A Stochastic Reach-Avoid Problem with Random Obstacles* †

Sean Summers
Automatic Control Laboratory
ETH Zürich
Physikstrasse 3
8092 Zürich, Switzerland
summers@control.ee.ethz.ch

Maryam Kamgarpour
Department of Mechanical Engineering
UC Berkeley
Berkeley CA 94720, USA
maryamka@eecs.berkeley.edu

Claire Tomlin
Department of Electrical Engineering and Computer Sciences
UC Berkeley
Berkeley CA 94720, USA
tomlin@eecs.berkeley.edu

John Lygeros
Automatic Control Laboratory
ETH Zürich
Physikstrasse 3
8092 Zürich, Switzerland
lygeros@control.ee.ethz.ch

ABSTRACT

We present a dynamic programming based solution to a stochastic reachability problem for a controlled discrete-time stochastic hybrid system. A sum-multiplicative cost function is introduced along with a corresponding dynamic recursion which quantifies the probability of hitting a target set at some point during a finite time horizon, while avoiding an obstacle set during each time step preceding the target hitting time. In contrast with earlier works which consider the reach and avoid sets as both deterministic and time invariant, we consider the avoid set to be both time-varying and probabilistic. Optimal reach-avoid control policies are derived as the solution to an optimal control problem via dynamic programming. A computational example motivated by aircraft motion planning is provided.

Categories and Subject Descriptors

I.6.4 [**Simulation and modeling**]: Model Validation and Analysis

*This work was partially supported by the European Commission under the project iFly, FP6-TREN-037180, and the MoVeS project, FP7-ICT-2009-257005.

†The work of M. Kamgarpour was supported by AFOSR under grant FA9550-06-1-0312 and by National Science and Engineering Research Council of Canada (NSERC).

HSCC'11, April 12–14, 2011, Chicago, Illinois, USA.

General Terms

Theory, Verification

Keywords

reachability, dynamic programming, Markov processes, random sets

1. INTRODUCTION

Reachability analysis of deterministic dynamical systems constitutes a practically important and intensely researched area in control theory. Over the years, methods and numerical tools for reachability of continuous time deterministic systems have been well researched (see [5, 17, 19, 28] and the references therein). In particular, the reachability problems considered are often solved via dynamic programming [17, 20]. Additionally, reachability problems for deterministic hybrid systems and uncertain hybrid systems have been addressed using computational methods based on dynamic programming [26] and nonsmooth analysis [13].

Stochastic Hybrid System (SHS) models have become a common mechanism for the analysis and design of complex systems given their ability to capture the variable temporal and spatial behavior often found in realistic systems. In the continuous time setting, early contributions to SHS theory include the works of [11, 14, 16], with [9] establishing a theoretical foundation for the measurability of events for reachability problems. Given that technical issues such as measurability are easier to resolve in the discrete-time setting, consideration of discrete-time stochastic hybrid systems (DTSHS) [4] has also attracted considerable attention. Based on a theoretical foundation for the solution of stochastic optimal control problems of general discrete-time systems of [7], probabilistic reachability of DTSHS has been addressed in [2, 22, 25].

In this paper we extend the recent results of [2, 25] in the area of probabilistic reachability of DTSHS. We consider a probabilistic reach-avoid problem where the objective is to maximize or minimize the probability that a system starting at a specific initial condition will hit a target set while avoiding an unsafe set over a finite time horizon. However, in contrast with [2,25], we consider that the unsafe set (or obstacle set) in the reach-avoid problem may be time-dependent and random. In particular, that it can be accurately modeled by a time-indexed sequence of random closed sets [18, 21].

Following the methods of [2, 25], we formulate the reach-avoid problem with random obstacles as a finite horizon stochastic optimal control problem with a sum-multiplicative cost-to-go function. Specifically, we consider two distinct possibilities for the random set-valued obstacle process. In the first, we consider the random set process to be an independent stochastic process, and thus decoupled from the evolution of the DTSHS. In the second case, we consider the obstacle process as a set-valued Markov process that can be expressed through an appropriate parameterization. In both cases, dynamic programming is used to compute the Markov control policy that maximizes or minimizes the cost of the optimal control problem. A numerical example motivated by aircraft motion planning under uncertain weather predictions is provided.

The rest of the work is arranged as follows. In Section 2.1 we briefly recall the DTSHS model of [2]. In Section 2.2, we recall the basic theory of random closed sets and define a set-valued stochastic process as a model for the dynamic obstacle set. In Section 3, we introduce the notion of probabilistic reach-avoid over a finite time horizon with dynamic and stochastic obstacle sets, and develop a mechanism to determine optimal Markov control policies based on dynamic programming. In Section 4 we provide a numerical example.

2. MATHEMATICAL BACKGROUND

Here we recall the DTSHS model and associated semantics introduced in [2] and aspects of the theory of random closed sets [18, 21].

2.1 DTSHS

Throughout, given a Borel set K, $\mathcal{B}(K)$ denotes the Borel $\sigma-$algebra of K.

Definition 1. *A discrete-time stochastic hybrid system,* $\mathcal{H} = (\mathcal{Q}, n, \mathcal{A}, T_v, T_q, R)$*, comprises*

- *A discrete state space* $\mathcal{Q} := \{q_1, q_2, ..., q_m\}$*, for some* $m \in \mathbb{N}$*;*
- *A map* $n : \mathcal{Q} \to \mathbb{N}$ *which assigns to each discrete state value* $q \in \mathcal{Q}$ *the dimension of the continuous state space* $\mathbb{R}^{n(q)}$*. The hybrid state space is then given by* $X := \bigcup_{q\in\mathcal{Q}}\{q\} \times \mathbb{R}^{n(q)}$*;*
- *A compact Borel space* $\mathcal{A}$ *representing the control space;*
- *A Borel-measurable stochastic kernel on* $\mathbb{R}^{n(\cdot)}$ *given* $X \times \mathcal{A}$*,* $T_v : \mathcal{B}(\mathbb{R}^{n(\cdot)}) \times X \times \mathcal{A} \to [0, 1]$*, which assigns to each* $x = (q, v) \in X$ *and* $a \in \mathcal{A}$ *a probability measure* $T_v(\cdot|x, a)$ *on the Borel space* $(\mathbb{R}^{n(q)}, \mathcal{B}(\mathbb{R}^{n(q)}))$*;*
- *A discrete stochastic kernel on* $\mathcal{Q}$ *given* $X \times \mathcal{A}$*,* $T_q : \mathcal{Q} \times X \times \mathcal{A} \to [0, 1]$*, which assigns to each* $x \in X$ *and* $a \in \mathcal{A}$ *a probability distribution* $T_q(\cdot|x, a)$ *over* $\mathcal{Q}$*;*
- *A Borel-measurable stochastic kernel on* $\mathbb{R}^{n(\cdot)}$ *given* $X \times \mathcal{A} \times \mathcal{Q}$*,* $R : \mathcal{B}(\mathbb{R}^{n(\cdot)}) \times X \times \mathcal{A} \times \mathcal{Q} \to [0, 1]$*, which assigns to each* $x \in X$*,* $a \in \mathcal{A}$*, and* $q' \in \mathcal{Q}$ *a probability measure* $R(\cdot|x, a, q')$ *on the Borel space* $(\mathbb{R}^{n(q')}, \mathcal{B}(\mathbb{R}^{n(q')}))$*.*

Consider the DTSHS evolving over the finite time horizon $k = 0, 1, \dots, N$ with $N \in \mathbb{N}$. We specify the initial state as $x_0 \in X$ at time $k = 0$, and define the notion of a Markov policy.

Definition 2. *A Markov Policy for a DTSHS,* $\mathcal{H}$*, is a sequence* $\mu = (\mu_0, \mu_1, ..., \mu_{N-1})$ *of universally measurable maps* $\mu_k : X \to \mathcal{A}$*,* $k = 0, 1, ..., N - 1$*. The set of all admissible Markov policies is denoted by* $\mathcal{M}_m$*.*

Let $\tau_v : \mathcal{B}(\mathbb{R}^{n(\cdot)}) \times X \times \mathcal{A} \times \mathcal{Q} \to [0, 1]$ be a stochastic kernel on $\mathbb{R}^{n(\cdot)}$ given $X \times \mathcal{A} \times \mathcal{Q}$, which assigns to each $x \in X$, $a \in \mathcal{A}$, and $q' \in \mathcal{Q}$, a probability measure on the Borel space $(\mathbb{R}^{n(q')}, \mathcal{B}(\mathbb{R}^{n(q')}))$ given by

$$\tau_v(dv'|(q, v), a, q') = \begin{cases} T_v(dv'|(q, v), a), & \text{if } q' = q \\ R(dv'|(q, v), a, q'), & \text{if } q' \neq q. \end{cases}$$

Based on τ_v we introduce the kernel $Q : \mathcal{B}(X) \times X \times \mathcal{A} \to [0, 1]$:

$$Q(dx'|x, a) = \tau_v(dv'|x, a, q')T_q(q'|x, a).$$

Definition 3. *Consider the DTSHS,* $\mathcal{H}$*, and time horizon* $N \in \mathbb{N}$*. A stochastic process* $\{x_k,\ k = 0, ..., N\}$ *with values in* X *is an execution of* $\mathcal{H}$ *associated with a Markov policy* $\mu \in \mathcal{M}_m$ *and an initial condition* $x_0 \in X$ *if and only if its sample paths are obtained according to the DTSHS Algorithm.*

Algorithm 1 DTSHS Algorithm

Require: Sample Path $\{x_k, k = 0, ..., N\}$
Ensure: Initial hybrid state $x_0 \in X$ at time $k = 0$, and Markov control policy $\mu = (\mu_0, \mu_1, ..., \mu_{N-1}) \in \mathcal{M}_m$

1: **while** $k < N$ **do**
2: Set $a_k = \mu_k(x_k)$
3: Extract from X a value x_{k+1} according to $Q(\cdot|x_k, a_k)$
4: Increment k
5: **end while**

Equivalently, the DTSHS $\mathcal{H}$ can be described as a Markov control process with state space X, control space $\mathcal{A}$, and controlled transition probability function Q. Further, given a specific control policy $\mu \in \mathcal{M}_m$ and initial state $x_0 \in X$, the execution $\{x_k, k = 0, ..., N\}$ is a time inhomogeneous stochastic process defined on the canonical sample space $\Omega = X^{N+1}$, endowed with its product $\sigma-$algebra $\mathcal{B}(\Omega)$. The probability measure $P^{\mu}_{x_0}$ is uniquely defined by the transition kernel Q, the Markov policy $\mu \in \mathcal{M}_m$, and the initial condition $x_0 \in X$ (see [7]). From now on, we will use interchangeably the notation $Q(\cdot|x, \mu_k(x))$ and $Q^{\mu_k}_x(\cdot|x)$ to represent the one-step transition kernel.

2.2 Random Sets

For the hybrid state space X one can select a metric d such that (X, d) becomes a complete separable metric space (see e.g. [11]). Let $\mathcal{K}$ denote the set of all closed subsets of the hybrid state space X and let d_H denote the Hausdorff

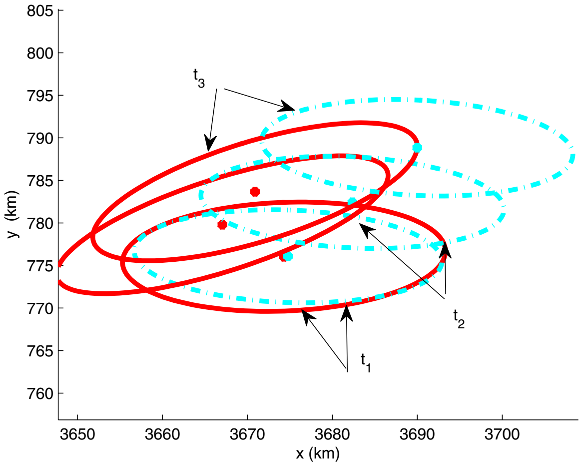

Figure 1: A trajectory of forecasted and realized weather obstacles over a 15 minute horizon according to [12] for a section of airspace centered at latitude 30^o and longitude 86.5^o, near the gulf coast of Florida, on 01/07/2009. The forecast storms are shown in dashed lines. The horizon is shown with the labels $t_1 = 5$ minutes, $t_2 = 10$ minutes, and $t_3 = 15$ minutes.

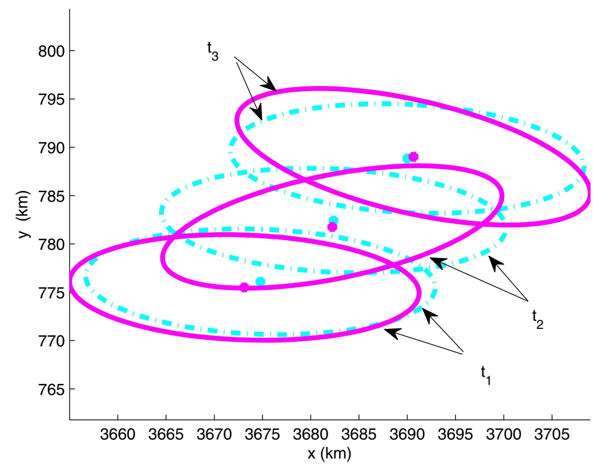

Figure 2: A trajectory of the random obstacles based on forecast data is shown. The forecasted ellipses are represented by dashed lines and the realization of the random ellipses is represented as solid lines.

metric. It follows that $(\mathcal{K}, d_H)$ is also a complete separable metric space where the corresponding open subsets generate a σ-algebra on $\mathcal{K}$ [21], i.e. the Borel σ-algebra $\mathcal{B}(\mathcal{K})$ corresponding to the Hausdorff metric of $\mathcal{K}$.

Definition 4. *A random closed set is a measurable function Ξ from a probability space $(\Omega, \mathcal{F}, P)$ into the measure space $(\mathcal{K}, \mathcal{B}(\mathcal{K}))$.*

The distribution of a random closed set Ξ, is given by the probabilities

$$P\{\Xi \cap F \neq \emptyset\}$$

for $F \in \mathcal{K}$. For $F = \{x\} \in X$, the probability $P\{x \in \Xi\} = P\{\omega \in \Omega : x \in \Xi(\omega)\}$ is obtained which satisfies the expression

$$P\{x \in \Xi\} = 1 - P\{x \notin \Xi\}.$$

We refer to the function $p_\Xi(x) = P\{x \in \Xi\}$ as the covering function. For some set $K \subseteq X$, let $\mathbf{1}_K(\cdot) : X \to \{0,1\}$ denote the indicator function. The covering function can also be interpreted as the mean of the indicator function $\mathbf{1}_\Xi$, i.e.

$$p_\Xi(x) = E[\mathbf{1}_\Xi(x)].$$

Note that the covering function is a universally measurable function [18] and takes values between 0 and 1.

We now define a stochastic set-valued process to be used as a model for obstacle movement. For $k = 0, 1, 2, \ldots, N$, let G_k be a Borel-measurable stochastic kernel on $\mathcal{K}$ given $\mathcal{K}$, $G_k : \mathcal{B}(\mathcal{K}) \times \mathcal{K} \to [0,1]$, which assigns to each $K \in \mathcal{K}$ a probability measure $G_k(\cdot|K)$ on the Borel space $(\mathcal{K}, \mathcal{B}(\mathcal{K}))$. That is, let G_k represent a collection of probability measures on $(\mathcal{K}, \mathcal{B}(\mathcal{K}))$ parametrized by the elements of $\mathcal{K}$ and indexed by time k. A discrete-time time-inhomogeneous set-valued Markov process $\mathbf{\Xi} = (\Xi_k)_{k\in\mathbb{N}_0}$ taking values in the Borel space $\mathcal{K}$ is described by the stochastic kernel G_k. This general model for the set-valued Markov process includes as special cases time-homogeneous set-valued Markov processes and independent distributions of random sets taking values according to time-indexed random closed set stochastic kernels.

In most cases the characterization of a stochastic set valued process and the computation of associated functions (e.g. the covering function) is difficult due to the size of $\mathcal{K}$. Yet, methods have been suggested in the literature that alleviate the complexity of these processes and the related functions [10, 24]. For example, random closed set processes are often characterized by families of closed subsets of $\mathcal{K}$ which are parametrized by real parameters (referred to as morphological transformations in [18]).

Definition 5. *A parameterization of a discrete-time set-valued Markov process $\mathbf{\Xi}$ is a discrete-time Markov process $\xi = (\xi_k)_{k\in\mathbb{N}_0}$ with parameter space $\mathcal{O}$ and transition probability function $T_k : \mathcal{B}(\mathcal{O}) \times \mathcal{O} \to [0,1]$ together with a function $\gamma : \mathcal{O} \to \mathcal{K}$ such that*

$$\mathbf{\Xi} = (\Xi_k)_{k\in\mathbb{N}_0} = (\gamma(\xi_k))_{k\in\mathbb{N}_0}.$$

From now on we restrict our attention to parameterized discrete-time set-valued Markov processes. Analysis of the associated functions is often completed via Monte Carlo methods. Consider the following example.

2.3 Example - Vertically Integrated Liquid

In aircraft path planning, the ability to identify and characterize regions of hazardous weather is vitally important. One factor that can be used to determine the safety of a region of the airspace for an aircraft to fly through is the Vertically Integrated Liquid (VIL) water content measurement [3], which represents the level of precipitation in a column of the airspace. This measurement has proven useful in the detection of severe storms and short-term rainfall forecasting [8], and hence can be used as an indicator for establishing a no-fly zone for aircraft: A region of the airspace

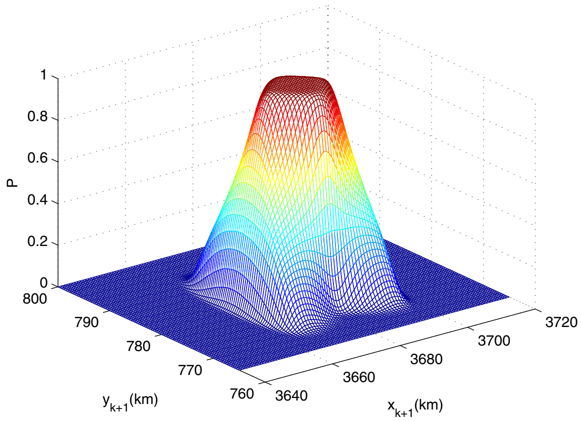

Figure 3: An example covering function for the set-valued Markov process of Section 2.2.

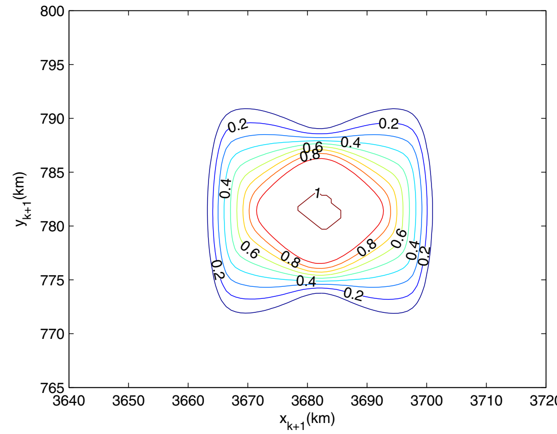

Figure 4: Contour plot of the covering function shown in Figure 3.

with a VIL measurement level above a certain threshold is recommended as a no-fly zone for aircraft.

Regions with VIL levels above the safety threshold come in numerous shapes and sizes. To simplify the expression of these regions, one method of representation is to enclose the no-fly region by minimum volume ellipsoids. Simplifying the problem further, we consider a constant flight level for the aircraft and consequently focus on ellipses in two dimensions. Consider the forecast data at time instance k, from which we obtain elliptical obstacles $\{\mathcal{E}_k^l\}$, where $l = 1, 2, \ldots, L_k$, and L_k denotes the number of obstacles at discrete time k. Each ellipse $\mathcal{E}_k^l$ is parameterized by its center $m_k^l \in \mathbb{R}^2$ and its positive definite eccentricity matrix M_k^l:

$$x \in \mathcal{E}_k^l(m_k^l, M_k^l) \iff (x - m_k^l)^T M_k^l (x - m_k^l) \leq 1 \qquad (6)$$

For simplicity, in the sequel we assume that there exists $l = 1$ elliptical no-fly region and denote the ellipse and its representative parameters at time k by $\mathcal{E}_k$, m_k, and M_k.

Uncertainties associated with forecast data naturally exist, and become more prominent with the horizon length of the forecast. For example, Figure 1 shows forecast storm regions versus actual storm regions over a 15 minute horizon at samples of 5 minutes. Figure 1 clearly illustrates the effect that uncertainties can have on aircraft path planning. We account for these uncertainties by modeling the hazardous regions (i.e. obstacles) as a sequence of random closed sets whose distribution is determined from forecast data and associated statistics. In this formulation, the subsets of the airspace that are unsafe to fly through become probabilistic.

As stated in Section 2.2, mathematically characterizing the uncertainty in the forecast, and hence modeling the hazardous regions as random closed sets, is difficult in general. However, given the ellipse parameterization introduced in (6), we can model the dynamic obstacles as random closed sets by considering the parameters of the ellipse representation (e.g. the ellipse center) as random variables distributed according to the forecast data. Assuming that the forecast data is available in 5 minute increments and is given in terms of the expected ellipse center m_k and eccentricity matrix M_k, we consider two mathematical formulations for the model of the random sets. In the first, we consider that the evolution of the ellipse centers is Markovian, and hence the obstacle process is a set-valued Markov process (i.e. the distribution of $\mathcal{E}_{k+1}$ is a function of $\mathcal{E}_k$). In the second, we consider the distribution of the obstacle at time $k+1$ to be independent of $\mathcal{E}_k$. In the current section we consider the Markov model for the set-valued process, although both special cases will be addressed in Sections 3 and 4.

Consider a storm region characterized by its minimum-volume ellipse and let m_k and m_{k+1} be the forecast center of the ellipse at time k and time $k+1$ respectively. Then the term $\mu_k \in \mathbb{R}^2$ defined as $\mu_k = m_{k+1} - m_k$, for $k \in \mathbb{N}$ denotes the incremental motion of the storm. Given that there are uncertainties in the forecast, we assume that the true center is a random variable whose motion is described as

$$\xi_{k+1} = \xi_k + \mu_k + \eta_k, \qquad (7)$$

where $\eta_k \sim \mathcal{N}(0, \Sigma)$. In addition, we approximate the variation in the ellipse eccentricity according to the expression

$$C_k = R(\theta_k)^T M_k R(\theta_k), \qquad (8)$$

where M_k is the eccentricity obtained from the forecast at time k, $R(\cdot)$ is a rotation matrix, and the angle of rotation θ_k at time k is a random variable with uniform distribution over an interval $[-\alpha, \alpha]$. Note that in this formulation only the ellipse centers are Markov, although one could also consider the angle of rotation to be Markov as well. The noise parameters Σ and α are best determined from the quality of the forecast and the rate of movement of the storms. A trajectory of the forecast ellipses and a realization of the random ellipses over a horizon of 15 minutes is shown in Figure 2. Based on the analysis of the storm movements [12], the noise parameters were set to $\Sigma = I_{2\times 2}$ and $\alpha = \frac{\pi}{6}$.

Given that the objective is to evaluate the safety of an aircraft path with respect to hazardous weather, the covering function for the random closed sets is of immediate interest. That is, for an aircraft position $x_k \in \mathbb{R}^2$ at some time k, the probability of being in the hazardous region $\Xi_k = \mathcal{E}_k(\xi_k, C_k)$ is given by

$$P\{x_k \in \Xi_k\} = P\{(x_k - \xi_k)^T C_k (x_k - \xi_k) \leq 1\}.$$

Unfortunately, the calculation of the above probability is not analytically possible. For the case in which the eccentricity of the ellipse is assumed to be deterministic, the above probability obeys a Chi-squared distribution and can be approximated using statistical computational tools. For the more general case, one can use Monte Carlo simulations to approximate the probabilities of hitting an obstacle. For example, consider the Markov set-valued obstacle process introduced above with $\xi_k = [3675 \;\; 775]^T$, $\mu_k = [7.1 \;\; 6.4]^T$, and $M_{k+1} = [0.0028 \;\, 0; \;\; 0 \;\, 0.0278]$. The covering function $p_{\Xi_{k+1}}(x_{k+1})$ at time $k+1$ is approximated over the region $[3640 \;\; 3720] \times [765 \;\; 800]$ using 10^5 Monte Carlo samples and a 101×101 grid discretization. The result is displayed in Figure 3 and Figure 4.

3. FINITE HORIZON REACH-AVOID

Let $K, K'_k \in \mathcal{B}(X)$, with $K \subseteq K'_k$ for all $k = 0, 1, \ldots, N$. We define the stopping time associated with hitting K as $\tau := \inf\{j \geq 0 | x_j \in K\}$, and the stopping time associated with hitting $X \setminus K'_k$ as $\tau' := \inf\{j \geq 0 | x_j \in X \setminus K'_j\}$; if a set is empty we set its infimum equal to $+\infty$. Our goal is to evaluate the probability that the execution of the Markov control process associated with the Markov policy $\mu \in \mathcal{M}_m$ and the initial condition x_0 will hit K before hitting $X \setminus K'_k$ during the time horizon N. We assume that the initial avoid set $\Xi_0 = X \setminus K'_0$ is known and $\Xi_k = X \setminus K'_k$ for $k = 1, \ldots, N$ is an execution of the stochastic set-valued process G. The probability that the system initialized at $x_0 \in X$, with control policy $\mu \in \mathcal{M}_m$ and initial avoid set $\Xi_0 \in \mathcal{K}$, reaches K while avoiding $X \setminus K'_k$ for all $k = 0, 1, \ldots, N$ is given by

$$\begin{aligned} r^{\mu}_{(x_0,\Xi_0)}(K) \;\; &:= \;\; P^{\mu}_{(x_0,\Xi_0)}\{\exists j \in [0, N] : x_j \in K \,\wedge \\ &\qquad\qquad \forall i \in [0, j-1] \;\; x_i \in K'_i \setminus K\}, \\ &= \;\; P^{\mu}_{(x_0,\Xi_0)}\{\{\tau < \tau'\} \wedge \{\tau \leq N\}\}, \end{aligned}$$

where $\wedge$ denotes the logical AND, and we operate under the assumption that the requirement on i is automatically satisfied when $x_0 \in K$; subsequently we will use a similar convention for products, i.e. $\prod_{i=k}^{j}(\cdot) = 1$ if $k > j$.

As in [25], consider

$$\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}\mathbf{1}_{K'_i\setminus K}(x_i)\right)\mathbf{1}_K(x_j) = \left\{\begin{array}{ll} 1, & \text{if } \exists j \in [0, N] : x_j \in K \wedge \\ & \quad \forall i \in [0, j-1] \;\; x_i \in K'_i \setminus K \\ 0, & \text{otherwise.} \end{array}\right.$$

Hence $r^{\mu}_{(x_0,\Xi_0)}(K)$ can be expressed as the expectation

$$r^{\mu}_{(x_0,\Xi_0)}(K) \;\; = \;\; E^{\mu}_{(x_0,\Xi_0)}\left[\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}\mathbf{1}_{K'_i\setminus K}(x_i)\right)\mathbf{1}_K(x_j)\right].$$

Analytical (and computational) evaluation of $r^{\mu}_{(x_0,\Xi_0)}(K)$ can be separated into two distinct classes of problems. In the first class of problems we consider the dynamics of the DTSHS and the set-valued obstacle process to be decoupled. We assume that the set-valued obstacle process is described (or can be fairly approximated) by a time-indexed independent distribution of random sets. Further, we assume the control actions (optimal or otherwise) do not depend on the state of the set-valued obstacle process at each step in time. It follows (in Section 3.1) that a property of this model class is the ability to separate the computational burden of the obstacle process from the computational burden of the DTSHS.

In the second class of problems, we consider the dynamics of the DTSHS and the set-valued obstacle process to be coupled. We assume that the set-valued obstacle process is modeled as a set-valued Markov process and that the control policy can depend on both the state of the DTSHS and the state of the obstacle process. The second class of systems subsumes the first class of systems, hence a more general set of models and problems is considered. However, the approach for the second class of systems (introduced in Section 3.2) requires that the state space of the Markov process be augmented, consequently restricting the size of problems that can be approximated numerically due to the Curse of Dimensionality [6]. As a result, it is sometimes necessary to approximate a coupled system with a decoupled system (in both the model dynamics and the space of control policies). While computationally prudent, this approximation will lead to inferior success rates since the available control policy cannot make use of the state of the obstacle at each time step.

In Sections 3.1 and 3.2 we assume that the obstacle process can be characterized according to Definition 5. In Section 3.1 this results in the ability to computationally evaluate the covering functions via Monte Carlo analysis. In Section 3.2 this results in the ability to augment the state space of the DTSHS with the parameters of the obstacle process.

3.1 Decoupled Markov Process

Assume that the set-valued obstacle process is described (or can be fairly approximated) by a time-indexed independent distribution of random sets. It follows that the product measure of the obstacle process is equal to (or well approximated by) the product measure of time-indexed independent stochastic kernels, i.e. for $N \in \mathbb{N}$

$$\prod_{j=0}^{N-1} G_j(d\Xi_j|\Xi_{j-1}) \approx \prod_{j=0}^{N-1} G_j(d\Xi_j).$$

Note that since the initial state of the obstacle Ξ_0 is assumed known, we define $G_0(d\Xi_0|\Xi_{0-1}) = G_0(d\Xi_0) = \delta_{\Xi_0}(d\Xi_0)$.

For a DTSHS with independent set-valued obstacle processes, it can be shown that

$$\begin{aligned} r^{\mu}_{(x_0,\Xi_0)}(K) \;\; &= \;\; E^{\mu}_{(x_0,\Xi_0)}\left[\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}\mathbf{1}_{K'_i\setminus K}(x_i)\right)\mathbf{1}_K(x_j)\right], \\ &= \;\; E^{\mu}_{x_0}\left[\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}p_{K'_i\setminus K}(x_i)\right)\mathbf{1}_K(x_j)\right], \end{aligned}$$

where the covering function notation is used liberally for the simplifying expression $p_{K'_i\setminus K}(x_i) = \mathbf{1}_{X\setminus K}(x_i) - p_{\Xi_i}(x_i)$ (since $K'_i\setminus K$ is not necessarily a random closed set). Clearly, the covering functions are defined

$$p_{\Xi_i}(x_i) = E\left[\mathbf{1}_{\Xi_i}(x_i)\right] = \int_{\mathcal{K}} \mathbf{1}_{\Xi_i}(x_i) G_i(d\Xi_i).$$

A proof (by Fubini's Theorem [23]) of the preceding claim follows:

$$
\begin{aligned}
r^{\mu}_{(x_0,\Xi_0)}(K) &= E^{\mu}_{(x_0,\Xi_0)}\left[\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}\mathbf{1}_{K_i'\setminus K}(x_i)\right)\mathbf{1}_K(x_j)\right] \\
&= \int_{X^N\times\mathcal{K}^N}\left[\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}(\mathbf{1}_{X\setminus K}(x_i)-\mathbf{1}_{\Xi_i}(x_i))\right)\right. \\
&\qquad \left.\mathbf{1}_K(x_j)\right]\prod_{j=0}^{N-1}Q^{\mu_j}(dx_{j+1}|x_j)G_j(d\Xi_j) \\
&= \int_{X^N}\int_{\mathcal{K}^N}\left[\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}(\mathbf{1}_{X\setminus K}(x_i)-\mathbf{1}_{\Xi_i}(x_i))\right)\right. \\
&\qquad \left.\mathbf{1}_K(x_j)\right]\prod_{j=0}^{N-1}G_j(d\Xi_j)\prod_{j=0}^{N-1}Q^{\mu_j}(dx_{j+1}|x_j) \\
&= \int_{X^N}\left[\sum_{j=0}^{N}\left(\int_{\mathcal{K}^j}\prod_{i=0}^{j-1}(\mathbf{1}_{X\setminus K}(x_i)-\mathbf{1}_{\Xi_i}(x_i))\right.\right. \\
&\qquad \left.\left.\prod_{i=0}^{j-1}G_i(d\Xi_i)\right)\mathbf{1}_K(x_j)\right]\prod_{j=0}^{N-1}Q^{\mu_j}(dx_{j+1}|x_j) \\
&= \int_{X^N}\left[\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}(\mathbf{1}_{X\setminus K}(x_i)-\right.\right. \\
&\qquad \left.\left.\int_{\mathcal{K}}\mathbf{1}_{\Xi_i}(x_i)G_i(d\Xi_i))\right)\mathbf{1}_K(x_j)\right]\prod_{j=0}^{N-1}Q^{\mu_j}(dx_{j+1}|x_j) \\
&= \int_{X^N}\left[\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}p_{K_i'\setminus K}(x_i)\right)\mathbf{1}_K(x_j)\right] \\
&\qquad \prod_{j=0}^{N-1}Q^{\mu_j}(dx_{j+1}|x_j) \\
&= E^{\mu}_{x_0}\left[\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}p_{K_i'\setminus K}(x_i)\right)\mathbf{1}_K(x_j)\right].
\end{aligned}
$$

For a fixed Markov policy $\mu \in \mathcal{M}_m$, let us define the functions $V_k^\mu : X \to [0,1]$, $k = 0,\dots,N$ as

$$
\begin{aligned}
V_N^\mu(x) &= \mathbf{1}_K(x)\ , \qquad (9)\\
V_k^\mu(x) &= \mathbf{1}_K(x) + \\
&\quad p_{K_k'\setminus K}(x)\int_{X^{N-k}}\sum_{j=k+1}^{N}\left(\prod_{i=k+1}^{j-1}p_{K_i'\setminus K}(x_i)\right) \\
&\quad \mathbf{1}_K(x_j)\prod_{j=k+1}^{N-1}Q^{\mu_j}(dx_{j+1}|x_j)Q^{\mu_k}(dx_{k+1}|x). \qquad (10)
\end{aligned}
$$

Note that

$$
\begin{aligned}
V_0^\mu(x_0) &= E^{\mu}_{x_0}\left[\sum_{j=0}^{N}\left(\prod_{i=0}^{j-1}p_{K_i'\setminus K}(x_i)\right)\mathbf{1}_K(x_j)\right] \\
&= r^{\mu}_{(x_0,\Xi_0)}(K).
\end{aligned}
$$

Let $\mathcal{F}$ denote the set of functions from X to $\mathbb{R}$ and define the operator $H : X \times \mathcal{A} \times \mathcal{F} \to \mathbb{R}$ as

$$H(x,a,Z) := \int_X Z(y)Q(dy|x,a). \qquad (11)$$

The following lemma shows that $r^{\mu}_{(x_0,\Xi_0)}(K)$ can be computed via a backwards recursion.

Lemma 12. *Fix a Markov policy $\mu = (\mu_0, \mu_1, \dots \mu_{N-1}) \in \mathcal{M}_m$. The functions $V_k^\mu : X \to [0,1]$, $k = 0,1,\dots,N-1$ can be computed by the backward recursion:*

$$V_k^\mu(x) = \mathbf{1}_K(x) + p_{K_k'\setminus K}(x)H(x,\mu_k(x),V_{k+1}^\mu), \qquad (13)$$

initialized with $V_N^\mu(x) = \mathbf{1}_K(x)$, $x \in X$.

proof 14. *By induction. First, due to the definition of (9) and (10), we have that*

$$
\begin{aligned}
V_{N-1}^\mu(x) &= \mathbf{1}_K(x) + \\
&\quad p_{K_{N-1}'\setminus K}(x)\int_X V_N^\mu(x_N)Q^{\mu_{N-1}}(dx_N|x),
\end{aligned}
$$

so that (13) is proven for $k = N-1$. For $k < N-1$ we can separate the terms associated with x_{k+1} as follows

$$
\begin{aligned}
V_k^\mu(x) &= \mathbf{1}_K(x) + \\
&\quad p_{K_k'\setminus K}(x)\int_X\Bigg(\mathbf{1}_K(x_{k+1}) + p_{K_{k+1}'\setminus K}(x_{k+1}) \\
&\quad \int_{X^{N-k-1}}\sum_{j=k+2}^{N}\left(\prod_{i=k+2}^{j-1}p_{K_i'\setminus K}(x_i)\right)\mathbf{1}_K(x_j) \\
&\quad \prod_{j=k+2}^{N-1}Q^{\mu_j}(dx_{j+1}|x_j)Q^{\mu_{k+1}}(dx_{k+2}|x_{k+1})\Bigg) \\
&\quad Q^{\mu_k}(dx_{k+1}|x) \\
&= \mathbf{1}_K(x) + \\
&\quad p_{K_k'\setminus K}(x)\int_X V_{k+1}^\mu(x_{k+1})Q^{\mu_k}(dx_{k+1}|x)
\end{aligned}
$$

which concludes the proof. □

Definition 15. *Let $\mathcal{H}$ be a Markov control process, $\mathbf{\Xi} = (\Xi_k)_{k\in\mathbb{N}_0}$ a random closed set stochastic process, $K \in \mathcal{B}(X)$, $K_k' \in \mathcal{B}(X)$, with $K \subseteq K_k'$ and $K_k' = X \setminus \Xi_k$ for all $k = 0,1,2,\dots,N$. A Markov policy μ^* is a maximal reach-avoid policy if and only if $r^{\mu^*}_{(x_0,\Xi_0)}(K) = \sup_{\mu\in\mathcal{M}_m} r^{\mu}_{(x_0,\Xi_0)}(K)$, for all $x_0 \in X$.*

Theorem 16. *Define $V_k^* : X \to [0,1]$, $k = 0,1,\dots,N$, by the backward recursion:*

$$V_k^*(x) = \sup_{a\in\mathcal{A}}\{\mathbf{1}_K(x) + p_{K_k'\setminus K}(x)H(x,a,V_{k+1}^*)\} \qquad (17)$$

initialized with $V_N^*(x) = \mathbf{1}_K(x)$, $x \in X$. Then, $V_0^*(x_0) = \sup_{\mu\in\mathcal{M}_m} r^{\mu}_{(x_0,\Xi_0)}(K)$, $x_0 \in X$ and $\Xi_0 \in \mathcal{K}$. If $\mu_k^* : X \to \mathcal{A}$, $k \in [0,N-1]$, is such that for all $x \in X$

$$\mu_k^*(x) = \arg\sup_{a\in\mathcal{A}}\{\mathbf{1}_K(x) + p_{K_k'\setminus K}(x)H(x,a,V_{k+1}^*)\} \qquad (18)$$

then $\mu^* = (\mu_0^*, \mu_1^*, \dots, \mu_{N-1}^*)$ is a maximal reach-avoid policy. A sufficient condition for the existence of μ^* is that $U_k(x,\lambda) = \{a \in \mathcal{A} | H(x,a,V_{k+1}^*) \geq \lambda\}$ is compact for all $x \in X$, $\lambda \in \mathbb{R}$, $k \in [0,N-1]$.

PROOF 19. *For all $k = 0, 1, ..., N$, the covering function p_{Ξ_k}, and therefore $p_{K'_k \setminus K}$, is a universally measurable function [18]. Hence, we apply the proof of Theorem 6 in [25] with $p_{K'_k \setminus K}$ replacing $\mathbf{1}_{K' \setminus K}$ everywhere.* □

Note that Theorem 16 gives a sufficient condition for the existence of an optimal nonrandomized Markov policy. While the consideration of randomized Markov policies is indeed interesting in the event that an optimal nonrandomized Markov policy does not exist, in most cases the "best" policy can be taken to be nonrandomized [7]. In light of this fact and in the interest of space, we do not consider randomized Markov policies in the present work and urge the interested reader to consider Chapter 8 in [7] for additional details.

3.2 Coupled Markov Process

Assume that the sequence of obstacles is modeled as a set-valued Markov process. It follows that the product measure of the obstacle process is equal to the product measure of the stochastic kernel G, i.e. for $N \in \mathbb{N}$ the product measure is

$$\prod_{j=0}^{N-1} G(d\Xi_j|\Xi_{j-1}).$$

Note that since the initial state of the obstacle Ξ_0 is assumed known, we define $G(d\Xi_0|\Xi_{0-1}) = \delta_{\Xi_0}(d\Xi_0)$.

By Definition 5, we have that an equivalent characterization of the set-valued Markov process $\mathbf{\Xi}$ with transition kernel G is given by the discrete-time Markov process $\xi = (\xi_k)_{k\in\mathbb{N}_0}$ with parameter space $\mathcal{O}$ and transition probability function T along with the function γ. Let $\bar{x} \in \bar{X}$ be the augmented state of the DTSHS, where $\bar{x} = \begin{bmatrix} x^T, & \xi^T \end{bmatrix}^T$ and $\bar{X} = X \times \mathcal{O}$ is the augmented state space of the DTSHS. Further, let us define the stochastic kernel $\bar{Q} : \mathcal{B}(\bar{X}) \times \bar{X} \times \mathcal{A} \to [0, 1]$:

$$\bar{Q}(d\bar{x}'|\bar{x}, a) = Q(dx'|x, a)T(d\xi'|\xi).$$

We call the resulting process an augmented DTSHS (ADTSHS) $\bar{\mathcal{H}}$.

DEFINITION 20. *A Markov Policy for an augmented DTSHS, $\bar{\mathcal{H}}$, is a sequence $\mu = (\mu_0, \mu_1, ..., \mu_{N-1})$ of universally measurable maps $\mu_k : \bar{X} \to \mathcal{A}$, $k = 0, 1, ..., N-1$. The set of all admissible Markov policies is denoted by $\bar{\mathcal{M}}_m$.*

Hence, the stochastic reach-avoid problem with time-varying probabilistic obstacles is transformed into a stochastic reach-avoid problem with deterministic obstacles, and thus can be solved using the nominal reach-avoid methods in [25].

4. AIRCRAFT PATH PLANNING

Here we consider a discrete-time stochastic hybrid model of aircraft motion inspired by the work [27]. We model the aircraft motion as a simple point mass unicycle with three modes of operation; straight flight, right turn, and left turn. The discrete-time continuous dynamics of the aircraft are given by

$$\begin{aligned} x^1_{k+1} &= x^1_k + t_e a^1_k \cos(x^3_k) + w^1_k \quad (21) \\ x^2_{k+1} &= x^2_k + t_e a^1_k \sin(x^3_k) + w^2_k \\ x^3_{k+1} &= x^3_k + t_e a^2_k + w^3_k \end{aligned}$$

where $x = \begin{bmatrix} x^1, & x^2, & x^3 \end{bmatrix}^T \in \mathbb{R}^3$ are the states of the system, $a = \begin{bmatrix} a^1, & a^2 \end{bmatrix}^T \in \mathcal{A}$ are the control variables for the system, $w = \begin{bmatrix} w^1, & w^2, & w^3 \end{bmatrix}^T \sim \mathcal{N}(0, \Sigma_w)$ is the process noise of the system, and t_e is the sampling time according to a Euler discretization of the continuous time model in [27]. In the model, $\begin{bmatrix} x^1, & x^2 \end{bmatrix}^T \in \mathbb{R}^2$ denotes the position of the aircraft in two dimensions and $x^3 \in [-\pi, \pi]$ denotes the heading angle of the aircraft. The linear velocity of the aircraft takes values between the minimum and maximum aircraft velocity, i.e. $a^1 \in [v_{\min}, v_{\max}]$, with $v_{\min} \in \mathbb{R}$ and $v_{\max} \in \mathbb{R}$. The angular velocity of the aircraft takes one of three possible values, corresponding to the three modes of operation of the DTSHS, i.e. $a^2 \in \{0, -u, u\}$ where $u \in \mathbb{R}$ is the angular velocity of the aircraft when in turning mode. In the following we consider a sampling time of $t_e = 1$ minute, aircraft speed $a^1 = 7.1$ km per minute, angular velocity $u = 0.3$ radians per minute, and disturbance variance $\Sigma_w \in \mathbb{R}^{3\times3}$ defined by $\Sigma_w(1, 1) = 0.25$, $\Sigma_w(2, 2) = 0.25$, $\Sigma_w(3, 3) = 0.05$, and $\Sigma_w(i, j) = 0$ if $i \neq j$. As in [27], the protected zone of the aircraft is an 8 km cylindrical block in the state space x that should not intersect the weather obstacle.

In the following examples, we consider the maximization and verification of aircraft trajectory safety and success given a probabilistic hazard forecast. In the first example, we model the probabilistic obstacles as a sequence of independent random closed sets and evaluate the probability of the aircraft attaining a target region while avoiding the hazardous regions. In the second example, we augment the state space of the aircraft DTSHS (21) with a parameterization of the hazardous regions and evaluate the probability of safety of the aircraft.

In both examples, set-valued obstacle processes are constructed according to the forecast product [12], which provides VIL numbers in a 1 km by 1 km gridded form for the entire United States over a time horizon of two hours. We consider a section of airspace centered at latitude 30^o and longitude 86.5^o, near the Florida gulf coast, on 01/07/2009, a day in which storms were observed in the region under consideration. Evaluating the data from [12] for the location and time above, we have extracted a thirty minute forecast comprising centers m_k and eccentricities M_k at one minute increments, i.e. $k \in \{0, \dots, 30\}$. Figure 5 represents the deterministic forecast over a thirty minute period. Executions of the random sets are a function of the deterministic forecast in both the decoupled and coupled examples. Figure 6 shows an aircraft path and obstacle location over a 10 minute period for a flight on the same day. While the aircraft path avoids the deterministic forecast, it intersects the hazardous region which is shown by the true obstacle location obtained from the weather data.

4.1 Decoupled Process

Consider the region $\bar{K} = [3600, 3800] \times [750, 850] \times [-\pi, \pi]$ with target set $K = [3742, 3768] \times [752, 778] \times [-\pi, \pi]$ and safe set

$$K'_k = \bar{K} \setminus \Xi_k.$$

Given forecast data extracted from [12] in the form of expected centers m_k and expected eccentricities M_k, we consider a time-indexed probabilistic model of the ellipse pa-

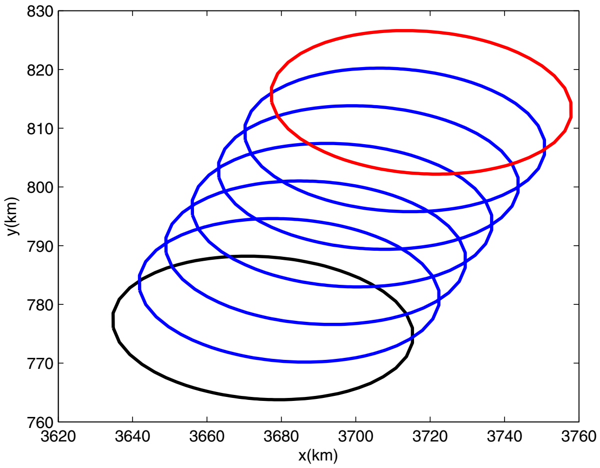

Figure 5: Deterministic weather forecast over a thirty minute period, given at 5 minute increments.

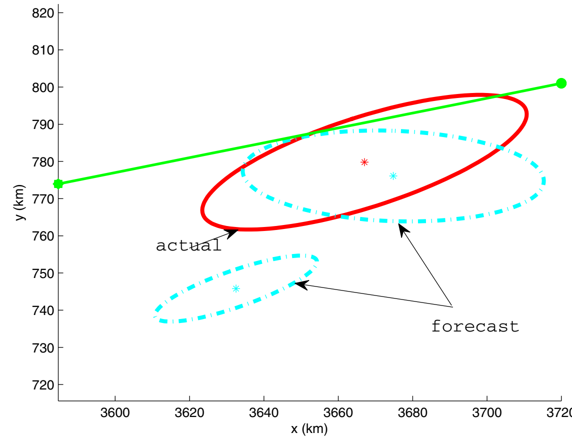

Figure 6: Path of aircraft around the forecasted and actual weather obstacles for a 5-minute portion of the flight. Notice that the path avoids the forecasted ellipse but intersects with the actual storm.

rameters

$$\xi_k \sim \mathcal{N}(m_k, \Sigma_k), \tag{22}$$

$$C_k = R(\theta_k)^T M_k R(\theta_k), \tag{23}$$

where Σ_k is a covariance matrix, $R(\cdot)$ is a rotation matrix, and the angle of rotation θ_k at time k is a random variable with uniform distribution over an interval $[-\alpha, \alpha]$. Σ_k and α are defined as in Section 2.3 and the initial parameter values are $\xi_0 = [3675, 776]^T$ and $\theta_0 = 0$ (see Figure 5). Accounting for the hazardous weather regions of Section 2.3 and the protected zone of the aircraft, the random closed set Ξ_k is consequently defined

$$\Xi_k = \mathcal{E}(\xi_k, C_k) \oplus \mathcal{C}(0, 8) \times [-\pi, \pi]$$

where $\mathcal{C}(c, r)$ is a circle (representing the protected zone of the aircraft) defined by its center c and radius r and $\oplus$ denotes the Minkowski sum.

Considering the DTSHS for aircraft motion (21) and the model of probabilistic obstacles, characterized through (22) and (23), we would like to evaluate (and subsequently maximize) the probability that an aircraft attains K while avoiding the hazardous regions over a horizon of thirty minutes (i.e. $k \in \{0, \dots, 30\}$). All numerical computations were performed on a $201 \times 101 \times 40$ grid according to the methods in [1]. The optimal value function, which represents the maximum probability of attaining the target region safely at some point during the time horizon, is shown in Figure 7 for an initial heading angle of $x_0^3 = -0.0785$ radians. For example, the DTSHS initialized at $x_0 = [3620,\ 830,\ -0.0785]^T$ has a maximum probability of success of 93.3 percent according to the optimal value function. An example execution of the process from $x_0 = [3620,\ 830,\ -0.0785]^T$ is shown in Figure 8.

4.2 Coupled Process

Here we consider the DTSHS for aircraft motion (21) and the set-valued Markov process for obstacle (hazardous weather) movement given in Section 2.3 by the equations (7) and (8). According to Section 3.2 we augment the state of the DTSHS with the state of the obstacle such that the state of the coupled Markov process is

$$[x^T, \xi^T, \theta]^T \in \mathbb{R}^6.$$

Solving a dynamic program in 6 dimensions is intractable due to the Curse of Dimensionality. We therefore make the following modifications. We assume that $\theta_k = 0$ and $M_k = M$ for all k, thereby removing θ as a state. Additionally, we form a new state corresponding to the relative coordinate of the aircraft and obstacle location

$$[x^1, x^2]^T - \xi \in \mathbb{R}^2$$

and remove the states x^1, x^2, ξ^1, and ξ^2. The resulting state of the coupled process is

$$\bar{x} = \begin{bmatrix} x^1 - \xi^1 \\ x^2 - \xi^2 \\ x^3 \end{bmatrix} = \begin{bmatrix} \bar{x}^1 \\ \bar{x}^2 \\ \bar{x}^3 \end{bmatrix}.$$

Combining equations (21) and (7), the process dynamics of the augmented system are given by the difference equations

$$\begin{aligned} \bar{x}_{k+1}^1 &= \bar{x}_k^1 + t_e a_k^1 \cos(\bar{x}_k^3) + w_k^1 - \mu_k^1 - \eta_k^1 \\ \bar{x}_{k+1}^2 &= \bar{x}_k^2 + t_e a_k^1 \sin(\bar{x}_k^3) + w_k^2 - \mu_k^2 - \eta_k^2 \\ \bar{x}_{k+1}^3 &= \bar{x}_k^3 + t_e a_k^2 + w_k^3. \end{aligned} \tag{24}$$

By defining an augmented system for the coupled process, which combines the dynamics of the DTSHS and the dynamics of the random obstacle process, we can now define a reach-avoid problem in the spirit of [25]. Consider $\bar{K}_1 = \mathbb{R} \times \mathbb{R} \times [-\pi, \pi]$ and $\bar{K}_2 = [-69, 89] \times [-24, 40] \times [-\pi, \pi]$. We define the target region $K = \bar{K}_1 \setminus \bar{K}_2$ and the safe set $K' = \bar{K}_1 \setminus \Xi$, where the obstacle set Ξ is static, deterministic, and defined

$$\Xi = \mathcal{E}(0, M) \oplus \mathcal{C}(0, 8) \times [-\pi, \pi]$$

where $\mathcal{C}(c, r)$ is a circle (representing the protected zone of the aircraft) defined by its center c and radius r and $\oplus$ denotes the Minkowski sum.

Considering the ADTSHS (24), we would like to evaluate (and subsequently maximize) the probability that an aircraft

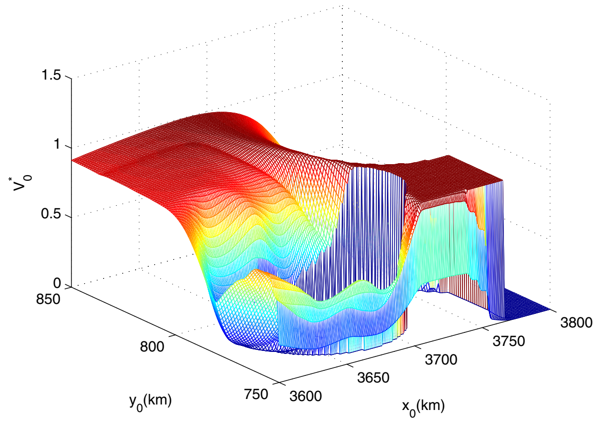

Figure 7: Optimal value function for the aircraft path planning problem with decoupled Markov processes (with initial heading angle $\bar{x}_0^3 = -0.0785$).

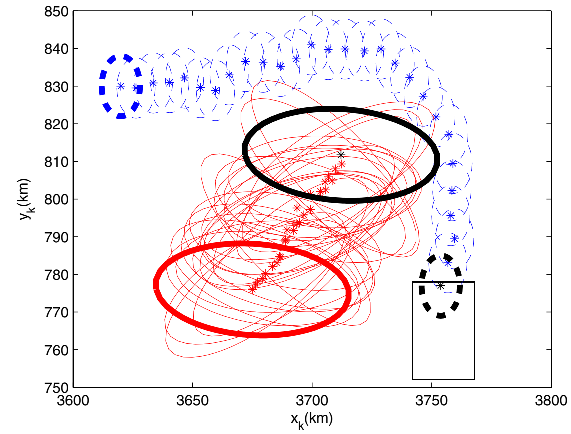

Figure 8: Example process execution for the aircraft path planning problem with decoupled Markov processes.

attains K (i.e. the aircraft escapes a region of the airspace that is considered too close to the hazardous weather region) while avoiding the hazardous region Ξ over a horizon of thirty minutes (i.e. $k \in \{0, \dots, 30\}$). All numerical computations were performed on a $161 \times 67 \times 20$ grid according to the methods in [1]. The optimal value function, which represents the maximum probability of attaining the target region safely at some point during the time horizon, is shown in Figure 9 for an initial aircraft heading angle of $\bar{x}_0^3 = -0.1571$ radians (note that the value function is shown as $1 - V_0^*$). For example, the DTSHS initialized at $\bar{x}_0 = [-60,\ -10,\ -0.1571]^T$ has a maximum probability of success of 83.08 percent according to the optimal value function. The $V_0^* = 0.95$ level set of the optimal value function is shown in Figure 10. Note that all initial conditions which start outside the level set have a success probability greater than 95 percent.

5. CONCLUSION

Extending the methods of [2, 25] and integrating the theory of random closed sets [18, 21], we formulated a reach-avoid problem with random obstacles as a finite horizon stochastic optimal control problem. We considered two possibilities for the random set-valued obstacle process. In the first, we considered the random set process to be an independent stochastic process, and thus decoupled from the evolution of the DTSHS. In the second case, we considered the obstacle process as a set-valued Markov process, equivalently expressed through parameterization. In both cases, it was shown that dynamic programming can be used to compute the Markov control policy that maximizes or minimizes the cost of the optimal control problem. A numerical example motivated by aircraft motion planning under uncertain weather predictions was used to illustrate the effectiveness of the methods introduced.

6. REFERENCES

[1] A. Abate, S. Amin, M. Prandini, J. Lygeros, and S. Sastry. Computational approaches to reachability analysis of stochastic hybrid systems. In A. Bemporad, A. Bicchi, and G. C. Buttazzo, editors, *Hybrid Systems: Computation and Control*, volume 4416 of *Lecture Notes in Computer Science*, pages 4–17. Springer, 2007.

[2] A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, November 2008.

[3] S. A. Amburn and P. L. Wolf. VIL density as a hail indicator. *Weather and Forecasting*, 12(3):473–478, 1997.

[4] S. Amin, A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Reachability analysis for controlled discrete time stochastic hybrid systems. In Hespanha and Tiwari [15], pages 49–63.

[5] J.-P. Aubin. *Viability theory*. Birkhauser Boston Inc., Cambridge, MA, USA, 1991.

[6] R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, New Jersey, USA, 1957.

[7] D. P. Bertsekas and S. E. Shreve. *Stochastic Optimal Control: The Discrete-Time Case*. Athena Scientific, February 2007.

[8] B. Boudevillain and H. Andrieu. Assessment of vertically integrated liquid (VIL) water content radar measurement. *Journal of Atmospheric and Oceanic Technology*, 20(6):807–819, 2003.

[9] M. L. Bujorianu and J. Lygeros. Reachability questions in piecewise deterministic markov processes. In O. Maler and A. Pnueli, editors, *Hybrid Systems: Computation and Control*, volume 2623 of *Lecture Notes in Computer Science*, pages 126–140. Springer, 2003.

[10] N. Cressie and G. M. Laslett. Random set theory and problems of modeling. *SIAM Review*, 29(4):pp. 557–574, December 1987.

[11] M. Davis. *Markov Models and Optimization*. Chapman & Hall, London, 1993.

[12] J. Evans, K. Carusone, M. Wolfson, B. Crowe, D. Meyer, and D. Klingle-Wilson. The Corridor Integrated Weather System (CIWS). In *10th*

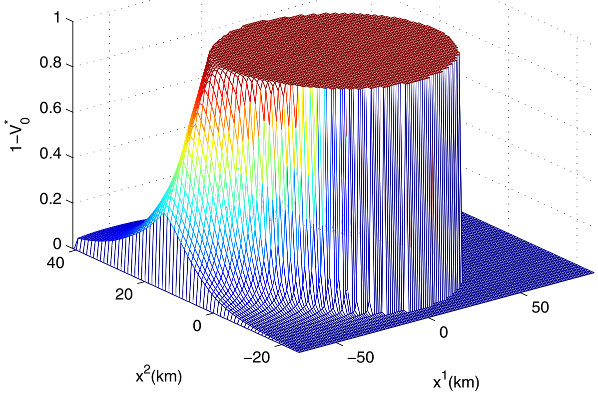

Figure 9: Optimal value function for the aircraft path planning problem with coupled Markov processes (portrayed as $1 - V_0^*$ with initial heading angle $x_0^3 = -0.1571$).

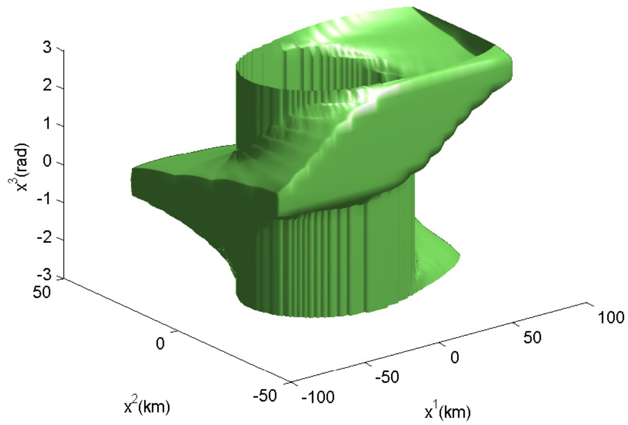

Figure 10: Level set corresponding to $V_0^* = 0.95$ for the aircraft path planning problem with coupled Markov processes.

Conference on Aviation, Range, and Aerospace Meteorology, 2001.

[13] Y. Gao, J. Lygeros, and M. Quincampoix. On the Reachability Problem for Uncertain Hybrid Systems. *IEEE Transactions on Automatic Control*, 52(9):1572–1586, September 2007.

[14] M. K. Ghosh, A. Arapostathis, and S. I. Marcus. Ergodic control of switching diffusions. *SIAM J. Control Optim.*, 35(6):1952–1988, November 1997.

[15] J. P. Hespanha and A. Tiwari, editors. *Hybrid Systems: Computation and Control*, volume 3927 of *Lecture Notes in Computer Science*. Springer, 2006.

[16] J. Hu, J. Lygeros, and S. Sastry. Towards a theory of stochastic hybrid systems. In N. A. Lynch and B. H. Krogh, editors, *Hybrid Systems: Computation and Control*, volume 1790 of *Lecture Notes in Computer Science*, pages 160–173. Springer, 2000.

[17] J. Lygeros. On reachability and minimum cost optimal control. *Automatica*, 40(6):917–927, 2004.

[18] G. Matheron. *Random Sets and Integral Geometry*. Wiley, New York, 1975.

[19] I. Mitchell. The flexible, extensible and efficient toolbox of level set methods. *J. Sci. Comput.*, 35(2-3):300–329, 2008.

[20] I. Mitchell and C. Tomlin. Level set methods for computation in hybrid systems. In *Hybrid Systems: Computation and Control*, pages 310–323, London, UK, 2000. Springer-Verlag.

[21] I. Molchanov. *Theory of Random Sets*. Springer, New York, 2005.

[22] F. Ramponi, D. Chatterjee, S. Summers, and J. Lygeros. On the connections between PCTL and dynamic programming. In *HSCC '10: Proceedings of the 13th ACM international conference on Hybrid systems: computation and control*, pages 253–262, New York, NY, USA, 2010. ACM.

[23] W. Rudin. *Real and complex analysis, 3rd ed.* McGraw-Hill, Inc., New York, NY, USA, 1987.

[24] D. Stoyan. Random sets: Models and statistics. *International Statistical Review*, 66(1):pp. 1–27, April 1998.

[25] S. Summers and J. Lygeros. Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica*, 46(12):1951 – 1961, 2010.

[26] C. Tomlin, J. Lygeros, and S. Sastry. A Game Theoretic Approach to Controller Design for Hybrid Systems. *Proceedings of IEEE*, 88:949–969, July 2000.

[27] C. Tomlin, G. J. Pappas, and S. Sastry. Conflict resolution for air traffic management: a study in multiagent hybrid systems. *Automatic Control, IEEE Transactions on*, 43(4):509–521, 1998.

[28] C. J. Tomlin, I. Mitchell, A. M. Bayen, and M. Oishi. Computational techniques for the verification of hybrid systems. *Proceedings of the IEEE*, 91(7):986–1001, July 2003.