Massachusetts Institute of Technology

Optimal, Robust Predictive Control using Particles

Lars Blackmore

November 27, 2006

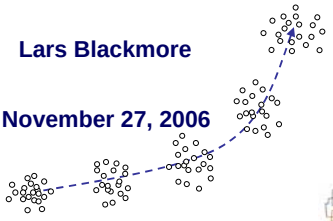

Model-based Embedded & Robotic Systems

CSAIL

Context

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

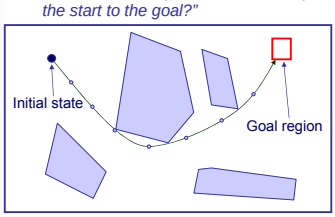

- Non-convex feasible region
- Non-holonomic dynamics
- Discrete-time

LB43

- Constrained optimization approaches have been successful *(Schouwenaars01,Maciejowski02,Dunbar02,Richards03)*

Optimal Controls are Not Robust to Uncertainty

- Uncertainty arises due to: LB5
  - Disturbances
  - Uncertain state estimation
  - Inaccurate modeling

  These are typically described **probabilistically**

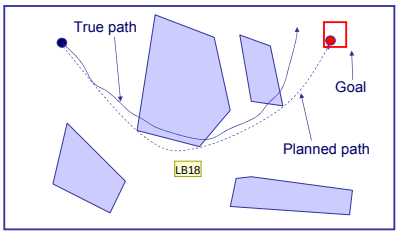

Why Probabilistic Uncertainty?

- Two principal ways to represent uncertainty:

LB44

1. Set-bounded uncertainty *(Zhou88,Gossner97,Bemporad99,Kerrigan01)*

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

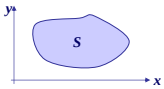

2. Probabilistic uncertainty *(Bertsekas78, Ma 2005)*

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

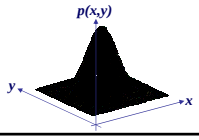

Why Probabilistic Uncertainty?

- 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 of Probabilistic Dynamic Systems

- Given probabilistic uncertainty, we want to plan **distribution** of future state in optimal, robust manner
- *Chance-constrained* formulation: Find optimal sequence of control inputs such that $p(failure) \leq \delta$

LB42

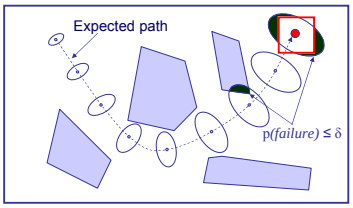

L1

Slide 2

LB43 Mention:

Example: UAV path planning using Mixed Integer Linear Programming
Lars Blackmore, 8/18/2006

Slide 3

LB5 Relate to UAV example
Lars Blackmore, 8/14/2006

LB18 stress optimal paths particularly bad
Lars Blackmore, 8/15/2006

Slide 4

LB44 Spend a little more time on this
Lars Blackmore, 8/18/2006

Slide 6

LB42 stress the trade of performance vs conservatism
Lars Blackmore, 8/18/2006

L1 mention the 3 challenges
Lars, 8/20/2006

Problem Statement

- Design a finite, optimal sequence of control inputs $\mathbf{u}_{0:T-1}$ such that system state trajectory leaves feasible region with probability at most δ
  - Cost function: $h(\mathbf{u}_{0:T-1}, \mathbf{x}_{1:T}) \rightarrow E[h(\mathbf{u}_{0:T-1}, \mathbf{x}_{1:T})]$
  - Feasible region: F
  - System dynamics: $\mathbf{x}_t = f_t(\mathbf{u}_{0:t-1}, \mathbf{x}_0, v_{0:t-1})$ LB37

Random variable with p.d.f. to be optimized

Random variables with **known** p.d.f.s (at least approximately)

Chance-constrained Control: Prior Work

- Prior work developed chance-constrained Model Predictive Control *(Li00, VanHessem04)*
  - Restricted to case of Gaussian uncertainty
  - Restricted to control within convex regions
- We extend this work to **arbitrary uncertainty distributions** and **non-convex regions**

➢ Chance-constrained particle control

Chance-constrained Particle Control: Intuition

- In estimation, Kalman Filters have been very successful for linear systems, Gaussian noise
  - State distribution given model and observations can be calculated analytically
- More recently, Particle Filters have been successful for nonlinear systems, with non-Gaussian noise
  - State distribution is approximated by a number of particles
  - Convergence of approximation to true distribution as number of particles tends to infinity
  - Number of particles used determined by available resources

LB45

- **Idea: Use particles for anytime robust control**
  - Control the distribution of particles to achieve a probabilistic goal

Particles: Probabilistic Properties

- Particles can approximate arbitrary distributions:
  - Draw N samples $x^{(i)}$ from a r.v. X with distribution $p(x)$
  - Distribution approximated with delta functions at samples:

$$p(x) \approx \frac{1}{N}\sum_{i=1}^{N} \delta_{x^{(i)}}(x)$$

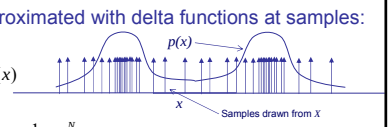

$$P(X \in S) = \int_S p(x)dx \approx \frac{1}{N}\int_S \sum_{i=1}^{N} \delta_{x^{(i)}}(x)dx = \text{fraction of particles in } S$$

- Convergence results:

$$\frac{1}{N}\sum_{i=1}^{N} f(x^{(i)}) \xrightarrow{\text{almost surely}} E[f(X)] \quad \text{as} \quad N \rightarrow \infty$$

$$\text{fraction of particles in set } S \xrightarrow{\text{almost surely}} p(X \in S) \quad \text{as} \quad N \rightarrow \infty$$

LB20

Technical Approach

- Question: How can particles be used to solve chance-constrained probabilistic control problem?
- Chance-constrained particle control:
  1. Use particles to sample random variables (noise, initial position, disturbances)
  2. Calculate future state trajectory for each particle **leaving explicit dependence** on control inputs $\mathbf{u}_{0:T-1}$
  3. Express **probabilistic** optimization problem **approximately** in terms of particles
  4. Solve approximate **deterministic** optimization problem for $\mathbf{u}_{0:T-1}$

➢ Approximate optimization goal **tends to true goal** as number of particles tends to infinity

Technical Approach

1. Use particles to sample random variables

$$\mathbf{x}_0^{(i)} \sim p(\mathbf{x}_0) \quad v_t^{(i)} \sim p(v_t) \qquad i = 1\text{K } N \quad t = 0\text{K } T$$

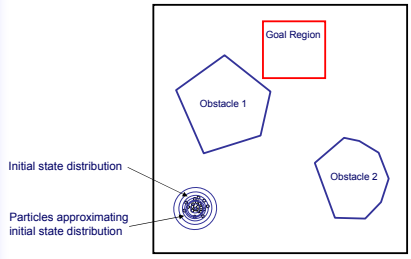

Slide 7

LB37 "state depends on x_0 and v which are...

because they are random variables, x_t is also."
Lars Blackmore, 8/16/2006

Slide 9

LB45 note anytime
Lars Blackmore, 8/18/2006

Slide 10

LB20 these are used in filtering --> I will use for control
Lars Blackmore, 8/15/2006

Technical Approach

2. Calculate future state trajectory for each particle **leaving explicit dependence** on control inputs $\mathbf{u}_{0:T-1}$

$$\mathbf{x}_t^{(i)} = f_t(\mathbf{u}_{0:t-1}, \mathbf{x}_0^{(i)}, \nu_{0:t-1}^{(i)}) \qquad \mathbf{x}_{1:T}^{(i)} = \begin{bmatrix} \mathbf{x}_1^{(i)} \\ \mathrm{M} \\ \mathbf{x}_T^{(i)} \end{bmatrix}$$

Particle 1 for $\mathbf{u} = \mathbf{u}_B$

Particle 1 for $\mathbf{u} = \mathbf{u}_A$

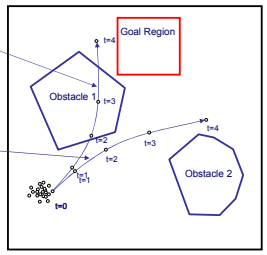

LB21

Technical Approach

2. Calculate future state trajectory for each particle **leaving explicit dependence** on control inputs $\mathbf{u}_{0:T-1}$

$$\mathbf{x}_t^{(i)} = f_t(\mathbf{u}_{0:t-1}, \mathbf{x}_0^{(i)}, \nu_{0:t-1}^{(i)}) \qquad \mathbf{x}_{1:T}^{(i)} = \begin{bmatrix} \mathbf{x}_1^{(i)} \\ \mathrm{M} \\ \mathbf{x}_T^{(i)} \end{bmatrix}$$

Particles $1 \ldots N$ for $\mathbf{u} = \mathbf{u}_B$

Particles $1 \ldots N$ for $\mathbf{u} = \mathbf{u}_A$

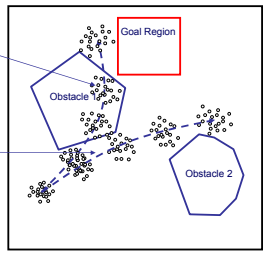

LB38

LB24

Technical Approach

3. Express **probabilistic** optimization problem **approximately** in terms of particles

LB46

Fraction of particles failing approximates probability of failure

Sample mean approximates state mean

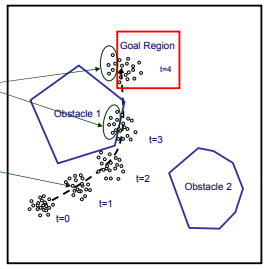

Sample mean of cost function approximates true expectation:

$$E[h(\mathbf{u}_{0:T-1}, \mathbf{x}_{1:T})] \approx \frac{1}{N}\sum_{i=1}^{N} h(\mathbf{u}_{0:T-1}, \mathbf{x}_{1:T}^{(i)})$$

Technical Approach

4. Solve approximate **deterministic** optimization problem for $\mathbf{u}_{0:T-1}$

$\delta = 0.1$

10% of particles fail in optimal solution

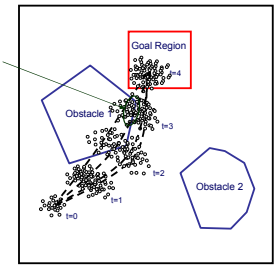

Convergence

– As N→∞, approximation becomes exact

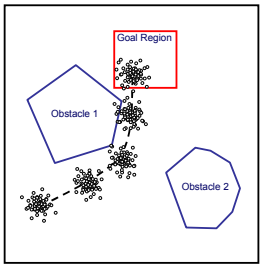

Convergence

– As N→∞, approximation becomes exact

10% probability of failure

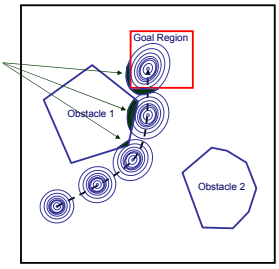

Slide 13

LB21 This is easy to do because all uncertainty has been removed
Lars Blackmore, 8/15/2006

Slide 14

LB24 highlight every particle has same control input
Lars Blackmore, 8/15/2006

LB38 stress a particle is a _trajectory_
Lars Blackmore, 8/16/2006

Slide 15

LB46 MAX failure rate
Lars Blackmore, 8/18/2006

Particle Control for Linear Systems

- For nonlinear systems, resulting optimization too complex for real-time operation
- In the case of:
  1. Linear systems
  2. Polygonal feasible regions
  3. Piecewise linear cost function

  global optimum can be found extremely efficiently using Mixed Integer Linear Programming (MILP)
- Important problems using linear systems, cost function:
  - Aircraft control about trim state
  - UAV path planning
  - Satellite control

LB33

Particle Control for Linear Systems

- Future state for each particle:

$$x_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + v_t \qquad \mathbf{x}_t^{(i)} = A^t\mathbf{x}_0^{(i)} + \sum_{j=0}^{t-1} A^{t-j-1}\left(B\mathbf{u}_j + v_j^{(i)}\right)$$

- Approximate expected state:

$$E[\mathbf{x}_T] \approx \frac{1}{N}\sum_{i=1}^{N}\mathbf{x}_T^{(i)}$$

- Approximate expected cost:

$$E\left[h(\mathbf{u}_{0:T-1}, \mathbf{x}_{1:T})\right] \approx \frac{1}{N}\sum_{i=1}^{N} h(\mathbf{u}_{0:T-1}, \mathbf{x}_{1:T}^{(i)})$$

- All of these are **linear** functions of control inputs

Particle Control for Linear Systems

- Approximate chance constraints:
  - Express feasible region in terms of constraints

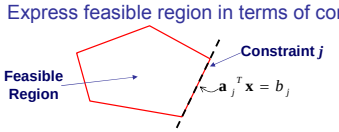

LB34

  - Introduce binary variables indicating particle success

$$\forall j \quad \mathbf{a}_j^T\mathbf{x}^{(i)} - b_j \le Cz_i \quad z_i \in \{0,1\} \qquad C \text{ large, positive}$$

z_i=0 **implies all constraints satisfied by particle** i

  - Constrain sum of binary variables

$$\frac{1}{N}\sum_i z_i \le \delta \quad \Rightarrow \quad \textit{at most a fraction } \delta \textit{ of particles fail}$$

Particle Control for Linear Systems

- Chance-constrained problem is a MILP
- Decision variables:
  - Control inputs (continuous) $\mathbf{u}_{0:T-1}$
  - State trajectory for each particle (continuous) $\mathbf{x}_{1:T}^{(i)}$
  - Success indicator for each particle (binary) z_i
- Can be solved to global optimality
- Extremely efficient MILP solvers available

Simulation Results

- Task A:
  - Control of Boeing 747 in heavy turbulence
  - Aircraft must remain within defined flight envelope
  - Longitudinal dynamics linearized about trim state
  - Assume inner-loop altitude hold controller
  - Minimize fuel use (elevator deflection)
- Uncertainty sources:
  - Disturbances due to heavy turbulence (MIL-F-8785C)
  - Sensor noise gives rise to attitude uncertainty

➢ **Highly non-Gaussian distributions**

LB47

Simulation Results

LB48

Num particles = 100, Planning horizon = 20 steps, dt = 2s, δ = 0.1

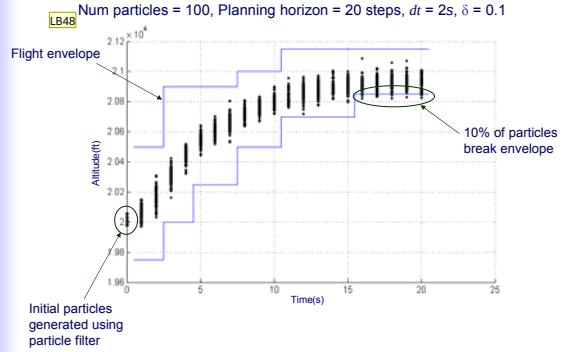

Slide 19

LB33 Mention min time, min fuel
Lars Blackmore, 8/16/2006

Slide 21

LB34 say more complex for nonconvex but won't go into details
Lars Blackmore, 8/16/2006

Slide 23

LB47 go faster over this
Lars Blackmore, 8/18/2006

Slide 24

LB48 Mention closed loop with particle filter
Lars Blackmore, 8/15/2006

Simulation Results

- Convergence of failure probability:

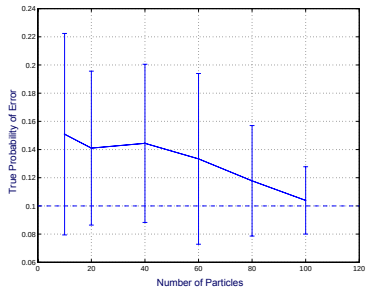

25

Simulation Results

- Conservatism vs. performance

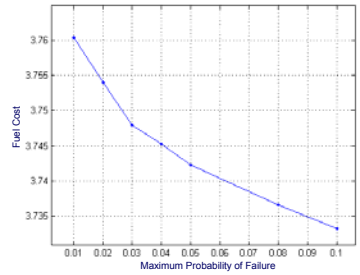

26

Simulation Results

- Solution time vs. number of particles

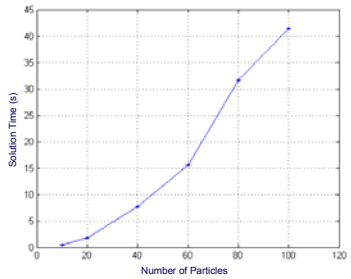

27

Simulation Results

- Task B:
  - 2-D Control of Unmanned Air Vehicle
  - Environment contains obstacles to be avoided
  - Linear dynamics model *(Schouwenaars01)*
  - Minimize fuel use
- Uncertainty sources:
  - Wind disturbances
  - Initial state uncertainty

➢ **Highly non-Gaussian distributions**

28

Simulation Results

- Maximum probability of failure: 0.04

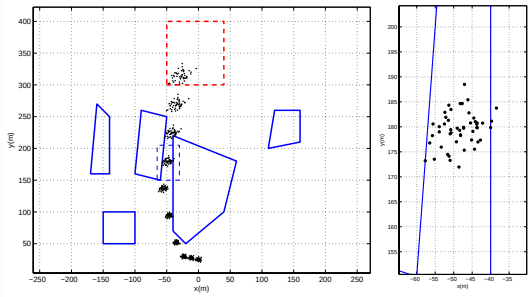

Num particles = 50, planning horizon = 10 steps, dt = 1s, δ = 0.04

29

Simulation Results

- Maximum probability of failure: 0.02

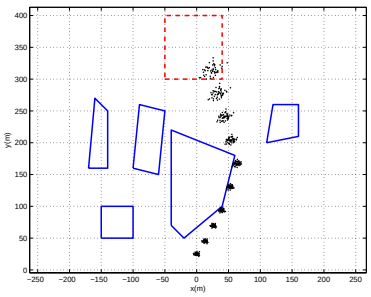

LB29

Num particles = 50, planning horizon = 10 steps, dt = 1s, δ = 0.02

30

Slide 30

LB29 maybe show a box here as well
Lars Blackmore, 8/16/2006

Future Work

- Applications
  - Production, operations LB23
  - Landing site selection
- Efficient solutions to more general systems
  - Switching (hybrid) systems → failure tolerant control
  - Nonlinear systems
- Receding horizon results
  - Robust probabilistic feasibility?
- Bias reduction/elimination

Conclusion

- Novel any-time approach to robust probabilistic control with arbitrary probability distributions
  - Very general formulation, challenge is tractable optimization
  - For linear systems, global optimum can be found efficiently

Questions?

Backup

Related Work Using Particles

- Particles have previously been used in decision making, variously referred to as:
  - Particles *(Doucet01,Greenfield03)*
  - Simulations *(Singh06)*
  - Scenarios *(Ng00, Yu03, Tarim06)*
- *(Ng00)* converts a stochastic MDP into deterministic MDP by representing all randomness in initial state
- *(Greenfield03)* proposed finite horizon control with expected cost approximated using particles
- *(Doucet01, Singh06)* use samples to approximate cost function value and gradient in local optimizer
- Key contributions of chance-constrained particle control:
  - Use particles to approximate the **probability** of failure
    - Optimization with constraints on **probabilities** is a novel and powerful tool (e.g. chance-constrained planning)
  - Efficient solution using MILP

Complexity

- Problem is worst-case exponential in:
  - Length of planning horizon
  - System order
  - Number of particles
  - Number of constraints
- MILP solution means that worst-case complexity is almost never realized
- With MILP optimization, typical solution time difficult to characterize
- Empirical results show for convex F, relatively large problems can be solved in less than a minute
- For non-convex F, even with heuristic pruning approach, medium-sized problems take many minutes
- MILP can use a large amount of time proving that solution is global optimum
  - Good feasible solutions can typically be found much more quickly

Slide 31

LB23 Lars Blackmore 8/15/2006

Highlight generality, appealing chance-constrained approach, only need to find ways to make optimisation tractable

Lars Blackmore, 8/15/2006

Bias

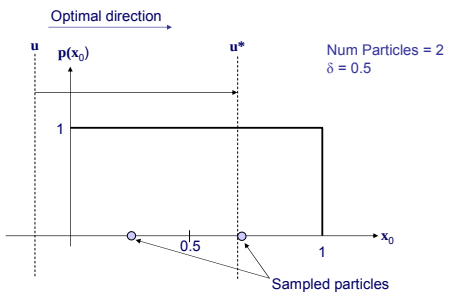

Expected position of furthest right particle = 0.7