omega)] d\omega &= \\ = \mathbf{H}(x^\infty, \lambda^\infty, u^\infty, t) \end{aligned}$$

Thus proving the theorem.

The relationship between these results and those derived for standard control is diagrammed in Figure 4.2. The convergence of the dual problems after covector mapping provides with the opportunity to extend the some of the features of the dual

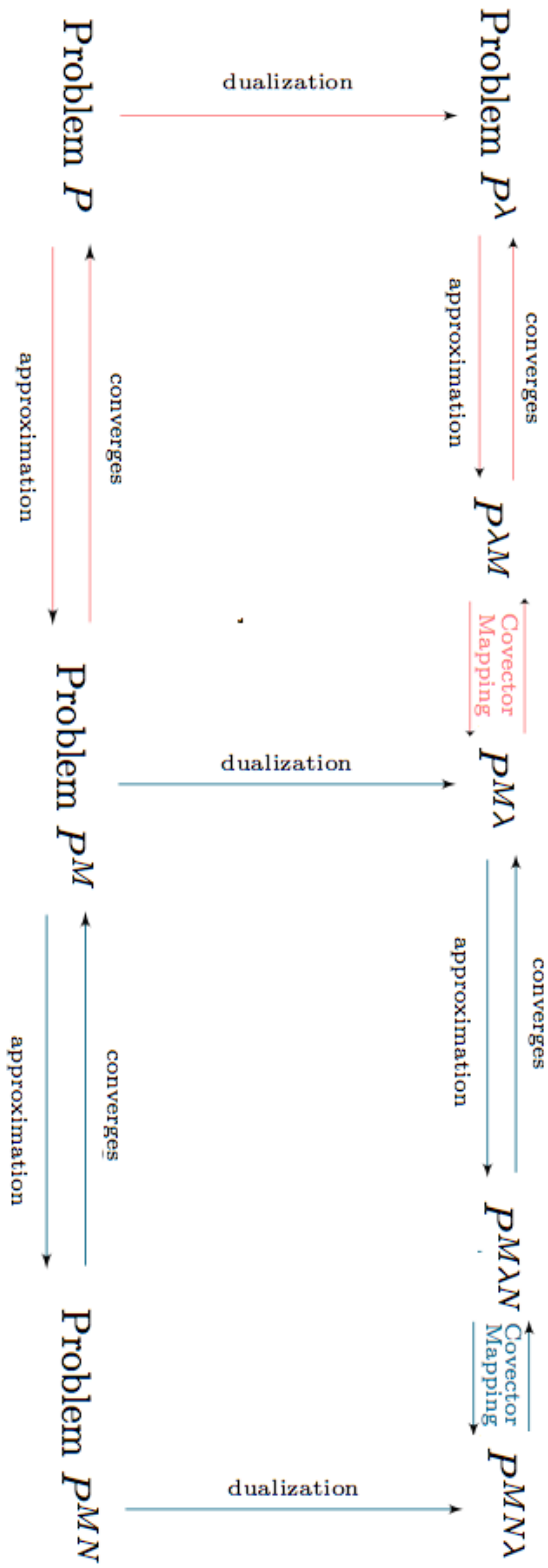


Figure 4.2: Diagram of Primal and Dual Relations for Standard Control Combined with Parameter Uncertainty Control

problem for standard control to the dual problem for Problem **P**.

Corollary 1 Hamiltonian Minimization Principle

Let $\{u_M^\}$ be a sequence of optimal controls for Problem $\mathbf{P}^{\mathbf{M}}$ with an accumulation point u^∞ for the infinite set $V \subset \mathbb{N}$. Let $(x^\infty, \lambda^\infty)$ be the primal and dual variables for Problem **P** created by the control u^∞ . Then*

$$\mathbf{H}(x^\infty, \lambda^\infty, u^\infty, t) \leq \mathbf{H}(x^\infty, \lambda^\infty, u, t)$$

for all feasible u .

Proof:

Let $\{(\tilde{X}_{*M}, \tilde{\Lambda}_{*M})\}$ be the sequence of solutions to Problem $\mathbf{P}^{\lambda\mathbf{M}}$ created by $\{u_M^*\}$. The minimization principle holds for the Hamiltonian of the standard control problem and thus

$$\tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}_M, u_M, t) \leq \tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}_M, u, t)$$

for all feasible u . By Theorem 4,

$$\mathbf{H}(x^\infty, \lambda^\infty, u^\infty, t) = \lim_{M \in V} \tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}_M, u_M, t)$$

The arguments of Theorem 4 also supply that

$$\lim_{M \in V} \tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}_M, u, t) = \mathbf{H}(x^\infty, \lambda^\infty, u, t).$$

Thus:

$$\begin{aligned}\mathbf{H}(x^\infty, \lambda^\infty, u^\infty, t) &= \lim_{M \in V} \tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}_M, u_M, t) \\ &\leq \lim_{M \in V} \tilde{H}^M(\tilde{X}_M, \tilde{\Lambda}_M, u, t) = \mathbf{H}(x^\infty, \lambda^\infty, u, t)\end{aligned}$$

which supplies the desired relation.

4.1.5 Linear Quadratic System

The following provides an opportunity to gain insight into the behavior of this algorithm, by supplying a problem in which both Problem **P** and Problem **P^M** are analytically solvable. Consider the following system:

$$\text{Problem P: } \left\{ \begin{array}{ll} \text{Minimize} & J = \int_{\Omega} \left(\int_0^1 \sum_{k=1}^K [(x_k - \omega_k)^2 + u_k^2] dt \right) p(\omega) d\omega \\ \text{Subject to} & \dot{x}(t) = u \\ & x(0) = 0 \end{array} \right. \quad (4.13)$$

where $x = (x_1, x_2, \dots, x_K) \in \mathbb{R}^K$, $u = (u_1, u_2, \dots, u_K) \in \mathbb{R}^K$. The parameters $\omega = (\omega_1, \omega_2, \dots, \omega_K) \in \Omega$ are independent random variables with joint probability density function $p(\omega)$. This simple objective function can represent the K-dimensional distance from a stationary target at position $(\omega_1, \omega_2, \dots, \omega_K)$ with a penalty function $\sum_{k=1}^K u_k^2$ meant to keep each control u_k within a reasonable range. Interchanging the order of

integration allows this to be formulated as a standard control problem:

$$\left\{ \begin{array}{l} \text{Minimize} \quad J = \int_0^1 \left(\int_{\Omega} \sum_{k=1}^K [(x_k - \omega_k)^2 + u_k^2] p(\omega) d\omega \right) dt \\ \text{Subject to} \quad \dot{x}(t) = u \\ \\ x(0) = 0 \end{array} \right.$$

and independence of the random variables allows for explicit integration over Ω , reducing the objective function to

$$J = \int_0^1 \sum_{k=1}^K (x_k^2 + u_k^2 - 2x_k E[\omega_k] + E[\omega_k^2]) dt.$$

The Hamiltonian of this system is then

$$\mathbf{H}(x, u, \lambda) = \lambda^T u + \sum_{k=1}^K (x_k^2 + u_k^2 - 2x_k E[\omega_k] + E[\omega_k^2])$$

with $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_K)$ and extremal solutions (x^*, u^*) satisfying the adjoint equations and Transversality conditions

$$\begin{aligned} \frac{\partial \mathbf{H}}{\partial u} \Big|_{u=u^*} &= 0 \\ \frac{\partial \mathbf{H}}{\partial x} \Big|_{x=x^*} &= -\dot{\lambda}^* \quad \lambda^*(1) = 0 \end{aligned}$$

This creates a linear system which can be solved analytically, yielding expressions for the extremal trajectories, adjoints, and optimal objective value:

$$x_k^* = E[\omega_k] \left(1 - \frac{e^t + e^{2-t}}{1 + e^2} \right), \quad u_k^* = -E[\omega_k] \left(\frac{e^t - e^{2-t}}{1 + e^2} \right), \quad \lambda_k^* = 2E[\omega_k] \left(\frac{e^t - e^{2-t}}{1 + e^2} \right),$$

$$J(x^*, u^*) = \sum_{k=1}^K \left(E[\omega_k]^2 \left(\frac{e^2 - 1}{e^2 + 1} - 1 \right) + E[(\omega_k)^2] \right).$$

The same procedure can be used to provide an analytical solution to the approximated problem, Problem P^M , created by discretizing the parameter space. Let $\{\omega_{k,i}^M\}_{i=1}^M$ be a set of nodes for Ω_k and let $\{a_{k,i}^M\}_{i=1}^M$ be a set of weights for a convergent quadrature scheme as described in Section 4.1 (Note that this is a slight modification of previous notation, as here M is the number of nodes used per dimension k , rather than the whole space Ω .)

Because ω_k are independent random variables, we can separate $p(\omega)$ into component probability distributions $p_k(\omega_k)$. Introducing the following useful constants

$$c_k^M = \sum_{i=1}^M p_k(\omega_{k,i}^M) a_{k,i}^M, \quad c_{-k}^M = \prod_{j \neq k} c_j^M, \quad c^M = \prod_j c_j^M$$

$$\tilde{E}[\omega_k] = \sum_{i=1}^M \omega_{k,i}^M p_k(\omega_{k,i}^M) a_{k,i}^M, \quad \tilde{E}[(\omega_k)^2] = \sum_{i=1}^M (\omega_{k,i}^M)^2 p_k(\omega_{k,i}^M) a_{k,i}^M$$

we can write Problem P^M as

$$\left\{ \begin{array}{l} \text{Minimize} \quad J^M = \sum_{k=1}^K c_{-k}^M \sum_{i=1}^M \left[\int_0^1 ((x_k^2 - \omega_{k,i}^M)^2 + u_k^2) dt \right] p_k(\omega_{k,i}^M) a_{k,i}^M \\ \\ \text{Subject to} \quad \dot{x}(t) = u \\ \\ \quad \quad \quad x(0) = 0. \end{array} \right.$$

Notice that c_k^M is the quadrature approximation of $\int_{\omega_k} p_k(\omega_k) d\omega_k$, $\tilde{E}[\omega_k]$ is the approx-

imation of $E[\omega_k]$, and $\tilde{E}[(\omega_k)^2]$ is the approximation of $E[(\omega_k)^2]$. Thus

$$\lim_{M \rightarrow \infty} c_k^M = 1 \quad \lim_{M \rightarrow \infty} \tilde{E}[\omega_k] = E[\omega_k] \quad \lim_{M \rightarrow \infty} \tilde{E}[(\omega_k)^2] = E[(\omega_k)^2].$$

The Hamiltonian for P^M is

$$\begin{aligned} \mathbf{H}(x, u, \lambda) &= \lambda^T u + \sum_{k=1}^K c_{-k}^M \sum_{i=1}^M \left[(x_k^2 - \omega_{k,i}^M)^2 + u_k^2 \right] p_k(\omega_{k,i}^M) a_{k,i}^M \\ &= \lambda^T u + \sum_{k=1}^K c_{-k}^M \left(c_k^M (x_k^2 + u_k^2) - 2x_k \tilde{E}[\omega_k] + \tilde{E}[(\omega_k)^2] \right) \end{aligned}$$

and extremal solutions (x_M^*, u_M^*) satisfy the adjoint equations

$$\begin{aligned} \left. \frac{\partial \mathbf{H}}{\partial u_k} \right|_{u_k = u_{k,M}^*} &= \lambda_{k,M}^* + 2c^M u_{k,M}^* = 0 \\ \left. \frac{\partial \mathbf{H}}{\partial x_k} \right|_{x_k = x_{k,M}^*} &= 2c^M u_{k,M}^* - 2\tilde{E}[\omega_k] = -\dot{\lambda}_{k,M}^*. \end{aligned}$$

Solving the system yields the following expressions for the extremal trajectories, adjoints, and optimal objective value:

$$\begin{aligned} x_{k,M}^* &= \frac{\tilde{E}[\omega_k]}{c^M} \left(1 - \frac{e^t + e^{2-t}}{1 + e^2} \right), \quad u_{k,M}^* = -\tilde{E}[\omega_k] \left(\frac{e^t - e^{2-t}}{1 + e^2} \right), \\ \lambda_{k,M}^* &= 2\tilde{E}[\omega_k] \left(\frac{e^t - e^{2-t}}{1 + e^2} \right), \\ J^M(x_M^*, u_M^*) &= \sum_{k=1}^K \left(\tilde{E}[\omega_k]^2 \left(\frac{e^2 - 1}{e^2 + 1} - 1 \right) + \tilde{E}[(\omega_k)^2] \right) \end{aligned}$$

Comparison of the solutions to Problem **P** and Problem **P^M** reveals that in this case the convergence of states, controls, costates, and cost are all determined by the convergence




Figure 4.3: Convergence of analytical solutions for $J(x^*, u^*)$ and $J^M(x_M^*, u_M^*)$ using Euler's method vs Legendre pseudospectral. The dimension of the problem has been set at $K = 2$, and a Beta(10,10) distribution over the domain $[0,1]$ has been given to each parameter.

of the quadrature schemes when applied to the expectation and variance of the uncertain parameter. Figure 4.3 shows the convergence rates of the cost function using Euler's method versus Legendre pseudospectral to discretize the parameter domain.

4.2 Implementations

To implement the following problems, the method of approximating parameter space described in Section 4.1.1 is used, as well as direct collocation for the time domain, as described in Section 4.1.3. This method creates a large finite-dimensional nonlinear programming problem (NLP), which can be solved using a variety of available software packages. For the numerical solutions in this paper we use the commercial

solver SNOPT, which runs on the SQP algorithm detailed in [41].

4.2.1 Ensemble Control of Nonholonomic Unicycles

Another studied ensemble system is the no-slip nonholonomic unicycle. This system models the movement of a single wheel across a two-dimensional surface, using the states $x = [x_1, x_2, x_3]^T$, where $[x_1, x_2]$ is the two-dimensional location of a wheel and x_3 is the heading angle. When no error is present this system has the following equations of motion:

$$\begin{cases} \dot{x}_1(t) = u_1(t) \cos x_3(t) \\ \dot{x}_2(t) = u_1(t) \sin x_3(t) \\ \dot{x}_3(t) = u_2(t) \end{cases} \quad , \quad x(0) = x_0$$

When uncertainty is introduced which scales the turning rate and velocity, this creates the following ensemble system, which has been studied in [28] and [45]:

$$\begin{cases} \dot{x}_1(t, \omega) = \omega \cdot u_1(t) \cos x_3(t, \omega) \\ \dot{x}_2(t, \omega) = \omega \cdot u_1(t) \sin x_3(t, \omega) \\ \dot{x}_3(t, \omega) = \omega \cdot u_2(t) \end{cases} \quad , \quad x(0, \omega) = x_0(\omega)$$

In [28] it was shown that while this overall system is not controllable, the final positions $x_1(T, \omega)$ and $x_2(T, \omega)$ are controllable (again in the sense that in a finite time interval $[0, T]$ the system can be steered to within a ball of radius ϵ around a desired point).

A control was then *analytically* derived to drive a group of robots to the origin, given various starting positions and a spectrum of values for ω . Here, *numerical* solutions to this problem are obtained by applying direct methods to the general parameter uncertainty problem created by minimizing the average final time distance, over all parameter values, from a desired point (α_1, α_2) .

Unicycle Ensemble Problem: Determine the control $u : [0, T] \rightarrow U \in \mathbb{R}^{n_u}$ that minimizes the objective

$$J = \frac{1}{2} \int_{\Omega} [(x_1(T, \omega) - \alpha_1)^2 + (x_2(T, \omega) - \alpha_2)^2] d\omega \quad (4.14)$$

subject to the dynamics

$$\begin{cases} \dot{x}_1(t, \omega) = \omega \cdot u_1(t) \cos x_3(t, \omega) \\ \dot{x}_2(t, \omega) = \omega \cdot u_1(t) \sin x_3(t, \omega) \\ \dot{x}_3(t, \omega) = \omega \cdot u_2(t) \end{cases} \quad , \quad x(0, \omega) = x_0(\omega)$$

and control constraint $|u_i(t)| \leq K_i$ for all $t \in [0, T]$, $i = 1, 2, 3..$

Figures 4.4 and 4.5 show an implementation of this problem for the values $\Omega = [1.0, 1.5]$, $x_0(\omega) = 0$, $K_i = 10$, and $T = 10$. In the time domain, 200 time nodes have been utilized, chosen using Euler's method, and 10 parameter nodes have been utilized, chosen using Legendre pseudospectral. Figure 4.4 shows the control solutions for this implementation and Figure 4.5 shows the trajectories. The final value of the objective,




Figure 4.4: Optimal Controls for Unicycle Ensemble Problem

i.e. the average end time distance from the goal point, is $J = 3.950704157621467 \cdot 10^{-05}$.

As a numerical verification method, we can examine these results in terms of the dual variables returned by the final NLP problem. Using the mappings provided by Covector Mapping Theorem for direct methods applied to standard control problems, [20], and the additional mapping provided by Lemma 3, then by Lemma 4 the discrete duals returned by the NLP solver should approximate the costates provided by Equations 3.4, 3.5, and 3.6 at the nodes of collocation. Equations 3.4, 3.5, and 3.6 provide the conditions:

$$\begin{aligned}
 H(x, \lambda, u, t, \omega) &= \lambda f(x(t, \omega), u(t), \omega) + r(x(t, \omega), u(t), t, \omega) \\
 &= \lambda_1 \omega u_1 \cos x_3 + \lambda_2 \omega u_1 \sin x_3 + \lambda_3 \omega u_2
 \end{aligned}$$

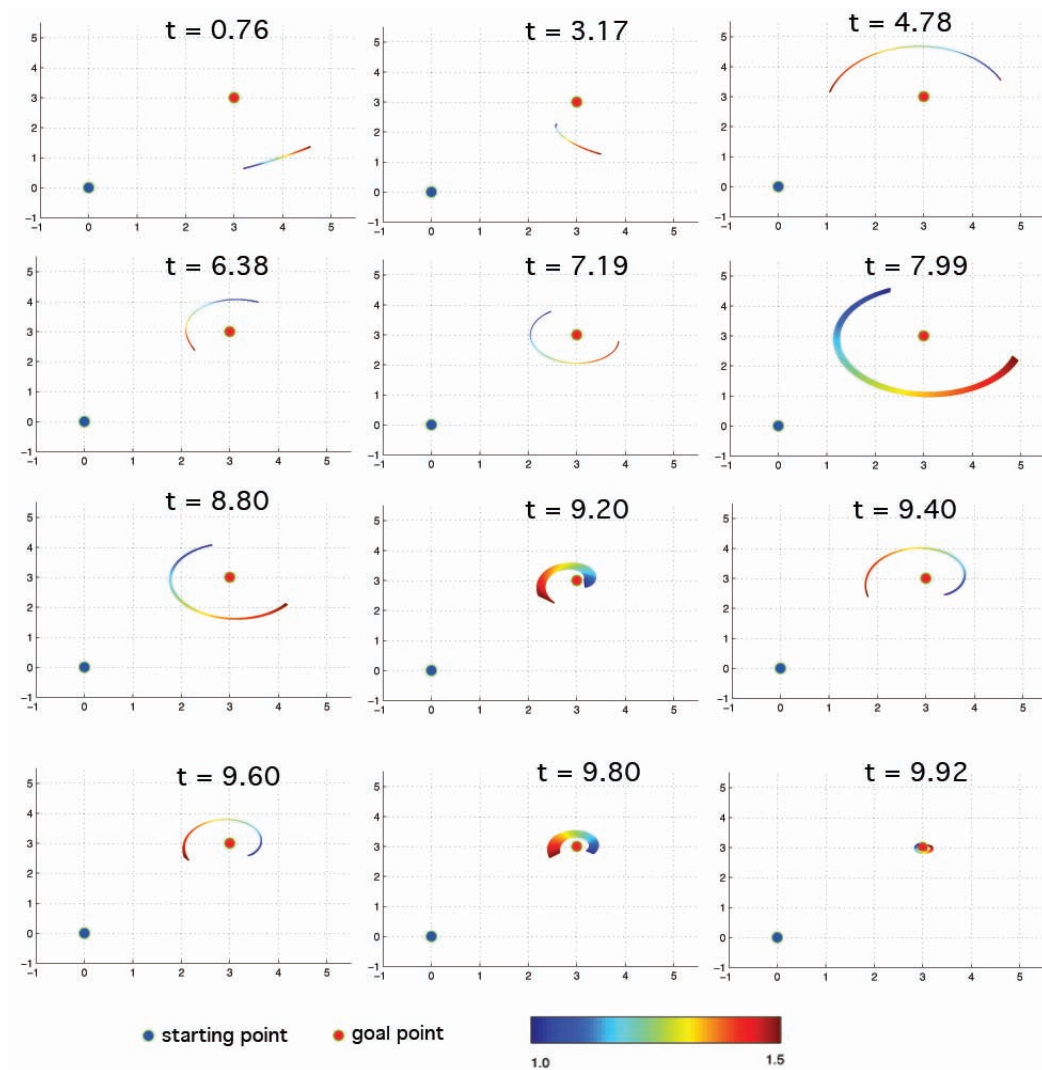


Figure 4.5: Optimal Trajectories for Unicycle Ensemble Problem

and

$$\begin{pmatrix} \dot{\lambda}_1(t, \omega) \\ \dot{\lambda}_2(t, \omega) \\ \dot{\lambda}_3(t, \omega) \end{pmatrix} = \begin{pmatrix} -\frac{\partial H(x, \lambda, u, t, \omega)}{\partial x_1} \\ -\frac{\partial H(x, \lambda, u, t, \omega)}{\partial x_2} \\ -\frac{\partial H(x, \lambda, u, t, \omega)}{\partial x_3} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \omega u_1 (\lambda_1 \sin x_3 - \lambda_2 \cos x_3) \end{pmatrix},$$

$$\begin{pmatrix} \lambda_1(T, \omega) \\ \lambda_2(T, \omega) \\ \lambda_3(T, \omega) \end{pmatrix} = \begin{pmatrix} \frac{\partial F(x(T, \omega), \omega)}{\partial x_1} \\ \frac{\partial F(x(T, \omega), \omega)}{\partial x_2} \\ \frac{\partial F(x(T, \omega), \omega)}{\partial x_3} \end{pmatrix} = \begin{pmatrix} x_1(T, \omega) - \alpha_1 \\ x_2(T, \omega) - \alpha_2 \\ 0 \end{pmatrix}$$

Thus we find that $\lambda_1(t, \omega)$ and $\lambda_2(t, \omega)$ are constant, with values given by the—in this case extremely small—final distances from the goal point. If the NLP solution has yielded a close approximation to the true optimal, the dual variables it returns should reflect this behavior. Figure 4.6 shows the values of the mapped duals corresponding to $\lambda_1(t, \omega)$ and $\lambda_2(t, \omega)$ returned by the NLP solver, plotted over all values of parameter nodes. Indeed, for all parameter node values, these dual variables return values near zero.




Figure 4.6: Values of the mapped duals corresponding to $\lambda_1(t, \omega)$ and $\lambda_2(t, \omega)$ returned by the NLP solver, plotted over all values of parameter nodes

4.2.2 Ensemble Control of Harmonic Oscillators

An area of research relevant to optimal trajectory design with parameter uncertainty is that of ensemble control, studied for instance in [28], [27], and [46]. In ensemble control, the goal is to control, with a single unified control, an ‘ensemble’ of agents which depend continuously on some parameter value. In other words, to steer a dynamical system of the form:

$$\dot{x}(t, \omega) = f(x(t, \omega), u(t), \omega)$$

with a control $u(t)$ which does not change relative to the parameter.

This problem has been studied in the context of applied quantum mechanics

applications, such as medical imaging, where a single electromagnetic pulse is the available input into a system with a spectrum of values (see [29]), and also in the context of robotics, where a minor part uncertainty may provide a range of possible systems to be instantiated by an autonomous robot (see [45]). A studied example application of the problem is a family of two-dimensional harmonic oscillators with states $x = [x_1, x_2]^T$ and state dynamics:

$$\begin{cases} \dot{x}_1(t, \omega) = -\omega x_2(t, \omega) + u_1(t) \\ \dot{x}_2(t, \omega) = \omega x_1(t, \omega) + u_2(t) \end{cases}, \quad x(0, \omega) = x_0(\omega)$$

where the frequency, ω , may lie in some range $[\omega_1, \omega_2]$.

In [27], this system is proved controllable in the sense that in a finite time interval $[0, T]$ the system can be steered to within a ball of radius ϵ around a desired point. Optimal control strategies are then derived for, first, steering to within a ball of radius ϵ around a final state $(x_1(T, \omega), x_2(T, \omega))$ through constraining the final states while minimizing the expended control through the objective:

$$J = \int_0^T u(t)^T u(t) dt$$

and, secondly, for minimizing the average final distance from the origin, with a magnitude constraint on $u(t)^T u(t)$, through the objective:

$$J = \int_{\Omega} x(T, \omega)^T x(T, \omega) d\omega$$

These results were derived analytically, for each of these specific cases. As a contrast, this section applies our numerical algorithm directly, to a quadratic cost function which averages the cumulative distance over time as well as the expended control.

Harmonic Oscillators with Frequency Uncertainty Determine the control $u :$

$[0, T] \rightarrow \mathbb{R}^{n_u}$ that minimizes the cost:

$$J[x, u] = \frac{1}{2} \int_{\Omega} \int_0^T \left(x(t, \omega)^T x(t, \omega) + u(t)^T u(t) \right) dt d\omega \quad (4.15)$$

subject to the dynamics:

$$\begin{cases} \dot{x}_1(t, \omega) = -\omega \cdot x_2(t, \omega) + u_1(t) \\ \dot{x}_2(t, \omega) = \omega \cdot x_1(t, \omega) + u_2(t) \end{cases}, \quad x(0, \omega) = x_0(\omega)$$

and control constraint $|u_i(t)| \leq 1$ for all $t \in [0, T]$, $i = 1, 2$.

Before solving this general system, it is illustrative to examine the behavior of this problem when dealt with in more traditional ways. This provides an opportunity to observe what implications of parameter uncertainty in such a system are. One classic numerical strategy for such a problem would be to solve the problem as a standard control problem, using a nominal value as an estimate for ω . For nominal value $\hat{\omega}$ of the parameter this reduces the problem to the standard control problem:

Standard Harmonic Oscillators Determine the control $u : [0, T] \rightarrow \mathbb{R}^{n_u}$ that mini-

mizes the cost:

$$J[x, u] = \frac{1}{2} \int_0^T \left(x(t)^T x(t) + u(t)^T u(t) \right) dt \quad (4.16)$$

subject to the dynamics:

$$\begin{cases} \dot{x}_1(t) = -\hat{\omega} \cdot x_2(t) + u_1(t) \\ \dot{x}_2(t) = \hat{\omega} \cdot x_1(t) + u_2(t) \end{cases}, \quad x(0) = x_0(\hat{\omega})$$

and control constraint $|u_i(t)| \leq 1$ for all $t \in [0, T]$, $i = 1, 2$.

This second problem can be solved using the necessary conditions provided by Pontryagin's Minimum Principle and has the optimal solution:

$$u^*(t) = \frac{e^{-t} (e^{2t} - e^{2T})}{1 + e^{2T}} \begin{pmatrix} x_{0,1} \cos(\hat{\omega}t) - x_{0,2} \sin(\hat{\omega}t) \\ x_{0,2} \cos(\hat{\omega}t) + x_{0,1} \sin(\hat{\omega}t) \end{pmatrix}$$

and

$$x^*(t) = \frac{e^{-t} (e^{2t} + e^{2T})}{1 + e^{2T}} \begin{pmatrix} x_{0,1} \cos(\hat{\omega}t) - x_{0,2} \sin(\hat{\omega}t) \\ x_{0,2} \cos(\hat{\omega}t) + x_{0,1} \sin(\hat{\omega}t) \end{pmatrix}$$

Figures 4.7 and 4.8 demonstrate the behavior of this solution for values of $T = 100$ and $\hat{\omega} = 2$. Both the trajectories and the control are rapidly driven to zero.

Though the optimal solution to the standard control problem drives the system quickly near the origin for the nominal parameter value, using this control input on systems with moderate differences in parameter values will not replicate that behavior.




Figure 4.7: Optimal Control Values, $u^*(t)$, for Nominal Oscillator System




Figure 4.8: Optimal State Values, $x_1^*(t)$, for Nominal Oscillator System

Figure 4.9 show the trajectories generated by the control input $u^*(t)$ for parameter values ω in the interval $[2, 2.5]$. Color is determined by $|\omega - \hat{\omega}|$.

In contrast, solving the general parameter uncertainty problem on the other hand creates a noticeably more robust solution, which drives the system near the origin for the entire range of parameter values. Figure 4.10 demonstrates the solution of the parameter uncertainty problem for $T = 100$ and $\Omega = [1.5, 2.5]$. This solution was computed using 200 time discretization nodes and 10 nodes for each parameter dimension. The discretization scheme in the time domain is Euler's method and the scheme in the parameter dimensions is Legendre pseudospectral.

Since the control constraints become inactive after the first few time steps, the conditions derived in Section 3 can provide a tool for numerical verification, to assess the viability of these numerical solutions as optimal solutions to the general problem by showing that the control satisfies the necessary conditions for optimality. Equations

3.4, 3.5, and 3.6 provide the conditions:

$$\begin{aligned} H(x, \lambda, u, t, \omega) &= \lambda f(x(t, \omega), u(t), \omega) + r(x(t, \omega), u(t), t, \omega) \\ &= \lambda_1 (-\omega x_2 + u_1) + \lambda_2 (\omega x_1 + u_2) + \frac{1}{2} [x_1^2 + x_2^2 + u_1^2 + u_2^2] \end{aligned}$$

and

$$\begin{aligned} \begin{pmatrix} \dot{\lambda}_1(t, \omega) \\ \dot{\lambda}_2(t, \omega) \end{pmatrix} &= \begin{pmatrix} -\frac{\partial H(x, \lambda, u, t, \omega)}{\partial x_1} \\ -\frac{\partial H(x, \lambda, u, t, \omega)}{\partial x_2} \end{pmatrix} = \begin{pmatrix} -\omega \lambda_2(t, \omega) - x_1 \\ \omega \lambda_1(t, \omega) - x_2 \end{pmatrix}, \\ \begin{pmatrix} \lambda_1(T, \omega) \\ \lambda_2(T, \omega) \end{pmatrix} &= \begin{pmatrix} \frac{\partial F(x(T, \omega), \omega)}{\partial x_1} \\ \frac{\partial F(x(T, \omega), \omega)}{\partial x_2} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \end{aligned}$$

With these values, after the control constraints become inactive, the time invariant Hamiltonian results of Theorem 2 apply. These conditions state that if we define:

$$\mathbf{H}(x^*, \lambda^*, u, t) = \int_{\Omega} H(x^*, \lambda^*, u, t, \omega) d\omega$$

then

$$\frac{d\mathbf{H}(x^*, \lambda^*, u^*, t)}{dt} = 0$$

Figure 4.13 shows that indeed the Hamiltonian $\mathbf{H}(x^*, \lambda^*, u^*, t)$ for this solution is approximately constant. Furthermore, Theorem 1 states that when control constraints are inactive:

$$\left(\frac{\partial}{\partial u} \int_{\Omega} H(x^*, \lambda^*, u, t, \omega) d\omega \right) \Big|_{u=u^*(t)} = 0$$

which gives us the switching structure:

$$u_1(t) = - \int_{\Omega} \lambda_1(t, \omega) d\omega$$

$$u_2(t) = - \int_{\Omega} \lambda_2(t, \omega) d\omega.$$

Figures 4.11 and 4.12 shows the switching behaviors of the controls and the costates for this solution..




Figure 4.9: Perturbed state values, $x_1^*(t, \omega)$, $x_2^*(t, \omega)$, created by the optimal control to the nominal system. Color is determined by $|\omega - \hat{\omega}|$.




Figure 4.10: State values created by the optimal control to the general parameter uncertainty problem. Color is determined by $|\omega - \hat{\omega}|$.




Figure 4.11: Values for u_1 and $\int_{\Omega} \lambda_1(t, \omega) d\omega$




Figure 4.12: Values for u_2 and $\int_{\Omega} \lambda_2(t, \omega) d\omega$




Figure 4.13: Hamiltonian values for general parameter uncertainty problem.

4.2.3 Channel Search

As another numerical example, we consider an instance of the classic channel search problem, created by [8], and studied by [12] and [7]. In this scenario, $K = 4$ searchers are tasked with surveying a channel of water of dimension $[0, 55] \times [0, 15]$ (the units for these values will remain unspecified). A target is floating down the surface of the channel from right to left in a straight line with a constant known velocity $v_a = .25$. Though there are four searchers and a single target, the channel of water is significantly larger than the sensor ranges of the searchers, which increases the difficulty of the search. The target's location in time is conditional on its unknown starting position, $\omega = [\omega_1, \omega_2]$, and is given by the function:

$$y(t, \omega) = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} \omega_1 - v_a t \\ \omega_2 \end{pmatrix}.$$

The searchers are constrained to the two-dimensional surface of the water. Their objective is to minimize the probability of not detecting the target in the given time interval $[0, 100]$. Each defender's state, x_k , is modeled as a Dubin's vehicle, with dynamics given by:

$$\dot{x}_k = \begin{pmatrix} \dot{x}_{k,1} \\ \dot{x}_{k,2} \\ \dot{x}_{k,3} \end{pmatrix} = \begin{pmatrix} v \sin x_{k,3} \\ v \cos x_{k,3} \\ u_k \end{pmatrix} = f_k(x_k, u_k), \quad k = 1, \dots, 4$$

and initial conditions $x_k(0) = x_{k,0} = [0, 3k, 0]^T$. The searchers' velocities are set as $v = 1$ and the searchers' turning rates, u_k , are constrained by $|u_k| \leq .5$. Parameter

space Ω is the entire rectangular region of the channel, $[0, 55] \times [0, 15]$, with a uniform probability density function.

A rate of detection model is provided by the Poisson Scan Model, descriptions of which can be found in [47] or [1]. The Poisson Scan Model designates that the rate of detection is given by:

$$r_k(x_k(t), y(t, \omega), t) = \lambda \Phi \left(\frac{F - a \|x_k(t) - y(t, \omega)\|^2}{\sigma} \right)$$

where Φ is the cumulative normal distribution. The values of λ , F , a , and σ are equipment specific constants which are set in this scenario to: $F = 0, a = 1, \lambda = 2, \sigma = 2.5$.

Applying the same methods as those used to derive the exponential detection probability for one searcher, the worst-case scenario probability, conditioned on ω , of none of the searchers detecting the target can be derived as:

$$P(t, \omega) = e^{-\int_0^t \sum_{k=1}^K r_k(x_k(\tau), y(\tau, \omega)) d\tau}.$$

Optimizing the expectation of this probability over domain of ω creates the following optimal control problem, in the form of problem **P**:

Channel Search Problem: Determine the control $u : [0, T] \rightarrow \mathbb{R}^4$ that minimizes the expectation

$$J = \int_{\Omega} \left(e^{-\int_0^T \sum_{k=1}^K r_k(x_k(\tau), y(\tau, \omega)) d\tau} \right) p(\omega) d\omega$$

Table 4.1: Parameter Values for Channel Search Problem

K	T	Ω	$p(\omega)$	$x_{k,0}$
4	100	$[0, 15] \times [0, 55]$	$\frac{1}{15} \cdot \frac{1}{55}$	$(0, \quad 3k, \quad 0)^T$

subject to the searcher dynamics $\dot{x}_k(t) = f_k(x_k, u_k)$, $k = 1, \dots, K$, with initial conditions $x_k(0) = x_{k,0}$, control constraint $|u_k(t)| \leq 0.5$, $\forall t \in [0, T]$, and the following values:

An illustration of a numerical solution to this problem is demonstrated in Figure 4.14. This solution was computed using 200 time discretization nodes and 25 nodes for each parameter dimension. The quadrature scheme in the time domain is Euler’s method and the scheme in the parameter dimensions is Legendre pseudospectral. To examine the effectiveness of the optimal control solution, we compare its final search probability with the probabilities generated by feasible trajectories created with heuristic methods. These trajectories are illustrated in Figure 4.15.

The first comparison trajectory is created by launching the searchers on horizontal search paths with constant velocities. When the searchers reach the end of parameter space in the x_1 direction, they reverse their direction while maintaining curvature constraints and continue back towards the left. The second comparison is a looping patrol pattern created by the parameterized curve $x_1 = 15 \sin(s(t))$, $x_2 = 3 \cos(3s(t))$. The parameterization by $s(t)$ maintains the constant searcher velocity $v = 1$ through satisfying the equation:

$$\frac{ds}{dt} \sqrt{\left(\frac{dx_1}{ds}\right)^2 + \left(\frac{dx_2}{ds}\right)^2} = 1.$$

Each patrol curve is furthermore translated to align with the initial positions of each searcher. The final comparison, meant to establish a base for poor-performing methods, is a sample of random walks created by sampling random headings from a uniform

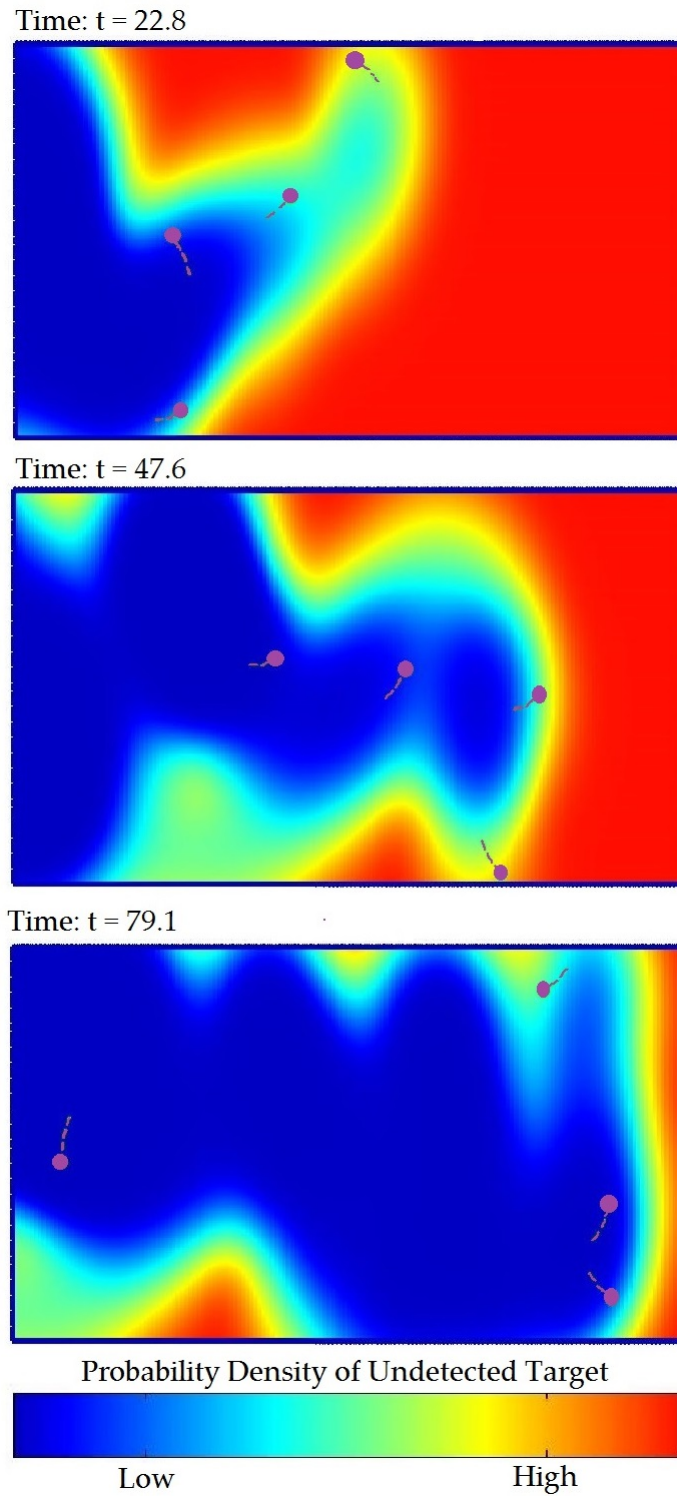


Figure 4.14: Snapshots of numerical solution to ‘Channel Search Problem’. Colors represent the log probability density value of an undetected target at a point at time t , i.e. the evolution over time of the probability density of target location given parameter value compounded with the probability of target non-detection for given parameter value. ‘Low’ = 4.14×10^{-6} , ‘High’ = 2.9317×10^{-4} .




Figure 4.15: Searching trajectories. Image (a) illustrates optimal control solutions, image (b) straight line patrols, and image (c) looping patrols.




Figure 4.16: Comparison of the performance of the optimal control solution vs heuristic methods.

distribution over $[-\frac{\pi}{2}, \frac{\pi}{2}]$. Performance results for these trajectories and the optimal control solution are compared in Figure 4.16.

A major issue in the numerical implementation of search problems has been the length of the computation time. Direct comparison of run times to previously published times is not possible in detail, due to computing platform differences and a paucity of published times. In [7] computation times are referred to in terms of days; in [1], algorithm times for single searchers searching over a two-dimensional parameter space range from 5,000 seconds to 20,000 seconds. The ability to now implement more

Table 4.2: Run Times Vs Nodes for Channel Search Problem

Time Nodes	Parameter Nodes	Run Time (seconds)
75	10×10	85.17
150	15×15	131.27
200	25×25	672.70

efficient methods (pseudospectral direct collocation with encoded sparsity and linearity) has reduced these times by an order of magnitude. Representative times are given in Table 4.2, as computed on a 2.3 GHz Intel Core i5 Macbook.

4.2.4 Kamikaze Swarm Scenario

We implement a kamikaze shooting problem as described in Section 2.2.2.1. An attacker is floating down the surface of a channel of dimension $[-20, 10] \times [0, 20]$ from right to left in a straight line with a constant known velocity $v_a = .25$. The attacker's location in time is conditional on its unknown starting position, $\omega = [\omega_1, \omega_2]$ with probability density function given by joint normalized beta distributions with parameters $\alpha = \beta = 3$ and the parameter space $[0, 10] \times [0, 20]$. There is one defender, moving as a Dubin's vehicle, with velocity set as $v = 1$ and turning rate, u , constrained by $|u| \leq .5$. The defender's initial state is $[-10, 15, 0]$. The firing rates of the defender and the attacker are modeled using the Poisson Scan Model, with identical calibration constants, given by: $F = 20, a = 1, \lambda = 2, \sigma = 10$. The objective is optimized over the time interval $[0, 75]$.

An illustration of a numerical solution to this problem is demonstrated in Figure 4.17. This solution was computed using 150 time discretization nodes and 25 nodes for each parameter dimension. The quadrature scheme in the time domain is




Figure 4.17: Snapshots of numerical solution to ‘Kamikaze Shooting Problem’. The magenta icon indicates the position of the HVU and the green icon is defender. Colors represent the log probability density value of a surviving attacker at a point at time t . ‘Low’ = 4.14×10^{-6} , ‘High’ = 2.9317×10^{-4} .

Euler’s method and the scheme in the parameter dimensions is Legendre pseudospectral. The final probability of HVU destruction in this implementation is 9.32%. This low percentage is achieved despite the fact that the probability of destroying the attacker in this case is only 7.69%. To gauge the efficacy of this trajectory, this result can be contrasted with the performance of a trajectory generated with the same numerical methods but using the objective of merely maximizing the probability of destruction of the attacker. In this case, the probability of destroying the attacker can be increased, to 15.31%. However, due to the dispersed attention, the resulting probability of HVU destruction comes out to 79.65%.

Chapter 5

Conclusions and Future Work

In the beginning of this thesis, three aspects were identified as necessary for progressing the research devoted to The Optimal Search Problem: the design of useful interactive performance models and kinematic models, the development of a general framework for dealing with parameter uncertainty in both the cost and dynamics of optimal control problems, and a numerical algorithm for generating solutions. Each of these aspects have now been addressed. Section 2 provides several new models for both kinematic and performance interaction in multi-agent systems. Section 3 puts forth a general mathematical framework along with useful necessary conditions for optimality. And finally, refComputational Approach Section provides a numerical algorithm for finding solutions to this class of optimal control problems. The following sections discuss future areas of research.

5.1 Increasing Computational Speed

One of the foremost challenges with this type of problem is the *curse of dimensionality*. Since the computational scheme described in Section 4 is based on a discretization of the parameter space Ω , the dimension of that space impacts the size of Problem \mathbf{P}^M . Both the number of operations necessary for a single evaluation of the cost function:

$$\bar{J}^M[\bar{X}_M, u] = \sum_{i=1}^M \left[F(\bar{x}_i^M(T), \omega_i^M) + \int_0^T r(\bar{x}_i^M(t), u(t), t, \omega_i^M) dt \right] \alpha_i^M$$

and the number of state variables present in Problem \mathbf{P}^M :

$$\begin{cases} \dot{\bar{x}}_i^M(t) = f(\bar{x}_i^M(t), u(t), \omega_i^M) \\ \bar{x}_i^M(0) = x_0(\omega_i^M), \end{cases} \quad i = 1, \dots, M. \quad (5.1)$$

grow proportionally to M , the number of nodes utilized in the approximation scheme. And for quadrature schemes which are formed using the full tensor product of the parameter space, the number of nodes, M , increases exponentially with refinement of the discretization grid. For a d -dimensional parameter space, M grows proportionally to m^d , where m is the number of nodes used in each dimension. This exponential growth makes computation using these types of quadrature methods infeasible for problems with parameter spaces with high (or sometimes even moderate) dimension. The issue of exponential growth in quadrature schemes is not unique to this problem, and many methods have been developed which may be fruitful for application.

5.1.0.1 Monte Carlo Methods

Monte Carlo methods approximate integrals of high dimension by sampling M random points from the domain Ω and approximating the integral as the average:

$$\int_{\Omega} h(\omega) d\omega \approx \frac{1}{M} \sum_{i=1}^M h(\omega_i)$$

[48] The convergence rate of Monte Carlo integration is proportional to the quantity $\frac{1}{\sqrt{M}}$. [48] Although this is a relatively slow convergence rate, it is notable for being independent of the dimension of Ω . Furthermore, the increasing refinement of the discretization grid is no longer subject to exponential growth. In application to Problem **P**, the M sample points also designate the collocation points for the state variables, in the style of Equation 4.1. Figure 5.1 shows an implementation of the Kamikaze Swarm Scenario of Section 2.2.2.1 using Monte Carlo methods and a sample size of $M = 200$ to approximate a swarm of ten attackers with an overall parameter domain of dimension $d = 20$. The consistency of applying these methods to Problem **P** under various regularity assumptions is analyzed in [49]. What are still needed, however, are methods for estimating the proximity of numerical answers to the true optimal solution. Since the convergence rate of Monte Carlo methods is slow, these tools are crucial as the order of magnitude of the required number of nodes can vary drastically.

5.1.0.2 Quasi-Monte Carlo Methods

The square root convergence rate of Monte Carlo methods is not a necessity, as demonstrated by the existence of much more rapidly converging quadrature meth-




Figure 5.1: High Dimensional Kamikaze Scenario Implemented with Monte Carlo; $d = 20$, $M = 200$

ods. One approach to bridging the gap between the convergence rates of random (but dimensionally independent) and the convergence rates of deterministic (but dimensionally dependent) methods is to construct a sampling scheme which is still independent of the domain dimension but which utilizes structurally deterministic features to construct sampling points. Sequences of sampling points which in which refinement is still independent of dimension but which more efficiently span domains can, it turns out, be constructed. These sequences are generated using number theoretic approaches and their use is referred to under the umbrella term ‘quasi-Monte Carlo methods.’ As an example of the convergence properties of quasi-Monte Carlo methods, one can consider the family of Sobol sequences. Sobol sequences are generated by considering the roots of a chosen polynomial in relation to the field of integers modulo two. It can be shown that their rate of convergence is proportional to $\frac{(\ln M)^d}{M}$, or in other words *asymptotically* approaching $\frac{1}{M}$. [48]

5.1.0.3 Sparse Grid Methods

Sparse grid methods, first developed by Smolyak in 1963 [50] are an alternate method for constructing refinement sequences which, in contrast to Monte Carlo related methods, originate from full tensor products of discretization grids and work to reduce them to sequences which grow with less than strictly exponential rates. For a space Ω of dimension d , polynomial approximation schemes such as the pseudospectral collocation scheme described in Section 4.1.3 rely on a number of basis functions equal to the full product of $M = \prod_{l=1}^d m_l$ total nodes, where m_l is the number of nodes taken in each dimension. Sparse grid methods construct sequences of subsets of these basis functions,




Figure 5.2: Cost and Error for Sparse Grid vs Full Grid System of [4] with $d = 4$

which grow in number at a slower rate as the grid is refined but at the cost of a reduced convergence rate. The tradeoff, however, is often beneficial.

For instance, [4] describes a scheme based interpolation with a hierarchical basis of piecewise linear basis functions for functions with mixed bounded second derivatives. For an initial full grid created with m nodes in each dimension the full tensor product of the basis functions in each dimension would yield $M = m^d$ basis functions overall with an error of order $O(m^{-2})$. However, applying sparse grid methods provides a subset sequence in which the number of basis functions grows with order $O(m(\log m)^{d-1})$ with error on the order of $O(m^{-2}(\log m)^{d-1})$ in the L_2 -norm. Figure 5.2 demonstrated the cost and error rates for this approach versus using a full grid. As [51] describes, this approach can also be generalized to basis sets built with piecewise polynomials of order p , in which error becomes on the order of $O(m^{-p}(\log m)^{d-1})$.

Since sparse grids are constructed by utilizing the underlying structure of the product space of interpolating functions in each dimension, one major possibility in the application of this method to the problems discussed in this thesis is the specification of

a sparse grid which caters to specific properties of the problem. For instance, cost sensitivity in relation to the cost function of the optimal control problem may be robust with respect to certain parameter dimensions and robust with respect to others, as can the uncertain state dynamics. Were this to be quantified, a sparse grid scheme can be constructed with a varyingly refined grid in the relative dimensions. A dimension-adaptive generalization of basic sparse grid schemes for numerical integration was developed in [52]

An area of future research would be to develop tools for quantitatively assessing the importance of different dimensions in this problem—where both the integration in the cost function and the progression of the ODE’s of the state variables are effected. Another area which holds possibility for the future is to examine the utilization of a sparse grid approach for intertwining the discretization of the parameter domain with the time domain. Currently, the algorithm in Section 4, is using the full product of time nodes by parameter nodes, leading to an algorithm which grows in size with order $O(MN)$ where M is the number of parameter nodes and N is the number of time nodes. Considering the time domain as an additional dimension from the start, rather than separating the problem into the two distinct steps denoted by Problem \mathbf{P}^M and Problem \mathbf{P}^{MN} , then applying sparse grid methods may yield further savings.

5.1.0.4 Parallelization

In addition to seeking schemes which decrease the size of the numerical problem, there are also many avenues for increasing the speed of computation for given sizes. One such avenue is that of parallelization of processing of various problem components

during the course of running the iterative NLP optimization algorithm.

For instance, the discretization of parameter space using M nodes creates M differential equations:

$$\begin{cases} \dot{\bar{x}}_i^M(t) = f(\bar{x}_i^M(t), u(t), \omega_i^M) \\ \bar{x}_i^M(0) = x_0(\omega_i^M), \end{cases} \quad i = 1, \dots, M. \quad (5.2)$$

which, notably, are uncoupled for the different values of ω_i^M . During each iteration of the NLP optimization algorithm which provides a guess control $u(t)$, each of these differential equations can be processed simultaneously.

In addition to this fundamental feature, many of the multi-agent models developed in Section 2 lend themselves to further decomposition. For instance, one can consider the mutual attrition scenario of Section 2.2.2.2 in the case where the distributions for the uncertain parameter influencing each attacker are independent of each other. For L attackers, the cost function is:

$$J = \int_{\Omega} [1 - p_o(T, \omega)] p(\omega) d\omega$$

where

$$\dot{p}_0(t, \omega) = -p_0(t, \omega) \sum_{l=1}^L q_l(t, \omega) r_{0,l}(y_l(t, \omega), x_0(t))$$

and with each l -th attacker dependent on parameter w_l with p.d.f. $p_l(w_l)$ over domain

Ω_l , this can be decomposed as:

$$\begin{aligned}
J &= \int_{\Omega} [1 - p_o(T, \omega)] p(\omega) d\omega \\
&= \int_{\Omega_1} \int_{\Omega_2} \dots \int_{\Omega_L} [1 - p_o(T, \omega_1, \dots, \omega_L)] p_1(\omega) p_2(\omega) \dots p_L(\omega) d\omega_1 d\omega_2 \dots d\omega_L \\
&= 1 - \int_{\Omega_1} \dots \int_{\Omega_L} \left[e^{-\int_0^T \sum_{l=1}^L q_l(t, \omega_l) r_{0,l}(y_l(t, \omega_l), x_0(t)) dt} \right] p_1(\omega) \dots p_L(\omega) d\omega_1 \dots d\omega_L \\
&= 1 - \int_{\Omega_1} \dots \int_{\Omega_L} \left[\prod_{l=1}^L e^{-\int_0^T q_l(t, \omega_l) r_{0,l}(y_l(t, \omega_l), x_0(t)) dt} \right] p_1(\omega) \dots p_L(\omega) d\omega_1 \dots d\omega_L \\
&= 1 - \prod_{l=1}^L \int_{\Omega_l} e^{-\int_0^T q_l(t, \omega_l) r_{0,l}(y_l(t, \omega_l), x_0(t)) dt} p_l(\omega) d\omega_l
\end{aligned}$$

Each component, $\int_{\Omega_l} e^{-\int_0^T q_l(t, \omega_l) r_{0,l}(y_l(t, \omega_l), x_0(t)) dt}$ can be calculated independently, and thus implemented in parallel. Similar decomposition possibilities for independent agents are available in The Optimal Search Problem for multiple agents and The Kamikaze Swarm Scenario of Section 2.2.2.1.

5.2 Expanding Mathematical Framework

5.2.0.5 Feedback Control

The current design of Problem **P** is as an open-loop control problem. An obvious engineering desire, however, is to incorporate an aspect of feedback control to increase the robustness of solutions. While the case of feedback driven by information is obviously preferable, the question remains of how to incorporate that information optimally in cases where, unlike this academic example, the feedback solution is not available a priori, and also in cases where such information may only be gathered at

discrete time points. One direction is to utilize repeated measurements to gradually tighten the estimated p.d.f. of the parameter. For instance, if we define the map $I(\omega) : \Omega \rightarrow X$ by:

$$I(\omega) = \int_0^t f(x(\tau, \omega), u(\tau), \omega) d\tau + x(0, \omega)$$

then for an unknown parameter ω with p.d.f. $p_0(\omega)$, each measurement $x(t_i, \omega) \in X$ constrains the possibilities for the values of ω to lie in $\omega \in I^{-1}(x(t_i, \omega))$, a fact from which an updated p.d.f. $p_1(\omega)$ can be calculated.

5.2.0.6 End Time State Constraints

Many engineering applications necessitate the use of end time state or mixed state/control constraints, including minimum time problems, which minimize the time it takes to satisfy some state/control constraint. In this case, however, where the state dynamics are subject to uncertainty, the design of what particular kind of end time constraints should be enforced is an open question. In ensemble control, a notion of ‘ensemble controllability’ has been established for some systems, in the sense that the final states of a system can be steered to within an ϵ ball of a goal.[27] However, this threshold of controllability has only been shown to be possible for a limited number of simple systems. It may be the more standard case that steering to within an ϵ ball of a goal for the *entire* distribution of end time values is infeasible, in which case other types of constraints need to be considered. A possible alternate condition is a combination of constraints on the expectation and variance of the end time distributions.

5.3 Furthering Applications

5.3.1 Protective Herding

Currently, in collaboration with Panos Lambrianides, the decomposition techniques and parallelization observations of Section 5.1 are in the process of being implemented in a manner which takes advantage of GPU computing with an end goal of solidifying a software architecture for parallelization in such problems. The current application for this methodology is a scenario which combines the performance metric of Section 2.2.2.2 with the herding dynamics of Section 2.1.1. In the scenario, the objective is to defend an HVU but with two inputs at a defender's disposal: firing power and the non-destructive force of herding. The goal in implementing this scenario is to aid the development of tactical simulations which explore the tradeoff between destructive force (firing power) and non-destructive force (herding).

Protective Herding Scenario: Given the probability density function $\phi : \Omega \rightarrow \mathbb{R}$ determine the controls $u_k : [0, T] \rightarrow U \in \mathbb{R}^{n_u}$, $k = 1, \dots, K$, that minimize:

$$J = \int_{\Omega} [1 - p_o(T, \omega)] p(\omega) d\omega$$

subject to:

$$\begin{aligned}
\dot{x}_k(t) &= f(x(t), u(t)), & x_k(0) &= x_k^0 \\
\dot{y}_{l,1}(t, \omega) &= \frac{H_{l,1}(y_l(t, \omega), x_1(t), \dots, x_k(t))}{\|H_l(y_l(t, \omega), x_1(t), \dots, x_k(t))\|} v_l, & y_{l,1}(0, \omega) &= y_{l,1}^0 \\
\dot{y}_{l,2}(t, \omega) &= \frac{H_{l,2}(y_l(t, \omega), x_1(t), \dots, x_k(t))}{\|H_l(y_l(t, \omega), x_1(t), \dots, x_k(t))\|} v_l, & y_{l,2}(0, \omega) &= y_{l,2}^0 \\
\dot{p}_0(t, \omega) &= -p_0(t, \omega) \sum_{l=1}^L q_l(t, \omega) r_{0,l}(y_l(t, \omega), x_0(t)), & p_0(0, \omega) &= 0 \\
\dot{p}_k(t, \omega) &= -p_k(t, \omega) \sum_{l=1}^L q_l(t, \omega) r_{a,l}(y_l(t, \omega), x_k(t)), & p_k(0, \omega) &= 0 \\
\dot{q}_l(t, \omega) &= -q_l(t, \omega) \sum_{k=1}^K p_k(t, \omega) r_{d,k}(x_k(t), y_l(t, \omega)), & q_l(0, \omega) &= 0
\end{aligned}$$

for all $l = 1, \dots, L$, $k = 1, \dots, K$, and $\omega \in \Omega$ with possible control constraint $g(u(t)) \leq 0$, $\forall t \in [0, T]$. Let H_l be defined as per Section 2.4.

Figure 5.3 shows the behavior of this model for a case with 2 defenders and two ‘swarms’ of attackers defined by a range of uncertain distance around a deterministically reacting centroid. Future work will focus on numerical validation, model refinement, and strategy analyses such as the tradedoff between raising firing power versus utilizing more herders. Additionally, the goal is to be able to implement this scenario for large scale multi-agent situations with independent attackers swarms.

5.3.2 Longterm Autonomous Flight with Solar-Powered Thermalizing Gliders

Currently in nature, longterm flight capabilities are observed in birds which modern aeronautics has not replicated. Large birds can fly for days at a time, expending

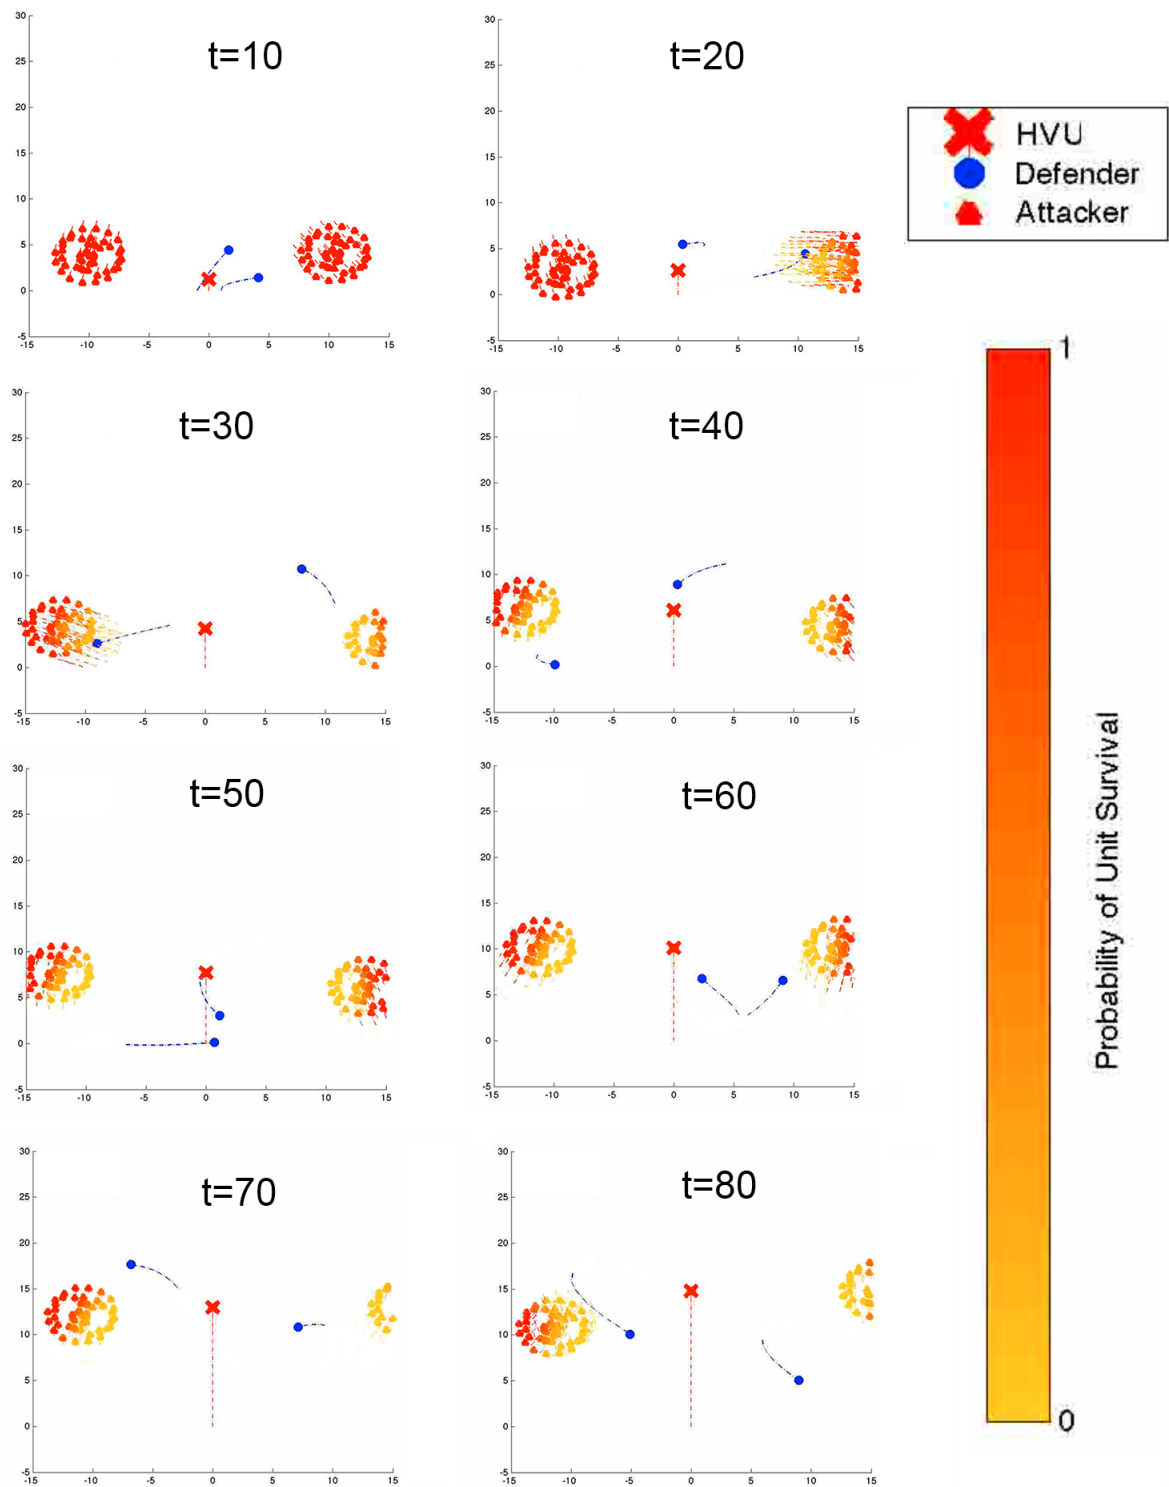


Figure 5.3: Snapshots of Protective Herding Scenario with 2 Defenders and 2 Dependent Attacker Swarms with 50 Members Each

little energy in the process. [53] These birds utilize so-called “dynamic soaring” [54], a technique which gains kinetic energy by exploiting thermal updrafts, the columns of rising air created by temperature gradients.

The ability to implement low energy, longterm flight in autonomous vehicles has numerous applications, including weather and earth surface monitoring and communication coverage. Recent work at the Naval Postgraduate School has progressed on developing routines for implementing dynamic soaring techniques on a selection of autonomous, solar-powered gliders.[55] These routines combine automated detection of thermal updrafts once encountered and the subsequent behavior needed to latch into them and utilize their kinetic energy.

The detection of thermal updrafts is a feature which requires both global and local planning. On the local level, recognition of the initial entry into an updraft is needed. On the global level, the design of flight paths which optimize in some fashion both the probability of encountering a updrafts as well as the accomplishment of some flight goal (for example, moving from point A to point B). Current global updraft search plans in [55] utilize a Traveling Salesman Approach, identifying mission objective locations and additionally ‘points of interest’, which are spots deemed likely to have updrafts to prior information such as elevation maps, infrared imagery, land use maps, and meteorological data. A goal of future work on this project is to develop more optimal strategies for designing global flight plans using this type of data and the prior probability estimates it provides for the location of updrafts.

Bibliography

- [1] Foraker, J., *Optimal Search for Moving Targets in Continuous Time and Space Using Consistent Approximations*, Ph.D. thesis, Naval Postgraduate School, September 2011.
- [2] Chung, T. H., Hollinger, G. A., and Isler, V., “Search and Pursuit-Evasion in Mobile Robotics,” *Autonomous Robots*, Vol. 31, 2011, pp. 299–316.
- [3] Benkowski, S. J., Monticino, M. G., and Weisinger, J. R., “A Survey of the Search Theory Literature,” *Naval Research Logistics Quarterly*, Vol. 38, 1991, pp. 469–494.
- [4] Gerstner, T. and Griebel, M., “Sparse grids,” *Encyclopedia of Quantitative Finance*, 2010.
- [5] Stone, L. D., “What’s Happened in Search Theory since the 1975 Lanchester Prize?” *Operations Research*, Vol. 37, No. 3, 1989, pp. 501–506.
- [6] Pursiheimo, U., “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*, Vol. 2, 1976, pp. 259–264.

- [7] Chung, H., Polak, E., Royset, J. O., and Sastry, S., “On the optimal detection of an underwater intruder in a channel using unmanned underwater vehicles,” *Naval Research Logistics (NRL)*, Vol. 58, No. 8, 2011, pp. 804–820.
- [8] Koopman, B., *Search and Screening*, Naval Historical Center, 1946.
- [9] Koopman, B., “The Theory of Search II: Target Detection,” *Operations Research*, Vol. 4, No. 5, October 1956, pp. 503–531.
- [10] Hellman, O., “On the Optimal Search for a Randomly Moving Target,” *SIAM Journal on Applied Mathematics*, Vol. 22, No. 4, June 1972, pp. 545–552.
- [11] Ohsumi, A., “Optimal Search for a Markovian Target,” *Naval Research Logistics Quarterly*, Vol. 38, 1991, pp. 531–554.
- [12] Washburn, A., “On Patrolling a Channel,” *Naval Research Logistics Quarterly*, Vol. 29, No. 4, December 1982, pp. 609–615.
- [13] Lukka, M., “On the Optimal Searching Tracks for a Moving Target,” *SIAM Journal on Applied Mathematics*, Vol. 31, No. 1, Jan. 1977, pp. 126–132.
- [14] Phelps, C., Gong, Q., Royset, J. O., Walton, C., and Kaminer, I., “Consistent approximation of a nonlinear optimal control problem with uncertain parameters,” *Automatica*, 2014.
- [15] Walton, C. L., Gong, Q., Kaminer, I., and Royset, J. O., “Optimal Motion Planning for Searching for Uncertain Targets,” *Proceedings of the 19th IFAC World Congress*, Vol. 19, 2014.

- [16] Pontryagin, L., Boltyanskii, V., Gamkrelidze, R., and Mishchenki, E., *The Mathematical Theory of Optimal Processes*, Interscience Publishers, 1962.
- [17] Pursiheimo, U., “On the Optimal Search for a Target Whose Motion is Conditionally Deterministic with Stochastic Initial Conditions on Location and Parameters,” *SIAM Journal on Applied Mathematics*, Vol. 32, No. 1, Jan. 1977, pp. 105–114.
- [18] Phelps, C., *Computational Optimal Control of Nonlinear Systems with Parameter Uncertainty*, Ph.D. thesis, UCSC, June 2014.
- [19] Walton, C., Phelps, C., Gong, Q., Royset, J. O., and Kaminer, I., “A Numerical Algorithm for Optimal Control of Systems with Parameter Uncertainty,” .
- [20] Gong, Q., Ross, I. M., Kang, W., and Fahroo, F., “Connections Between the Covector Mapping Theorem and Convergence of Pseudospectral Methods for Optimal Control,” *Computational Optimization and Applications*, Vol. 41, No. 3, 2008, pp. 307–335.
- [21] Phelps, C., Gong, Q., Royset, J. O., and Kaminer, I., “Consistent approximation of an optimal search problem,” *CDC*, 2012, pp. 630–637.
- [22] Lu, Z., *Cooperative Optimal Path Planning for Herding Problems*, Master’s thesis, Texas A&M, December 2006.
- [23] Shedio, S. A., *Optimal Control for a Two Player Dynamic Pursuit Evasion Game; The Herding Problem*, Ph.D. thesis, Virginia Polytechnique Institute, January 2002.
- [24] Washburn, A. and Kress, M., *Combat Modeling*, Springer, 2009.

- [25] Hess, J., Beinhofer, M., and Burgard, W., “A Probabilistic Approach to High-Confidence Cleaning Guarantees for Low-Cost Cleaning Robots,” *IEEE International Conference on Robotics and Automation*, 2014.
- [26] Choset, H., “Coverage for robotics - A survey of recent results,” *Annals of Mathematics and Artificial Intelligence*, Vol. 31, 2001, pp. 113–126.
- [27] Li, J.-S. and Khaneja, N., “Ensemble Control of Linear Systems,” *Proceedings of the 46th IEEE Conference on Decision and Control*, December 2007.
- [28] Becker, A. T., *Ensemble Control of Robotic Systems*, Ph.D. thesis, University of Illinois at Urbana-Champaign, 2012.
- [29] Ruths, J. and Li, J.-S., “Optimal Control of Inhomogeneous Ensembles,” *Transactions on Automatic Control*, Vol. 57, No. 8, 2012.
- [30] Terwiesch, P., Ravemark, D., Schenker, B., and Rippin, D. W. T., “Semi-batch process optimization under uncertainty: Theory and experiments,” *Computers and Chemical Engineering*, Vol. 22.1, 1998, pp. 201–213.
- [31] Ruppen, D., Benthack, C., and Bonvin, D., “Optimization of batch reactor operation under parametric uncertainty- computational aspects,” *Journal of Process Control*, Vol. 5, No. 4, 1995, pp. 235–240.
- [32] A.G. Butkovsky, A. E. and Lur’e, K., “Optimal Control of Distributed Systems (A Survey of Soviet Publications),” *SIAM J. Control*, Vol. 6, No. 3, 1968.
- [33] Lur’e, K., “The Mayer-Bolza Problem for Multiple Integrals and the Optimization

- of the performance in Systems with Distributed Parameters,” *PMM*, Vol. 27, No. 5, May 1963, pp. 842–853.
- [34] Armand, J.-L. P., “Applications of the Theory of Optimal Control of Distributed-Parameter Systems to Structural Optimization,” *NASA Contractor Report*, 1972.
- [35] Courant, R. and Hilbert, D., *Methods of Mathematical Physics, Chapter IV, The Calculus of Variations*, Vol. I, 1953.
- [36] Akhiezer, N. I., *The Calculus of Variations*, Blaisdell Publishing Company, 1962.
- [37] Darlington, J., Pantelides, C., Rustem, B., and Tanyi, B., “Decreasing the Sensitivity of Open-Loop Optimal Solutions in Decision Making under Uncertainty,” *European Journal of Operational Research*, Vol. 121, 2000, pp. 343–362.
- [38] Fisher, J. and Bhattacharya, R., “Optimal Trajectory Generation with Probabilistic System Uncertainty Using Polynomial Chaos,” *Journal of Dynamic Systems, Measurement, and Control*, Vol. 133, January 2011.
- [39] Hover, F. S. and Triantafyllou, M. S., “Application of Polynomial Chaos in Stability and Control,” *Automatica*, Vol. 42, 2006, pp. 789–795.
- [40] Phelps, C., Gong, Q., Royset, J. O., and Kaminer, I., “Consistent Approximation of an Optimal Search Problem,” *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, Dec. 2012, pp. 630–637.
- [41] Gill, P. E., Murray, W., and Saunders, M. A., “SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization,” *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131.

- [42] Ross, I. M., “A Roadmap for Optimal Control: The Right Way to Commute,” *Annals of the New York Academy of Sciences*, Vol. 1065, January 2006, pp. 210–233.
- [43] Gong, Q., Kang, W., Bedrossian, N. S., Fahroo, F., Sekhavat, and Bollino, K., “Pseudospectral Optimal Control for Military and Industrial Applications,” *Decision and Control, 2007 46th IEEE Conference on*, 2007.
- [44] Ross, M. and Fahroo, F., *Lecture Notes in Control and Information Sciences*, Vol. 295, chap. Legendre Pseudospectral Approximations of Optimal Control Problems, Springer-Verlag, 2003.
- [45] Becker, A. and Bretl, T., “Motion Planning under Bounded Uncertainty Using Ensemble Control,” *Robotics: Science and Systems*, Citeseer, 2010.
- [46] Zlotnik, A. and Li, J.-S., “Synthesis of Optimal Ensemble Controls for Linear Systems using the Singular Value Decomposition,” .
- [47] Kim, K., *Approximating the Poisson Scan and Lambda-Sigma Acoustic Detection Model with a Random Search Formula*, Master’s thesis, Naval Postgraduate School, December 2009.
- [48] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P., *Numerical Recipes in Fortran: The Art of Scientific Computing*, Cambridge University Press Cambridge, 2nd ed., 1996.
- [49] Phelps, C., Royset, J. O., and Gong, Q., “Sample average approximations in op-

- timal control of uncertain systems.” *Proceedings of the 52nd IEEE Conference on Decision and Control (CDC 2013)*, 2013.
- [50] Smolyak, S. A., “Quadrature and interpolation formulas for tensor products of certain classes of functions,” *Dokl. Akad. Nauk SSSR*, Vol. 4, 1963, pp. 240–243.
- [51] Bungartz, H.-J. and Griebel, M., “Sparse grids,” *Acta numerica*, Vol. 13, 2004, pp. 147–269.
- [52] Gerstner, T. and Griebel, M., “Dimension–adaptive tensor–product quadrature,” *Computing*, Vol. 71, No. 1, 2003, pp. 65–87.
- [53] Arnould, J., Briggs, D., Croxall, J., Prince, P., and Wood, A., “The foraging behaviour and energetics of wandering albatrosses brooding chicks,” *Antarctic Science*, Vol. 8, No. 03, 9 1996, pp. 229–236.
- [54] Pennycuik, C., “The flight of petrels and albatrosses (Procellariiformes), observed in South Georgia and its vicinity,” *Philosophical Transactions of the Royal Society of London*, Vol. B300, 1982, pp. 75–106.
- [55] Dobrokhodov, V., Jones, K., and Camacho, N., “Cooperative Autonomy of Multiple Solar-Powered Thermalizing Gliders,” *World Congress*, Vol. 19, 2014, pp. 1222–1227.