Control Barrier Functions in Sampled-Data Systems

60th IEEE Conference on Decision and Control
Austin, Texas, United States, December 17th 2021

Joseph Breeden, Kunal Garg, Dimitra Panagou

Department of Aerospace Engineering

University of Michigan, Ann Arbor, MI, USA

Background - Safety Critical Control

- System state should always lie within a predefined "safe set"

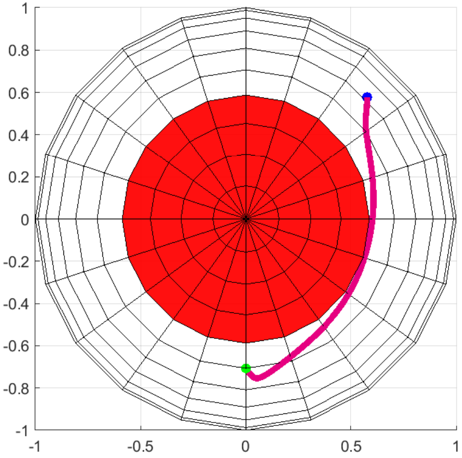

- [Ames et al., ECC 2019] A $\mathcal{C}^1$ function $h : \mathbb{R}^n \to \mathbb{R}$ is a Control Barrier Function (CBF) for the set $S \triangleq \{x \in \mathbb{R}^n \mid h(x) = 0\}$ under the control set U if there exists $\alpha \in \mathcal{K}$ such that

$$\inf_{u \in U} \dot{h}(x, u) \leq \alpha(-h(x)), \ \forall x \in S\,.$$

- S is called the "safe set"
- The pointwise condition

$$\dot{h}(x, u) \leq \alpha(-h(x)), \ \forall x \in S \tag{1}$$

is sufficient to guarantee forward invariance of S, so that the system remains "safe" for all time ("continuous-time CBF condition")

  - This condition is enforced online

- Method for continuous systems:
  - Fix $\alpha \in \mathcal{K}$, measure $x(t)$ continuously, and enforce $\dot{h}(x(t), u(t)) \leq \alpha(-h(x(t)))$ for all times t
- Method for sampled-data systems:
  - Fix $\alpha \in \mathcal{K}$, measure $x_k = x(t_k)$ at t_k, and enforce $\dot{h}(x_k, u_k) \leq \alpha(-h(x_k)) - \nu_0^g$ for all samples k
  - Choose ν_0^g such that the above condition is sufficient to ensure $\dot{h}(x(t), u_k) \leq \alpha(-h(x(t)))$ for all $t \in [t_k, t_{k+1}]$.

Contribution

- Existing sampled-data CBF approaches are over-conservative
- Contributions
  - Two metrics of conservatism
  - Three methods of decreasing conservatism in ν_0^g while still guaranteeing safety
  - Distinction between local and global techniques for choosing margins

	Global	Local
Prior Work	ϕ_0^g	
Method 1	ϕ_1^g	ϕ_1^l
Method 2	ϕ_2^g	ϕ_2^l
Method 3	ϕ_3^g	ϕ_3^l

Problem Formulation

- System $\dot{x} = f(x) + g(x)u,\ x \in \mathbb{R}^n, u \in U \subset \mathbb{R}^m$ for compact U and f, g locally Lipschitz continuous
- System sampling at t_k, where $t_{k+1} - t_k = T$ for fixed time-step T
- $u(t) = u_k, \forall t \in [t_k, t_{k+1})$

Problem 1

Given a $\mathcal{C}^1$ function $h : \mathbb{R}^n \to \mathbb{R}$ with locally Lipschitz derivatives and the set $S = \{x \in \mathbb{R}^n \mid h(x) \leq 0\}$, design a function $\phi : \mathbb{R}_{>0} \times \mathbb{R}^n \to$ such that the condition

$$\dot{h}(x_k, u_k) = L_f h(x_k) + L_g h(x_k) u_k \leq \phi(T, x_k), \forall k \in \mathbb{N} \tag{2}$$

is sufficient to guarantee forward invariance of S.

Defining Metrics of Conservatism

- All ϕ are of the form $\phi(T, x) = \alpha(-h(x)) - \nu(T, x)$

Definition 1

The function $\nu : \mathbb{R}_{>0} \times \mathbb{R}^n \to \mathbb{R}$ is the *controller margin*.

- Controller margin is instantaneous margin
- Large ν leads to large u_k to satisfy $\dot{h}(x_k, u_k) \leq \phi(T, x_k)$ in (2)

Defining Metrics of Conservatism

Definition 2

The *physical margin* is a function $\delta : \mathbb{R}_{>0} \to \mathbb{R}$ defined as

$$\delta(T) = \sup_{\substack{x \in S \\ \phi(T,x)=0}} -h(x)$$

- $\phi(T, x_k)$ in (2) is potentially negative at states x where $h(x) > -\delta$ (unlike with (1))
- This makes states $x \in S_\delta \triangleq \{x \in \mathbb{R}^n \mid -\delta \leq h(x) \leq 0\}$ potentially inaccessible

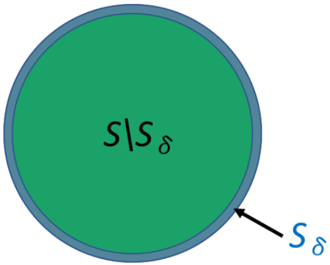

Prior Work

- Define $u_{\max} \triangleq \max_{u \in U} \|u\|$

Lemma 1 (Thm. 2 in Cortez, Oetomo, Manzie, Choong, TCST 2021)

Let the set S be compact and $\alpha \in \mathcal{K}$ be locally Lipschitz continuous. Let $l_{L_f h}, l_{L_g h}, l_{\alpha(h)}$ be the Lipschitz constants of $L_f h, L_g h, \alpha(-h)$, respectively. Then the function $\phi_0^g : \mathbb{R}_{>0} \times \mathbb{R}^n$, defined as

$$\phi_0^g(T,x) \triangleq \alpha(-h(x)) - \underbrace{\frac{l_1 \Delta}{l_2}\left(e^{l_2 T} - 1\right)}_{\nu_0^g(T)},$$

solves Problem 1, where $l_1 = l_{L_f h} + l_{L_g h} u_{max} + l_{\alpha(h)}, l_2 = l_{L_f h} + l_{L_g h} u_{max}$, and $\Delta = \sup_{x \in S, u \in U} \|f(x) + g(x)u\|$.

(similar approaches in [Singletary, Chen, Ames, CDC 2020] & [Usevitch, Panagou, ACC 2021])

Corollary 1 (to Theorem 1)

Under the assumptions of Lemma 1, and with l_1, Δ *as in Lemma 1, the function* $\phi_1^g : \mathbb{R}_{>0} \times \mathbb{R}^n$, *defined as*

$$\phi_1^g(T,x) \triangleq \alpha(-h(x)) - \underbrace{l_1 T \Delta}_{\nu_1^g(T)},$$

solves Problem 1. Furthermore, for the same α, *it holds that* $\nu_1^g(T) < \nu_0^g(T), \forall T \in \mathbb{R}_{>0}$.

- Linear in T rather than exponential

Method 1 Local Margin

- Define $\mathcal{R}(x,T)$ as the set of all states $x(t)$ reachable from $x = x(\tau)$ in times $t \in [\tau, \tau + T]$ OR an over-approximation of this set
  - $\mathcal{R}(x,T)$ is bounded because f, g locally Lipschitz and U compact

Theorem 1

Let $\alpha \in \mathcal{K}$ be locally Lipschitz. Let $l_{L_f h}(x), l_{L_g h}(x), l_{\alpha(h)}(x)$ be the Lipschitz constants of $L_f h, L_g h, \alpha(-h)$ over the set $\mathcal{R}(x,T)$, respectively. Then the function $\phi_1^l : \mathbb{R}_{>0} \times \mathbb{R}^n$, defined as

$$\phi_1^l(T,x) \triangleq \alpha(-h(x)) - \underbrace{l_1(x) T \Delta(x)}_{\nu_1^l(T,x)},$$

solves Problem 1, where $l_1(x) = l_{L_f h}(x) + l_{L_g h}(x) u_{max} + l_{\alpha(h)}(x)$, and $\Delta(x) = \sup_{z \in \mathcal{R}(x,T), u \in U} ||f(z) + g(z)u||$. Furthermore, for the same α, it holds that

$$\nu_1^l(T,x) \leq \nu_1^g(T) < \nu_0^g(T), \forall x \in S, \forall T \in \mathbb{R}_{>0}.$$

Method 3

- Define $\psi(x,u) \triangleq \nabla[\dot{h}(x)]\,(f(x)+g(x)u)$ (second derivative of h)
- Define

$$\eta(T,x) \triangleq \max\left\{\left(\sup_{z\in\mathcal{R}(x,T)\setminus\mathcal{Z},u\in U} \psi(z,u)\right),0\right\}$$

where $\mathcal{Z}$ is any set of Lebesgue measure zero (to account for CBFs that are not twice differentiable everywhere).

Method 3

Theorem 3

The function $\phi_3^l : \mathbb{R}_{>0} \times \mathbb{R}^n$, *defined as*

$$\phi_3^l(T,x) \triangleq -\frac{\gamma}{T}h(x) - \underbrace{\frac{1}{2}T\eta(T,x)}_{\nu_3^l(T,x)}$$

solves Problem 1, for any $\gamma \in (0,1]$.

- γ controls rate of convergence to boundary of S, similar to α in (1)

Remarks on Method 3

- Comparison between Method 1 and Method 3:

Theorem 4

The controller margins for ϕ_3^l, ϕ_3^g *and* ϕ_1^l, ϕ_1^g *satisfy* $\nu_3^l(T,x) \le \frac{1}{2}\nu_1^l(T,x)$ *and* $\nu_3^g(T) \le \frac{1}{2}\nu_1^g(T)$, $\forall x \in S, \forall T \in \mathbb{R}_{>0}$.

- The physical margin for Methods 1, 2 decrease linearly in T
- The physical margin for Method 3 decreases quadratically in T. This occurs because ν_3^l, ν_3^g do not depend on α (in this case, $\alpha(\lambda) = \frac{\gamma}{T}\lambda$)

Unicycle:

$$\dot{x}_1 = u_1 \cos(x_3),\ \dot{x}_2 = u_1 \sin(x_3),\ \dot{x}_3 = u_2,$$

$$h = \rho - \sqrt{x_1^2 + x_2^2 - (\mathrm{wrap}_\pi(x_3 - \sigma \arctan 2(x_2, x_1)))^2}\,,$$

where ρ is the radius to be avoided, and σ is a shape parameter.

Spacecraft attitude:

$$\dot{p} = \omega \times p,\ \dot{\omega} = u,$$

$$h_1 = s \cdot p - \cos(\theta) + \mu(s \cdot (\omega \times p))\left|s \cdot (\omega \times p)\right|,$$

$$h_2 = ||\omega||_\infty - w_{\max}$$

where $s \in \mathbb{R}^3$, $||s|| = 1$, is a constant vector pointing to an object to be avoided, θ is the smallest allowable angle, and μ is a shape parameter.

Simulations - Physical Margins

	Unicycle			Spacecraft		
T	0.1	0.01	0.001	0.1	0.01	0.001
$\delta_0^{g,\inf}$	$1.2(10)^{42}$	420	0.010	9.8	0.23	0.021
$\delta_1^{g,\inf}$	0.54	0.054	0.0054	2.0	0.20	0.020
$\delta_2^{g,\inf}$	0.53	0.053	0.0053	0.81	0.082	0.0082
$\delta_3^{g,\inf}$	0.013	$1.3(10)^{-4}$	$1.3(10)^{-6}$	0.013	$1.3(10)^{-4}$	$1.3(10)^{-6}$

Table: Global physical margins for selected time-steps T

- Based on the above, we expect the results under Method 3 to approach much closer to the edge of the safe set then the other methods
- All subsequent simulations used $T = 0.1$

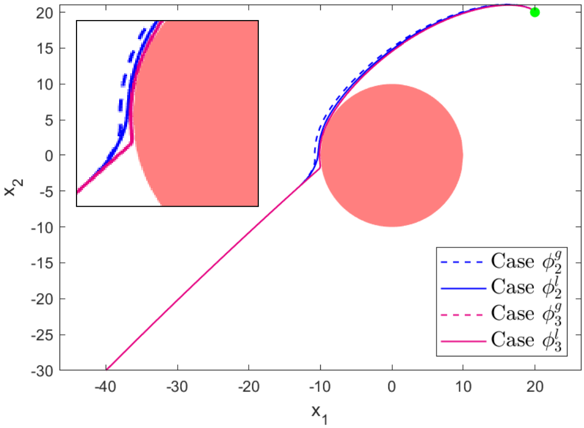

Figure: Trajectories of the unicycle system

- The trajectories under $\phi_0^g, \phi_1^g, \phi_1^l$ immediately turned away from the green target

Simulated controller margins

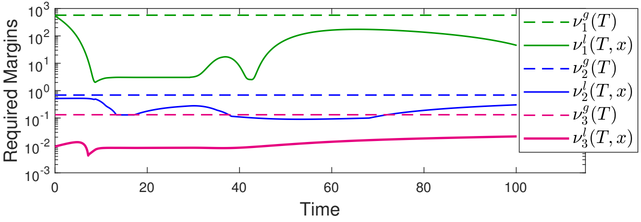

Figure: Controller margins for the unicycle system (ν_0^g omitted because it is larger than 10^{40})

Simulated physical margins

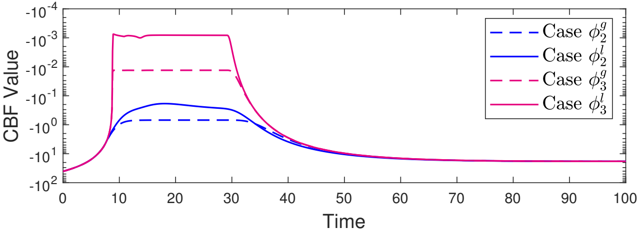

Figure: CBF values along the 4 unicycle trajectories

Simulations - Unicycle

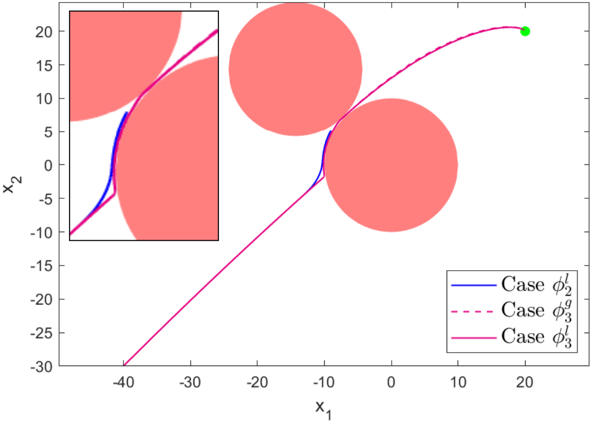

Figure: A simulation with two tightly-spaced obstacles, in which controllers using margins ϕ_3^l and ϕ_3^g permit passage through the obstacles, while the other functions force the agent to stop.

Simulations - Spacecraft

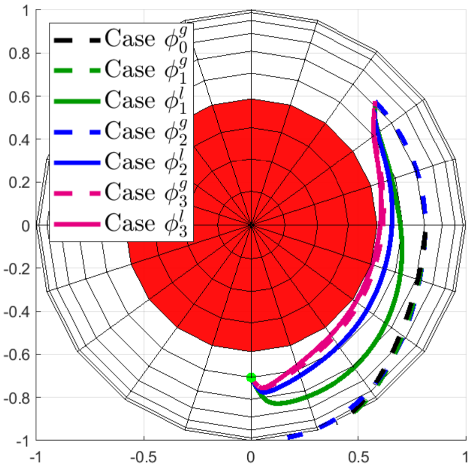

Figure: Trajectories of the spacecraft attitude system for all 7 margin functions

Conclusion

- To better approximate continuous results and improve performance, we have introduced three new ways of generating sufficient sampled-data margins when using CBFs
- Controller sampling should be taken into account when implementing a safety-critical system
- The above two systems could not be provably controlled with CBFs at $T = 0.1$ using results from prior literature
- Future work
  - Adaptively adjusting margins to further decrease conservatism
  - Input constraints $+$ sampled-data CBFs (see upcoming AIAA paper)

Acknowledgements

The authors would like to acknowledge the support of the U.S. National Science Foundation and Air Force Office of Scientific Research

Control Barrier Functions in Sampled-Data Systems

60th IEEE Conference on Decision and Control
Austin, Texas, United States, December 17th 2021

Joseph Breeden, Kunal Garg, Dimitra Panagou

Department of Aerospace Engineering

University of Michigan, Ann Arbor, MI, USA

UNIVERSITY OF MICHIGAN