Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Optimal Motion Planning for Searching for Uncertain Targets ⋆

Claire L. Walton * **Qi Gong** ** **Isaac Kaminer** *** **Johannes O. Royset** ****

* *University of California Santa Cruz, Santa Cruz, CA 95060 USA (e-mail: cwalton @ soe.ucsc.edu).*
** *University of California Santa Cruz, Santa Cruz, CA 95060 USA (e-mail: qigong @ soe.ucsc.edu).*
*** *Naval Postgraduate School, Monterey, CA 93943 USA (e-mail: kaminer @ nps.edu)*
**** *Naval Postgraduate School, Monterey, CA 93943 USA (e-mail: joroyset @ nps.edu)*

Abstract: This paper explores the potential for applying newly available numerical methods in optimal control to solve motion planning problems created by the search for targets with motion uncertainty characterized by constant but unknown parameters. These recent developments enable the efficient computation of numerical solutions for search problems with multiple searchers, nonlinear dynamics, and a broad class of objectives. We demonstrate the efficacy of these methods through implementing a multi-agent optimal search problem. We then derive an expansion of the optimal search modeling framework which facilitates the consideration of multi-agent searching problems with more general strategic objectives and utilize this expanded framework to implement an example combat defense scenario.

Keywords: trajectory planning, autonomous vehicles, nonlinear control, numerical methods

1. INTRODUCTION

An area of interest in robotics and autonomous vehicle research is the development of optimal or sub-optimal motion planning 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.

In this paper we will visit the question of modeling and numerically implementing motion planning problems created by the task of optimally searching for a target, given uncertainty in target motion. We will focus on the case where target motion is characterized by a set of uncertain parameters. We demonstrate that this gives rise to a class of optimal control problems where the constant but unknown parameter values are incorporated in the cost function. Recently, algorithms for obtaining numerical solutions to this class of problems have been developed–in [Foraker (2011)], [Chung et al. (2011)], and [Phelps et al. (2012)]. Due to these new developments we are able to use more efficient numerical methods, such as pseudospectral computational optimal control methods [Gong et al. (2006)], than have been previously implemented. We are furthermore able to flexibly incorporate a larger class of objectives, state and control constraints, and nonlinear searcher and target dynamics. This has created a new potential which can be utilized to address a greater variety of problems.

In the first part of this paper, we describe a framework for modeling basic search problems, derived in [Koopman (1956)], which was originally developed as part of the antisubmarine warfare efforts of WWII and has been widely utilized in industrial and tactical applications since [Chudnovsky and Chudnovsky (1988)]. We elaborate on the mathematical problem this model creates and the new numerical methods now available for solving this problem. We then utilize these methods to implement a high-dimensional search problem with multiple searchers, nonlinear searcher dynamics, and control constraints.

In the second part of the paper, we consider a more 'realistic' motion planning problem, where the act of searching for targets is done in the service of a strategic objective of protecting assets from the targets. This is a goal which can arise in many applications similar to but distinct from search; for instance, during an oil spill a valuable at-risk ecosystem may need to be prioritized during clean up. In these situations, optimization of a basic search objective may be valuable, however this approach would be indirect– it would not necessarily provide an optimization of the primary goal. Direct optimization of the primary strategic goal requires the modeling of additional directions of heterogeneous influences, which is not a feature of current search models. We build a framework for modeling this additional direction of target influence and demonstrate that the new numerical methods established described in

⋆ This work is supported by US Office of Naval Research under Grant N0001412WX21229.

the first section can be applied to these scenarios as well. We then use these methods to numerically implement this framework on an example tactical simulation.

2. A BASELINE SCENARIO: THE OPTIMAL SEARCH PROBLEM

A basic design for motion planning when searching for uncertain targets is provided by 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 uncertain features, given detection equipment capabilities and some limitations on searcher trajectories. The problem has been studied extensively in the fields of applied mathematics and operations research; a review is given by [Stone (1989)].

In order to construct a performance criterion for the optimal search problem, the probability of target detection must be modeled. In [Koopman (1956)] an exponential probability of detection model is derived which has since become ubiquitous in optimal search literature. The exponential detection model follows from the assumption that the *instantaneous rate of detection* of a target can be effectively modeled. Given the position of a searcher at $x(t)$ and a target at $y(t)$, this instantaneous rate of detection is a function $r(x(t), y(t), t)$ such that the probability of detection in a sufficiently small interval $[t, t + \Delta t]$ is independent from previous time intervals and given by the quantity $r(x(t), y(t), t)\Delta t$. The rate function $r(x(t), y(t), t)$ 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.

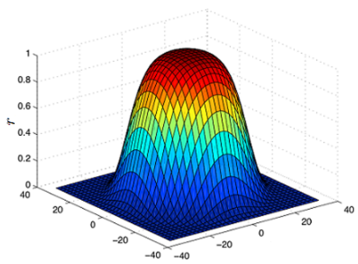

Fig. 1. Example detection rate function: Poisson Scan Model

Proceeding with these assumptions, an explicit formula for the probability of target detection is derived. If we denote the probability of *non-detection* with the function $p(t)$, the independence of the time intervals creates the following difference equation:

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

As $\Delta t \to 0$ this has the exponential solution:

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

A variety of probabilities for multiple searchers and targets can be derived using similar methods. These include the probability of detecting all targets and the worst-case scenario probability of detecting none of the targets. In general, this can be expressed as:

$$P(t) = G\left(\int_0^t r(x(\tau), y(\tau), \tau)d\tau\right).$$

This probability function is further conditioned on the uncertain target motion, which can be modeled in a variety of ways, including as a diffusion process with stochastic parameters, [Hellman (1972)], and Markovian motion, [Ohsumi (1991)]. In this paper, we focus exclusively on the model of conditionally deterministic motion, which assumes that the motion of the targets is given entirely by a function of time and a parameter ω. This parameter is an element of a bounded parameter space $\Omega \subset \mathbb{R}^n$ and has a known probability density function $p : \Omega \to \mathbb{R}$.

Incorporating conditionally deterministic target motion $y(t|\omega) = h(t, \omega)$ leads to a probability quantity which is itself a random variable, i.e. $P(t, \omega)$. A natural performance measure is to minimize the expectation of this random variable over a time interval $[0, T]$. This gives rise to the class of cost functions:

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

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

2.1 Computational Framework

The optimal search problem as formulated above is an instance of the following class of optimal control problems:

Problem P: Given probability density function $p : \Omega \to \mathbb{R}$, determine the control $u : [0, T] \to U \in \mathbb{R}^{n_u}$ 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: $\dot{x}(t) = f(x(t), u(t))$ with initial condition $x(0) = x_0$ and control constraint $g(u(t)) \le 0$, $\forall t \in [0, T]$.

This is a nonstandard optimal control problem due to the integration over parameter space. Numerical methods for this problem focus on creating an approximation of parameter space in coordination with approximations of state and control spaces. In [Phelps et al. (2012)] it was shown that, under minimal assumptions on compactness and continuity, a suitable approximation of parameter space may consist of a set of nodes, $\{\omega_i^M\}_{i=1}^M \in \Omega$, along with a set of weights, $\{a_i^M\}_{i=1}^M \in \mathbb{R}$ such that the approximation of integration over parameter space is convergent; that is, such that for any continuous function $h : \Omega \to \mathbb{R}$, the following convergence holds:

$$\int_\Omega h(\omega)d\omega = \lim_{M\to\infty} \sum_{i=1}^M h(\omega_i^M)a_i^M$$

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

Problem $\mathbf{P^M}$: Given probability density function $p : \Omega \to \mathbb{R}$, determine the control $u : [0, T] \to U \in \mathbb{R}^{n_u}$ that minimizes the cost:

$$J^M = \sum_{i=1}^M a_i^M \left[F(x(T), \omega_i^M) + G\left(\int_0^T r(x(t), u(t), t, \omega_i^M)dt\right)\right] p(\omega_i^M)$$

subject to: $\dot{x}(t) = f(x(t), u(t))$ with initial condition $x(0) = x_0$ and control constraint $g(u(t)) \le 0$, $\forall t \in [0, T]$.

The approximate problem is a standard control problem, which can be addressed with a variety of established methods. To implement the scenarios in this paper, we utilize the method of direct collocation, as described in [Gong et al. (2006)]. 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 [Gill et al. (2005)]

2.2 An Optimal Search Problem with Multiple Searchers

As a numerical example, we consider an instance of the classic channel search problem, created by [Koopman (1946)], and studied by [Washburn (1982)] and [Chung et al. (2011)]. 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:

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

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 [Kim (2009)] or [Foraker (2011)]. 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 $u : [0, T] \to \mathbb{R}^4$ that minimizes the expectation

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

subject to the searcher dynamics $\dot{x}_k(t) = f_k(x_k, u_k), k = 1, \ldots, K$, 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:

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

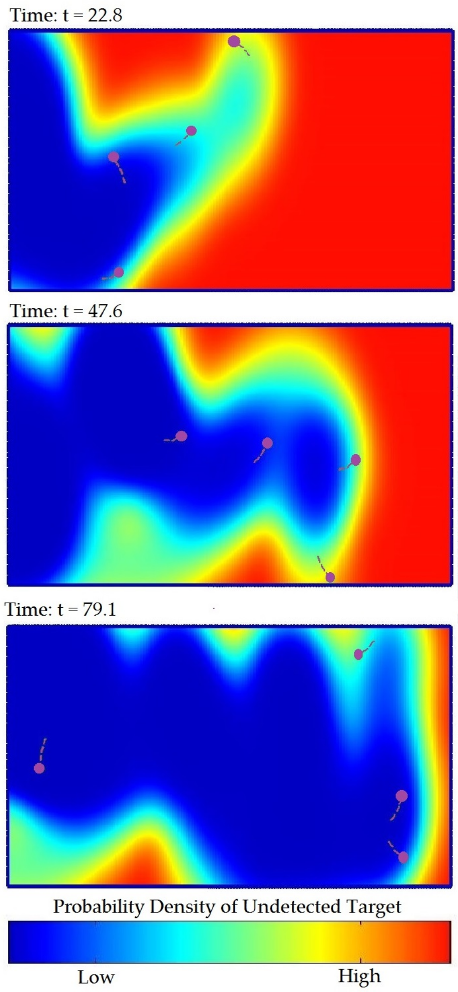

Fig. 2. Snapshots of numerical solution to 'Channel Search Problem' in Sec. 2.2. Colors represent the log probability density value of an undetected target at a point at time t, i.e. the evolution over time of the probability density of target location given parameter value compounded with the probability of target non-detection for given parameter value . 'Low' = $4.14 * 10^{-6}$, 'High' = $2.9317 * 10^{-4}$.

An illustration of a numerical solution to this problem is demonstrated in Figure [2]. 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 [3]. The first com-

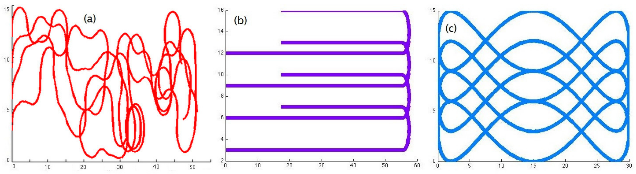

Fig. 3. Searching trajectories. Image (a) illustrates optimal control solutions, image (b) straight line patrols, and image (c) looping patrols.

parison 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

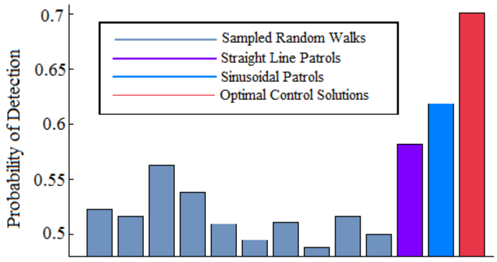

Fig. 4. Comparison of the performance of the optimal control solution vs heuristic methods.

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 distribution over $[-\frac{\pi}{2}, \frac{\pi}{2}]$. Performance results for these trajectories and the optimal control solution are compared in Figure [4]. 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 [Chung et al. (2011)] computation times are referred to in terms of days; in [Foraker (2011)], 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 efficient methods (pseudospectral direct collocation with encoded sparsity and linearity) has reduced these times by an order of magnitude. Representative times are given in the table below, as computed on a 2.3 GHz Intel Core i5 Macbook:

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

3. EXPANDING THE OPTIMAL SEARCH FRAMEWORK

The exponential detection model provides us with a performance criterion which can be used as the objective function in the evaluation of basic optimal search problems. The objective in these problems accounts for one direction of influence–i.e. it accounts for the emission by a searcher of search effort. In order to evaluate a strategic objective of protecting a high-valued asset, heterogeneous influences need to be modeled between multiple agents in a scenario.

To construct a framework for such objectives, we will begin by adopting a more general terminology. Rather than detection, we will use the more general term '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 previous section would be interpreted as the rate of attrition of an attacker's probability of remaining undetected. This is an action directed *against* the attacker *by* the defender.

The concept of attrition is applicable to many relevant strategic gauges. Among other options, it can be used to describe reductions in:

- unmapped information
- survival probabilities
- equipment effectiveness
- the number of agents

We now imagine such a rate of attrition function directed against a target by a searcher and acting as a force of gradual reduction on a quantity $Q(t)$. This quantity could be the target's weapon's effectiveness for instance. As a reference, we will say that this is an action directed by a searcher at position $x(t)$ against a target at position $y(t)$ and that the rate of attrition is given by the function $r_d(x(t), y(t))$. This rate of attrition can be a proportional attrition rate, satisfying:

$$\dot{Q}(t) = -Q(t)r_d(x(t), y(t))$$

or an absolute rate of attrition, satisfying

$$\dot{Q}(t) = -r_d(x(t), y(t)).$$

In both cases, we get solutions of the form:

$$Q(t) = F\left(\int_0^t r_d(x(\tau), y(\tau))d\tau\right).$$

We compose this with a second direction of influence, that of the target on a third agent, a high-valued asset, when the target acts as a force of reduction of the asset's quantity $P(t)$. As an element of effective design, it is assumed that the quantity $Q(t)$ which the searcher is diligently working to reduce is a quantity which has bearing on the target's effectiveness at reducing $P(t)$. That is, if we assume that a target has a baseline rate of attrition against the asset, given by the function $r_a(y(t), x_0(t))$ (where $x_0(t)$ is the trajectory of the asset) , then this baseline rate is then ameliorated by the quantity $Q(t)$. This amelioration yields an effective rate of $\bar{r}_a(Q(t), y(t), x_0(t))$. The quantity $P(t)$ is thus reduced by the equation:

$$\dot{P}(t) = -\bar{r}_a(Q(t), y(t), x_0(t))$$

if the attrition is absolute, or:

$$\dot{P}(t) = -P(t)\bar{r}_a(Q(t), y(t), x_0(t))$$

if it is proportional.

In both cases, a closed form solution is possible due to the hierarchical nature of these relations (a searcher influences a target which in turn influences an asset) which creates a decoupled system of differential equations. Solutions are of the form

$$P(t) = G\left\{\int_0^t \bar{r}_a\left(F\left[\int_0^\tau r_d(x(s),y(s))ds\right], y(t), x_0(t)\right)d\tau\right\}.$$

The existence of this closed form solution enables the application of the methods developed for problem **P**. The following problem will demonstrate how these methods can be utilized to model a multi-agent tactical scenario.

3.1 The 'Kamikaze Shooting Problem'

We consider the scenario where a swarm of L attackers is headed towards a moving high-valued unit (HVU). The trajectory of the HVU is predetermined and given by the function $x_0(t) \in \mathbb{R}^{N_x}$. Attacker trajectories are modeled as conditionally deterministic motion with trajectories given by $y_l(t,\omega)$, for parameters $\omega \in \Omega$ with probability density function $p : \Omega \to \mathbb{R}$ The positions for all L attackers are given by the vector $y = [y_1, \ldots, y_{N_y}]^T \in \mathbb{R}^{N_y}$.

K defenders are dispatched to neutralize the attackers before they destroy the HVU. The objective 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 dynamics: $\dot{x}_k = f_k(x,u)$, $x_k(0) = x_{k,0}$. These quantities are referenced with the vectors $x = [x_1, \ldots, x_{N_x}]^T \in \mathbb{R}^{N_x}$, $u = [u_1, \ldots, u_{N_u}]^T \in \mathbb{R}^{N_u}$, and $f(x,u) = [f_1(x,u), \ldots, f_{N_x}(x,u)] \in \mathbb{R}^{N_x}$.

The following situation is referred to as a 'kamikaze' scenario, because we will for now 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:

$$Q_l(t+\Delta t,\omega) =$$

$$Q_l(t,\omega)\left(1 - \sum_{l=1}^{K} r_{d,k}(x_k(t), y_l(t,\omega))\Delta t\right) + O\left((\Delta t)^2\right).$$

As $\Delta t \to 0$ we find:

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

which 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:

$$P(t,\omega) = e^{-\int_0^t (\sum_{l=1}^{L} Q_l(\tau,\omega) r_{a,l}(y_l(\tau,\omega),x_0(\tau)))d\tau}$$

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 probability density function $p : \Omega \to \mathbb{R}$ and conditionally deterministic attacker trajectories $y(t,\omega)$, determine the control $u : [0,T] \to U \in \mathbb{R}^{n_u}$ that minimizes the expectation

$$J = \int_\Omega (1 - P(T,\omega))\, p(\omega) d\omega$$

where

$$P(T,\omega) = e^{-\int_0^T (\sum_{l=1}^{L} Q_l(\tau,\omega) r_{a,l}(y_l(\tau,\omega),x_0(\tau)))d\tau}$$

and

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

subject to the dynamics $\dot{x}(t) = f(x(t), u(t))$ with initial condition $x(0) = x_0$ and control constraint $g(u(t)) \leq 0$, $\forall t \in [0,T]$.

3.2 Numerical Implmentation

We implement a kamikaze shooting problem with similar physical features to the search problem of Section [2.2]. 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 [5]. This solution was computed using 150 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. The final

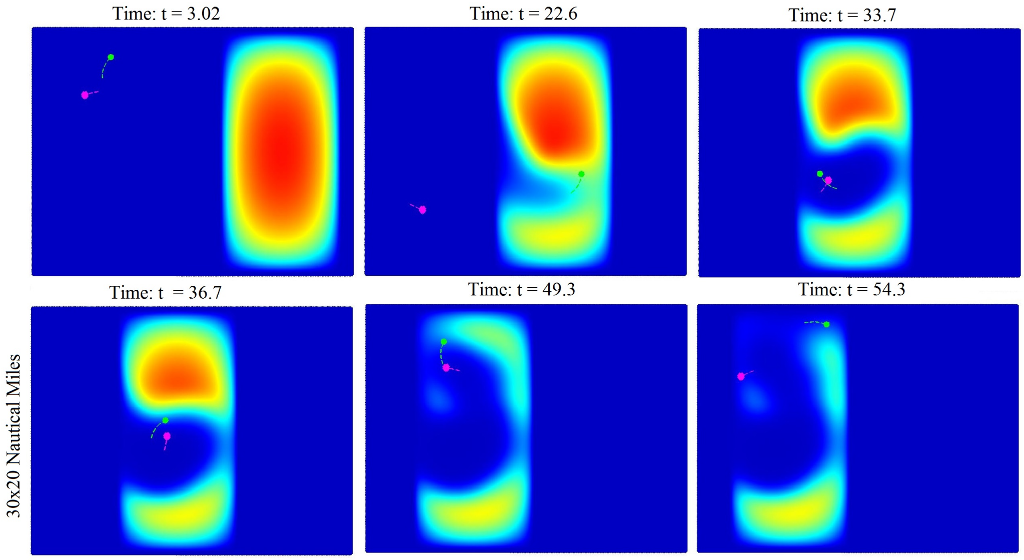

Fig. 5. Snapshots of numerical solution to 'Kamikaze Shooting Problem' in Section 3.2. The magenta icon indicates the position of the HVU and the green icon is defender. Colors represent the log probability density value of a surviving attacker at a point at time t. 'Low' $= 4.14 * 10^{-6}$, 'High' $= 2.9317 * 10^{-4}$.

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–this example serves as a good basis of comparison because it is equivalent in form to the optimal search problem of Section [2]. 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%.

4. CONCLUSIONS

Though solutions to search problems have historically been limited to simple cases, the development of more efficient numerical methods has now made complex problems approachable. Due to these developments we are 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. We have exhibited the effectiveness of this by example–through implementing nonlinear multi-agent search problems and presenting their tractability and functionality. These problems are implemented in seconds, rather than hours or days, and demonstrate a potential for these methods to be applied to realtime situations.

REFERENCES

Chudnovsky, D. and Chudnovsky, G. (1988). *Search Theory: Some Recent Developments*, volume 112 of *Lecture Notes in Pure and Applied Mathematics*. CRC Press.

Chung, H., Polak, E., Royset, J.O., and Sastry, S. (2011). On the optimal detection of an underwater intruder in a channel using unmanned underwater vehicles. *Naval Research Logistics (NRL)*, 58(8), 804–820.

Foraker, J. (2011). *Optimal Search for Moving Targets in Continuous Time and Space Using Consistent Approximations*. Ph.D. thesis, Naval Postgraduate School.

Gill, P.E., Murray, W., and Saunders, M.A. (2005). Snopt: An sqp algorithm for large-scale constrained optimization. *SIAM Review*, 47(1), 99–131.

Gong, Q., Kang, W., and Ross, I.M. (2006). A pseudospectral method for the optimal control of constrained feedback linearizable systems. *IEEE Transactions on Automatic Control*, 51(7), 1115–1129.

Hellman, O. (1972). On the optimal search for a randomly moving target. *SIAM Journal on Applied Mathematics*, 22(4), 545–552.

Kim, K. (2009). *Approximating the Poisson Scan and Lambda-Sigma Acoustic Detection Model with a Random Search Formula*. Master's thesis, Naval Postgraduate School.

Koopman, B. (1946). *Search and Screening*. Naval Historical Center.

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

Ohsumi, A. (1991). Optimal search for a markovian target. *Naval Research Logistics Quarterly*, 38(4), 531–554.

Phelps, C., Gong, Q., Royset, J.O., and Kaminer, I. (2012). Consistent approximation of an optimal search problem. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, 630–637.

Stone, L.D. (1989). What's happened in search theory since the 1975 lanchester prize? *Operations Research*, 37(3), 501–506.

Washburn, A. (1982). On patrolling a channel. *Naval Research Logistics Quarterly*, 29(4), 609–615.