Robust, Optimal Predictive Control of Jump Markov Linear Systems using Particles

Lars Blackmore[1], Askar Bektassov[2], Masahiro Ono[1], and Brian C. Williams[1]

[1] Massachusetts Institute of Technology, Cambridge, MA 02139
larsb@mit.edu, hiro_ono@mit.edu, williams@mit.edu
[2] Università degli Studi Roma Tre, Rome, Italy 00146
askar.bektassov@gmail.com

Abstract. Hybrid discrete-continuous models, such as Jump Markov Linear Systems, are convenient tools for representing many real-world systems; in the case of fault detection, discrete jumps in the continuous dynamics are used to model system failures. Stochastic uncertainty in hybrid systems arises in both the continuous dynamics, in the form of uncertain state estimation, disturbances or uncertain modeling, and in the discrete dynamics, which are themselves stochastic.
In this paper we present a novel method for optimal predictive control of Jump Markov Linear Systems that is robust to both continuous and discrete uncertainty. The approach extends our previous 'particle control' approach, which approximates the predicted distribution of the system state using a finite number of particles. Here, we present a *weighted* particle control approach, which uses importance weighting to ensure that low probability events such as failures are considered. We demonstrate the method with a car braking scenario.

1 Introduction

Hybrid discrete-continuous models, such as Jump Markov Linear Systems (JMLS), are convenient tools for representing many real-world systems[1, 2]. In the case of fault detection and fault-tolerant control, discrete jumps in the continuous dynamics are used to model component failures[3]. Stochastic uncertainty in hybrid systems arises in both the continuous dynamics, in the form of uncertain state estimation, disturbances or uncertain modeling, and in the discrete dynamics, which are themselves stochastic.

Control of stochastic systems has received a great deal of attention in recent years, see [4] for a survey. Much work has been done in the area of feedback control for JMLS, see [5] for a survey. By contrast, *predictive* optimal stochastic control takes into account probabilistic uncertainty in dynamic systems and aims to control the predicted distribution of the system state in some optimal manner. In the case of stochastic linear dynamic systems, recent work has developed tractable algorithms for optimal, *robust* predictive control[6–10]. These methods are robust in the sense that they ensure the system state leaves a given feasible region with probability at most δ. This *chance-constrained* formulation is a powerful one, as it enables the user to specify a desired level of conservatism, which can be traded against performance.

Chance-constrained optimal stochastic control of JMLS has a number of important applications. In the case of fault-tolerant control, we would like to be able to control a system in an optimal manner while taking into account both continuous disturbances and the possibility of system failures, such that task failure is below a certain threshold. For example, in controlling an autonomous ground vehicle we would like to ensure that, despite having brakes that may fail, collision with obstacles or other vehicles happens with low probability. Recent work developed a Model Predictive Control approach for JMLS which imposes constraints on the mean and covariance of the system state[11]. These are not the same as chance constraints even when all forms of uncertainty are Gaussian since the state distribution in JMLS is multimodal.

In this paper we develop a tractable algorithm for chance-constrained optimal predictive control of JMLS that extends our previous work on particle-based control of continuous systems[12]. The key idea behind this approach is to approximate all probability distributions using a finite number of samples, or 'particles'. In doing so, we approximate the stochastic predictive control problem as a deterministic one, with the property that as the number of particles tends to infinity, the approximation tends to the original stochastic problem. The resulting optimization problem can be solved efficiently using Mixed Integer Linear Programming (MILP). The approach generalizes previous work by [17] by handling stochastic uncertainty with general distributions, in both the continuous and discrete dynamics.

In this paper we present first a straightforward extension of the particle control method to JMLS. This extension uses particles to represent uncertainty in the discrete mode sequences as well as the continuous variables. An empirical validation with a ground vehicle braking scenario shows that the method is effective, but is prone to neglect low-probability events such as failures. We therefore develop a new *weighted* particle control approach that overcomes these difficulties by drawing on the idea of importance weighting from particle filtering[13–16]. The key idea is that by sampling mode sequences from a proposal distribution, and representing the discrepancy between the proposal distribution and the true distribution by an analytic weight, an increase in sampling efficiency can be achieved. The resulting optimization can be solved efficiently and to global optimality using MILP. We demonstrate empirically that a dramatic improvement in performance is achieved by employing the weighted particle control approach.

2 Problem Statement

In this paper we are concerned with the following stochastic control problem:

> *Design a finite, optimal sequence of control inputs $\mathbf{u}_{0:T-1}$, taking into account probabilistic uncertainty, which ensures that the continuous state trajectory $\mathbf{x}_{c,1:T}$ of a JMLS leaves a defined feasible region F with probability at most δ, and satisfies constraints on the expected system state.*

We consider four sources of stochastic uncertainty; initial state uncertainty; system model uncertainty; disturbances, modeled as stochastic processes; and random mode transitions. These transitions can model component failures, for example. We assume that the p.d.f.s of the uncertainty mentioned here are known at least approximately, but we make no assumptions about the form the distributions take. We assume a cost function that is piecewise linear in the control inputs; previous work has shown that minimum control effort and minimum time problems can be posed using such functions[18]. Finally, we assume that the feasible region F is a polytope, and that the control inputs $\mathbf{u}_t$ are subject to interval constraints.

We define a Jump Markov Linear System as a system with hybrid discrete-continuous state $\mathbf{x} = \langle \mathbf{x}_c, x_d \rangle$. The discrete state x_d is a Markov chain that can take one of M values and evolves according to:

$$p(x_{d,t+1} = j | x_{d,t} = i) = T_{ij}. \tag{1}$$

The continuous state $\mathbf{x}_c$ evolves according to:

$$\mathbf{x}_{c,t+1} = A(x_{d,t})\mathbf{x}_{c,t} + B(x_{d,t})\mathbf{u}_t + \nu_t. \tag{2}$$

The initial hybrid discrete-continuous state is random, with a known distribution $p(\mathbf{x}_{c,0}, x_{d,0})$. The variable ν_t is a random disturbance process distributed according to $p(\nu_t | x_{d,t})$, which we assume independent from the initial state. Modeling errors can be modeled as an additional stochastic disturbance. For notational simplicity we assume a single disturbance process.

The key idea behind solving this stochastic control problem is to approximate all distributions using samples, or particles, and then solve the resulting deterministic problem. In Section 3 we review some results relating to sampling from random variables. In Section 4 we review the chance-constrained particle control approach introduced in [12] for systems with continuous state. We then extend this approach to JMLS in Section 5 and show that the resulting problem can be solved using MILP. In Section 6 we introduce a novel weighted particle control method that gives a dramatic increase in performance by using a proposal distribution to sample from discrete mode sequences. In Section 7 we provide empirical results.

3 Sampling from Random Variables

Previous work has shown that approximating the probability distribution of a random variable using samples drawn from that distribution, or particles, can lead to tractable algorithms for estimation and control[19]. Here we review some properties of samples drawn from random variables.

Suppose that we have a multivariate random variable X with p.d.f. $p(\mathbf{x})$. We draw N independent, identically distributed random samples $\mathbf{x}^{(1)}, \cdots, \mathbf{x}^{(N)}$ from this distribution. Often, we would like to calculate an expectation involving this random variable:

$$E_X[f(X)] = \int_X f(\mathbf{x})p(\mathbf{x})d\mathbf{x} \tag{3}$$

In many cases this integral cannot be evaluated in closed form. Instead it can be approximated using the sample mean:

$$\hat{E}_X[f(X)] = \frac{1}{N}\sum_{i=1}^{N} f(\mathbf{x}^{(i)}). \tag{4}$$

From the strong law of large numbers, the sample mean converges to the true expectation as N tends to infinity. This can be used to approximate the probability of a certain event, such as the event $f(\mathbf{x}) \in A$. This is given exactly by:

$$P_A = \int_{f(\mathbf{x})\in A} p(\mathbf{x})d\mathbf{x} = E_X[g(\mathbf{x})] \quad \text{where} \quad g(\mathbf{x}) = \begin{cases} 1 & f(\mathbf{x}) \in A \\ 0 & f(\mathbf{x}) \notin A. \end{cases} \tag{5}$$

We can therefore approximate P_A as:

$$\hat{P}_A = \frac{1}{N}\sum_{i=1}^{N} g(\mathbf{x}^{(i)}) \qquad \text{where} \qquad \hat{P}_A \longrightarrow P_A \text{ as } N \longrightarrow \infty. \tag{6}$$

Note that $\sum_{i=1}^{N} g(\mathbf{x}^{(i)})$ is simply the number of particles for which $f(\mathbf{x}^{(i)}) \in A$. Assuming that evaluating $f(\cdot)$, and checking whether a given value is in A, are both straightforward, calculating $\hat{P}_A$ is also; we simply need to count how many of the propagated particles, $f(\mathbf{x}^{(i)})$ fall within A. By contrast, evaluating P_A as in (5) requires a finite integral over an arbitrary probability distribution, where even calculating the bounds on the integral may be intractable. Hence the particle-based approximation is extremely useful, especially given the convergence property in (6). In Section 4 we use this property to approximate the stochastic control problem defined in Section 2.

3.1 Importance Weighting

In certain situations, drawing samples from the distribution $p(\mathbf{x})$ may be intractable. In such cases, previous work proposed sampling from an alternative *proposal* distribution and using *importance sampling* to correct for the discrepancy between the desired distribution and the proposal distribution[19]. We review relevant results here.

The proposal distribution $q(\mathbf{x})$ is chosen so that sampling from $q(\mathbf{x})$ is easy, and so that $p(\mathbf{x}) > 0$ implies $q(\mathbf{x}) > 0$. We draw N independent, identically distributed random samples $\mathbf{x}^{(1)}, \cdots, \mathbf{x}^{(N)}$ from $q(\mathbf{x})$. To each sample we assign an importance weight w_i, where $w_i = p(\mathbf{x}^{(\mathbf{i})})/q(\mathbf{x}^{(\mathbf{i})})$. In order to approximate the expectation of the function $f(\cdot)$ we now use the *weighted* sample mean:

$$\hat{E}_X[f(X)] = \frac{1}{N}\sum_{i=1}^{N} w_i f(\mathbf{x}^{(i)}). \tag{7}$$

From the strong law of large numbers, we have the convergence property as N tends to infinity:

$$\hat{E}_X[f(X)] \longrightarrow E_X[f(X)]. \tag{8}$$

In order to approximate the probability of the event $f(\mathbf{x}) \in A$ we use the *weighted* number of propagated particles that fall within A:

$$\hat{P}_A = \frac{1}{N} \sum_{i=1}^{N} w_i g(\mathbf{x}^{(i)}), \tag{9}$$

where $g(\cdot)$ is as defined in (5). As in (6) we have the convergence property $\hat{P}_A \longrightarrow P_A$ as $N \rightarrow \infty$.

4 Review of Particle Control Approach

In this section we review the chance-constrained particle control approach introduced in [12] for robust control of systems with continous state.

The key observation behind the method is that, by approximating all probabilistic distributions using particles, an intractable stochastic optimization problem can be approximated as a tractable deterministic optimization problem. By solving this deterministic problem we obtain an approximate solution to the original stochastic problem, with the additional property that as the number of particles used tends to infinity, the approximation becomes exact.

The method is outlined as follows:

1. Generate N samples from the joint distribution of initial state and disturbances.
2. Express the distribution of the future state trajectories approximately as a set of N *analytic* particles, where each particle $\mathbf{x}_{1:T}^{(i)}$ corresponds to the state trajectory given a particular set of samples. Each particle *depends explicitly on the control inputs* $\mathbf{u}_{0:T-1}$, which are yet to be generated.
3. Approximate the chance constraints in terms of the generated particles; the probability of $\mathbf{x}_{1:T}$ falling outside of the feasible region is approximated as the *fraction* of particles $\mathbf{x}_{1:T}^{(i)}$ that do so.
4. Approximate the cost function in terms of particles.
5. Solve the deterministic constrained optimization problem for control inputs $\mathbf{u}_{0:T-1}$.

The method is illustrated in Fig. 1. The general particle control problem results in a deterministic optimization problem that is intractable, except for very small problems. However in [12] we showed that for a polytopic feasible region F, piecewise linear cost function h and linear system dynamics $\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t$, the deterministic optimization can be solved to global optimality in an efficient manner using MILP. This relies on the fact that each particle $\mathbf{x}_{1:T}^{(i)}$ is a *linear* function of the control input sequence $\mathbf{u}_{0:T-1}$.This is also true for time-varying linear systems. In Section 5 we use this to extend the method to JMLS.

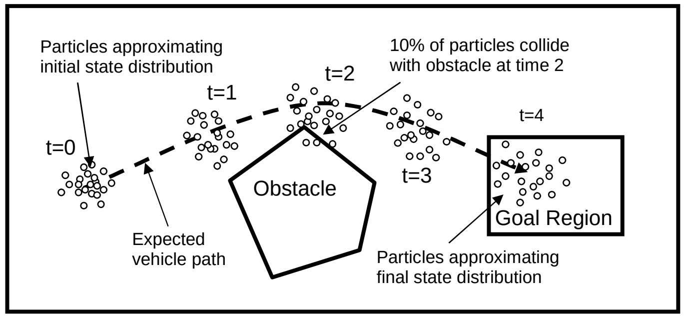

Fig. 1. Illustration of chance constrained particle control method for continuous systems. For this vehicle path planning scenario, the feasible region is defined so that the plan is successful if the vehicle avoids the obstacles at all time steps and is in the goal region at the final time step. The objective is to find the optimal sequence of control inputs so that the plan is successful with probability at least 0.9. The particle control method approximates this so that at most 10% of the particles fail.

5 Straightforward Extension of Particle Control to JMLS

For JMLS, we approximate the stochastic control problem by sampling from *discrete mode sequences* as well as disturbances. Given a discrete mode sequence and samples for all of the disturbance variables, the future system state trajectory is a known deterministic function of the control inputs. Hence each particle provides a sample of the future state trajectory corresponding to a sample of the discrete mode sequence and disturbances.

Note that the mode sequence is independent of the control inputs and the continuous state in JMLS, and hence:

$$p(\mathbf{x}_{c,1:T}, x_{d,1:T}|\mathbf{u}) = p(\mathbf{x}_{c,1:T}|x_{d,1:T}, \mathbf{u})p(x_{d,1:T}). \tag{10}$$

We therefore first generate samples of the mode sequence $x_{d,1:T}$, and for each sample $x^{(i)}_{d,1:T}$, we generate samples of the disturbance variables. While there are M^T different mode sequences, sampling from $p(x_{d,1:T})$ is straightforward due to the Markov property. The algorithm is described in full in Table 1. From the results in Section 3 we have convergence of the approximated problem to the original stochastic problem as the number of particles tends to infinity.

5.1 MILP Solution of JMLS Particle Control

We now show that the approximated problem can be solved efficiently using MILP. For a given particle, the mode at each time step in the horizon is known, as are the disturbances at each time step. From the definition of JMLS in Section 1 we obtain the following expression for each particle:

$$\mathbf{x}^{(i)}_{c,t} = \sum_{j=0}^{t-1}\Big(\prod_{l=1}^{t-j-1} A(x^{(i)}_{d,l})\Big)\Big(B(x^{(i)}_{d,j})\mathbf{u}_j + \nu^{(i)}_j\Big) + \Big(\prod_{l=1}^{t} A(x^{(i)}_{d,l})\Big)\mathbf{x}^{(i)}_{c,0}. \tag{12}$$

1) Generate N samples of the initial discrete mode $\{x_{d,0}^{(1)}, \ldots, x_{d,0}^{(N)}\}$ according to the distribution $p(x_{d,0})$. 2) For each sample $x_{d,0}^{(i)}$, generate a sample of the initial continuous state $\{\mathbf{x}_{c,0}^{(1)}, \ldots, \mathbf{x}_{c,0}^{(N)}\}$ according to $p(\mathbf{x}_{c,0}\|x_{d,0}^{(i)})$. 3) For each sample $x_{d,0}^{(i)}$ generate a sample of the discrete mode sequence $x_{d,1:T}^{(i)}$ according to $p(x_{d,1:T}\|x_{d,0})$. 4) For each sample $x_{d,0:T}^{(i)}$ generate a sample of the disturbances $\{\nu_0^{(i)}, \ldots, \nu_{T-1}^{(i)}\}$ from the distribution $p(\nu_0, \ldots, \nu_{T-1}\|x_{d,0:T}^{(i)})$. 5) Express the distribution of the future state trajectories approximately as a set of N particles, where each particle $\mathbf{x}_{c,1:T}^{(i)}$ corresponds to the continuous state trajectory given a particular set of samples $\{\mathbf{x}_0^{(i)}, x_{d,1:T}^{(i)}, \nu_0^{(i)}, \cdots, \nu_{T-1}^{(i)}\}$. Each particle depends explicitly on the control inputs $\mathbf{u}_0, \cdots, \mathbf{u}_{T-1}$, which are yet to be generated. 6) Approximate the expected state constraints and chance constraints in terms of the generated particles. $E[\mathbf{x}_{1:T}^{(i)}] \approx \frac{1}{N}\sum_{i=1}^{N} \mathbf{x}_{1:T}^{(i)} = \mathbf{x}_{1:T}^{equality} \quad p(\mathbf{x}_{c,1:T} \notin F) \approx \frac{1}{N}\sum_{i=1}^{N} g(\mathbf{x}_{1:T}^{(i)}) \leq \delta, \quad (11)$ where $g(\cdot)$ is as defined in (5). 7) Approximate the cost function in terms of particles. 8) Solve deterministic constrained optimization problem for inputs $\mathbf{u}_{0:T-1}$.

Table 1. Straightforward Particle Control Approach for JMLS

Note that this is a linear function of the control inputs, and that $\mathbf{x}_{c,0}^{(i)}$, $\nu_j^{(i)}$ and $x_{d,l}^{(i)}$ are all known values. Hence each particle $\mathbf{x}_{c,1:T}^{(i)}$ is linear in the control inputs.

In accordance with (11), we need to constrain the number of particles that fall outside of the feasible region. In the same manner as described in [12], we define a set of N binary variables $z_1, \cdots, z_N$, where $z_i \in \{0, 1\}$. These binary variables are defined so that $z_i = 0$ implies that particle i falls inside the feasible region. We then constrain the sum of these binary variables:

$$\frac{1}{N}\sum_{i=1}^{N} z_i \leq \delta. \tag{13}$$

This constraint ensures that the fraction of particles falling outside of the feasible region is at most δ. In [12] we showed how to impose constraints such that $z_i = 0 \implies \mathbf{x}_{1:T}^{(i)} \in F$ for convex and non-convex polygonal feasible regions. We do not repeat this here, but we do note that the linearity of (12) and piecewise linearity of the cost function h ensures that the encoding results in a MILP, which can be solved efficiently to global optimality. We have therefore introduced a new method for robust optimal control of JMLS, where the probability distributions of uncertain variables can take an arbitrary form.

1) Generate N samples of the initial discrete mode $\{x_{d,0}^{(1)}, \dots, x_{d,0}^{(N)}\}$ according to the distribution $p(x_{d,0})$.
2) For each sample $x_{d,0}^{(i)}$, generate a sample of the initial continuous state $\{\mathbf{x}_{c,0}^{(1)}, \dots, \mathbf{x}_{c,0}^{(N)}\}$ according to $p(\mathbf{x}_{c,0}|x_{d,0}^{(i)})$.
3) For each sample $x_{d,0}^{(i)}$ generate a sample of the discrete mode sequence $x_{d,1:T}^{(i)}$ *according to the proposal distribution* $q(x_{d,1:T})$.
4) For each sample $x_{d,0:T}^{(i)}$ generate a sample of the disturbances $\{\nu_0^{(i)}, \dots, \nu_{T-1}^{(i)}\}$ from the distribution $p(\nu_0, \dots, \nu_{T-1}|x_{d,0:T}^{(i)})$.
5) For each sample $x_{d,0:T}^{(i)}$ calculate $p(x_{d,0:T}^{(i)})$ and assign weight w_i as in (16).
6) Express the distribution of the future state trajectories approximately as a set of N particles, where each particle $\mathbf{x}_{c,1:T}^{(i)}$ corresponds to the continuous state trajectory given a particular set of samples $\{\mathbf{x}_0^{(i)}, x_{d,1:T}^{(i)}, \nu_0^{(i)}, \cdots, \nu_{T-1}^{(i)}\}$.
7) Approximate the chance constraints using the *weighted fraction* of particles outside of the feasible region:

$$p(\mathbf{x}_{1:T} \notin F) \approx \frac{1}{N} \sum_{i=1}^{N} w_i g(\mathbf{x}_{1:T}^{(i)}) \leq \delta. \tag{14}$$

8) Approximate the expected state constraints using the *weighted* sample mean approximation, for example:

$$E[\mathbf{x}_{1:T}] = \mathbf{x}_{1:T}^{equality} \quad \text{becomes} \quad \frac{1}{N} \sum_{i=1}^{N} w_i \mathbf{x}_{1:T}^{(i)} = \mathbf{x}_{1:T}^{equality}. \tag{15}$$

9) Approximate the cost function in terms of weighted particles.
10) Solve the deterministic constrained optimization problem for inputs $\mathbf{u}_{0:T-1}$.

Table 2. Weighted Particle Control Approach for JMLS

6 Weighted Particle Control for JMLS

We now extend the method described in Section 5 to deal more efficiently with low probability mode transitions such as failures. The key idea behind the extension is to sample mode sequences from a proposal distribution designed to ensure that low probability events such as failures are more likely to be taken into consideration. Drawing on the idea of importance weighting in particle filtering[19], the discrepancy between the actual distribution over mode sequences and the proposal distribution is represented using an analytical weighting. In doing so, we retain the convergence property that the approximate problem converges to the original stochastic problem as the number of particles tends to infinity. The algorithm is described in Table 2.

We now show how to calculate the weights w_i. For the approximated problem to converge to the original stochastic problem as the number of particles tends to infinity, weights must be assigned according to[19]:

$$w_i = \frac{p(\mathbf{x}_{c,1:T}^{(i)}, x_{d,1:T}^{(i)}|\mathbf{u}_{0:T-1})}{q(\mathbf{x}_{c,1:T}^{(i)}, x_{d,1:T}^{(i)}|\mathbf{u}_{0:T-1})}. \tag{16}$$

Since we sample the disturbances from their true distributions, the joint proposal $q(\mathbf{x}_{c,1:T}, x_{d,1:T})$ can be written in terms of the proposal over mode sequences to give:

$$w_i = \frac{p(\mathbf{x}^{(i)}_{c,1:T}|x^{(i)}_{d,1:T}, \mathbf{u}_{0:T-1})p(x^{(i)}_{d,1:T})}{p(\mathbf{x}^{(i)}_{c,1:T}|x^{(i)}_{d,1:T}, \mathbf{u}_{0:T-1})q(x^{(i)}_{d,1:T})} = \frac{p(x^{(i)}_{d,1:T})}{q(x^{(i)}_{d,1:T})}. \tag{17}$$

Since calculating both the true probability of a given mode sequence and its probability according to the proposal distribution is straightforward, calculating the weight to assign to a sampled mode sequence is also.

We now show that the weighted particle control problem for JMLS can be solved using MILP. The key insight is that, since the weights do not depend on the control inputs $\mathbf{u}_{0:T-1}$, incorporating weighted particles does not affect the form of the optimization problem.

The weighted particle control problem can be formulated in exactly the same manner as the unweighted approach described in Section 5, except for the approximate chance constraint and the approximate cost function. We now must constrain the *weighted* fraction of particles that fall outside of the feasible region. Defining again binary variables z_i such that $z_i = 0 \implies \mathbf{x}^{(i)}_{c,1:T} \in F$, we constrain the *weighted* sum of the binary variables:

$$\frac{1}{N}\sum_{i=1}^{N} w_i z_i \leq \delta. \tag{18}$$

The weights w_i do not depend on the control inputs, as shown in (17). Hence (18) is a linear constraint on the binary variables z_i. The expected cost is now approximated using the *weighted* sample mean as follows:

$$E[h] \approx \hat{h} = \frac{1}{N}\sum_{i=1}^{N} w_i h(\mathbf{u}_0, \cdots, \mathbf{u}_{T-1}, \mathbf{x}^{(i)}_{c,1:T}). \tag{19}$$

As the number of particles tends to infinity, we have the convergence result $\hat{h} \longrightarrow E[h]$. Furthermore, since the weights w_i do not depend on the control inputs, the approximate value $\hat{h}$ is piecewise-linear in the control inputs assuming a piecewise-linear cost function h. Similarly, the expected state is approximated using the *weighted* sample mean. This weighted sample mean is a linear function of the control inputs, hence expected state constraints such as (15) are linear.

In summary, therefore, the weighted particle control problem for JMLS can be posed as a MILP. It now remains to choose a proposal distribution $q(x^{(i)}_{d,1:T})$.

6.1 Choosing a Proposal Distribution

The convergence of the approximate problem to the original deterministic problem applies for any choice of the proposal distribution $q(x_{d,1:T})$ subject to the constraint that $q(x_{d,1:T}) > 0$ wherever $p(x_{d,1:T}) > 0$. However for a finite number of particles the performance of the weighted particle control approach is affected greatly by the choice of $q(x_{d,1:T})$. As in particle filtering, the appropriate choice of $q(x_{d,1:T})$ depends on the application, and a great deal of work has focussed on developing proposal distributions for specific applications, for example [14, 20]. We now introduce a proposal distribution designed to improve the performance of the particle control approach for JMLS when dealing with low-probability transitions such as faults.

Consider first a proposal distribution equal to the true mode sequence distribution:

$$q(x_{d,1:T}) = p(x_{d,1:T}). \tag{20}$$

In a JMLS with low-probability transitions such as faults, there is a high probability that no fault transitions will be sampled if this proposal is used.

Next consider a proposal equal to the pseudo-uniform distribution $q(x_{d,1:T}) = U(x_{d,1:T})$, where $U(\cdot)$ assigns an equal probability to each mode sequence for which $p(x_{d,1:t}) > 0$. More precisely:

$$U(x_{d,1:T}) = \begin{cases} 1/n_p & p(x_{d,1:T}) > 0 \\ 0 & p(x_{d,1:T}) = 0, \end{cases} \tag{21}$$

where n_p is the number of mode sequences for which $p(x_{d,1:T}) > 0$. Using this proposal ensures that sequences involving faults are sampled with the same likelihood as the mode sequence without failures, which in reality has much higher probability. This means that the control algorithm is more likely to take into account the sequences involving faults in the control design. The drawback in using this proposal is that there is a significant likelihood that the nominal mode sequence is not sampled. If this occurs, the deterministic optimization will typically be infeasible, since achieving most control tasks requires nominal operation of the system components with non-zero probability.

We therefore choose a proposal distribution $q^*(x_{d,1:T})$ that increases the probability of sampling failure sequences, while ensuring that the nominal mode sequence is sampled at least once with a probability λ:

$$q^*(x_{d,1:T}) = \begin{cases} P_{nom} & x_{d,1:T} = x_{d,1:T}^{nom} \\ \frac{1-P_{nom}}{n_p-1} & x_{d,1:T} \neq x_{d,1:T}^{nom} \end{cases} \quad \text{where} \quad P_{nom} = 1-(1-\lambda)^{1/N}. \tag{22}$$

The proposal distribution $q^*(x_{d,1:T})$ therefore ensures a minimum probability of sampling the nominal mode sequence and shares the remaining probability space evenly among the remaining mode sequences.[3]

In Section 7 we give an empirical analysis that shows that using this proposal distribution the weighted particle control algorithm significantly outperforms straightforward particle control for JMLS when there are low-probability transitions such as failures.

[3] For simplicity of exposition, $q^*(x_{d,1:T})$ described here assumes a single nominal mode sequence. The extension to multiple nominal mode sequences is straightforward.

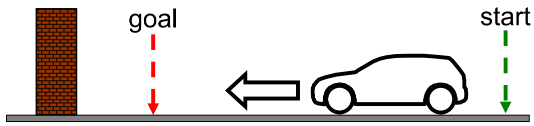

Fig. 2. Illustration of ground vehicle brake failure scenario. The expected vehicle position must arrive at the goal in the minimum possible time, while avoiding collision with the wall.

7 Results

In order to illustrate the new particle control approach for JMLS we use a simple ground vehicle braking example. In this example the system to be controlled is a ground vehicle that can accelerate and brake along a one-dimensional track. The brakes however, can be in one of two modes; mode 1 = *ok* and mode 2 = *faulty*. In the *ok* mode, accelerations and decelerations can be applied to the vehicle, however when the brakes are in the *faulty* mode, decelerations cannot be applied. The continuous system state $\mathbf{x}_c$ is comprised of the position along the track y and the velocity $\dot{y}$. The continuous state evolves according to:

$$\dot{\mathbf{x}}_c = \begin{bmatrix} \dot{y} \\ \ddot{y} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & -b_{fric} \end{bmatrix} \begin{bmatrix} y \\ \dot{y} \end{bmatrix} + B(x_{d,t}) \begin{bmatrix} u_{power} \\ u_{brake} \end{bmatrix} + \nu_t, \tag{23}$$

where the control inputs u_{power} and u_{brake} are both constrained to be greater than or equal to zero (in other words neither negative power nor negative braking can be applied). The term b_{fric} represents a damping term due to friction. Random disturbances ν_t act on the vehicle. The matrix $B(x_{d,t})$ is defined as follows:

$$B(x_{d,t}) = \begin{cases} \begin{bmatrix} 1 & -1 \end{bmatrix} & x_{d,t} = ok \\ \begin{bmatrix} 1 & 0 \end{bmatrix} & x_{d,t} = faulty. \end{cases} \tag{24}$$

The discrete state evolves according to the transition matrix:

$$T = \begin{bmatrix} 0.999 & 0.001 \\ 0.0 & 1.0 \end{bmatrix}. \tag{25}$$

We consider the problem where the car is initially at rest and must travel to the goal and stop, as illustrated in Fig. 2. Task failure is defined as collision with the wall.

Fig. 3 compares two typical solutions generated by the weighted particle control approach for a maximum probability of failure of 0.01 and 10^{-6} respectively. The more conservative solution takes $9s$, while the aggressive one takes only $6s$. We now demonstrate that the weighted particle control approach enables the controller to take into account the low probability brake failures. Fig. 4 compares two typical solutions generated with and without weighting respectively.

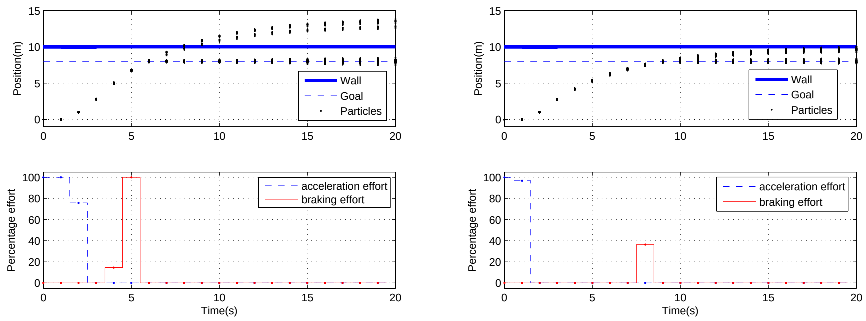

Fig. 3. Two typical solutions with 100 particles. Left: Maximum probability of task failure set to 0.01. The vehicle arrives at the goal within $6s$, but will collide with the wall if a brake failure occurs at or before $5s$. This particular solution gives a true probability of task failure of approximately 0.006. Right: Maximum probability of task failure set to 10^{-6}. The vehicle travels more slowly and arrives later than with the more aggressive solution. In the case of brake failure, however, friction brings the vehicle to rest before collision with the wall. This solution is therefore robust to brake failure, giving a probability of task failure of approximately 1.0×10^{-6}

In the unweighted case, the algorithm did not sample any of the failure transitions and so has generated an inappropriately aggressive control policy that does not take into account the possibility of brake failure. By increasing the probability of sampling failure transitions, the weighted algorithm by contrast has taken into account brake failure, generating an appropriately conservative plan.

Fig. 5 compares the weighted particle control approach against the unweighted particle control approach in terms of the true probability of task failure. In this example the desired probability of task failure was 10^{-6}. The weighted approach achieves a true probability of failure dramatically closer to the desired value than the unweighted approach. Notice also that for larger particle sets the unweighted case approaches the weighted one, except that the variance is much greater in the unweighted case. This is because on the rare occasion that brake failure transitions are sampled, the solution is very different from the average case. This variance is particularly undesirable for control. Fig. 5 also shows the solution time as a function of the number of weighted particles used. Solutions were found in seconds even for relatively large particle sets.

8 Conclusion

In this paper we have presented a novel approach to optimal stochastic control for Jump Markov Linear Systems that takes into account probabilistic uncertainty due to disturbances, uncertain state estimation, modeling error and stochastic mode transitions. The new weighted particle control method is robust in ensur-

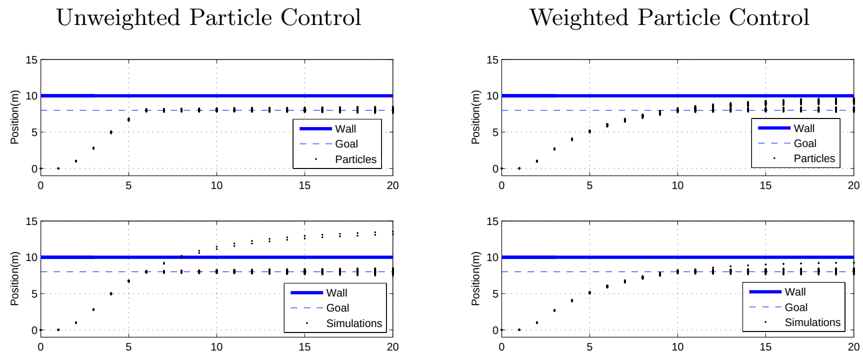

Fig. 4. Typical solutions with and without weighting for $\delta = 10^{-6}$ and 100 particles. The top row shows the particles used for planning, while the bottom row shows Monte-Carlo simulations of the true state trajectory. **Left:** *Without* weighting, no particles have sampled the brake failure so the controller plans aggressively. In reality, there is a probability of approximately 0.0050 that a brake failure occurs at or before $t = 5s$, causing the vehicle to collide with the wall. **Right:** *With* weighting, many particles have sampled brake failures, hence the controller plans taking brake failures into account. The controller is less aggressive, giving a collision probability of approximately 1.0×10^{-6}.

ing that the probability of task failure is less than a defined threshold δ. By approximating the original stochastic problem as a deterministic one using a number of importance-weighted particles, the approach is able to handle arbitrary probability distributions. Furthermore the approximation error tends to zero as the number of particles tends to infinity. Importance weighting is used in conjunction with sampling from a proposal distribution to make sure that the method takes into account low probability events such as component failures.

References

1. Hofbaur, M.W., Williams, B.C.: Hybrid estimation of complex systems. In: IEEE Trans. on Systems, Man, and Cybernetics - Part B: Cybernetics. (2004)
2. Oh, S.M., Rehg, J.M., Balch, T., Dallaert, F.: Learning and inference in parametric switching linear dynamic systems. In: Proc. IEEE Int. Conf. Comp. Vision. (2005)
3. Blackmore, L., Rajamanoharan, S., Williams, B.C.: Active estimation for switching linear dynamic systems. In: Proc. of the CDC. (2006)
4. Pham, H.: On some recent aspects of stochastic control and their applications. Probability Surveys **2** (2005) 506–549
5. Costa, O.L.V., Fragoso, M.D., Marques, R.P.: Discrete-Time Markovian Jump Linear Systems. Springer-Verlag (2005)
6. Schwarm, A., Nikolaou, M.: Chance-constrained model predictive control. AlChE Journal **45** (1999) 1743–1752
7. Li, P., Wendt, M., Wozny, G.: A probabilistically constrained model predictive controller. Automatica **38** (2002) 1171–1176

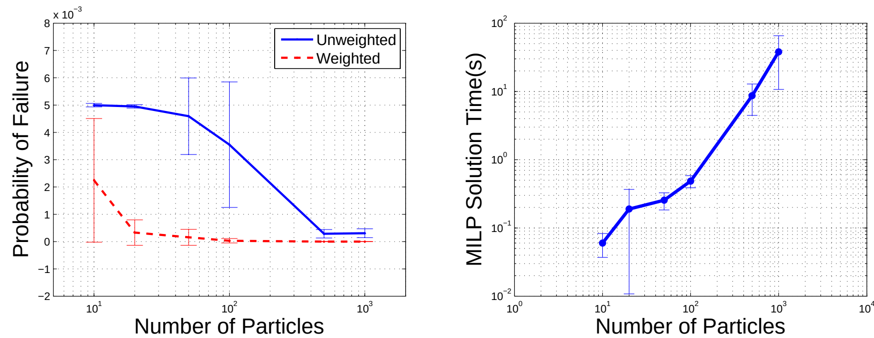

Fig. 5. Left: True probability of failure against number of particles for weighted method and unweighted method. The desired probability of failure was 10^{-6}. The weighted approach achieves a true probability of failure dramatically closer to the desired value than the unweighted approach. With a very small particle set, the effect of weighting is diminished since the probability of sampling the nominal sequence must be high in order to satisfy constraints on the probability of a feasible solution. **Right:** MILP solution time for weighted particle control with ground vehicle scenario using ILOG CPLEX 9.0 on Intel Pentium 4 2.8GHz machine with 1GB RAM. The specified maximum probability of task failure was 0.01.

8. Batina, I.: Model Predictive Control for Stochastic Systems by Randomized Algorithms. PhD thesis, TU Eindhoven, The Netherlands (2004)
9. Hessem, D.H.V.: Stochastic Inequality Constrained Closed-loop Model Predictive Control. PhD thesis, Technische Universiteit Delft, Delft, The Netherlands (2004)
10. Blackmore, L., Li, H.X., Williams, B.C.: A probabilistic approach to optimal robust path planning with obstacles. In: Proceedings of the ACC. (2006)
11. Vargas, A.N., Furloni, W., do Val, J.B.R.: Constrained MPC of jump linear systems with noise and non-observed markov state. In: Proc. ACC. (2006)
12. Blackmore, L.: A probabilistic particle control approach to optimal, robust predictive control. In: Proc. of the AIAA GNC Conference. (2006)
13. Metropolis, N., Ulam, S.: The Monte Carlo method. American Statistical Association **44** (1949) 335–341
14. Morales-Menendez, R., de Freitas, N., Poole, D.: Real-time monitoring of complex industrial processes with particle filters. In: Proceedings of NIPS. (2002)
15. Doucet, A., de Freitas, N., Murphy, K., Russell, S.: Rao-Blackwellised particle filtering for dynamic Bayesian networks. In: Proceedings of UAI. (2000)
16. Gordon, N.J., Salmond, D.J., Smith, A.F.M.: Novel approach to nonlinear/non-Gaussian Bayesian state estimation. IEE Proceedings - F **140**(2) (1993) 107–113
17. Bemporad, A., Cairano, S.D.: Optimal control of discrete hybrid stochastic automata. In Morari, M., Thiele, L., eds.: Lecture Notes in Computer Science 3414. Springer, Berlin (2005) 151–167
18. Schouwenaars, T., Moor, B.D., Feron, E., How, J.: Mixed integer programming for multi-vehicle path planning. In: Proc. European Control Conference. (2001)
19. Doucet, A., de Freitas, N., Gordon, N.J.: Sequential Monte Carlo Methods in Practice. Springer Verlag (2001)
20. de Freitas, N.: Rao-Blackwellised particle filtering for fault diagnosis. IEEE Aero. (2002)