Stochastic receding horizon control for robots with probabilistic state constraints

Shridhar K. Shah, Chetan D. Pahlajani, Nicholaus A. Lacock and Herbert G. Tanner

***Abstract*— This paper presents a receding horizon control design for a robot subject to stochastic uncertainty, moving in a constrained environment. Instead of minimizing the expectation of a cost functional while ensuring satisfaction of probabilistic state constraints, we propose a two-stage solution where the path that minimizes the cost functional is planned deterministically, and a local stochastic optimal controller with exit constraints ensures satisfaction of probabilistic state constraints while following the planned path. This control design strategy ensures boundedness of errors around the reference path and collision-free convergence to the goal with probability one under the assumption of unbounded inputs. We show that explicit expressions for the control law are possible for certain cases. We provide simulation results for a point robot moving in a constrained two-dimensional environment under Brownian noise. The method can be extended to systems with bounded inputs, if a small nonzero probability of failure can be accepted.**

Keywords - stochastic receding horizon, exit time, stochastic optimal control, stochastic path following

I. Introduction

Uncertainty plagues almost all robotic systems, especially those deployed in the real world. Sources of uncertainty are un-modeled dynamics, measurement errors, unknown environment or environmental effects and component failures. The uncertainty is often expressed stochastically, in which case, one obtains a dynamic system in the form of a stochastic differential equation (SDE). The problem of robot navigation with guarantees of collision avoidance and convergence becomes challenging in the stochastic framework due to the unbounded nature of the applied disturbances.

In this paper, we consider robots with stochastic dynamics moving in constrained environments, and we want to design controllers which can *guarantee* that even in the presence of uncertainty, the system does not collide with obstacles, does not deviate too much from a nominal path, and converges to its goal configuration while satisfying a given formal probabilistic convergence criterion. Our approach is to construct a reference path to the goal using the drift (known) term of the robot's dynamics and use stochastic controllers in receding horizon manner to keep the system close to the path and force the robot to follow it asymptotically to the goal.

In a stochastic framework, state constraints are popularly treated as *chance-constraints* [1]–[6]. Almost all the existing chance constraint formulations are for linear discrete-time systems. In a recent work, the *chance-constraint optimization* problem for finite horizon planning with obstacles was solved using mixed integer linear programming (MILP) [6]. Our intent is to implement stochastic receding horizon control on platforms with limited computational resources and nonlinear dynamics. The solution of the aforementioned chance-constraint problems is currently too computationally demanding for real-time implementation on the robotic platforms of interest. Hence, this paper deals with the receding horizon control framework for nonlinear continuous-time systems and we look for closed form solutions of control laws to reduce computation burden.

There exist multiple efforts toward a solution of stochastic model predictive control and it is difficult to provide a comprehensive review of literature due to limited space. We refer readers to [6] and references there in for the same. In addition, there exists model predictive control method based on path integral [7], which is difficult to extend for practical problems while method using density propagation [8] faces curse of dimensionality for extension to practical systems.

Methods based on dynamic programming often express the optimal value for the cost as the solution of a Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE). Viscosity solutions of the HJB in the context of stochastic optimal control are treated in [9]. The latter reference [9] (see also [10]) showed that the logarithm of some exit probability represents the solution of a particular HJB equation. In this setting, one can derive stochastic optimal control laws that guarantee that the system does not hit a pre-specified region of the workspace boundary (obstacles) with probability one. The control law can be found analytically in certain simple cases. What motivates us is the realization that the complement of this exit region can be mapped to the boundary of a region around the goal (or intermediate waypoints), in which case the formulation can be exploited to ensure obstacle avoidance for the robot navigation scenario we consider.

Toward this end, we adopt an approach to the considered stochastic robot navigation problem in which a reference path from initial to final configuration is planned using the drift term of the dynamics only, and then a finite sequence of overlapping regions cover this path. An exit time optimal controller [10] can then be used to guide the system from one region to the next in a cascaded fashion all the way to its destination. The overall control strategy is a *switched* one, where a series of different control laws are implemented, one for each local region.

Such a two-stage approach has a long history in deterministic robot navigation—see [11]–[13] for indicative applications. This type of switching strategy allows simpler

This work is supported by ARL MAST CTA # W911NF-08-2-0004
Shridhar K. Shah and Herbert G. Tanner are with Department of Mechanical Engineering, University of Delaware, Newark, DE, USA {shridhar, nlacock, btanner}@udel.edu
Chetan D. Pahlajani is with Department of Mathematical Sciences, University of Delaware, Newark, DE, USA chetan@math.udel.edu

control design, since the local potential fields the system switches over are easier to construct and tune compared to a global one. While these deterministic approaches will fail to ensure collision avoidance in a stochastic setting, the idea of computational and analytical savings from the local construction carries over. Thus, instead of using a general chance constrained optimization approach, we opt for the sequential solution of local exit problems which by appropriate choice of workspace boundaries might even afford closed form solutions.

As each subsystem's evolution is truncated at an exit time, the complete solution can be thought of as a collection of truncated Markov processes, one for each (local) exit problem. To make sure such solutions of our closed loop system are well defined as the stochastic optimal controllers switch, we use the notion of Markov strings [14]. In this setting, concatenations of such truncated processes provide a complete solution and can be shown to possess the strong Markov and càdlàg properties.

The contribution of this paper is a planning and control framework for receding horizon control of continuous time stochastic systems that provides closed form control laws to drive the system towards the goal in finite time and avoid obstacles with probability one. Since the uncertain term in the SDE can take arbitrary large values with a small but nonzero probability, such performance guarantees can only be achieved under assumption of unbounded inputs. Moreover, the receding horizon framework requires that the probability of reaching intermediate way-points is one, otherwise if there is a large number of such way-points, the probability of reaching the final goal becomes arbitrary small. For practical purposes, the assumption of unbounded inputs can be relaxed by allowing a small probability of collision and restricting the number of way-points to a comfortable finite natural.

The work presented in this paper is organized in the following way. Section II states the problem formally. Section III is a brief introduction to computing exit probabilities, which is utilized in Section IV for control design. The latter section presents our main results on the design of the navigation controller followed by an example in Section V. A comment on existence of solutions for our closed loop system is presented in Section VI. The stability and convergence aspects are covered in Section VII followed by conclusion in Section VIII.

II. PROBLEM STATEMENT

We consider an open bounded region $\mathcal{W} \subseteq \mathbb{R}^n$ which represents our workspace. A closed set $\mathcal{O} \subset \mathcal{W}$ represents forbidden regions (obstacles) in the workspace. Thus, the free workspace is $\mathcal{P} \triangleq \mathcal{W} \backslash \mathcal{O}$. Let $W = \{W(t), \mathcal{F}_t : 0 \leq t < \infty\}$ be an m-dimensional Brownian motion on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ where Ω is the sample space, $\mathcal{F}$ is a σ-algebra on Ω, $\mathbb{P}$ is the probability measure and $\{\mathcal{F}_t : t \geq 0\}$ is the filtration (i.e., an increasing family of sub-σ-algebras of $\mathcal{F}$) that is right continuous and $\mathcal{F}_0$ contains all $\mathbb{P}$-null sets [15]. Consider a representation of a robot with stochastic uncertainty in the form of an SDE within the free workspace $\mathcal{P}$

$$dx(t) = b(x(t))dt + u_i(x(t))dt + \sigma(x(t))dW(t) \ , \quad (1)$$

with $x(0) = x_0$, where $x \in \mathbb{R}^n$ is the state, $b : \mathbb{R}^n \to \mathbb{R}^n$ is the drift term, $u_i : \mathbb{R}^n \to \mathbb{R}^n$ is an affine control input and $\sigma : \mathbb{R}^n \to \mathbb{R}^{n\times m}$ is the diffusion term of the system. Subscript $i \in \mathbb{N}^+$ is a switching index that indicates which of the vector fields from a family $\{u_i : i \in \mathbb{N}^+\}$ is active at a given time $t \geq 0$. We define $a : \mathbb{R}^n \to \mathbb{R}^{n\times n}; x \mapsto a(x) \triangleq \sigma(x)\sigma^T(x)$.

The problem is to find a control law to drive a robot to the goal in the constrained environment $\mathcal{P}$ without hitting the obstacles $\mathcal{O}$, i.e. satisfying the state constraint:

$$\mathbb{P}[x(t) \in \mathcal{O}] = 0, \quad \forall t \geq 0,$$

and minimizing the following discrete finite horizon cost:

$$\min_{x_i} \mathcal{J} = \sum_{i=0}^{n} L(x_i) + V(x_n); \text{ s.t. } \|x_{i-1} - x_i\| \leq d$$

and continuous local cost:

$$\min_{u_i} J_i = \int_{t_{i-1}}^{t_i} (u_i)^T a(x)(u_i)dt,$$

s.t. $\mathbb{P}[\|x(t_i) - x_i\| \leq \varepsilon_i] = 1$. Here, x_i's are discrete control horizon way-points along the receding horizon trajectory, x_n represents the point at the end of prediction horizon, t_i represents the time at i^{th} control horizon, the ε_i's are arbitrary small numbers, $V(x)$ is a positive definite function and d is a given constant. Minimization of the above costs involves first selecting the series of way-points to minimize $\mathcal{J}$ and then designing control inputs to navigate between these way-points while minimizing J_i.

III. PRELIMINARIES

Consider a stochastic system similar to (1),

$$dx(t) = b(x(t))dt + \sigma(x(t))dW(t); \ x(0) = x_0,$$

where $b : \mathbb{R}^n \to \mathbb{R}^n$ is the drift and $\sigma : \mathbb{R}^n \to \mathbb{R}^{n\times m}$ is the diffusion term. Let $\mathcal{L}$ be the second-order partial differential operator

$$\mathcal{L} \triangleq \sum_{i=1}^{n} b_i(x)\frac{\partial}{\partial x_i} + \frac{1}{2}\sum_{i=1}^{n}\sum_{k=1}^{n} a_{ik}(x)\frac{\partial^2}{\partial x_i \partial x_k}. \quad (2)$$

Consider an open bounded domain $\mathcal{D} \subset \mathbb{R}^n$, the closure $\overline{\mathcal{D}}$ and boundary $\partial\mathcal{D}$. Also assume that $b(x)$ and $\sigma(x)$ are Lipschitz continuous in $\overline{\mathcal{D}}$.

Under the assumption $\min_{x\in\overline{\mathcal{D}}} a_{ll}(x) > 0$ for some $1 \leq l \leq n$, one can show that $\mathbb{E}^{x_0}[\tau_{\mathcal{D}}] < \infty$[1], $\forall x_0 \in \overline{\mathcal{D}}$ [15, Lemma 7.4], where $\tau_{\mathcal{D}}$ is the first exit time from $\mathcal{D}$.

Moreover, expectations of the form $\mathbb{E}^{x_0}[f(x(\tau_{\mathcal{D}}))]$ for any continuous function $f : \partial\mathcal{D} \to \mathbb{R}$ can be computed by finding

[1]The representation $\mathbb{E}^{x_0}$ means the expectation with respect to initial condition being x_0

a function $h \in C(\overline{\mathcal{D}}) \cap C^2(\mathcal{D})$ which solves the Dirichlet problem

$$\mathcal{L}h = 0 \qquad \text{in } \mathcal{D} \tag{3a}$$
$$h = f \qquad \text{on } \partial\mathcal{D}, \tag{3b}$$

where, $\mathcal{L}$ is assumed to be uniformly elliptic in $\mathcal{D}$, that is, for some $\delta > 0$,

$$\sum_{i=1}^{n}\sum_{k=1}^{n} a_{ik}(x)\zeta_i\zeta_k \geq \delta\|\zeta\|^2; \quad \forall x \in \mathcal{D}, \zeta \in \mathbb{R}^n$$

It can be shown [15, Proposition 7.2], [15, Lemma 7.4] that, if $h \in C(\overline{\mathcal{D}}) \cap C^2(\mathcal{D})$ is a solution to the Dirichlet problem (3) with $f : \partial\mathcal{D} \to \mathbb{R}$ a continuous function, then

$$h(x_0) = \mathbb{E}^{x_0}[f(x(\tau_\mathcal{D}))].$$

In particular, if[2]

$$f_{\partial\mathcal{D}_1}(x) = \begin{cases} 1 & \text{if } x \in \partial\mathcal{D}_1 \\ 0 & \text{if } x \in \partial\mathcal{D}_2 \end{cases}$$

where $\partial\mathcal{D}_1 \cup \partial\mathcal{D}_2 = \partial\mathcal{D}$ then

$$h(x_0) = \mathbb{P}^{x_0}(x(\tau_\mathcal{D}) \in \partial\mathcal{D}_1)$$

represents the probability of hitting a particular part of boundary for initial condition $x(0) = x_0$.

IV. Stochastic Receding Horizon Control

Consider the switched stochastic system (1) defined in Section II,

$$dx(t) = b(x(t))dt + u_i(x(t))dt + \sigma(x(t))dW(t) \tag{4}$$

with $x(0) = x_0$, $x_0 \in \mathcal{P}$. The system undergoes forced switching where the switch occurs upon the state hitting a particular part of boundary of a domain. In between each switch, the system is a Markov process represented by an SDE and its exit time is represented as τ_i. The switching times are thus described as $t_i \triangleq \sum_{j=1}^{i} \tau_j$.

In receding horizon framework one plans a trajectory over $(0, T]$ where T is the prediction horizon, executes a part of the trajectory over time $(0, \delta]$, $\delta < T$ and recomputes the prediction horizon trajectory for $(\delta, T + \delta]$. In our implementation, we use drift term to plan a discrete time trajectory for the length of the prediction horizon, $\{x_i\}_{i=0}^{n}$, implement it until the first way-point x_1 and then repeat. The discrete sequence of points $\{x_i\}$ satisfies

$$\min_{x_k} \mathcal{J} = \sum_{k=i}^{n+i} L(x_k) + V(x_{n+i}); \;\; \|x_{i-1} - x_i\| \leq d$$

where $\{x_k\}_{k=i}^{n+i} \cap \mathcal{O} = \emptyset$ and straight line segments connecting these points do not intersect $\mathcal{O}$.

The objective of the local navigation is to steer the stochastic system through the succession of these waypoints making the state of (4) visit arbitrarily small neighborhood of these points. For that purpose, we construct overlapping domains the union of which contain all waypoints and such that each domain contains two consecutive points in the sequence. The overlapping domains are defined such that their union defines a collision-free safe corridor in which the system will be steered to its goal configuration.

A. Stochastic Navigation Controller

Assume that the system is ε-close to waypoint x_{i-1} and the goal is to reach an ε neighborhood of x_i with probability one, before hitting any obstacles. For the purpose, we construct a bounded region $\mathcal{D}_i$, the closure of which is denoted $\overline{\mathcal{D}}_i$ and its boundary $\partial\mathcal{D}_i$. The boundary $\partial\mathcal{D}_i$ consists of two disjoint parts: $\mathcal{N}_i \cup \mathcal{N}_i^c = \partial\mathcal{D}_i$, where $\mathcal{N}_i = \{x \in \partial\mathcal{D}_i \mid \|x - x_i\| \leq \varepsilon_i\}$. Assume that

$$\mathcal{D}_i \cap \mathcal{O} = \emptyset, \;\; \mathcal{N}_{i-1} \subset \mathcal{D}_i \;\text{ and }\; \mathcal{N}_{i-1} \cap \mathcal{N}_i = \emptyset \tag{5}$$

A controller u_{i-1} acts on the system (1) while $x(t) \in \mathcal{D}_{i-1}$ and switches to controller u_i when the state $x(t)$ hits $\mathcal{N}_{i-1}$ and acts for time $t \in [t_{i-1}, t_i)$. When applied, $u_i(x)$ should satisfy the following probabilistic requirements:

$$\mathbb{E}^x[\tau_i] < \infty, \quad \tau_i = t_i - t_{i-1} \tag{6}$$
$$\mathbb{P}^x(x(t_i) \in \mathcal{N}_i^c) = 0 \Leftrightarrow \mathbb{P}^x(x(t_i) \in \mathcal{N}_i) = 1. \tag{7}$$

Condition (6) requires that each waypoint is reached in finite time. Condition (7) requires that the system reaches an ε-neighborhood of x_i with probability one before hitting the obstacles.

We design a local optimal stochastic controller u_i in an exit problem framework which ensures satisfaction of above state constraints with optimal inputs within the local domain $\mathcal{D}_i$. Intuitively, boundary $\mathcal{N}_i$ acts as an attraction region while boundary $\mathcal{N}_i^c$ acts as a repulsive region. The method presented is based on the results of [9], [10].

The Markov process $x(t)$ within the local domain $\mathcal{D}_i$ is described by an SDE (1)

$$dx(t) = b(x(t))dt + u_i(x(t))dt + \sigma(x(t))dW(t) \tag{8}$$

with initial condition $x(t_{i-1})$, within the local domain $\mathcal{D}_i \subset \mathbb{R}^n$, where $\mathcal{D}_i$ is assumed a bounded domain with $\mathcal{C}^2$ boundary. The functions $\sigma(x)$ and $b(x) + u_i(x)$ are assumed to be Lipschitz on $\mathcal{D}_i$ and $\sigma^{-1}(x)$ along with $\sigma(x)$ and $b(x)+u_i(x)$ are bounded. The local controller is built with an objective to find a solution to the following stochastic optimal control problem such that $u_i(x(t))$ is bounded, progressively measurable[3] and achieves,

$$\min \mathbb{E}^{x(t_{i-1})}\left[\int_{t_{i-1}}^{t_i} L(x(s), u_i(s))ds + \Phi(x(t_i))\right] \tag{9}$$

where t_i is the time when the system hits boundary $\mathcal{N}_i$ and $\tau_i = t_i - t_{i-1}$ is the exit time for the SDE evolving in $\mathcal{D}_i$. There exists an explicit solution to such an optimization problem for a particular choice of cost using a log transformation method [9]:

$$I(x) = -\log g(x)$$

[2]Such a function f will be continuous if $\partial\mathcal{D}_1$ and $\partial\mathcal{D}_2$ are closed and disconnected.

[3]A stochastic process is progressively measurable when it is non-anticipative.

where

$$g(x) = \mathbb{E}^x \left[\exp(-\Phi(\zeta(\tau_{\partial \mathcal{D}_i}))) \right].$$

In the above, $\tau_{\partial \mathcal{D}_i}$ is the first time when the system hits boundary $\partial \mathcal{D}_i$ and $g(x)$ is the solution of the PDE

$$\mathcal{L}g = 0 \qquad \text{in } \mathcal{D}_i \tag{10a}$$
$$g = \exp(-\Phi(\cdot))) \qquad \text{on } \partial \mathcal{D}_i, \tag{10b}$$

with $\mathcal{L}$ being the generator defined in (2) for an auxiliary Markov process $\zeta(t)$ on the same bounded open set $\mathcal{D}_i \subset \mathbb{R}^n$. The Markov process $\zeta(t)$ is given as

$$d\zeta(t) = b(\zeta(t))dt + \sigma(\zeta(t))dW(t).$$

Then, according to [9], the optimal control law $u^*(x(t))$ for (8) which minimizes (9) can be given as

$$u_i{}^*(x) = -a(x)\nabla I(x), \tag{11}$$

where the incremental cost is defined as

$$L(x, u) = \frac{1}{2}(u_i)' a(x)^{-1}(u_i).$$

Choosing Φ to be

$$\Phi = +\infty \cdot \mathcal{X}_{\mathcal{N}_i^c} \tag{12}$$

where

$$\mathcal{X}_{\mathcal{N}_i^c} = \begin{cases} 0 & \text{on } \mathcal{N}_i \\ 1 & \text{on } \mathcal{N}_i^c \end{cases}$$

yields a solution [10]

$$I(x) = -\log g(x)$$

$$g(x) = \mathbb{P}^x \left[\zeta(\tau_i) \in \mathcal{N}_i \right]$$

where $g(x)$ is computed as discussed in Section III. It can be shown that $I(x)$ is the optimal solution of the optimization problem [10] and u_i^* is the optimal feedback control which minimizes the cost,

$$\min \mathbb{E}^x \left[\int_{t_{i-1}}^{t_i} L(x(s), u_i(s)) ds \right]$$

In (12), Φ can be interpreted as an infinite penalty for state hitting the forbidden part of the boundary at the first exit time $x(\tau_i) \in \mathcal{N}_i^c$, $\tau_i = t_i - t_{i-1}$. The optimal control input u^* needs to be unbounded in order to satisfy the constraints with probability one; in the case of bounded inputs the system has some positive probability to exit at either part of the boundary of $\mathcal{D}_i$.

The controller u_i^* is recursively applied in receding horizon manner until the system reaches an ε-neighborhood of the goal. The control law (11) is computable analytically in simple cases and the method lends itself to real time applications. By simple case, we mean that the vector field $b(\cdot)$ and $\sigma(\cdot)$ are such that the boundary value problem for PDE (10) is solvable explicitly.

V. EXAMPLE

We present an example for a simple system describing a point robot with stochastic uncertainty in a 2-D space. Consider the system without drift and with an identity diffusion term,

$$d\mathbf{x}(t) = u_i(\mathbf{x}(t))dt + dW(t); \;\; \mathbf{x}(0) = \mathbf{x}_0, \tag{13}$$

where $\mathbf{x} = (x, y)$ and $W(t)$ is a 2-dimensional Brownian motion. The first step is to find a reference trajectory for (13). We use the nominal dynamics $\dot{\mathbf{x}} = u(\mathbf{x}(t))$ and the model predictive control approach presented in [16] to find a continuous trajectory for $t \in [0, T]$, where T is the prediction horizon.

Let us assume that a feasible continuous state trajectory is given as $\mathbf{x}^*(t) \in \mathcal{P}$. The control horizon points $\{\mathbf{x}_k\}_{k=i}^{n+i}$, are found along the solution trajectory $\hat{\mathbf{x}}^*(t)$, $t \in [0, T]$ such that given a positive definite function V,

$$\max_{y \in \mathcal{N}_i} \{V(y)\} - \min_{z \in \mathcal{N}_{i-1}} \{V(z)\} \leq -\gamma(\|\mathbf{x}_{i-1}\|)$$
$$\text{and } \; \|\mathbf{x}_{i-1} - \mathbf{x}_i\| \leq d$$

where d gives an upper bound on the distance between waypoints. The next step is to construct bounded domains around the waypoints. We construct these domains in the form of a set of concentric circles.

Consider notation for a closed ball $\overline{\mathcal{B}}(p, q) \triangleq \{\mathbf{x} \in \mathbb{R}^n : \|p - \mathbf{x}\| \leq q\}$. The boundary is represented as $\partial \mathcal{B}$ and the interior of the ball is represented as $\mathcal{B}(p, q) \triangleq \overline{\mathcal{B}}(p, q) \setminus \partial \mathcal{B}$.

Consider two concentric balls $\overline{\mathcal{B}}(\mathbf{x}_i, R_i)$ and $\overline{\mathcal{B}}(\mathbf{x}_i, \varepsilon_i)$ where R_i and ε_i are radius of concentric circles centered at $\mathbf{x}_i$, $\varepsilon_i < R_i$. The local domains $\mathcal{D}_i$ are defined as

$$\overline{\mathcal{D}}_i \triangleq \overline{\mathcal{B}}(\mathbf{x}_i, R_i) \setminus \mathcal{B}(\mathbf{x}_i, \varepsilon_i)$$
$$\partial \mathcal{D}_i \triangleq \partial \mathcal{B}(\mathbf{x}_i, \varepsilon_i) \cup \partial \mathcal{B}(\mathbf{x}_i, R_i)$$
$$\mathcal{D}_i = \overline{\mathcal{D}}_i \setminus \partial \mathcal{D}_i \;\; .$$

Here, $\mathcal{N}_i = \partial \mathcal{B}(\mathbf{x}_i, \varepsilon_i)$ and $\mathcal{N}_i^c = \partial \mathcal{B}(\mathbf{x}_i, R_i)$.

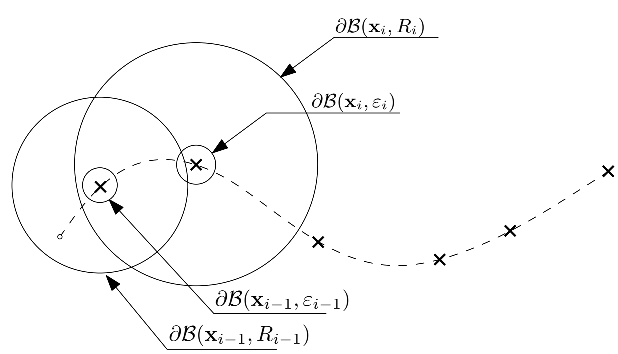

Fig. 1. The construction of local domains. The dotted line shows deterministically planned trajectory and cross marks are the waypoints. The bigger circles show the boundary $\partial \mathcal{B}(\mathbf{x}_i, R_i)$ while the smaller circles show the boundary $\partial \mathcal{B}(\mathbf{x}_i, \varepsilon_i)$.

We pose the following constraints for determining the values of radius R_i and ε_i,

$$\mathcal{B}(\mathbf{x}_i, R_i) \cap \mathcal{O} = \emptyset \quad (14)$$

$$\mathcal{B}(\mathbf{x}_{i-1}, \varepsilon_{i-1}) \subset \mathcal{B}(\mathbf{x}_i, R_i) \setminus \overline{\mathcal{B}}(\mathbf{x}_i, \varepsilon_i) \quad (15)$$

$$\varepsilon_i < \|\mathbf{x}_i - \mathbf{x}(t)\| < R_i, \quad t \in [t_{i-1}, t_i) \quad (16)$$

where, t_{i-1} and t_i are the exit times of the system (1) and correspond to hitting at boundary $\partial\mathcal{B}(\mathbf{x}_{i-1}, \varepsilon_{i-1})$ and $\partial\mathcal{B}(\mathbf{x}_i, \varepsilon_i)$ respectively. There is no unique solution for the above constraints and one can further define an arbitrary upper bound for R_i and an arbitrary lower bound for ε_i for the system. The construction of such domains for a two dimensional system is shown in Fig. 1. The local control laws for such a system can be constructed as explained in Section IV-A, achieving

$$\min \mathbb{E}^{\mathbf{x}(t_{i-1})}\left[\frac{1}{2}\int_{t_{i-1}}^{t_i} u^2\right]$$

and is expressed as $u^*(\mathbf{x}) = -\nabla I(\mathbf{x})$ where $I(\mathbf{x}) = -\log g(\mathbf{x})$ and $g(\mathbf{x}) = \mathbb{P}^{\mathbf{x}(t_{i-1})}\left[\zeta(\tau_i) \in \mathcal{N}_i\right]$. The auxiliary process ζ is given as $d\zeta(t) = dW(t)$ and $g(x)$ is the solution of the exit time problem

$$\mathcal{L}g = 0 \qquad \text{in } \mathcal{D}_i$$
$$g = 0 \quad \text{on } \mathcal{N}_i^c, \quad g = 1 \quad \text{on } \mathcal{N}_i$$

and $\mathcal{L} = \frac{1}{2}\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right)$.

Using the derivations of Section III, we can find $g(\mathbf{x})$ to be

$$g(\mathbf{x}) = \frac{R_i - \|\mathbf{x} - \mathbf{x}_i\|}{R_i - \varepsilon_i} \ .$$

Hence:

$$I(\mathbf{x}) = -\log \frac{R_i - \|\mathbf{x} - \mathbf{x}_i\|}{R_i - \varepsilon_i}$$

and the control law is given as

$$u_i(\mathbf{x}) = -\nabla I(\mathbf{x}) = \frac{-(\mathbf{x} - \mathbf{x}_i)}{(R_i - \|\mathbf{x} - \mathbf{x}_i\|)\|\mathbf{x} - \mathbf{x}_i\|} \ . \quad (18)$$

Control input $u_i(\mathbf{x})$ switches to $u_{i+1}(\mathbf{x})$ upon hitting the boundary $\mathcal{N}_i$ for $i = 1, 2, \ldots$ until the state is in ε-neighborhood of the goal configuration.

A. Simulation Results

Simulations are performed for a point robot moving in $\mathbb{R}^2$. The overall bounded domain is considered to be $\mathcal{W} = \mathcal{B}(0, 10)$, and the robot initial condition is taken as $\mathbf{x}_0 = (x, y) = (-3.0, -3.0)$ m. The goal is to drive the system to the origin. The environment includes two obstacles at $(-3.0, -1.0)$ m and $(-2.0, -2.0)$ m with radius 0.2 m. Simulation results are shown in Fig. 2.

A navigation function $V(x)$ is first constructed on $\mathbb{R}^2$ and a randomized algorithm is used to generate a trajectory for $\dot{\mathbf{x}} = u(\mathbf{x}(t))$ [16]. The navigation function is depicted in the form of a contour plot while the discrete control horizon waypoints are center of red circles in Fig. 2. The local navigation boundary is then decided based on (14)-(16) represented by dotted black circles. The robot is steered using the control law described as (18) under the effect of noise in the form of (13).

To compare our method with existing chance constraint approaches [6], we simulated two obstacle setup of Fig. 2. That method found an 8 way-point path in 1.19 seconds, while an 20 way-point path required 70.0 seconds. That path intersected with obstacles. In contras, an implementation of [16] yielded in a collision-free *continuous* path in 0.79 seconds. It is to be noted, however, that the contribution of the present paper is not the use of randomized receding horizon optimization, but on the construction of a two-stage planning strategy, using any deterministic path planning method to achieve closed-form stochastic receding horizon control with performance guarantees.

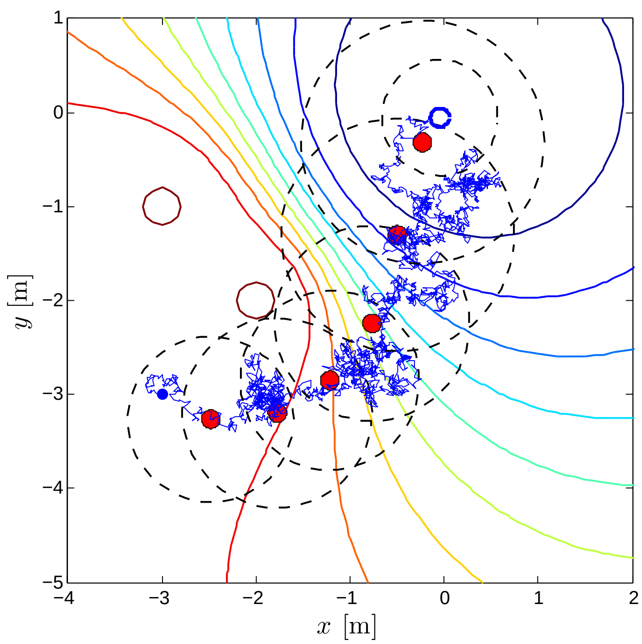

Fig. 2. Simulation of a stochastic receding horizon control for a point robot in a two obstacle environment. The blue trajectory shows the actual stochastic path taken by the robot. The initial condition of the robot is marked with a blue dot. The black dashed circles represent the outer boundary of domains $\mathcal{D}_i$ while red circles represent the region around control horizon with its boundary $\mathcal{N}_i$ and the blue circle is the boundary around the final goal. The simulation was generated using MATLAB™ Econometrics toolbox.

VI. Existence of Solutions

The solution of the system (1) is a collection of Markov processes truncated at exit time. Such a collection of truncated Markov processes can be represented as a Markov string. A Markov string is a hybrid state jump Markov process, where the continuous state evolves according to some evolution modes which is truncated at a given stopping time, jumps to a new mode according to a renewal kernel, while a discrete state defines the modes which could have either random jumps or forced jumps [14]. Jumps in the continuous state are also allowed at mode switching. A detailed construction of such Markov strings as in the form of an execution of *General Stochastic Hybrid System* is presented in [14]. There it is also shown that the Markov string enjoys the strong Markov and càdlàg properties. These properties are in fact inherited from the component processes.

The system uses only forced transitions while the renewal kernel is a simple Dirac measure. The continuous process under control law $u_i(x)$ is truncated at the exit time τ_i which is the exit time when x is close to waypoint x_i. Without state reset, the system updates its control law to a precomputed u_{i+1}. The new process starts as soon as u_{i+1} is implemented, and the procedure repeats until a neighborhood of the goal is reached. The use of Markov strings to represent this switched system requires that drift terms are bounded, in order for local solutions to exist; but if $\Phi = +\infty \cdot \mathcal{X}_{\mathcal{N}_i^c}$ in (12), control law $u_i^*(x)$ is unbounded on $\mathcal{N}_i^c$. To overcome this problem, one can assume $\Phi = +M \cdot \mathcal{X}_{\mathcal{N}_i^c}$ where M is a large known positive number representing an arbitrarily large penalty of hitting the obstacle boundary. The resulting probability of hitting $\mathcal{N}_i$ is close to one. The modification allows the inputs to be bounded and guarantees the existence of solution in the form of a Markov string. Note also that since the system path is defined through a sequence of finite waypoints, the system undergoes only a finite number of transitions, and the switching condition ensures that the system evolves forward; therefore, the system does not exhibit Zeno behavior.

VII. Convergence and Stability

Proposition 1: Consider the switched stochastic system (1) in a open bounded domain $\mathcal{W} \subset \mathbb{R}^n$, where $i \in \mathbb{N}^+$ is the switching index and $W(t)$ is a Wiener process. Let $V(x)$ be a $\mathcal{C}^2$, positive definite function in a closure of a bounded domain $\mathcal{W}$ which contains origin. For some class $\mathcal{K}$ function γ defined on $\mathcal{W}$ and for every solution of the stochastic switched system, if there exists a sequence of points $\{x_i\}_{i=1}^n \in \mathcal{W}$, such that,

$$\max_{y \in \mathcal{N}_i}\{V(y)\} - \min_{z \in \mathcal{N}_{i-1}}\{V(z)\} \le -\gamma(\|x_{i-1}\|) \tag{19}$$

and there exist bounded domains $\mathcal{D}_i$ that satisfy condition (5) then the switched stochastic system (1) converges an ϵ-neighborhood of origin in finite time using control law (11).

Proof: Given bounded domains $\mathcal{D}_i$ that satisfy condition (5), the application of control (11) ensures that

$$\forall x(t_{i-1}) \in \mathcal{N}_{i-1}, \quad \mathbb{P}\{x(t_i) \in \mathcal{N}_i\} = 1 \ ,$$

which together with the assumption (19) implies that sequence $\{V(x(t_i))\}$ is strictly decreasing. In addition, $\{V(x(t_i))\}$ is lower bounded by zero and therefore it converges. Since the sequence is converging, the Cauchy criterion implies that as $i \to \infty$ and $\varepsilon_i \to 0$, the difference between two consecutive terms $\|V(x(t_i)) - V(x(t_{i-1}))\|$ converges. Now it is not possible that $\gamma(\|x_i\|) \to 0$, without $(\|x_i\|) \to 0$ as $\gamma(\cdot)$ is a class- $\mathcal{K}$ function. Which implies that $\lim_{i\to\infty,\varepsilon_i\to 0} V(x(t_i)) = 0$. Therefore $\{V(x(t_i)) \to 0\}$ as $i \to \infty$, and with $\gamma(\cdot)$ being a class-$\mathcal{K}$ function, $\|x_i\| \to 0$ as $i \to \infty$ and $\varepsilon_i \to 0$. The fact the convergence to a finite neighborhood of origin in finite time is ensured by $\mathbb{E}[t_{i+1} - t_i] < \infty$ and $\varepsilon_i > 0$. ■

If so desired, one can construct a stabilizing controller (in classical sense) [17] on an $n+1^{th}$ domain once the system reaches an ϵ-neighborhood of origin. The conditions (6)-(7) ensure the reachability of this neighborhood in finite time.

VIII. Conclusions and Future Work

The proposed method allows to design a navigation controller for systems governed by stochastic differential equations. If a feasible optimal path is given in the form of a finite sequence of waypoints, then an explicit solution for an optimal control law can be constructed, steering the system along these waypoints while avoiding the obstacles with probability one. This method can be applied to aerial or ground mobile robots subject to stochastic disturbances.

References

[1] A. T. Schwarm and M. Nikolaou, "Chance-constrained model predictive control," *AIChE Journal*, vol. 45, no. 8, pp. 1743–1752, 1999.

[2] "Robust model predictive control under chance constraints," *Computers and Chemical Engineering*, vol. 24, no. 2-7, pp. 829 – 834, 2000.

[3] D. Chatterjee, P. Hokayem, and J. Lygeros, "Stochastic receding horizon control with bounded control inputs: a vector space approach," *IEEE Transactions on Automatic Control*, vol. PP, no. 99, p. 1, 2011.

[4] E. Cinquemani, M. Agarwal, D. Chatterjee, and J. Lygeros, "On convex problems in chance-constrained stochastic model predictive control," *Automatica*, vol. To appear, 2011.

[5] L. Blackmore, M. Ono, A. Bektassov, and B. Williams, "A probabilistic particle-control approximation of chance-constrained stochastic predictive control," *IEEE Transactions on Robotics*, vol. 26, no. 3, pp. 502 –517, June 2010.

[6] L. Blackmore, M. Ono, and B. Williams, "Chance-constrained optimal path planning with obstacles," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1080 –1094, dec. 2011.

[7] H. J. Kappen, "An introduction to stochastic control theory, path integrals and reinforcement learning," *AIP Conference Proceedings*, vol. 887, no. 1, pp. 149–181, 2007.

[8] F. Weissel, M. Huber, and U. Hanebeck, "A nonlinear model predictive control framework approximating noise corrupted systems with hybrid transition densities," in *46th IEEE Conference on Decision and Control*, dec. 2007, pp. 3661 –3666.

[9] W. H. Fleming and H. M. Soner, *Controlled Markov Processes and Viscosity Solutions (Stochastic Modelling and Applied Probability)*, 2nd ed. Springer, Nov. 2005.

[10] M. Day, "On a stochastic control problem with exit constraints," *Applied Mathematics and Optimization*, vol. 6, pp. 181–188, 1980.

[11] R. Burridge, A. Rizzi, and D. Koditschek, "Sequential composition of dynamically dexterous robot behaviors," *The International Journal of Robotics Research*, vol. 18, pp. 534–555, 1999.

[12] F. Lamiraux and J.-P. Laumond, "Smooth motion planning for car-like vehicles," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 4, pp. 498–502, 2001.

[13] K. Pathak and S. Agrawal, "An integrated path-planning and control approach for nonholonomic unicycles using switched local potentials," *IEEE Transactions on Robotics*, vol. 21, no. 6, pp. 1201 – 1208, dec. 2005.

[14] M. Bujorianu and J. Lygeros, "Toward a general theory of stochastic hybrid systems," *Lecture Notes in Control and Information Sciences (LNCIS)*, vol. 337, pp. 3–30, 2006.

[15] I. Karatzas and S. E. Shreve, *Brownian Motion and Stochastic Calculus (Graduate Texts in Mathematics)*, 2nd ed. Springer, aug 1991.

[16] H. Tanner and J. Piovesan, "Randomized receding horizon navigation," *IEEE Transactions on Automatic Control*, vol. 55, no. 11, pp. 2640 –2644, nov. 2010.

[17] S. Shah, C. Pahlajani, and H. Tanner, "Probability of success in stochastic robot navigation with state feedback," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems 2011*, Sep. 2011.