Convexification of Motion Planning with Temporal and Logical Specifications and Control through Liftings and Hypercomplex Numbers

Behçet Açıkmeşe (Autonomous Control Laboratory)
Mehran Mesbahi (Robotics, Information, and Networks Laboratory)

William E. Boeing Department of Aeronautics & Astronautics
University of Washington

2024 AFOSR Dynamical Systems and Control Theory

Program Review

August 28, 2024

in collaboration with ...

  

 

from left to right: Samet Uzun, Taewan Kim, Spencer Kraisler, Behcet Acikmese, and Mehran Mesbahi

Outline

underlying theme: role of parameterization–and the corresponding geometry– in developing efficient algorithms for trajectory optimization and control

We cover a subset of our contributions over the past year, including:

1. *high level decision-making/guidance:* Optimization with temporal and logical specifications/successive convexification/prox-linear algorithms
2. *guidance to control interface:* Constrained funnel synthesis/optimal control on PSD matrices/succesive convexification
3. *feedback control:* Policy optimization on quotient Riemannian manifolds/Riemmanian first order methods induced by system theoretic metrics

each topic covered today corresponds to a distinct level in the general autonomy stack: from higher-level decision making, to funnels, to control

let us start with STL ...

Optimization with Temporal and Logical Specifications

Signal Temporal Logic

- A formal language for modeling logical and temporal specifications
- Effectively used in trajectory generation to specify high-level mission constraints

Examples:

- (Or) Final position of the vehicle should be either point A **or** point B.
- (Implication) **If** the vehicle enters a certain zone, **then** it's speed should not exceed 20 m/s.
- (Eventually) When the vehicle moves from point A to point B, the camera on the vehicle should **eventually** capture an image of a target point.
- (Until) The drone's speed should not exceed 20 m/s **until** it visits one of the battery charging stations.

Optimization with Temporal and Logical Specifications

STL Syntax $\mu := (f(x) \geq 0) \quad \varphi ::= \mu \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \boldsymbol{U}_I \varphi_2$

STL specification		Formula
disjunction	$\varphi_1 \vee \varphi_2$	$\neg(\neg\varphi_1 \wedge \neg\varphi_2)$
implication	$\varphi_1 \implies \varphi_2$	$\neg\varphi_1 \vee \varphi_2$
eventually	$\boldsymbol{F}_I \varphi$	$\top \boldsymbol{U}_I \varphi$
always	$\boldsymbol{G}_I \varphi$	$\neg \boldsymbol{F}_I \neg\varphi$

complete D-SR goes like this ...

$$
\begin{aligned}
&\rho^{\mu}(x,k) := f(x_k)\\
&\rho^{\neg\varphi}(x,k) := -\rho^{\varphi}(x,k)\\
&\rho^{\varphi_1 \wedge \varphi_2}(x,k) := \min((\rho^{\varphi_1}(x,k), \rho^{\varphi_2}(x,k)))\\
&\rho^{\varphi_1 \vee \varphi_2}(x,k) := \max((\rho^{\varphi_1}(x,k), \rho^{\varphi_2}(x,k)))\\
&\rho^{\varphi_1 \implies \varphi_2}(x,k) := \max((-\rho^{\varphi_1}(x,k), \rho^{\varphi_2}(x,k)))\\
&\rho^{\boldsymbol{F}_{[a:b]}\varphi}(x,k) := \max((\rho^{\varphi}(x,k+a), \rho^{\varphi}(x,k+a+1), \ldots, \rho^{\varphi}(x,k+b)))\\
&\rho^{\boldsymbol{G}_{[a:b]}\varphi}(x,k) := \min((\rho^{\varphi}(x,k+a), \rho^{\varphi}(x,k+a+1), \ldots, \rho^{\varphi}(x,k+b)))\\
&\ldots
\end{aligned}
$$

Optimization with Temporal and Logical Specifications

Autonomous Rocket Landing

Speed-triggered constraints: Speed upper-bound, 3 engines → 1 engine

Altitude-triggered constraints: Speed, tilt angle, glideslope angle, gimbal angle, angular velocity

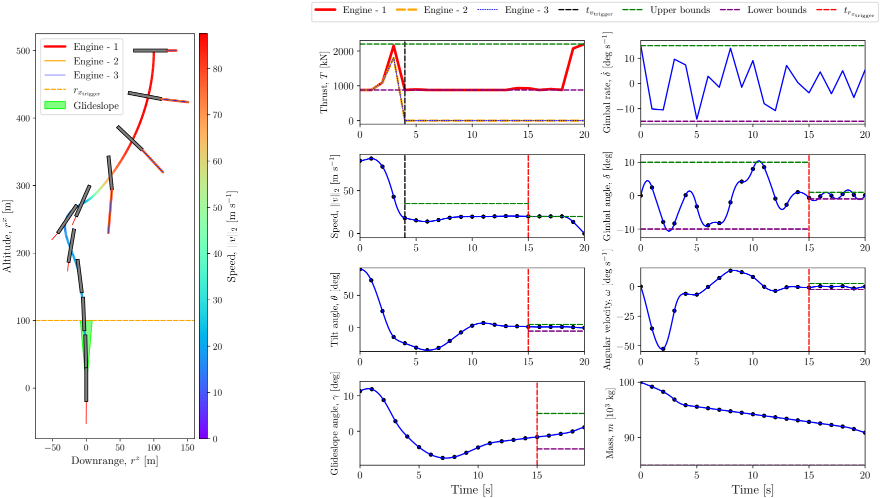

Optimization with Temporal and Logical Specifications

$$\text{maximize} \quad \tilde{\rho}^{\varphi_1 \wedge \varphi_2 \wedge \ldots \wedge \varphi_5}(x, k) \qquad \text{maximize} \quad \tilde{\rho}^{\varphi_1 \vee \varphi_2 \vee \ldots \vee \varphi_5}(x, k)$$

$$\tilde{\rho}^{\varphi_1 \wedge \varphi_2 \wedge \ldots \wedge \varphi_5}(x, k) = \widetilde{\min}_\kappa((f_1(x_k), f_2(x_k), \ldots, f_5(x_k)))$$
$$\tilde{\rho}^{\varphi_1 \vee \varphi_2 \vee \ldots \vee \varphi_5}(x, k) = \widetilde{\max}_\kappa((f_1(x_k), f_2(x_k), \ldots, f_5(x_k)))$$

features that we want to capture ...

- designer's flexibility for including a "degree" of focus on enforcing the spatial/temporal specifications
- preserving soundness, completeness, and monotonicity properties, avoiding locality/masking
- embeddable in smooth optimization, e.g., prox-linear methods for optimization of composite objectives

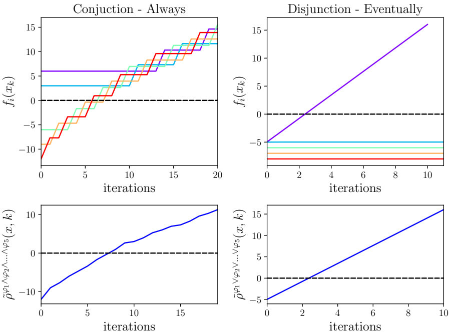

Robust semantic - Generalized Mean based Smooth Robustness (GMSR)

SR	GMSR
min	${}^{\wedge}h^{c}_{p,w}$
max	${}^{\vee}h^{c}_{p,w}$

$$ {}^{\wedge}h^{c}_{p,w}(y) := \left(M^{c}_{0,w}(|y|^{2}_{+})\right)^{\frac{1}{2}} - \left(M^{c}_{p,w}(|y|^{2}_{-})\right)^{\frac{1}{2}}, \text{ where} $$

$$ M^{c}_{0,w}(z) := \left(c^{\mathbf{1}^T w} + \prod_{i=1}^{n} z_i^{w_i}\right)^{1/\mathbf{1}^T w}, $$

$$ M^{c}_{p,w}(z) := \left(c^{p} + \frac{1}{\mathbf{1}^T w} \sum_{i=1}^{n} w_i z_i^{p}\right)^{1/p}, $$

$y \in \mathbb{R}^n$, $z \in \mathbb{R}^n_+$, $c \in \mathbb{R}_{++}$, $p \in \mathbb{Z}_{++}$, $w \in \mathbb{Z}^n_{++}$.

$$ {}^{\vee}h^{c}_{p,w}(y) := -{}^{\wedge}h^{c}_{p,w}(-y) $$

intuition behind GMSR

$(f_1(x) \geq 0)$ **and** $(f_2(x) \geq 0)$ $c = 1e - 8$ and $w = \mathbf{1}$

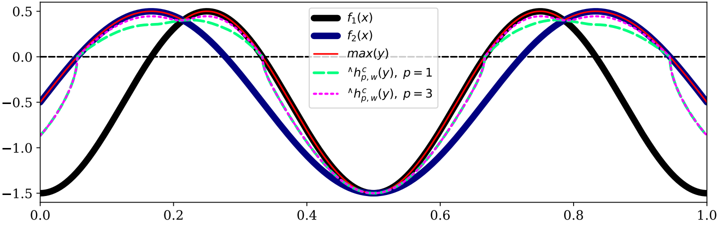

$(f_1(x) \geq 0)$ **or** $(f_2(x) \geq 0)$ $c = 1e - 8$ and $w = \mathbf{1}$

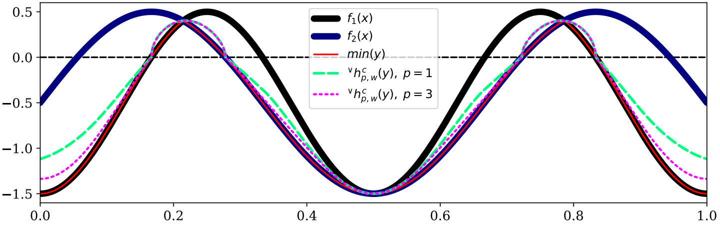

Optimization with Temporal and Logical Specifications

How ${}^{\wedge}h^{c}_{p,w}$ and ${}^{\vee}h^{c}_{p,w}$ functions work

$$\text{maximize} \;\; \Gamma^{\varphi_1 \wedge \varphi_2 \wedge \ldots \wedge \varphi_5}_{\boldsymbol{c},\boldsymbol{p},\boldsymbol{w}}(x,k) \qquad \text{maximize} \;\; \Gamma^{\varphi_1 \vee \varphi_2 \vee \ldots \vee \varphi_5}_{\boldsymbol{c},\boldsymbol{p},\boldsymbol{w}}(x,k)$$

$$\Gamma^{\varphi_1 \wedge \varphi_2 \wedge \ldots \wedge \varphi_5}_{\boldsymbol{c},\boldsymbol{p},\boldsymbol{w}}(x,k) = {}^{\wedge}h^{c}_{p,w}((f_1(x_k), f_2(x_k), \ldots, f_5(x_k)))$$
$$\Gamma^{\varphi_1 \vee \varphi_2 \vee \ldots \vee \varphi_5}_{\boldsymbol{c},\boldsymbol{p},\boldsymbol{w}}(x,k) = {}^{\vee}h^{c}_{p,w}((f_1(x_k), f_2(x_k), \ldots, f_5(x_k)))$$

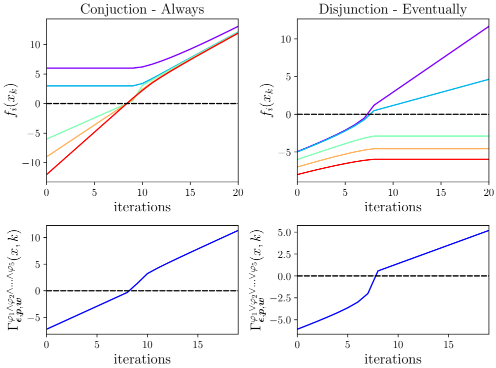

Soundness $\rho^{\varphi}(x,k) \geq 0 \implies (x,k) \models \varphi$

Completeness $(x,k) \models \varphi \implies \rho^{\varphi}(x,k) \geq 0$

Monotonicity ${}^{\wedge}h^{c}_{p,w}$ and ${}^{\vee}h^{c}_{p,w}$ functions are non-decreasing for each of their variables.

Locality & Masking If ${}^{\wedge}h^{c}_{p,w}(y)$ and ${}^{\vee}h^{c}_{p,w}(y)$ are negative for $y = (y_1, y_2, \ldots, y_n)$, then their derivatives with respect to all the y_i variables causing the negative yield are positive.

	Donze et al.	Pant et al.	Gilpin et al.	Mehdipour et al.	Varnai et. al.	D-GMSR
$\mathcal{C}^1$-smoothness	×	✓	✓	×	×	✓
Soundness	✓	○	✓	✓	✓	✓
Completeness	✓	○	○	✓	✓	✓
Monotonicity	✓	✓	✓	✓	×	✓
Locality & Masking	×	△	△	✓	✓	✓

Table: Comparison with the previous robustness measures
(○: Satisfied only for very large smoothing parameters
△: Satisfied only for small smoothing parameters)

Ongoing Works

- Continuous-time modeling of STL specifications and a successive convexification based solution method to ensure the continuous-time satisfaction of STL specifications will be the subject of future studies.

- Future work involves the design of optimization algorithms that exploit the inherent structure of STL specifications to achieve faster convergence.

we now consider our next contribution pertaining to funnel synthesis ...

- funnel synthesis becomes relevant after the guidance development, e.g., including STL specifications ... the key idea is how tight can the nominal specification/mission requirement can actually be followed using feedback

Funnel synthesis

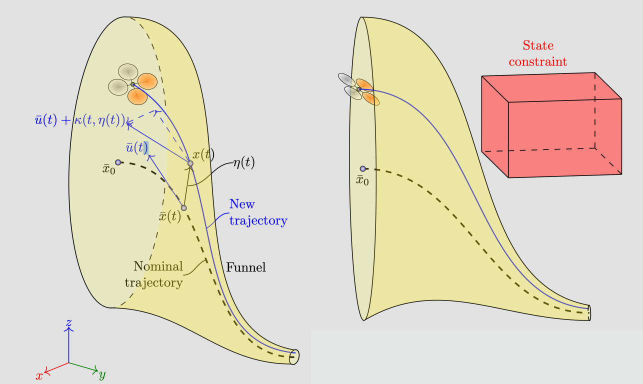

A procedure computing time-varying controlled-invariant set and associated feedback control law such that

$$(x(t), u(t)) \in (\bar{x}(t), \bar{u}(t)) \oplus \mathcal{F}(t) \subset \mathcal{X}(t) \times \mathcal{U}(t)$$

invariance feasibility

$\mathcal{F}$: state and input funnel, $\mathcal{X}, \mathcal{Y}$: feasible state and input space

Example: trajectory generation for the unicycle model with obstacles and input constraints

$$\begin{bmatrix} \dot{r}_x \\ \dot{r}_y \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} u_v \cos\theta \\ u_v \sin\theta \\ u_\theta \end{bmatrix}$$

State (r_x, r_y, θ) : x, y positions and yaw angle
Input (u_v, u_θ) : velocity, angular velocity

- source of uncertainty: nonlinearity modeled by pointwise incremental quadratic inequalities
- objective: maximize the initial funnel
- avoid obstacles
- input constraints: $0 \leq u_v \leq 0, |u_\theta| \leq 2$

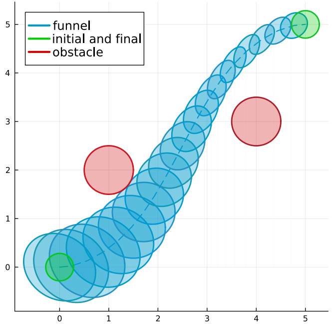

Incremental dynamical system

Define deviation variables

$$\eta := x - \bar{x}, \quad \xi := u - \bar{u}, \quad \delta p := p - \bar{p}, \quad \delta q := q - \bar{q}$$

With linear feedback gain $\xi = K(t)\eta$ and $A_{cl} = A + BK$, $C_{cl} = C + DK$,

Incremental dynamics

$$\begin{aligned}
\dot{\eta}(t) &= f(t, x, u, w, p) - f(t, \bar{x}, \bar{u}, 0, \bar{p}), \\
&= A_{cl}(t)\eta(t) + F(t)w(t) + E\delta p(t), \\
\delta p(t) &= \phi(t, q(t)) - \phi(t, \bar{q}(t)), \\
\delta q(t) &= C_{cl}\eta(t) + Gw(t),
\end{aligned}$$

Incremental dynamics is a finite-horizon uncertain LTV system

- w represents (external) bounded disturbance
- Pair $(\delta q, \delta p)$ represents state-, input-, disturbance-dependent uncertainty/nonlinearity

Uncertainty models

- External disturbance

$$w(t) \in \mathcal{W}(t), \quad \mathcal{W}(t) : \text{compact set}$$

- State/input-dependent uncertainty

$$\langle M(\delta q, \delta p), (\delta q, \delta p) \rangle \geq 0,$$

where M is a bounded operator

This includes uncertainties satisfying

1. pointwise quadratic inequalities (pQI): norm-bounded, polytopic, conic uncertainties
2. pointwise incremental QIs (piQIs): Lipschitz

Funnel Dynamics

Lemma (Lyapunov condition (dissipation inequality))

Consider a time-varying Lyapunov function $V(t,\eta) = \eta^\top Q(t)^{-1}\eta$ *such that*

$$\dot{V}(t) + \alpha V(t) - \lambda(t) w(t)^\top w(t) \leq 0, \quad 0 \leq \lambda(t) \leq \alpha \tag{2}$$

Then, the sublevel set (funnel) $\{\eta \mid \eta^\top Q(t)^{-1}\eta \leq w_{max}^2\}$ *is invariant.*

Theorem (Invariance condition by differential LMI)

Suppose there exist $Q(t), Y(t), Z(t), \nu(t)$ *and* $\lambda(t)$ *such that* $0 < \lambda(t) \leq \alpha$

$$\dot{Q}(t) = Q(t)A(t)^\top + Y(t)^\top B(t)^\top + A(t)Q(t) + B(t)Y(t) + \alpha Q(t) + Z(t),$$

$$\begin{bmatrix} -Z(t) & * & * & * \\ \nu(t)E^\top & -\nu(t)I & * & * \\ F^\top & 0 & -\lambda(t)I & * \\ CQ(t) + DY(t) & 0 & G & -\nu(t)\frac{1}{\gamma^2}I \end{bmatrix} \preceq 0, \forall t \in [t_0, t_f]$$

Then, the Lyapunov condition (2) *holds with* γ*-Lipschitz nonlinearity.*

Recent and ongoing works

- Joint synthesis of trajectory and funnel
  - optimizing feedforward and feedback controllers together
- Uncertainty models by (incremental) integral quadratic constraints
  - general than pointwise (incremental) quadratic constraint
  - time-delay, H_∞ or H_2-norm bounded uncertainties
- Computational methods for solving optimal control over PSD cone
  - preserving positive definiteness of Q
  - guaranteeing continuous-time invariance between node points
- Customized-SDP solvers for computational efficiency and reliability

let us conclude with our work on dynamic policy optimization ...

LQG Direct Policy Optimization using Riemannian Optimization

- Direct bridge between model based design and RL with explicit emphasis on stabilization guarantees (Talebi et al. [2024], Hu et al. [2023])
  - feedback control is directly represented in terms of the policy optimization, e.g.,

$$\min J(K) \quad K \text{ stabilizing}$$

- Offers an interesting twists on control design yet forces revisiting fundemental system theoretic issues
- Offers an approach to go beyond typical performance metrics, clarify the role of robustness, and synthesizing dynamic controllers under structural constraints
- Facilitates developing first order type algorithms for feedback design (for high dimensional systems); when an orcale model for gradient estimates are allowed, one can implement these algorithms efficiency and in a model-free setting, while also providing guarantees

Problem setup

Consider

$$\dot{x}(t) = Ax(t) + Bu(t) + w(t), \; y(t) = Cx(t) + v(t)$$

and a stabilizing, controllable + observable, dynamic output-feedback controller

$$\dot{\xi}(t) = A_K\xi(t) + B_K y(t), \; u(t) = C_K\xi(t)$$

represented as an element $K = (A_K, B_K, C_K) \in \widetilde{\mathcal{C}}_n \subset \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times p} \times \mathbb{R}^{m\times n}$; the space stabilizing full order dynamic output feedback controllers

LQG cost $J : \widetilde{\mathcal{C}}_n \to \mathbb{R}$ is analytic

Minimize J over $\widetilde{\mathcal{C}}_n$ via direct policy gradient

Policy Gradient

Gradient descent (GD):

$$K_{t+1} = K_t - s_t \,\mathrm{grad} J(K_t)$$

issues:

- No local or global theoretical convergence guarantees for LQG case
  - i.e., we are not guaranteed $(K_t) \to K^*$, where K^* is the LQG controller; no stabilization guarantee
- Redundancy due to coordinate transformation: $(A_K, B_K, C_K) \sim (SA_K S^{-1}, SB_K, C_K S^{-1})$, where $S \in \mathbb{R}^{n\times n}$ is invertible
- Often, we have a sub-linear convergence rate
  - i.e., $\lim_{t\to\infty} \frac{\|e_{t+1}\|}{\|e_t\|} = 1$, where $e_t := \|K_t - K^*\|$
- $\widetilde{\mathcal{C}}_n$ is too large; $n^2 + nm + np = \mathcal{O}(n^2)$ dimensions
- $\widetilde{\mathcal{C}}_n$ is not path-connected

Orbit Geometry

since stabilizing full order dynamic feedback controllers $\widetilde{\mathcal{C}}_n$ have a natural quotient geometry, it make sense to instead consider policy optimization over the quotient/orbit space induced by similarity transformation:

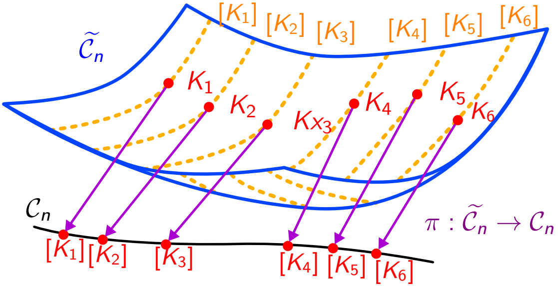

Figure: $\widetilde{\mathcal{C}}_n$ and its orbit space

- orbit of $K = (A_K, B_K, C_K)$ is $[K] = \{(SA_K S^{-1}, SB_K, C_K S^{-1}) : S \in GL(n)\}$
- orbit space is the set of all orbits: $\mathcal{C}_n := \widetilde{\mathcal{C}}_n / GL(n) := \{[K] : K \in \widetilde{\mathcal{C}}_n\}$
- dimension of $\mathcal{C}_n$ is $nm + np = \mathcal{O}(n)$
- minimal controllers: $\mathcal{C}_n^{\min}$ path-connected smooth manifold (locally Euclidean)

Riemannian Gradient Descent

Now comes the issue of choosing the metric: it does make sense to examine $\mathcal{C}_n^{\min}$ (minimal) that also admit a nice coordinate invariant metric introduced by Krishnaprasad-Martin in 1980s! KM metric makes $\widetilde{\mathcal{C}}_n^{\min}$ a Riemannian manifold and by embedding $\mathcal{C}_n^{\min}$ a Riemannian quotient manifold: we now develop a Riemannian gradient descent (RGD) direct policy optimization for LQG:

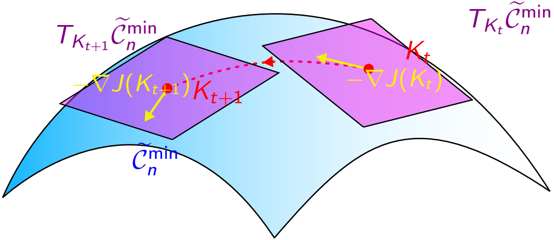

- RGD is an intrinsic 1st-order method for optimizing over minimal stabilizing dynamic feedback controllers
- Although $\mathcal{C}_n^{\min}$ is an orbit space, it isometrically embeds into $\widetilde{\mathcal{C}}_n^{\min}$ via a coordinate-invariant Riemannian metric; this implies RGD over the total manifold $\widetilde{\mathcal{C}}_n^{\min}$ coincides with RGD over the quotient manifold $\mathcal{C}_n^{\min}$

Theoretical results ...

under a non-degenerecy assumption:

- Local convergence guarantee
  - there exists a neighborhood U about K^* such that $K_0 \in U$ implies $(K_t) \to K^*$
  - we are working on extending this to a *global* guarantee
- linear rate of convergence
- much faster than ordinary GD (sub-linear rate)

Numerical observations:

- More robust at escaping saddle points

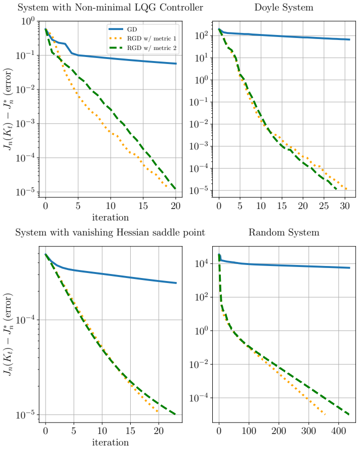

Figure: Comparison of RGD vs. GD for LQG PO for four distinct systems

Next steps and summary

- Apply the PO methodology to structurally constrained dynamic output-feedback controllers/connections to rank constrained LMIs and BMIs
- 2nd-order optimization procedures to generate a superlinear convergence
- deeper understanding of the interplay between th control performance and dynamic stabilization, particularly as it relates to conditioning and algorithmic performance

hence, in a nutshell, we have been examining the role of

- parameterization
- geometry
- algorithms

at distinct levels of the autonomy stack: high order planning, guidance/control interaction, and policy optimization ... this line of work is actively being pursued in our groups to be not only theoretically interesting but also highly relevant for next generation aerospace systems ...

thank you ...

References

Alexandre Donzé and Oded Maler. Robust satisfaction of temporal logic over real-valued signals. In *International Conference on Formal Modeling and Analysis of Timed Systems*, pages 92–106. Springer, 2010. URL `https://doi.org/10.1007/978-3-642-15297-9_9`.

Yash Vardhan Pant, Houssam Abbas, and Rahul Mangharam. Smooth operator: Control using the smooth robustness of temporal logic. In *2017 IEEE Conference on Control Technology and Applications (CCTA)*, pages 1235–1240. IEEE, 2017. URL `https://doi.org/10.1109/CCTA.2017.8062628`.

Yann Gilpin, Vince Kurtz, and Hai Lin. A smooth robustness measure of signal temporal logic for symbolic control. *IEEE Control Systems Letters*, 5(1):241–246, 2020. URL `https://doi.org/10.1109/LCSYS.2020.3001875`.

Noushin Mehdipour, Cristian-Ioan Vasile, and Calin Belta. Arithmetic-geometric mean robustness for control from signal temporal logic specifications. In *2019 American Control Conference (ACC)*, pages 1690–1695. IEEE, 2019. URL `https://doi.org/10.23919/ACC.2019.8814487`.

Peter Varnai and Dimos V Dimarogonas. On robustness metrics for learning STL tasks. In *2020 American Control Conference (ACC)*, pages 5394–5399. IEEE, 2020. URL `https://doi.org/10.23919/ACC45564.2020.9147692`.

Shahriar Talebi, Yang Zheng, Spencer Kraisler, Na Li, and Mehran Mesbahi. Policy optimization in control: Geometry and algorithmic implications. *arXiv preprint arXiv:2406.04243*, 2024.

Bin Hu, Kaiqing Zhang, Na Li, Mehran Mesbahi, Maryam Fazel, and Tamer Başar. Toward a theoretical foundation of policy optimization for learning control policies. *Annual Review of Control, Robotics, and Autonomous Systems*, 6(1):123–158, 2023.

Oliver Sheridan and Behçet Açıkmeşe. Equivalent linear programming formulations for robust trajectory planning under input dependent uncertainties. In *2022 American Control Conference (ACC)*, pages 1873–1878. IEEE, 2022.