Massachusetts Institute of Technology

Optimal, Robust Path Planning – a Probabilistic Approach

Lars Blackmore, Hui Li
and Brian Williams

November 27, 2006

Model-based Embedded & Robotic Systems

Context

- Optimal path planning for dynamic systems
  - *"What is best sequence of control inputs that takes system state from the start to goal?"*

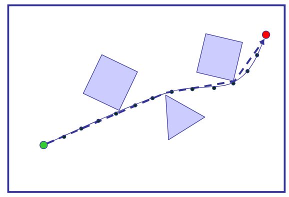

- Non-convex feasible region
- Non-holonomic dynamics
- Discrete-time

- Prior work has solved problem using Disjunctive LP

Optimal Paths are not Robust

- Uncertainty arises due to:
  - Disturbances
  - Uncertain state estimation
  - Inaccurate modeling
- Optimal paths are not robust to uncertainty

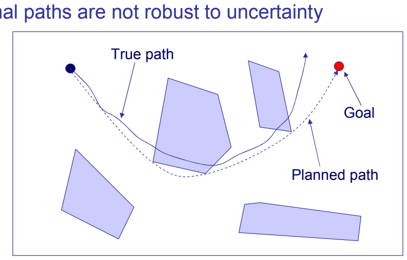

Representing Uncertainty

- Two principal ways to represent uncertainty:

1. Set-bounded uncertainty

$$\mathbf{x}_0 \in S$$

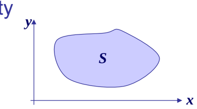

2. Probabilistic uncertainty

$$p(\mathbf{x}_0) = N(\hat{\mathbf{x}}_0, \mathbf{P}_0)$$

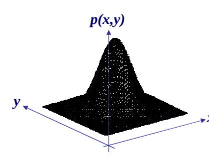

Representing Uncertainty

- Probabilistic representations much richer
  - Set-bounded representation subsumed by p.d.f.
- Probabilistic representations often more realistic
  - What is the absolute maximum possible wind velocity?
- Probabilistic representations readily available in many cases
  - Disturbances ← **e.g. Dryden turbulence model**
  - Uncertain state estimation ← **e.g. Particle filter, Kalman filter, SLAM**
  - Inaccurate modeling ← **e.g. Parameter estimation**
- We deal with **probabilistic uncertainty**

Robust Control under Probabilistic Uncertainty

- Robustness formulated using chance constraints:
  - *"Ensure that failure occurs with probability at most δ"*
- Prior work developed chance-constrained MPC
  - Stochastic problem converted to deterministic problem
  - Deterministic problem solved using LP or QP
  - **Restricted to control within convex feasible region**
- We extend this work to control within non-convex regions → robust path planning with obstacles
  - Resulting problem solved using Disjunctive Linear Programming (DLP) with same complexity as problem without uncertainty

Problem Statement

- Design a finite, optimal sequence of control inputs $\mathbf{u}_0 \ldots \mathbf{u}_{k-1}$ such that the expected final vehicle position is the goal
  - Take into account uncertainty such that collision with any obstacle at a given time step occurs with probability at most δ

Problem Statement

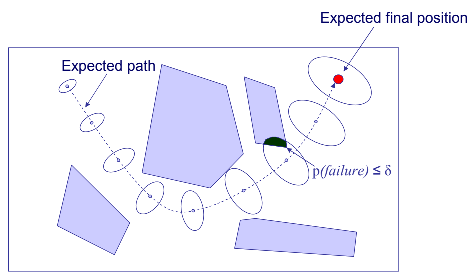

Technical Approach: Assumptions

- Assume discrete-time linear **stochastic** system:

$$x_{t+1} = Ax_t + Bu_t + w_t + v_t$$

Noise due to disturbances (w_t) **Noise due to uncertain model** (v_t)

- Assume that $p(w_t)$ and $p(v_t)$ are known.
- Initial state x_0 is unknown, but assume $p(x_0)$ is known
- Polytopic obstacles
- **All uncertainty is Gaussian**

Technical Approach: Summary

1. Convert stochastic problem into deterministic one
2. Show that deterministic problem can be solved using Disjunctive Linear Programming

Linear Chance Constraints

- Consider linear chance constraint on an uncertain multivariate variable X: "$p(\mathbf{Ax} > \mathbf{b}) \leq \delta$"

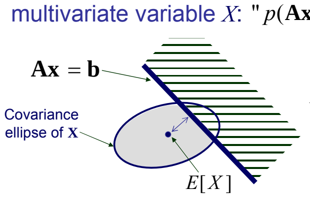

$$p(\mathbf{Ax} > \mathbf{b}) = \int_{\mathbf{Ax}>\mathbf{b}} p(X)d\mathbf{x}$$

- Probability of constraint violation depends on:
  1. Covariance of X
  2. Distance of $\mathrm{E}[X]$ from constraint

Linear Chance Constraints

- So for given covariance $\boldsymbol{P}$, linear chance constraint is equivalent to deterministic linear constraint on $\mathrm{E}[X]$

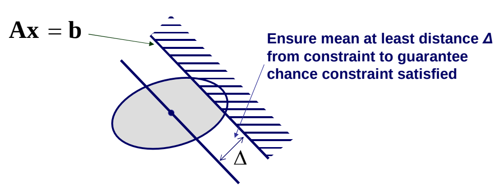

$$p(\mathbf{Ax} > \mathbf{b}) \leq \delta \quad \Leftrightarrow \quad \mathbf{A}E[\mathbf{x}] \leq \mathbf{b} - \Delta$$

Linear Chance Constraints

- How is Δ calculated?
- Find vector normal $\hat{\mathbf{n}}$ to constraint line $\mathbf{Ax} = \mathbf{b}$
- The distribution of $\mathbf{x}$ projected along $\hat{\mathbf{n}}$ is a univariate Gaussian with variance $\hat{\mathbf{n}}^T \mathbf{P} \hat{\mathbf{n}}$

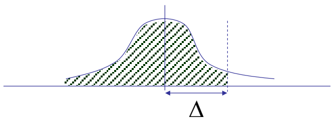

- So Δ can be calculated using a simple lookup of *erf*, the Gaussian c.d.f. function

Obstacle Chance Constraints

- Extension: use this to ensure that an **obstacle** is hit with probability at most δ

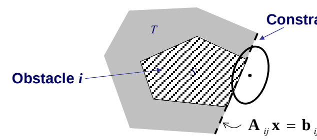

$$\int_S p(X)d\mathbf{x} \le \int_{S \cup T} p(X)d\mathbf{x}$$

- Notice that:

  p(collision with obstacle) $\le$ *p(constraint violated)*

- Simply need to ensure that expected state is Δ away from **at least one** of obstacle's constraints
  - **Conservatism introduced**

Multiple Obstacles

- Must ensure that probability of hitting **any** obstacle is at most δ
  - Probabilities of collision are **not** mutually exclusive

    p(collision with obstacle 1 or 2) $\neq$ *p(collision w. obs. 1) + p(collision w. obs. 2)*

  - But can bound probability of collision with any obstacle

    p(collision with obstacle 1 or 2) $\le$ *p(collision w. obs. 1) + p(collision w. obs. 2)*

- Constrain probability of hitting each of N obstacles to be at most δ/N
  - Then probability of collision with **any** obstacle guaranteed to be at most δ
  - Additional conservatism introduced LB2

Robust Path Planning as DLP

- Important analytic properties:
  - Future state distribution is Gaussian
  - Future state mean is **linear** function of control inputs
  - Future state covariance is **not a function** of control inputs
- Hence the constraint:

  "Ensure expected state is Δ away from **at least one** of obstacle's constraints"
  - is a **disjunctive linear** constraint on control inputs $u_{1\ldots t}$
- Cost functions such as fuel use can be expressed as piecewise linear functions of control inputs

➢ Problem can be posed as a Disjunctive Linear Program

Robust Path Planning as DLP

- Summary:
  - Calculate covariance P_t of predicted state at each t in horizon
  - Calculate required margin Δ_t for each t in horizon to ensure probability of failure less than δ
  - Pose disjunctive linear program to ensure margin satisfied

$$\text{minimize} \quad J = \sum_{i=1}^{T} |\mathbf{u}_i|$$

$$\text{subject to} \quad \bigvee_i \mathbf{A}_{ij} E[X_t] \le \mathbf{b}_{ij} - \Delta_t$$

$$\text{for each time step } t \text{ and each obstacle } j$$

  - Solve using efficient, readily available solvers

Results

- Trade off performance against plan conservatism

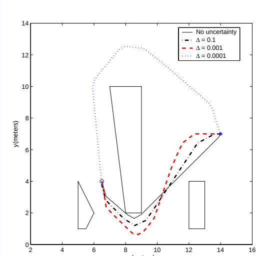

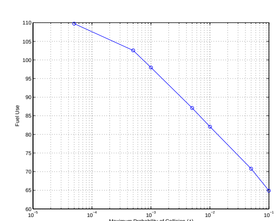

Slide 15

LB2 Work out how to link into the next slide - see next slide

Lars Blackmore, 6/10/2006

Conclusion

- Robust path planning problem can be solved as DLP
  - **Essentially same complexity** as DLP that does not take into account uncertainty (one lookup, matrix multiplication per constraint)
- The catch?
  - The resulting plan is **excessively conservative**
    - We guarantee *p(collision)* is less than δ, but in practice, *p(collision)* is much less than δ
    - Hence there exists a better solution that still satisfies chance constraint
    - If we try to constrain the probability of collision at **any** time step, we get very conservative plans
- Solution: Ongoing research
  1. Particle Control approach approximates distributions using samples
     - Approximate approach instead of conservative approach
  2. Ellipsoidal approximations have been used in the literature to solve analogous problems

Questions?

Backup

- Conservatism plot