Stochastic receding horizon control: application to an octopedal robot

Shridhar K. Shah[a] and Herbert G. Tanner[a]

[a]Department of Mechanical Engineering, University of Delaware, Newark, DE, USA

ABSTRACT

Miniature autonomous systems are being developed under ARL's Micro Autonomous Systems and Technology (MAST). These systems can only be fitted with a small-size processor, and their motion behavior is inherently uncertain due to manufacturing and platform-ground interactions. One way to capture this uncertainty is through a stochastic model. This paper deals with stochastic motion control design and implementation for MAST-specific eight-legged miniature crawling robots, which have been kinematically modeled as systems exhibiting the behavior of a Dubin's car with stochastic noise. The control design takes the form of stochastic receding horizon control, and is implemented on a Gumstix Overo Fire COM with 720 MHz processor and 512 MB RAM, weighing 5.5 g. The experimental results show the effectiveness of this control law for miniature autonomous systems perturbed by stochastic noise.

Keywords: legged robots, OctoRoACH, stochastic optimal control, receding horizon control

1. INTRODUCTION

This paper is motivated by the objective of controlling miniature bio-inspired crawling robots which are being developed under ARL's Micro Autonomous Systems and Technology (MAST) (Fig. 1(a)). The dynamics of these robots is inherently stochastic due to a variety of reasons, including size, manufacturing process, compliance, ground interactions, and battery charge fluctuations.[1, 2] Due to this stochastic uncertainty, many deterministic control approaches may not produce acceptable results. For miniature systems, such as one shown in Fig. 1(a), it is also desirable to have controllers that are input-optimal and computationally efficient, because of severe limitations in terms of payload and available energy storage capacity. Our goal in this work is to demonstrate application of input-optimal stochastic receding horizon controllers[3, 4] for autonomous navigation of these miniature robots.

The OctoRoACH platform (Fig. 1(a)), the specific robot used for this work, is developed at University of California, Berkley[1] under MAST. The controller on-board this platform can receive left-right motor pulse width modulation (PWM) signal commands through 802.15.4 radio and a lower-level cyclic rate control is implemented on-board using motor back EMF feedback.[1] The robot turns through differential steering. The turning rate is observably inconsistent due to the uncertainty in the mechanics of the robot and the variation in friction on different surfaces. To increase the turning rate, a dynamic turning control was demonstrated using a tail attachment.[1] Due to stochastic nature of this platform, a lot of effort has been on modeling of the stochastic kinematics[2, 5] and on the performance analysis of this platform on different surfaces.[6] There has also been work on different steering mechanisms for similarly constructed hexapods.[7, 8] To our knowledge, there does not exist any work on higher level control of this OctoRoACH platform, for deliberative navigation with obstacle avoidance.

For the purpose of achieving autonomous control of these 8-legged walking robots, one needs have a mathematical model. It was shown that under certain conditions, the dynamics of similar type of walking robots with differential drive mechanisms can be reasonably approximated in the form of a unicycle.[9] We model the OctoRoACH as a stochastic Dubin's car type model, as this OctoRoACH robot walks only in forward direction. Figure 1 shows a comparison of experimentally observed OctoRoACH trajectories for a straight, a left and a

---

Further author information: (Send correspondence to Herbert G. Tanner)
Shridhar K. Shah: E-mail: shridhar@udel.edu
Herbert G. Tanner: E-mail: btanner@udel.edu

right motion command with simulated trajectories of a stochastic Dubin's car equation. The system parameters are identified by observing a set of open loop trajectories. In addition to actual stochastic uncertainty of the platform, the mis-match of the model and actual robot is captured in the stochastic noise.

$$
\begin{bmatrix} \mathrm{d}x \\ \mathrm{d}y \\ \mathrm{d}\theta \end{bmatrix} = \begin{bmatrix} v\cos\theta \\ v\sin\theta \\ 0 \end{bmatrix} \mathrm{d}t + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \big(\omega\,\mathrm{d}t + \Sigma(\mathrm{q})\,\mathrm{d}W(t) \big)
$$

Equation of a stochastic Dubin's car

(a) OctoRoACH platform[1]

(b) Simulated stochastic differential equation (SDE)

(c) OctoRoACH experimental trajectories
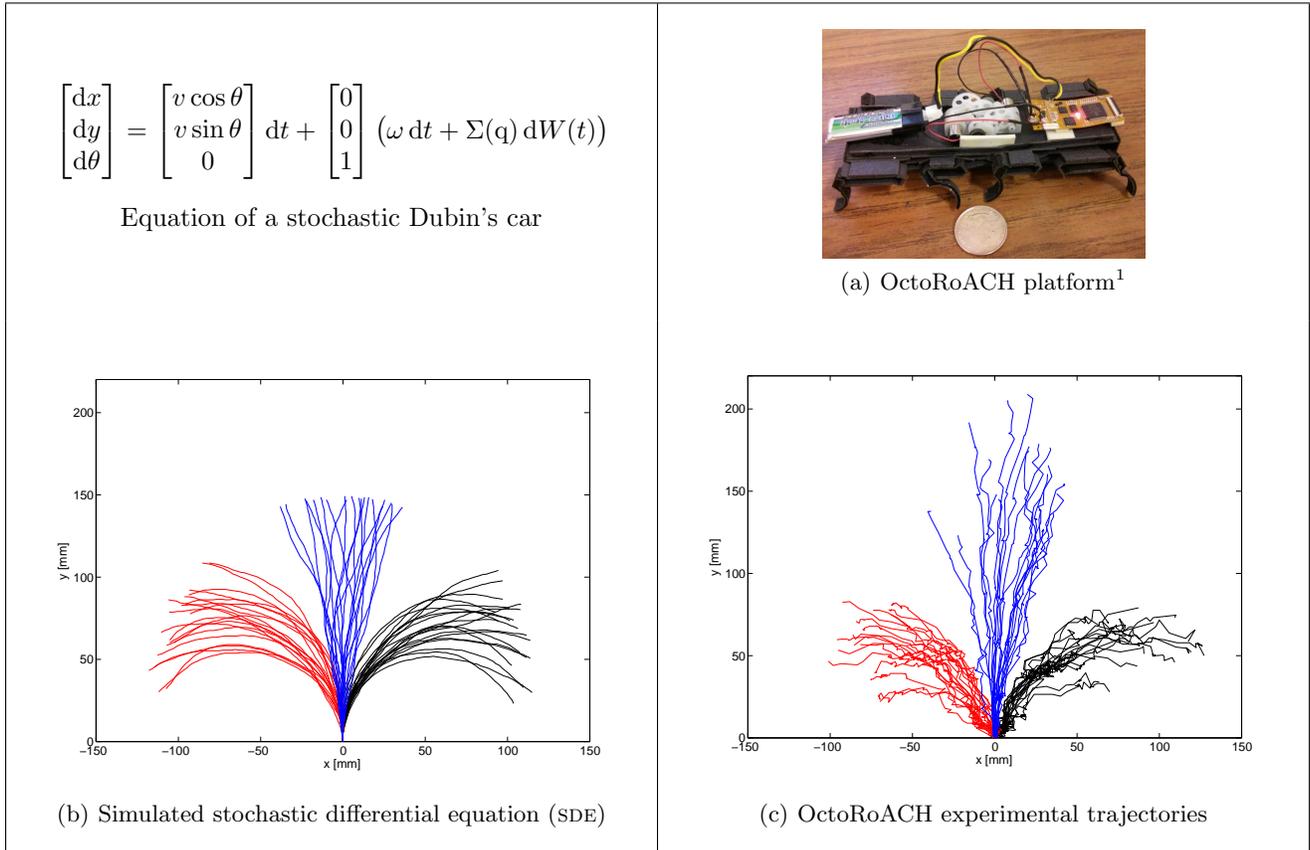
Figure 1. Comparison of simulated trajectories of a stochastic Dubin's car model with observed trajectories of an OctoRoACH platform.

1.1 Related Work and Scope

There exist deterministic control methods for Dubin's car type vehicles,[10] but it is unclear whether these control methods are robust under the effect of stochastic uncertainty. On the other hand, there has been recent efforts on directly considering stochastic uncertainty during control design.[3,4,11,12] For stochastic Dubin's car type vehicles, there are minimum-expected-time controllers[11] computed using the Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE), however obstacle avoidance is not addressed. A stochastic target tracking problem has also been considered for Dubin-type vehicles.[13] The application of obstacle avoidance for stochastic Dubin-type vehicles is considered within the problem of maximizing the probability of satisfying a temporal logic specification,[12] however this method produces an open loop trajectory and not a feedback motion plan. To construct a feedback motion plan capable of achieving obstacle avoidance and convergence to the goal, we apply solution of a local stochastic optimal control problem in a receding horizon manner. This is achieved by first computing a set of intermediate goals (or waypoints) that lead to the global goal set. The stochastic optimal controller is computed off-line and applied on-line to achieve convergence to each intermediate goal set.

For general nonlinear systems, there exist formulations of the stochastic optimal control problem.[14–17] For a specific stochastic optimal control problem, a logarithm transformation[14] can be used that linearizes the associated HJB PDE, and enables solutions using path integral methods.[16] Path integrals have been used for control design in applications ranging from reinforcement learning,[17] to variable stiffness control,[18] risk sensitive

control,[19] and more recently, receding horizon-based robot motion planning.[3,4] Although theoretically possible, global solutions of stochastic optimal control problem become intractable or inaccurate as the workspace becomes non-convex and large. Our receding horizon-based framework can provide sub-optimal solutions[3,4] even in case of large non-convex workspaces by using appropriate planning method. The theoretical foundations of stochastic optimal control with exit constraints[15] was used in a recursive fashion[3,4] to achieve probabilistic guarantees of reaching each intermediate goal sets. While the aforementioned work[3,4] applies stochastic receding horizon control to holonomic systems and tests it in simulation, in this paper, we demonstrate through experimentation, the application of the same method to nonholonomic stochastic Dubin's car type vehicles and achieve convergence to a goal while avoiding workspace boundary. The associated HJB PDE is computed off-line using a novel computation method based on an application of the Feynman-Kac Lemma[20] (see also[15,17]). The solution of the PDE yields a numerical potential field that steers the stochastic system away from the workspace boundary and toward the target set. The experiments were performed on a OctoRoACH platform while the receding horizon controller was implemented on a Gumstix computer-on-module (COM) with 720 MHz processor and 512 MB RAM, weighting 5.5 g. The objective was to travel a narrow wall like structure without hitting the boundary (which will result in fall) while achieving convergence to intermediate waypoints.

1.2 Organization

The paper is organized in the following way. Section 2 states the problem. It is followed by Section 3, which gives technical details of our input-optimal stochastic receding horizon control framework. Section 4 presents details of experimental evaluation of our controller on OctoRoACH platform followed by Section 5 which concludes this paper.

2. PROBLEM STATEMENT

Let $\mathcal{D} \subseteq \mathbb{R}^2 \times \mathbb{S}$ is the obstacle free workspace. Let $\mathcal{D}_i \subset \mathcal{D}$ be open bounded domains with each set having a closed set $\overline{\mathcal{E}}_i \subset \mathcal{D}_i$, where $i = 1, \ldots, N$. Assume that the boundary of $\mathcal{D}_i$, denoted $\partial \mathcal{D}_i$ and the boundary of $\overline{\mathcal{E}}_i$, denoted $\partial \mathcal{E}_i$, are described by a twice differentiable function (i.e., in $\mathcal{C}^2$), and consider a Dubin's car perturbed with stochastic noise evolving in $\mathcal{D}_i \setminus \overline{\mathcal{E}}_i$

$$\underbrace{\begin{bmatrix} \mathrm{d}x \\ \mathrm{d}y \\ \mathrm{d}\theta \end{bmatrix}}_{\mathrm{dq}} = \underbrace{\begin{bmatrix} v\cos\theta \\ v\sin\theta \\ 0 \end{bmatrix}}_{b(\mathrm{q})} \mathrm{d}t + \underbrace{\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}}_{G(\mathrm{q})} \left(\underbrace{\omega}_{u_i(\mathrm{q})} \mathrm{d}t + \Sigma(\mathrm{q})\,\mathrm{d}W(t) \right) \tag{1}$$

where $\mathrm{q} = [x, y, \theta]^\intercal \in \mathbb{R}^2 \times \mathbb{S}$ is the state, v is a fixed linear velocity, ω is the angular velocity input, $b(\mathrm{q}) : \mathbb{R}^2 \times \mathbb{S} \to \mathbb{R}^3$ is the uncontrolled drift term, $G(\mathrm{q}) : \mathbb{R}^2 \times \mathbb{S} \to \mathbb{R}^{3\times 1}$ is the control matrix, $u_i(\mathrm{q}) : \mathbb{R}^2 \times \mathbb{S} \to \mathbb{R}^1$ is the control input, and $\Sigma(\mathrm{q}) : \mathbb{R}^2 \times \mathbb{S} \to \mathbb{R}^1$ is a state-dependent variance. Let $W = \{W(t), \mathcal{F}(t) : 0 \leq t < \infty\}$ model a 1-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 a probability measure, and $\{\mathcal{F}(t) : t \geq 0\}$ is a filtration (i.e., an increasing family of sub-σ-algebras of $\mathcal{F}$), assumed right-continuous and such that $\mathcal{F}_0$ contains all $\mathbb{P}$-null (of measure zero) sets.[20]

PROBLEM 1 (STOCHASTIC RECEDING HORIZON CONTROL). *Find a sequence of feedback control laws* $\{u_i(\mathrm{q})\}_{i=1}^N$ *for each set* $\mathcal{D}_i$ *for system (1), such that if given a set of waypoints* $\{\mathrm{q}_i\}_{i=1}^N, \mathrm{q}_0 = \mathrm{q}(0),\ \mathrm{q}_i \in \mathcal{D}_i \cap \mathcal{D}_{i+1}$ *the application of* $\{u_i(\mathrm{q})\}$ *to (1) results in sample paths* $\mathrm{q}(t)$ *that achieve*

(i) $\mathbb{P}\left[\mathrm{q}(t) \in \bigcup_{i=1}^N \mathcal{D}_i\right] = 1,\ \forall t > 0$ *(almost-sure safety);*

(ii) $\mathbb{P}\left[\exists\, t_i < \infty : \|\mathrm{q}_i - \mathrm{q}(t_i)\| < \varepsilon,\ i = 1 \ldots N\right] = 1$ *(almost-sure convergence with accuracy* $\varepsilon > 0$*);*

(iii) $\mathbb{E}\left[\int_{t_{i-1}}^{t_i} L\big(\mathrm{q}(s), u_i(s)\big)\,\mathrm{d}s + \Phi\big(\mathrm{q}(t_i)\big)\right]$ *is minimized, where* t_{i-1} *and* t_i *are the first times* $\mathrm{q}(t)$ *enters an* ε*-neighborhood of* q_{i-1} *and* q_i *($t_i - t_{i-1} = \tau_i$), respectively, and* $\Phi(\mathrm{q}) : \mathbb{R}^n \to \mathbb{R}_+$ *is a terminal cost function (local optimality).*

In other words, the problem is to find N intermediate goal sets, and then design local stochastic optimal controllers acting on small sub-domains of the workspace defined as $\mathcal{D}_i$, each containing previous and current intermediate goal set, to make the system transition in finite time from one goal set to the next. The controllers minimize a cost functional which penalizes input effort and provide a probabilistic guarantee of convergence at an intermediate goal set in finite time.

3. STOCHASTIC RECEDING HORIZON CONTROL DESIGN

The following section describes the main mathematical framework of stochastic receding horizon control. We refer to our previous work[3, 4] for more rigorous treatment of this framework.

3.1 Planning & waypoint generation

The stochastic receding horizon algorithm begins by computing an obstacle free reference path which can be re-planned at specific time intervals. The reference path is computed without consideration of noise and hence deterministic planning methods can be used. In our previous work[3, 4] we have used a receding horizon approach[21] to obtain such a reference path, however a multitude of alternative methodologies, including potential field methods,[22] rapidly exploring random trees RRTs,[23] or cell decomposition methods[10] can also be used. The planning problem is assumed to be solved here, and for the purpose of this paper, the path planning stage is not described. Instead, we assume that through one of the above methods we are able to determine N successive waypoints leading to the goal $\mathbf{q}_N$ where, $\mathbf{q}_0 = \mathbf{q}(0)$. For waypoints $\mathbf{q}_i = [x_i, y_i, \theta_i]^\intercal$ we define the intermediate goal sets as

$$\overline{\mathcal{E}}_i \triangleq \{\mathbf{q} \in \mathcal{D}_i : \sqrt{(x - x_i)^2 + (y - y_i)^2} \leq \varepsilon, \forall \theta, \theta_i \in \mathbb{S}\}, i = 0, \ldots, N .$$

We want to design controller u_i, that can drive the system within domain $\mathcal{D}_i$ and ensure convergence to $\overline{\mathcal{E}}_i$, and the controller is switched to u_{i+1} on the state hitting the boundary $\partial \mathcal{E}_i$. This requires the condition that $\overline{\mathcal{E}}_i, \overline{\mathcal{E}}_{i+1} \in \mathcal{D}_i,\ i = 0, \ldots, N-1$. Moreover, we also impose a condition that the intermediate goal sets are disjoint, which leads to a condition on waypoints, $\sqrt{(x_{i-1} - x_i)^2 + (y_{i-1} - y_i)^2} > 2\varepsilon,\ i = 1, \ldots, N$. We define domains $\mathcal{D}_i \subset \mathcal{D}$ around these waypoints that satisfy $\overline{\mathcal{E}}_i, \overline{\mathcal{E}}_{i+1} \in \mathcal{D}_i$. It is important to remember that definition of domains $\mathcal{D}_i$ and waypoints can be a parallel process if there are constraints of keeping domains $\mathcal{D}_i$ free of obstacles as discussed in our previous work.[3, 4] Figure 2 shows an illustrative example of definition of domains $\mathcal{D}_i$ and intermediate goal sets $\overline{\mathcal{E}}_i$ around waypoints $\mathbf{q}_i$. Once the set of waypoints are determined, a set of stochastic optimal controllers are used in a recursive way to achieve convergence to each waypoint $\mathbf{q}_i$. The stochastic optimal controllers are computed as described in the following section.
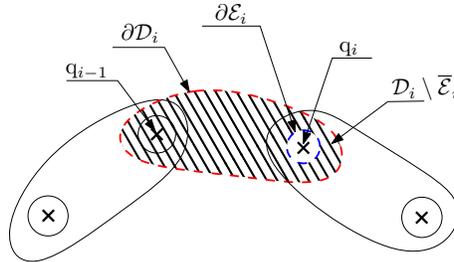
Figure 2. Illustration showing domains $\mathcal{D}_i$ and intermediate goal set $\overline{\mathcal{E}}_i$.

3.2 Stochastic optimal controllers

The system state is a Markov process $\mathbf{q}(t)$ that evolves between way-points according to the SDE

$$\mathrm{d}\mathbf{q}(t) = b\big(\mathbf{q}(t)\big)\,\mathrm{d}t + G\big(\mathbf{q}(t)\big)\big[u_i\big(\mathbf{q}(t)\big)\,\mathrm{d}t + \Sigma\big(\mathbf{q}(t)\big)\,\mathrm{d}W(t)\big]$$
$$\text{given } \mathbf{q}(0) = \mathbf{q}_0 . \tag{2}$$

where we assume that $b(\mathbf{q})$, $G(\mathbf{q})$, $\Sigma(\mathbf{q})$, and $\Sigma^{-1}(\mathbf{q})$ are bounded and Lipschitz continuous on $\mathcal{D}_i$ and together with u_i, are all bounded in $\mathcal{D}_i$. The latter is the control input responsible for taking the state from $\partial \mathcal{E}_{i-1}$ to $\partial \mathcal{E}_i$ while avoiding $\partial \mathcal{D}_i$. Let t_i is the first time $\mathbf{q}(t)$ hits the boundary of $\partial(\mathcal{D}_i \setminus \overline{\mathcal{E}}_i)$

When (2) under u_i hits $\partial\mathcal{E}_i$ at some time t_i, it undergoes a forced transition with u_i switching to u_{i+1}, and the switch occurs upon the state hitting a part of the boundary $\partial\mathcal{E}_i$. Control law u_i gives a solution to the stochastic optimal control problem

$$\min_{u_i} \mathbb{E}\left[\int_{t_{i-1}}^{t_i} \frac{1}{2}u_i^\mathsf{T}(q(s))\, a^{-1}(q(s))\, u_i(q(s))\, \mathrm{d}s + \Phi\big(q(t_i)\big) \,\Big|\, q(t_{i-1}) = q\right] =: V_i(q) \tag{3}$$

where $a(q) = \Sigma(q)\Sigma^\mathsf{T}(q)$ and the terminal cost Φ is chosen in theory as

$$\Phi = \begin{cases} 0 & \text{on } \partial\mathcal{E}_i \\ \infty & \text{on } \partial\mathcal{D}_i \end{cases} \quad.$$

This terminal cost Φ imposes an infinite penalty for hitting the boundary of $\mathcal{D}_i$. The HJB equation associated with (3) is[24]

$$\min_{u_i(q)} \left\{ \mathcal{A}V_i(q) + \frac{1}{2}u_i^\mathsf{T}(q)a^{-1}(q)u_i(q) \right\} = 0 \tag{4}$$

where $\mathcal{A}$ is the second-order partial differential operator

$$\mathcal{A} \triangleq \big(b(q) + G(q)u_i(q)\big)^\mathsf{T}\frac{\partial}{\partial q} + \frac{1}{2}\mathrm{tr}\left[a(q)\frac{\partial^2}{\partial q^2}\right] \quad.$$

The optimal control law u_i^* that solves (4) is

$$u_i^*(q) = -a(q)G^\mathsf{T}(q)\,\frac{\partial V_i(q)}{\partial q} \quad. \tag{5}$$

Substituting (5) in (4) and applying the logarithmic transformation[14] $V_i(q) = -\log g_i(q)$, yields

$$b^\mathsf{T}(q)\frac{\partial g_i(q)}{\partial q} + \frac{1}{2}\mathrm{tr}\left\{\frac{\partial^2 g_i(q)}{\partial q^2}G(q)\,\Sigma(q)\,\Sigma^\mathsf{T}(q)\,G^\mathsf{T}(q)\right\} = 0 \tag{6}$$

with boundary conditions

$$\begin{cases} g_i\big(q(t_i)\big) = 1 & \text{if } q(t_i) \in \partial\mathcal{E}_i \\ g_i\big(q(t_i)\big) = 0 & \text{if } q(t_i) \in \partial\mathcal{D}_i \end{cases} \quad.$$

Using the Feynman-Kac formula,[20] the solution of (6) becomes

$$g_i(q) = \mathbb{P}\left[\zeta_i(\tau) \in \partial\mathcal{E}_i \mid \zeta_i(0) = q\right] \tag{7}$$

where $\zeta_i(t)$ is the Markov process

$$\mathrm{d}\zeta_i(t) = b\big(\zeta_i(t)\big)\,\mathrm{d}t + G\big(\zeta_i(t)\big)\,\Sigma\big(\zeta_i(t)\big)\,\mathrm{d}W(t) \tag{8}$$

evolving on $\mathcal{D}_i \setminus \overline{\mathcal{E}}_i$ and τ is the first exit time for process $\zeta_i(t)$.

The optimal control (5) can thus be computed as

$$u_i^*(q) = -a(q)G^\mathsf{T}(q)\,\frac{\partial}{\partial q}\log\mathbb{P}\left[\zeta_i(\tau) \in \partial\mathcal{E}_i \mid \zeta_i(0) = q\right] \quad. \tag{9}$$

Since there is infinite penalty for hitting the boundary of $\mathcal{D}_i$, if an optimal control law exists, it ensures that the system exits at the goal boundary $\overline{\mathcal{E}}_i$ with probability one. This exit occurs in finite time, i.e., $\mathbb{E}[t_i \mid q(t_{i-1}) = q] < \infty \ \forall q \in \mathcal{D}_i \setminus \overline{\mathcal{E}}_i$, if [20, Lemma 7.4]

$$\min_{q\in\mathcal{D}_i\setminus\overline{\mathcal{E}}_i} a_{ll}(q) > 0 \tag{10}$$

for some $1 \leq l \leq m$. The assumption that Σ and Σ^{-1} are bounded, ensures a_{ll} is nonzero and positive for some $1 \leq l \leq m$, hence satisfies (10). Assuming that an optimal control input u_i^* exists, imposing infinite penalty on exit through $\partial \mathcal{D}_i$ is equivalent to a constraint (see[15])

$$\mathbb{P}\big[\mathrm{q}(\tau) \in \partial \mathcal{D}_i \mid \mathrm{q}(0) = \mathrm{q}\big] = 0 \ .$$

The above constraint means that the stochastic optimal controller ensures that the robot avoids obstacles and converges to the goal set with probability one in finite time; hence it satisfies all the requirements of problem statement, namely, almost-sure safety, almost-sure convergence in finite time and minimization of cost. The caveat is that infinitely large inputs may be required arbitrarily close to the workspace or obstacle boundary ($\partial \mathcal{D}_i$). Since, realistically, control inputs are always bounded, depending on the magnitude of noise and control input bounds, the true probability of success of convergence to goal is actually smaller—and can be estimated using methods in.[25]

Given a sequence of waypoints $\{\mathrm{q}_i\}_{i=0}^N$, the process of transitioning from waypoint q_{i-1} to waypoint q_i under (9) is repeated. By the time a new waypoint is reached, the reference path can be recomputed in a receding horizon manner, and the waypoint sequence $\{\mathrm{q}_i\}_{i=0}^N$ can be redefined with the initial element q_0 being the waypoint just reached. What is important for real-time implementation is that for predetermined domains $\mathcal{D}_i$, the optimal control law u_i^* (9) can be precomputed off-line.

4. EXPERIMENTAL EVALUATION

Figure 3 shows the experimental setup used in this paper. Figure 3(a) shows the coordinate system and OctoRoACH dimensions, while Fig. 3(b) shows the workspace boundary in thick black line. For this experiment, it is required to have a map of workspace, which can be either observed using onboard sensors such as a camera or laser or it can be known a priori. Here it is assumed that the map was provided a priori, and a sequence of waypoints with initial condition $\{\mathrm{q}_i\}_{i=0}^5, q_0 = q(0)$ is given. Figure 3(b) also shows the domains $\mathcal{D}_i$ constructed around these waypoints along with the intermediate goal sets $\overline{\mathcal{E}}_i$.

(a) Coordinate system (b) Experimental setup (c) Experimental setup
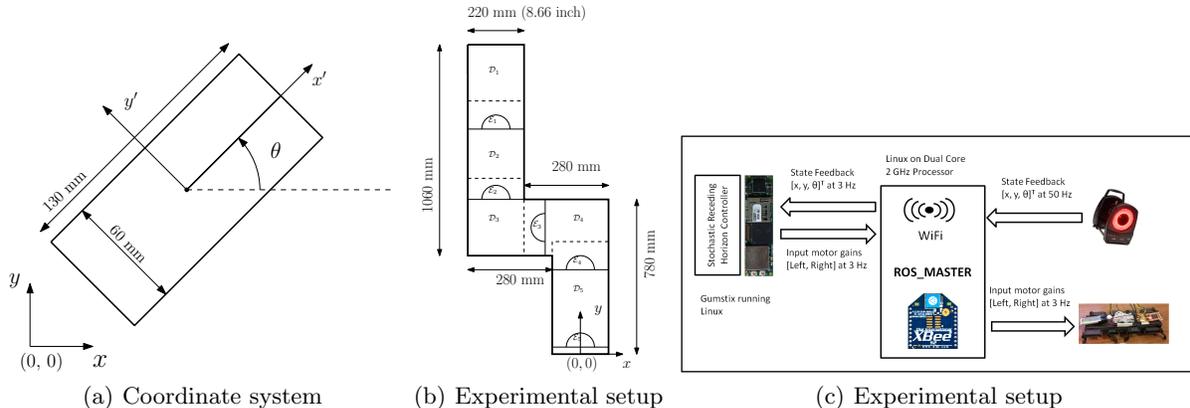
Figure 3. (a) Platform dimensions and the coordinate system (b) Workspace divided into domains $\mathcal{D}_i$ and intermediate goal sets $\overline{\mathcal{E}}_i$, dotted lines show overlapping boundary of the next domain (c) Experimental setup showing different components involved. The complete experiment was performed using ROS (Robot Operating System).

Two parameters in the stochastic Dubin's car model (1) are required to be identified, which are linear velocity v and variance Σ. These parameters are approximated based on open loop observed trajectories. We approximated the constant linear velocity in our model to be 20 mm/s and constant variance $\Sigma = 0.25$ mm/s. With these model parameters, the stochastic optimal controller is computed in an off-line manner by simulating (8) and computing $g(\mathrm{q})$ using (7), which is part of the control input. A grid is created on the set $\mathcal{D}_i \setminus \overline{\mathcal{E}}_i$ and the process (8) is simulated at each grid point for 500 sample paths. Based on the outcome of these sample paths, $g(\mathrm{q})$ is computed using (7) for each grid point. Finally, the control input is computed by taking numerical gradient of $g(\mathrm{q})$. The optimal control input is computed off-line for a set of bounded domains with different objectives

and boundary definitions, which can be used repetitively. We identified three different shapes of domain $\mathcal{D}_i$, and corresponding functions $g(\text{q})$ are shown in Fig. 4. This figure shows a three slices in a four-dimensional data space, with color representing the forth dimension—the value of the function $g(\text{q})$. These controllers are used in a receding horizon manner to achieve each waypoint as an intermediate goal.

(a) Straight controller   (b) Left turn controller   (c) Right turn controller
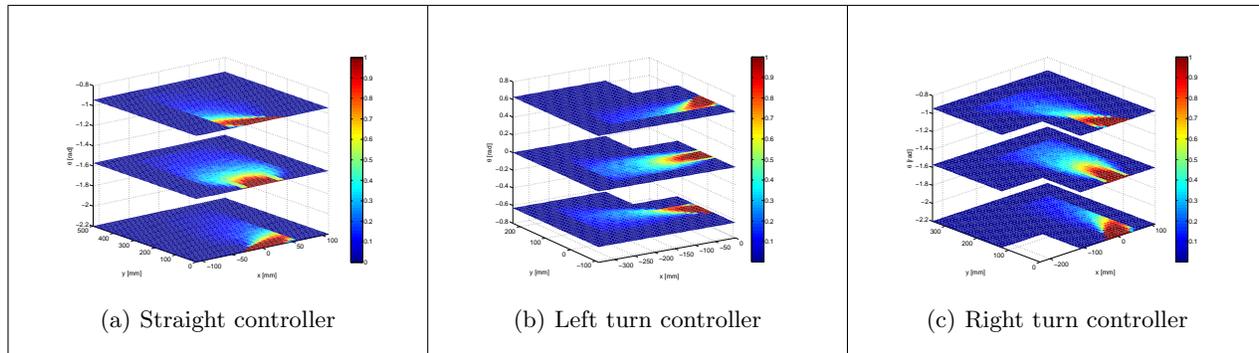
Figure 4. The function $g(\text{q})$ for three different domain shapes. Function $g(\text{q})$ was computed off-line for each of these domain. Value of $g(\text{q})$ is shown using a color-map at three different angles for each of the controllers.

Figure 3(c) shows the flow of information for the experimental setup. The stochastic receding horizon controller was executed on a Gumstix Overo Fire COM board running at 720 MHz with 500 MB RAM. While linear velocity is fixed, the stochastic controller generated optimal angular velocity commands which were then converted into six discrete motion commands to be executed by OctoRoACH as shown in Table 4. A 2 GHz

Table 1. Continuous angular velocity commands mapping to discrete input motor gains.

Command Name	Required angular velocity (ω [rad/s])	Applied input motor gains
Stop	0	[0, 0]
Straight	$-0.0625 \leq \omega \leq 0.0625$	[80,80]
Clockwise Turn	$-0.125 \leq \omega \leq -0.0625$	[80,40]
Sharp Clockwise Turn	$\omega \leq -0.125$	[90, 0]
Counter-clockwise Turn	$0.125 \geq \omega \geq 0.0625$	[40,80]
Sharp Counter-clockwise Turn	$\omega \geq 0.125$	[90, 0]

linux machine was used as a communication server. The state feedback was obtained at 50 Hz through a Vicon® motion capture system; however the stochastic controller used feedback only at 3 Hz due to constraints imposed by the robot's firmware, although the Gumstix processor is able to execute stochastic controller at a frequency higher than 20 Hz. For the given five waypoints, we concatenated controllers in the sequence {Straight Controller, Straight Controller, Left Controller, Right Controller, Straight Controller} and the resulting path of the OctoRoACH is shown in Fig. 5(a) along with the input motor gains applied to the robot. Figures 5(b) and 5(c) shows the time trajectories for $x-y$ and θ components. The goal was to reach the final waypoint $[x, y, \cdot]^{\mathsf{T}} = [0, 50, \cdot]^{\mathsf{T}}$ mm with any orientation.

The theoretical framework of stochastic receding horizon control provides probability one of reaching goal set without hitting the outer (or obstacle) boundary. However, due to various approximations (state-time discretization, input discretization, approximate linear velocity and noise variance) used for experimental realization of the controller, the actual probability of reaching goal set will be less than one. However, depending on the available control authority and size of workspace, it is possible to get realization of probability of reaching goal close to one.

5. CONCLUSION

We have shown that the stochastic receding horizon framework[3, 4] described in this paper can be applied for control of miniature bio-inspired robots that exhibit inherent stochastic behavior. Existing methods on path

(a) Path of OctoRoACH and input motor gains applied

(b) $x - y$ trajectory
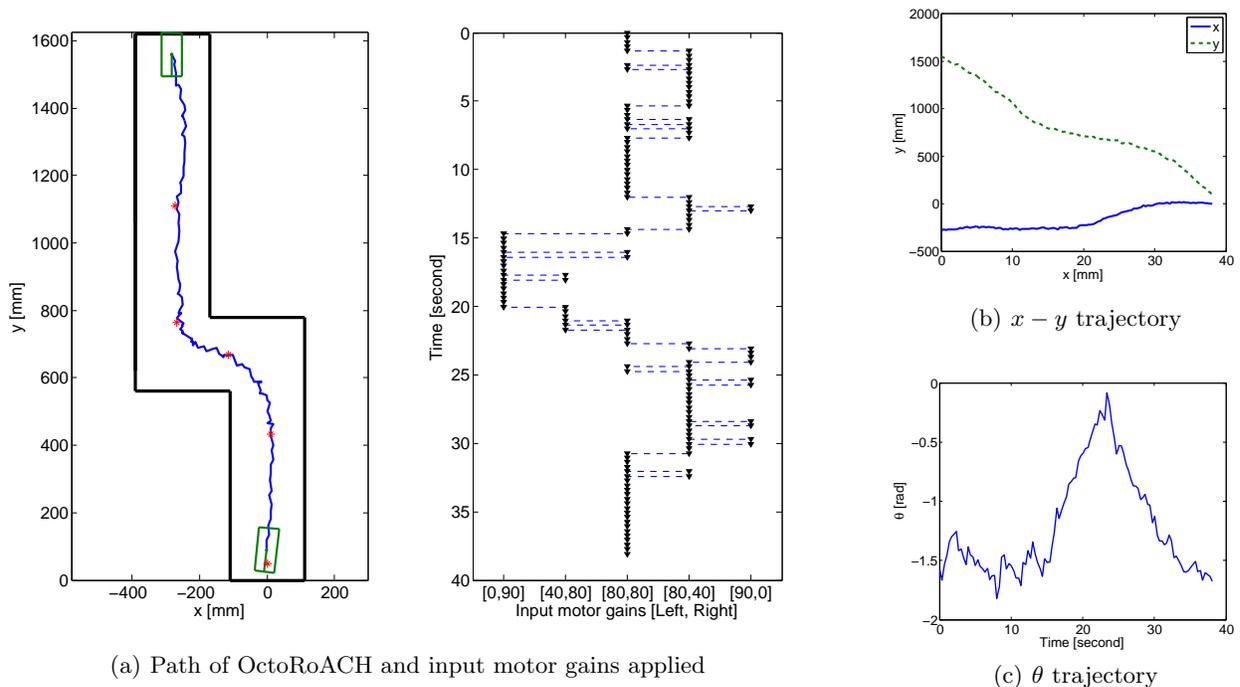
(c) θ trajectory

Figure 5. Observed experimental data (a) The left figure is the observed path superimposed on the workspace. Black solid lines is the outer boundary which is to be avoided at all time, blue line is the observed $x - y$ trajectory and red $*$ are the waypoints used within the stochastic receding horizon control. The figure on the right shows input motor gains applied during the experiment.(b) Observed $x - y$ trajectory for the experiment with goal $x = 0$ mm, $y = 50$ mm (c) Observed θ trajectory of the experiment.

planning can be combined with stochastic controllers (which are computed off-line) to create a switching strategy and achieve convergence to the goal in a receding horizon manner. Moreover the framework provides theoretical guarantee of convergence with probability one.

ACKNOWLEDGMENTS

REFERENCES

[1] Pullin, A., Kohut, N., Zarrouk, D., and Fearing, R., "Dynamic turning of 13 cm robot comparing tail and differential drive," in [*Proceedings of the IEEE International Conference on Robotics and Automation*], 5086–5093 (2012).

[2] Karydis, K., Poulakakis, I., and Tanner, H. G., "Probabilistic validation of a stochastic kinematic model for an eight-legged robot," in [*Proceedings of the IEEE International Conference on Robotics and Automation*], (to appear) (2013).

[3] Shah, S., Pahlajani, C., Lacock, N., and Tanner, H., "Stochastic receding horizon control for robots with probabilistic state constraints," in [*Proceedings of the IEEE International Conference on Robotics and Automation*], 2893 –2898 (2012).

[4] Shah, S. K., Tanner, H. G., and Pahlajani, C. D., "Stochastic receding horizon control of nonlinear stochastic systems with probabilistic state constraints," *arXiv:1211.4038 [cs.SY]* (2012).

[5] Karydis, K., Poulakakis, I., and Tanner, H., "A switching kinematic model for an octapedal robot," in [*IEEE/RSJ International Conference on Intelligent Robots and Systems*], 507–512 (2012).

[6] Bermudez, F., Julian, R., Haldane, D., Abbeel, P., and Fearing, R., "Performance analysis and terrain classification for a legged robot over rough terrain," in [*IEEE/RSJ International Conference on Intelligent Robots and Systems*], 513–519 (2012).

[7] Zarrouk, D. and Fearing, R., "Compliance-based dynamic steering for hexapods," in [*IEEE/RSJ International Conference on Intelligent Robots and Systems*], 3093–3098 (2012).

[8] Hoover, A., Burden, S., Fu, X.-Y., Sastry, S., and Fearing, R., "Bio-inspired design and dynamic maneuverability of a minimally actuated six-legged robot," in [*3rd IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*], 869–876 (2010).

[9] Panagou, D. and Tanner, H., "Modeling of a Hexapod Robot; Kinematic Equivalence to a Unicycle," tech. rep., University of Delaware, Department of Mechanical Engineering (04 2009).

[10] LaValle, S. M., [*Planning Algorithms*], Cambridge University Press, Cambridge, U.K. (2006).

[11] Anderson, R. P., Bakolas, E., Milutinovic, D., and Tsiotras, P., "The markov-dubins problem in the presence of a stochastic drift field," in [*Proceedings of the IEEE Conference on Decision and Control*], 130–135 (2012).

[12] Cizelj, I. and Belta, C., "Probabilistically safe control of noisy dubins vehicles," in [*Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*], 2857 –2862 (2012).

[13] Anderson, R. and Milutinovic, D., "A stochastic approach to dubins feedback control for target tracking," in [*IEEE/RSJ International Conference on Intelligent Robots and Systems*], 3917–3922 (2011).

[14] Fleming, W. H., "Exit probabilities and optimal stochastic control," *Applied Mathematics and Optimization* **4**, 329–346 (1977).

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

[16] Kappen, H. J., "Path integrals and symmetry breaking for optimal control theory," *Journal of Statistical Mechanics: Theory and Experiment* **2005**(11), P11011 (2005).

[17] Theodorou, E., Stulp, F., Buchli, J., and Schaal, S., "Iterative path integral stochastic optimal control for learning robotic tasks," in [*The 18th World Congress of The International Federation of Automatic Control, Milan, Italy*], (2011).

[18] Buchli, J., Stulp, F., Theodorou, E., and Schaal, S., "Learning variable impedance control," *The International Journal of Robotics Research* **30**(7), 820–833 (2011).

[19] van den Broek, B., Wiegerinck, W., and Kappen, B., "Stochastic optimal control of state constrained systems," *International Journal of Control* **84**(3), 597–615 (2011).

[20] Karatzas, I. and Shreve, S. E., [*Brownian Motion and Stochastic Calculus*], Springer, 2nd ed. (1991).

[21] Tanner, H. and Piovesan, J., "Randomized receding horizon navigation," *IEEE Transactions on Automatic Control* **55**(11), 2640 –2644 (2010).

[22] Rimon, E. and Koditschek, D., "Exact robot navigation using artificial potential functions," *IEEE Transactions on Robotics and Automation* **8**(5), 501 –518 (1992).

[23] Lavalle, S. M., Kuffner, J. J., and Jr., "Rapidly-exploring random trees: Progress and prospects," in [*Algorithmic and Computational Robotics: New Directions*], 293–308 (2000).

[24] Fleming, W. H. and Soner, H. M., [*Controlled Markov Processes and Viscosity Solutions*], Springer, 2nd ed. (2005).

[25] Shah, S., Pahlajani, C., and Tanner, H., "Probability of success in stochastic robot navigation with state feedback," in [*Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*], 3911–3916 (2011).