A SAMPLING APPROACH TO MODELING AND CONTROL OF A LARGE-SIZE ROBOT POPULATION

Dejan Milutinović *
Applied Mathematics and Statistics Department
University of California
Santa Cruz, California 95064
Email: dejan@soe.ucsc.edu

Devendra P. Garg
Mechanical Engineering and Materials Science Department
Duke University
Durham, North Carolina, 27708
dpgarg@pratt.duke.edu

ABSTRACT

Motivated by the close relation between estimation and control problems, we explore the possibility to utilize stochastic sampling for computing the optimal control for a large-size robot population. We assume that the individual robot state is composed of discrete and continuous components, while the population is controlled in a probability space. Utilizing a stochastic process, we can compute the state probability density function evolution, as well as use the stochastic process samples to evaluate the Hamiltonian defining the optimal control. The proposed method is illustrated by an example of centralized optimal control for a large-size robot population.

NOMENCLATURE

$\mathbb{R}, \mathbb{R}^n$ set of real numbers, set of real number vectors of dimension n.
Q set of discrete states, i.e., integer indexes $\{1,2,3,\ldots\}$
U_{ad} set of admissible control.
$\rho(x,t)$ probability density function (PDF) of the hybrid state at time t. This variable is a vector of functions; it depends on $x \in X$ and $t \in \mathbb{R}$, but x and t are frequently omitted in expressions.
$\rho_i(x,t)$ the PDF component corresponding to the discrete state i, $i \in Q$.
$\pi(x,t)$ the adjoint state distribution.
$\bar{\pi}$ the discrete approximation of the adjoint state distribution.
$\phi(x,t)$ the adjoint state PDF.
$P_i(x,t)$ the probability of the discrete state i, $i \in Q$.
$P_i^{\pi}(x,t)$ the probability of the discrete adjoint state i, $i \in Q$.
F_t the transition rate matrix.
λ_{ij} the component of the transition rate matrix $[F_t]_{ij} = \lambda_{ij}$.
F_u the transition rate matrix that depends on control vector u, $u \in U_{ad}$.
F_∂ the component of the linear operator F corresponding to the vector fields f_i of discrete states $i \in Q$.
$H(\rho,u,t)$ the PDF, the control and time dependent Hamiltonian, ρ, u and t are frequently omitted.
u^* the optimal control
E_ρ the expected value with respect to the state PDF ρ.

INTRODUCTION

Solutions of multi-robot control problems may be of enormous complexity because of the operating environment uncertainties, or a large number of redundant states and robots. For many years it has been known that the optimal control and optimal estimation problems are closely related [1]. For example, the linear quadratic regulator (LQR) and the Kalman filter (KF) estimator can be derived in the same optimal control framework. Having this in mind, it is expected that estimation methods based on statistical sampling, and employed for solving complex estimation problems [2], can contribute to solving complex control problems for a single and, more importantly, multi-robot systems under the presence of uncertainty.

Along this idea, Kappen et. al. [3, 4] used stochastic differential equations to model individual agents. Based on this description, it is possible to relate the Hamilton-Jacobi-Bellman partial differential equation with samples of the stochastic process trajectories and use the samples to define the stochastic optimal control of multi-agent systems. In this framework, the state of individual robots is continuous. However, the state of real robots is generally described by a combination of contin-

*Address all correspondence to this author.

uous and discrete variables, i.e., by a hybrid state. Therefore, it is more natural to describe the robot behavior using a hybrid automaton [5]. The automaton describes the discrete and continuous variables change in time, which depends on events influencing the robot behavior. In the case of a large-size multi-robot population, it becomes highly complex to predict events from the robot local environment. Because of that, we model a large-size robot population considering the stochastic hybrid model and study how it can be controlled.

In this paper, we consider a problem in which the presence of a large-size robot population in a desired region of operating space is maximized. This problem is formulated in a hybrid system framework in [6]. Its solution, based on the minimum principle for partial differential equations, is presented in [7, 8], and it is solved numerically when the presence of the robots is maximized along one dimension (1D).

The Hamiltonian, which defines the optimal control, includes integral terms that depend on the solution of a system of partial differential equations (PDE). This system of PDEs is in general difficult to evaluate and the numerical evaluation of integrals is prone to errors. However, we recognize that the problem solution can be simplified and propose to use samples of the stochastic processes to evaluate the Hamiltonian components from the expected values of the adjoint state distribution.

The direction of the research we are pursuing is considerably different from the stochastic optimal control work presented in [9]. There, stochastic processes have been used only as an analytical tool to map the stochastic process to be controlled into the finite state space, in which the optimization is performed. The benefit of using a solution based on sampling, i.e, computational statistical methods, is that control problems in robotics could be solved faster. This possibility also depends on the ability to implement sampling and computations with samples into the processor computing the control.

MODELING AND CONTROL FRAMEWORK

In the modeling framework we consider, the state of an individual robot at time t is uniquely defined by the couple $(x(t), q(t))$, $x \in X$, $X \in \mathbb{R}^n$, $q \in Q, Q = \{1, 2, ...N\}$. While in the discrete state (mode) $k \in Q$, the continuous state of a robot obeys the differential equation $\dot{x} = f_k(x,t)$. We also assume that switching among the discrete states, say from the state $k \in Q$ to the state $j \in Q$, $(k \neq j)$, is described by stochastic transition rates λ_{kj}, and that $x(t)$ is a continuous function of time. In other words, the continuous state just before the discrete state transition $x(t^-)$ is equal to the state $x(t^+)$ after the state transition. This very general model of an individual robot is illustrated in Fig. 1 and the modeling framework we are applying here is detailed in [8].

Recognizing that the state of an individual robot is composed of discrete and continuous components, the state probability density function (PDF) is a vector of functions $\rho(x,t) = [\rho_1(x,t)\ \rho_2(x,t) \ldots \rho_N(x,t)]'$ [8]. Each component $\rho_i(x,t)$ corresponds to the discrete state i, and the symbol (') denotes the vector transpose. The state PDF satisfies

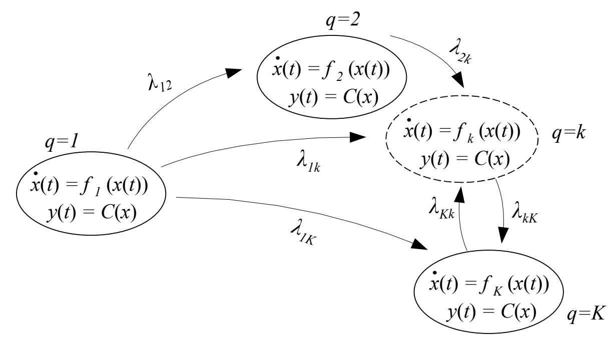

Figure 1. Stochastic hybrid automaton model of a robot in a probabilistic framework: discrete state q; continuous state x vector field f_k, $k \in Q$ describes the change of the continuous state; stochastic transition rates $\lambda_{kj}, k, j \in Q$ describe the mode switching; y is the measurable output; if the full continuous state of the robot is measurable, C is the unity matrix.

$$\sum_{i\in Q}\int_X \rho_i(x,t)dx = \sum_{i\in Q} P_i(t) = 1, \text{ where } P_i(t) = \int_X \rho_i(x,t)dx \quad (1)$$

where $P_i(t)$ is the probability of the discrete state i at the time point t. Let us define the vector of discrete probabilities $P(t) = [P_1(t), P_2(t), \ldots P_N(t)]'$, then evolution of the probability vector is given by:

$$\dot{P}(t) = F_t(t)P(t), \text{ where } [F_t]_{ij} = \lambda_{ij}(t) \quad (2)$$

with matrix F_t defining the transition rates among the discrete states. In general, the correspondence between the matrix F_t members $[F_t]_{ij}$ and the transition rates λ_{ij} is not one-to-one. Assuming that the transition rates depend on the vector $u(t) = [u_1(t)\ u_2(t)\ \ldots u_M(t)]'$ of variables $u_i, i = 1, 2\ldots M$, we can define the transition rate matrix as a function of the vector $u(t)$, i.e., $F_t(t) = F_u(u(t))$. Consequently, the vector of the discrete state probabilities obeys [10]:

$$\dot{P}(t) = F_u(u)P(t) \quad (3)$$

Moreover, it can be proven [8] that the state PDF obeys the following system of partial differential equations (PDE):

$$\frac{\partial \rho(x,t)}{\partial t} = F(u)\rho(x,t) = (F_u(u(t)) + F_\partial)\rho(x,t) \quad (4)$$

where F_∂ is the diagonal linear differential operator. When the operator F_∂ is applied to $\rho(x,t)$, it results in:

$$[F_\partial \rho(x,t)]_{ij} = \begin{cases} -\nabla\cdot(f_i\rho_i(x,t)), i = j \\ 0, i \neq j \end{cases} \quad i, j = 1, 2\ldots N \quad (5)$$

Taking into account that the state PDF evolution $\rho(x,t)$ depends on the vector $u(t)$, we can formulate the optimal control problem in the probability space using the cost function:

$$J = \int_X w'(x)\rho(x,T)dx \quad (6)$$

In this respect, the optimal control problem is the optimization problem:

$$u^*(t) = \max_{u(t)\in U_{ad}} J = \max_{u(t)\in U_{ad}} \int_X w(x)\rho(x,T)dx \quad (7)$$

Alternatively, to avoid the singular control problems [8], we can consider the optimal control that includes the term penalizing the control:

$$u^*(t) = \max_{u(t)\in U_{ad}} J = \max_{u(t)\in U_{ad}} \int_X w(x)\rho(x,T)dx + \varepsilon \int_0^T u'(t)u'(t) \quad (8)$$

Anyway, the solution of this problem is a sequence of the optimal control $u^*(t)$, from the set of admissible control U_{ad}, such that the cost function is maximized. By a suitable choice of the weighting function $w(x)$, the cost function can be used to find the optimal control maximizing probability of the robot presence in the desired region of the robots' operating space.

The optimal control maximizing the criterion (6) is a special case of a more general optimal control problem of the evolution equation [11]. Under the condition that the operator $F(u)$ is bounded, i.e., $\|F(u(t))\| < \infty$, the minimum principle for PDEs can be applied [11]. According to the minimum principle, the optimal control $u^*(t)$ satisfies:

$$u^*(t) = \arg\min_{u\in U_{ad}} H(\rho(x,t),u(t),t) \quad (9)$$

In other words, for the optimal state PDF trajectory $\rho^*(x,t)$, the optimal control minimizes the Hamiltonian at each time point. The Hamiltonian is:

$$H(\rho(x,t),u,t) = \langle \pi(x,t), F(u)\rho(x,t)\rangle \quad (10)$$

where brackets $\langle\cdot,\cdot\rangle$ denote the scalar product of function vectors defined as:

$$\langle p(x),q(x)\rangle = \int_X p'(x)q(x)dx = \int_X \sum_i p_i(x)q_i(x)dx \quad (11)$$

The function vector $\pi(x,t)$ is the so-called adjoint state distribution and satisfies:

$$\frac{\partial\pi(x,t)}{\partial t} = -F'(u)\pi(x,t) \quad (12)$$

$$\pi(x,T) = -w(x) \quad (13)$$

The major difficulty in computing the optimal control is in evaluation of integrals (11) and corresponding PDE system solutions (4) and (12). Based on the definition of the scalar product (11), the Hamiltonian can be expressed as:

$$H(\rho(x,t),u,t) = \langle \rho(x,t), F'(u)\pi(x,t)\rangle \quad (14)$$

,i.e., $$H(\rho(x,t),u,t) = \sum_i \int_X \rho_i(x,t)[F'(u)\pi(x,t)]_i dx \quad (15)$$

where $[\cdot]_i$ denotes the ith row of the vector. In the following section, we will explain how the evolution of the state PDF ρ, as well as expression (15) can be computed using the stochastic sampling propagator.

STOCHASTIC SAMPLING PROPAGATOR

The evolution of the large-size population probability density function $\rho(x,t)$ is described by the PDE system (see Eq.4). One way to obtain the evolution $\rho(x,t)$ is to solve the PDE system forward in time starting from the initial condition $\rho(x,0) = \rho^0(x)$. We propose an approach to computing the evolution $\rho(x,t)$ based on stochastic trajectories of the hybrid state (x,q) evolution resulting from the model presented in Fig. 1. To account for the fact that the transition rates can change in time, we assume that the control is a piecewise constant function of time discretized with the sample time ΔT. The basis for the proposed algorithm is the Gillespie's stochastic simulation algorithm [12].

To generate the trajectory of (x,q), we need to generate the initial state $(x(0),q(0))$ from the state PDF $\rho(x,0) = \rho^0(x)$. Probability $P_i(t)$ of $q(t) = i$ is:

$$P_i(t) = \int_X \rho_i(x,t)dx \quad (16)$$

Therefore, the random variable $q(0) = i$ should be generated from the discrete state probability distribution represented by the vector of discrete state probabilities $P(0) = [P_1(0)\ P_2(0)\ \ldots P_N(0)]'$. Symbolically, we will represent it as:

$$q(0) = i \sim P(0) \quad (17)$$

Once the initial discrete $q(0)$ state is defined, the continuous variable $x(0)$ can be generated from the corresponding $\rho_i(x,0)$ component of the state PDF ,i.e., from the probability $\mathcal{P}$ of $x(t)$ given that $q(t) = i$ and $t = 0$:

$$x(0) \sim \mathcal{P}\{x|q(t) = i, t = 0\} = \rho_i(x,0)/P_i(0) \quad (18)$$

Whenever the discrete state is $q(t) = i$, the evolution of the continuous state x obeys $\dot{x} = f_i(x)$. Therefore, generating trajectory $(x(t),q(t))$ reduces to the problem of generating the state transitions of the discrete state $q(t)$. Let us assume that at time $t = t_s$, $t_s \in [(k-1)\Delta T, k\Delta T)$, the hybrid state is $(x(t), q(t))$; then, the time instant at which the state changes t_c can be generated based on the following two rules:

(a) $t_c = t_s + t_t$, $t_t \sim e^{-t\sum_j \lambda_{ij}(k-1)}$, under the condition that $t_c < k\Delta T$. If the condition is not satisfied, apply rule (b).
(b) $t_c = k\Delta T + t_t$, $t_t \sim e^{-t\sum_j \lambda_{ij}(k)}$, under the condition that $t_c < (k+1)\Delta T$. If the condition is not satisfied, increase k by 1. Apply rule (b) until the condition is satisfied.

These two rules define the time point t_c at which the jump from the discrete state i to the discrete state j happens, but do not specify the variable j. The state j needs to be sampled from the discrete state probability density function, i.e., from the probability $\mathcal{P}$ of $q(t^+) = j$, given that $q(t) = i$ provided in the vector of the discrete probability distribution with $N-1$ elements:

$$j \sim \mathcal{P}\{q(t)|q(t^-) = i\} = \underbrace{\left[\frac{\lambda_{i1}}{\sum_{n=1}^{N}\lambda_{in}}, \frac{\lambda_{i2}}{\sum_{n=1}^{N}, \lambda_{in}} \cdots \frac{\lambda_{iN}}{\sum_{n=1}^{N}\lambda_{in}}\right]}_{N-1} \tag{19}$$

The above algorithm can be used to generate a single trajectory for the stochastic model shown in Fig. 1. In the limit of a large number of samples, the normalized density of trajectory points will correspond to the solution of the PDE system given by Eq. 4. In this respect, the stochastic simulation is a computational propagator of the evolution $\rho(x,t)$ and we can denote it as:

$$\frac{\partial \rho}{\partial t} = F_{sim}(u(t))\rho \tag{20}$$

HAMILTONIAN EVALUATION

Let us assume that the total number of the trajectories we use to propagate the state PDF ρ is N_{samp}, that $x_k(t)$ and $q_k(t)$ denote the continuous state and discrete states of trajectory k at time t, $k = 1 \dots N_{samp}$. At a given time point t, among N_{samp} trajectories only $N_i(t)$ trajectories are in the discrete state i, and naturally, $\sum_i N_i(t) = N_{samp}$, $\forall t$. From the state PDF normalizing condition (1), we can conclude that:

$$\frac{1}{P_i(t)} \int_X \rho_i(x,t) = 1 \tag{21}$$

Following, the expected value of $[F'(u)\pi(x,t)]_i$ under the condition that the discrete state $q(t) = i$ is

$$E\{F'(u)\pi(x,t)|q(t) = i\} = \frac{1}{P_i(t)} \int_X \rho_i(x,t)[F'(u)\pi(x,t)]_i dx \tag{22}$$

and can be approximated as:

$$E\{F'(u)\pi(x,t)|q(t) = i\} \approx \frac{1}{N_i(t)} \sum_k [F'(u)\pi(x_k,t)]_i \delta(q_k(t) - i) \tag{23}$$

where $\delta(q_k(t) - i) = 1$ if $q_k(t) = i$, and zero elsewhere. Conseqently, the Hamiltonian (15) can be expressed as:

$$\begin{aligned} H\ (\rho(x,t), u, t) &= \textstyle\sum_i P_i(t) E\{F'(u)\pi(x,t)|q(t) = i\} \\ &\approx \textstyle\sum_i P_i(t) \frac{1}{N_i(t)} \sum_k [F'(u)\pi(x_k,t)]_i \delta(q_k(t) - i) \end{aligned} \tag{24}$$

In the limit of a large number of samples $P_i(t) \approx N_i(t)/N_{samp}$; therefore, the Hamiltonian value can be estimated using the following expression:

$$\hat{H}(\rho(x,t), u, t) = \frac{1}{N_{samp}} \sum_i \sum_k [F'(u)\pi(x_k,t)]_i \delta(q_k(t) - i) \tag{25}$$

where $\sum_k$ denotes the sum over all trajectories and $\sum_i$ over all discrete states. This expression is exact in the limit of a large number of samples N_{samp}. To illustrate and verify the algorithm for generating stochastic trajectories $(x(t), q(t))$ and computing the Hamiltonian components, we use a 1D example in the next section.

1D EXAMPLE

The stochastic model presented in Fig. 2*b* illustrates the state PDF evolution of a large-size robot population along one dimension (Fig. 2*a*), in which u_1, u_2 and u_3 correspond to stochastic rates of the commands: move-left, move-right and stop. In this example, $k_1 = -0.5$ and $k_2 = 0.25$. The control $u(t) = [u_1(t)\ u_2(t)\ u_3(t)]$ is computed as the optimal control based on the minimum principle and Hamiltonian presented in the previous section.

The cost function is:

$$J = \int_X w'(x)\rho(x,t)dx + \varepsilon \int_0^T u_1^2(t) + u_2^2(t) + u_3^2(t) dt \tag{26}$$

where $\varepsilon = 10^{-7}$, the weighting $w(x) = [0\ 0\ w_3(x)]'$ and the initial

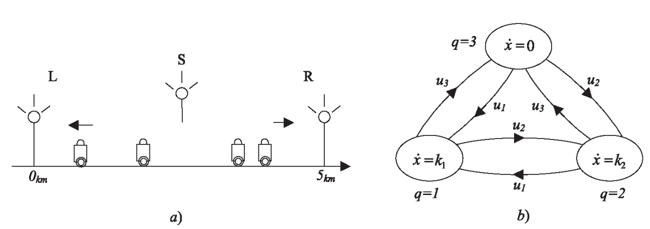

Figure 2. 1D example [7, 8]

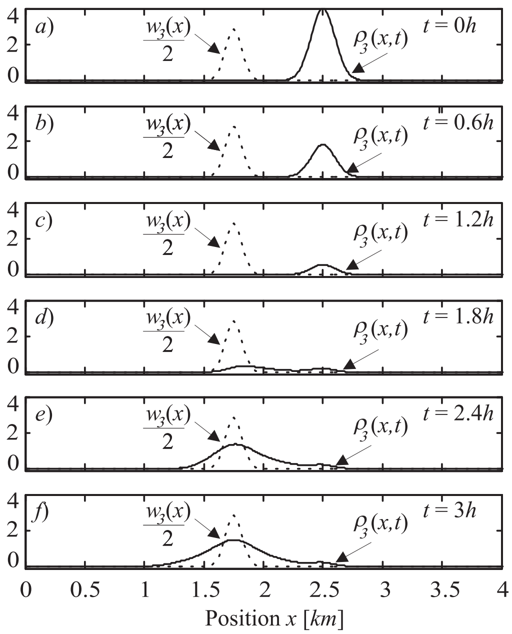

Figure 3. The finite element solution of the state PDF evolution for the 1D example under the optimal control $u^*(t)$, 500 points [7,8]

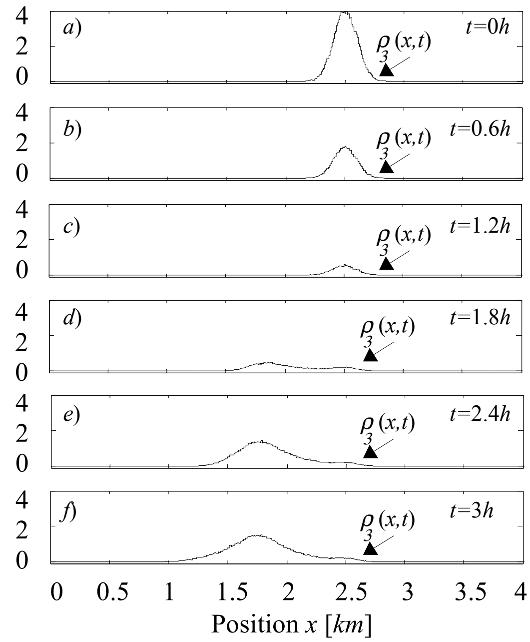

Figure 4. The stochastic simulation solution of the state PDF evolution for the 1D example under the optimal control $u^*(t)$, 10^5 samples

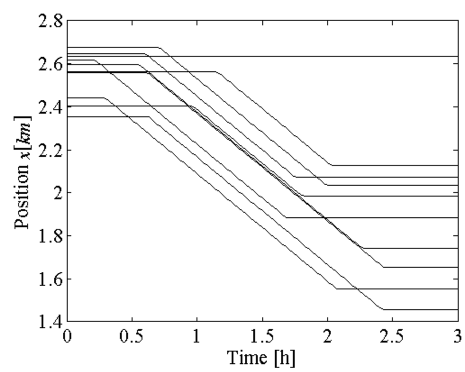

Figure 5. A Random set of 10 trajectories resulting from the stochastic simulation under the optimal control $u^*(t)$

condition $\rho(x,0) = [0\ \ 0\ \rho_3(x,0)]'$ are defined by:

$$w_3(x) = \begin{cases} \frac{1}{\sqrt{0.01}} exp(-\frac{(x-1.75)^2}{0.01}),\ \ 1.25 < x < 2.25 \\ 0,\ otherwise \end{cases} \tag{27}$$

$$\rho_3(x,0) = \begin{cases} \frac{1}{\sqrt{0.02\pi}} exp(-\frac{(x-2.5)^2}{0.02}),\ \ 2 < x < 3 \\ 0,\ otherwise \end{cases} \tag{28}$$

The optimal control sequence $u^*(t) = [u_1^*(t)\ u_2^*(t)\ u_3^*(t)]$ in the time interval $0 < t < 3$ is defined by:

$$u_1^*(t) = \begin{cases} 2,\ 0.21 < t < 1.74 \\ 0, elsewhere \end{cases},\ \ u_2^*(t) = 0 \tag{29}$$

$$u_3^*(t) = \begin{cases} 2,\ 1.71 < t < 3 \\ 0, elsewhere \end{cases} \tag{30}$$

The evolution of the state PDF for this system under the control $u^*(t)$ is presented in Fig. 3. We present only $\rho_1(x,t)$ and $\rho_3(x,t)$ because under this control $\rho_2(x,t) = 0,\ \forall t$.

For the illustration, we generated 10 stochastic trajectories of the continuous variable x (see Fig. 5) under the control $u^*(t)$. The evolution of the discrete state q can be observed from the trend in x. When x decreases, the discrete state is 1, and when it remains constant, the state is $q = 3$. It is worth mentioning that, among these 10 trajectories, there is one for which $x(t)$ is constant. The small pick around the point 2.5 in the right panel of the Fig.3 at $t = 3$, confirms that the probability of such trajectories is non-zero, but it is small.

To obtain the state PDF $\rho(x,t)$, i.e., its components $\rho_i(x,t)$ at a specific time point t, we need to collect points $x(t)$ and estimate components $\rho_i(x,t)$. It is obvious that 10 trajectories cannot provide a good estimate of $\rho(x,t)$ For this reason, we generated 10^5 trajectories and computed the histogram probability density function estimate. That means that we discretized the x axis into intervals of the length Δx and counted how many points fell into

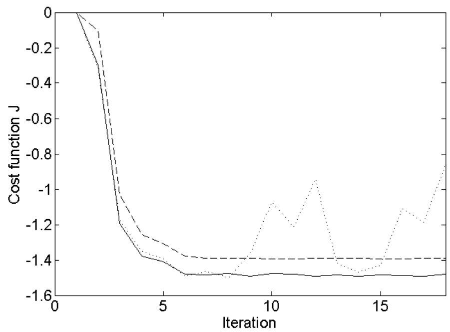

Figure 6. Cost function J iterations computed with the algorithm [8] with the Hamiltonian evaluated based on the PDE solutions (dashed), and estimated based on 10^4 (dotted) and 10^5 (solid) stochastic samples.

a specific region. Finally, we normalized the histogram so that the estimated $\rho(x,t)$ is normalized to 1. Our results are presented in Fig. 4. As expected, the match between the numerical PDE system solution and the result obtained from stochastic trajectories is exact. There are only negligible discrepancies due to data sampling from a finite number of trajectories.

In Table 1, we provide the time our MATLAB code takes to compute the state PDF evolution based on the stochastic simulation approach (Fig. 4) and the time it takes to compute the evolution based on the finite element (FE) approach with 500 points (Fig. 3). We can see that the time for the stochastic simulation approach increases roughly linearly with the number of samples and that the time of the FE approach can be reached if we use approximately $3 \cdot 10^5$ samples.

For the purpose of computing the optimal control, the number of the samples we need depends on the convergence of the optimization algorithm to the solutions. Therefore, as the final test of the stochastic sampling propagator, we computed the optimal control based on the algorithm presented in [8], but using the stochastic samples for the Hamiltonian evaluation (25) instead of the Hamiltonian which is completely based on PDE systems solutions [8]. We can see from Fig. 6 that for 10^4 samples, the Hamiltonian (25) fluctuations are at such a level that the optimization algorithm does not converge, and we see it as large fluctuations of the cost function. While the fluctuations are intrinsic property of the stochastic samples evaluation, they can be smaller if we use a larger number of samples.

When we use 10^5 samples, the fluctuations are much smaller and the optimization algorithm converges. It is interesting to notice that the stochastic sampling evaluations result into the cost function value which is smaller that the value resulting from the PDE system solutions (see Fig. 6). This is because of discrete approximations involved into solutions of PDE systems, as well as approximation of integrals contributing to the cost function.

The average time that our MATLAB code takes for a single iteration based on the finite element method (500 points) is $1629s$ (Fig. 6, dashed line). Using the same optimization code, a single iteration in the stochastic-based method with 10^5 samples is 350s (Fig. 6, solid line). Under the conditions presented above, the stochastic-based method is in average $4.5x$ faster than its deterministic counterpart.

Table 1. Speed comparison between the stochastic simulation approach and the finite element (FE) solution with 500 points

Number of samples	10^3	10^4	10^5	FE
Time (s)	0.48s	4.8s	44.8s	1363s

CONCLUSION

In this paper we considered a large-size robot population control problem that had been previously formulated and solved in a probability space utilizing systems of PDEs. Solving these PDEs is computationally expensive; therefore, having in mind that the PDEs are in close connection with the stochastic process to be controlled, we explore an opportunity to utilize the stochastic process samples to compute the control.

Our paper describes an algorithm for generating the stochastic process that can be used to propagate the state PDF of the robot population. We show that the algorithm predicts exactly the state PDF evolution and we derive expression for the Hamiltonian evaluation which involves the stochastic process samples. The Hamiltonian evaluated in this way can be used in iterations computing the optimal control as if it was computed based on the PDE system solutions. We also notice that the cost function resulting from utilizing stochastic processes has smaller values than the cost function computed based on PDEs. This means that the evaluations involving the stochastic process samples are closer to the true values.

In summary, we can conclude that utilizing stochastic processes for computing control of multi-robot systems considering discrete, as well as continuous robot states is possible. By embedding stochastic process generators into analog circuits and utilizing them in dedicated processors for computing control, complex stochastic optimal control problems can be solved efficiently and potentially exploited for real-time multi-robot systems control.

REFERENCES

[1] Ho, Y., and Bryson, A. E., 1975. *Applied Optimal Control: Optimization Estimation and Control*. John Wiley & Sons.

[2] Robert, C., and Casella, G., 2004. *Monte Carlo Statistical Methods*. Springer.

[3] Kappen, H. J., 2005. "Linear theory for control of nonlinear stochastic systems". *Physical Review Letters,* **95**(20), p. 200201.

[4] van den Broek, B., Wiegerinck, W., and Kappen, B., 2008. "Graphical model inference in optimal control of stochastic multi-agent systems". *Journal of Artificial Intelligence Research,* **32**, pp. 95–122.

[5] der Schaft, A. V., and Schumacher, J., 2000. *An Introduction to Hybrid Dynamical Systems*. Springer-Verlag.

[6] Milutinović, D., Lima, P., and Athans, M., 2003. "Biologically inspired stochastic hybrid control of multi-robot systems". In Proceedings of the 11th International Conference on Advanced Robotics.

[7] Milutinović, D., and Lima, P., 2006. "Modeling and optimal centralized control of a large-size robotic population". *IEEE Transactions on Robotics,* **22**, pp. 1280–1285.

[8] Milutinović, D., and Lima, P., 2007. *Cells and Robots : Modeling and Control of Large-Size Agent Populations*. Springer.

[9] Kushner, H. J., 1967. *Stochastic Stability and Control*. Academic Press, New York.

[10] Cassandras, C., and Lafortune, S., 1999. *Introduction to Discrete Event Systems*. Kluwer Academic Publ.

[11] Fattorini, H., 1999. *Infinite Dimensional Optimization and Control Theory*. Cambridge University Press.

[12] Gillespie, D., 1977. "Exact stochastic simulation of coupled chemical reactions". *Journal of Physical Chemistry,* **81**(25), pp. 2340–2361.