Covariance Steering as a Tool for Planning and Control in the Presence of Uncertainty

Panagiotis Tsiotras

School of Aerospace Engineering
School of Electrical and Computer Engineering
School of Interactive Computing
Institute for Robotics and Intelligent Machines
Georgia Institute of Technology

Control of Distributions Workshop
May 24, 2021


x(T)
x(0)

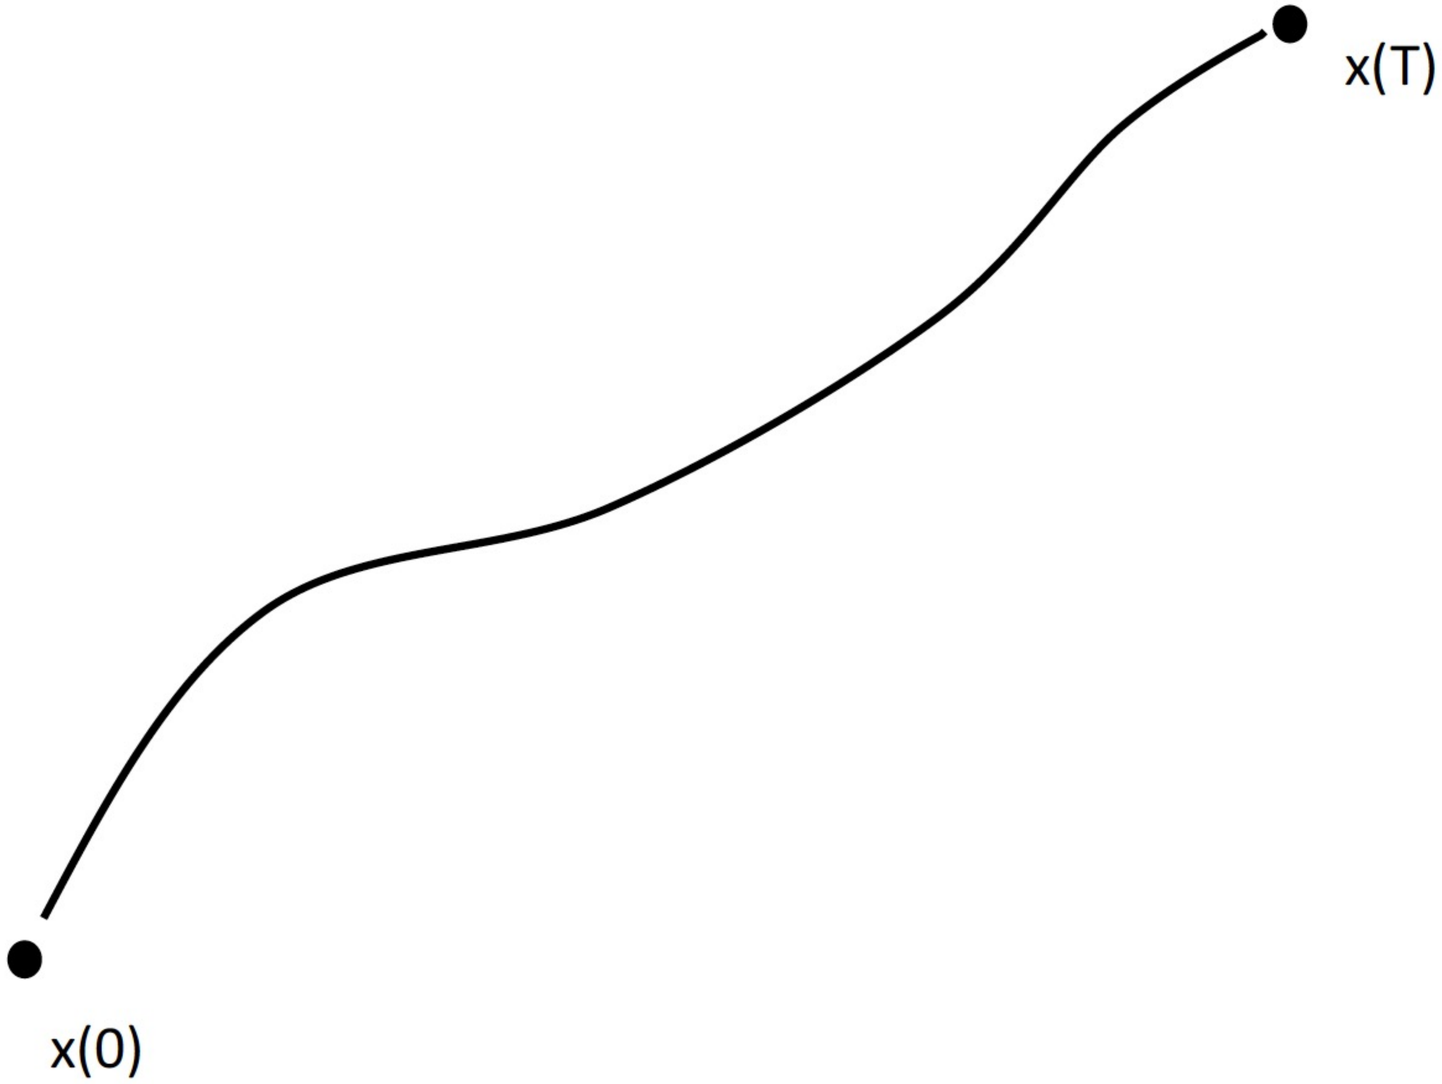
x(T)
x(0)

A Fundamental Question

- Let the LTV system

$$\dot{x} = A(t)x + B(t)u$$

with boundary conditions $x(0) = x_0$ and $x(T) = x_T$.

- Then the control

$$u(t) = B^\top(t)\Theta^\top(0,t)\left(\Theta(T,\tau)B(\tau)B^\top(\tau)\Theta^\top(0,\tau)\,\mathrm{d}\tau\right)^{-1}\left(x(T) - \Theta(T,0)x(0)\right)$$

minimizes

$$\int_0^T u^\top(\tau)u(\tau)\,\mathrm{d}\tau$$

and satisfies the given boundary conditions.

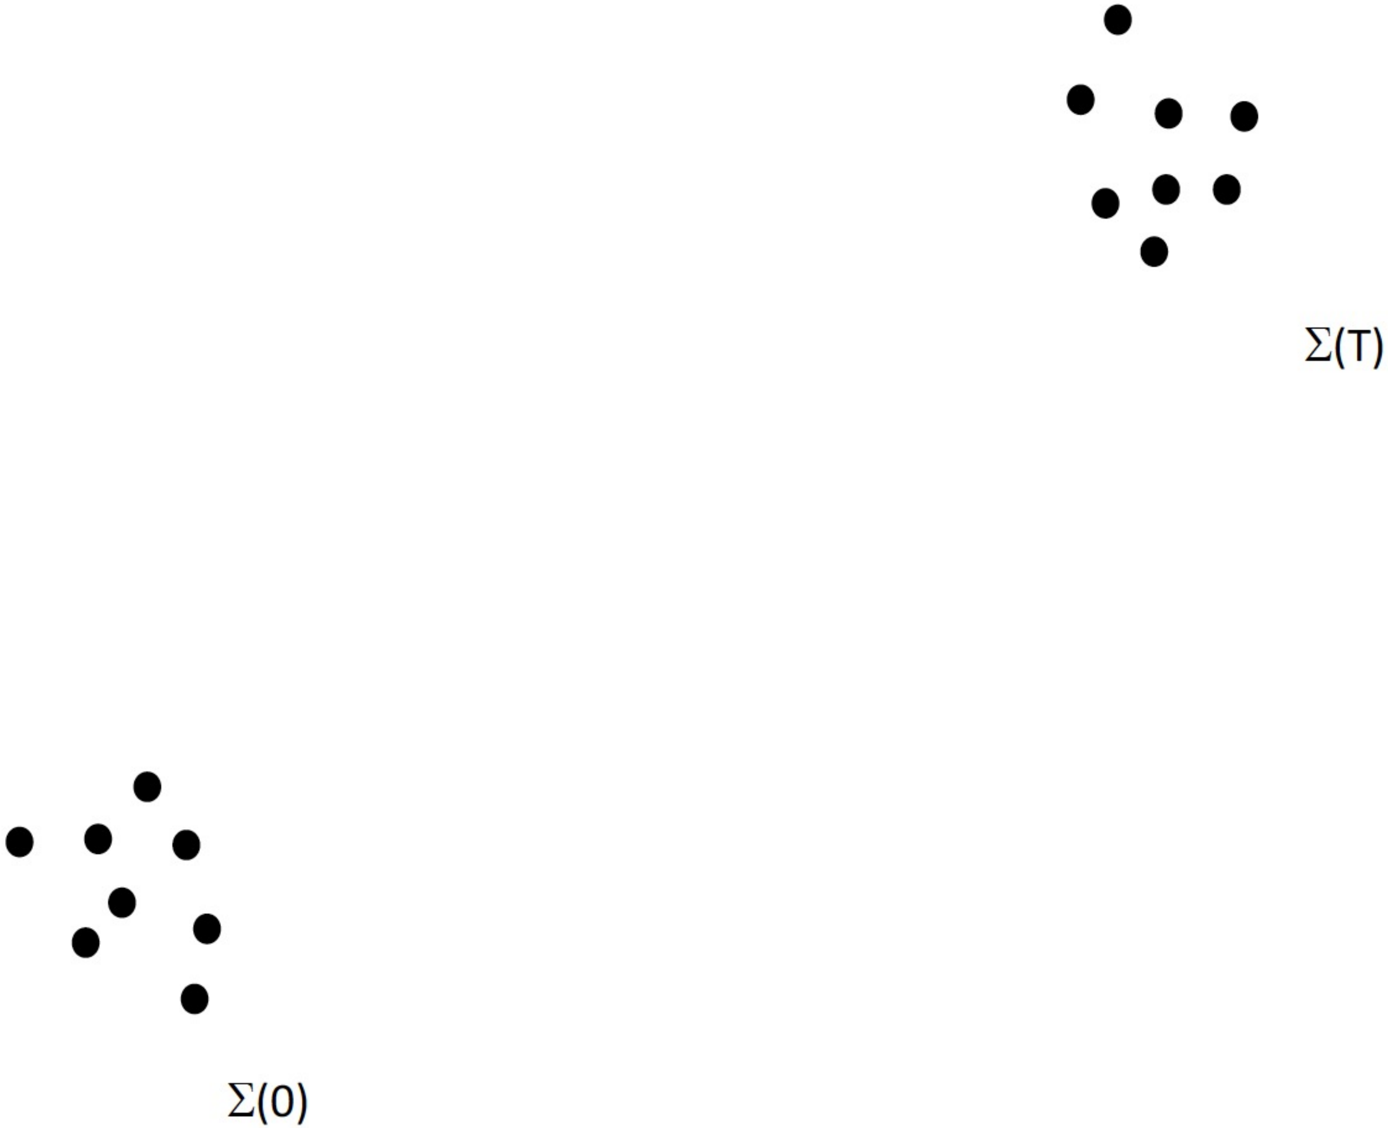
Σ(T)
Σ(0)

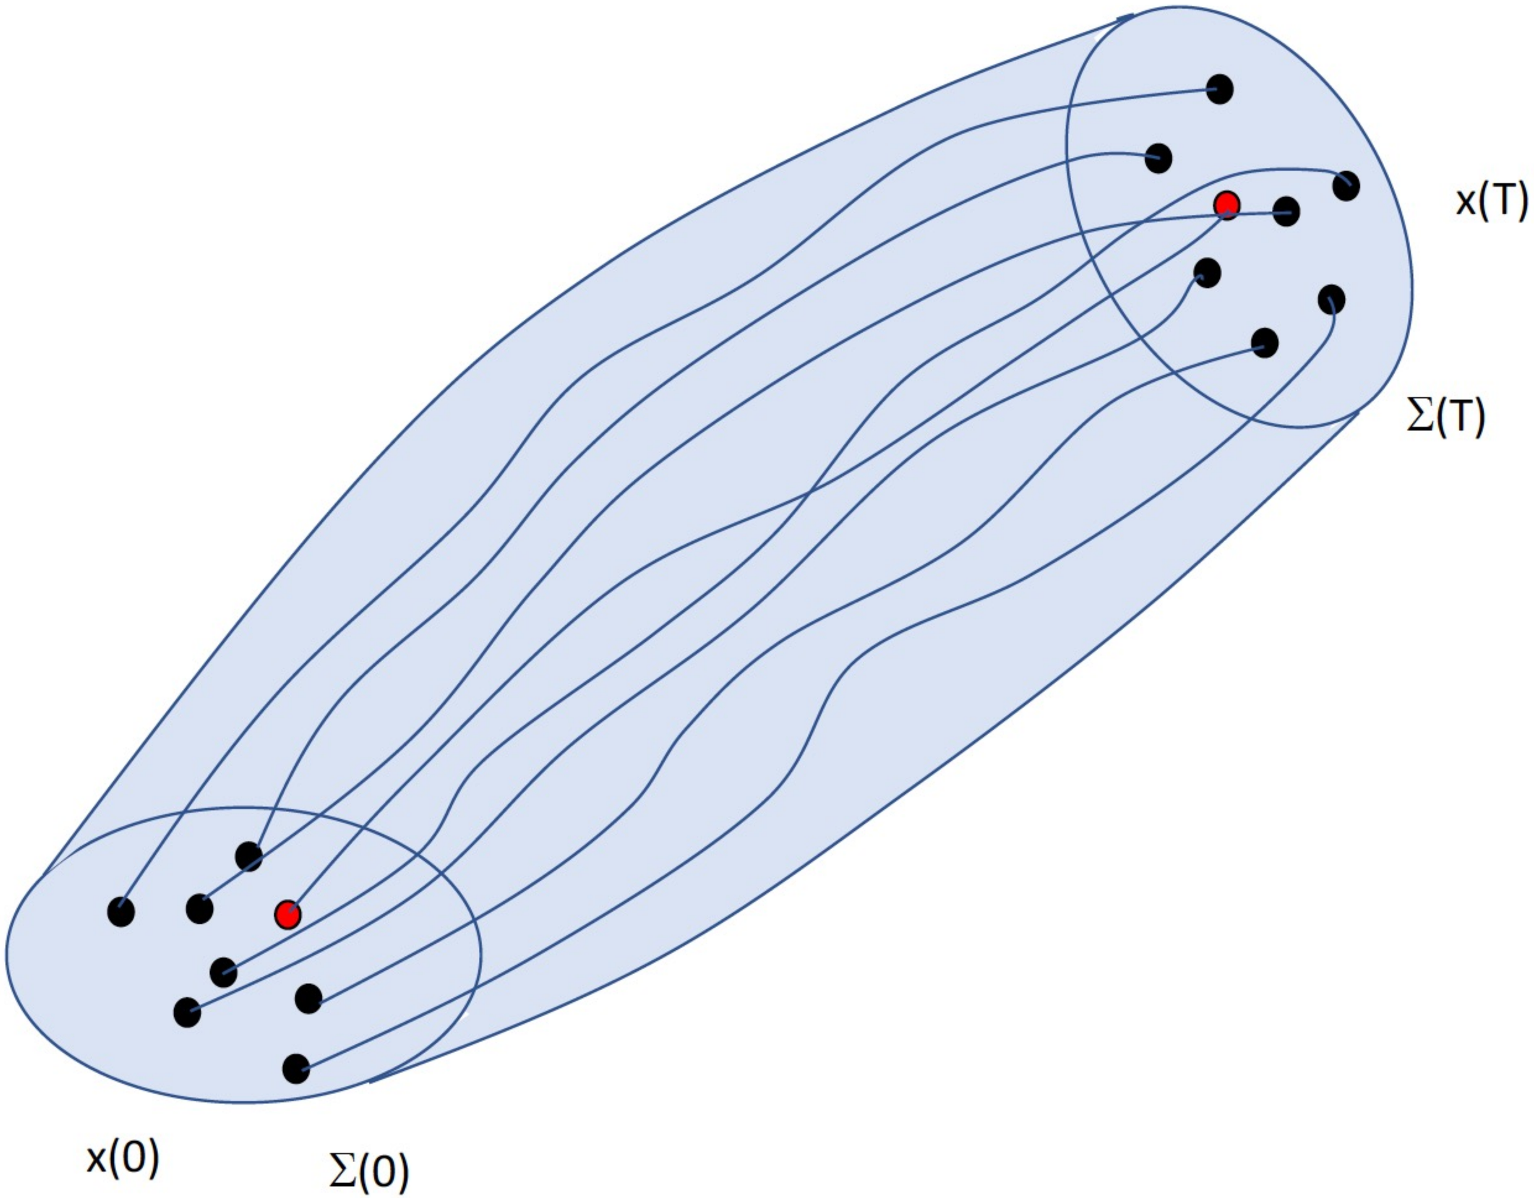
x(T)
Σ(T)
x(0)
Σ(0)

A Controllability Result

Theorem (Brockett, 2012)

Consider the system

$$\dot{x} = Ax + Bu + Dw$$

with the ontrol law

$$u(t) = K(t)x(t) + v(t)$$

Let (A, B) be controllable and let Σ denote the (co)variance, which satisifes

$$\dot{\Sigma} = (A + BK(t))\Sigma + \Sigma(A + BK(t))^\top + DD^\top$$

With K as a control, any $\Sigma_1 \succ 0$ can be reached from any $\Sigma(0) \succ 0$.

Problem Formulation

Consider discrete-time stochastic linear system

$$x_{k+1} = A_k x_k + B_k u_k + D_k w_k$$

- We wish the initial and final states to be distributed according to

$$x_0 \sim \mathcal{N}(\mu_0, \Sigma_0), \qquad x_N \sim \mathcal{N}(\mu_N, \Sigma_N)$$

where $\mu_0, \Sigma_0, \mu_N, \Sigma_N$ given, while minimizing the cost function

$$J(x,u) = \mathbb{E}\left[\sum_{k=0}^{N-1} x_k^\top Q_k x_k + u_k^\top R_k u_k\right] \cancel{+ x_N^\top Q_N x_N}$$

where $Q_k \succeq 0$ and $R_k \succ 0$ for all $k = 0, 1, \ldots, N-1$.

- Assume that $\Sigma_0 \succeq 0$ and $\Sigma_N \succ 0$, The system at time step $k = N$ can be written as

$$x_N = E_N X = \mathcal{A} x_0 + \mathcal{B} U + \mathcal{D} W$$

where

$$X = \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_N \end{bmatrix}, \qquad U = \begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_{N-1} \end{bmatrix}, \qquad W = \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_{N-1} \end{bmatrix}$$

Covariance Controller

- Let the control sequence

$$u_k = v_k + K_k y_k$$

where y_k is given by

$$y_{k+1} = A_k y_k + D_k w_k$$
$$y_0 = x_0 - \mu_0$$

and let the control law

$$U = V + KY$$

Theorem (Okamoto & PT, 2018)

The cost function takes the form

$$J(V,K) = \mathrm{tr}\left(((I+\mathcal{B}K)^\top \bar{Q}(I+\mathcal{B}K) + K^\top \bar{R}K)(\mathcal{A}\Sigma_0\mathcal{A}^\top + \mathcal{D}\mathcal{D}^\top)\right) + (\mathcal{A}\mu_0 + \mathcal{B}V)^\top \bar{Q}(\mathcal{A}\mu_0 + \mathcal{B}V) + V^\top \bar{R}V$$

In addition, the terminal state constraints can be written as

$$\mu_N = E_N\left(\mathcal{A}\mu_0 + \mathcal{B}V\right),$$
$$\Sigma_N = E_N(I+\mathcal{B}K)(\mathcal{A}\Sigma_0\mathcal{A}^\top + \mathcal{D}\mathcal{D}^\top)(I+\mathcal{B}K)^\top E_N^\top$$

Note that V steers the mean and K steers the covariance, respectively.
Letting

$$\Sigma_N \succeq E_N(I+\mathcal{B}K)(\mathcal{A}\Sigma_0\mathcal{A}^\top + \mathcal{D}\mathcal{D}^\top)(I+\mathcal{B}K)^\top E_N^\top$$

yields a **convex problem**.

State Constraints

- Can handle convex **chance constraints** of the form

$$\mathtt{Pr}(x_k \notin \chi) \leq P_{\text{fail}}, \qquad k = 0, \dots, N-1$$

where

$$\chi = \bigcap_{j=1}^{M} \{x : \alpha_j^\top x \leq \beta_j\}$$

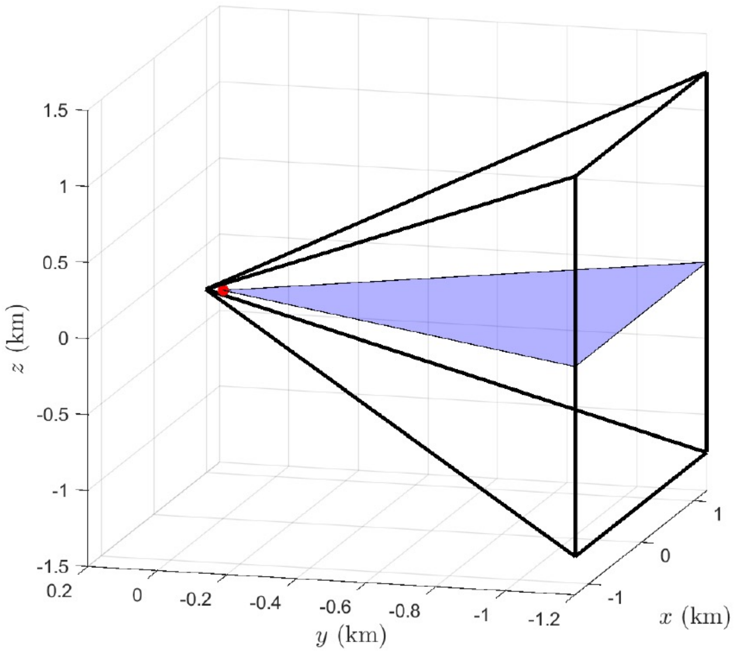

the chance constraint can be formulated as

$$\alpha_j^\top(\mathcal{A}\mu_0 + \mathcal{B}V) + \|(\mathcal{A}\Sigma_0\mathcal{A}^\top + \mathcal{D}\mathcal{D}^\top)^{1/2}(I + \mathcal{B}K)^\top \alpha_j\| \Phi^{-1}(1 - p_{j,\text{fail}}) - \beta_j \leq 0$$

Second order cone (convex) constraint in K and V.

Example

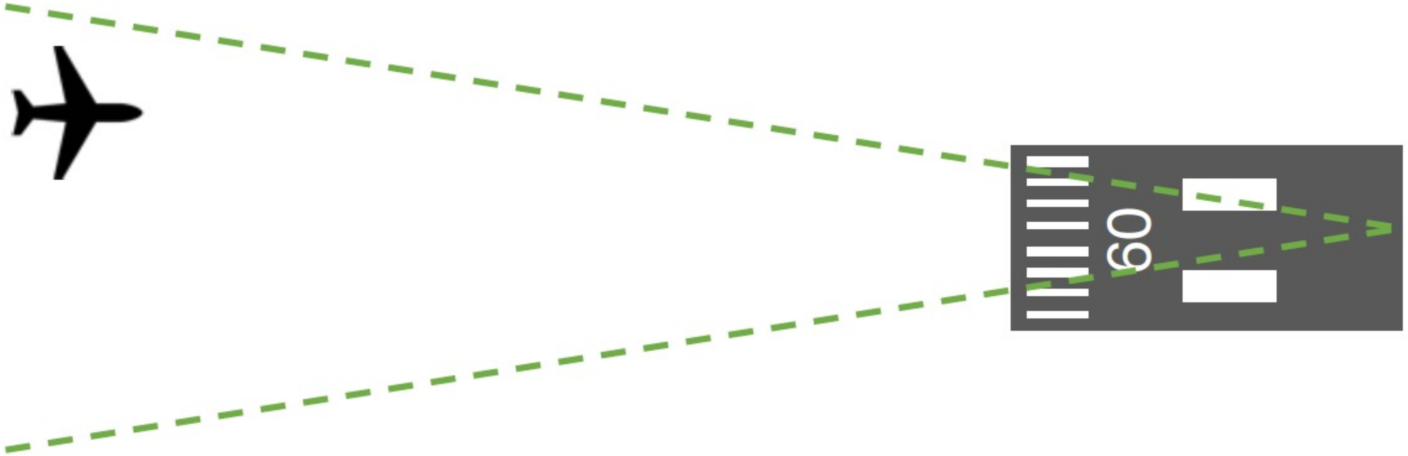

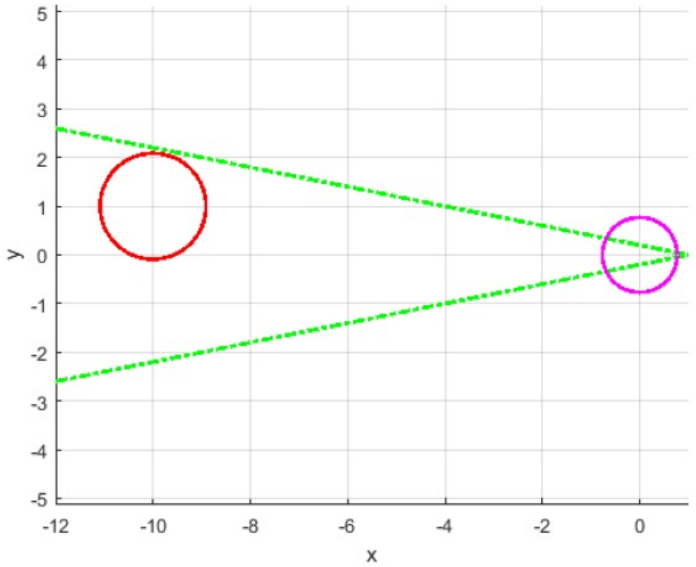

Example

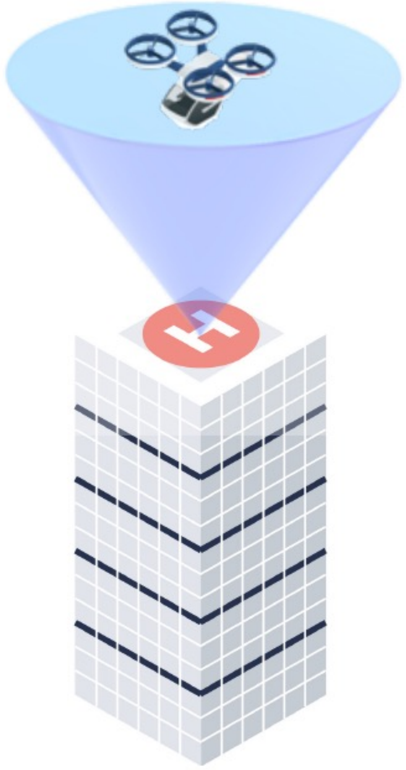

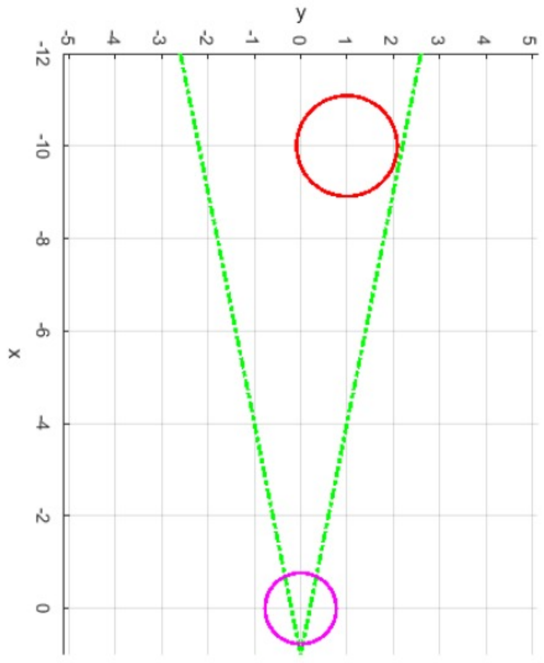

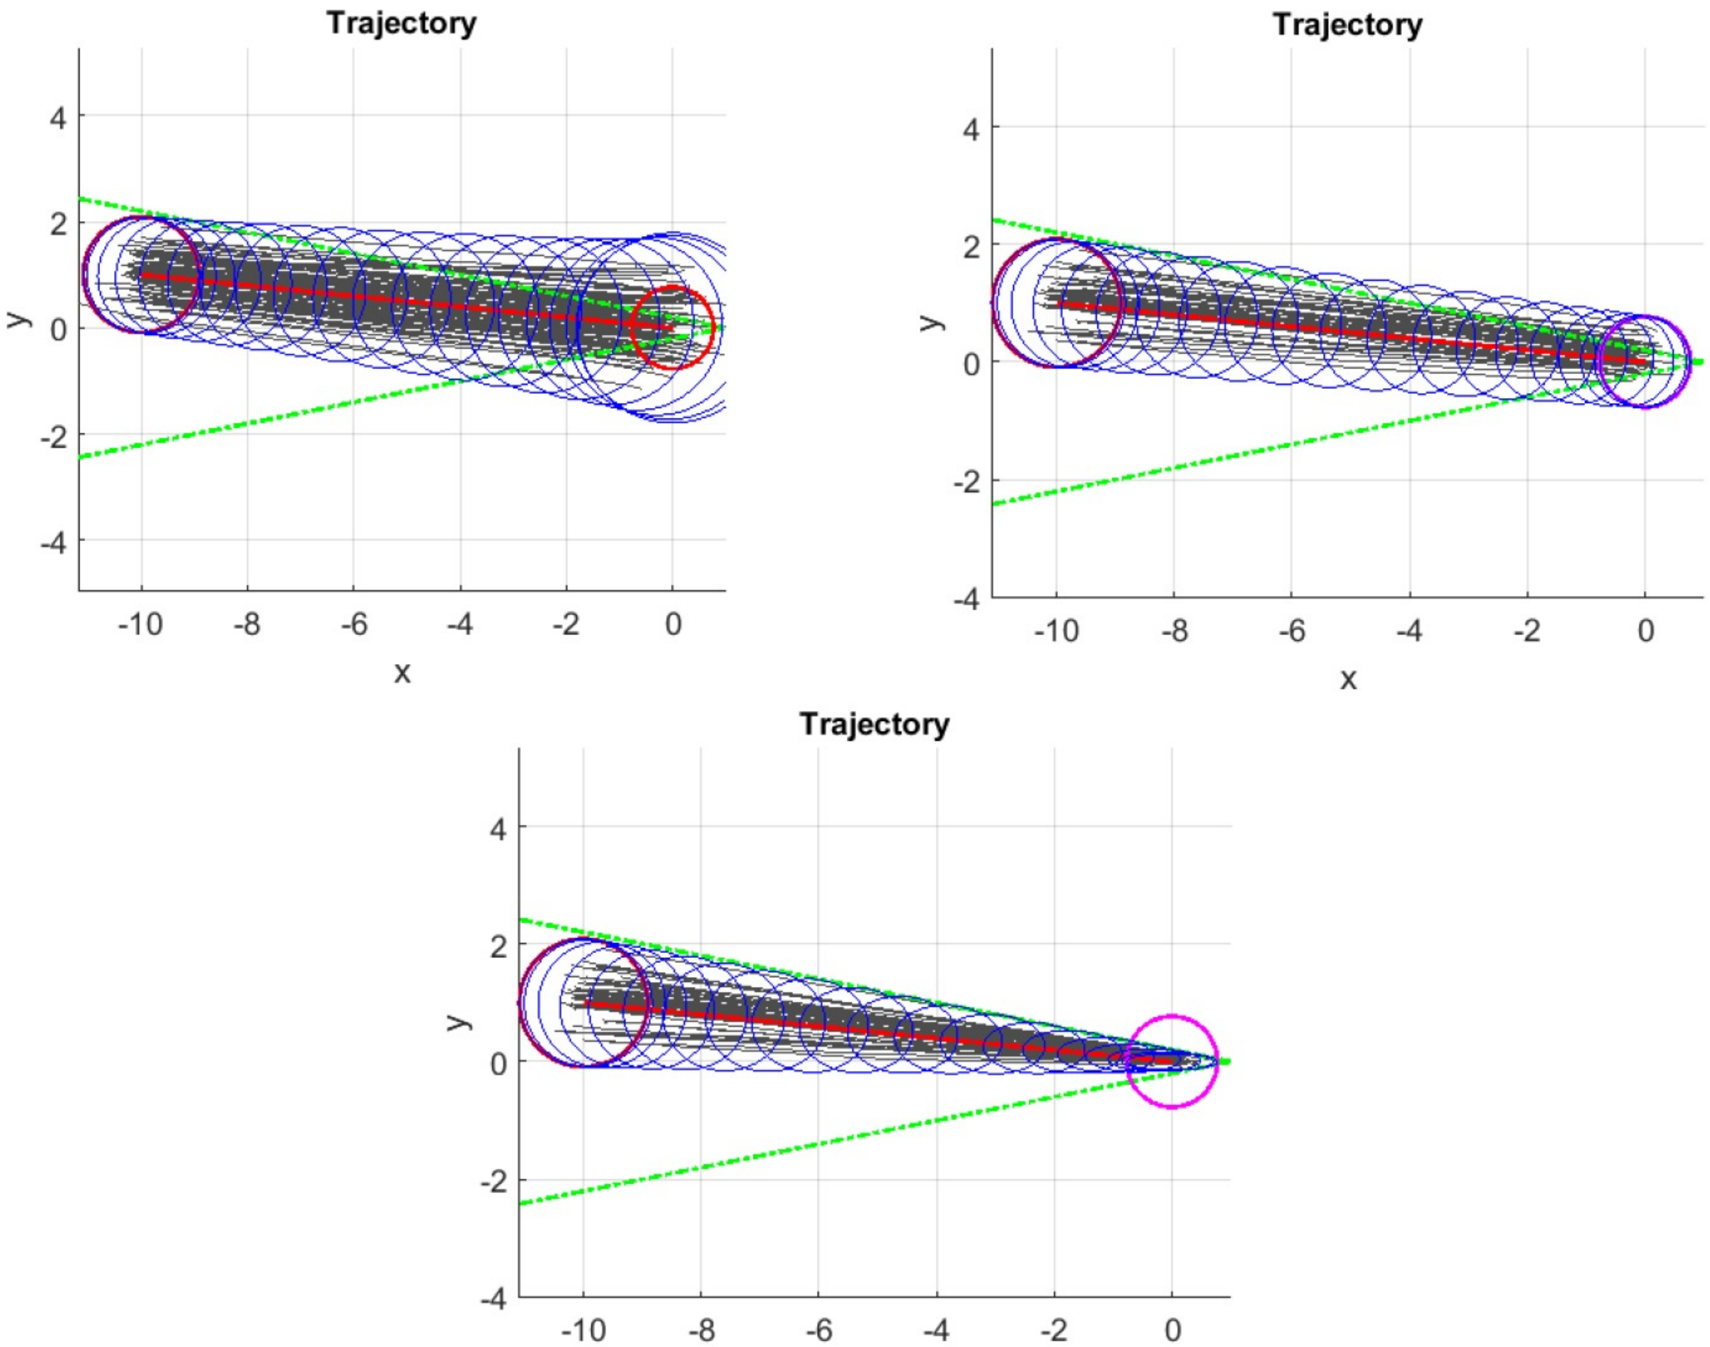
Trajectory
Trajectory
Trajectory
x
y
x
y
y

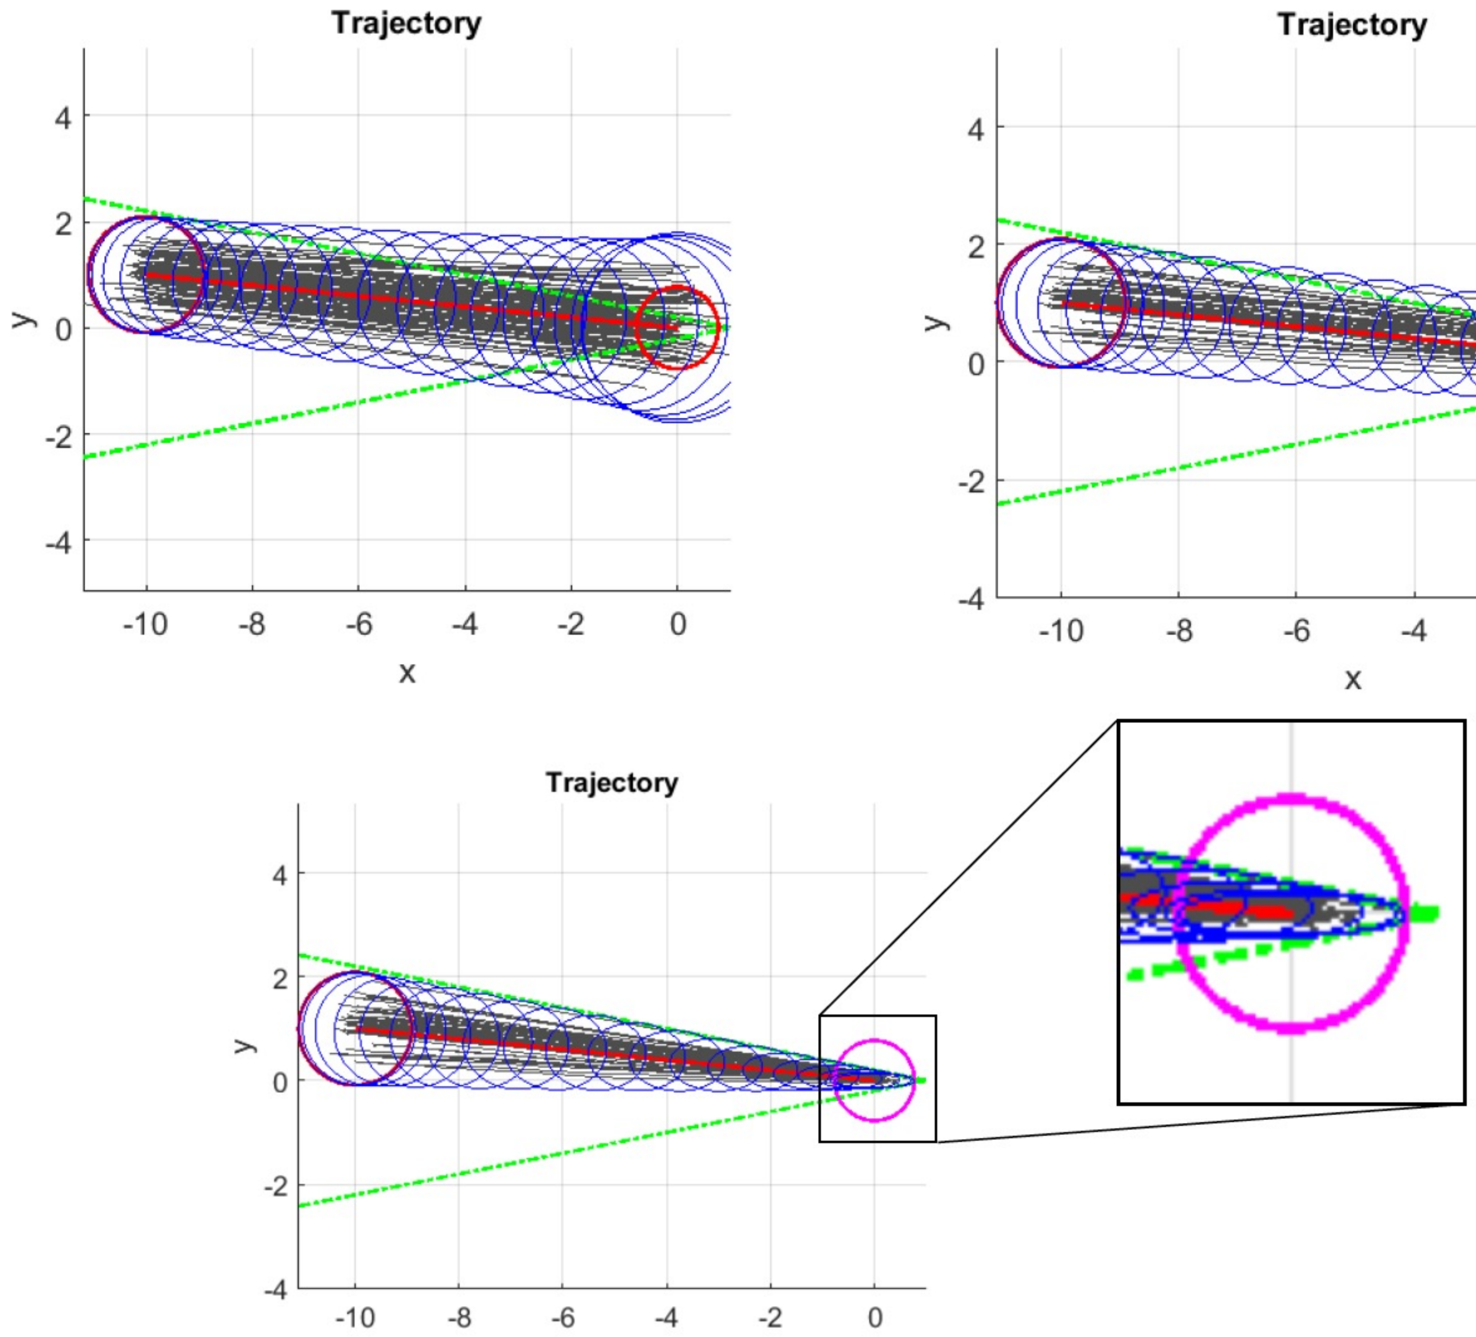
Trajectory
4
2
0
-2
-4
y
-10
-8
-6
-4
-2
0
x
Trajectory
4
2
0
-2
-4
y
-10
-8
-6
-4
-2
0
x
Trajectory
4
2
0
-2
-4
y
-10
-8
-6
-4
-2
0

Iterative Risk Allocation

- Chance constraint

$$\alpha_j^\top (\mathcal{A}\mu_0 + \mathcal{B}V) + \|(\mathcal{A}\Sigma_0\mathcal{A}^\top + \mathcal{D}\mathcal{D}^\top)^{1/2}(I + \mathcal{B}K)^\top \alpha_j\| \Phi^{-1}(1 - p_{j,\text{fail}}) - \beta_j \leq 0$$

- Iterate on risk thresholds $p_{j,\text{fail}}$

$$\texttt{Pr}(\alpha_j^\top X > \beta_j) \leq p_{j,\text{fail}}, \qquad \sum_{j=1}^{M} p_{j,\text{fail}} \leq P_{\text{fail}}$$

- Tends to maximize terminal covariance, while still satisfying the chance constraints
- Less conservative solutions than previous methodologies.
- Both polygonal and cone constraints

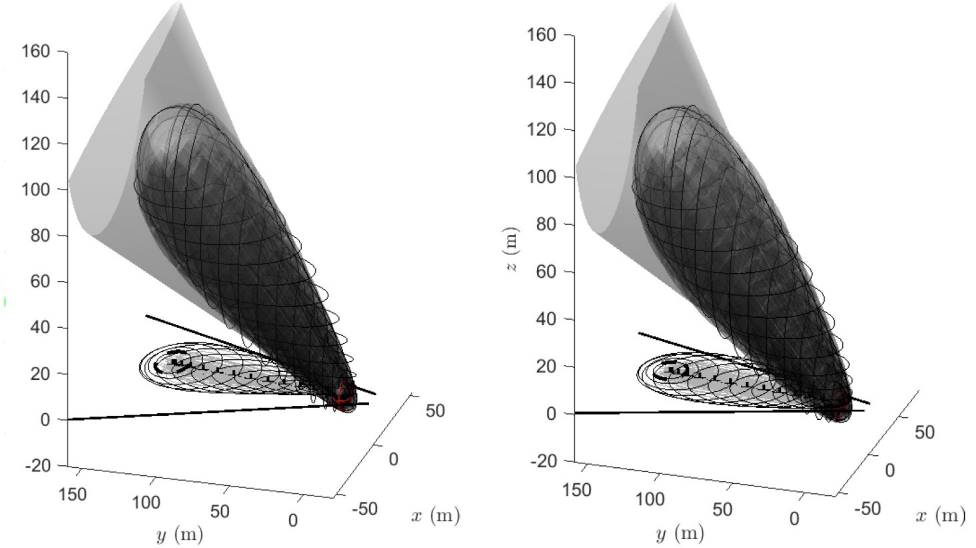

State and Control Constraints

- Assume that

$$x_k \in \mathcal{X}, \qquad u_k \in \mathcal{U}$$

where $\mathcal{X}$ and $\mathcal{U}$ are convex sets containing the origin.

$$\mathcal{X} = \bigcap_{j=0}^{N_s-1} \left\{ x : \alpha_{x,j}^\top x \leq \beta_{x,j} \right\}, \qquad \mathcal{U} = \bigcap_{s=0}^{N_c-1} \left\{ u : \alpha_{u,s}^\top u \leq \beta_{u,s} \right\},$$

- Since state is possibly unbounded we impose as chance constraints,

$$\Pr(x_k \notin \mathcal{X}) \leq \epsilon$$

- Keep $u_k \in \mathcal{U}$ since **hard constraints** for input are preferable.

Main Result

Theorem

The control law

$$u_k = v_k + K_k z_k$$

where z_k is governed by the dynamics

$$z_{k+1} = Az_k + \varphi(w_k)$$
$$z_0 = \varphi(\zeta_0), \quad \zeta_0 = x_0 - \mu_0$$

where $\varphi(\cdot) : \mathbb{R}^d \mapsto \mathbb{R}^d$ is an element-wise symmetric saturation function

$$\varphi_i(\zeta) = \max(-\zeta_i^{\max}, \min(\zeta_i, \zeta_i^{\max}))$$

converts the problem to a **convex programming problem**.

Numerical Example

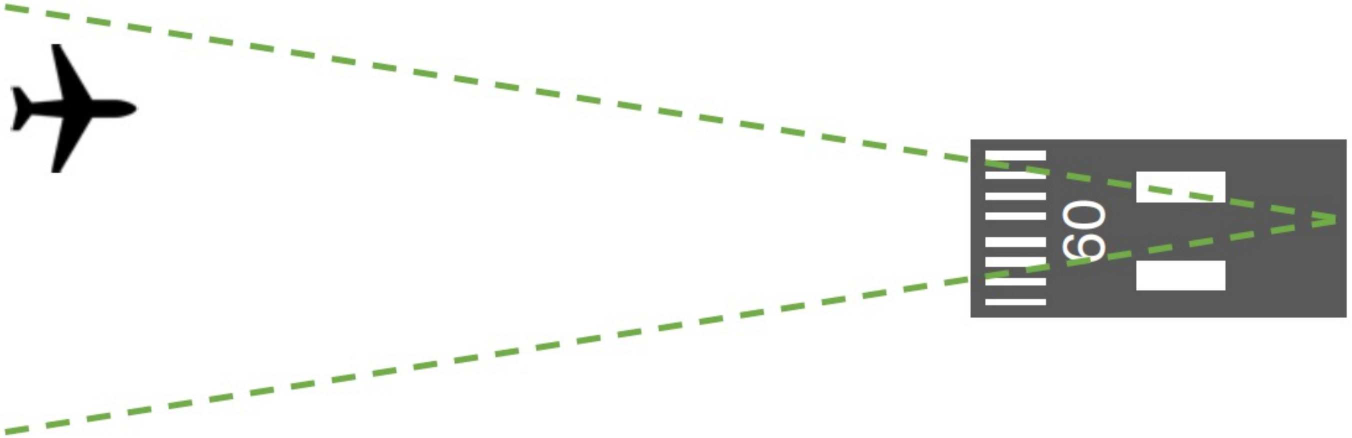

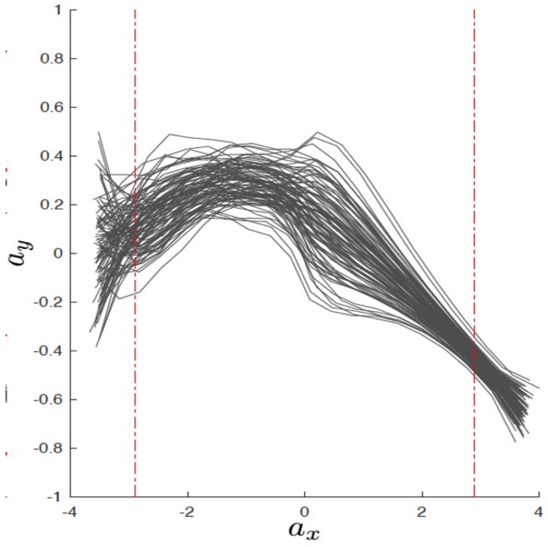

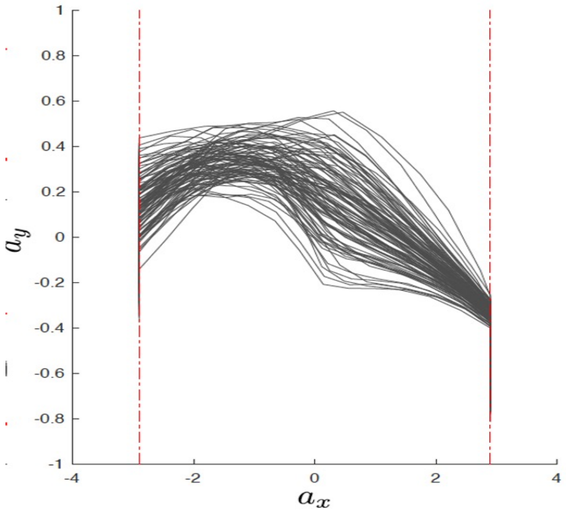

Acceleration limits: $a_x, a_y \leq 2.9\,\mathrm{m/s^2}$

Non-Convex Constraints

- For non-convex polytopic constraints, write

$$\chi = \bigcup_{r=0}^{N_R-1} \underbrace{\bigcap_{q=0}^{M_r-1} \{x : \alpha_{r,q}^\top x \leq \beta_{r,q}\}}_{\mathcal{R}_r}$$

  and enforce $\mathtt{Pr}\left(x_k \notin \mathcal{R}_r\right) < \epsilon$ and $\mathtt{Pr}\left(x_{k+1} \notin \mathcal{R}_r\right) < \epsilon$

Lemma

Given $\mathcal{R}_r$, the condition

$$\mathtt{Pr}\left(x_k \notin \mathcal{R}_r\right) < \epsilon \quad \text{and} \quad \mathtt{Pr}\left(x_{k+1} \notin \mathcal{R}_r\right) < \epsilon,$$

is a second-order cone constraint in V and K.

Example

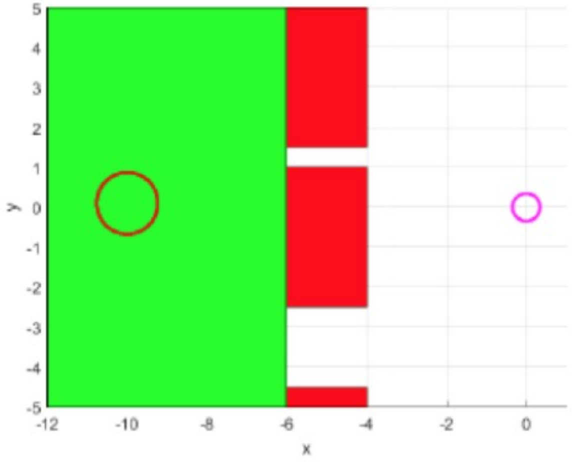

(a) Region 1

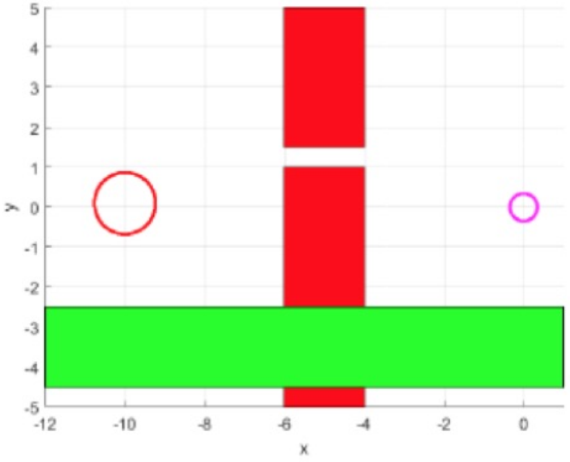

(b) Region 2

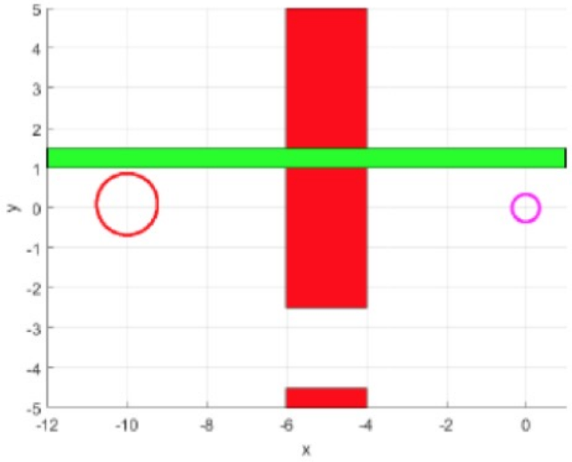

(c) Region 3

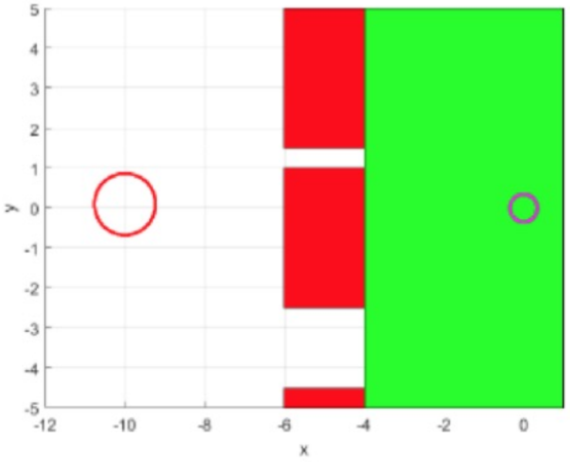

(d) Region 4

$N = 20$

$\varepsilon = 1e\text{-}3$

Example

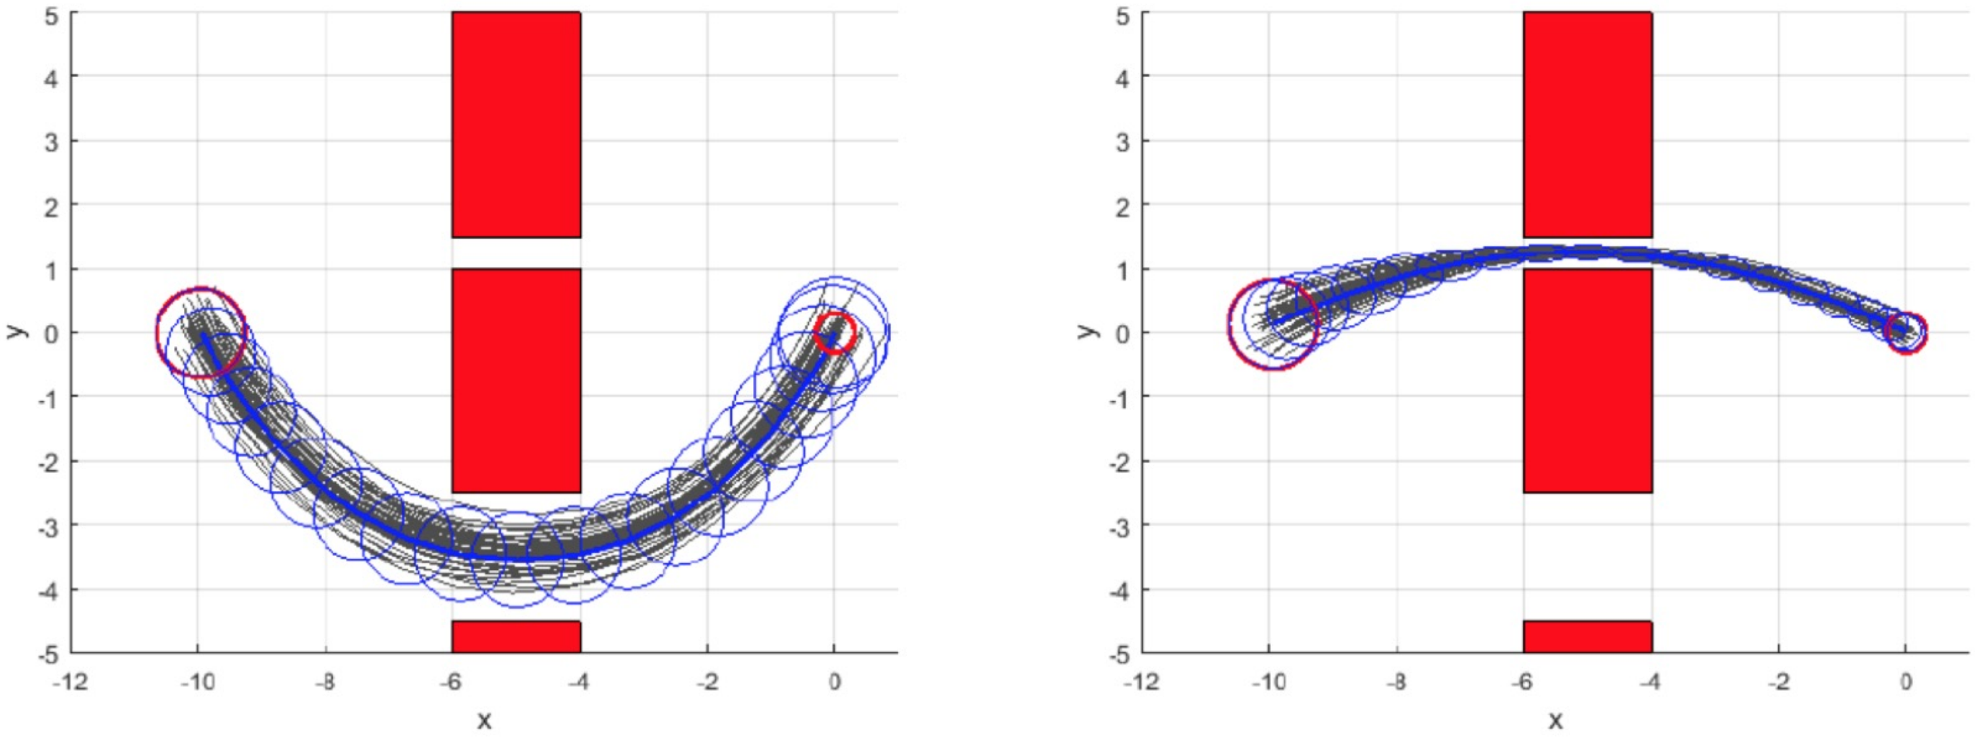

Reference trajectory depends on uncertainty

Sampling-Based Planning

- High-dimensional spaces
- Many variants: RRT, PRM, RRT*, RRT#, BIT*,...
- Do not handle uncertainty directly
- Can do planning in **belief** space: BRM, FIRM, ...
  - Uncertainty handled indirectly
  - Nodes have to be stabilized
- Can we do better?

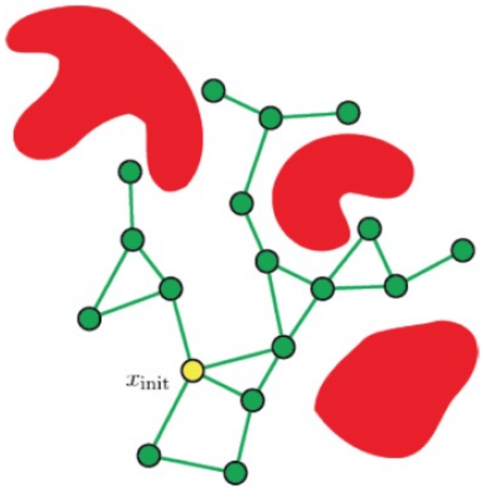

CS-BRM Motion Planning

Idea: Use CS controller as edge controller to perform planning in belief space

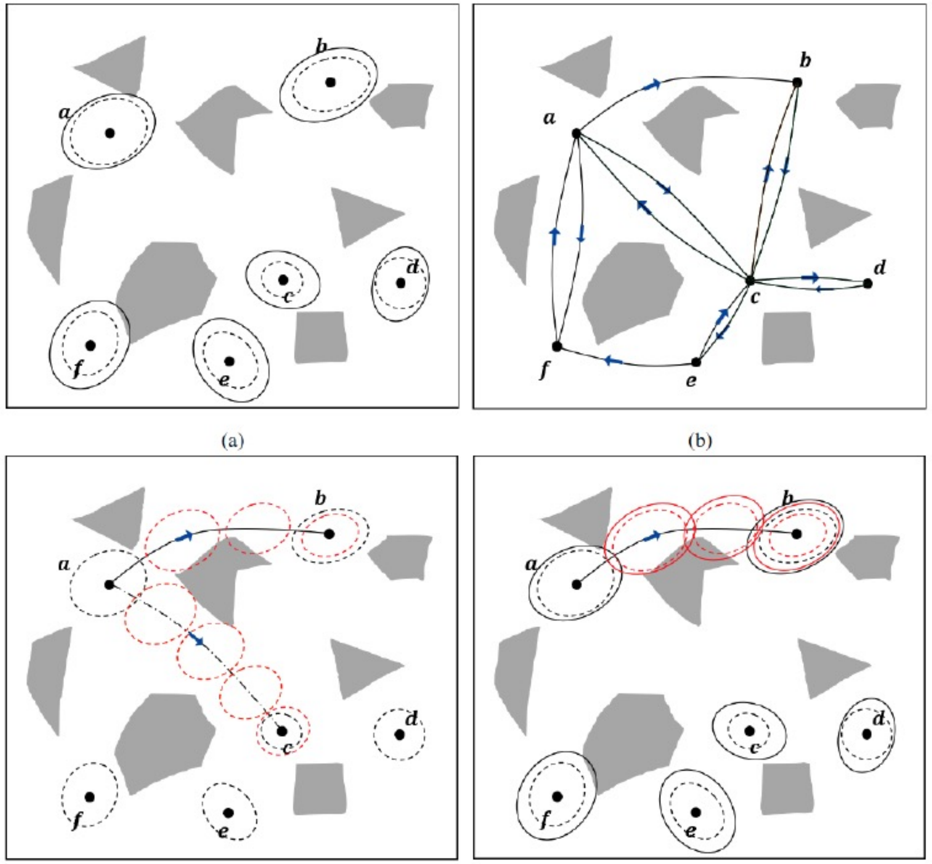

(a) (b) (c) (d)

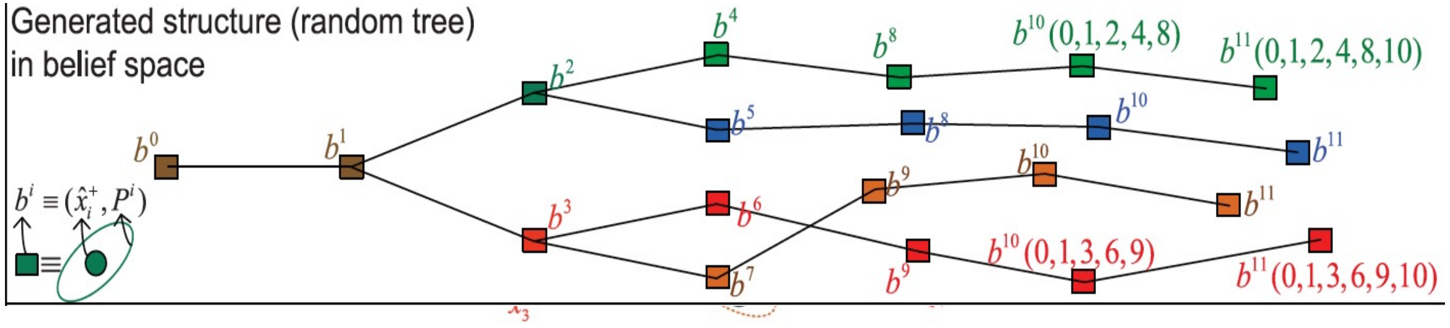

"FIRM: Sampling-based feedback motion-planning under motion uncertainty and imperfect measurements" A. Agha-mohammadi, S. Chakravorty, N. Amato

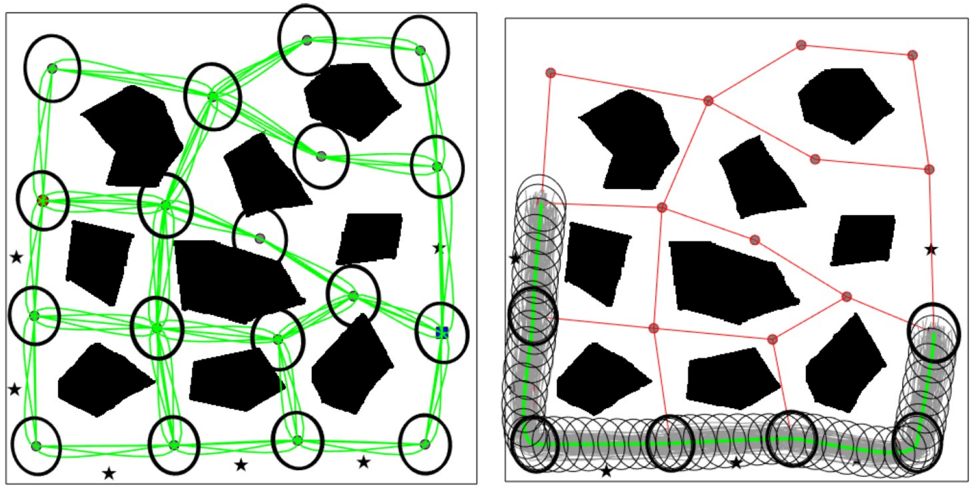

CS-Belief Space Planning (CS-BRM)

- Enables guaranteed satisfaction of terminal belief constraints in finite-time.
- The CS-BRM algorithm allows the sampling of non-stationary belief nodes, and thus is able to explore the velocity space and find efficient motion plans
- Addresses "node reachability" and "curse of history" problems of traditional BRMs

Shortest path has a higher probability of collision

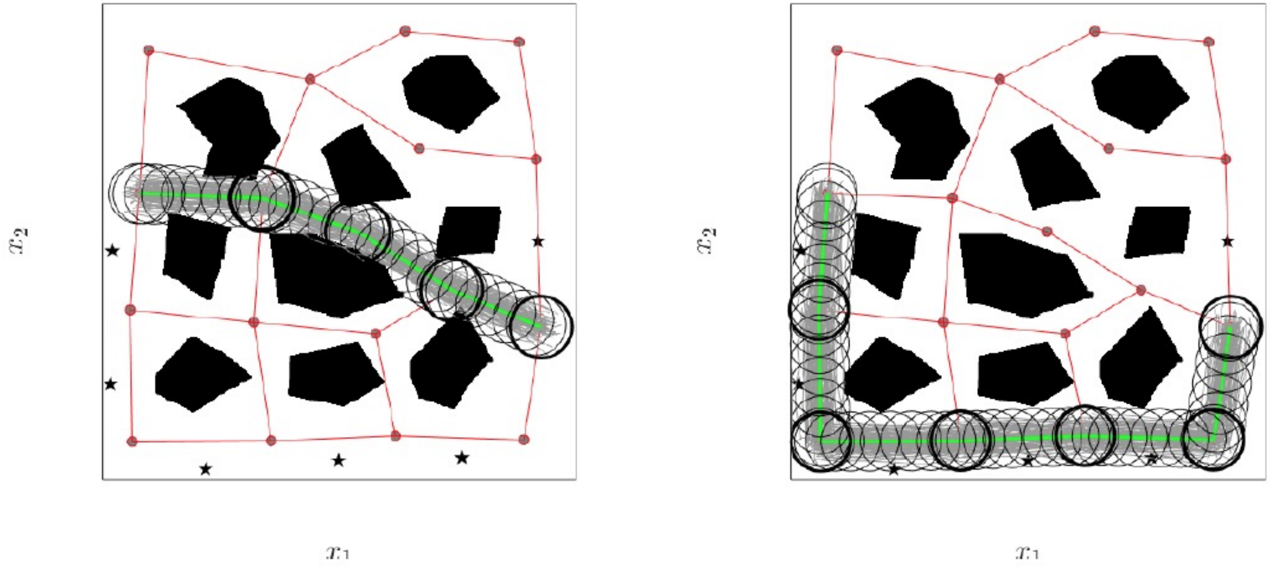

- Explicitly incorporate motion and observation uncertainties
- Directly control the belief between the BRM nodes
- Nodes do not have to be stationary
- Better explore the velocity space and find paths with lower cost.

Covariance Steering Stochastic MPC

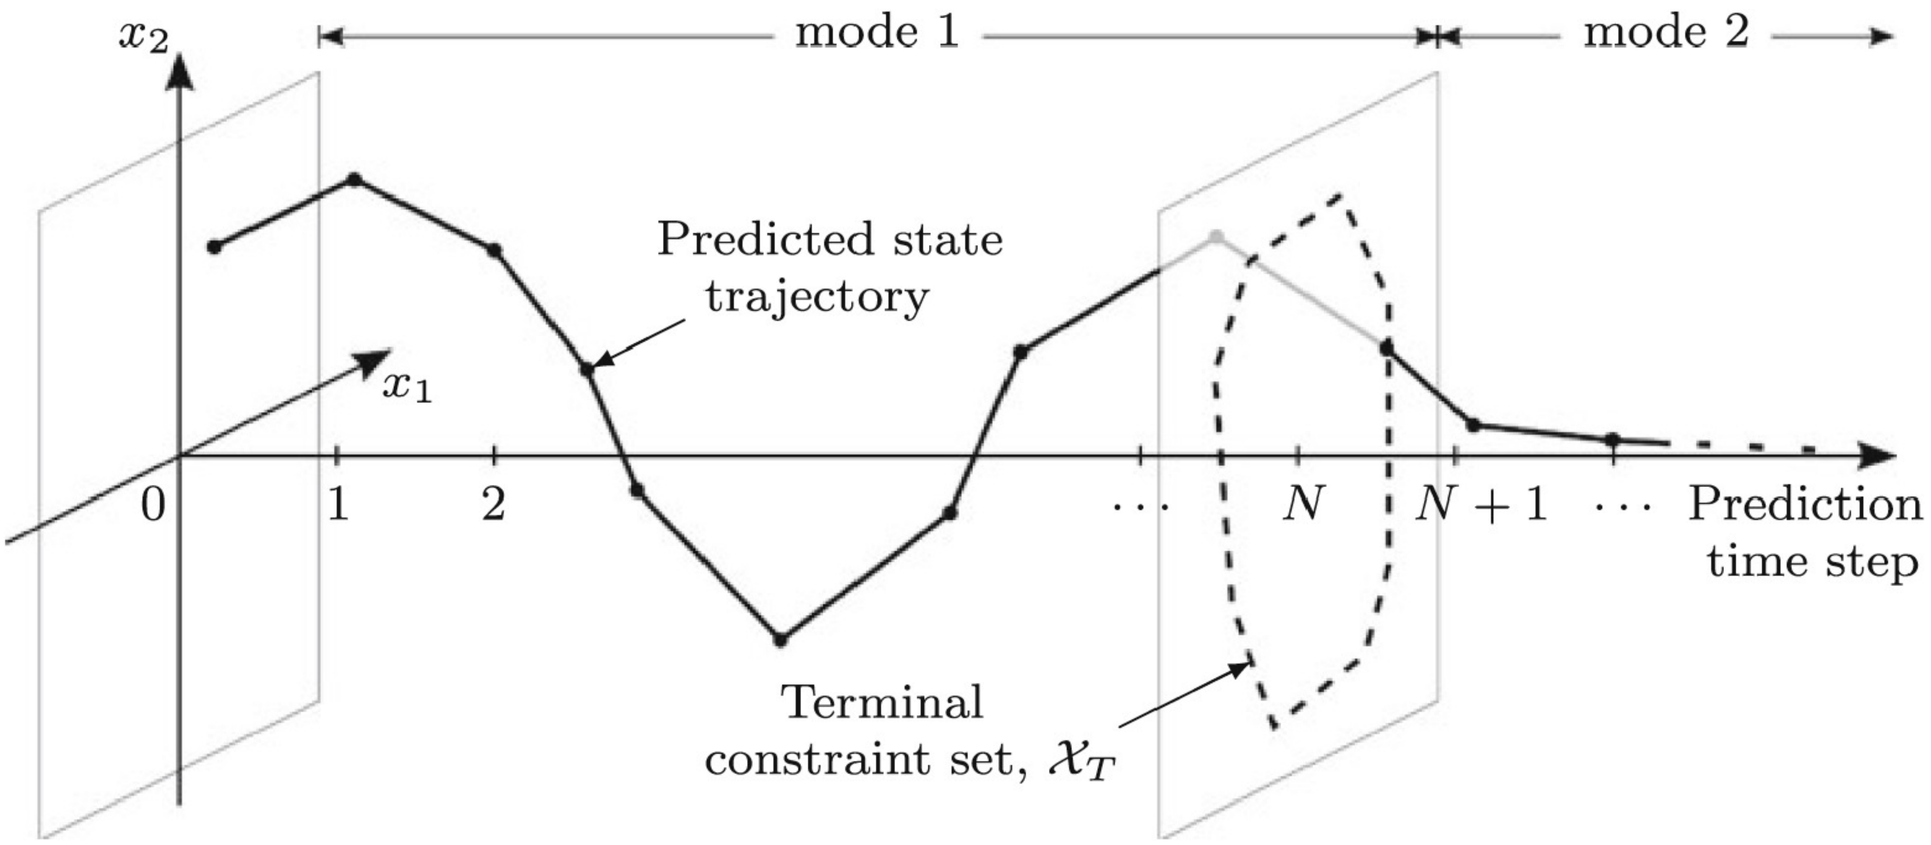

(*Model Predictive Control: Classical, Robust and Stochastic*, B. Kouvaritakis and M. Cannon)

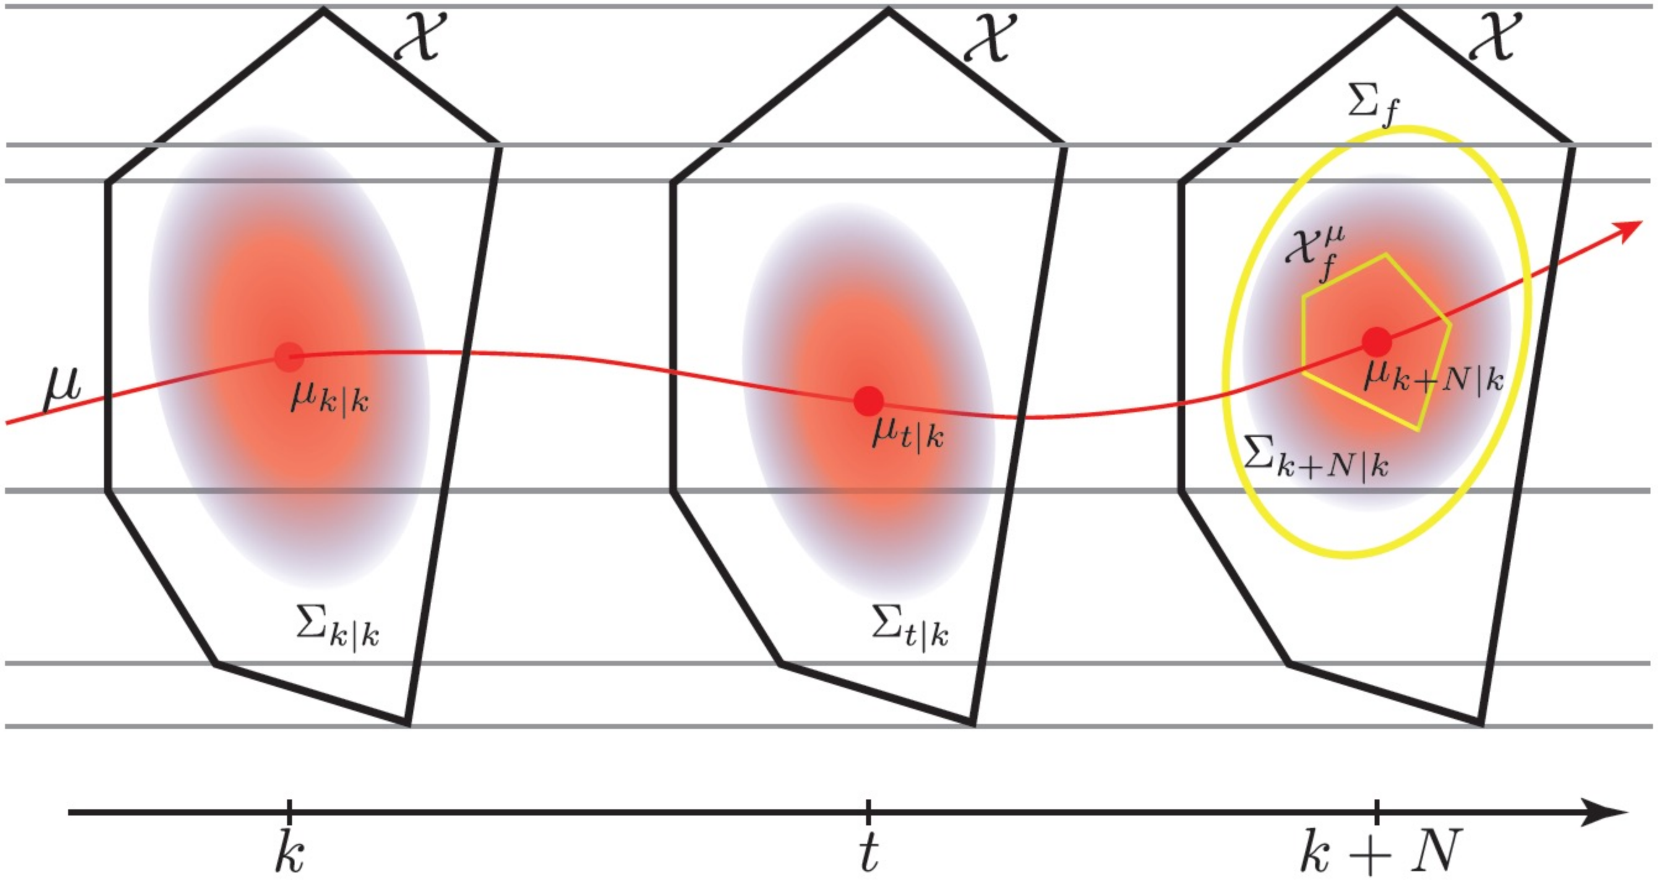

$\mathcal{X}$
$\mathcal{X}$
$\mathcal{X}$
Σ_f
$\mathcal{X}_f^\mu$
μ
$\mu_{k|k}$
$\mu_{t|k}$
$\mu_{k+N|k}$
$\Sigma_{k+N|k}$
$\Sigma_{k|k}$
$\Sigma_{t|k}$
k
t
$k+N$

Stochastic MPC

$$\min_{u_{k|k}, u_{k+1|k}, \ldots, u_{k+N-1|k}} J_N(x_k; u_{k|k}, u_{k+1|k}, \ldots, u_{k+N-1|k}) =$$

$$\mathbb{E}_k\left[\sum_{t=k}^{k+N-1} x_{t|k}^\top Q x_{t|k} + u_{t|k}^\top R u_{t|k}\right] + \mathbb{E}_k[x_{k+N|k}]^\top P_{\text{mean}} \mathbb{E}_k[x_{k+N|k}]$$

subject to

$$x_{t+1|k} = Ax_{t|k} + Bu_{t|k} + Dw_t, \qquad x_{k|k} = x_k \sim \mathcal{N}(\mu_k, \Sigma_k)$$

$$\Pr_k\left(\alpha_{x,i}^\top x_{t|k} \leq \beta_{x,i}\right) \geq 1 - p_{x,i}, \quad i = 0, \ldots, N_s - 1$$

$$\Pr_k\left(\alpha_{u,j}^\top u_{t|k} \leq \beta_{u,j}\right) \geq 1 - p_{u,j}, \quad j = 0, \ldots, N_c - 1$$

$$\mathbb{E}_k\left[x_{k+N|k}\right] \in \mathcal{X}_f^\mu$$

$$\mathbb{E}_k\left[(x_{k+N|k} - \mathbb{E}[x_{k+N|k}])(x_{k+N|k} - \mathbb{E}[x_{k+N|k}])^\top\right] \preceq \Sigma_f$$

Theorem

Suppose that Σ_f *is assignable,* $\mu_f \in \mathcal{X}_f^\mu$*, such that for all* $\mu \in \mathcal{X}_f^\mu$

$$
\begin{aligned}
&(A + B\tilde{K})\mu \in \mathcal{X}_f^\mu \\
&\alpha_{x,i}^\top \mu + \|{\Sigma_f}^{1/2}\alpha_{x,i}\|\Phi^{-1}(1 - p_{x,i}) - \beta_{x,i} \leq 0 \\
&\alpha_{u,j}^\top \tilde{K}\mu + \|{\Sigma_f}^{1/2}\tilde{K}^\top \alpha_{u,j}\|\Phi^{-1}\left(1 - p_{u,j}\right) - \beta_{u,j} \leq 0
\end{aligned}
$$

where $\tilde{K}$ *is from corresponding assignability gain matrix, and* P_{mean} *is the solution of the discrete-time Lyapunov equation*

$$
(A + B\tilde{K})^\top P_{\text{mean}}(A + B\tilde{K}) - P_{\text{mean}} + Q + \tilde{K}^\top R\tilde{K} = 0
$$

Then, the solution ensures recursive feasibility and stability.

Uncontrolled

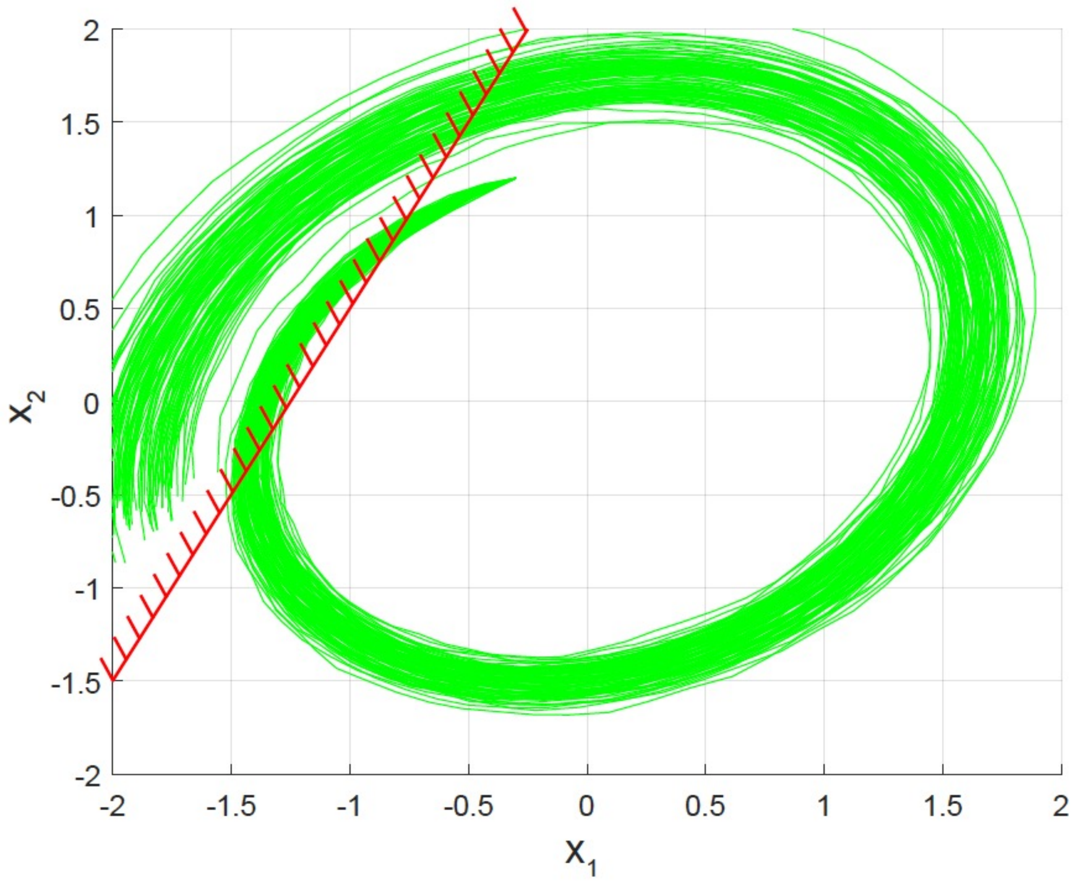

Infinite Horizon LQR

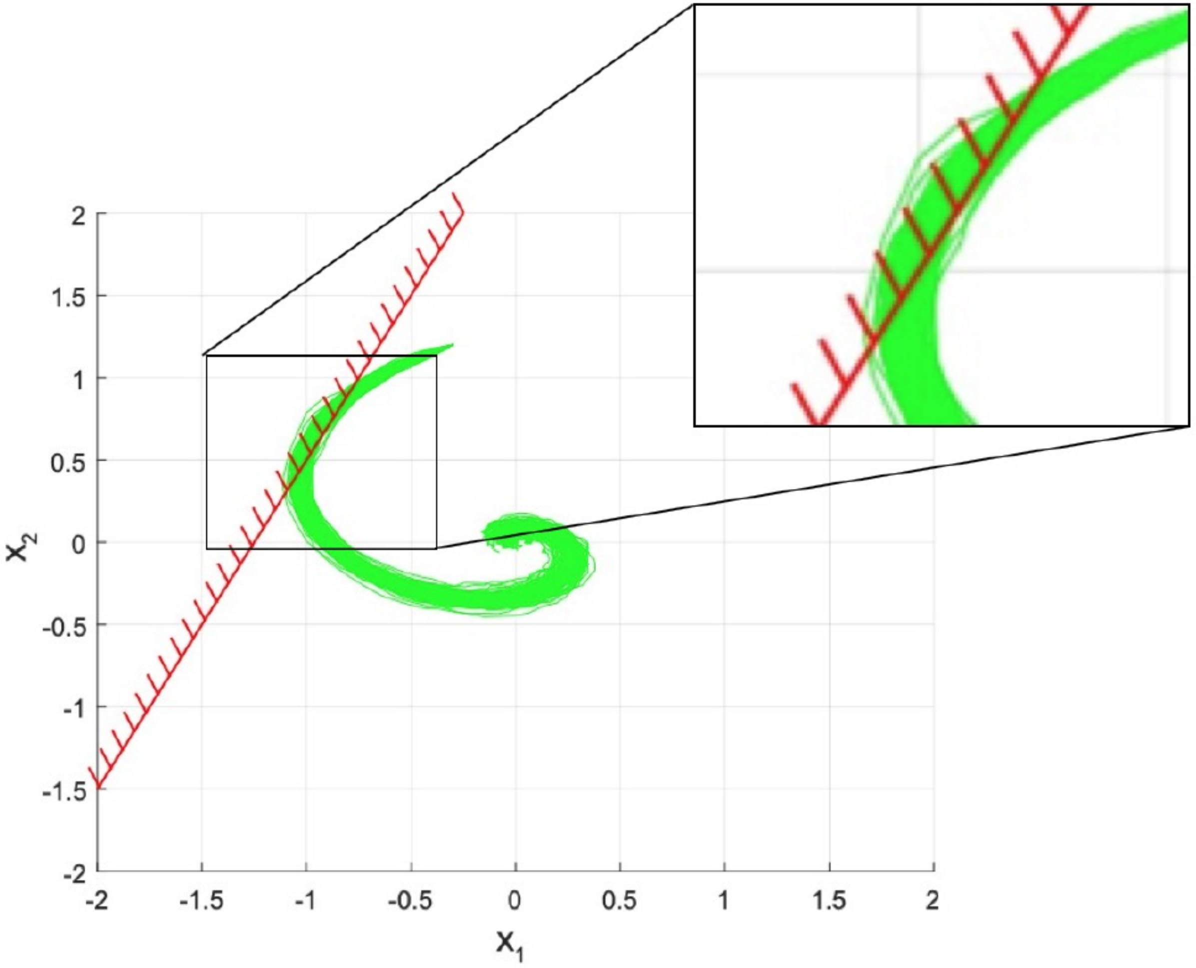

CS-SMPC

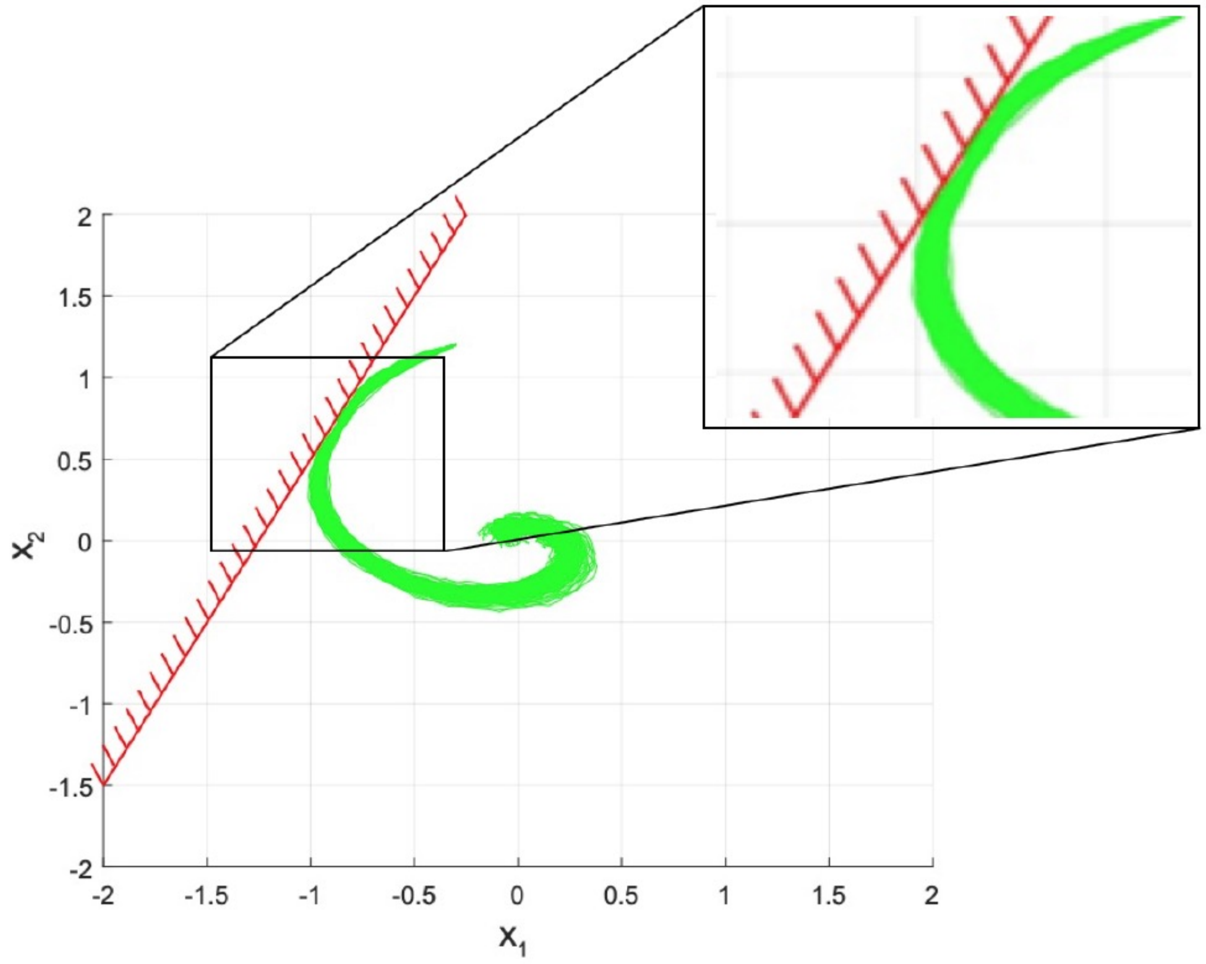

AutoRally

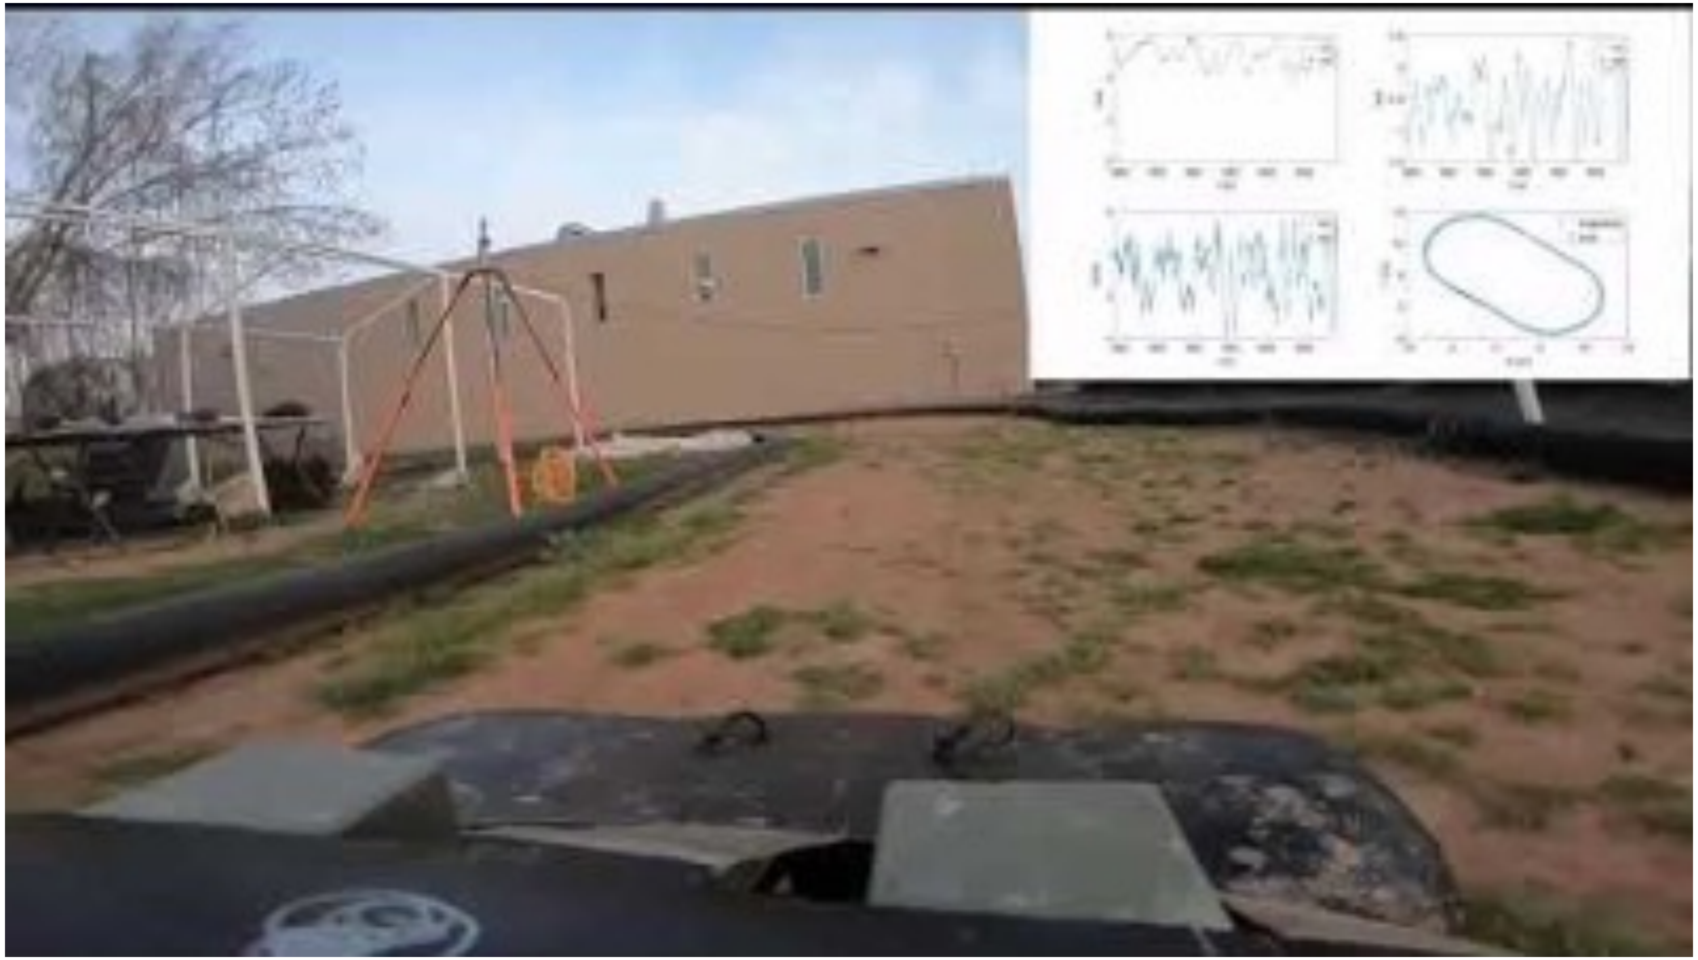

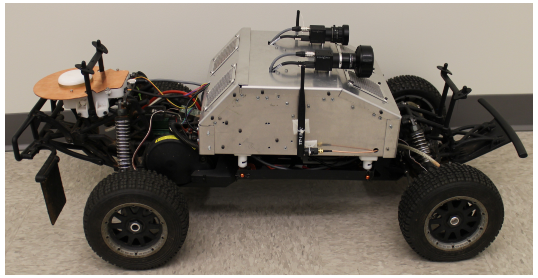

Take-Aways

- Directly controlling a **distribution of trajectories** results in strict performance guarantees.
- Eliminate the need for extensive Monte Carlo analysis.
- Many, many applications.
- For linear systems with Gaussian noise, theory well-developed.

Some Extensions

- Output feedback (Bakolas, 2019; Ridderhof and PT, 2020; Maity and PT, 2021)
- Nonlinear systems (Caluya and Halder 2019; Ridderhof, Okamoto, and PT, 2019)
- Differential games (Makkapatti, Okamoto and PT, 2020)
- Non-Gaussian noise (Sivaramakrishanan, Oishi, Pilipovsky, and PT, 2021)
- Hybrid & Switched Systems (Pakniyat and PT, 2021)