

A Probabilistic Particle Control Approach to Optimal, Robust Predictive Control

Lars Blackmore

Autonomous vehicles need to be able to plan trajectories to a specified goal that avoid obstacles, and are robust to the inherent uncertainty in the problem. This uncertainty arises due to uncertain state estimation, disturbances and modeling errors. Previous solutions to the robust path planning problem solved this problem using a finite horizon optimal stochastic control approach. This approach finds the optimal path subject to chance constraints, which ensure that the probability of collision with obstacles is below a given threshold. This approach is limited to problems where all uncertain distributions are Gaussian, and typically result in highly conservative plans. In many cases, however, the Gaussian assumption is invalid; for example in the case of localization, the belief state about a vehicle's position can consist of highly non-Gaussian, even multimodal, distributions. In this paper we present a novel method for finite horizon stochastic control of dynamic systems subject to chance constraints. The method approximates the distribution of the system state using a finite number of particles. By expressing these particles in terms of the control variables, we are able to approximate the original stochastic control problem as a deterministic one; furthermore the approximation becomes exact as the number of particles tends to infinity. For a general class of chance constrained problems with linear system dynamics, we show that the approximate problem can be solved using efficient Mixed-Integer Linear Programming techniques. We apply the new method to aircraft control in turbulence, and show simulation results that demonstrate the efficacy of the approach.

I. Introduction

Stochastic control of dynamic systems has received a great deal of attention in recent years.¹⁻⁴ Predictive stochastic control takes into account probabilistic uncertainty in dynamic systems and aims to control the predicted distribution of the system state in some manner.

One important application of predictive stochastic control is robust path planning for vehicles under uncertainty. Uncertainty arises due to uncertain state estimation, disturbances, and modeling errors. Predictive stochastic control can be used to plan the predicted distribution of the vehicle state, in order to ensure that the probability of failure is below a given threshold; this is known as a *chance constraint*.³ Failure can be defined as collision with an obstacle, or failure to reach a goal region. Optimal predictive stochastic control aims to find the *best* plan, in order to minimize fuel use, or time to completion, for example.

Predictive optimal stochastic control is a very challenging problem, since we must optimize over the space of possible future state distributions. This space is typically very large indeed. Previous approaches aim to make this problem tractable by discretizing this space, assuming a finite number of possible actions, and applying Dynamic Programming to compute the value of each action for each belief state.^{5,6} This typically leads to a vast Dynamic Programming problem that is impractical for on-line solution even with the discretization assumptions. Furthermore, the discretization assumptions, in particular the assumption that the number of possible actions is finite, are overly restrictive for a system such as an aircraft operating in the real world.

Previous work has developed alternative approaches that use constrained optimization to solve chance constrained stochastic control problems.⁷⁻⁹ These methods assume that the system under control is linear and that all sources of uncertainty are Gaussian. Extending this work, (10) showed that robust vehicle path planning can be posed as a chance-constrained stochastic control problem and solved using efficient Disjunctive Linear Programming techniques.¹¹ This method can optimize over continuous, i.e. non-discretized, actions and belief states, while being efficient enough to generate control actions online.

These methods, however, rely on the assumption that all uncertainty is described using Gaussian distributions. In this work we describe a new approach that does not rely on the Gaussian assumption. The key idea behind the new approach is to approximate all probability distributions using a finite set of samples, or ‘particles’.¹² We then 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. This method can handle arbitrary, even multimodal, distributions, and in principle can deal with nonlinear systems. However we show that in the case of linear systems, the resulting optimization problem can be solved efficiently using Mixed Integer Linear Programming (MILP).¹³

Particle-based methods have been used extensively for estimation.^{14–18} These particle filtering methods have been shown to be superior to traditional Kalman Filtering methods for non-Gaussian probability distributions.¹² Control methods based on sampling have been investigated by a number of authors.^{19–22} For example, (20) proposed a method for finite horizon control that approximates the expected cost of a given control sequence using a finite number of samples. (21) and (22) use particles to approximate the value function and its gradient in an optimal stochastic control problem, and use a gradient search method to find a locally optimal solution. We extend the prior art in two significant ways. First, we incorporate constraints on the probability of failure, enabling *robust* stochastic control, and second, we show that in the case of stochastic linear dynamic systems, the resulting robust control problem can be solved to *global* optimality using Mixed Integer Linear Programming in an extremely efficient manner.

We demonstrate the new method in simulation using two scenarios. In the first scenario we consider altitude control of an aircraft within a convex region, subject to turbulence and uncertain localization. In the second scenario we consider path planning for an aircraft in two dimensions, subject to obstacles, wind disturbances and uncertain localization. The results show that the method is effective in solving the approximated stochastic control problem, and that for a large number of particles the approximation error becomes small.

II. Problem Statement

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

Design a finite, optimal sequence of control inputs, taking into account probabilistic uncertainty, which ensures that the state of the system leaves a given feasible region with probability at most δ .

Here, optimality can be defined in terms of minimizing control effort, for example. In the case of vehicle path planning, the feasible region can be defined so that the system state remains in a goal region at the final time step, and avoids a number of obstacles at all time steps. We consider three sources of uncertainty:

1. The initial state is specified as a probabilistic distribution over possible states. This uncertainty arises due to partial observability of the system state, which means that we must estimate a distribution over the system state from noisy observations.
2. The system model is not known exactly. This may arise due to modeling errors or linearization.
3. Disturbances act on the system state. These are modeled as stochastic processes. In the case of aircraft path planning, this may represent wind or turbulence disturbances.

We assume that the distributions of the uncertainty mentioned here are known at least approximately, but we make no assumptions about the form the distributions take.

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 III we review some results relating to sampling from random variables. In Section IV we then describe the new approach in detail.

III. 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.¹² Here we review some properties of samples drawn from random variables.

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

$$E_X[f(X)] = \int f(\mathbf{x})p(\mathbf{x})d\mathbf{x} \quad (1)$$

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

$$\hat{E}_X[f(X)] = \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}^{(i)}) \quad (2)$$

From the strong law of large numbers, the sample mean converges to the true expectation as N tends to infinity.

$$\hat{E}_X[f(X)] \longrightarrow E_X[f(X)] \quad (3)$$

Hence the expectation, which could not be evaluated exactly in closed form, can be approximated as a summation over a number of particles. 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:

$$P_A = \int_{f(\mathbf{x}) \in A} p(\mathbf{x})d\mathbf{x} \quad (4)$$

This expression is equivalent to the expectation:

$$P_A = E_X[g(\mathbf{x})], \quad (5)$$

where:

$$g(\mathbf{x}) = \begin{cases} 1 & f(\mathbf{x}) \in A \\ 0 & f(\mathbf{x}) \notin A. \end{cases} \quad (6)$$

We can therefore approximate P_A as:

$$\hat{P}_A = \frac{1}{N} \sum_{i=1}^N g(\mathbf{x}^{(i)}) = \frac{1}{N} |\{f(\mathbf{x}^{(i)}) \in A, i = 1, \dots, N\}|, \quad (7)$$

where $|\{f(\mathbf{x}^{(i)}) \in A, i = 1, \dots, N\}|$ denotes 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 (4) 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:

$$\hat{P}_A \longrightarrow P_A, \quad (8)$$

as N tends to infinity. In Section IV we use this property to approximate the stochastic control problem defined in Section II.

IV. Particle Optimization for Robust Stochastic Control

In this section we describe a new method for solving the robust stochastic control problem described in Section II. In Section A we summarize the general method and explain the key ideas, while in Section B we show that for the case of linear dynamic systems and polygonal constraints, the resulting optimization can be solved efficiently using MILP.

A. Outline of General Method

The key observation behind the new 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. Consider a discrete-time dynamic system where the future states $\mathbf{x}_1, \dots, \mathbf{x}_T$ are functions of the control inputs $\mathbf{u}_0, \dots, \mathbf{u}_{T-1}$, the initial state \mathbf{x}_0 , and disturbances ν_0, \dots, ν_{T-1} .

$$\begin{aligned} \mathbf{x}_1 &= f_1(\mathbf{x}_0, \mathbf{u}_0, \nu_0) \\ \mathbf{x}_2 &= f_2(\mathbf{x}_0, \mathbf{u}_0, \mathbf{u}_1, \nu_0, \nu_1) \\ &\vdots \\ \mathbf{x}_T &= f_T(\mathbf{x}_0, \mathbf{u}_0, \dots, \mathbf{u}_{T-1}, \nu_0, \dots, \nu_{T-1}) \end{aligned} \quad (9)$$

The initial state and disturbances are uncertain, but are modeled as random variables with known distributions. Hence the future states are also random variables, whose distributions depend on the control inputs. We assume that the initial state and disturbances are independent. Modeling errors can be modeled as an additional stochastic disturbance process.²³ For notational simplicity we assume for the rest of the development a single disturbance process; however the method applies equally to multiple disturbance processes, and hence to modeling errors as well as external disturbances.

Consider a stochastic control problem where the objective is to minimize $h(\mathbf{u}_0, \dots, \mathbf{u}_{T-1}, \mathbf{x}_{1:T})$ while ensuring that the future state falls outside of the feasible region F , which need not be convex, with a probability at most δ . Here, F is the feasible region for a state *trajectory*, and $\mathbf{x}_{1:T}$ is a block vector describing the state trajectory:

$$\mathbf{x}_{1:T} = \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_T \end{bmatrix} \quad (10)$$

The state trajectory falls outside of the feasible region if and only if $\mathbf{x}_{1:T} \notin F$. The new particle control method is as follows:

1. Generate N samples from the joint distribution of initial state and disturbances.

Since these are independent, this corresponds to generating two separate sets of samples $\{\mathbf{x}_0^{(1)}, \dots, \mathbf{x}_0^{(N)}\}$ and $\{\nu_0^{(1)}, \dots, \nu_0^{(N)}, \dots, \nu_{T-1}^{(1)}, \dots, \nu_{T-1}^{(N)}\}$ drawn from the initial state and disturbance distributions respectively.

2. Express the distribution of the future state trajectories approximately as a set of N particles, where each particle $\mathbf{x}_{1:T}^{(i)}$ corresponds to the state trajectory given a particular set of samples $\{\mathbf{x}_0^{(i)}, \nu_0^{(i)}, \dots, \nu_{T-1}^{(i)}\}$. Each particle *depends explicitly on the control inputs* $\mathbf{u}_0, \dots, \mathbf{u}_{T-1}$, which are yet to be generated.

$$\mathbf{x}_{1:T}^{(i)} = \begin{bmatrix} \mathbf{x}_1^{(i)} \\ \mathbf{x}_2^{(i)} \\ \vdots \\ \mathbf{x}_T^{(i)} \end{bmatrix} \quad \mathbf{x}_t^{(i)} = f_t(\mathbf{x}_0^{(i)}, \mathbf{u}_0, \dots, \mathbf{u}_{t-1}, \nu_0^{(i)}, \dots, \nu_{t-1}^{(i)}) \quad (11)$$

In Equation 11, $\mathbf{x}_0^{(i)}$ and $\nu_0^{(i)}, \dots, \nu_{t-1}^{(i)}$ are *known values* sampled from random variables, whereas $\mathbf{u}_0, \dots, \mathbf{u}_{t-1}$ are decision variables over which to optimize.

3. Approximate the chance constraints in terms of the generated particles.

The original chance constraint is that the state trajectory $\mathbf{x}_{1:T}$ must fall outside of the feasible region with a probability at most δ :

$$p(\mathbf{x}_{1:T} \notin F) \leq \delta \quad (12)$$

Using the result in Equation 7 the probability of failure is approximated as follows:

$$p(\mathbf{x}_{1:T} \notin F) = \int_{\mathbf{x}_{1:T} \notin F} p(\mathbf{x}_{1:T}) d\mathbf{x}_{1:T} \approx \frac{1}{N} |\{\mathbf{x}_{1:T}^{(i)} \notin F, i = 1, \dots, N\}|, \quad (13)$$

where $\mathbf{x}_{1:T}^{(i)}$ is defined in Equation 11. From Equation 12 and Equation 13, the approximated chance constraint then becomes:

$$\frac{1}{N} |\{\mathbf{x}_{1:T}^{(i)} \notin F, i = 1, \dots, N\}| \leq \delta, \quad (14)$$

where $|\{\mathbf{x}_{1:T}^{(i)} \notin F, i = 1, \dots, N\}|$ denotes the number of particles outside of the feasible region. In other words, the approximate chance constraint is that a fraction of no more than δ of the particles can fall outside of the feasible region. Note that a particle represents a state *trajectory* over the entire planning horizon.

4. Approximate the cost function in terms of particles

$$\hat{h}(\mathbf{u}_0, \dots, \mathbf{u}_{T-1}, \mathbf{x}_{1:T}^{(1)}, \dots, \mathbf{x}_{1:T}^{(N)}) \approx h(\mathbf{u}_0, \dots, \mathbf{u}_{T-1}, \mathbf{x}_{1:T}) \quad (15)$$

5. Solve the deterministic constrained optimization problem for control inputs $\mathbf{u}_0, \dots, \mathbf{u}_{T-1}$:

$$\begin{aligned} & \text{Minimize } \hat{h}(\mathbf{u}_0, \dots, \mathbf{u}_{T-1}, \mathbf{x}_{1:T}^{(1)}, \dots, \mathbf{x}_{1:T}^{(N)}) \\ & \text{Subject to:} \end{aligned}$$

$$\frac{1}{N} |\{\mathbf{x}_{1:T}^{(i)} \notin F, i = 1, \dots, N\}| \leq \delta \quad (16)$$

The method is illustrated in Figure 1. This is a general formulation that can encompass a very broad range of chance-constrained problems. It is not necessarily true, however, that the optimization problem that results from this formulation is tractable. In Section B we show that in the case of linear system dynamics and polygonal feasible regions, the optimization can be solved using efficient Mixed-Integer Linear Programming methods.

B. Robust Control of Linear Systems

We now restrict our attention to the case of linear system dynamics and polygonal feasible regions. Furthermore, we assume that the cost function h is piecewise linear. Previous work has shown that optimal path planning with obstacles for vehicles such as aircraft or satellites can be posed as finite horizon control design for linear systems in polygonal feasible regions.^{24,25} Optimality can be defined in terms of fuel use or time, for example. We extend this work by showing that the particle control method outlined in Section A can be used to design control inputs for linear systems that are robust to probabilistic uncertainty in the initial state and disturbances.

We consider the linear discrete time system model:

$$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \nu_t \quad (17)$$

Substituting this system model into Equation 11 we obtain the following equation for $\mathbf{x}_t^{(i)}$:

$$\mathbf{x}_t^{(i)} = \sum_{j=0}^{t-1} A^{t-j-1} B(\mathbf{u}_j + \nu_j^{(i)}) + A^t \mathbf{x}_0^{(i)} \quad (18)$$

Note that this is a linear function of the control inputs, and that $\mathbf{x}_0^{(i)}$ and $\nu_j^{(i)}$ are known values. Hence each particle $\mathbf{x}_{1:T}^{(i)}$ is a known linear function of the control inputs.

In accordance with Equation 14, we need to constrain the number of particles that fall outside of the feasible region. We define a set of N binary variables z_1, \dots, 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:

$$\frac{1}{N} \sum_{i=1}^N z_i \leq \delta \quad (19)$$




Figure 1. Illustration of new particle control method in a 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.

This constraint ensures that the fraction of particles falling outside of the feasible region is at most δ . We now describe how to impose constraints such that:

$$z_i = 0 \implies \mathbf{x}_{1:T}^{(i)} \in F, \quad (20)$$

first for convex polygonal feasible regions, then for non-convex polygonal feasible regions.

1. Convex Feasible Regions

An example of control within a convex feasible region is given in Figure 2. In this example, the state at each time step, \mathbf{x}_t , is constrained to lie within a convex feasible region F_t . Together, these convex regions defined for $t = 1, \dots, T$ form a convex feasible region in the space of state trajectories. The feasible region for the problem in Figure 2 is therefore convex in the space of state trajectories. As shown in Figure 2, defining feasible regions at each time step can be used to ensure that the system state is taken from its initial value to a defined goal region, and stays within specified bounds at all times.

A convex polygonal feasible region F_t defined for \mathbf{x}_t can be defined as a conjunction of linear constraints $\mathbf{a}_{tl}^T \mathbf{x}_t \leq b_{tl}$ for $l = 1, \dots, N_t$, where \mathbf{a}_{tl} is defined as pointing outwards from the polygonal region. Then \mathbf{x}_t lies within F_t if and only if all of the constraints are satisfied:

$$\mathbf{x}_t \in F_t \iff \bigwedge_{l=1, \dots, N_t} \mathbf{a}_{tl}^T \mathbf{x}_t \leq b_{tl}. \quad (21)$$

This is illustrated in Figure 3. We now impose the following constraint:

$$\mathbf{a}_{tl}^T \mathbf{x}_t^{(i)} - b_{tl} \leq M z_i \quad \forall t \forall l, \quad (22)$$




Figure 2. Control within a convex polygonal feasible region. At each time step the state must remain within a convex polygonal feasible region. In the space of state trajectories, the feasible region is a convex polyhedron. This type of control can be used to ensure that the state ends up in the goal region and places hard constraints on the state until then.




Figure 3. Left: Two-dimensional convex polygonal feasible region F_t defined for the state at time t . The vectors $\mathbf{a}_{t1}, \dots, \mathbf{a}_{tN_t}$ are the unit outward normals to the N_t line segments that define the feasible region. Right: Two-dimensional obstacle O_j modeled as a convex polyhedron. The vectors $\mathbf{a}_{j1}, \dots, \mathbf{a}_{jN_j}$ are the unit outward normals to the N_j line segments that define the obstacle.

where M is a large positive constant. A value of $z_i = 0$ implies that every constraint is satisfied for every time step for particle i , whereas for large enough M , a value of $z_i = 1$ leaves particle i unconstrained. We therefore have:

$$z_i = 0 \implies \mathbf{x}_t^{(i)} \in F_t \quad \forall t \implies \mathbf{x}_{1:T}^{(i)} \in F, \quad (23)$$

as required.

2. Non-convex Feasible Regions

Predictive control of linear systems within a non-convex feasible region is a much more challenging problem than control within a convex feasible region.²⁵ However, as shown by (24), vehicle path planning with obstacles can be posed as such a problem, hence it is of great interest.

A polygonal non-convex feasible region can be described as the complement of a number of polygonal infeasible regions, or obstacles. In other words, the state trajectory $\mathbf{x}_{1:T}$ is in the feasible region if and only if all obstacles are avoided for all time steps.

As noted by (24), avoidance of a polygonal obstacle can be expressed in terms of a disjunction of linear constraints. That is, the system state at time t , \mathbf{x}_t , avoids the obstacle O_j shown in Figure 3 if and only if:

$$\bigvee_{l=1, \dots, N_j} \mathbf{a}_{jl}^T \mathbf{x}_t \geq b_{jl} \quad (24)$$

In a similar manner to (24), we introduce binary variables $d_{ijtl} \in \{0, 1\}$ that indicate whether a given constraint l for a given obstacle O_j is satisfied by a given particle i at a given time step t . The constraint:

$$\mathbf{a}_{jl}^T \mathbf{x}_t^{(i)} - \mathbf{b}_{jl} + M d_{ijtl} \geq 0, \quad (25)$$

means that $d_{ijtl} = 0$ implies that constraint l in obstacle O_j is satisfied by particle i at time step t . Again M is a large positive constant.

We now introduce binary variables $e_{ijt} \in \{0, 1\}$ that indicate whether a given obstacle O_j is avoided by a given particle i at a given time step t . The constraint:

$$\sum_{l=1}^{N_j} d_{ijtl} - (N_j - 1) \leq M e_{ijt}, \quad (26)$$

ensures that $e_{ijt} = 0$ implies that at least one constraint in obstacle O_j is satisfied by particle i at time step t . This in turn implies that obstacle O_j is avoided by particle i at time step t .

Next, we introduce binary variables $g_{ij} \in \{0, 1\}$ that indicate whether a given obstacle O_j is avoided by particle i at *all* time steps. The constraint:

$$\sum_{t=1}^T e_{ijt} \leq M g_{ij}, \quad (27)$$

ensures that $g_{ij} = 0$ implies that obstacle j is avoided at all time steps by particle i .

Finally, we introduce the constraint:

$$\sum_{j=1}^L g_{ij} \leq M z_i. \quad (28)$$

This ensures that $z_i = 0$ implies that all obstacles are avoided at all time steps by particle i . This in turn ensures that, for non-convex feasible regions F_t ,

$$z_i = 0 \implies \mathbf{x}_t^{(i)} \in F_t \quad \forall t \implies \mathbf{x}_{1:T}^{(i)} \in F, \quad (29)$$

as required.

3. Cost Function

The cost function h can be defined to be a function of both the control inputs u_0, \dots, u_{T-1} and the system state trajectory $\mathbf{x}_{1:T}$. Since the system state is uncertain, however, a cost function involving the system state must be defined in terms of expected cost. In this case we approximate the expectation using the sample mean as in Equation 2. This is evaluated using the particle population as follows. The true expectation is given by:

$$E[h] = \int h(\mathbf{u}_0, \dots, \mathbf{u}_{T-1}, \mathbf{x}_{1:T}) p(\mathbf{x}_{1:T}) d\mathbf{x}_{1:T} \quad (30)$$

Since $p(\mathbf{x}_{1:T})$ can be an arbitrary distribution, this integral is intractable in most cases. The approximated expectation is given by:

$$\hat{h} = \frac{1}{N} \sum_{i=1}^N h(\mathbf{u}_0, \dots, \mathbf{u}_{T-1}, \mathbf{x}_{1:T}^{(i)}), \quad (31)$$

and this expression can be evaluated without integration. As the number of particles tends to infinity, we have that:

$$\hat{h} \longrightarrow E[h] \quad (32)$$

Furthermore, since we assume that h is a piecewise linear function of the state and control inputs, the expression for \hat{h} in Equation 31 is also piecewise linear. Hence the approximate cost function \hat{h} is suitable for use within a Mixed Integer Linear Program.

C. Summary

To summarize, the approximated stochastic predictive control problem defined in Section A can be posed as a Mixed Integer Linear Program in the case of linear system dynamics, linear cost function, and a polygonal feasible region. The resulting optimization finds the best sequence of control inputs such that at most a fraction δ of the particles falls outside of the feasible region. This fraction approximates the probability of the future state trajectory falling outside of the feasible region, and as the number of particles tends to infinity, the approximation becomes exact. The optimality criterion is also approximated in terms of particles, and the approximation becomes exact as the number of particles tends to infinity.

We have therefore introduced a new method for robust optimal control of stochastic linear systems, where the probability distributions of uncertain variables can take an arbitrary form. We approximate the probability of failure and constrain this value to be below a specified threshold. By adjusting the desired probability of failure, it is possible to adjust the level of conservatism in the plan; this is demonstrated in Section V.

In the case of vehicle path planning, we can specify obstacles to be avoided and goal regions in which the state must fall at a given time by defining convex and non-convex feasible regions over the system state \mathbf{x}_t at time t . We define failure as collision with an obstacle, or failure to reach a goal region. The new particle control method can then find the optimal path for the vehicle, while taking into account probabilistic uncertainty in the initial state, disturbances, and modeling error, so that the probability of failure is below a specified value. Note that this is in contrast to our previous work,¹⁰ where only the probability of collision with obstacles at a *given* time step could be constrained, and even this led to highly conservative plans.

In Section A we assumed that the probability distributions of all uncertain variables are known. In fact, an approximation of the probability distribution described using a number of samples, is sufficient. This is the case, for example, with vehicle localization, where particle filtering¹² can be used to determine an approximate probability distribution over the vehicle's position. This lends the new particle control method to being used in closed loop with a particle-based estimator, the estimator providing the set of particles describing the initial state for the predictive controller. The particle controller then plans a sequence of future control inputs in order to control the trajectory of these particles.

V. Results

We now present simulation results for the particle control method applied to two different aircraft control scenarios. In Section A the method is used to control an aircraft within a flight envelope in heavy turbulence, while in Section B the method is used to generate robust, optimal paths for an aircraft operating in an environment containing obstacles.

A. Robust Predictive Control in a Convex Feasible Region

The new particle control method was used to generate robust predictive control inputs for an aircraft performing an altitude change maneuver in turbulence. In this scenario, successful execution means that the aircraft remains within a defined flight envelope, which forms a convex region in the space of state trajectories. An example is shown in Figure 4.

Disturbances due to turbulence have been studied extensively and are modeled as stochastic noise processes that are far from Gaussian.²⁶ In this work we use the Dryden turbulence model described in Military Specification MIL-F-8785C.²⁷ We assume heavy turbulence, as defined in MIL-F-8785C, with a low-altitude wind velocity of $25m/s$. Designing control inputs robust to such a non-Gaussian disturbance process is a highly challenging problem, however the particle control formulation is able to deal with this example naturally.

For the aircraft model, we use the linearized, discrete time longitudinal dynamics of a Boeing 747 travelling at Mach 0.8. Time increments of two seconds were used. Since the angle of attack of most aircraft is low in normal operation, linearizing the dynamics about the equilibrium state or *trim state* of the aircraft yields a good approximation to the true dynamics, which can be used to develop control algorithms.²⁸ Consistent with prior approaches to robust predictive control,^{29,30} we assume that there is an inner-loop controller issuing elevator commands, which is an altitude-hold autopilot. The aim of the predictive controller then, is to issue altitude commands to the autopilot in an optimal, robust manner, so that the aircraft breaks the flight envelope with a probability of at most δ . Optimality is defined in terms of fuel consumption, and we

assume the following relationship between fuel consumption F and elevator angle at time t , a_t :

$$F = \sum_{t=0}^{T-1} |a_t| \quad (33)$$

Since we assume an inner loop controller issues elevator angle commands, a_t depends on the disturbances that act on the aircraft; for example, if a fixed altitude is commanded by the predictive controller, the autopilot will issue larger elevator commands in order to reject large disturbances than for smaller ones. Since the disturbances are stochastic, we cannot directly optimize the fuel consumption defined in Equation 33. We can, however, optimize the expected fuel consumption, as described in Section IV.

We assume a maximum elevator deflection of 0.5 radians, due to actuator saturation. Again, since elevator deflection depends on stochastic disturbances, we cannot prevent actuator saturation with absolute certainty. We can, however, define a chance constraint that, approximated using the particle control method, ensures that actuator saturation occurs with at most a given probability. In the results shown here we define this probability to be zero, thereby ensuring that saturation occurs with approximately zero probability.

The initial state distribution was generated using a particle filter. The particle filter tracked the system state for ten time steps leading up to time $t = 0$ while the aircraft held altitude. Observations of pitch rate and altitude were made subject to additive Rayleigh noise.³¹ This non-Gaussian noise means that a particle filter will typically outperform a Kalman Filter. The number of particles used for estimation was the same as that used for control.

Figure 4 shows two typical solutions generated by the particle control algorithm for 100 particles. In the first solution, the desired probability of error is 0.1, while in the second, the desired probability of error is 0.01. It can be seen that in the case of the lower value, the path taken by the particles is further away from the edges of the flight envelope, causing one particle only to fall outside of the envelope. This more conservative plan has a greater fuel cost of 3.82, compared to the less conservative one, which has a fuel cost of 3.62. Hence conservatism can be traded off against fuel efficiency; the empirical relationship between the two for this scenario is shown in Figure 5.

The particle control method solves an approximated stochastic control problem. The accuracy of the approximation was assessed by calculating the true probability of failure for a given plan. This probability was calculated by carrying out a very large number of random simulations. Since the generated plan depends on the values sampled from the various probability distributions, 20 plans were generated for each scenario. Figure 6 shows the results for a desired probability of failure of 0.1. It can be seen that the mean probability of error gets closer to the desired value as the number of particles increases, and that the variance decreases. For 100 particles, the approximation is close; the mean is 0.104 and the standard deviation is 0.024. Hence the particle control algorithm can generate optimal solutions to problems that are very close to the full stochastic control problem.

B. Robust Vehicle Path Planning with Obstacles

The new particle control method was also applied to a UAV path planning scenario with obstacles, wind and uncertain localization. In this scenario, successful execution means that the UAV is in the goal region at the end of the time horizon, and that the UAV avoids all obstacles at all time steps within the horizon.

Previous work has shown that an aircraft operating in a two-dimensional environment can be modeled as a discrete-time linear system in the form of Equation 18.²⁴ We use the same aircraft model and assume a maximum aircraft velocity of $50m/s$, time steps of $1s$, and a planning horizon of $10s$.

As in Section A, the goal of the robust control algorithm is to design a sequence of control inputs so that the probability of failure is at most δ . Uncertain disturbances, due to wind, act on the UAV. We use the Dryden wind model with a low-altitude wind speed of $15m/s$ and light turbulence, as defined in MIL-F-8785C. We assume an inner-loop controller that acts to reject disturbances. Also, as in Section A, uncertainty in localization leads to uncertainty in the initial position of the UAV. The obstacle map used is shown in Figure 7. Optimality was again defined in terms of fuel consumption, which we assume is related to input acceleration as follows:

$$F = \sum_{t=0}^{T-1} |u_{x,t}| + |u_{y,t}|. \quad (34)$$

Here $u_{x,t}$ and $u_{y,t}$ are the commanded accelerations at time t in the x and y directions respectively. In order to reduce the complexity of the resulting MILP problem we employed heuristic pruning techniques to reduce the number of particles and number of obstacles considered that are considered, while still guaranteeing an optimal, feasible solution to the original problem. In this paper we do not describe these techniques in detail; a principled development of this topic is the subject of future research.

Results for the scenario are shown in Figures 7 and 8. 50 particles were used for these examples. Figure 7 shows that if a probability of failure of 0.04 or above is acceptable, the planned path of the UAV can go through the narrow corridor at $(-50, 200)$. It can be seen that exactly two particles collide with the obstacles as expected. For a lower probability of failure, however, the planned path is more conservative as shown in Figure 8. This path avoids the narrow corridor at the expense of fuel efficiency.

These results show that the new particle control method is effective in designing optimal paths for an aircraft that are robust to probabilistic uncertainty.

VI. Future Work

There are a number of promising areas for future research. First, while the particle control approach can typically find optimal controls quickly enough for real-time operation of a system, such as an aircraft, in a *convex* region, this is not the case for control within a *non-convex* region. This is because the resulting MILP problem contains a number of binary variables exponential in the number of particles; this makes the MILP problem intractable for a large number of particles. At the same time, the accuracy of the method depends on having enough particles to approximate the distributions of the random variables accurately enough. Future work will therefore draw upon ideas in particle filtering, where techniques such as importance sampling¹⁴ and Rao-Blackwellisation^{15,32} are used to reduce the number of particles needed to cover the probability space effectively. Second, the true probability of failure is biased towards being greater than the desired probability of failure, as shown in Figure 6. The reason for this bias is now understood and will be addressed in future work. Third, there are many areas other than robust control where considering non-Gaussian uncertainty is typically intractable, for example system identification. Future work will investigate the application of the method developed in this paper to these problems. Finally, the extension of this work to non-linear and time-varying systems is a topic for future research.

VII. Conclusion

In this paper we have presented a novel approach to optimal, robust stochastic control that takes into account probabilistic uncertainty due to disturbances, uncertain state estimation and modeling error so that the probability of failure is less than a defined threshold δ . The new method approximates the original stochastic problem as a deterministic one using a number of particles. By controlling the trajectories of these particles in a manner optimal with regard to the approximated problem, the method generates approximately optimal solutions to the original stochastic problem. Furthermore the approximation error tends to zero as the number of particles tends to infinity. By using a particle-based approach, the new particle control method is able to handle arbitrary probability distributions. We demonstrate the method in simulation with two aircraft control scenarios and show that the true probability of failure tends to the desired probability of failure as the number of particles used increases.

References

- ¹H. Pham, "On some recent aspects of stochastic control and their applications," *Probability Surveys*, vol. 2, pp. 506–549, 2005.
- ²D. Nikovski and M. Brand, "Non-linear stochastic control in continuous state spaces by exact integration in bellman's equations," in *Proceedings of the Workshop on Planning under Uncertainty and Incomplete Information of ICAPS*, 2003.
- ³A. Prekopa, *Stochastic Programming*. Kluwer, 1995.
- ⁴A. Aurnhammer and K. Marti, *Dynamic Stochastic Optimization*. Berlin: Springer, 2003, pp. 133–154.
- ⁵N. Roy, G. Gordon, and S. Thrun, "Planning under uncertainty for reliable health care robotics," in *Proceedings of the International Conference on Field and Service Robotics*, 2003.
- ⁶P. Lommel, "An extended kalman filter extension of the augmented markov decision process," Master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2004.
- ⁷D. H. V. Hessem, "Stochastic inequality constrained closed-loop model predictive control," Ph.D. dissertation, Technische Universiteit Delft, Delft, The Netherlands, 2004.
- ⁸Y. Yan and R. R. Bitmead, "Incorporating state estimation into model predictive control and its application to network traffic control," *Automatica*, vol. 41, pp. 595–604, 2005.
- ⁹P. Li, M. Wendt, and G. Wozny, "Robust model predictive control under chance constraints," *Computers and Chemical Engineering*, vol. 24, pp. 829–834, 2000.
- ¹⁰L. Blackmore, H. X. Li, and B. C. Williams, "A probabilistic approach to optimal robust path planning with obstacles," in *Proceedings of the American Control Conference*, 2006.
- ¹¹E. Balas, "Disjunctive programming," *Annals of Discrete Math*, vol. 5, pp. 3–51, 1979.
- ¹²A. Doucet, N. de Freitas, and N. J. Gordon, *Sequential Monte Carlo Methods in Practice*. Springer Verlag, 2001.
- ¹³ILOG, "Ilog cplex user's guide," 1999.
- ¹⁴N. Metropolis and S. Ulam, "The monte carlo method," *American Statistical Association*, vol. 44, pp. 335–341, 1949.
- ¹⁵R. Morales-Menendez, N. de Freitas, and D. Poole, "Real-time monitoring of complex industrial processes with particle filters," in *Proceedings of Neural Information Processing Systems (NIPS)*, 2002.
- ¹⁶A. Doucet, N. de Freitas, K. Murphy, and S. Russell, "Rao-Blackwellised particle filtering for dynamic Bayesian networks," in *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2000.
- ¹⁷N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proceedings - F*, vol. 140, no. 2, pp. 107–113, Apr. 1993.
- ¹⁸L. Blackmore, S. Funiak, and B. C. Williams, "Combining stochastic and greedy search in hybrid estimation," in *Proceedings of 20th National Conference on Artificial Intelligence (AAAI05)*, 2005.
- ¹⁹A. Y. Ng and M. Jordan, "Pegasus: A policy search method for large mdps and pomdps," in *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, 2000.
- ²⁰A. Greenfield and A. Brockwell, "Adaptive control of nonlinear stochastic systems by particle filtering," in *Proceedings of the Conference on Control and Automation*, 2003.
- ²¹C. Andrieu, A. Doucet, S. S. Singh, and V. B. Tadic, "Particle methods for change detection, system identification, and control," *Proceedings of the IEEE*, vol. 92, 2004.
- ²²S. S. Singh, N. Kantas, B. Vo, A. Doucet, and R. Evans, "Simulation-based optimal sensor scheduling with application to observer trajectory planning," *To appear in Automatica*.
- ²³L. Ljung, *System Identification: Theory for the user*. New York: Prentice Hall, 1999.
- ²⁴T. Schouwenaars, B. D. Moor, E. Feron, and J. How, "Mixed integer programming for multi-vehicle path planning," in *Proc. European Control Conference*, 2001.
- ²⁵A. Richards, J. How, T. Schouwenaars, and E. Feron, "Plume avoidance maneuver planning using mixed integer linear programming," in *Proc. AIAA Guidance, Navigational and Control Conference*, 2001.
- ²⁶N. M. Barr, D. Gangsaas, and D. R. Schaeffer, "Wind models for flight simulator certification of landing and approach guidance and control systems," Paper FAA-RD-74-206, 1974.
- ²⁷"US military handbook MIL-HDBK-1797," 1997.
- ²⁸D. McRuer, I. Ashkenas, and D. Graham, *Aircraft Dynamics and Automatic Control*. Princeton University Press, 1990.
- ²⁹A. Richards, "Robust constrained model predictive control," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2005.
- ³⁰J. Gossner, B. Kouvaritakis, and J. Rossiter, "Stable generalized predictive control with constraints and bounded disturbances," *Automatica*, vol. 33, p. 551, 1997.
- ³¹J. G. Proakis, *Digital Communications*. New York, 1995.
- ³²H. Akashi and H. Kumamoto, "Random sampling approach to state estimation in switching environments," *Automatica*, vol. 13, pp. 429–434, 1977.





Figure 4. Path of particles for two typical solutions to the flight envelope scenario. 100 particles were used. Top: The desired probability of failure is 0.1, and ten particles fall outside of the flight envelope. Bottom: The desired probability of failure is 0.01, and one particle falls outside of the flight envelope, at $t = 18$ s.




Figure 5. Mean fuel cost against desired probability of failure. The mean is taken over 20 different solutions to the same problem. As the allowed probability of failure increases, the plan becomes less conservative and the cost decreases.




Figure 6. True probability of failure against number of particles used to design control inputs. The desired probability of error was 0.1, shown as the dashed line. The results shown are for 20 sets of designed control inputs, with the dots denoting the mean and the error bars denoting the standard deviation of the probability of error. As the number of particles increases, the mean probability of error approaches the desired probability of error and the variance decreases.




Figure 7. Path of particles for typical solution to UAV path planning problem for a probability of failure of 0.04. Obstacles are in blue, while the goal is in red and dashed. At this level of conservatism, the aircraft is able to pass through the narrow corridor. The thin blue dashed box is shown in more detail on the right. The particle control method ensures that at most two particles collide with the obstacles. This solution has a fuel cost of 73.98.




Figure 8. Path of particles for typical solution to UAV path planning problem for a probability of failure of 0.02. Obstacles are in blue, while the goal is in red and dashed. At this level of conservatism, the aircraft no longer passes through the narrow corridor, but goes around the largest obstacle. This solution has a fuel cost of 74.72. Hence the reduced probability of failure comes at the expense of fuel.