Enforcing Safety under Uncertainty: Control Barrier Functions for Adaptive Flight Control Systems

Johannes Autenrieb, Peter A. Fisher, Anuradha M. Annaswamy

Control Barrier Functions in Aerospace: From Foundations to Real-World Applications, CEAS EuroGNC, 04 April 2026

Why do we need safe controllers?

Source: Ministry of National Defense (South Korea)

Source: Lockheed Martin

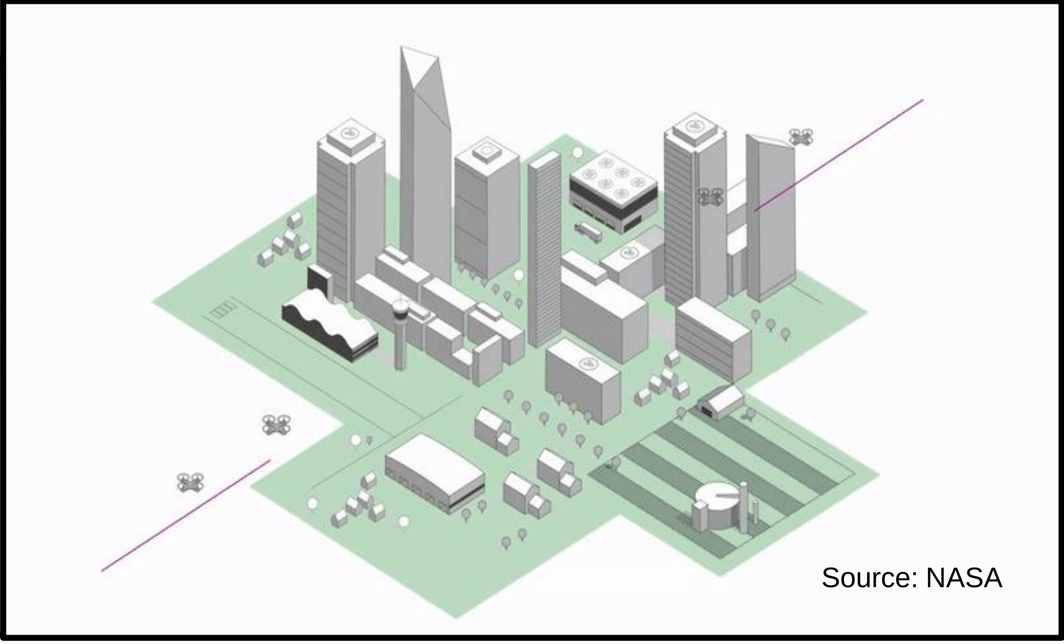

Source: NASA

Source: theworacle/Boeing Phantom Works

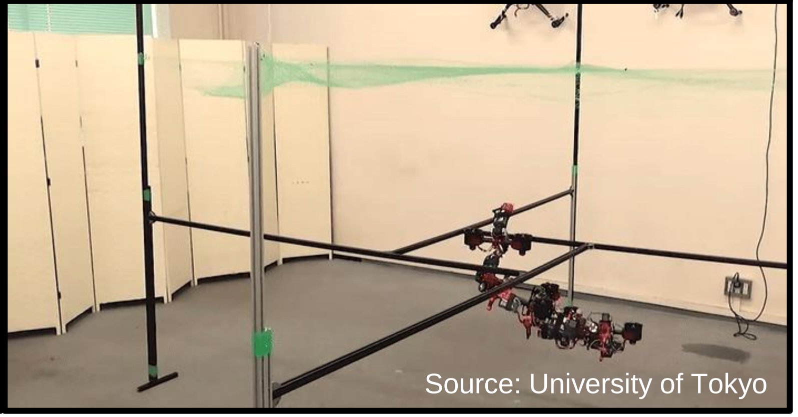

Source: University of Tokyo

Outline

1. Concept of Control Barrier Functions (CBFs)
2. Safety Filter Design for Adaptive Closed-Loop Systems
3. Outlook & Conclusions

Outline

1 Concept of Control Barrier Functions (CBFs)

2 Safety Filter Design for Adaptive Closed-Loop Systems

3 Outlook & Conclusions

Recap: How do we define safety?

We define a nonlinear continuous system:

$$\dot{x}(t) = f(x(t)) + g(x(t))u(t)$$

where $x(t) \epsilon \chi \subset \mathbb{R}^n$ and $u(t) \epsilon \mathbb{R}^m$.

We consider a continuously differentiable function $h: \mathbb{R}^n \rightarrow \mathbb{R}$ and and a closed set $S \epsilon \mathbb{R}^n$ defined as the zero-superlevel set of $h(x(t))$ defined as:

$$S \triangleq \{x(t) \epsilon \chi \mid h(x(t)) \geq 0\}$$

$$\partial S \triangleq \{x(t) \epsilon \chi \mid h(x(t)) = 0\}$$

$$int(S) \triangleq \{x(t) \epsilon \chi \mid h(x(t)) > 0\}$$

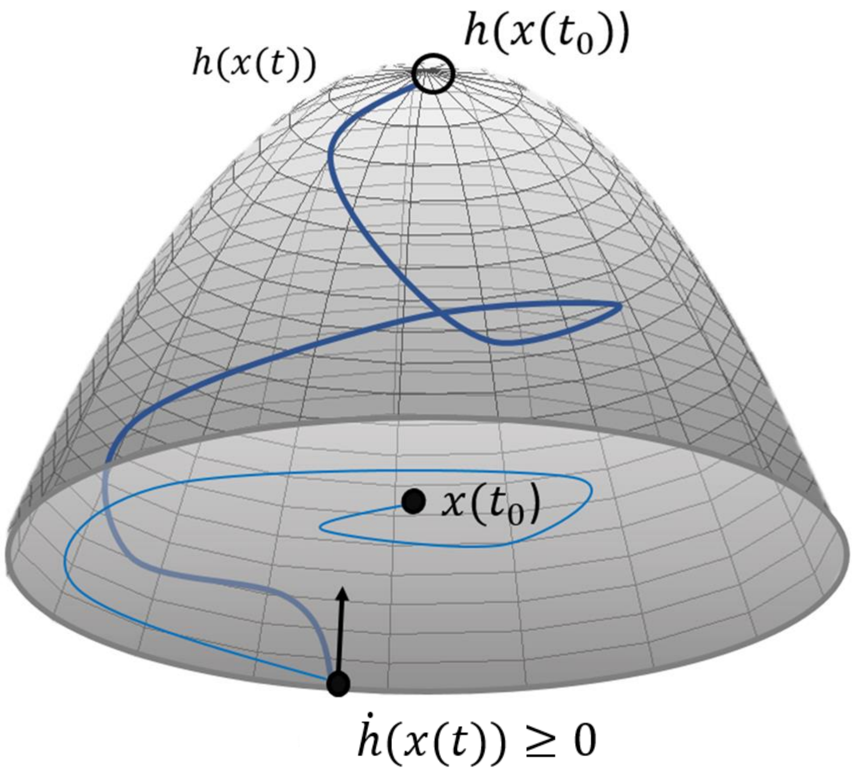

Recap: How can we render a closed-loop system safe?

Definition: Control Barrier Function (CBF)

The function $h(x(t))$ is a CBF for S, such that for the considered control affine system we obtain:

$$\sup_{u(t)\epsilon\,\mathbb{R}^m} \frac{\partial h}{\partial x}\left[f(x(t)) + g(x(t))u(t)\right] \geq -\alpha(h(x(t)))$$

for all $x(t)\epsilon\, S$ and $\alpha(h(x(t)))$ being a class K_∞ function.

<u>Popular approach</u>:

Exponential lower bound:

$$\alpha(h(x(t))) = \boxed{\gamma} h(x(t))$$

Parameter $\gamma > 0$ is a design choice.

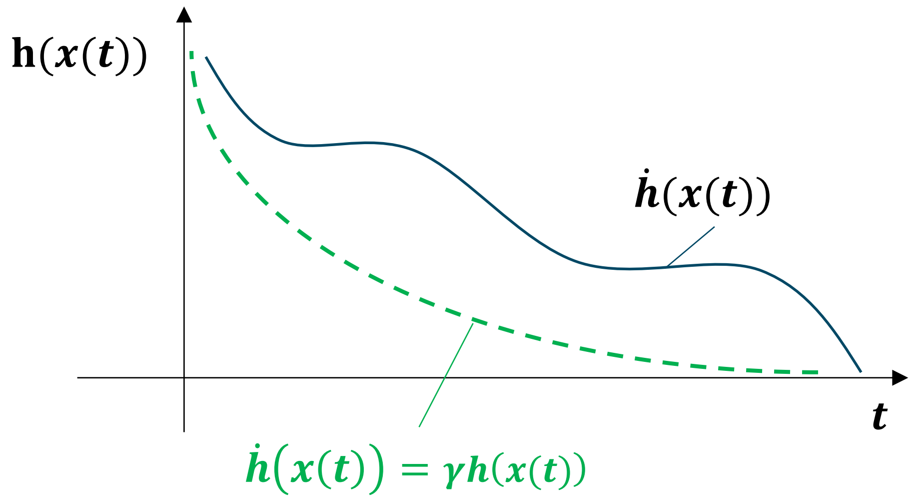

Recap: How CBFs help to ensure safety for closed-loop systems

We consider a performance-oriented controller:

$$u^*(t) = K(x(t), r(t))$$

Standard closed-loop system:

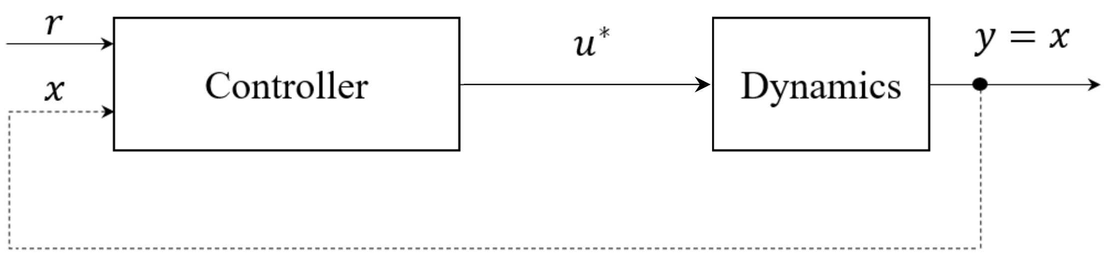

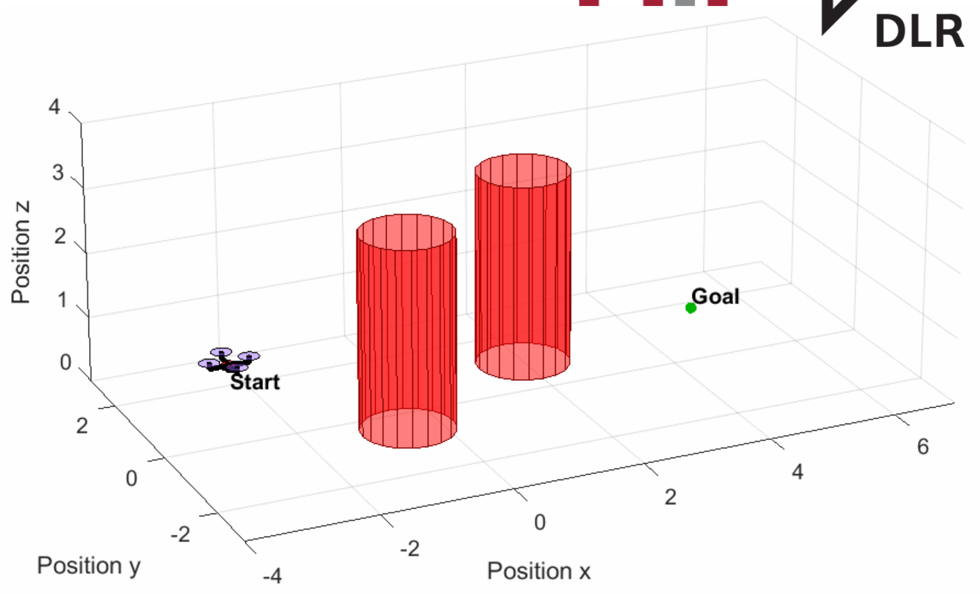

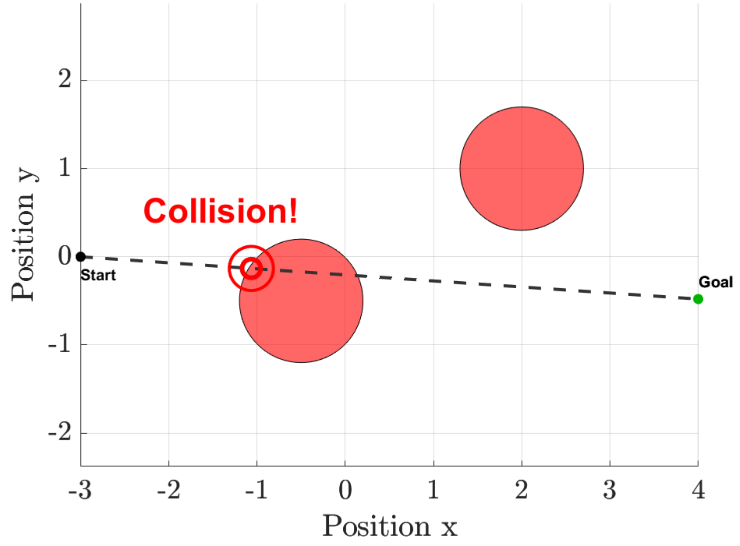

Controller designed to ensure desired closed-loop dynamics:

- Robust stability and reference tracking
- **<u>Safety in design neglected!</u>**

Recap: How CBFs help to ensure safety for closed-loop systems (cont.)

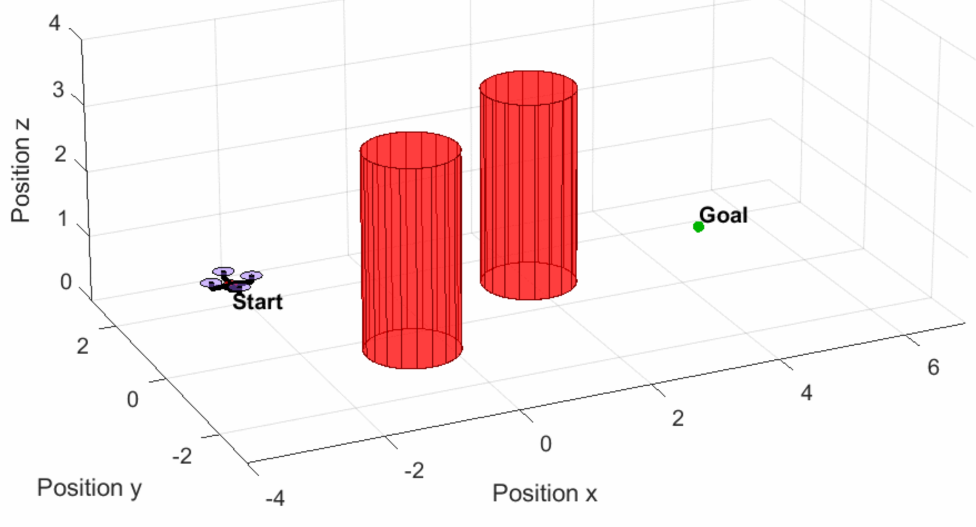

We consider a performance-oriented controller:

$$u^*(t) = K(x(t), r(t))$$

Closed-loop system with safety filter:

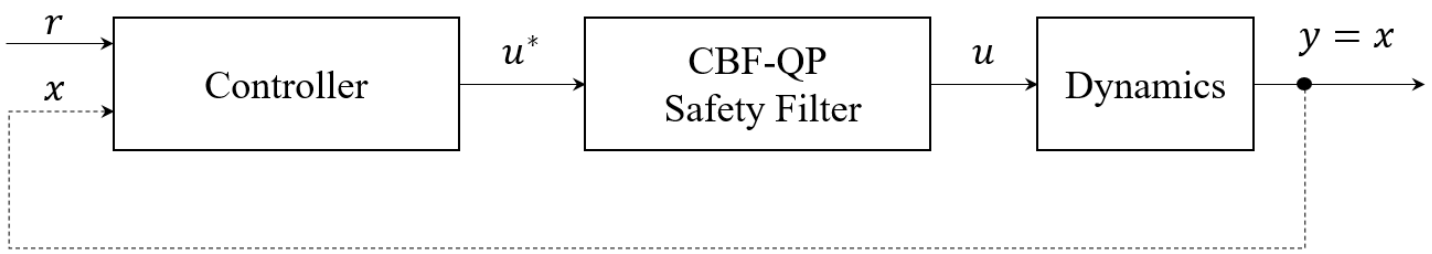

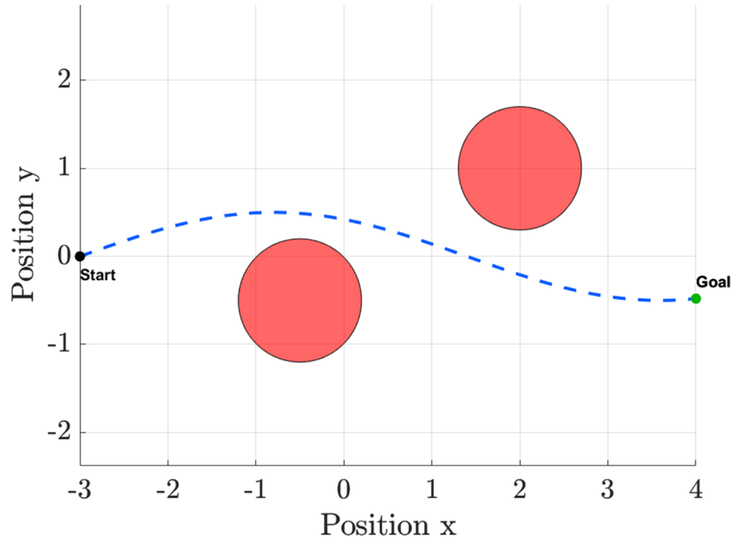

Safety filter often chosen to be quadratic program:

$$\min_{u \in \mathbb{R}^m} \|u - u^*\|_2^2$$

Cost function balances between safety and performance!

s.t.

$$L_f h(x(t)) + L_g h(x(t)) u(t) \geq -\gamma h(x(t))$$

Recap: What is the influence of the design parameter?

I sad γ can be freely chosen as long > 0?

- It's correct, but it's just part of the truth!
- Choice on γ has implications on how the system is rendered safe (balance between performance and safety).

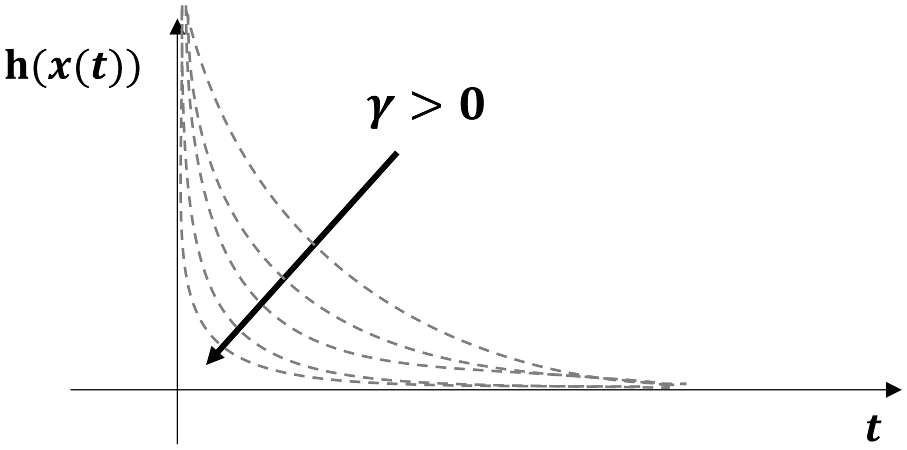

Example (first-order system):

- If $\gamma \to \infty$, we would essentially only impose safety at the boundary (Nagumo's Theorem).
- When not carefully chosen actuator limits could lead to infeasibilities.

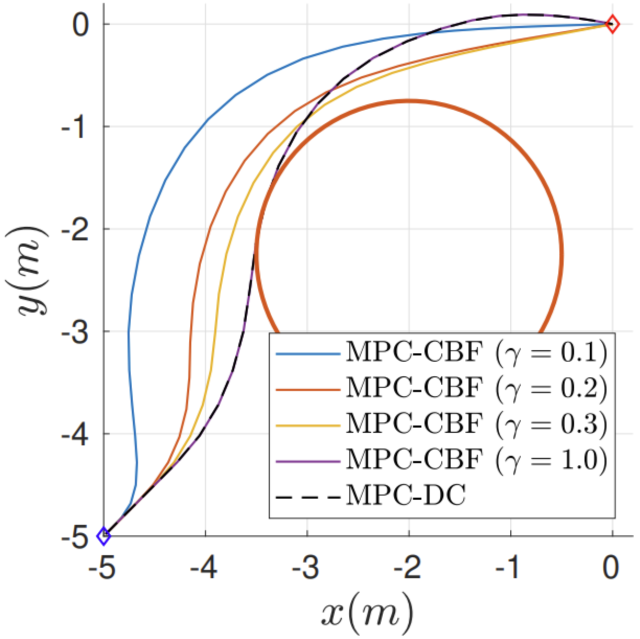

Further details: J. Zeng, B. Zhang, and K. Sreenath, "Safety-Critical Model Predictive Control with Discrete-Time Control Barrier Function," American Control Conference (ACC) 2021

Recap: Anything else I need to know before I can use CBFs on aerospace problems?

Presented CBF approach is only valid for systems with a relative degree of one. We consider:

$$L_f h\big(x(t)\big) + L_g h\big(x(t)\big) u(t) + \alpha\big(h\big(x(t)\big)\big) \geq 0,$$

If $L_g h\big(x(t)\big) = 0$ we got no influences on the dynamics/inequality.

To cope with that we use a backstepping–like recursive approach to derive **high-order control barrier functions (HOCBFs)**. We construct:

$$\psi_0(x(t), t) = h(x(t), t)$$
$$\psi_1(x(t), t) = \dot{\psi}_0(x(t), t) + \alpha_1(\psi_0(x(t), t))$$
$$\psi_2(x(t), t) = \dot{\psi}_1(x(t), t) + \alpha_2(\psi_1(x(t), t))$$
$$\vdots$$
$$\psi_d(x(t), t) = \dot{\psi}_{d-1}(x(t), t) + \alpha_d(\psi_{d-1}(x(t), t))$$

→ Final inequality for imposing safety: $\psi_d(x(t), t) \geq 0$

Case study 1: Flight envelope protection of missile

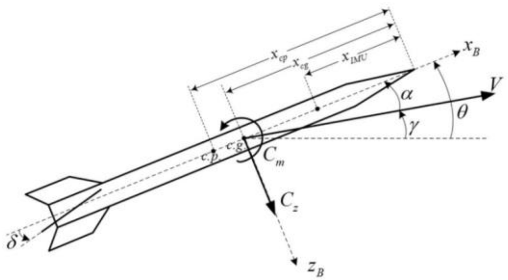

Considered dynamics

$$\dot{\alpha}(t) = \frac{QS}{mV}\left(C_{Z,\alpha}(M)\alpha(t) + C_{Z,\delta}(M)\delta(t)\right) + q(t)$$

$$\dot{q} = \frac{QSd}{I_{yy}}\left(C_{M,\alpha}(M)\alpha(t) + C_{M,q}(M)\frac{d}{2V}q(t) + C_{M,\delta}(M)\delta(t)\right)$$

$$n_z = \frac{QS}{mg}\left(C_{Z,\alpha}(M)\alpha(t) + C_{Z,\delta}(M)\delta(t)\right)$$

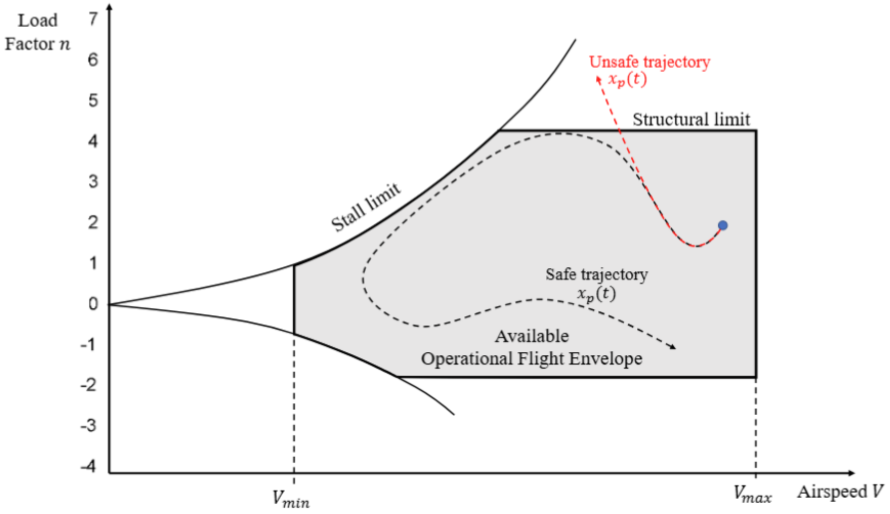

Nonlinear Attitude Controller Design

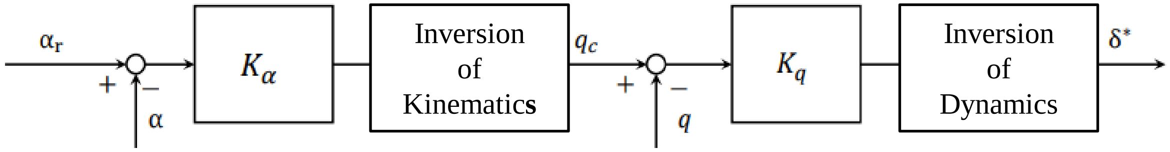

Outer-loop controller:

$$\dot{\alpha}_c(t) = K_\alpha\left(\alpha_r(t) - \alpha(t)\right)$$

$$q_c(t) = \dot{\alpha}_{c(t)} - \frac{QS}{mV}\, C_{Z,\alpha}(M)\alpha(t) + C_{Z,\delta}(M)\,\delta(t)$$

Inner-loop controller:

$$\dot{q}_c(t) = K_q\left(q_c(t) - q(t)\right)$$

$$\delta^{*(t)} = C_{M,\delta}^{-1}(M)\dot{q}_c(t) - C_{M,\delta}^{-1}(M)\frac{QSd}{I_{yy}}\left(C_{M,\alpha}(M)\alpha(t) + C_{M,q}(M)\frac{d}{2V}q(t)\right)$$

Further details: J. Autenrieb, A Quadratric Programming Approach to Flight Envelope Protection using Control Barrier Functions, Journal of Guidance, Control, and Dynamics

Case study 1: Flight envelope protection of missile (cont.)

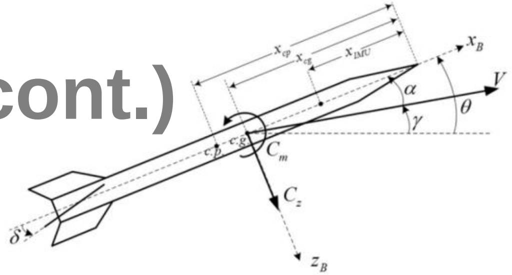

We focus on the upper load limits and we assume that we can impose them by limiting the maximum angle of attack α!

To impose an upper-limit for α we propose the following CBF candidate $h(x(t))$:

$$h\big(x(t)\big) = \alpha_{max} - \alpha(t)$$

Considering that the dynamics have a relative degree of two, it follows:

$$\psi_0(x(t)) = h(x(t))$$

$\alpha_{max} = const.,\ \dot{\alpha}_{max} = 0$

$$\psi_1\big(x(t)\big) = \dot{\psi}_0\big(x(t)\big) + \gamma_1\psi_0\big(x(t)\big) = -\dot{\alpha}(t) - \gamma_1(\alpha_{max} - \alpha(t))$$

Can be used in QP!

$$\psi_2\big(x(t)\big) = \dot{\psi}_1\big(x(t)\big) + \gamma_2\psi_1\big(x(t)\big) = -\ddot{\alpha}(t) - (\gamma_1 + \gamma_2)\dot{\alpha}(t) + \gamma_1\gamma_2(\alpha_{max} - \alpha(t)) \geq 0$$

Estimation via numerical construction and assumption on model properties, see paper!

Case study 1: Flight envelope protection of missile (cont.)

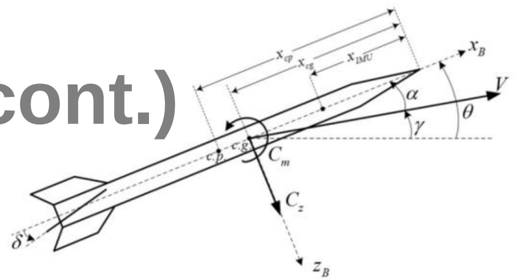

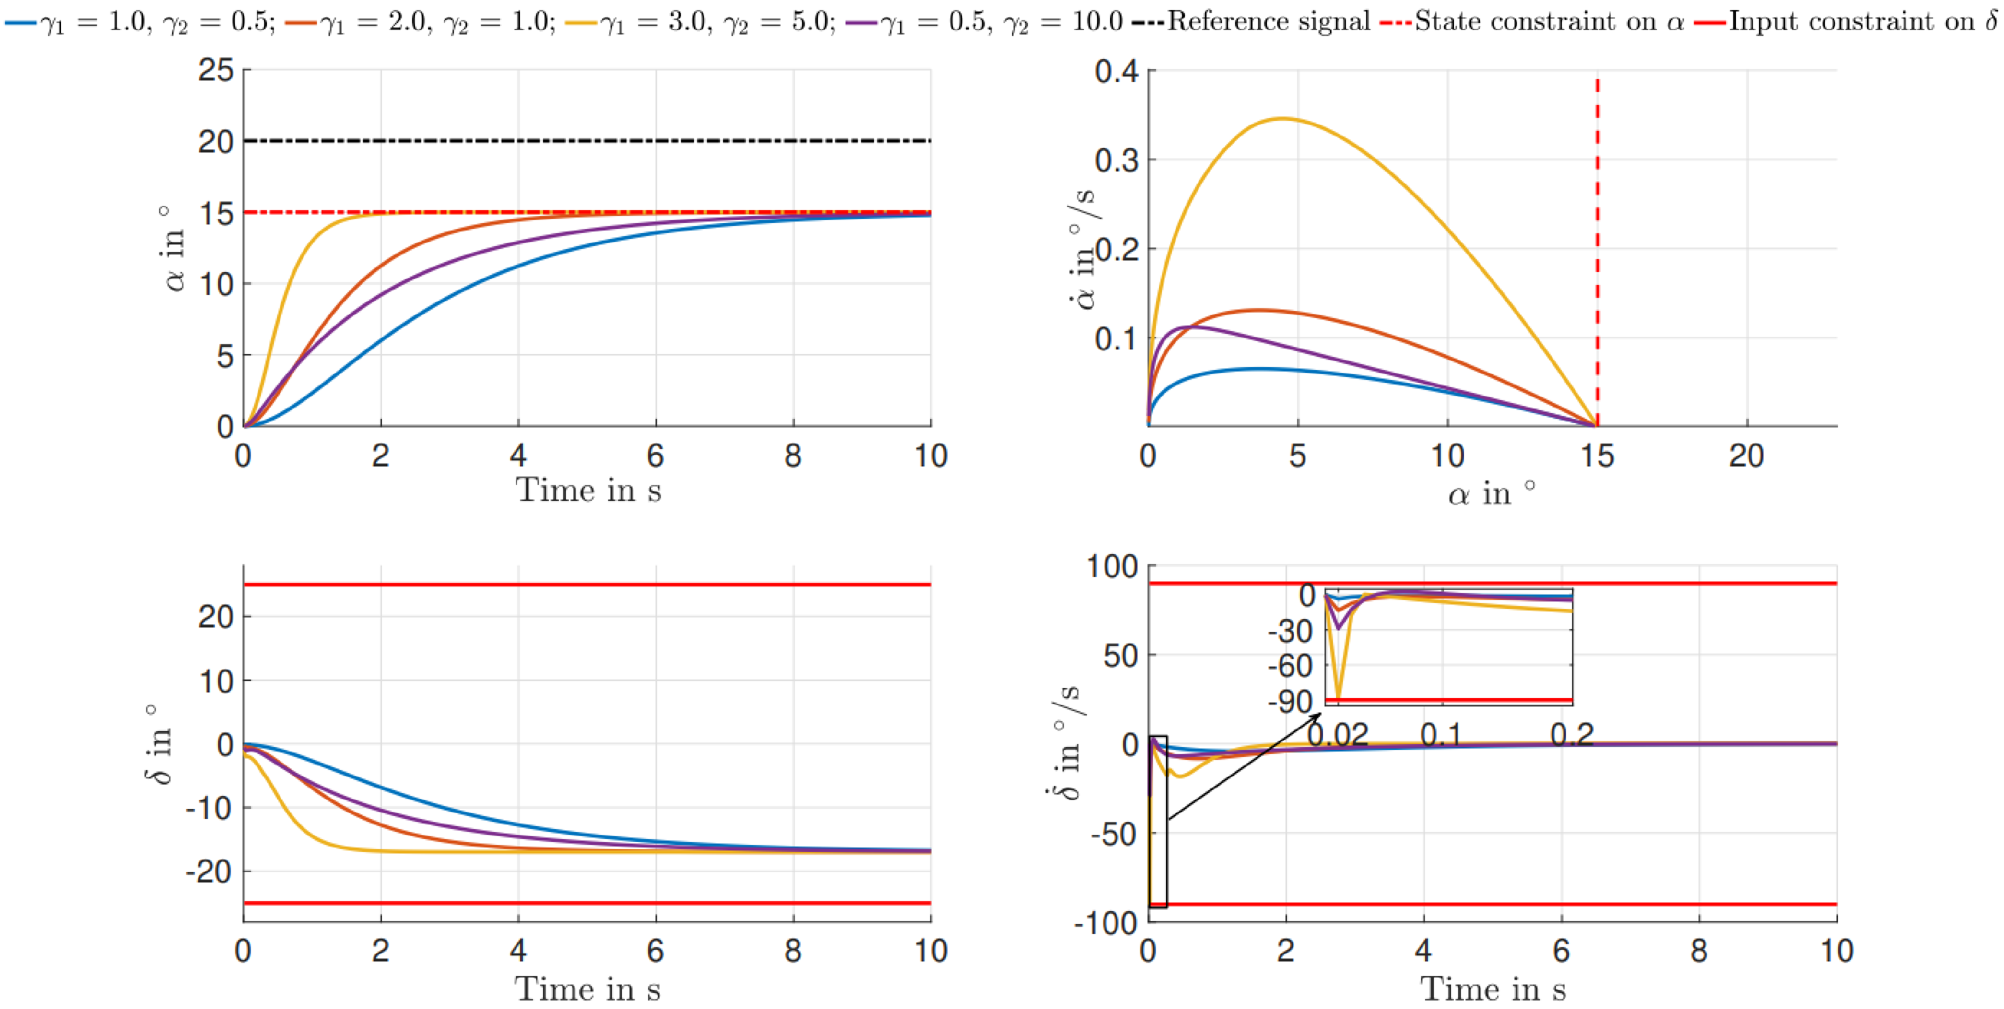

Further example: Obstacle avoidance for unmanned helicopter

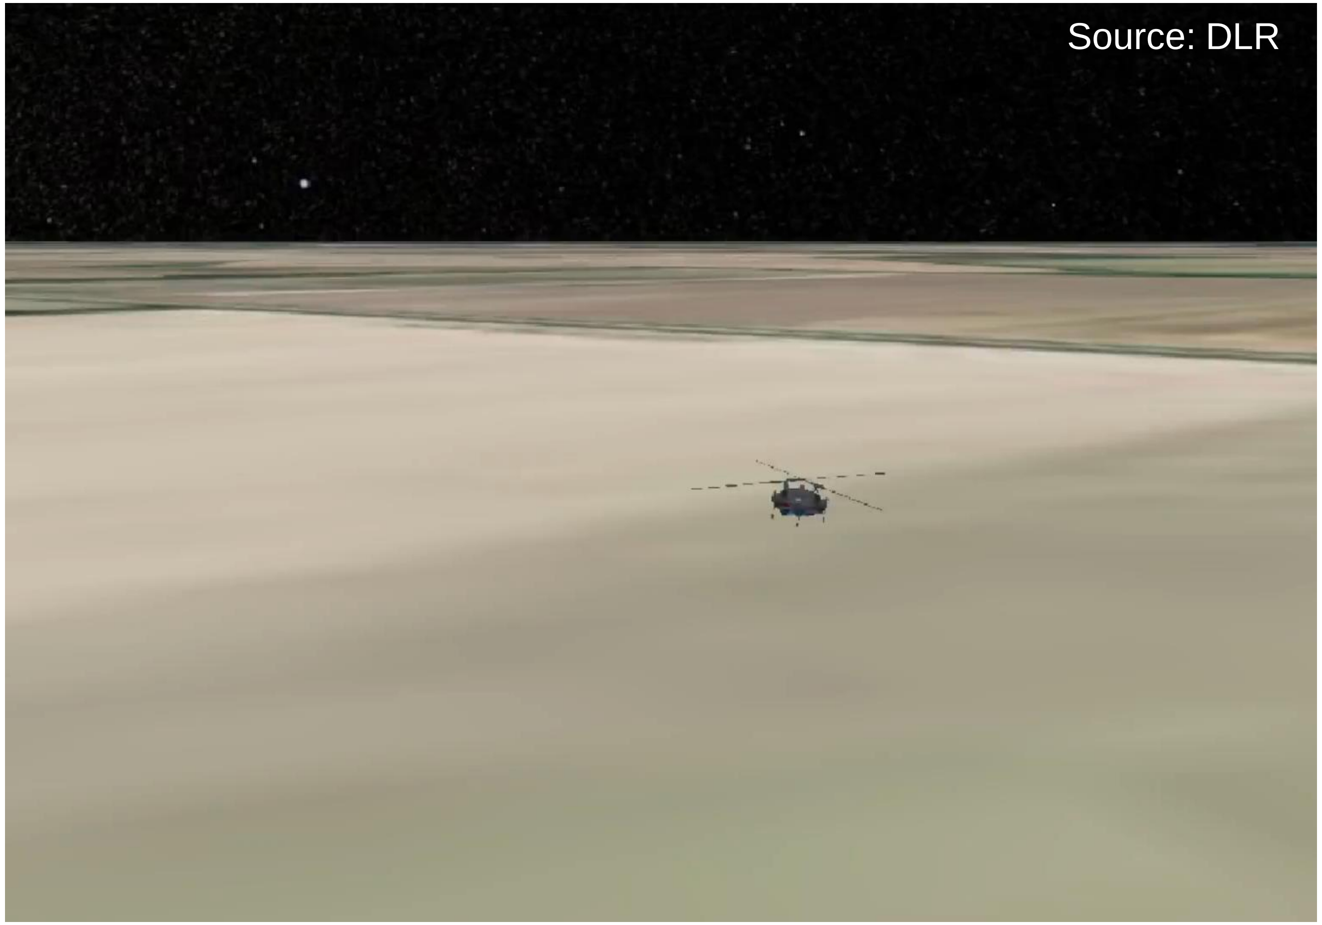

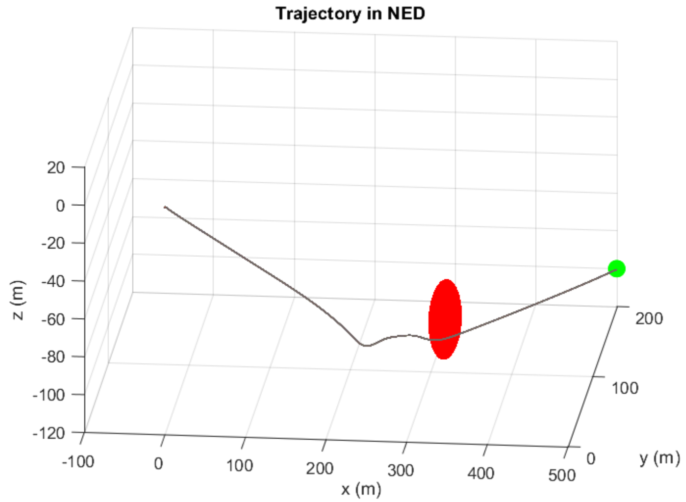

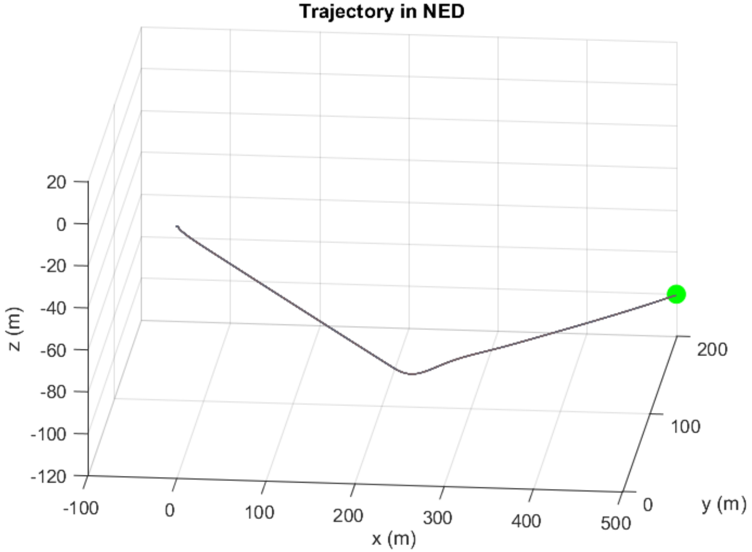

Further details: M. Spiller , E. Isbono, and P. Schitz, Feasibility of multiple robust control barrier functions for bounding box constraints, American Control Conference (ACC) 2025

Further example: Collision avoidance for cooperative missiles

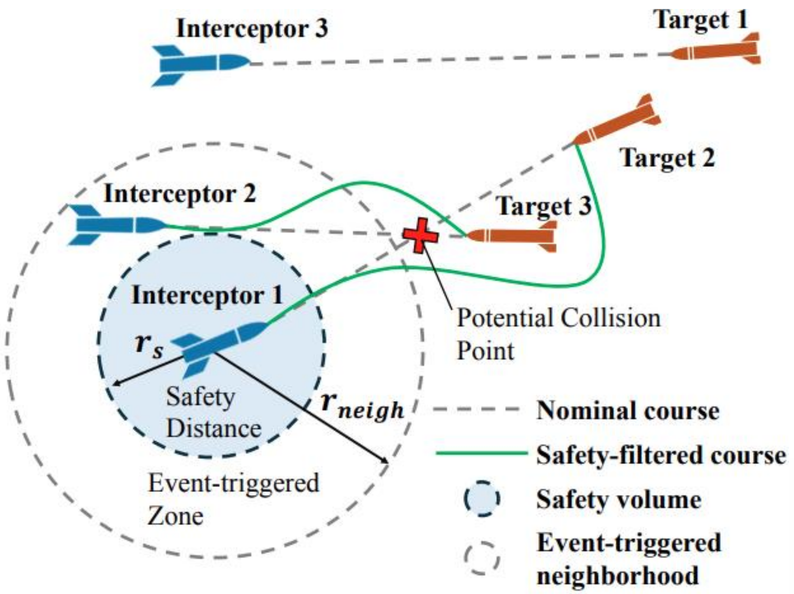

Further details: J. Autenrieb and M. Spiller, Decentralized CBF-Based Safety Filters for Collision Avoidance of Cooperative Missile Systems with Input Constraints, American Control Conference 2026

Cooperative systems w/ safety filters

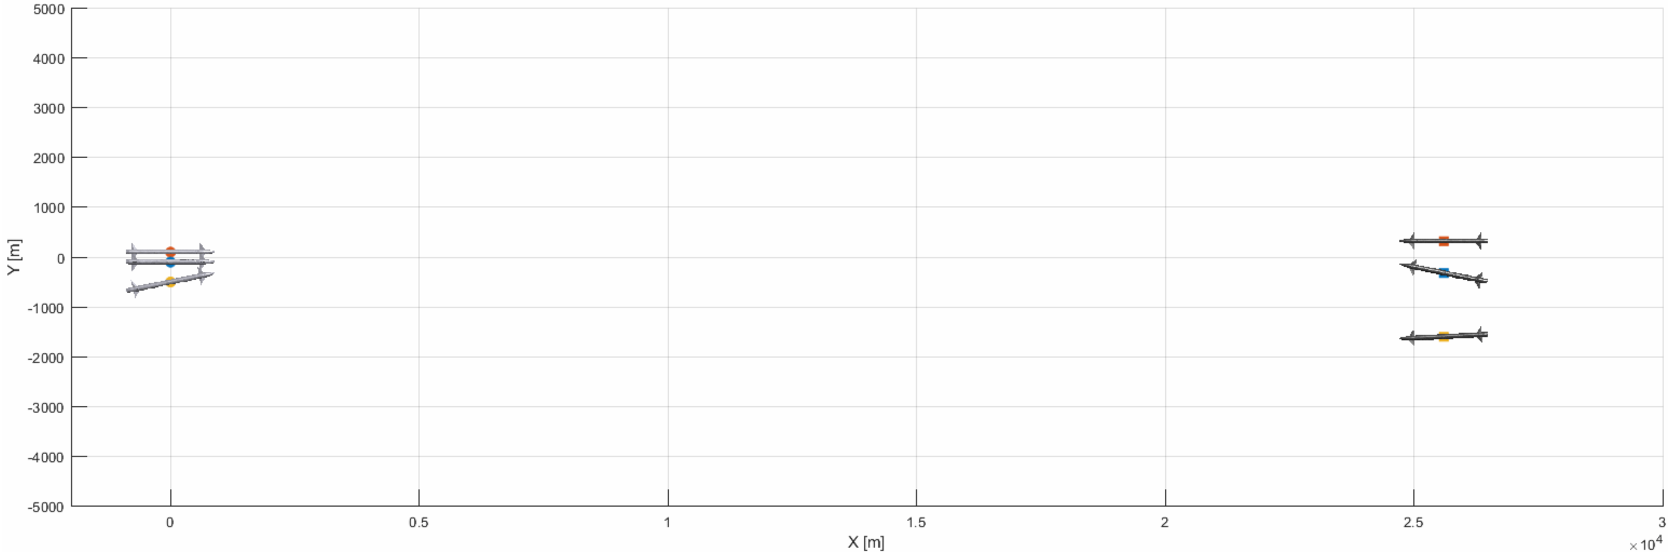

Cooperative systems w/o safety filters

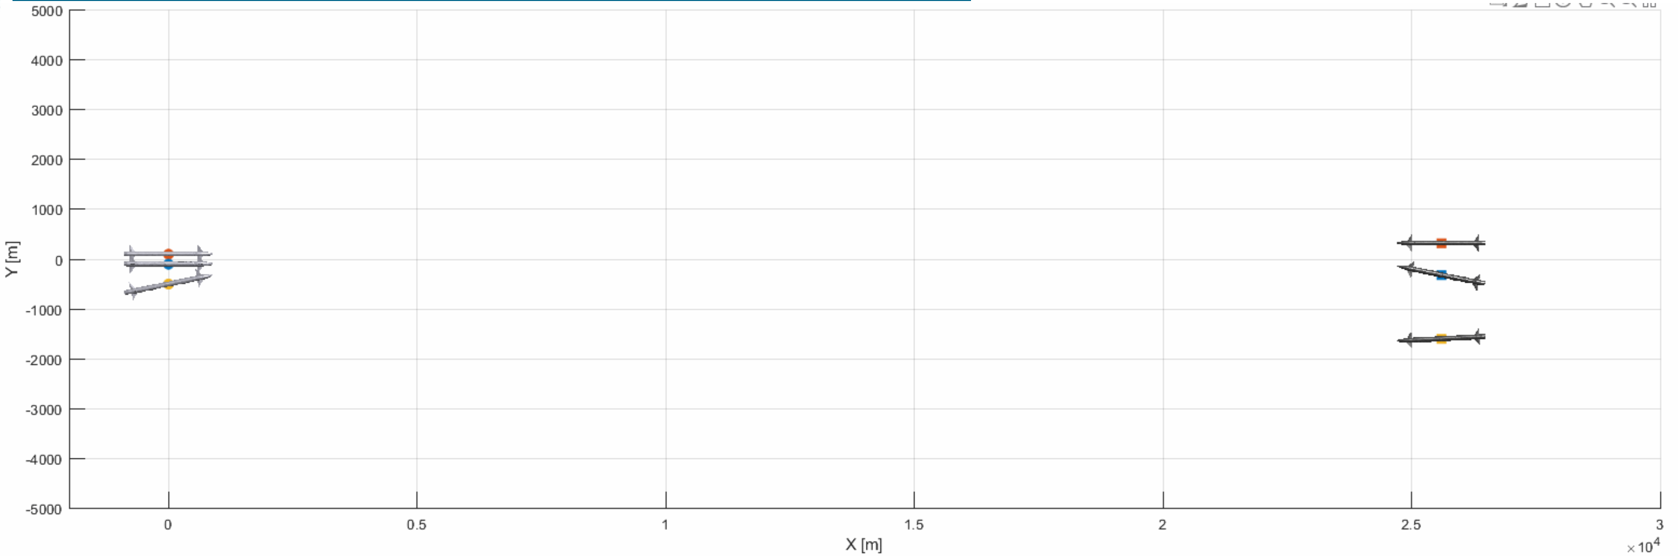

Further Examples in MATLAB

QR Code

MATLAB code can be found on GitHub!

Username: JohannesAutenrieb

https://github.com/JohannesAutenrieb

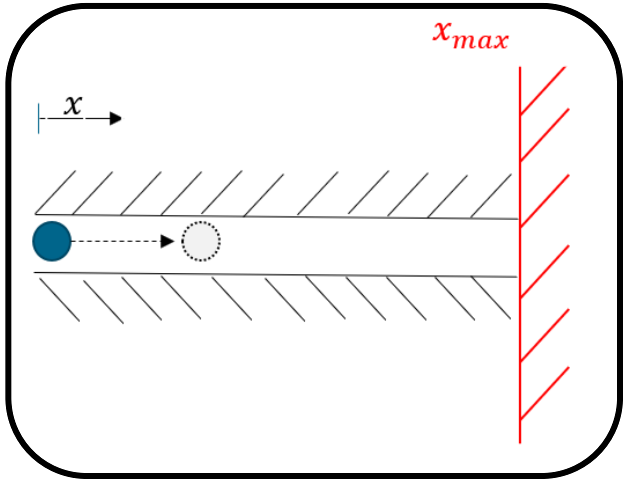

1-D State Limit

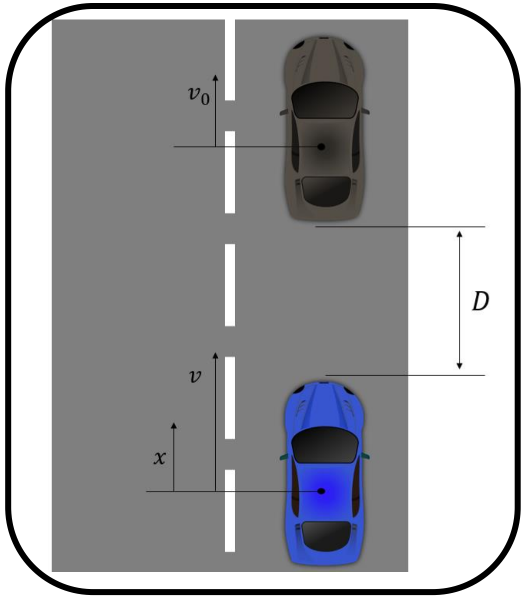

Adaptive Cruise Control

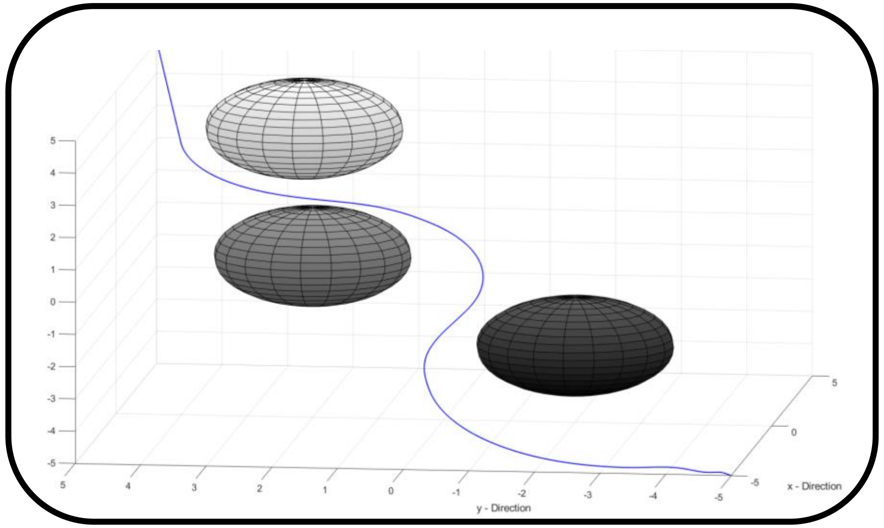

Obstacle Avoidance

What are the key takeaways so far?

Intermediate Summary:

- CBFs provide a powerful framework to enforce state constraints and safety guarantees for aerospace systems.
- Performance-oriented controllers can be designed, while additional CBF-based safety filters ensure constraint satisfaction.
- Model uncertainties and actuator limitations, can lead to possible problems and loss of safety guarantees.

Any Question so far?

Outline

1 Concept of Control Barrier Functions (CBFs)

2 Safety Filter Design for Adaptive Closed-Loop Systems

3 Outlook & Conclusions

Outline

1. Concept of Control Barrier Functions (CBFs)
2. Safety Filter Design for Adaptive Closed-Loop Systems
3. Outlook & Conclusions

Even though CBF-based safety filters have benefits, we got major issues to address!

Remember architecture:

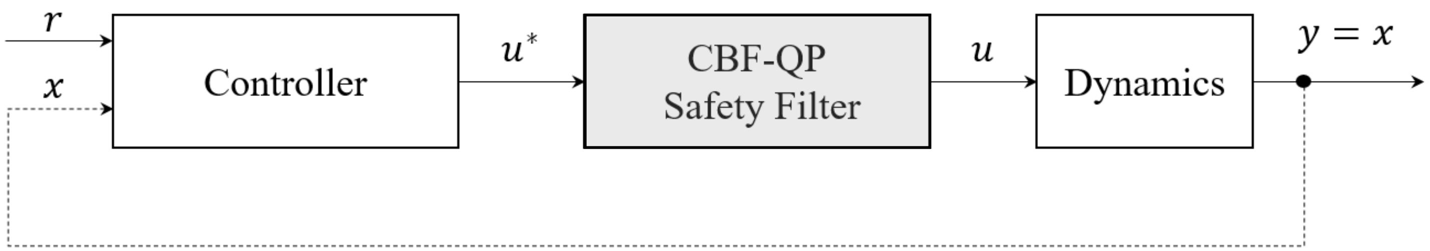

Possible challenges:

- Windup
- Robustness Guarantees

Further possible problems:

- Input Limits
- Uncertainties on the mathematical model
- Unknown external disturbances on the system

Why is it difficult to ensure safety for systems with uncertainties?

- **CBFs are a model-based method**,
  - deviation from the assumed model, could lead to safety violations.
- Origins for model mismatches:
  - wrong model assumptions, system failures or unknown external disturbances on the system.
- It is possible to address uncertainties via
  - **robust,** **adaptive** or **sensor-based approaches**.

Example 1: F-16 with damaged wing.

Example 2: Airbus A380 with engine failure.

How does adaptive control work?

One should distinguish between:

- **Direct adaptive control**: Estimated parameters are those directly used in the adaptive controller.
  - Error between desired and real dynamics is used to tune controller gains.

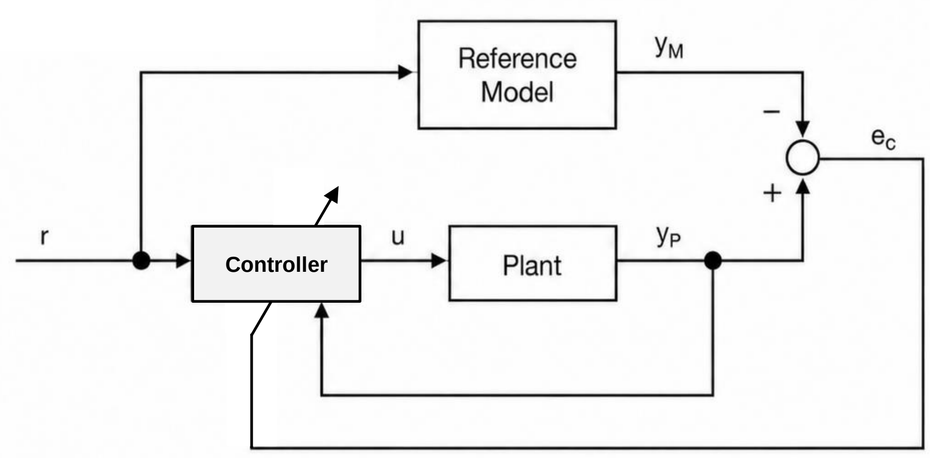

- **Indirect adaptive control**: Estimated parameters are used to calculate controller parameters
  - Error between assumed plant and real dynamics is used to identify plant parameters → plant parameters are then used to update controller.

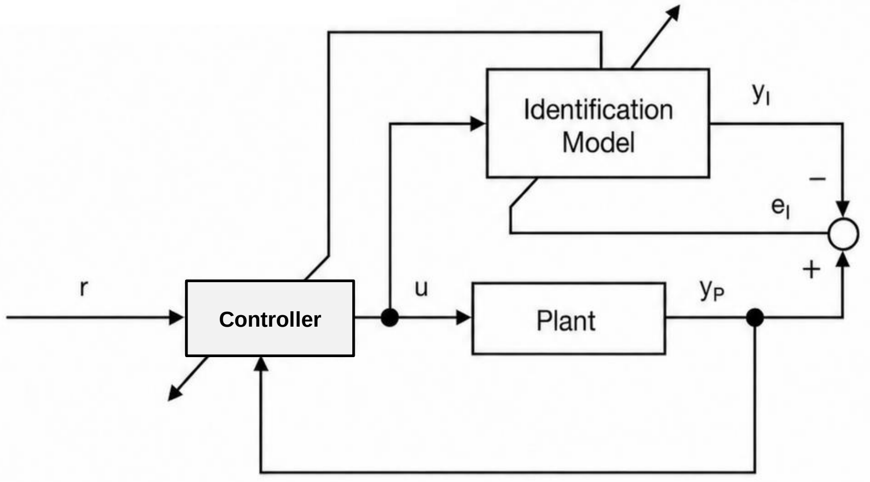

How does adaptive control work? (cont.)

Indirect Adaptive Control

- Controller performance depends on accuracy of the estimates.
  - Needs exciting input signals to converge correctly.
- Problem: Most systems do not provide sufficient excitation during operation.
- This leads potentially to inaccurate parameter estimates and with that:
  - Weaker guarantees
  - Risk of poor performance or instability

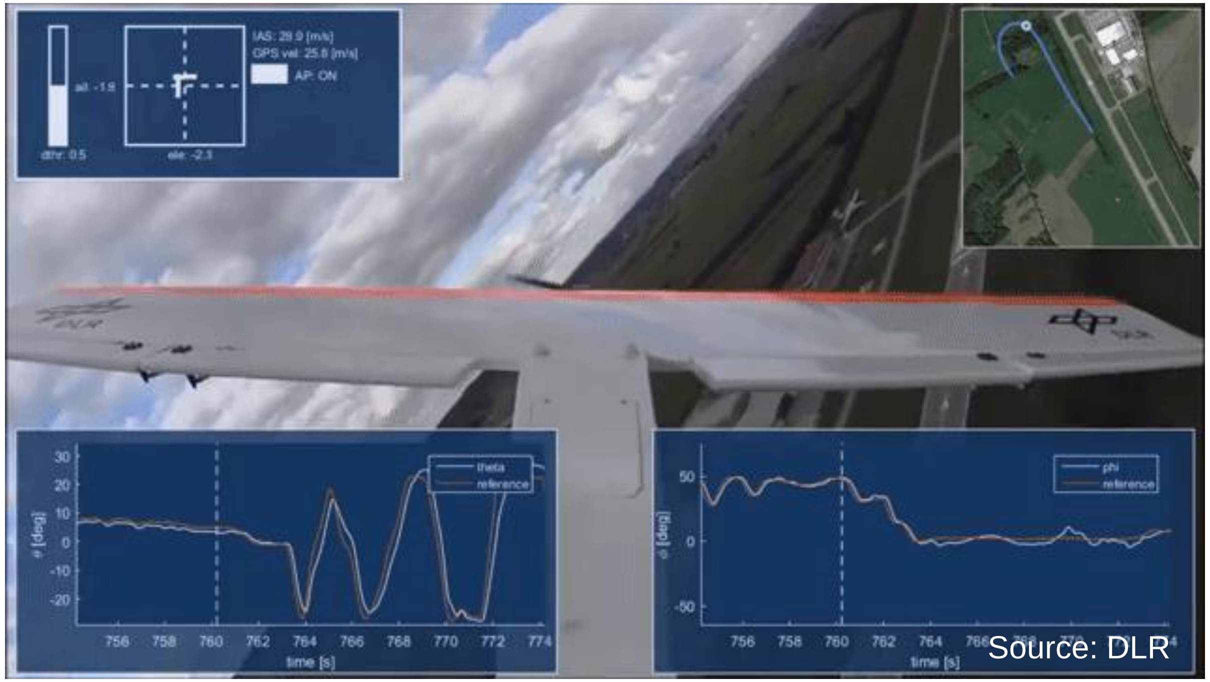

Example for strongly maneuvering fixed-wing drone.

How does adaptive control work? (cont.)

Direct Adaptive Control

- Controller performance depends on direct adaptation from tracking error
  - Does not require accurate parameter estimation for stability/performance guarantees → Smaller dependence on exciting input signals to work correctly
- This allows:
  - Stronger stability guarantees
  - More reliable closed-loop behavior

Example for quadcopter with damage (w/ & w/o direct adaptive control).

Problem: Safe Direct Adaptive Control

Plant: $\dot{x}_p(t) = A_p x_p(t) + B_p \Lambda u(t)$

Unknown: $A_p \in \mathbb{R}^{n\times n}, \Lambda \in \mathbb{R}^{m\times m}$, $\theta_x^* \in \mathbb{R}^{n\times n}$ and $\theta_r^* \in \mathbb{R}^{p\times p}$

Ideal controller: $u^*(t) = \theta_x^* x(t) + \theta_r^* r(t)$

Matching conditions:

$$A_m = A_p + B_p \Lambda \theta_x^*$$

$$B_m = B_p \Lambda \theta_r^*$$

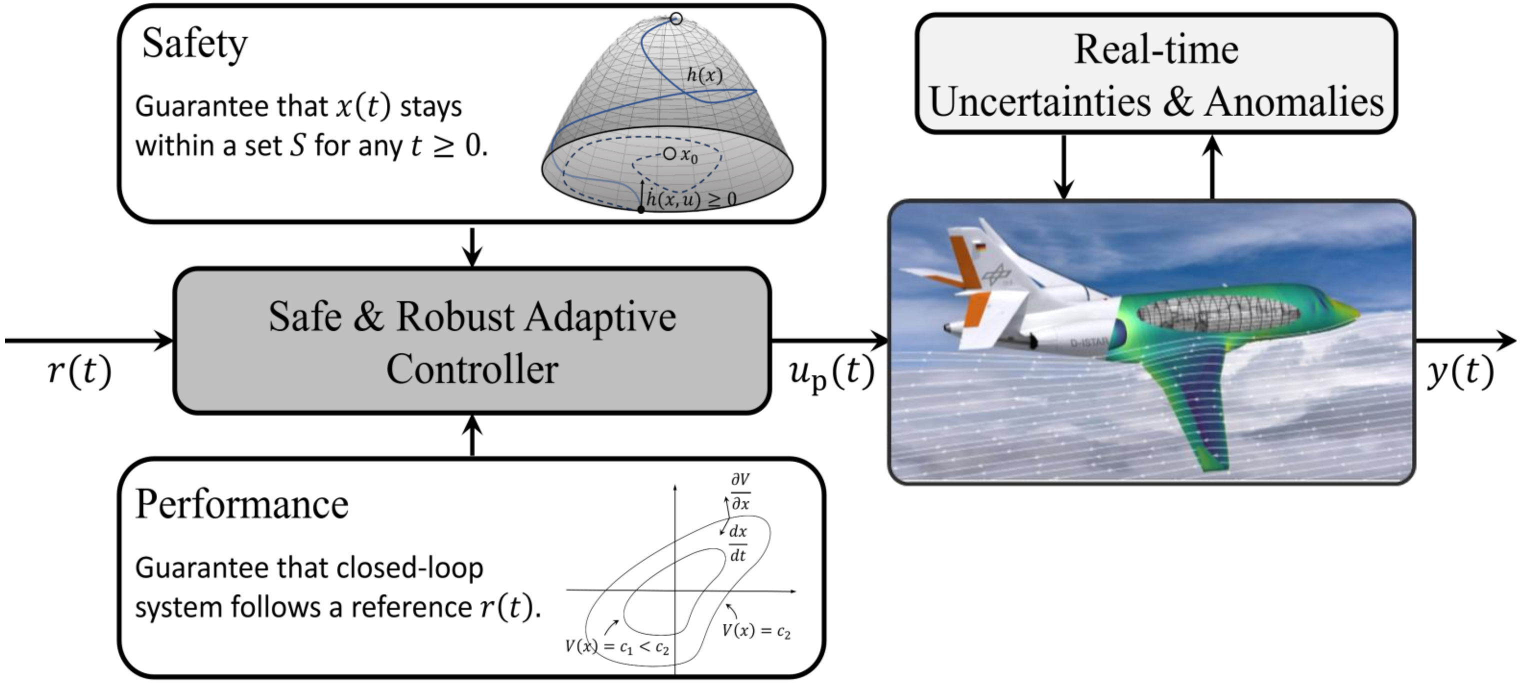

Design of Model Reference Adaptive Controller (MRAC)

Considered reference model:

$$\dot{x}_m(t) = A_m x_m(t) + B_m r(t)$$

Adaptive laws to enforce the desired closed-loop dynamics:

$$\dot{\hat{\theta}}_x(t) = -\Gamma_x x_p(t) e_x^T(t) P B_P$$

$$\dot{\hat{\theta}}_r(t) = -\Gamma_r r(t) e_x^T(t) P B_P$$

Integrated adaptive control law:

$$u(t) = \hat{\theta}_x(t)\, x(t) + \hat{\theta}_r(t)\, r(t)$$

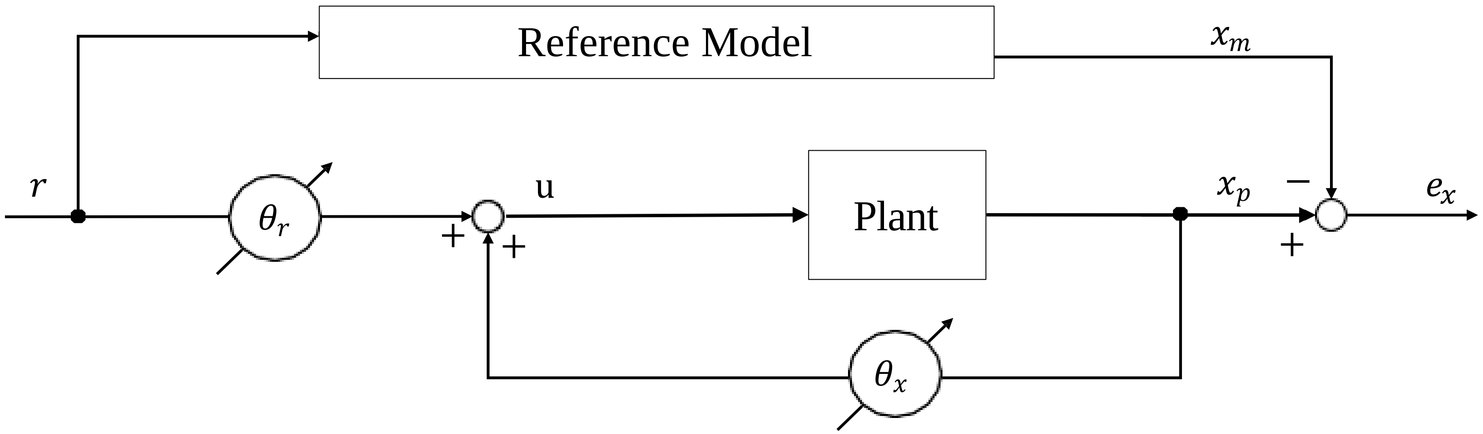

Is the closed-loop system with MRAC stable?

Theorem 1 (Simplified Version)

The closed-loop adaptive system has globally bounded solutions for any initial conditions $x_p(t_0)$, $\tilde{\theta}_x(t_0)$, *and* $\tilde{\theta}_r(t_0)$ *and both the errors* $e_x(t)$ *and* $e_u(t)$ *converge to zero as* $t \to \infty$.

Error terms:

$$e_x(t) = x_p(t) - x_m(t)$$

$$e_u(t) = u(t) - u^*(t)$$

Matching conditions::

$$A_m = A_p + B_p \Lambda \theta_x^*$$

$$B_m = B_p \Lambda \theta_r^*$$

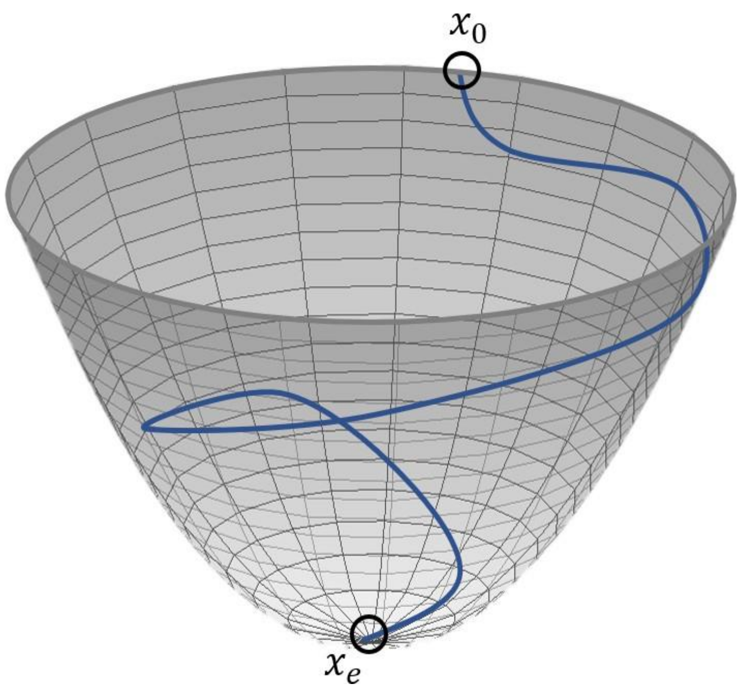

Lyapunov candidate:

$$V(x) = \frac{1}{2} e_x(t) P e_x(t) + \frac{1}{2} Tr\left[\tilde{\theta}_x \Gamma_x^{-1} \tilde{\theta}_x^T(t) \Lambda\right] + \frac{1}{2} Tr\left[\tilde{\theta}_r \Gamma_r^{-1} \tilde{\theta}_r^T(t) \Lambda\right]$$

Proof of stability:

$$\dot{V}(x) = -\frac{1}{2} e_x^T Q_0 e_x \leq 0$$

Further details: Kumpati S. Narendra & A. M. Annaswamy, Stable Adaptive Systems

Is it enough to just use adaptive laws to ensure safety?

Stability is good, right? Is this now enough?

→ No, stability does not imply safety!

We got two major questions to address:

1. Where to place the safety filter?
2. How to guarantee safety during the transient?

We can't just place the safety filter on the input-level!

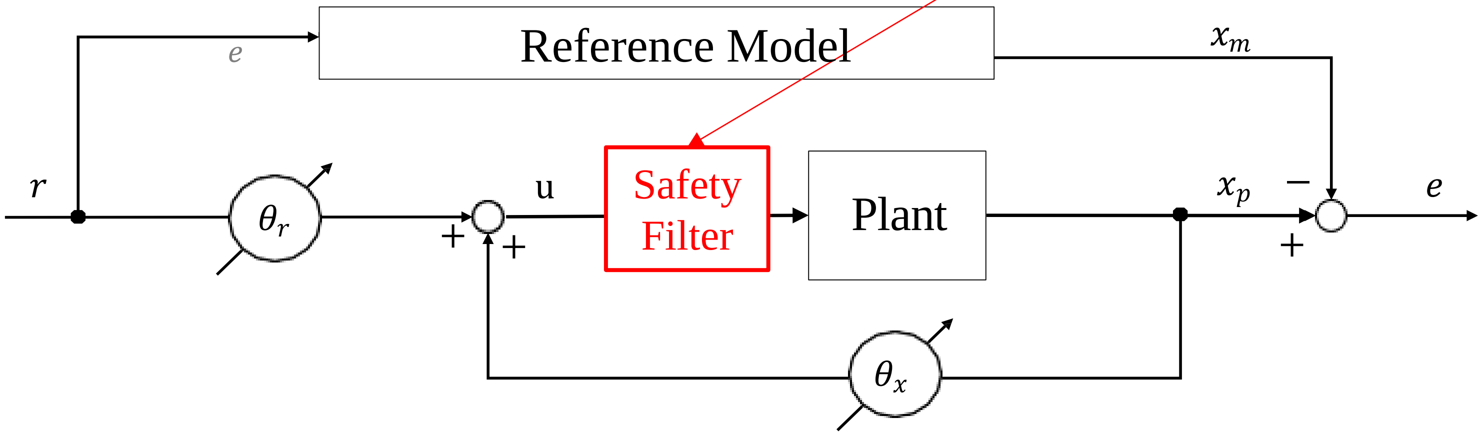

Further details: J. Autenrieb & A. M. Annaswamy, Safe and Stable Adaptive Control for a Class of Dynamic Systems, IEEE CDC 2023

Safe Adaptive controller with open-Loop reference model

- Standard CBF-based safety filters operate directly on the control input, we propose a CBF module that acts on the reference signal-level.
- Forming an **indirect safety-critical controller** that modifies the reference command.
  - Unsafe reference trajectories are rejected before reaching the controller.
  - The filtered reference is then tracked by the adaptive controller, which compensates for model uncertainties.

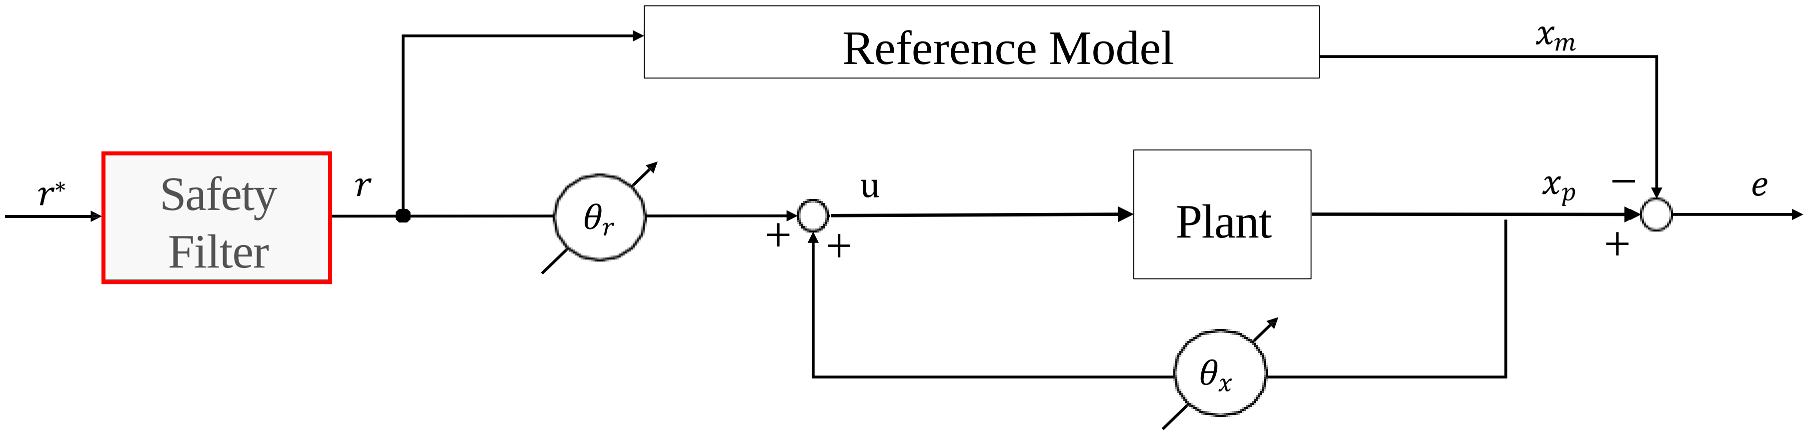

How do we render the adaptive closed-loop system safe? (cont.)

Simplified concept of indirect safety design

$x_m(t)$ is safe $\forall\ t \geq t_0 \Rightarrow x_p(t)$ *is safe for* $\forall\ t \geq t_0$

Standard QP-CB formulation on the reference signal r:

$$\min_{r(t)\epsilon\mathcal{R}} (r(t) - r^*(t))^2$$

s.t.

$$\frac{\partial h}{\partial x}[A_m x_m(t) + B_m r(t)] \geq -\alpha\left(h\big(x_m(t)\big)\right)$$

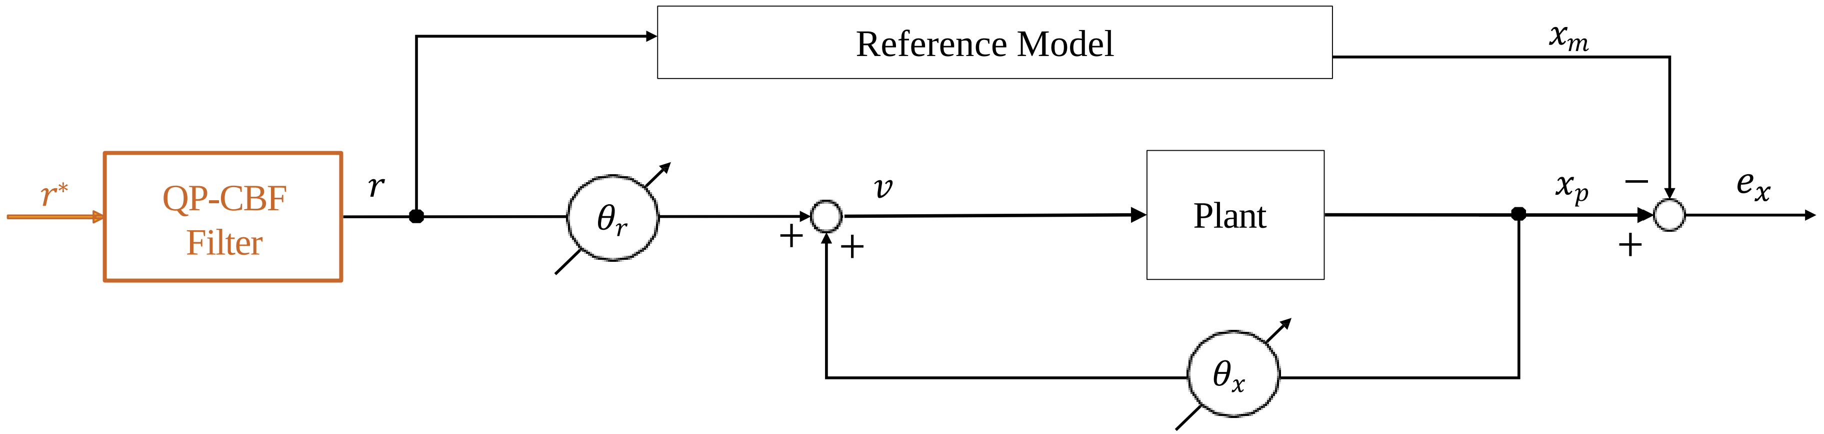

Is it enough to combine adaptive control with the standard CBF methods?

No, it is not enough to ensure safety for a system with uncertainties!

- There is a lack for safety guarantees during the transient.

For guarantees a more conservative formulation would be needed:

$$\dot{h}(x(t), r(t)) \geq -\alpha\big(h(x(t))\big) + \Delta$$

with Δ being a positive scalar constant that introduces a safety buffer.

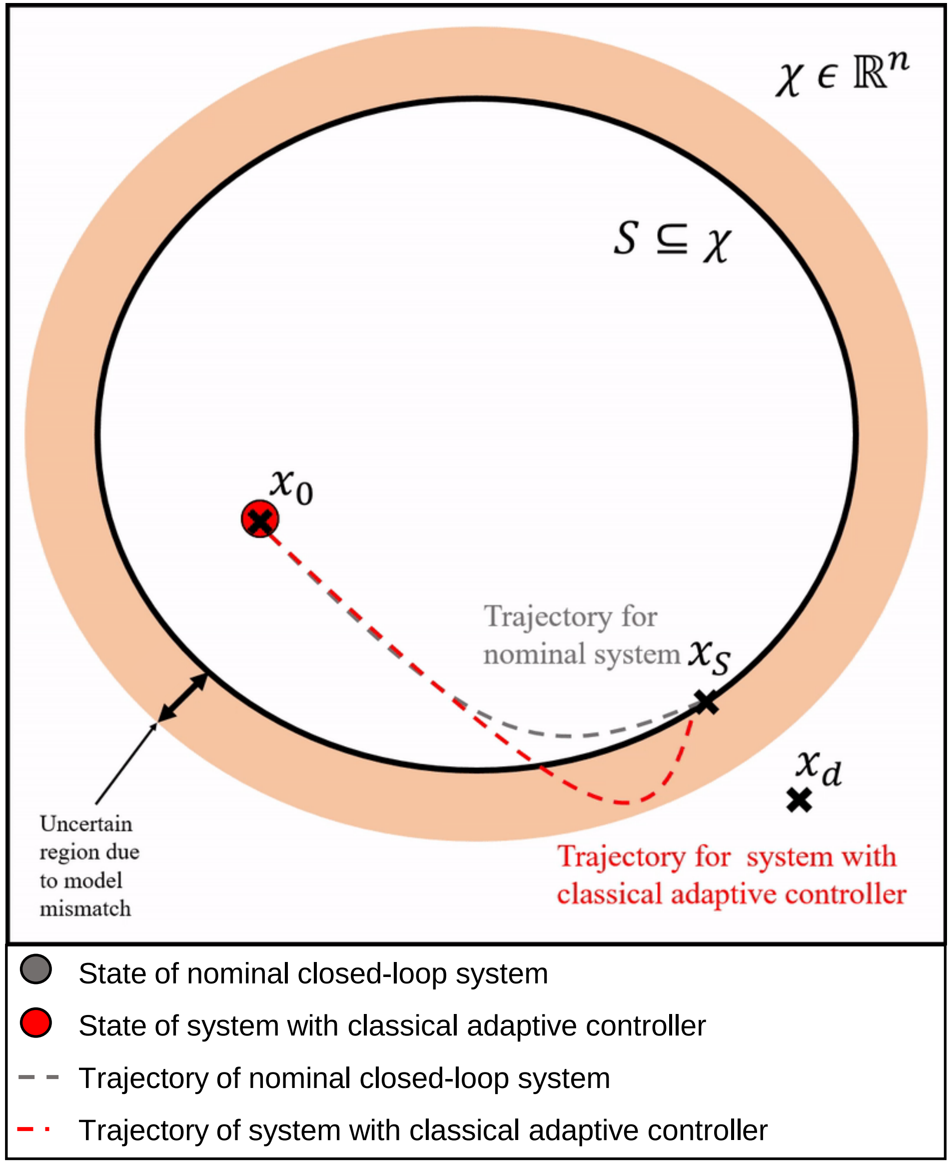

Further details: J. Autenrieb & A. M. Annaswamy, Safe and Stable Adaptive Control for a Class of Dynamic Systems, IEEE CDC 2023

How do we render the adaptive closed-loop system safe? (cont.)

Theorem 2 (Informal Version)

The overall closed-loop adaptive system will remain in the safe set S for $\forall\ t \geq t_0$, if $h\big(x_p(t)\big) \geq h_0$ is satisfied.

We define the barrier function error :

$$e_h(t) = h(x_m(t)) - h(x_p(t))$$

The error is bounded by a maximum value:

$$E = \sup_{t \geq 0} |e_h(t)|$$

Combining the error definition and the bound yields:

$$h\left(x_p(t)\right) = h\big(x_m(t)\big) - e_h(t) \geq h\big(x_m(t)\big) - E \geq 0,$$

We can show that choosing Δ such $\boldsymbol{h\big(x_{m(t)}\big) \geq E}$ holds, guarantees safety!

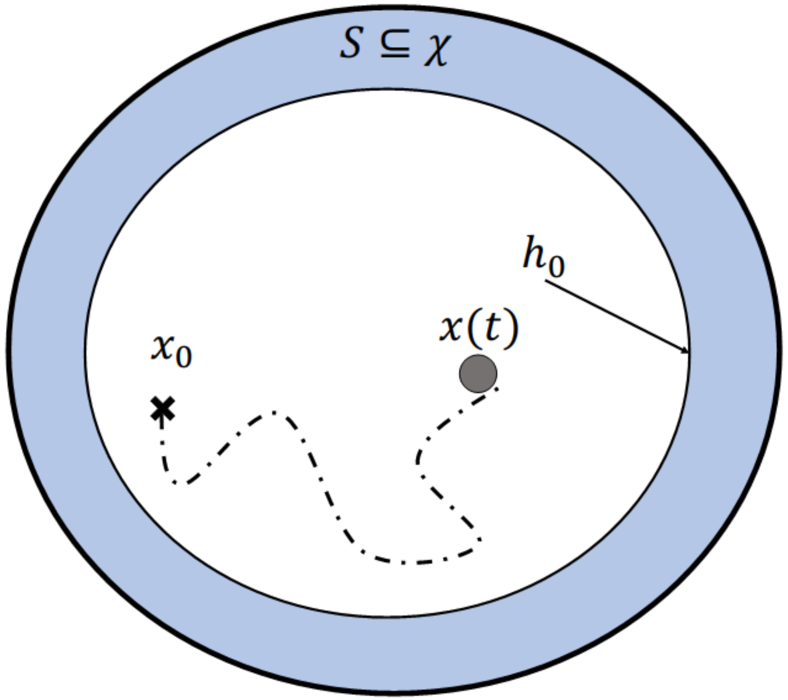

Can we be less conservative when we are confident?

The price of safety is to carry the static safety buffer at all times.

$$\dot{h}(x(t), r(t)) \geq -\gamma h(x) + \Delta$$

<u>However, once adaptation is completed the buffer is not needed anymore.</u>

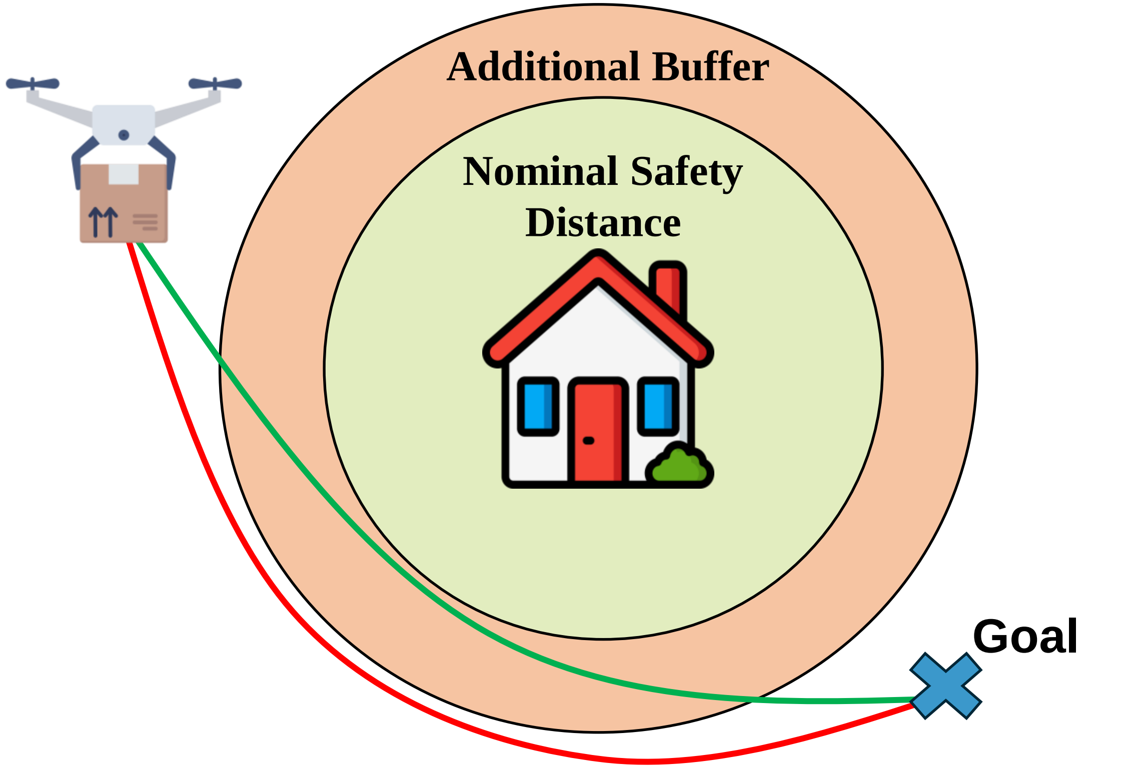

Question: Can we find an adaptive formulation, which is conservative when the model is not well known and less conservative when we are confident?

An intuitive approach: Error-based relaxation

Simple idea: We use $e_h(t)$ to adapt the relaxation parameter $\gamma(e_h(t))$ as needed!

RBF-based $\gamma(e_h(t))$ function:

$$\gamma(e_h(t)) = \gamma_0 e^{-(\epsilon e_h(t))} + \underline{\gamma}$$

with $\gamma_0, \underline{\gamma} > 0$ and $\epsilon \geq 0$ being a design parameter.

The standard CBF constraint be extended with the idea of error-based relaxation via,

$$\sup_{r(t) \epsilon \mathcal{R}} \frac{\partial h}{\partial x} [A_m x_m(t) + B_m r(t)] \geq -\gamma(e_h(t)) h(x_m(t))$$

for all $x_m(t) \epsilon S$.

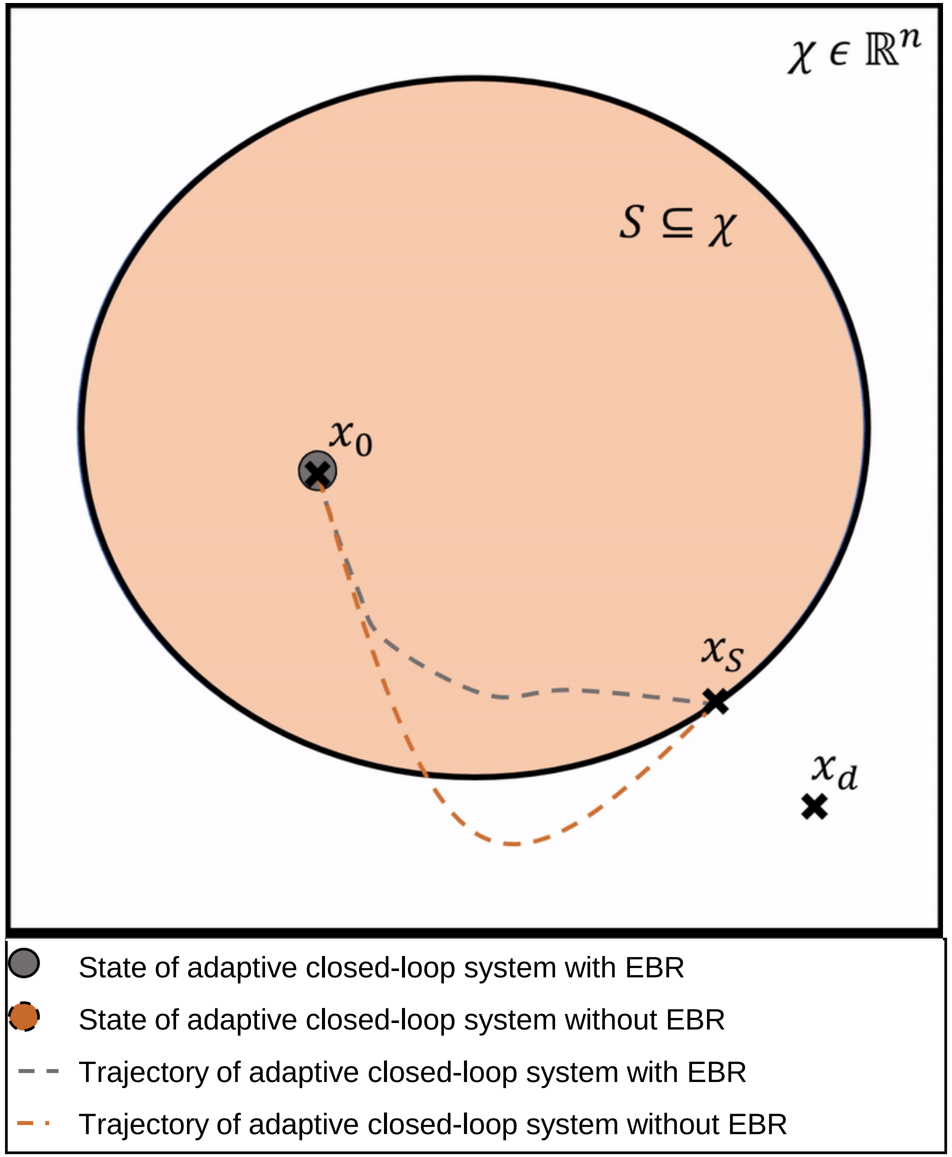

Example: Formation Control with Autonomous Multi-Agents

Only CBF

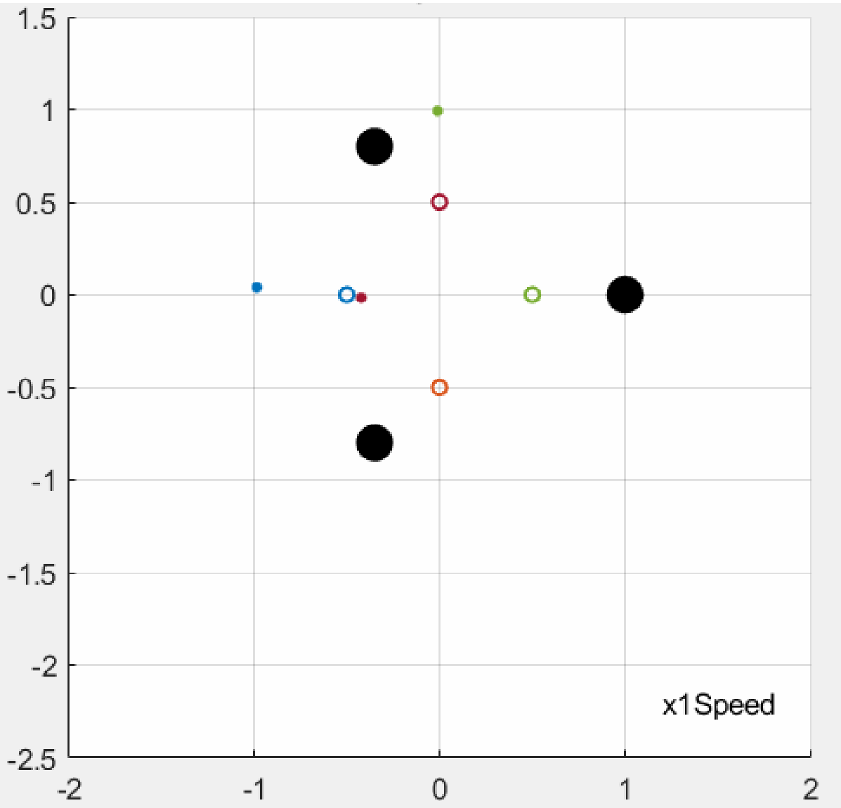

Only Adaptive Control

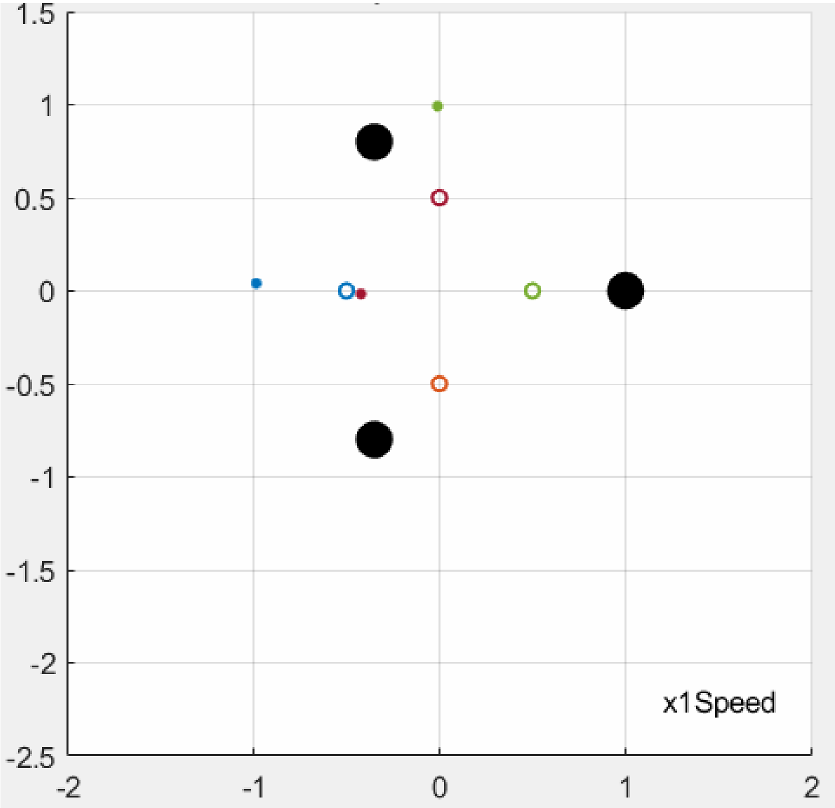

Further details: J.A. Solano-Castellanos, P.A. Fisher, and A. M. Annaswamy, Safe and Stable Formation Control with Autonomous Multi-Agents Using Adaptive Control, Conference on Modeling, Estimation and Control MECC 2025

Example: Formation Control with Autonomous Multi-Agents (cont.)

Only Adaptive Control

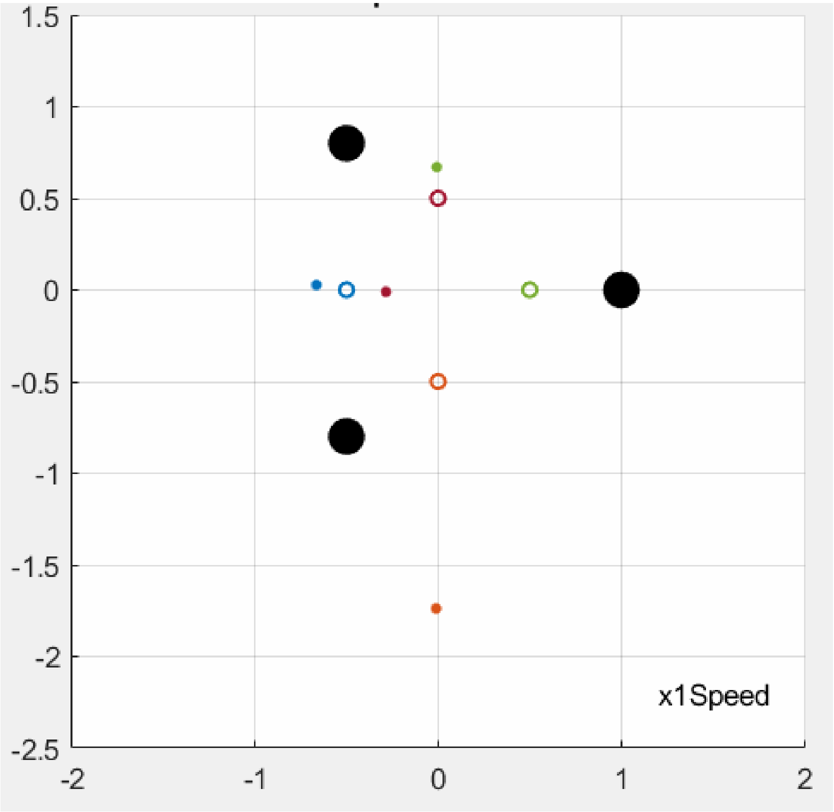

Adaptive Control + CBF (w/ EBR)

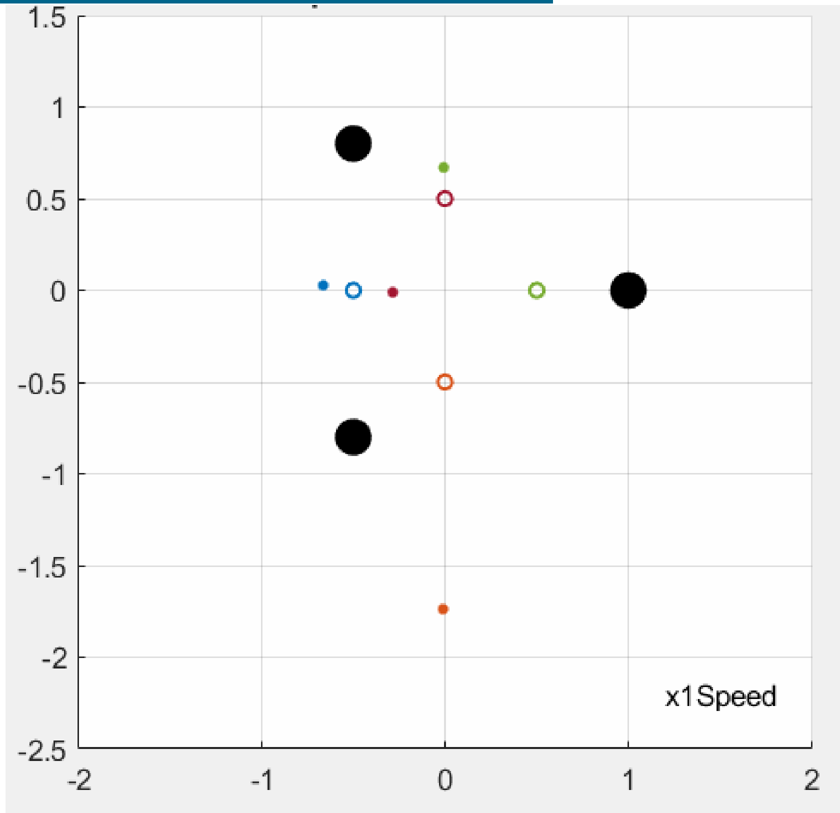

Further details: J.A. Solano-Castellanos, P.A. Fisher, and A. M. Annaswamy, Safe and Stable Formation Control with Autonomous Multi-Agents Using Adaptive Control, Conference on Modeling, Estimation and Control MECC 2025

Formal Safety Guarantees with an Error-Based Safety Buffer (EBSB)

Idea:
- $x_p \in S$ if distance from x_m to ∂S is greater than $\|e_x\|$.
- Let $\|\nabla h(x)\| \leq \kappa \; \forall x$. Then, $x_p \in S$ if $h(x_m) \geq \kappa \|e_x\|$.

We replace the static safety buffer with a function of $e_x := x_p - x_m$

$$\dot{h}(x_m(t), r(t)) \geq -\gamma h(x_m) + \Delta_{ebsb}(x_m, e_x)$$

Design Δ_{ebsb} such that:

1. $h\big(x_m(t)\big) \geq \kappa \|e_x(t)\| \; \forall t \geq t_0$
2. $\Delta_{ebsb} \to \delta$ as $e_x \to 0$ for any small $\delta > 0$

<u>Conservatism vanishes as controller adapts!</u>

Theorem (Informal): Safety Under EBSB

Let $\|\nabla h(x)\| \leq \kappa \; \forall x$ for a constant $\kappa > 0$, S bounded, $x_p(t_0), x_m(t_0) \in S$, and $\dot{h}(x_m\, r) \geq -\gamma h(x_m) + \Delta_{ebsb}(x_m, e_x)$.

Then:

1. $x_p(t) \in S \; \forall t \geq t_0$,
2. all signals of the closed-loop adaptive system are bounded,
3. $\lim_{t \to \infty} e_x = 0$, and
4. $\lim_{t \to \infty} \Delta_{ebsb} = \delta$.

Note: EBSB requires no input uncertainty (i.e. Λ known). For Λ unknown, need more sophistication…

Further details: P.A. Fisher, J. Autenrieb, and A.M. Annaswamy. *An Error-Based Safety Buffer for Safe Adaptive Control (Extended Version)*, arXiv, 2026

Formal Safety Guarantees with an Error-Based Safety Filter (EBSF)

Idea:

- When x_p near ∂S, choose r so that $\dot{h}(x_p, r) \geq 0$ for worst-case θ_x^*, Λ.
- When x_p far from ∂S, choose r as though adaptation was complete.
- For x_p in between, interpolate constraints.
- Width of interpolation window $\to 0$ as $e_x \to 0$.

Theorem (Informal): Safety Under EBSF

Let S be bounded, $x_p(t_0) \in S$, and . Then:

1. $x_p(t) \in S \;\forall t \geq t_0$,
2. all signals of the closed-loop adaptive system are bounded,
3. $\lim_{t\to\infty} e_x = 0$, and
4. $\lim_{t\to\infty} \beta_{ebsf} = 0$.

$$\min_{r \epsilon \mathcal{R}} (r - r^*)^2$$

s.t.

$$\beta_{ebsf}\, z_p(x,r) + (1 - \beta_{ebsf})\, z_m(x,r) \geq -\gamma h(x) + \delta$$

$$\frac{\partial h}{\partial x}\left[A_m x + B_m r\right] \geq -\gamma h(x) + \delta$$

$$\min_{\theta_x \in \Theta_x,\, \Lambda \in L} \frac{\partial h}{\partial x}\left[B_p \Lambda\left((\theta_x^* - \hat{\theta}_x)x + \theta_r r\right)\right] \geq -\gamma h(x) + \delta$$

$z_m(x,r)$ $z_p(x,r)$

Design $\beta_{ebsf}(x, e_x)$ so that $\beta_{ebsf} \to 0$ as $e_x \to 0$

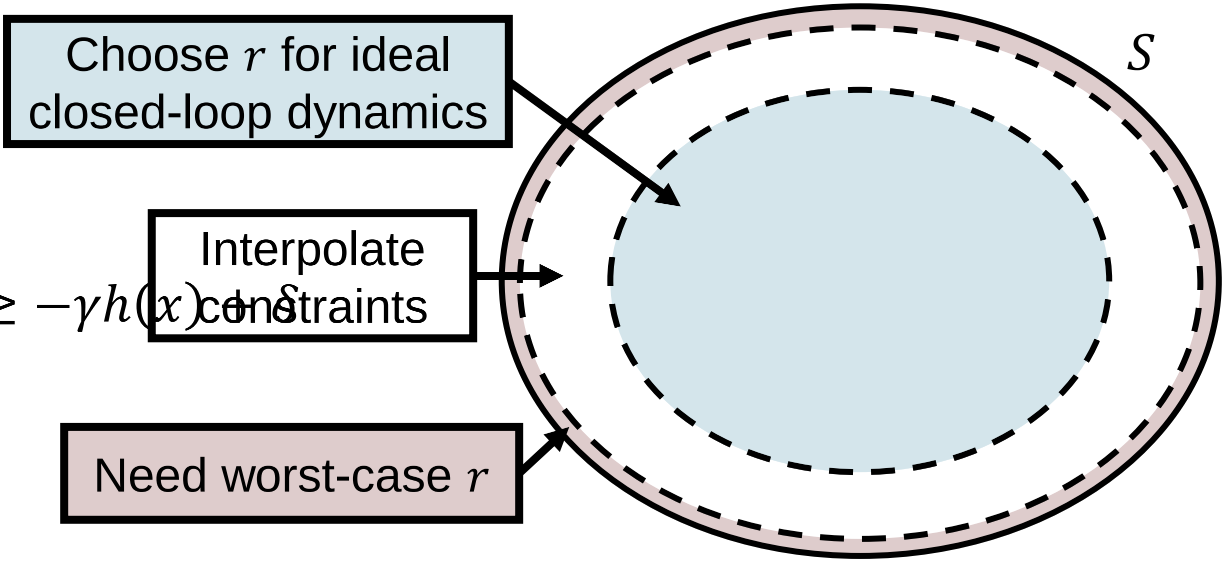

Example: Simple Academic System

- Simple academic system: planar point mass tethered by a spring/damper
  - Mass m and coeffs k, b unknown
  - Control: track desired trajectory $(x_*(t), y_*(t))$
  - Safety: avoid pillar centered at (x_p, y_p)

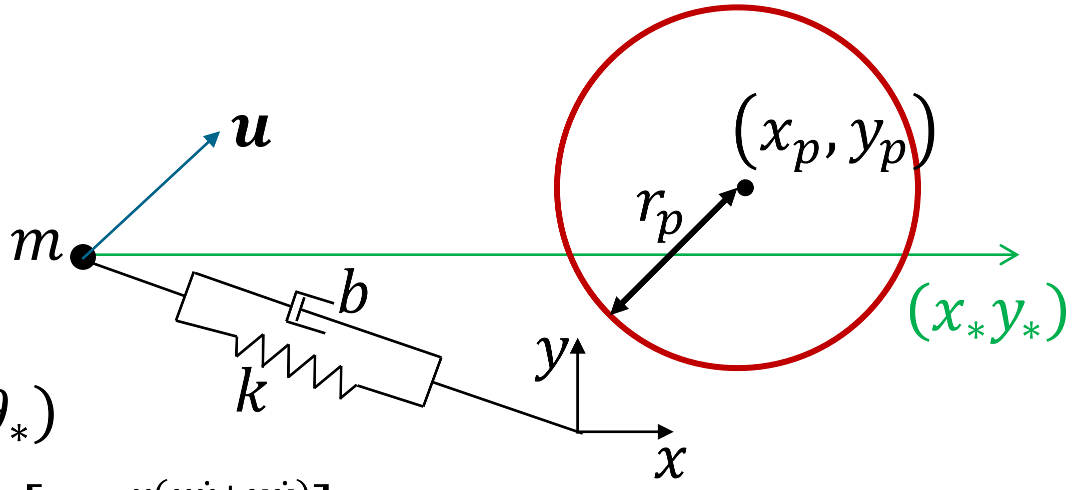

- Plant (nonlinear): $\dot{x} = Ax + B(\lambda_* u - F(x)\theta_*)$
  - $x = \begin{bmatrix} x \\ y \\ \dot{x} \\ \dot{y} \end{bmatrix}$, $A = \begin{bmatrix} 0_{2\times2} & I_2 \\ 0_{2\times2} & 0_{2\times2} \end{bmatrix}$, $B = \begin{bmatrix} 0_{2\times2} \\ I_2 \end{bmatrix}$, $F(x) = \begin{bmatrix} x & \frac{x(x\dot{x}+y\dot{y})}{x^2+y^2} \\ y & \frac{y(x\dot{x}+y\dot{y})}{x^2+y^2} \end{bmatrix}$, $\theta_* = \begin{bmatrix} k/m \\ b/m \end{bmatrix}$, $\lambda_* = \frac{1}{m}$
- Safe set: $S = \left\{ x : h(x) = \sqrt{(x - x_p)^2 + (y - y_p)^2} - r_p \geq 0 \right\}$
  - Relative degree 2 → employ backstepping and high-order CBFs (see paper for details)

Example: Simple Academic System (Cont.)

EBSF (m Unknown): State Trajectories

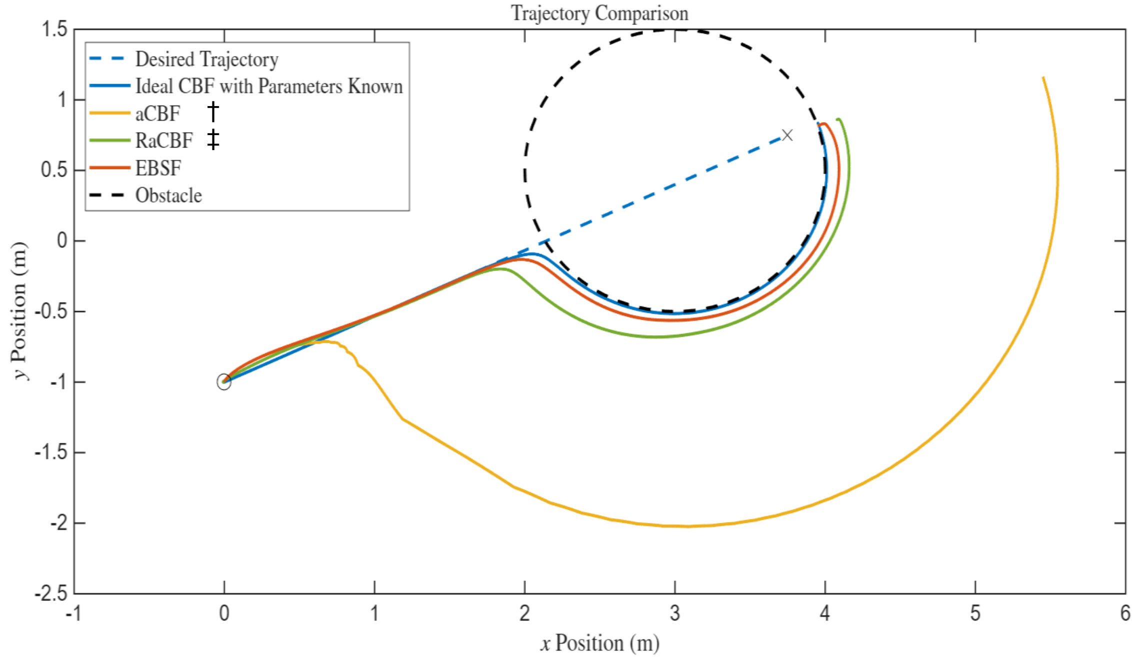

EBSF (m Unknown): Safe Reference Inputs $r(t)$

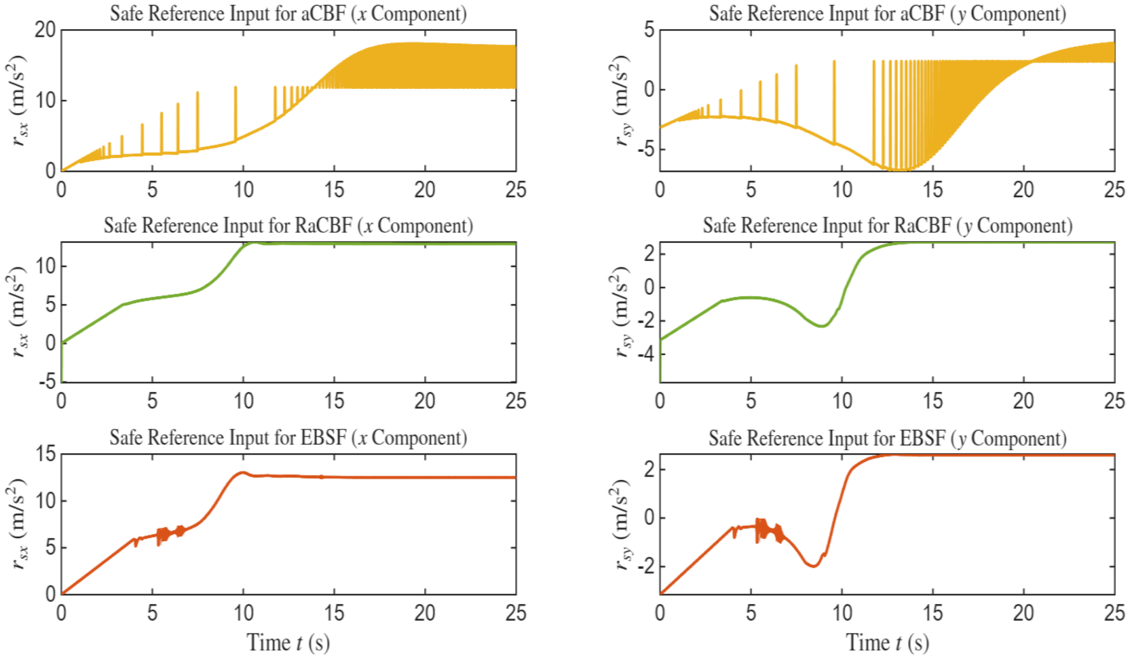

† A. Taylor and A. Ames. *Adaptive safety with control barrier functions*. In 2020 American Control Conference (ACC), 2020.
‡ B. Lopez, J.-J. Slotine, and J. How. *Robust adaptive control barrier functions: An adaptive and data-driven approach to safety*. IEEE Control Systems Letters, 2021.

Further details: P.A. Fisher, J. Autenrieb, and A.M. Annaswamy. *An Error-Based Safety Buffer for Safe Adaptive Control (Extended Version),*

Robust Safety Filters for Lipschitz-Bounded Adaptive Closed-Loop Systems

Idea: Known bounds on the closed-loop dynamics and uncertainties enable a robust adaptive safety condition that accounts for transient plant–reference mismatch.

We assume the following holds:

There exist known constants $\bar{\Theta}_x, \bar{\Theta}_r, \bar{\Lambda} > 0$ such that,

$$|\tilde{\Theta}_x| \leq \bar{\Theta}_x, \qquad |\tilde{\Theta}_x| \leq \bar{\Theta}_r, \qquad |\Lambda| \leq \bar{\Lambda}.$$

There exist constants $L_1, L_2 > 0$ such that,

$$|e_h(t)| = |h(x_p(t)) - h(x_m(t))| \leq L_1 |e_x|,$$
$$|\dot{e}_h(t)| = |\dot{h}(x_p(t)) - \dot{h}(x_m(t))| \leq L_2 |\dot{e}_x|.$$

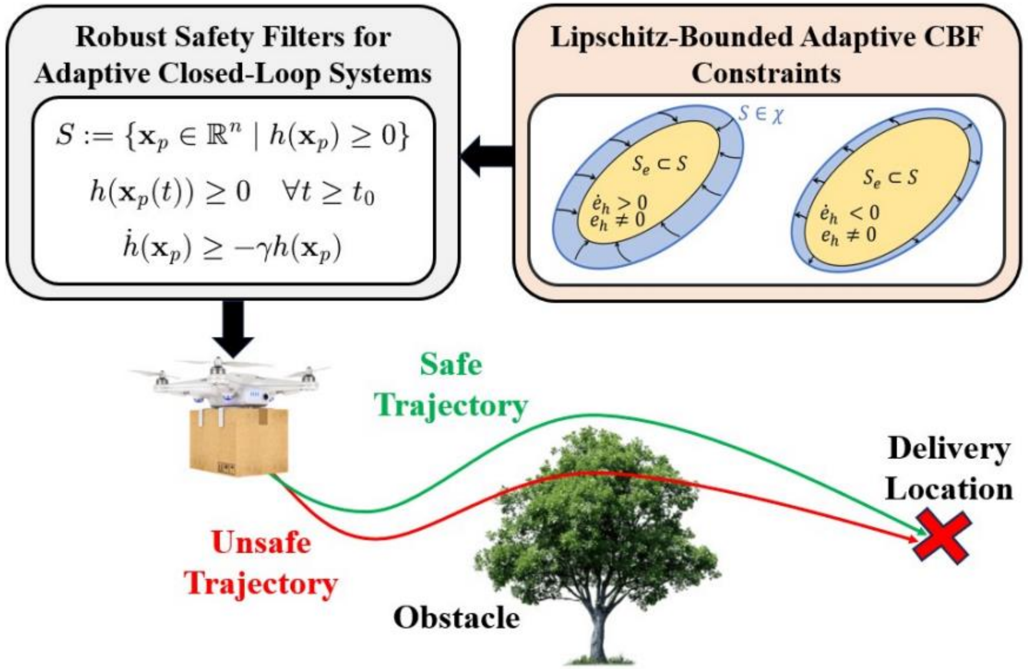

Further details: J J. Autenrieb, P. A. Fisher, A. M. Annaswamy, Robust Safety Filters for Lipschitz-Bounded Adaptive Closed-Loop Systems with Structured Uncertainties

Robust Safety Filters for Lipschitz-Bounded Adaptive Closed-Loop Systems (cont.)

Simplified Definition: Safety Constraint for Lipschitz-Bounded Adaptive Closed-Loop Systems

The reference signal $r(t)$ keeps the plant state $x_p(t)$ in S as long as the following inequality holds:

$$\begin{aligned}\dot{h}(x_m(t)) \geq\ & -\gamma h(x_m(t)) + \gamma L_1 \|e_x(t)\| + L_2 \|A_m\| \|e_x(t)\| \\ & + L_2 \|B_p\| \overline{\Lambda} \left(\overline{\Theta}_x \|x_p(t)\| + \overline{\Theta}_r \|r(t)\| \right)\end{aligned}$$

for all $t \geq t_0$ and $x_p(t_0) \epsilon S$.

Further details: J. Autenrieb, P. A. Fisher, A. Annaswamy, Robust Safety Filters for Lipschitz-Bounded Adaptive Closed-Loop Systems with Structured Uncertainties (Submitted to CDC)

Robust Safety Filters for Lipschitz-Bounded Adaptive Closed-Loop Systems (cont.)

However, the structure of the problem leads to a second-order cone program that is defined as:

$$\begin{aligned} \min_{r(t),\eta(t),v(t)} \quad & v(t) + \rho\,\eta(t) \\ \text{s.t.} \quad & \|r(t) - r^{\star}(t)\| \leq v(t) \\ & \|r(t)\| \leq \eta(t) \\ & -a(t)^{\top} r(t) - c\,\eta(t) \leq -\beta(t) \end{aligned}$$

where

$$a(t) := \left(\left.\frac{\partial h}{\partial x}\right|_{x_m(t)} B_m \right)^{\top}$$

$$c := L_2 \|B_p\|\, \overline{\Lambda}\, \overline{\Theta}_r$$

$$\beta(t) := \left.\frac{\partial h}{\partial x}\right|_{x_m(t)} A_m x_m(t) + \gamma h(x_m(t)) - \Delta(x_p(t), e_x(t))$$

$$\Delta(x_p(t), e_x(t)) := \gamma L_1 \|e_x(t)\| + L_2 \|A_m\|\, \|e_x(t)\| + L_2 \|B_p\|\, \overline{\Lambda}\, \overline{\Theta}_x\, \|x_p(t)\|$$

Robust Safety Filters for Lipschitz-Bounded Adaptive Closed-Loop Systems (cont.)

Considered reference model:

$$\dot{x}_m(t) = A_m x_m(t) + B_m r(t)$$

Adaptive laws to enforce the desired closed-loop dynamics:

$$\dot{\hat{\theta}}_x(t) = -\Gamma_x x_p(t) e_x^T(t) P B_P$$

$$\dot{\hat{\theta}}_r(t) = -\Gamma_r r(t) e_x^T(t) P B_P$$

Integrated control law:

$$u(t) = \hat{\theta}_x(t)\, x(t) + \hat{\theta}_r(t)\, r(t)$$

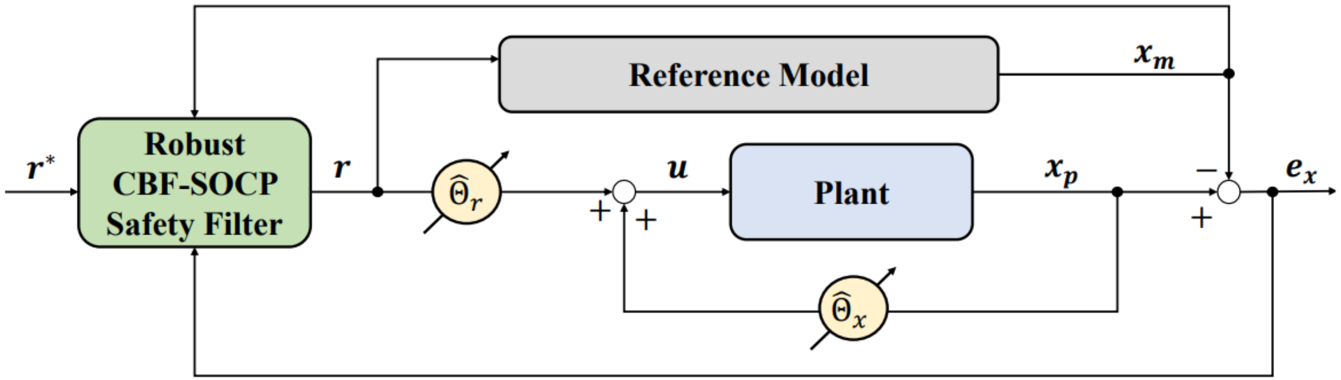

Example: Quadcopter Collision Avoidance

Adaptive Control + Standard QP-CBF

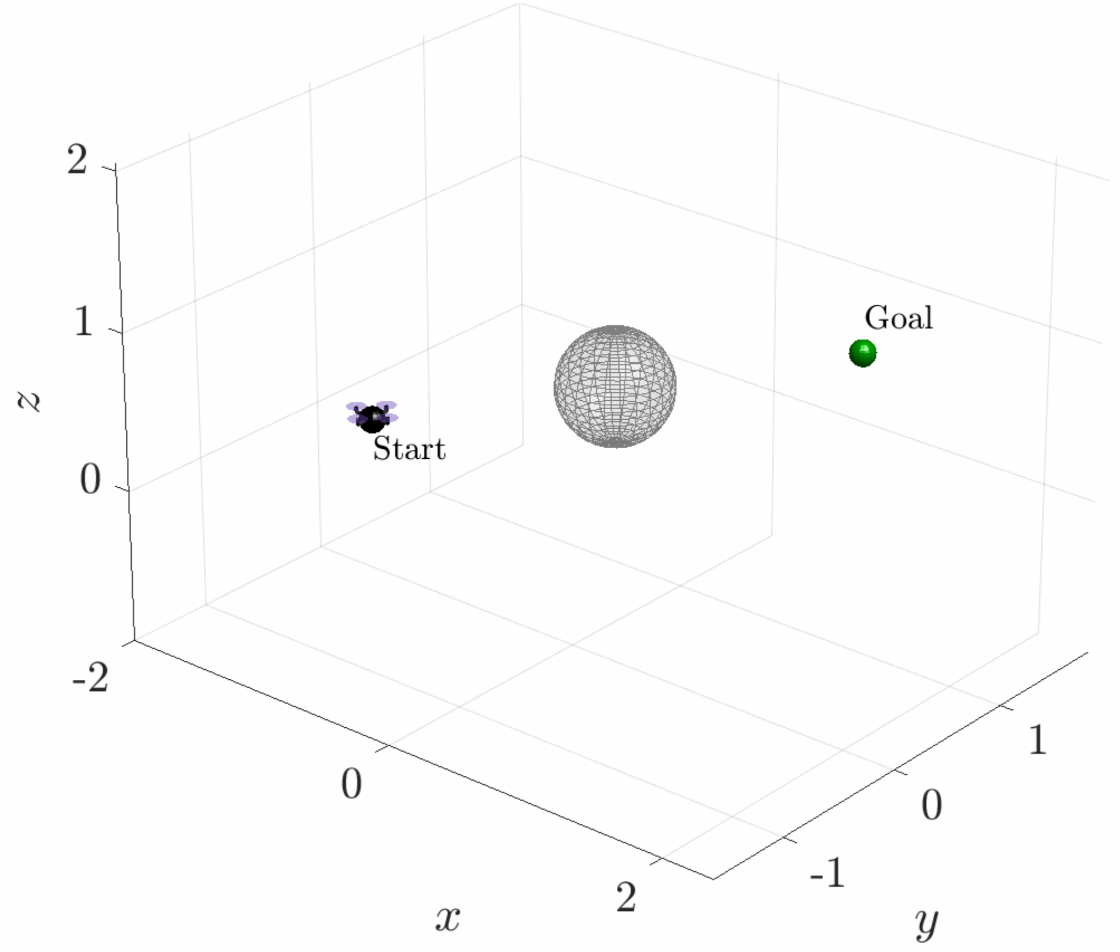

Adaptive Control + Robust SOCP-QP

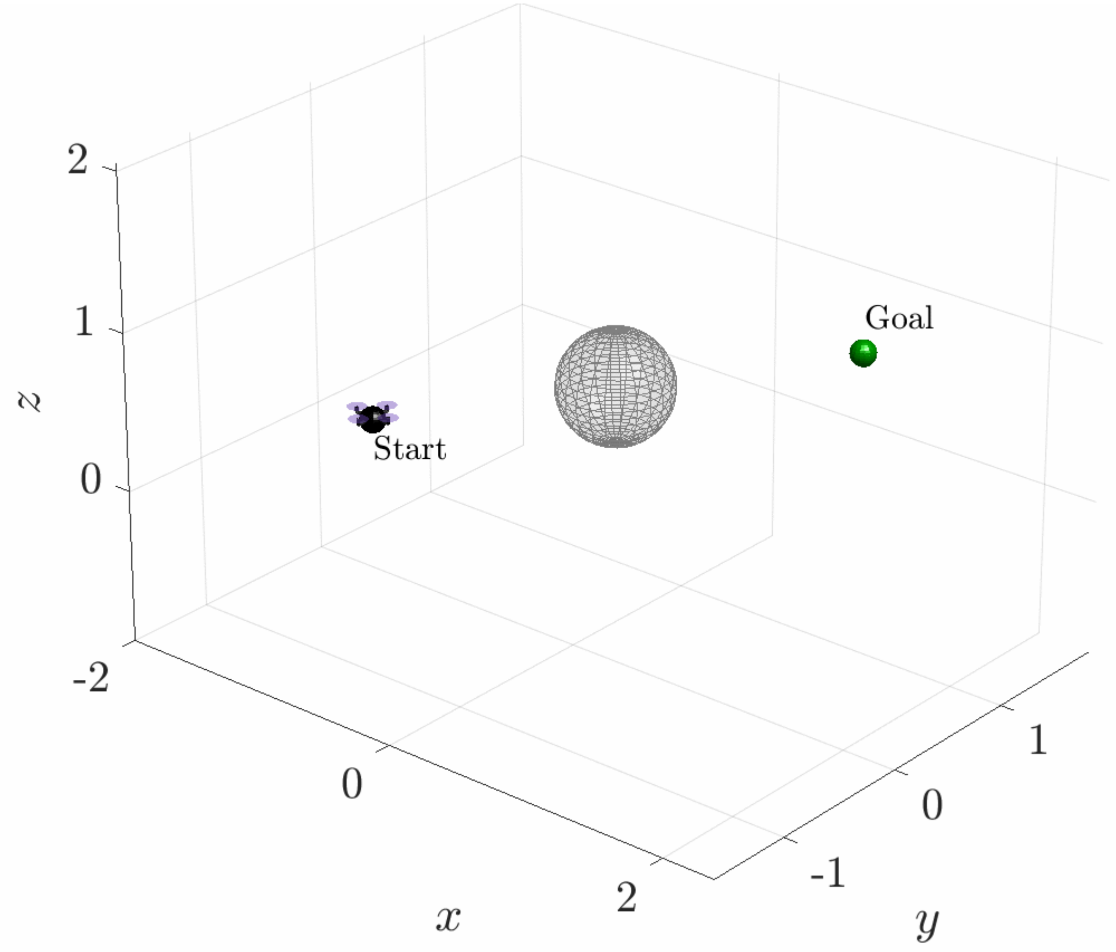

What are the key takeaways from the case studies?

<u>**Intermediate Summary**</u>**:**

- The integration of adaptive control with standard Control Barrier Function (CBF) approaches is not straightforward, since CBFs rely on accurate model knowledge, while adaptive systems exhibit transient model mismatch.
- This mismatch can lead to loss of safety guarantees, especially during the adaptation phase, where the plant and reference model may significantly differ.
- To address these challenges, modified CBF formulations are required that explicitly account for uncertainties, model mismatches, and transient deviations in the adaptive closed-loop system.

Any Question so far?

Outline

1 Concept of Control Barrier Functions (CBFs)

2 **Safety Filter Design for Adaptive Closed-Loop Systems**

3 Outlook & Conclusions

Outline

1 Concept of Control Barrier Functions (CBFs)

2 Safety Filter Design for Adaptive Closed-Loop Systems

3 Outlook & Conclusions

Outlook: Extension to multi-agent systems

Challenge:

- Agents operate in spatial proximity
- Each agent pursues individual mission objective → Trajectories could conflict!
- Only partial and delayed information of neighboring are available

Additional problem:

- Each action to ensure safety for one agent has implications on other agents
- Uncertainties on the dynamics are coupled → Finding guarantees is much harder.

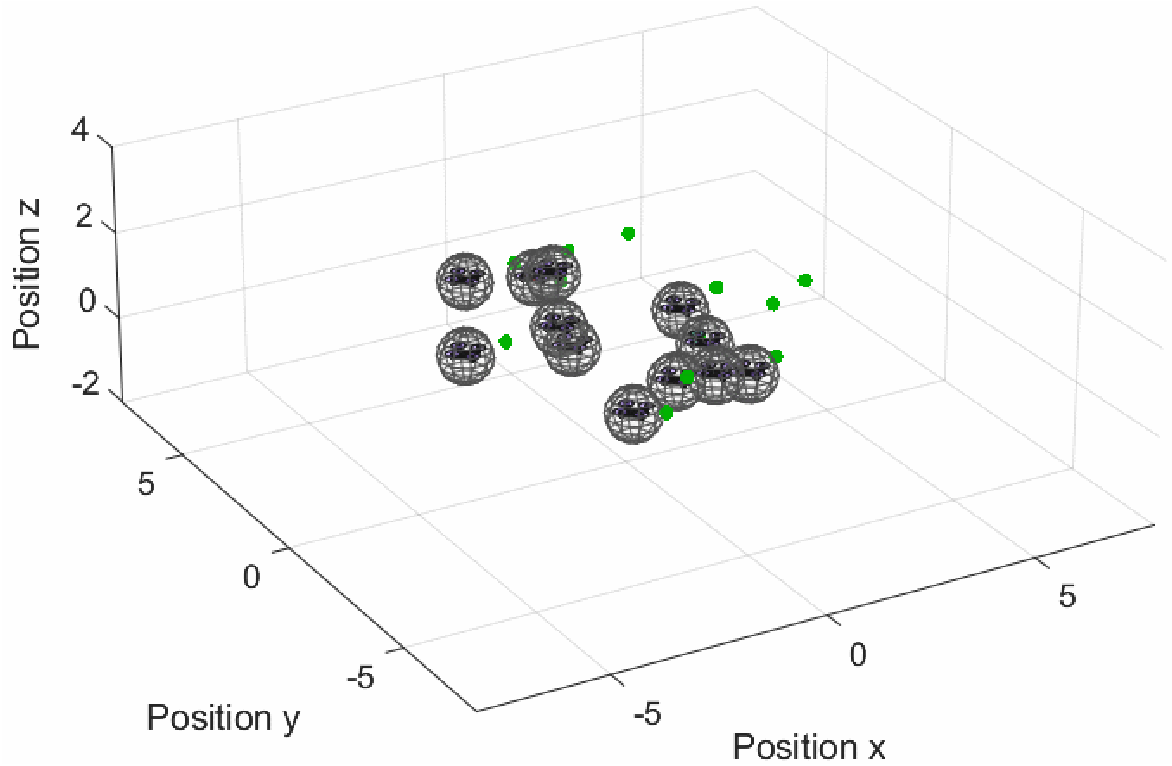

Outlook: Further applications to aerospace systems

Application of methods to different domains:

- Hypersonic systems, drones, aircraft and cooperative systems

The algorithms need stepwise real-world validation:

- First campaigns in DLR flight simulators
- Afterward real flight tests with DLR aircraft fleet

Source: DLR

Conclusions

- Control Barrier Functions (CBFs) provide a powerful framework to enforce state constraints and safety guarantees in aerospace systems.
- A key advantage is the separation between performance and safety: Performance controllers can be designed using conventional techniques, while safety filters ensure constraint satisfaction.
- However, the integration with advanced control strategies, such as adaptive control, is not straightforward, as model uncertainties and transient deviations can compromise safety guarantees.
- Recent research addresses these challenges through extended CBF formulations that explicitly account for uncertainties, enabling safe and stable closed-loop behavior.
- This proposed methods opens up promising applications in safety-critical aerospace systems, with significant potential for future research.

Impactful and helpful references on CBFs

[1] M. Nagumo, Über die lage der integralkurven gewöhnlicher differentialgleichungen, Proceedings of the Physico-Mathematical Society of Japan, 1942

[2] P. Wieland and F. Allgöwer, Constructive Safety Using Control barrier functions, 7th IFAC Symposium on Nonlinear Control Systems, 2007

[3] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, Control barrier function based quadratic programs for safety critical systems, IEEE Transactions on Automatic Control, 2017

[4] J. Zeng, B. Zhang, and K. Sreenath, "Safety-Critical Model Predictive Control with Discrete-Time Control Barrier Function," American Control Conference (ACC), 2021

[5] W. Xiao and C. Belta, "High-Order Control Barrier Functions," IEEE Transactions on Automatic Control, 2022

[6] M. Spiller , E. Isbono, and P. Schitz, Feasibility of multiple robust control barrier functions for bounding box constraints, American Control Conference (ACC), 2025

[7] Kumpati S. Narendra & A. M. Annaswamy, Stable Adaptive Systems

Checkout our work on CBFs

[8] J. Autenrieb, *A Quadratic Programming Approach to Flight Envelope Protection using Control Barrier Functions (in print)*, AIAA Journal of Guidance, Control, and Dynamics, 2025

[9] J. Autenrieb and A. M. Annaswamy, Safe and stable adaptive control for a class of dynamic systems, IEEE Conference on Decision and Control (CDC), 2023

[10] J. Autenrieb and A. M. Annaswamy, *Safe and stable adaptive control for a class of dynamic systems*, IEEE Conference on Decision and Control (CDC), 2023

[11] J.A. Solano-Castellanos, P.A. Fisher, and A. M. Annaswamy, Safe and Stable Formation Control with Autonomous Multi-Agents Using Adaptive Control, Conference on Modeling, Estimation and Control MECC 2025

[12] P. A. Fisher, J. Autenrieb. A. M. Annaswamy, An Error-Based Safety Buffer for Safe Adaptive Control, arXiv:2510.23491

[13] J. Autenrieb, P. A. Fisher, A. M. Annaswamy, Robust Safety Filters for Lipschitz-Bounded Adaptive Closed-Loop Systems with Structured Uncertainties, arXiv:2603.14403

Thank you for your attention!

Johannes Autenrieb, DLR

johannes.autenrieb@dlr.de

Peter A. Fisher, MIT

pafisher@mit.edu

Anuradha M. Annaswamy, MIT

aanna@mit.edu

Acknowledgment:

Many thanks to our colleagues and collaborators.

Jose A. Solano-Castellanos, MIT

jsolanoc@mit.edu

Mark Spiller, DLR

mark.spiller@dlr.de