
Efficient, yet safe, robot motion planning in dynamic, uncertain environments

RSS 2011 Workshop: Guaranteeing Motion Safety
for Robots

June 27, 2011


Noel du Toit and **Joel Burdick**
California Institute of Technology

Dynamic, Cluttered, Uncertain Env's



Safety vs Conservatism

- Probabilistic uncertainty => conservative



- Conservatism: incorporate anticipated measurements
 - Previous works: static environments [Prentice 09], [Censi 08], [Yan 05]
 - Enabling capability: complicated agent behaviors, more clutter, agent information gathered
- Safety: chance constraint conditioning

Formulation: Stochastic RHC (SRHC)

- Stochastic system: $x_i = f(x_{i-1}, u_{i-1}, \omega_{i-1})$
 $y_i = h(x_i, \nu_i)$

- Expected cost, chance constraints

- Belief space

- Belief state: $\zeta_k = p(x_k | \eta_k)$

- Transition function:

$$\zeta_{k+1} = f_\zeta(\zeta_k, u_k, y_{k+1})$$

- Disturbance: y_{k+1}


$$\min_{u_{k:M}} \left[E_{y_{k+1:M}} \left[l_M(x_M) + \sum_{i=k}^{M-1} l_i(x_i, u_i) \mid I \right] \right]$$

$$s.t. \quad \zeta_{k+1} = f_\zeta(\zeta_k, u_k, y_{k+1})$$

$$g_i(u_i, \zeta_i) \leq 0$$

$$P_u((x_i) \in \mathbb{X}_{free}) \leq \alpha$$

$$P(c_x(x_i) \leq 0) \geq \alpha$$



SRHC Approximations

- Open-loop approach [Bertsekas 05], [Yan & Bitmead 05], [Blackmore 07]


- Ignore future information: restricted information set

$$\bar{\eta}_i = (y_1, \dots, y_k, u_0, \dots, u_{i-1})$$

- Belief state:

$$\bar{\zeta}_i = p(x_i | y_{1:k}, u_{0:i-1})$$

- Deterministic in terms of the belief state:



$$\begin{aligned} \min_{u_k: u_M \in U} & E \bar{J}_{y_k}(\zeta_k^M, u_k^M) \left[\bar{l}_M(\zeta_M^M) + \bar{l}_i(\bar{\zeta}_i^M, \bar{u}_i^M) \right] \\ s.t. & \zeta_{k+1}^M = f(\zeta_k^M, u_k^M, y_{k+1}^M) \\ & c_u(u_i) \leq 0 \\ & P(x_i \in \mathbb{X}_{free}) \geq \alpha \end{aligned}$$

Approximation: Most Likely Measurement

- Effect of measurements

- Update covariance

$$\Sigma_{k|k} = f(\Sigma_{k-1|k-1}, W, V)$$

- Update mean


$$\hat{x}_{k|k} = f(y_k, \hat{x}_{k-1|k-1}, \Sigma_{k-1|k-1}, W, V)$$

- Most likely measurement:

$$\tilde{y}_j = E[y_j | \bar{\eta}_{j-1}]$$

- Same computational benefits of OLRHC approach

- Approximation: relies on RHC formulation to correct for assumed information to reduce conservatism in solutions



Chance Constraints

- Constrain uncertain state

$$c(x_i) \leq 0 \longrightarrow P(c(x_i) \leq 0) \geq \alpha$$

- Linear constraint, normally distributed state

$$P(a^T x_i^R \leq b) \geq \alpha \iff a^T \hat{x} \leq b - r(\alpha) \times \sqrt{a^T \Sigma a}$$

- Probabilistic Collision Avoidance $P(C) \leq \delta$


$$C(x_R, x_A) : \mathbb{X}_R(x_R) \cap \mathbb{X}_A(x_A) \neq \{\emptyset\}$$

- Robot and obstacle uncertainty

- Joint distribution and indicator function

$$P(C) = \int_{x_R} \int_{x_A} I_C(x_A, x_R) p(x_R, x_A) dx_R dx_A$$

$$I_C(x_A, x_R) = \begin{cases} 1 & \text{if } \mathbb{X}_R(x_R) \cap \mathbb{X}_A(x_A) \neq \{\emptyset\} \\ 0 & \text{otherwise.} \end{cases}$$



Collision Constraints: Evaluation

■ Monte-Carlo Simulation

$$P(C) \approx \frac{1}{N} \sum_{i=1}^N I_C(\tilde{x}^i)$$

- Small level of confidence: $\delta = 0.01$ requires ~ 10000 samples
- $\sim 5\text{ms}$ per evaluation

■ Approximate: small disk objects


$$P(C) \approx V_D \int_{x_R} p(x_A = x_R | x_R) p(x_R) dx_R$$

- Independent, Gaussian distributions:

$$(\hat{x}_R - \hat{x}_A)^T (\Sigma_R + \Sigma_A)^{-1} (\hat{x}_R - \hat{x}_A) \geq \kappa(\Sigma_R, \Sigma_A, \delta, V_D)$$

- Dependent, Gaussian distributions

$$(\hat{x}_R - \hat{x}_A)^T (\Sigma_R + \Sigma_A - \Sigma_M - \Sigma_M^T)^{-1} (\hat{x}_R - \hat{x}_A) \geq \kappa(\Sigma_R, \Sigma_A, \Sigma_M, \delta, V_S)$$




Safety: Reaction Horizon

- Quantify time (# of stages) to react to changes in environment
 - Robot dynamics
 - Environment
- Reaction horizon, r : react to modeled disturbances


- Chance constraint conditioning:

$$P(c(x_i) \leq 0 | \bar{\eta}_{i-r}) \geq \alpha$$

- Use r -stage open-loop predicted distribution
- Anticipated information: leverage PCL reduction in conservatism





Safety vs Conservatism




Uncertainty grows over reaction horizon

Next stage: new information + re-solve problem (RHC)




PCL: leverage new information

OL: uncertainty growth results in conservative solutions



1-D example:

- Car following:
 - Collision constraint (maintain some separation distance)



- Velocity controlled random-walk model


$$x_i = \begin{bmatrix} x_i^{(1)} \\ x_i^{(2)} \end{bmatrix} = \begin{bmatrix} \text{position} \\ \text{velocity} \end{bmatrix}$$

$$x_i = Ax_{i-1} + Bu_{i-1} + F\omega_{i-1} \quad A = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix} \quad B = F = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$


- Reaction horizon: $r=2$ (to influence position)

1-D Example (cont'd)



Reaction horizon = 1



Plot separation distance



Reaction horizon = 2



PCLRHC



Summary

- Practical systems: trade off conservatism and safety
- PCLRHC
 - Reduce conservatism through anticipated information
 - RHC: resolve problem to incorporate actual measurements
- Chance constraint conditioning
 - Allow for modeled disturbances
 - Can still leverage anticipated information

Thank you

ndutoit@caltech.edu

jwb@robotics.caltech.edu

Questions?

Publications:

Du Toit, N.E. and Burdick, J.W., “Robot Motion Planning in Dynamic, Cluttered, Uncertain Environments”, *accepted to IEEE Transactions on Robotics*

Du Toit, N.E. and Burdick, J.W., “Probabilistic Collision Checking with Chance Constraints”, *accepted to IEEE Transaction on Robotics*

Du Toit, N.E., “Robot Motion Planning in Dynamic, Cluttered, Uncertain Environments: the Partially Closed-Loop Receding Horizon Control Approach”, *Ph.D. thesis, Caltech, 2010*

Conferences:

Workshop on Motion Planning: From Theory to Practice (RSS) 2010

Du Toit, N.E. and Burdick, J.W., “Robotic Motion Planning in Dynamic, Cluttered, Uncertain Environments”, *ICRA 2010*

Problem Definition

Robot:

$$x_i^R = f_R(x_{i-1}^R, u_{i-1}^R, \omega_{i-1}^R)$$

$$\omega_i^R \sim p(\omega_i^R | x_i^R, u_i^R)$$

$$y_i^R = h_R(x_i^R, \nu_i^R)$$

$$\nu_i^R \sim p(\nu_i^R | x_i^R)$$

Agent:

$$x_i^A = f_A(x_{i-1}^A, x_{i-1}^R, u_{i-1}^A, \omega_{i-1}^A)$$

$$\omega_i^A \sim p(\omega_i^A | x_i^R, x_i^A, u_i^A)$$


$$y_i^A = h_A(x_i^R, x_i^A, \nu_i^A)$$

$$\nu_i^A \sim p(\nu_i^A | x_i^A, x_i^R)$$

Constraints: $c(x_i^R, x_i^A, u_i^R) \leq 0$

Objective function:

$$L(x_{0:N}^R, x_{0:N}^A, u_{0:N-1}^R, \omega_{0:N-1}^R) = l_N(x_N^R, x_N^A) + \sum_{i=0}^{N-1} l_i(x_i^R, x_i^A, u_i^R, \omega_i^R)$$



Related Work

■ Stochastic systems: Probabilistic vs. non-deterministic

	Deterministic		Probabilistic	
	Static	Dynamic	Static	Dynamic
Control	OC [Friedland 05] RHC [Mayne 00]	OC with augmented states	OC with separation [Friedland 05] Stochastic RHC (later)	PCLRHC Probabilistic velocity obstacles [Fulgenzi et al. 07]
Robotics	Graph search, roadmap methods, etc. [LaValle 06], [Choset et al. 07]	Dynamic window [Fox et al. 97] Velocity obstacles [Fiorini & Shiller 98]	Graph search in extended state space [LaValle 06], [Censi 08]	
AI	MDPs [LaValle 06]	Extended state space (time x pose) [LaValle 06]	DP [Bertsekas 05] MDPs, POMDPs [Thrun et al. 05]	

AI: Artificial Intelligence

OC: Optimal Control

RHC: Receding Horizon Control

DP: Dynamic Programming

MDP: Markov Decision Process

POMDP: Partially Observable MDP

Partially Closed-loop RHC

- Most likely measurement:

$$\tilde{y}_j = E[y_j | \bar{\eta}_{j-1}]$$

- Restricted information:

$$\bar{\eta}_i = (y_1, \dots, y_k, \tilde{y}_{k+1}, \dots, \tilde{y}_i, u, \dots, u_{i-1})$$

- Resulting belief state:

$$\bar{\zeta}_i = p(x_i | \bar{\eta}_i) = p(x_i | u_{0:i-1}, y_{1:k}, \tilde{y}_{k+1:i})$$

- Deterministic in belief state:

$$\begin{aligned} \min_{u_k: u_M \in U} \quad & E_{y_k} \left[\bar{\zeta}_k^T u_k \left[\bar{\zeta}_M^T \left(\sum_{i=k+1}^{M-1} \bar{\zeta}_i \right) + \bar{l}_i \left(\sum_{i=k}^{M-1} \bar{\zeta}_i, \bar{u}_i \right) \right] \right] \\ \text{s.t.} \quad & \bar{\zeta}_{k+1} = f(\bar{\zeta}_k, u_k, y_{k+1}) \\ & c_u(u_i) \leq 0 \\ & P(x_i \in \mathbb{X}_{free}) \leq \alpha \end{aligned}$$

Simulation Setup

- Robot: Linear model, linear measurements

$$\begin{aligned}x_i^R &= Ax_{i-1}^R + Bu_{i-1}^R + F\omega_{i-1}^R \\y_i^R &= Cx_i^R + H\nu_i^R\end{aligned}$$

$$\begin{aligned}\omega_{i-1}^R &\sim N(0, W) \\ \nu_i^R &\sim N(0, V)\end{aligned}$$

- Velocity constraints:


$$|v_{x,i}^R| \leq 2, |v_{y,i}^R| \leq 2$$

- Control constraints:

$$|u_{x,i}^R| \leq 1, |u_{y,i}^R| \leq 1$$

- Collision constraints:

$$P(C) \leq 0.01$$



- Agent: Linear constant velocity model, linear measurements

$$\begin{aligned}x_i^A &= A_A x_{i-1}^A + B_A u_{i-1}^A + F_A \omega_{i-1}^A \\y_i^A &= C_A x_i^A + H_A \nu_i^A\end{aligned}$$

$$\begin{aligned}\omega_{i-1}^A &\sim N(0, W_A) \\ \nu_i^A &\sim N(0, V_A)\end{aligned}$$

PCLRHC Approximation

- Information gain: relative entropy

$$D_{KL} = \frac{1}{2} \left(\ln \left(\frac{|\Sigma_C|}{|\Sigma_{BL}|} \right) + Tr(\Sigma_C^{-1} \Sigma_{BL}) + (\mu_C - \mu_{BL})^T \Sigma_C^{-1} (\mu_C - \mu_{BL}) - n_x \right)$$

- Baseline: $p_{BL}(x_i|y_{1:k}) = N(\hat{x}_{i|k}, \Sigma_{i|k})$

- PCL distribution: $p_C(x_i|y_{1:k}, \tilde{y}_{k+1:i}) = N(\hat{x}_{i|k}, \Sigma_{i|i})$

$$D_{pcl} = \frac{1}{2} \left(\ln \left(\frac{|\Sigma_{i|i}|}{|\Sigma_{i|k}|} \right) + Tr(\Sigma_{i|i}^{-1} \Sigma_{i|k}) - n_x \right)$$

- Executed distribution: $p_C(x_i|y_{1:k}, y_{k+1:i}) = N(\hat{x}_{i|i}, \Sigma_{i|i})$

$$D_{exec} = \frac{1}{2} \left(\ln \left(\frac{|\Sigma_{i|i}|}{|\Sigma_{i|k}|} \right) + Tr(\Sigma_{i|i}^{-1} \Sigma_{i|k}) + (\hat{x}_{i|i} - \hat{x}_{i|k})^T \Sigma_{i|i}^{-1} (\hat{x}_{i|i} - \hat{x}_{i|k}) - n_x \right)$$

$$D_{pcl} \leq D_{exec}$$