Carnegie Mellon University
Electrical & Computer Engineering

Long-term Safety for Autonomous Systems

Zhuoyuan (Jacob) Wang

3/4/2022

Autonomous control systems

autonomous vehicles

industrial robots

drones

Why robust safety in stochastic systems?

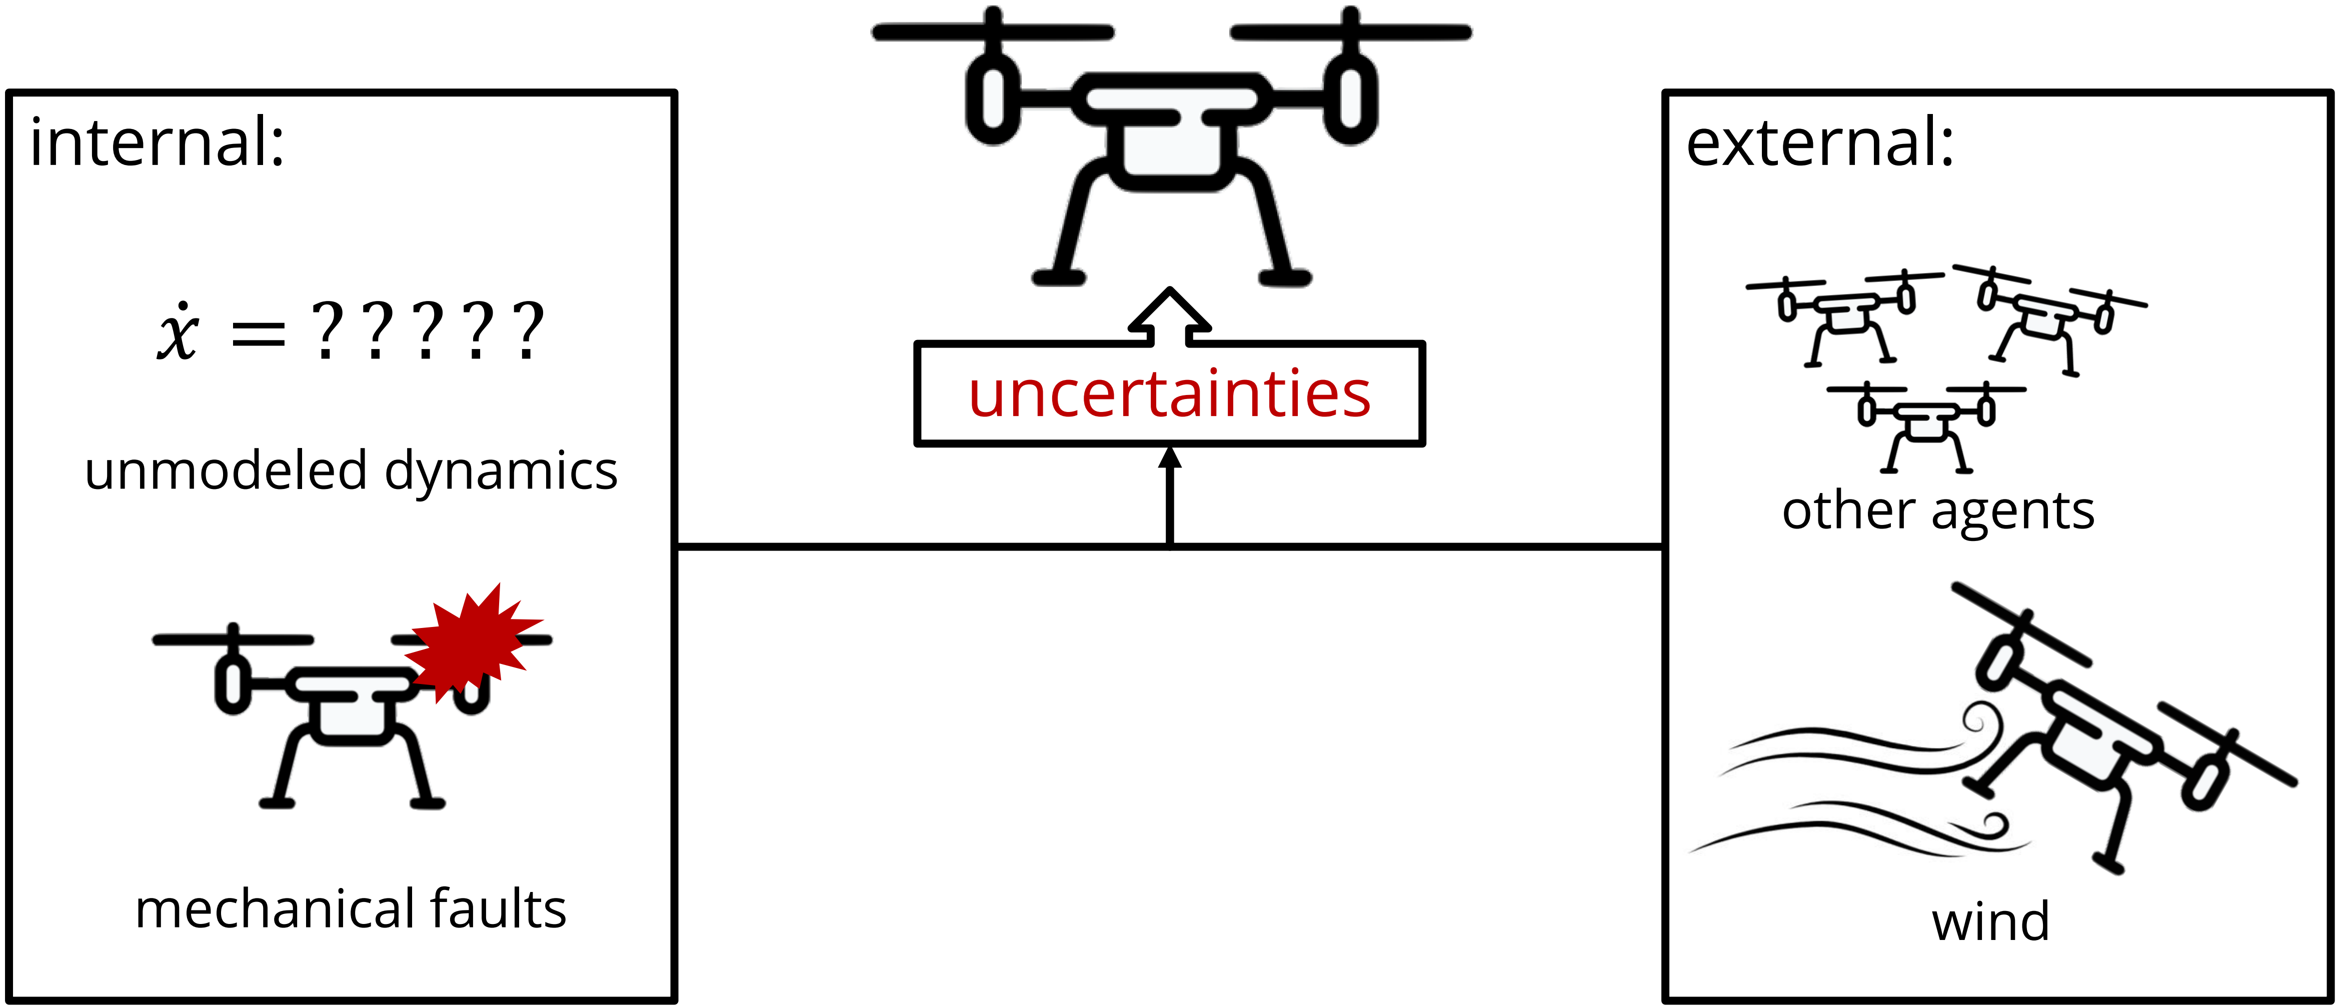

Existing methods

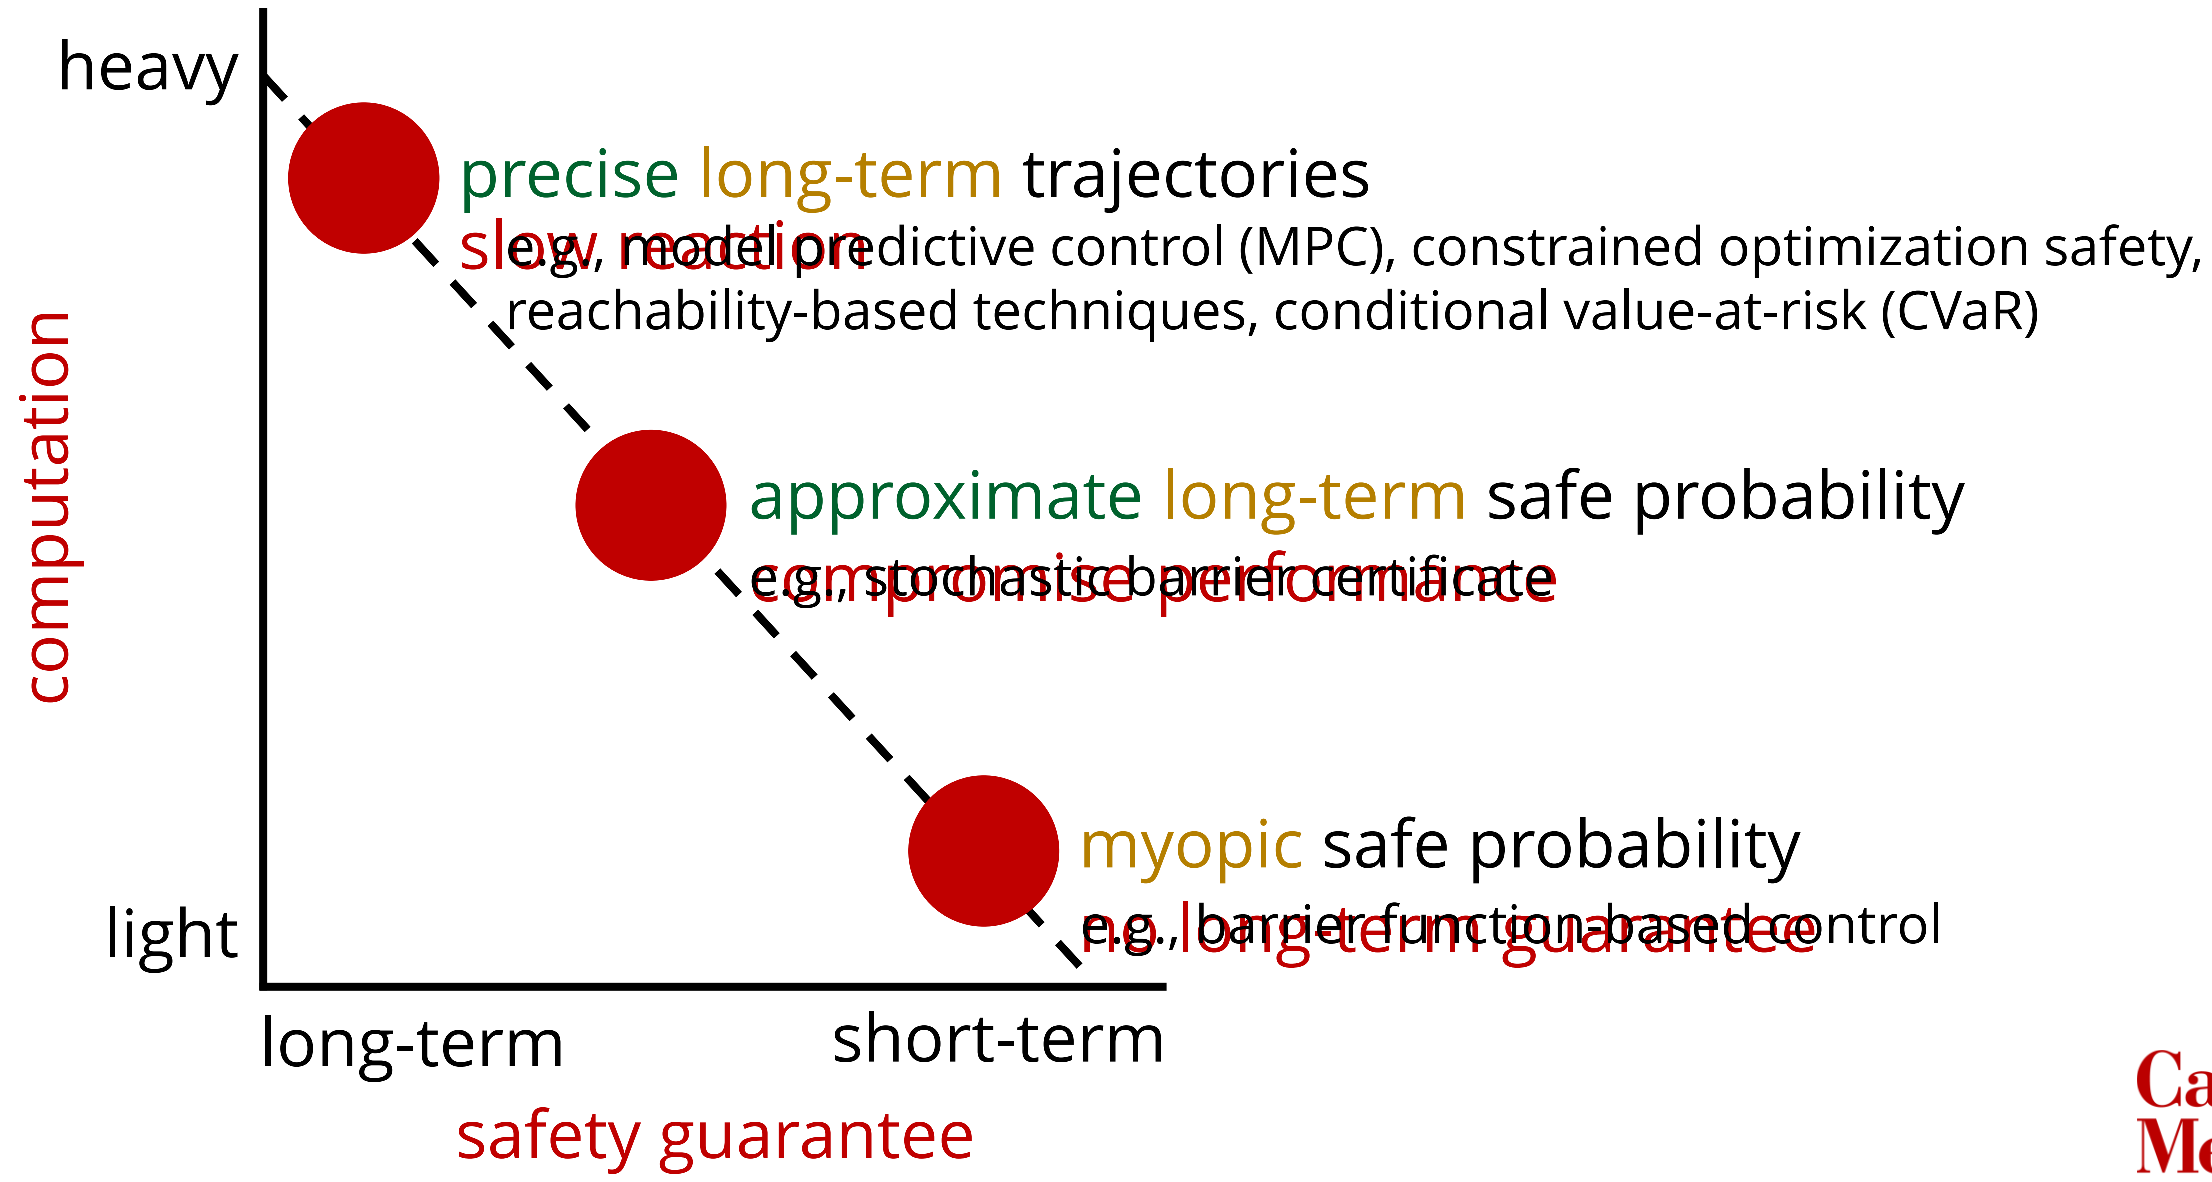

Carnegie Mellon University

Project aims

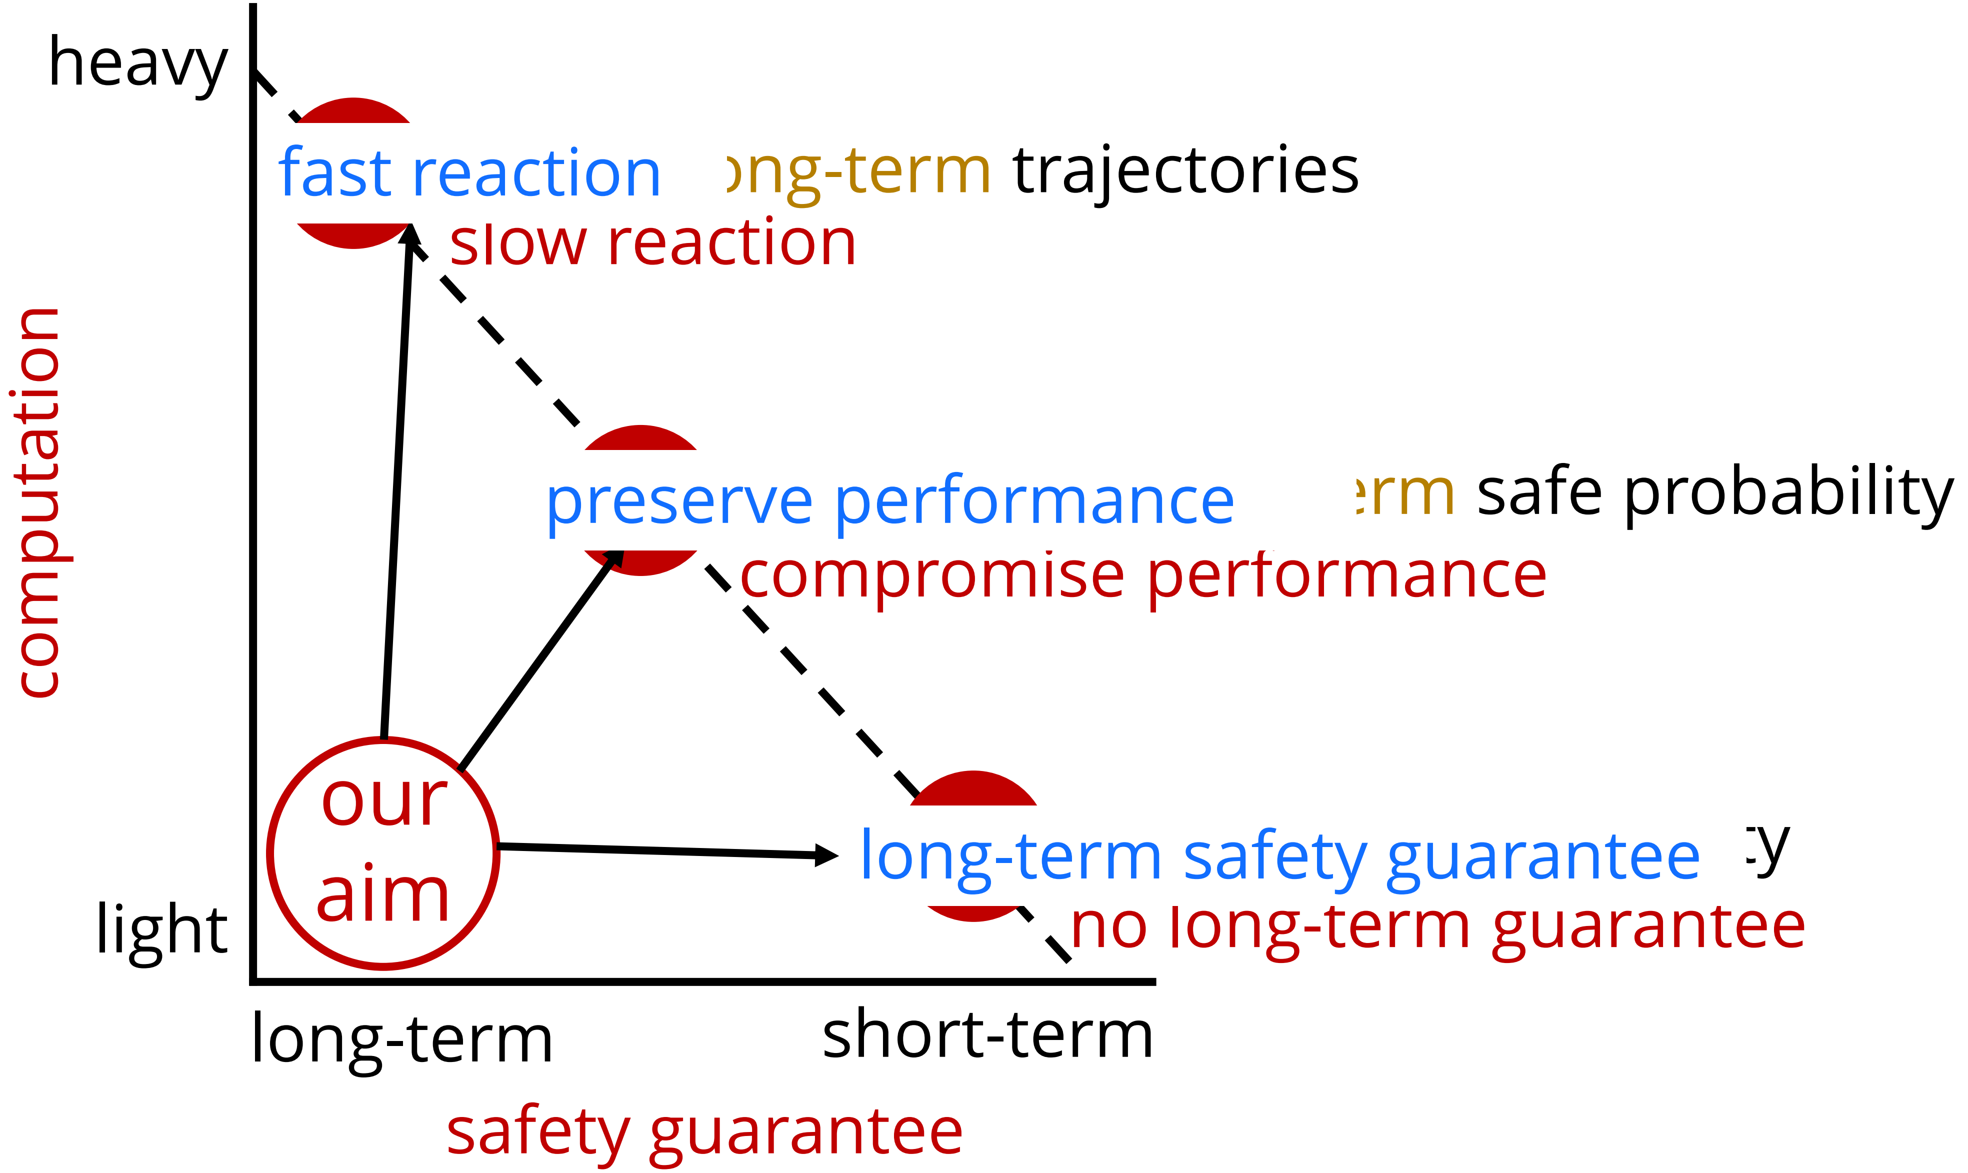

System description

X_t : state of the system

U_t : control input

W_t : system uncertainty

$\phi(x)$: barrier function

$\mathcal{C}(0)$: safe set

$\mathcal{C}(\underbrace{L}_{\text{Safety margin}})$: L-safe set

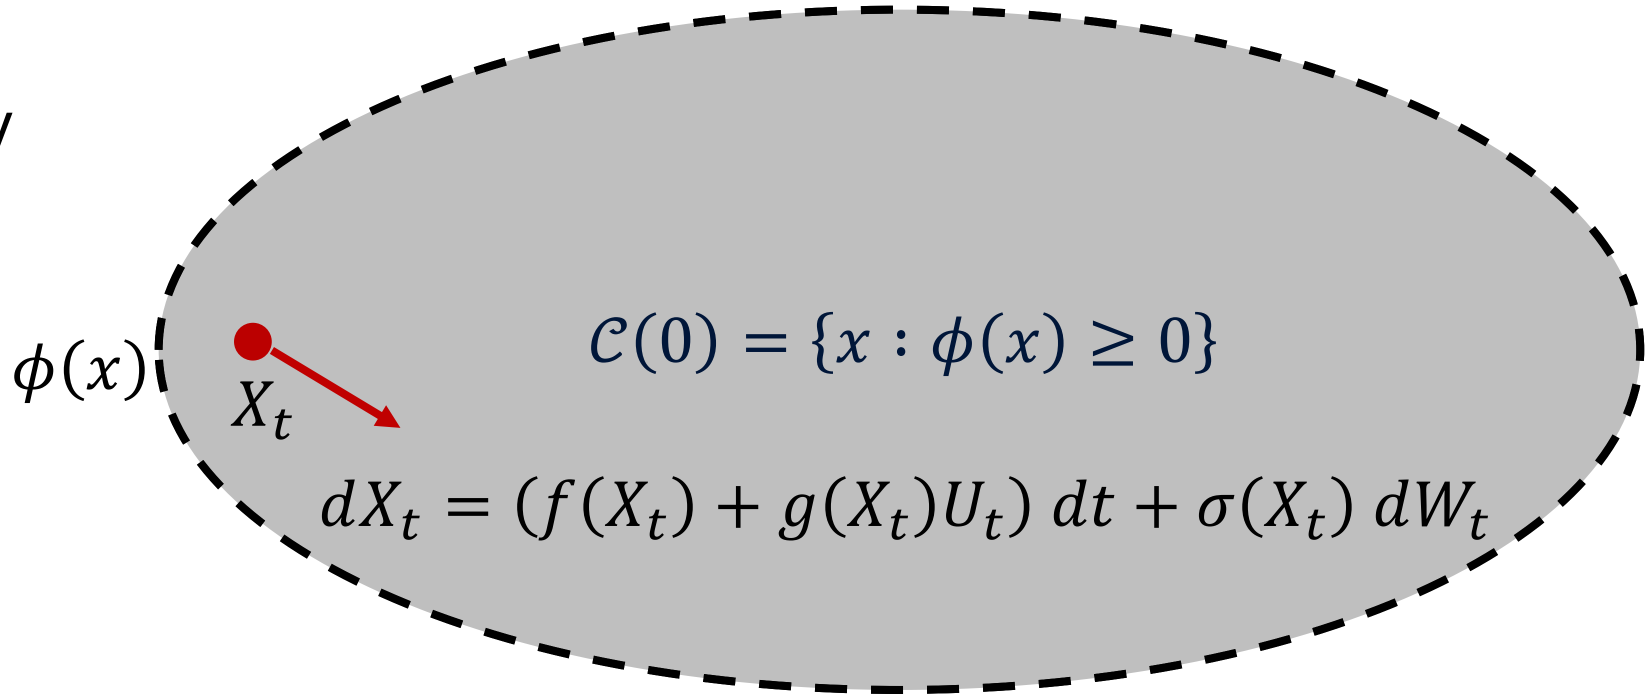

Safety in stochastic system: time horizon and safety margin

We consider fixed time horizon:

$T_t = H$, safety evaluated at t for $[t, t+H]$

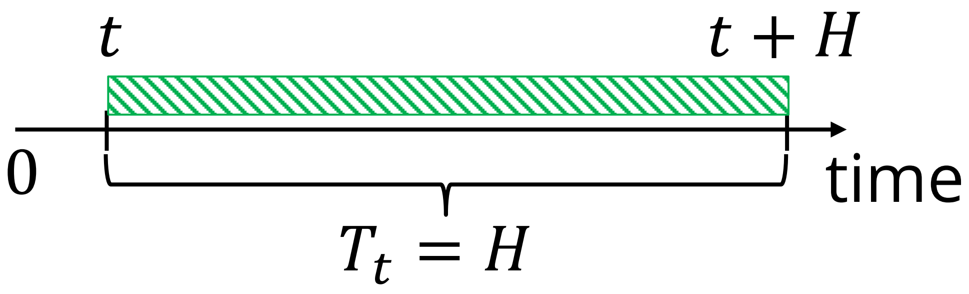

And receding time horizon:

$T_t = H - t$, safety evaluated at t for $[t, H]$

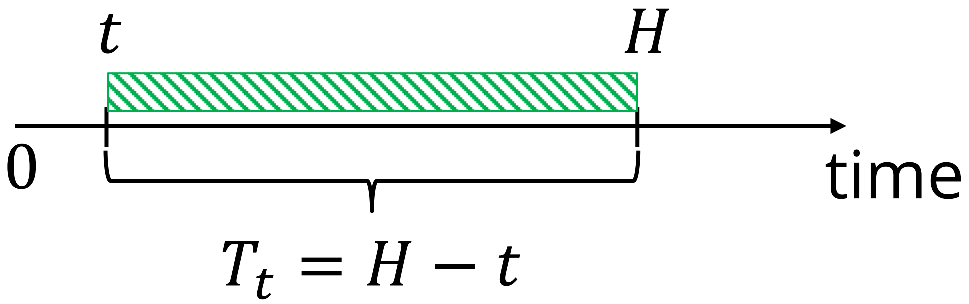

Deterministic systems (noiseless)

safe at next time => safe at all time

Myopic evaluation

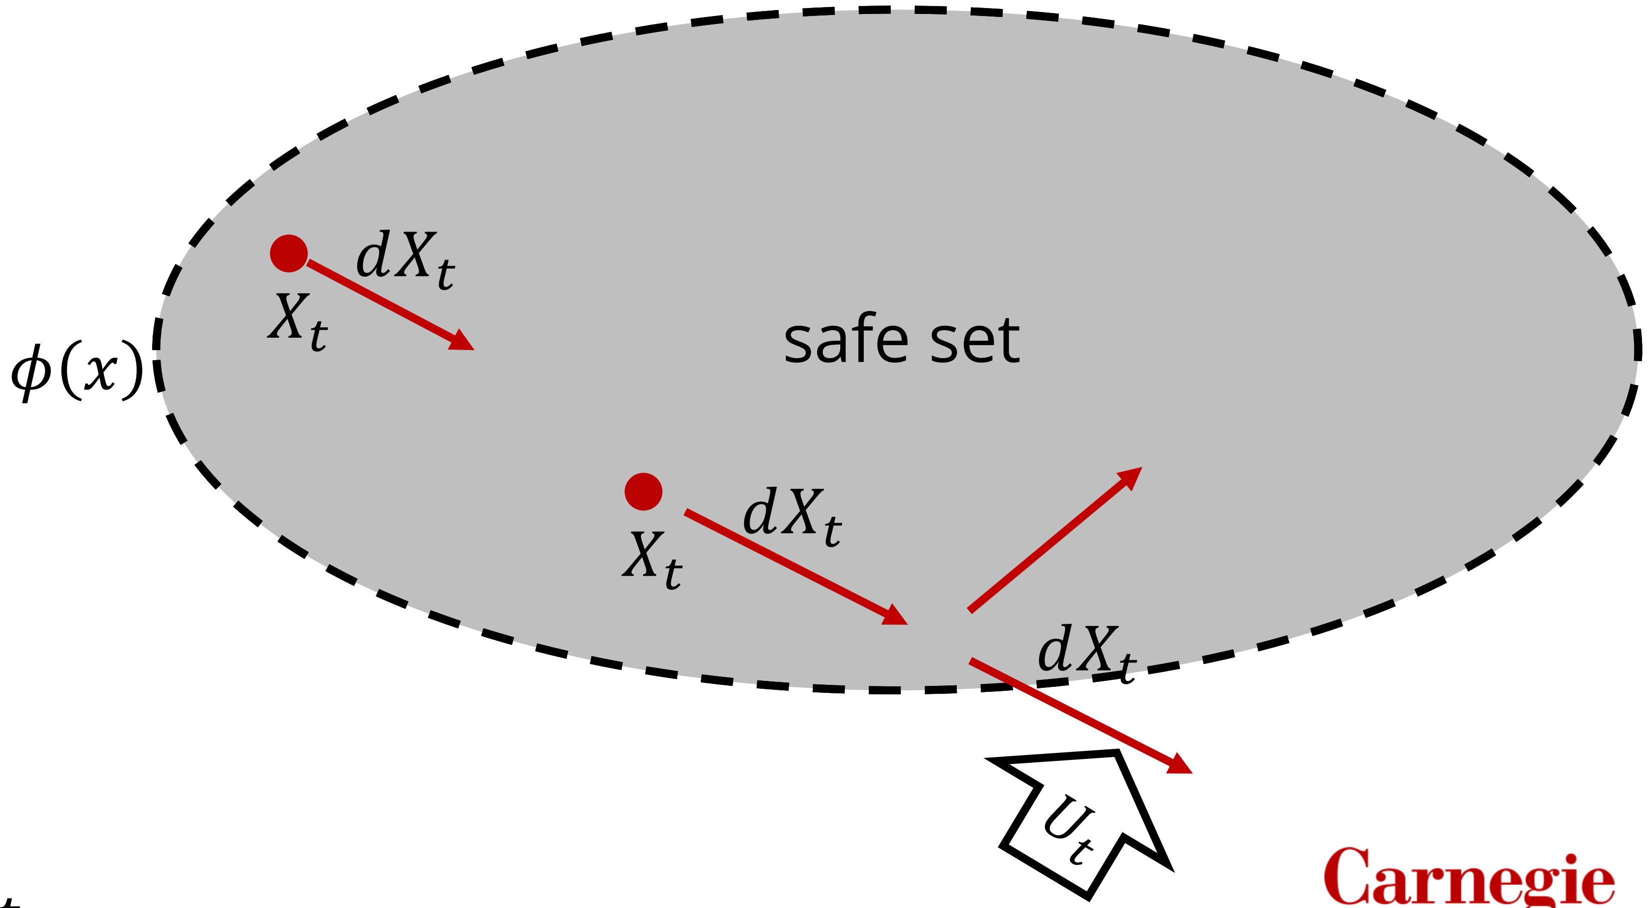

$$dX_t = (f(X_t) + g(X_t)U_t)\, dt$$

Stochastic settings

safe at next time => safe at all time

safe with probability
$1 - \delta$ at each step

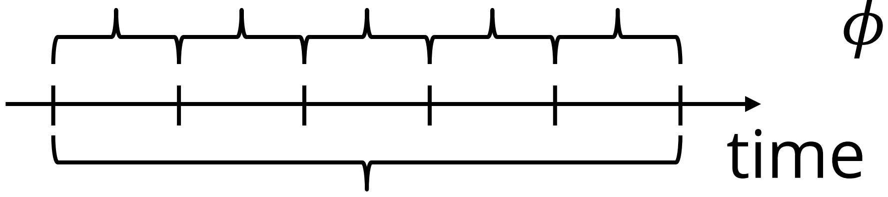

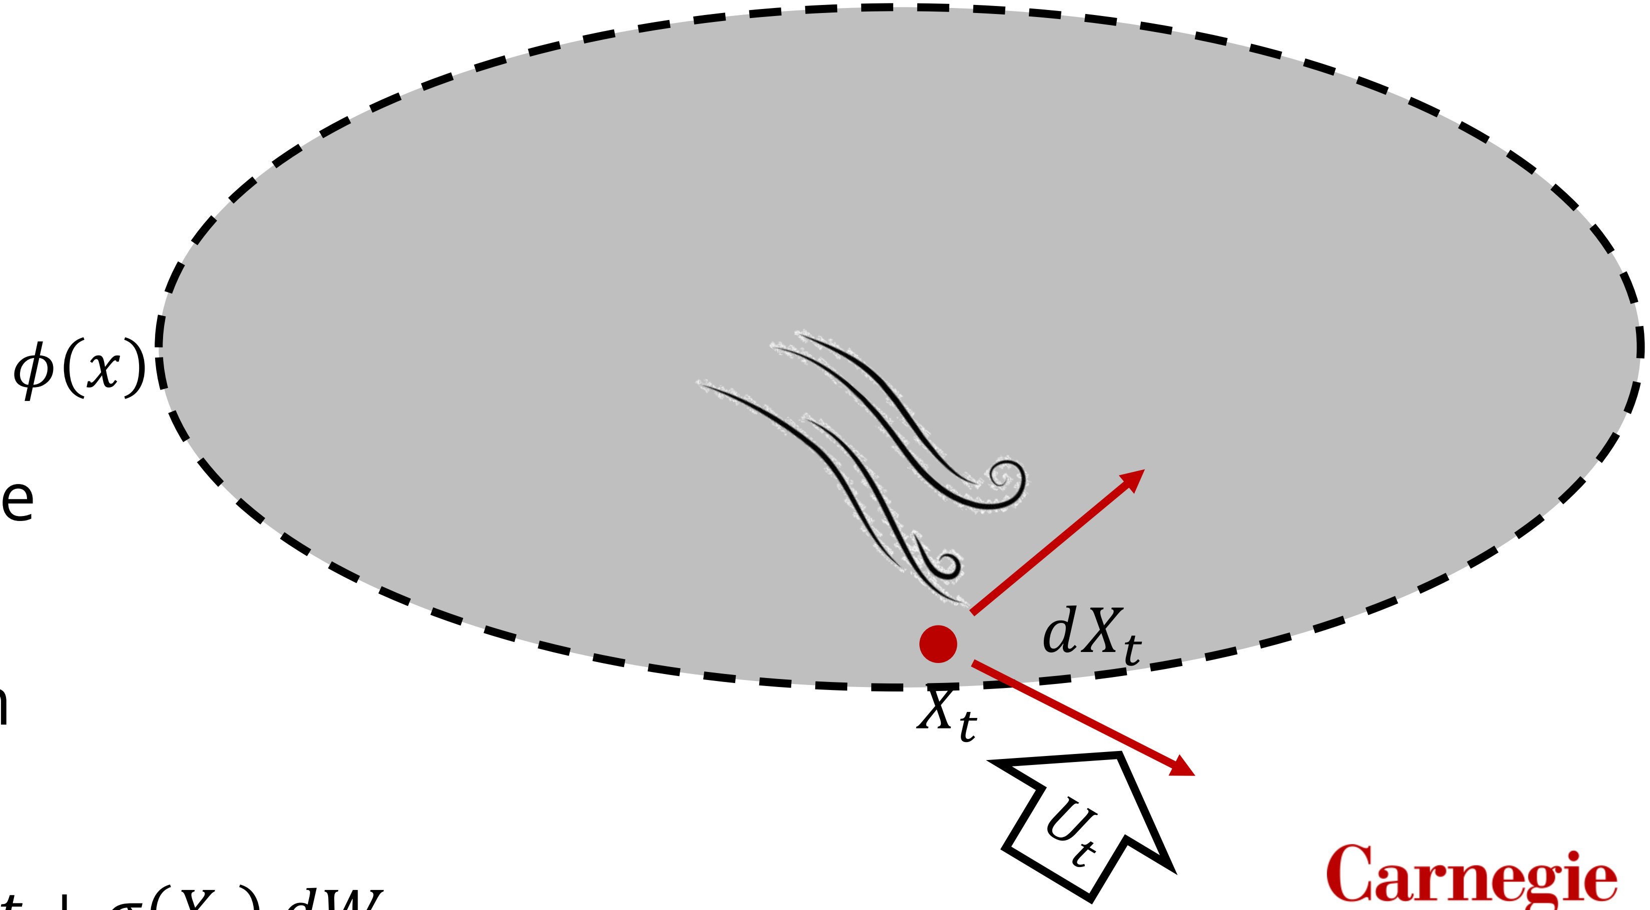

unsafe with high
probability in a long term

$$dX_t = (f(X_t) + g(X_t)U_t)\,dt + \sigma(X_t)\,dW_t$$

Carnegie Mellon University

Proposed method: intuitions

time derivative of safe probability

$$\frac{dF(X)}{dt} = \left(\frac{dF(X)}{dX}\right)\left(\frac{dX}{dt}\right) \geq 0$$

long-term guarantee

myopic evaluation

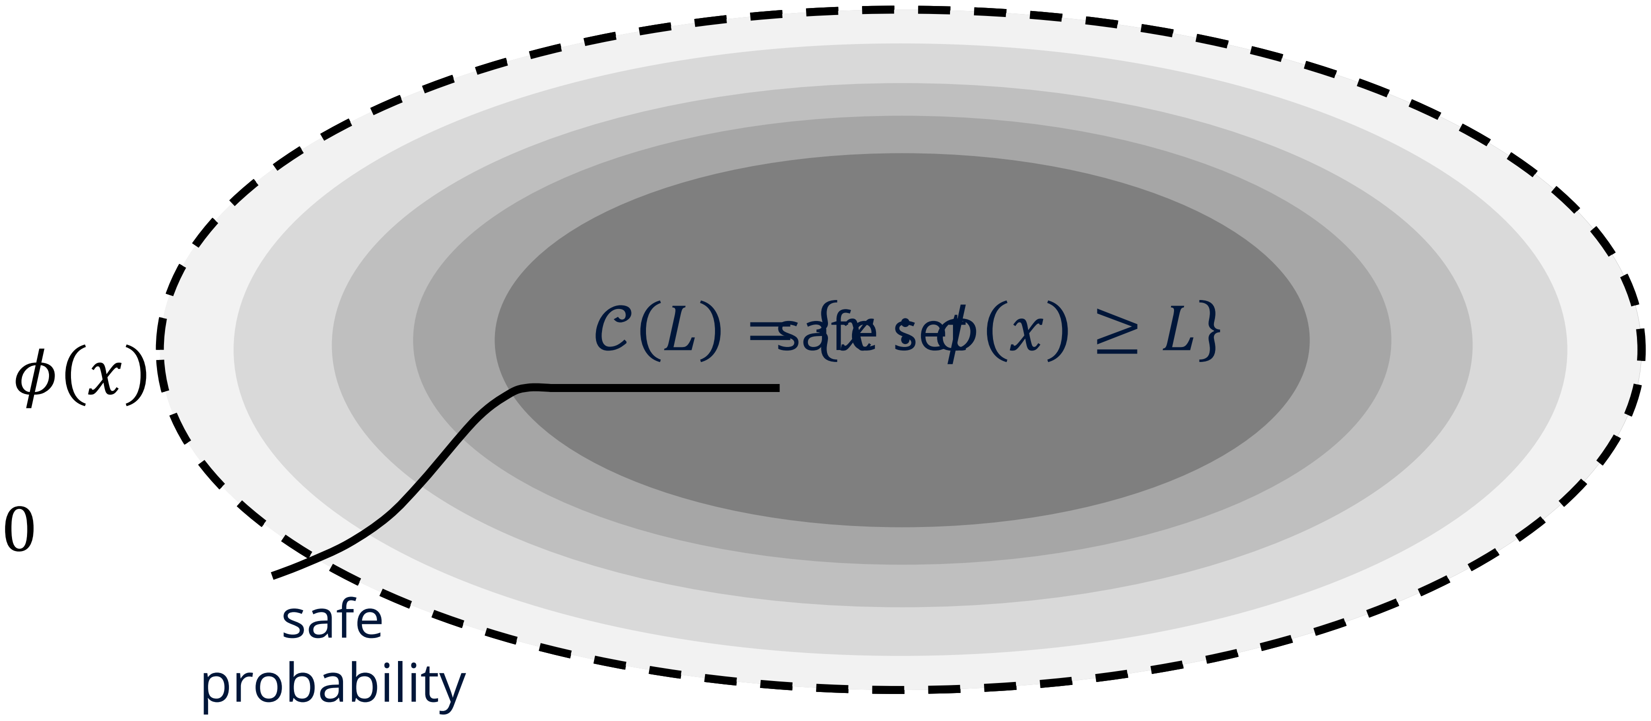

Proposed method: benefits

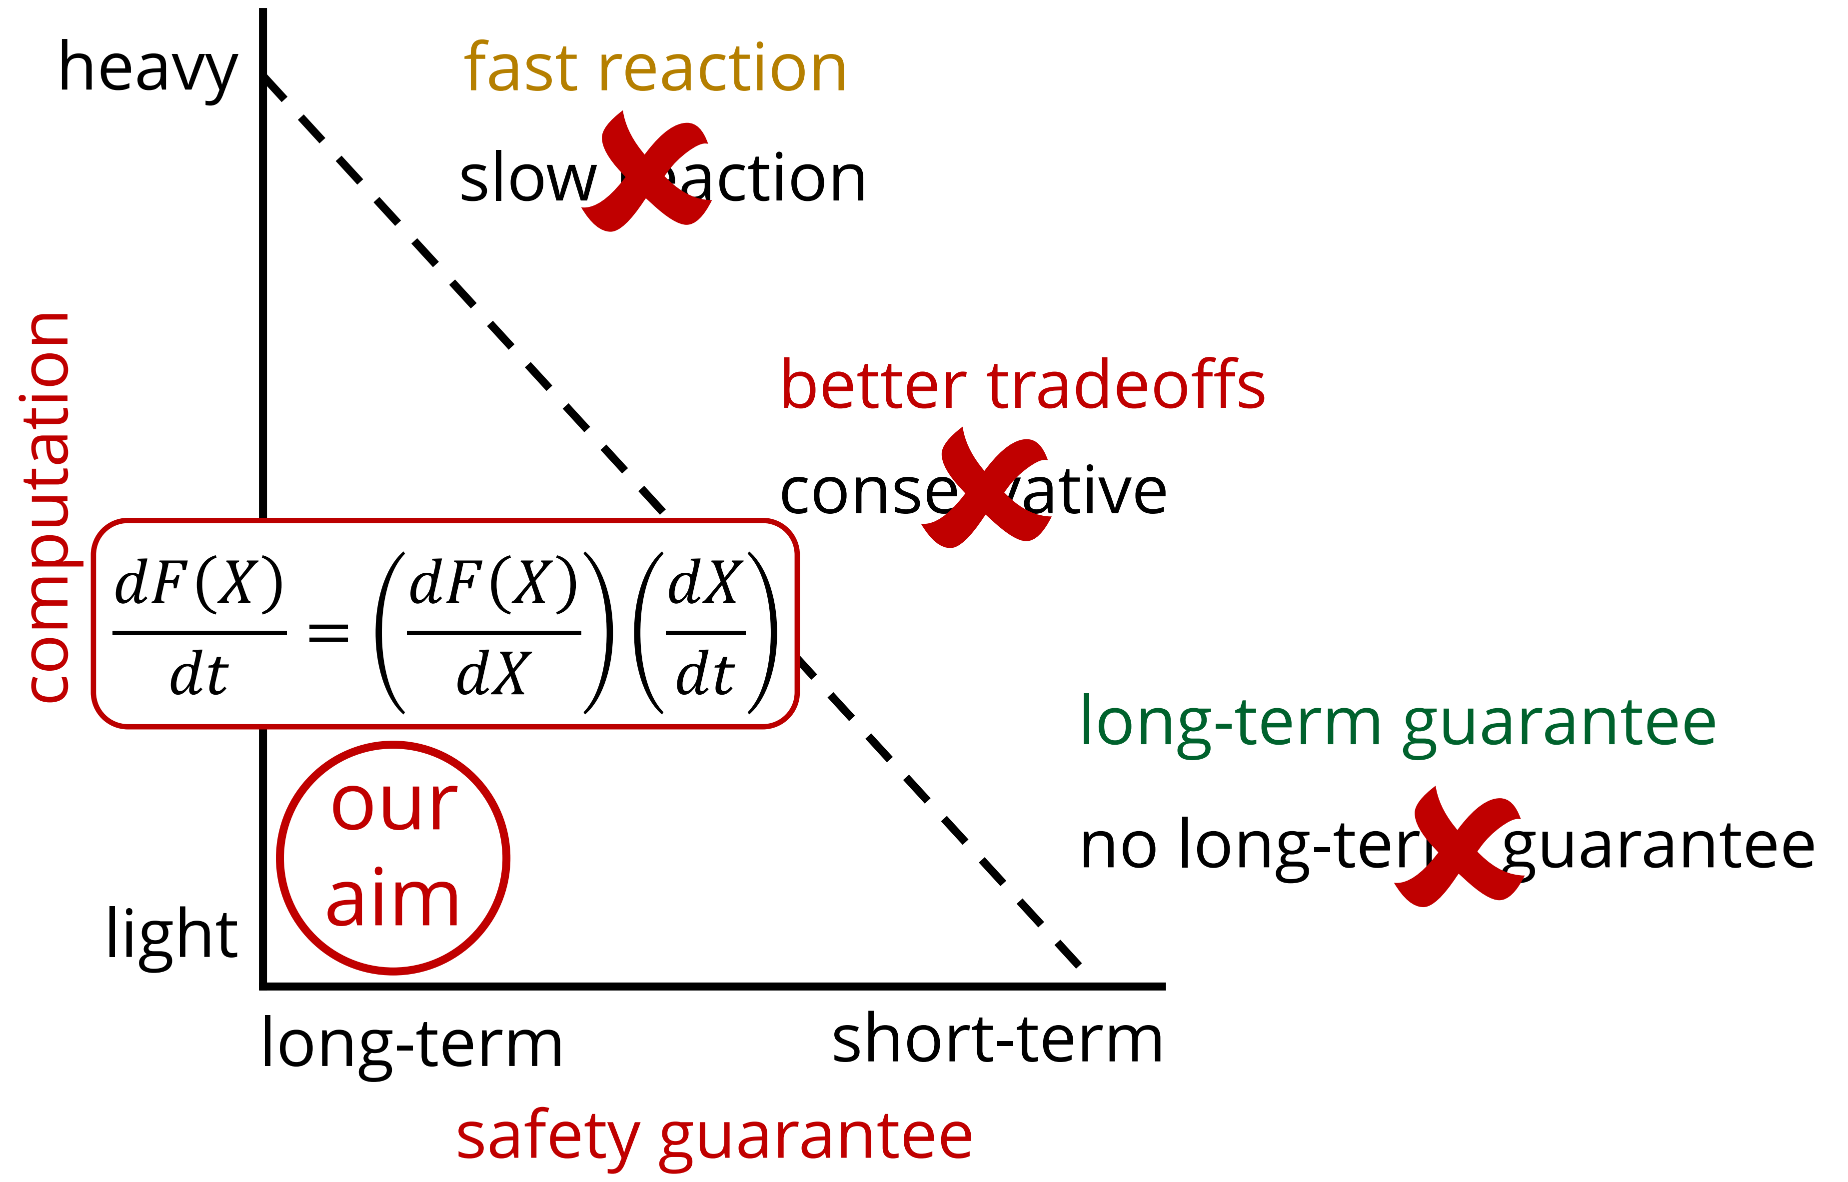

Proposed method: safety condition

The control action satisfies

$$D_{\boldsymbol{F}}(Z_t, U_t) \geq -\alpha(\boldsymbol{F}(Z_t) - (1-\epsilon))$$

where

$$D_{\boldsymbol{F}}(Z_t, U_t) \coloneqq A\boldsymbol{F}(Z_t)$$
$$= \mathcal{L}_{\tilde{f}}\boldsymbol{F}(Z_t) + \left(\mathcal{L}_{\tilde{g}}\boldsymbol{F}(Z_t)\right)U_t + \frac{1}{2}\mathrm{tr}([\tilde{\sigma}(Z_t)][\tilde{\sigma}(Z_t)]^{\intercal}\mathrm{Hess}\boldsymbol{F}(Z_t))$$

and $\alpha: \mathbb{R} \to \mathbb{R}$ is a monotonically increasing concave function that satisfies $\alpha(0) \leq 0$.

Proposed safe condition

Theorem [1]: If we choose the control action to satisfy

$$D_{\boldsymbol{F}}(Z_t, U_t) \geq -\alpha(\boldsymbol{F}(Z_t) - (1 - \epsilon)) \text{ for } t > 0,$$

then we have

$$\boldsymbol{F}(Z_0) > 1 - \epsilon \Rightarrow \mathbb{E}[\boldsymbol{F}(Z_t)] \geq 1 - \epsilon \text{ for } t > 0$$

$\alpha: \mathbb{R} \to \mathbb{R}$ is a monotonically increasing concave function that satisfies $\alpha(0) \leq 0$.

$\mathbb{E}(\boldsymbol{F}(Z_t))$: the probability of staying within $C(L_t)$ during $[t, t + T_t]$

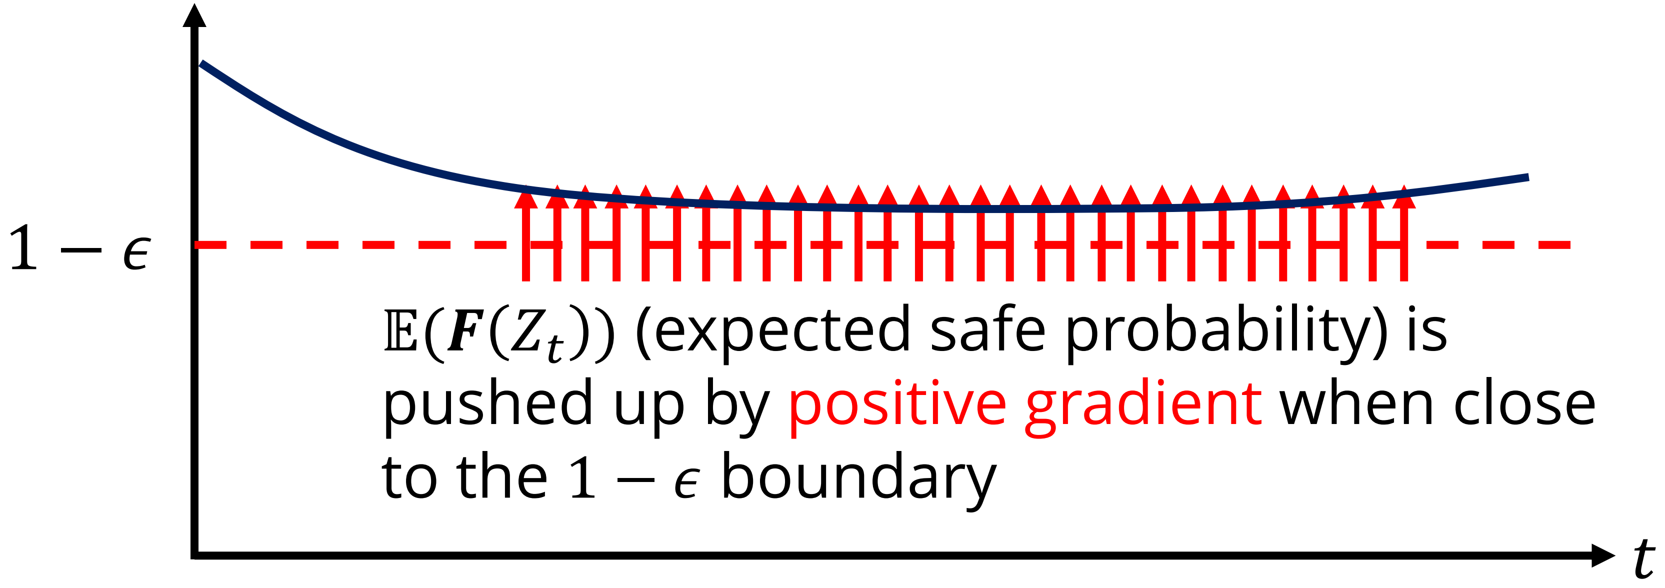

[1] Wang, Z., et al. "Myopically Verifiable Probabilistic Certificate for Long-term Safety." *arXiv preprint arXiv:2110.13380* (2021).

Proposed safe condition

Theorem [1]: If we choose the control action to satisfy

$$D_{\boldsymbol{F}}(Z_t, U_t) \geq -\alpha(\boldsymbol{F}(Z_t) - (1 - \epsilon)) \text{ for } t > 0,$$

then we have

$$\boldsymbol{F}(Z_0) > 1 - \epsilon \Rightarrow \mathbb{E}[\boldsymbol{F}(Z_t)] \geq 1 - \epsilon \text{ for } t > 0$$

$\alpha: \mathbb{R} \to \mathbb{R}$ is a monotonically increasing concave function that satisfies $\alpha(0) \leq 0$.

$\mathbb{E}(\boldsymbol{F}(Z_t))$: the probability of recovering to $C(L_t)$ within $[t, t + T_t]$

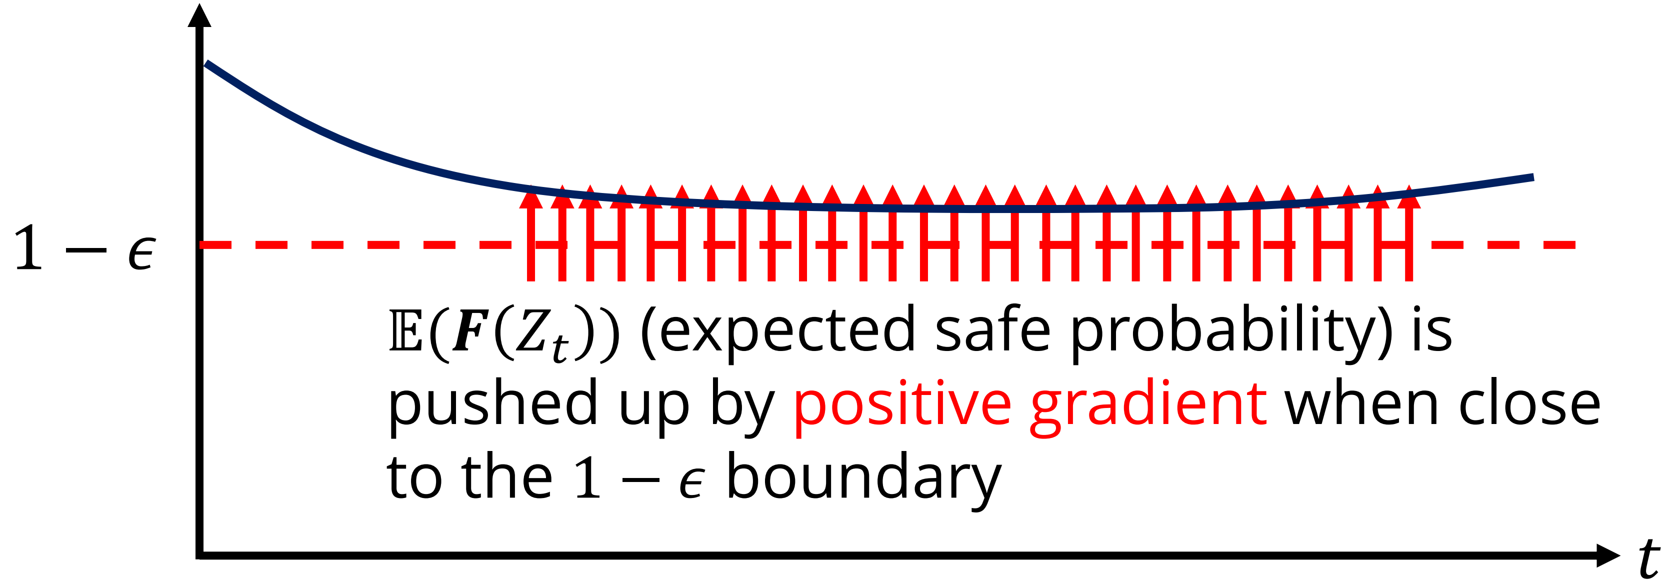

[1] Wang, Z., et al. "Myopically Verifiable Probabilistic Certificate for Long-term Safety." *arXiv preprint arXiv:2110.13380* (2021).

Proposed safe controllers

Assumption: nominal controller $U_t = N(X_t)$ ensures desired performance without considering safety.

Approach 1: additive modification

$$U_t = N(X_t) + \underbrace{\kappa(Z_t)}_{\text{Nonnegative function that pushes } D_{\boldsymbol{F}}(Z_t, U_t) \text{ to satisfy the proposed control policy}} \left(\mathcal{L}_{\tilde{g}}\boldsymbol{F}(Z_t)\right)^{\top}$$

Approach 2: conditioning

$$U_t = \arg\min_{u} \underbrace{J(N(X_t), u)}_{\text{Objective function that penalizes derivation from desired performance}} \text{ such that } A\boldsymbol{F}(Z_t) \geq -\alpha(\boldsymbol{F}(Z_t) - (1-\epsilon))$$

Simulation

system dynamic:

$$dX_t = (f(X_t) + g(X_t)U_t)\,dt + \sigma(X_t)\,dW_t$$

where $X \in \mathbb{R}$, $f(X) \equiv 2, g(X) \equiv 2.5, \sigma(t) \equiv 2$, dW is a standard Weiner process with 0 initial value.

safe set:

$$\mathcal{C}(0) = \{x \in \ \mathbb{R}^n : \phi(x) > 0\} = \{x \in \ \mathbb{R}^n : x - 1 > 0\}$$

The initial state is $x_0 = 3$. We consider fixed time horizon setting with $H = 1s$.

nominal controller:

$$N(X_t) = KX_t$$

with $K = 2.5$. This means the nominal controller tends to drag the system state to unsafe regions.

Carnegie Mellon University

Simulation

Worst case control

Impose safe controller all the time to exam the safety enforcement power of different safety constraints.

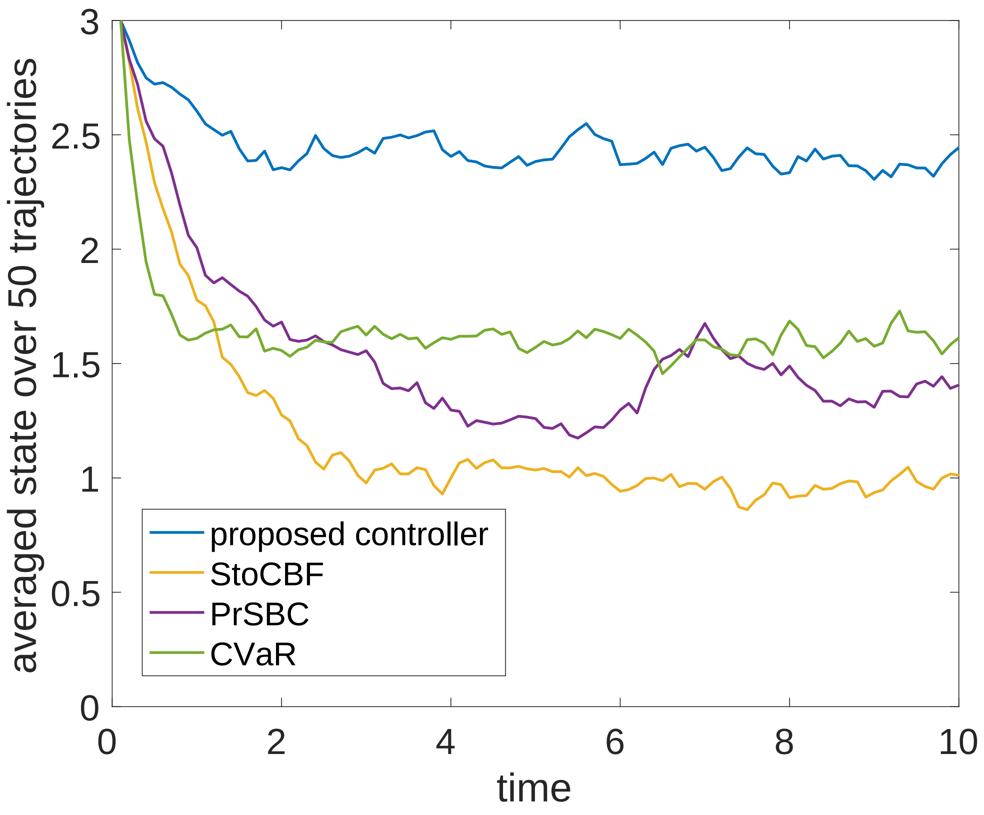

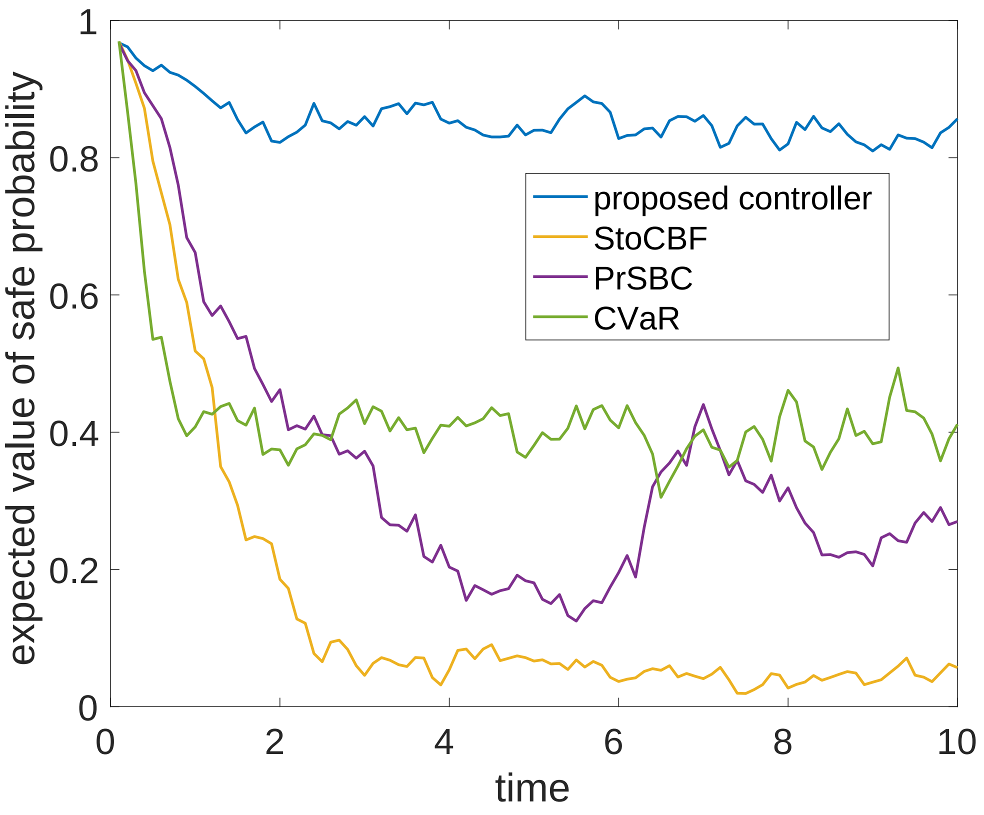

Carnegie Mellon University

Simulation

Switching control

Impose safe controller only when the nominal controller does not satisfy the safety constraint, to test the performance in a more practical setting.

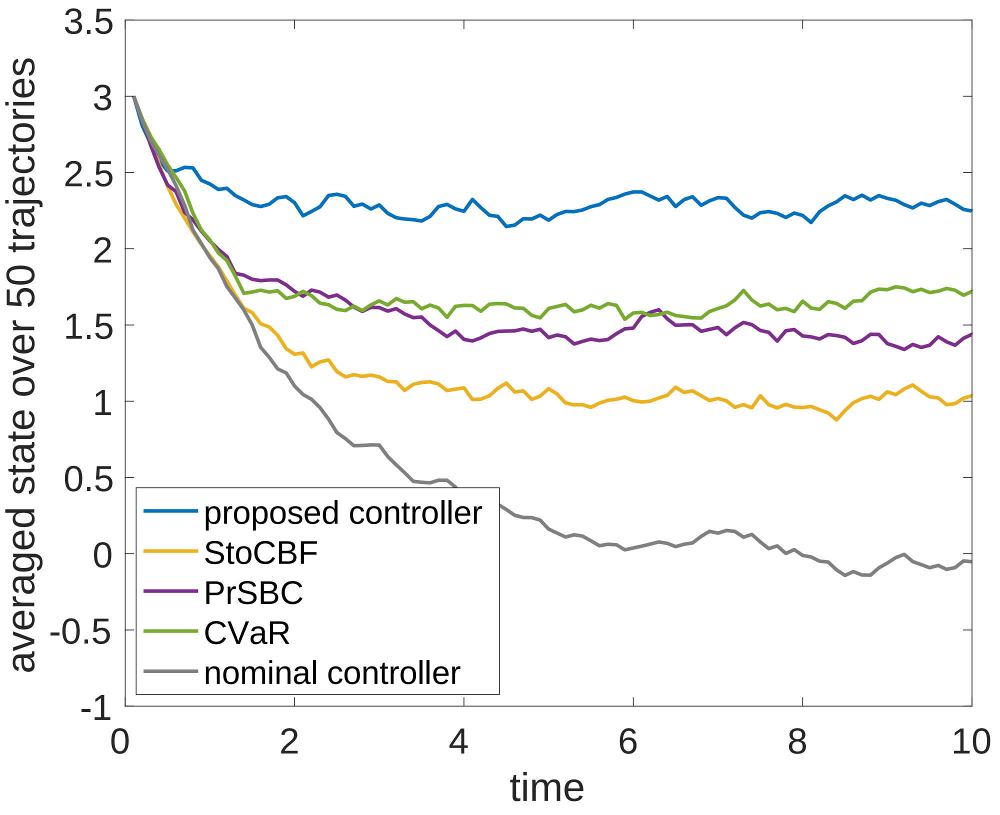

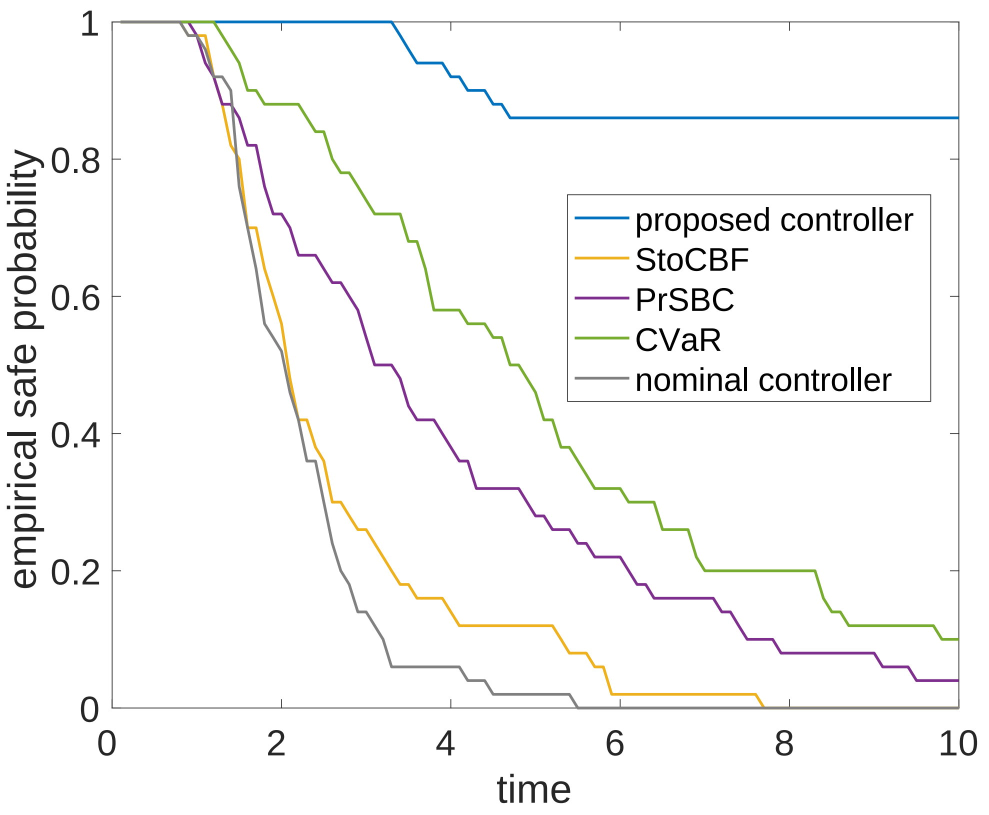

Carnegie Mellon University

Observations and insights

Myopic problems of CBF:

Example 1: violation of safety due to **nonlinear traps** in the system

car on slippery surface

To simulate similar behaviors, we add **nonlinear dynamics** to the system, when $X_t < 1.5$ the system becomes totally **uncontrollable**, with the new dynamics:

$$dX_t = f'(X_t)\, dt + \sigma(X_t)\, dW_t,$$

where $f'(X) \equiv -3$.
The safe set and the nominal controller are not changed.

Carnegie Mellon University

Simulation

Worst case control

Impose safe controller all the time to exam the safety enforcement power of different safety constraints.

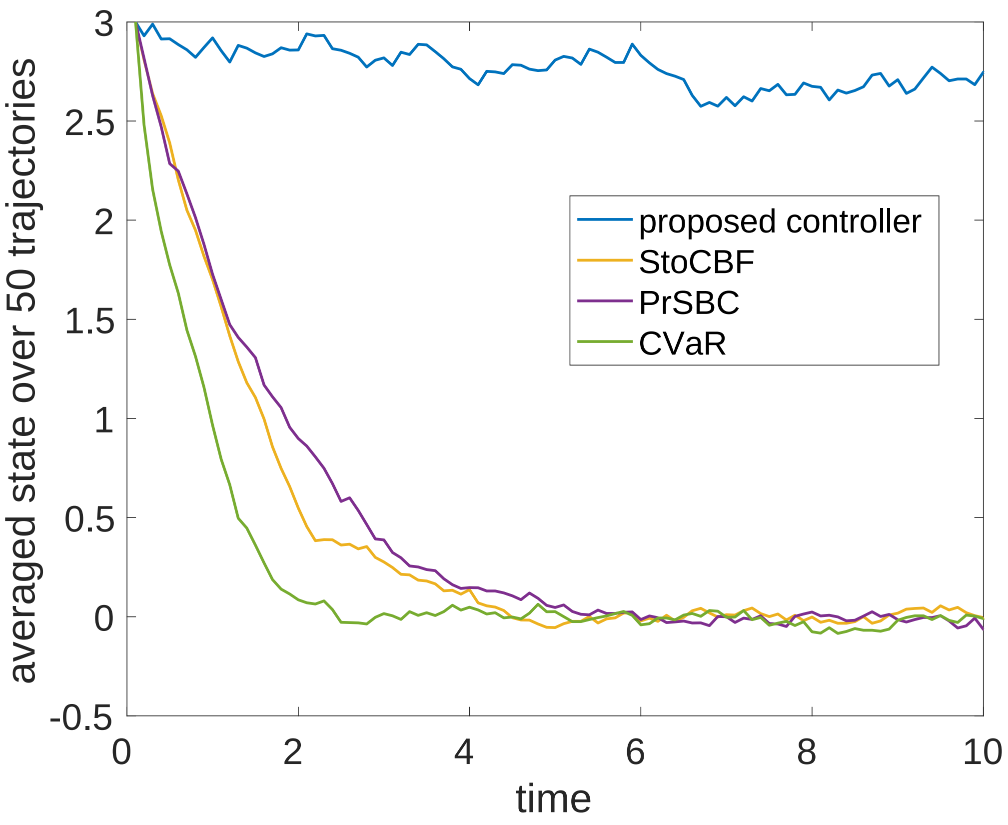

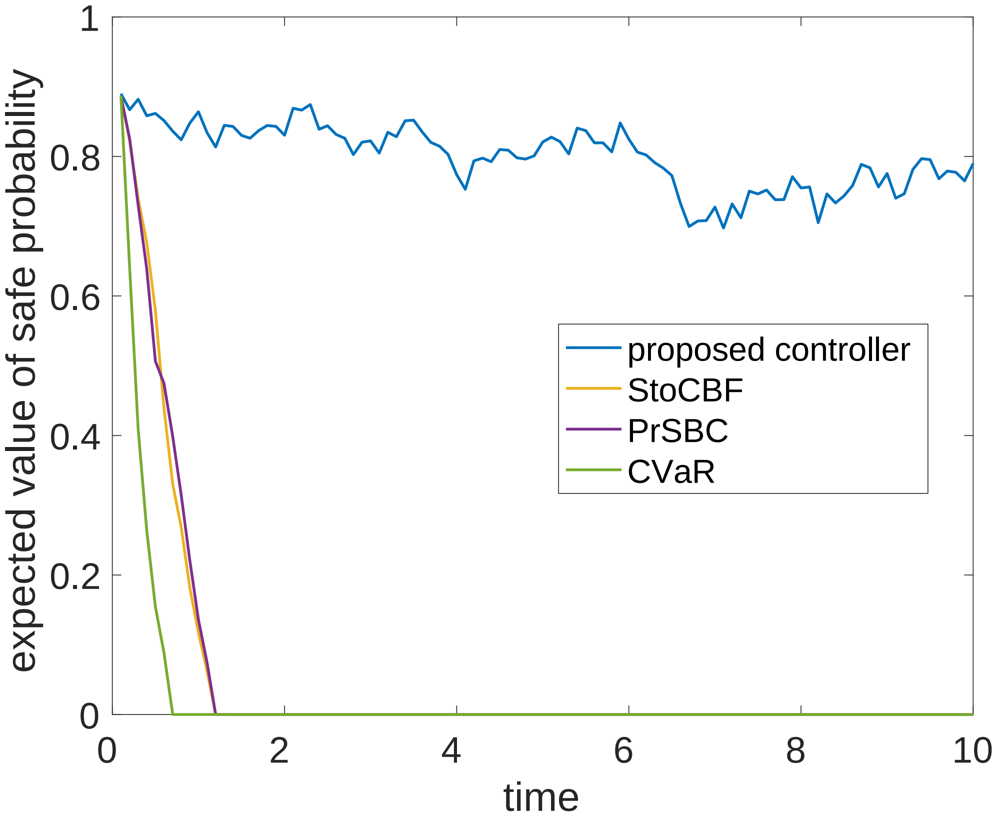

Simulation

Switching control

Impose safe controller only when the nominal controller does not satisfy the safety constraint, to test the performance in a more practical setting.

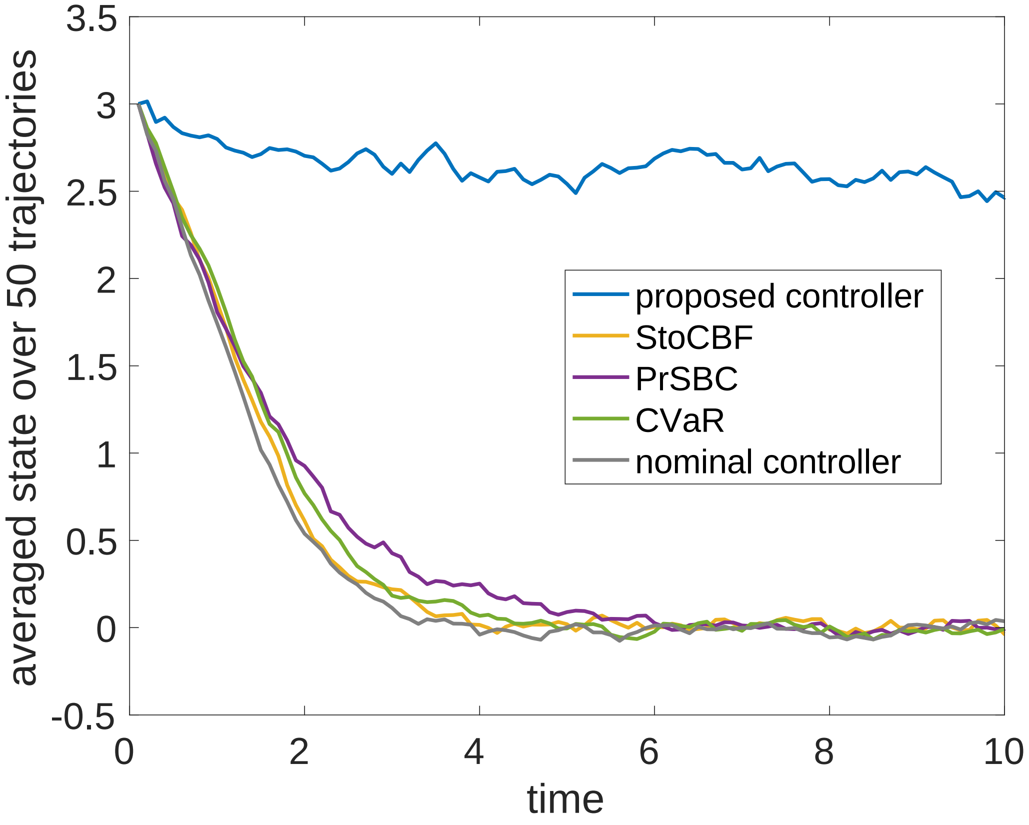

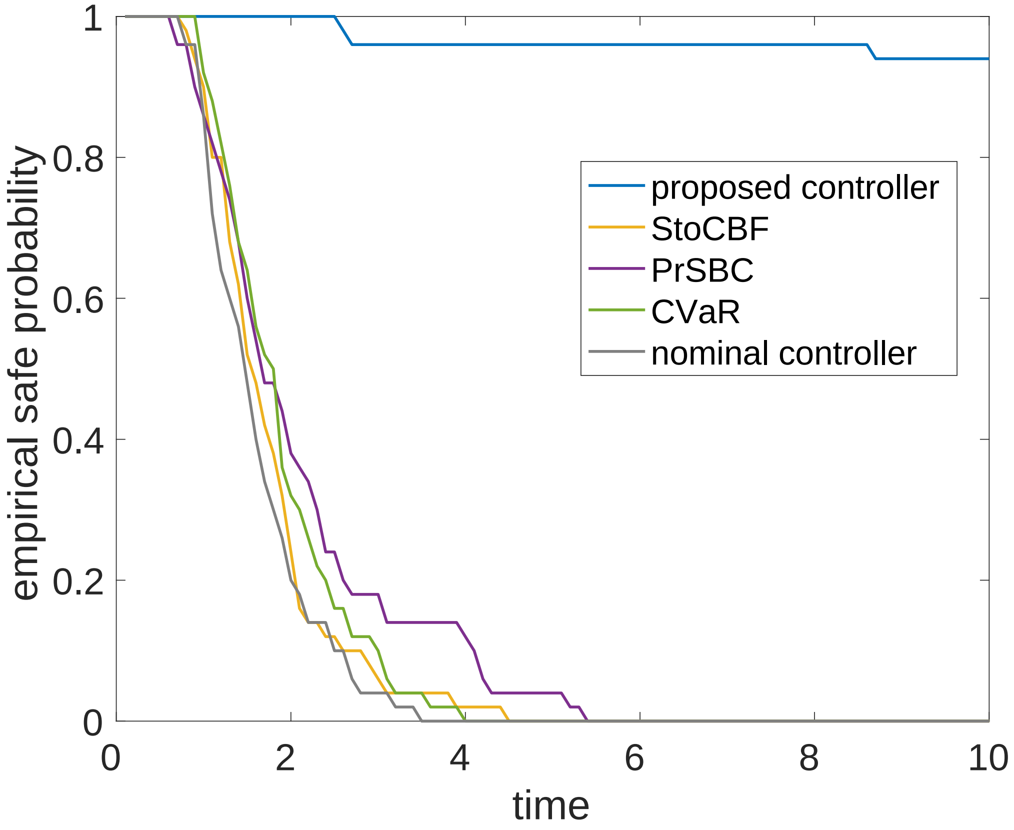

Carnegie Mellon University

Observations and insights

Myopic problems of CBF:

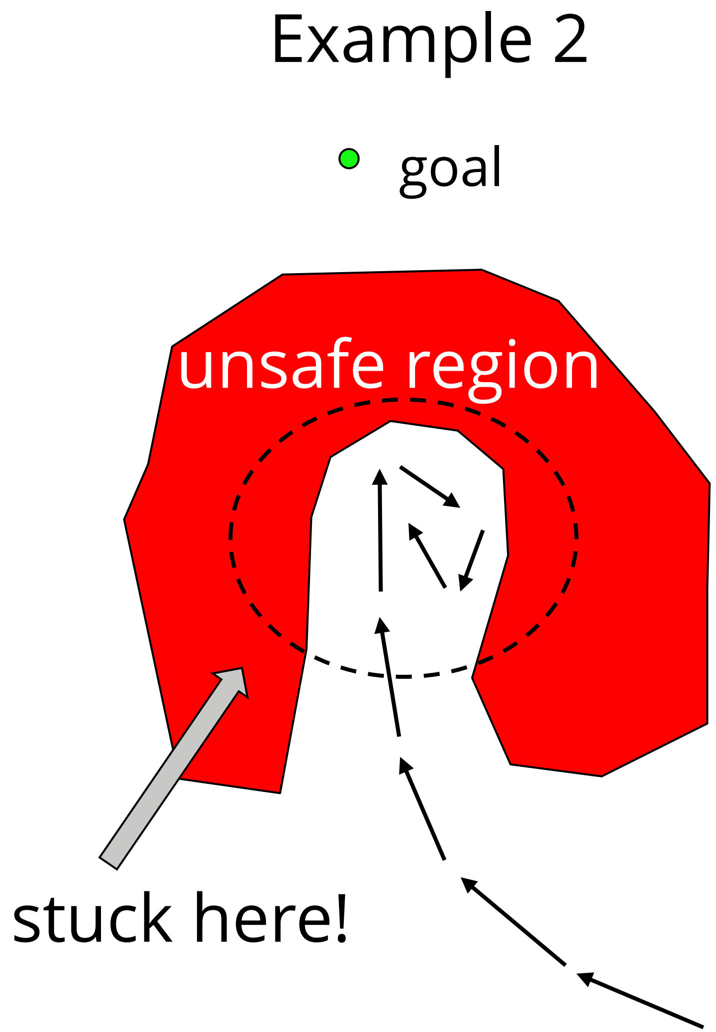

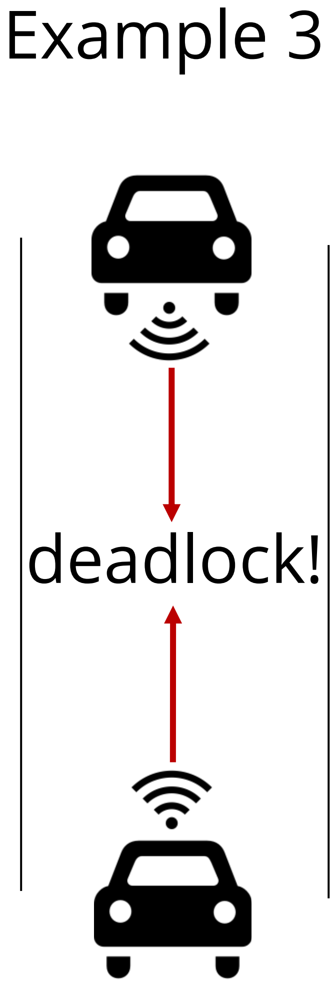

Carnegie Mellon University

Simulation

Traction control of 4-wheel vehicle

Performance goal: track a reference trajectory

Safety requirement: vehicle's tires do not slip

i.e., the total force of each tire do not exceed a certain percentage η of the maximum tire grip force.

Vehicle model: four-wheel 3-DoF vehicle [2]

Tire model: Burckhardt's tire model [3]

[2] Isaksson Palmqvist, Mia. "Model predictive control for autonomous driving of a truck." (2016).
[3] Kiencke, Uwe, and Lars Nielsen. "Automotive control systems: for engine, driveline, and vehicle." (2000): 1828.

Carnegie Mellon University

Simulation

Traction control of 4-wheel vehicle

B-spline planner and a Linear Time-Varying MPC (LTV-MPC) [4] as the baseline nominal controller with steering limits and lane constraints

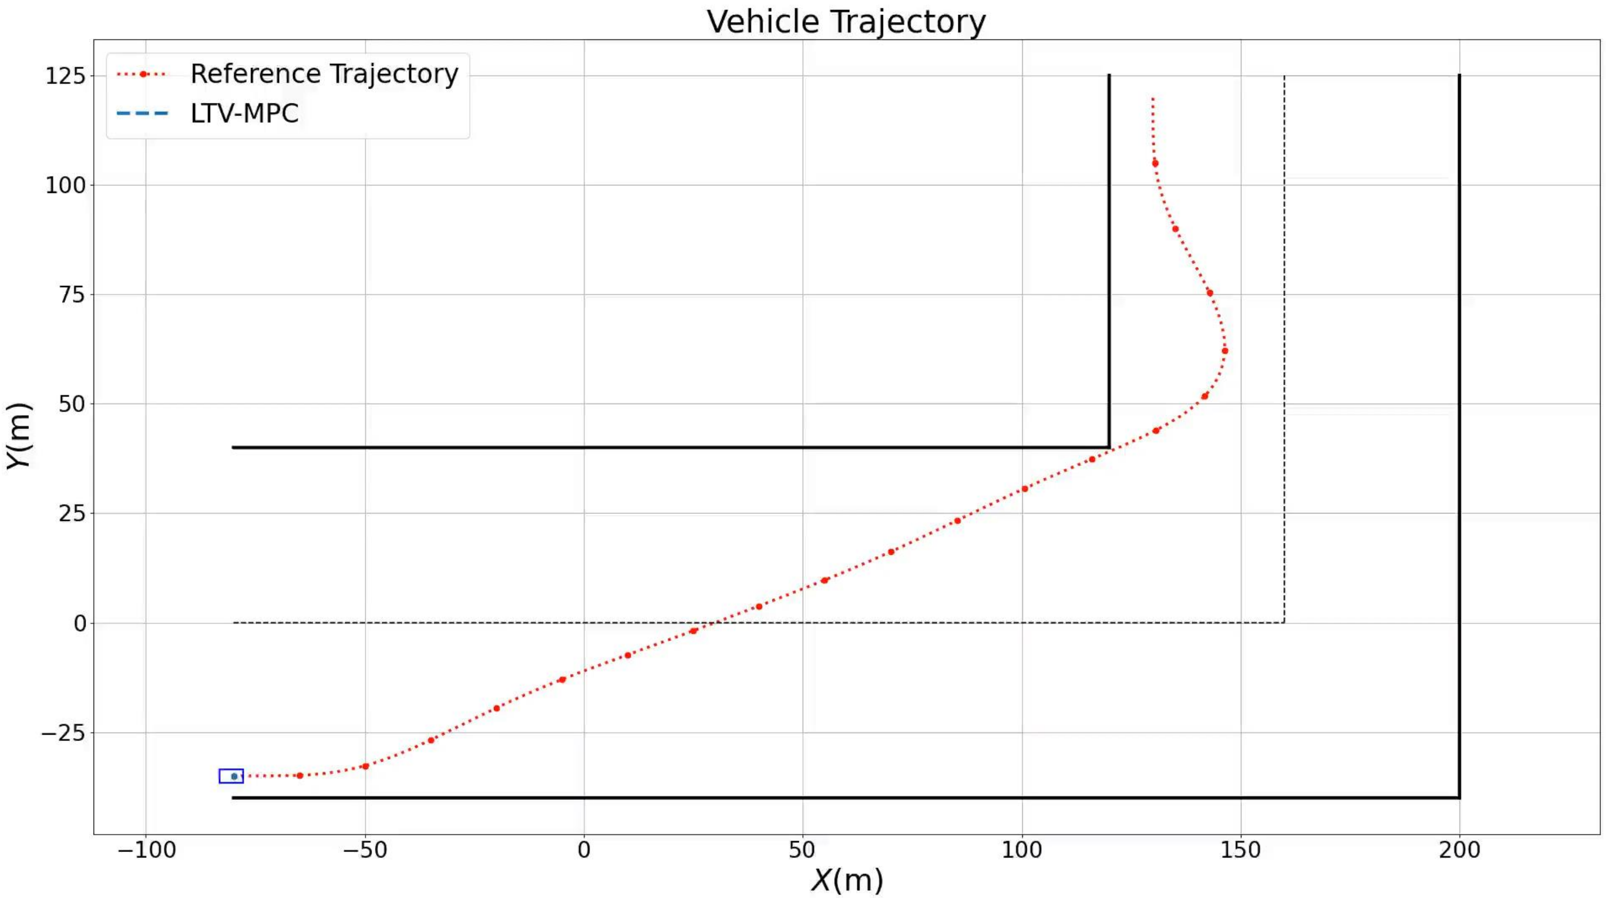

[4] Falcone, Paolo, et al. "A linear time varying model predictive control approach to the integrated vehicle dynamics control problem in autonomous systems." 2007 46th IEEE Conference on Decision and Control. IEEE, 2007.

Carnegie Mellon University

Simulation

Traction control of 4-wheel vehicle

Proposed controller [5]

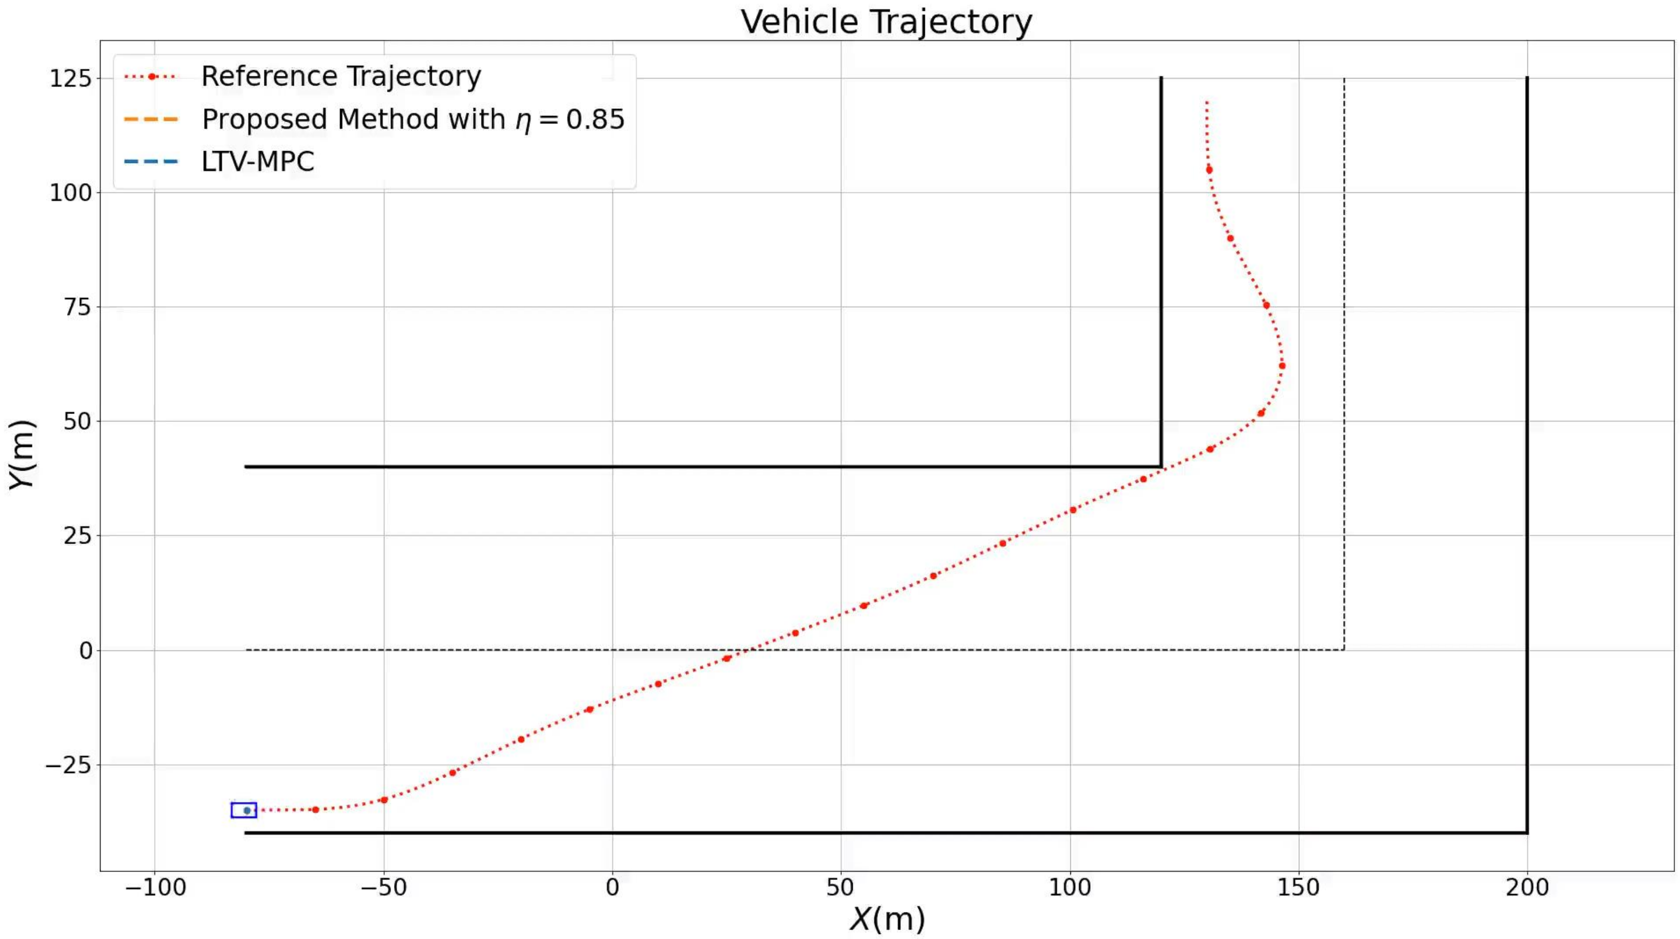

[5] Gangadhar, Siddharth, et al. "Dealing with Stochastic Uncertainty and Prediction in Extreme Driving." https://github.com/haomingj/Dealing-with-Stochastic-Uncertainty-and-Prediction-in-Extreme-Driving.

Simulation

Traction control of 4-wheel vehicle

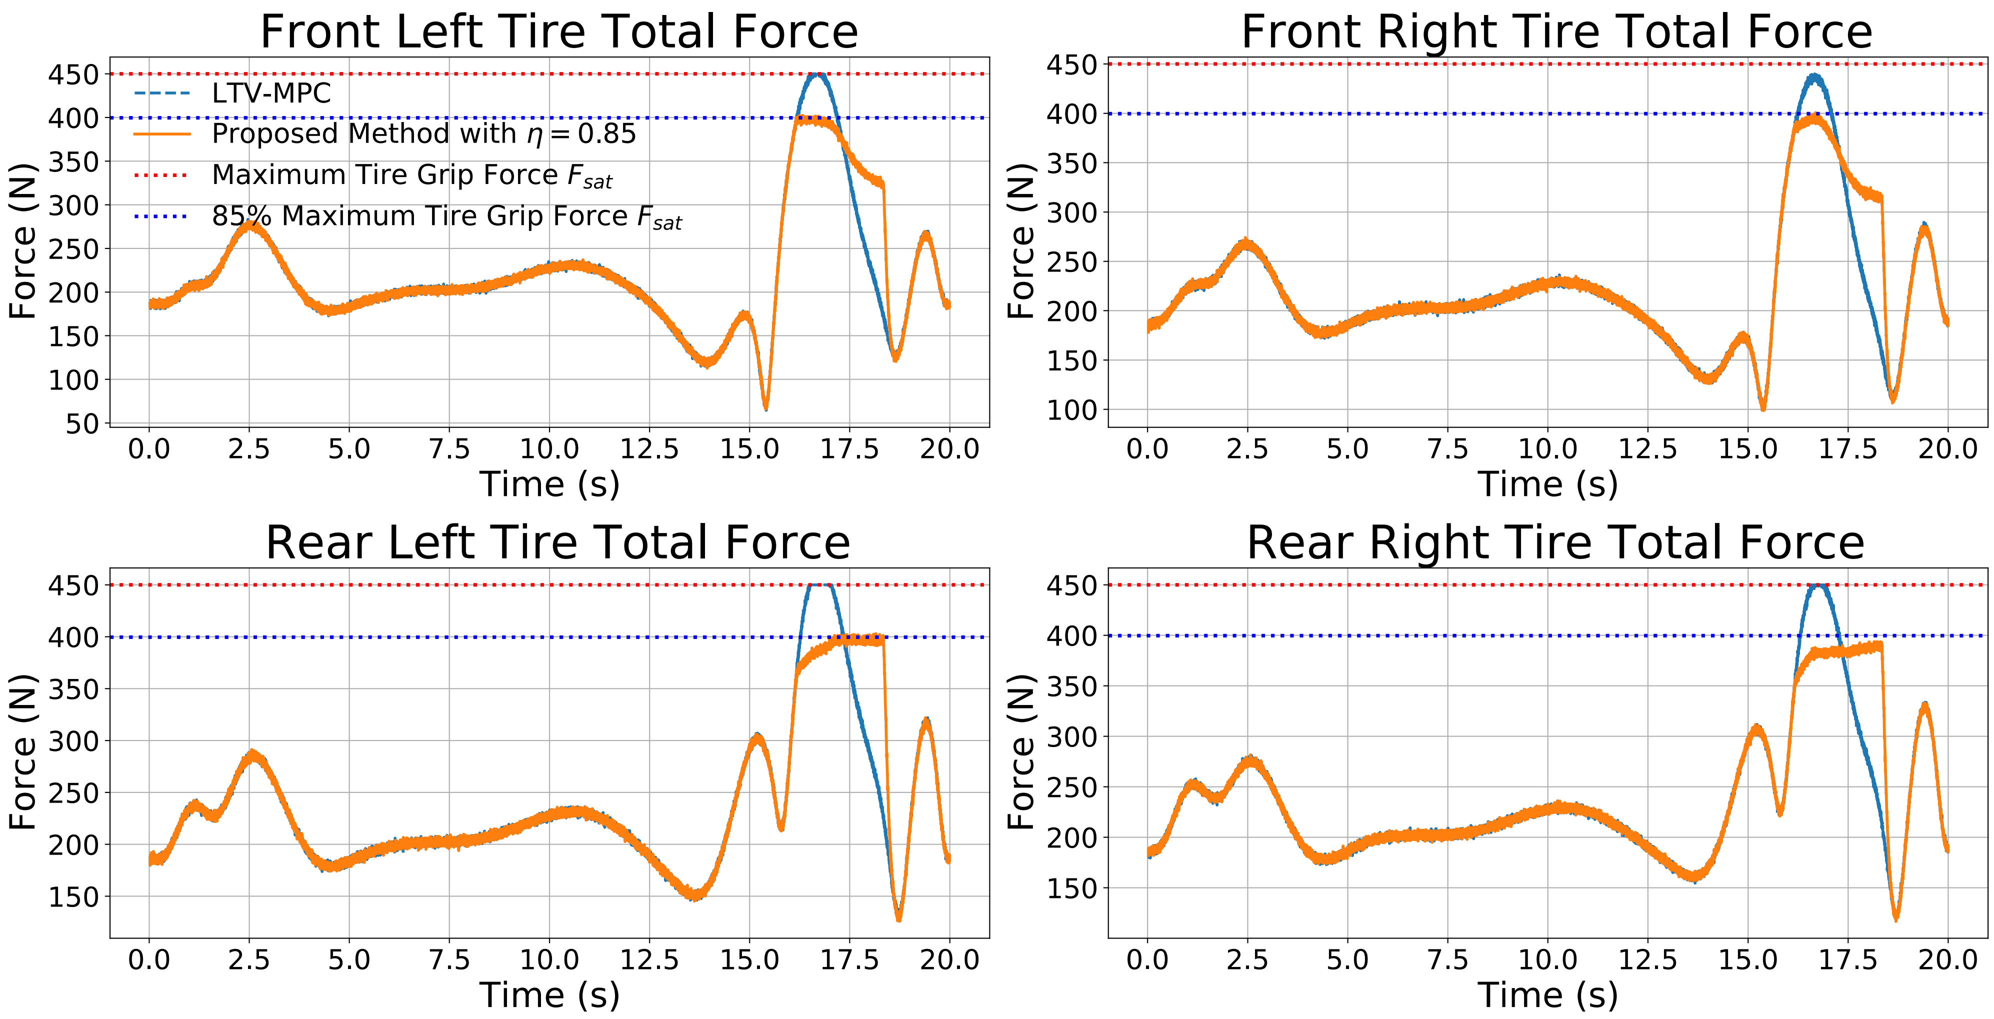

Carnegie Mellon University

Simulation

Traction control of 4-wheel vehicle

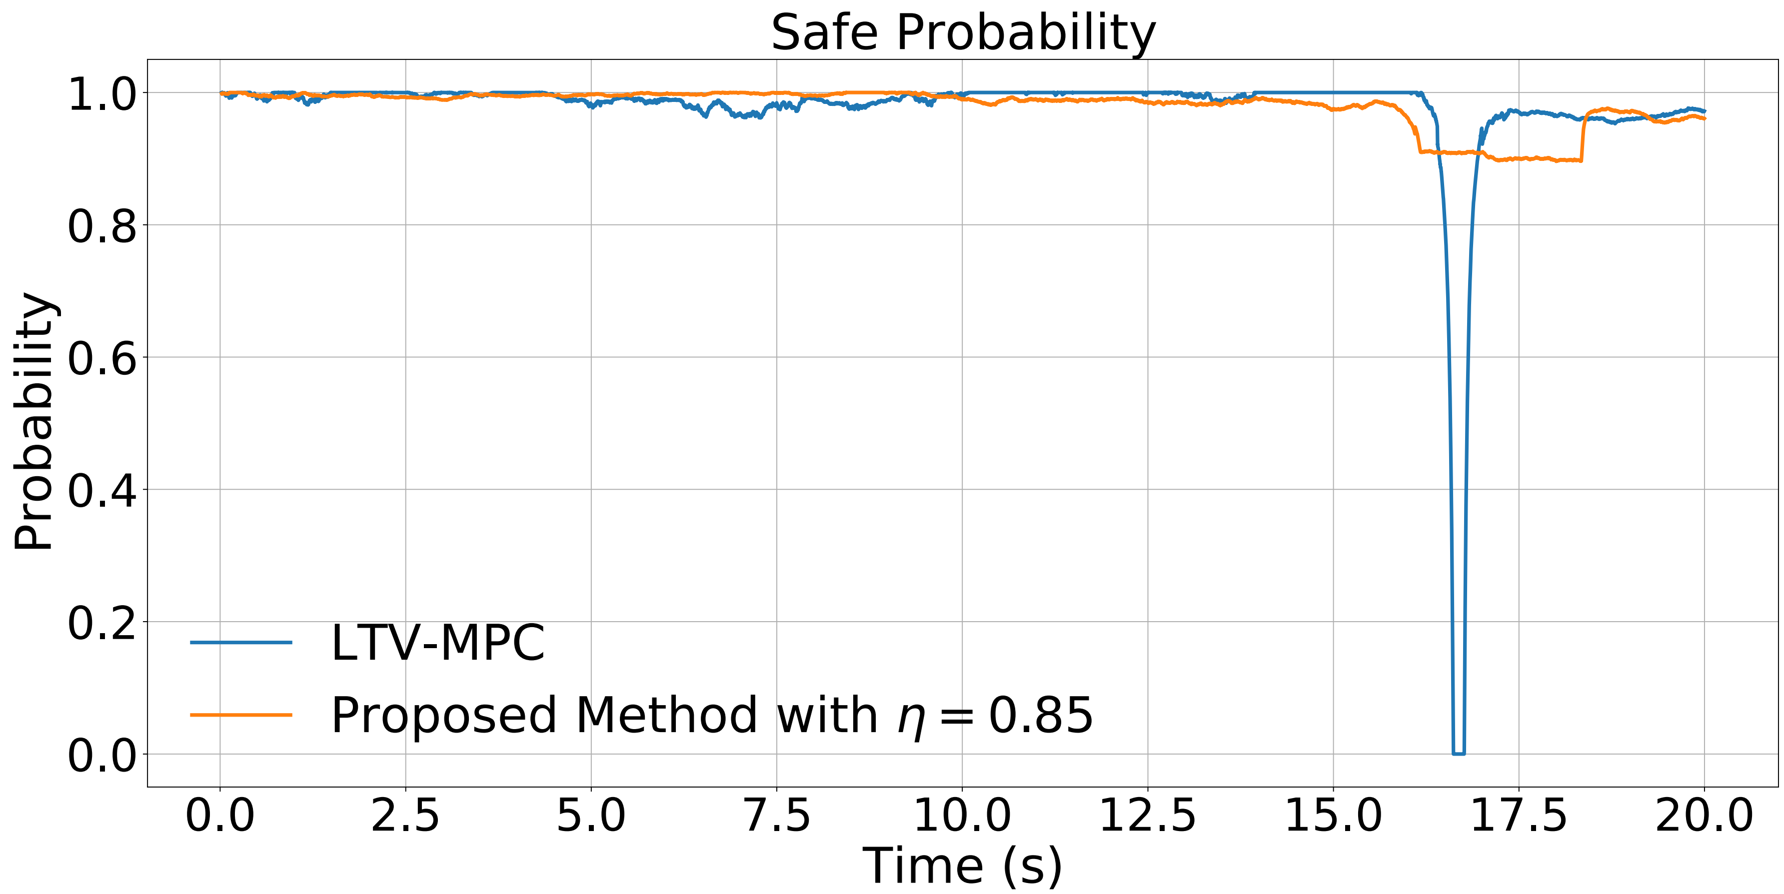

Conclusion

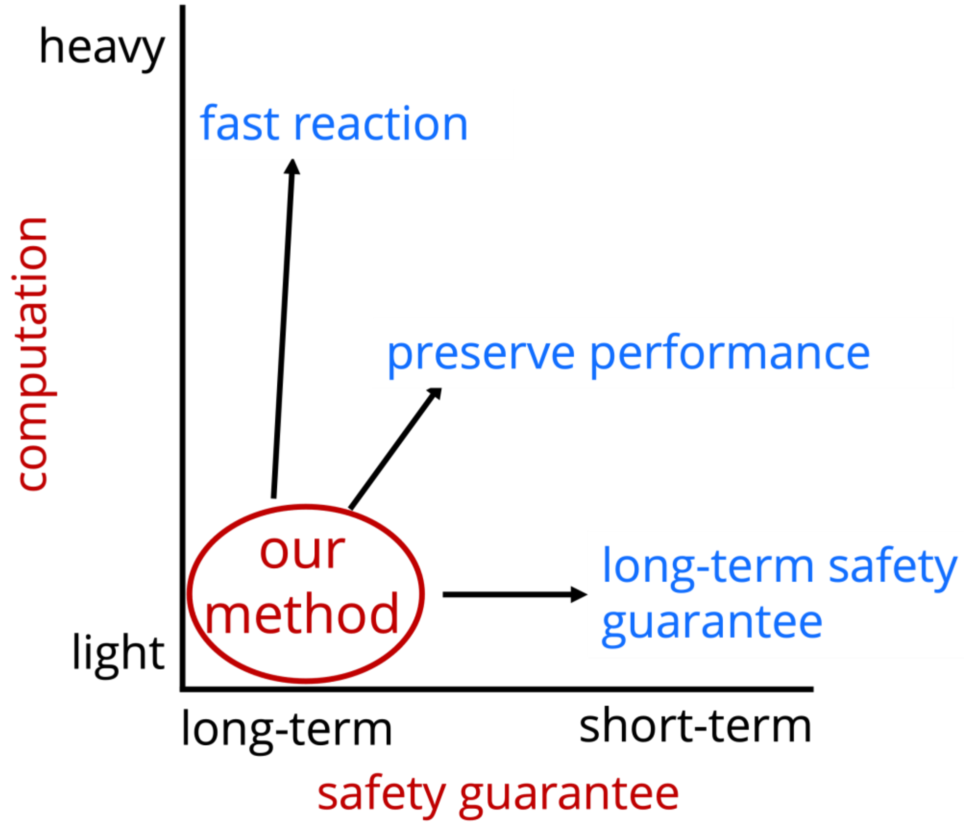

- **Provable long-term safety guarantee**
- **Fast reaction with reduced computation**
- **Controllable safety and performance trade-off**
- **Easy implementation with plug-in usage**

Carnegie Mellon University

Potential Future works

- Multi-agent and distributed version of the safe control strategy
- Online learning of safe probability with high-dimensional state space
- Online adaption of barrier function for out-of-distribution data
- Gradient-based methods for safe control in RL framework
- ...

Acknowledgement

Yorie Nakahira

Haoming Jing

Christian Kurniawan

Siddharth Gangadhar

Carnegie Mellon University

Thanks for listening!

Reference

[1] Wang, Zhuoyuan, et al. "Myopically Verifiable Probabilistic Certificate for Long-term Safety." *arXiv preprint arXiv:2110.13380* (2021).

[2] Isaksson Palmqvist, Mia. "Model predictive control for autonomous driving of a truck." (2016).

[3] Kiencke, Uwe, and Lars Nielsen. "Automotive control systems: for engine, driveline, and vehicle." (2000): 1828.

[4] Falcone, Paolo, et al. "A linear time varying model predictive control approach to the integrated vehicle dynamics control problem in autonomous systems." 2007 46th IEEE Conference on Decision and Control. IEEE, 2007.

[5] Gangadhar, Siddharth, et al. “Dealing with Stochastic Uncertainty and Prediction in Extreme Driving.” https://github.com/haomingj/Dealing-with-Stochastic-Uncertainty-and-Prediction-in-Extreme-Driving.

[6] Chern, Albert, et al. "Safe Control in the Presence of Stochastic Uncertainties." *arXiv preprint arXiv:2104.01259* (2021).

Carnegie Mellon University