Analysis of the Optimization Landscape of Linear Quadratic Gaussian Control

Yujie Tang, Yang Zheng and Na Li

Harvard John A. Paulson School of Engineering and Applied Sciences

UC San Diego
JACOBS SCHOOL OF ENGINEERING
Electrical and Computer Engineering

Linear Quadratic Gaussian Control

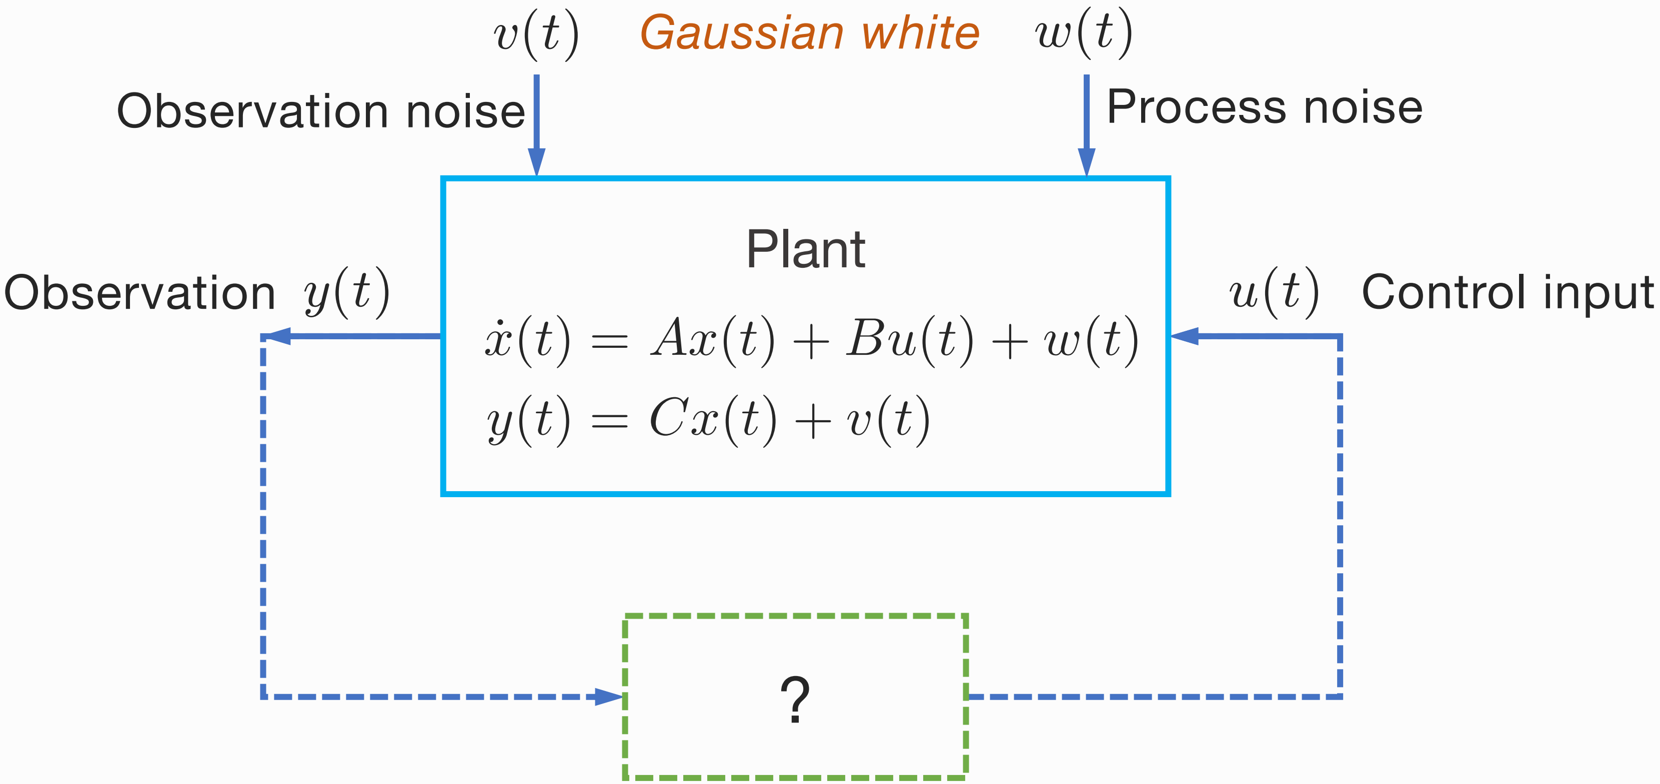

minimize $\displaystyle \lim_{T\to\infty} \frac{1}{T} \mathbb{E} \int_0^T \big(x(t)^\top Q x(t) + u(t)^\top R u(t)\big) dt$

Linear Quadratic Gaussian Control

- A classical control problem, rich theory in classical control
- Allows **partial observation** of the state
  - Perfect state observation is often not available
  - Wider range of applications than LQR

- Existing works on RL for partially observed LQ control mostly focus on **model-based** methods
  - [Tu 2017] [Boczar 2018] [Simchowitz 2020] [Zheng 2021]
- **Model-free** RL for LQG is substantially challenging
  - [Venkataraman 2019]
- Lack of understanding of LQG's **optimization landscape**

Optimization Landscape of LQR

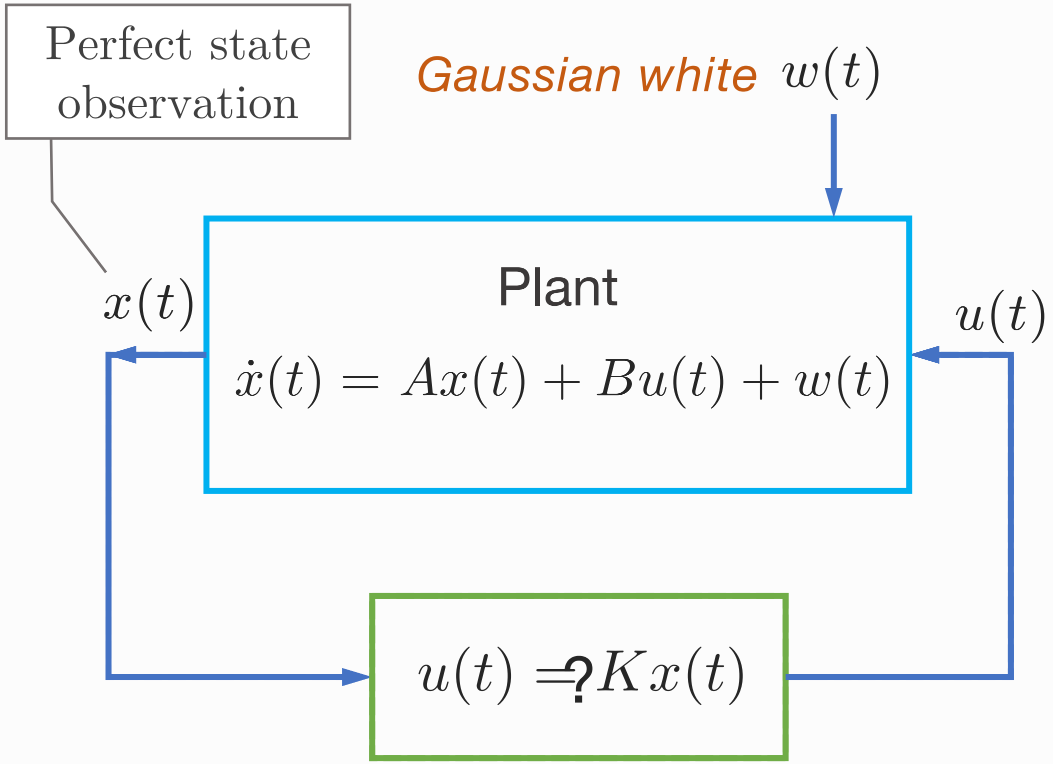

$$\min \quad \lim_{T\to\infty} \frac{1}{T}\mathbb{E}\int_0^T (x^\top Q x + u^\top R u)\,dt$$

Optimization Landscape of LQR

LQR cost

$$\min_{K} \quad f(K)$$
$$\text{s.t.} \quad K \in \mathcal{K}$$

Set of stabilizing feedback gains

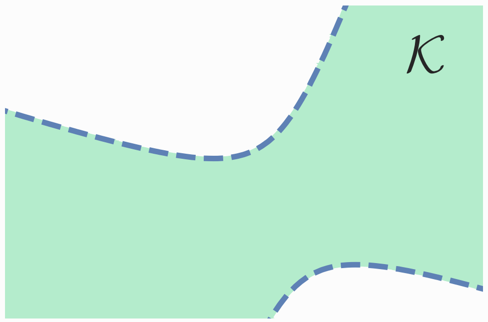

Open, connected, possibly nonconvex

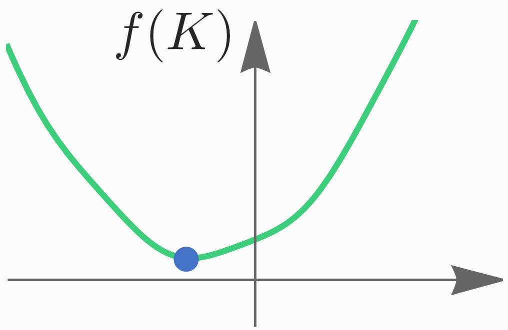

Unique stationary point, coercive, gradient dominance

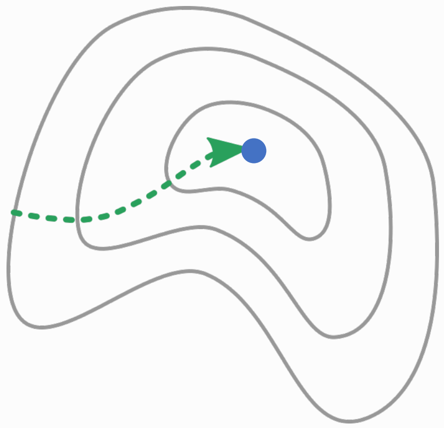

✓ Fast convergence to global optimum for gradient-based methods

[Fazel 2018] [Malik 2019] [Mohammadi 2019] [Bu 2021]

Optimization Landscape of LQG

- Landscape of LQG is fundamental for model-free RL of LQG

- Extension from LQR to LQG is highly nontrivial
  - Classical LQG control theory is more sophisticated
  - Some results of LQR may not hold for LQG anymore
  - The domain consists of **dynamic controllers**, leading to more complex landscape structure

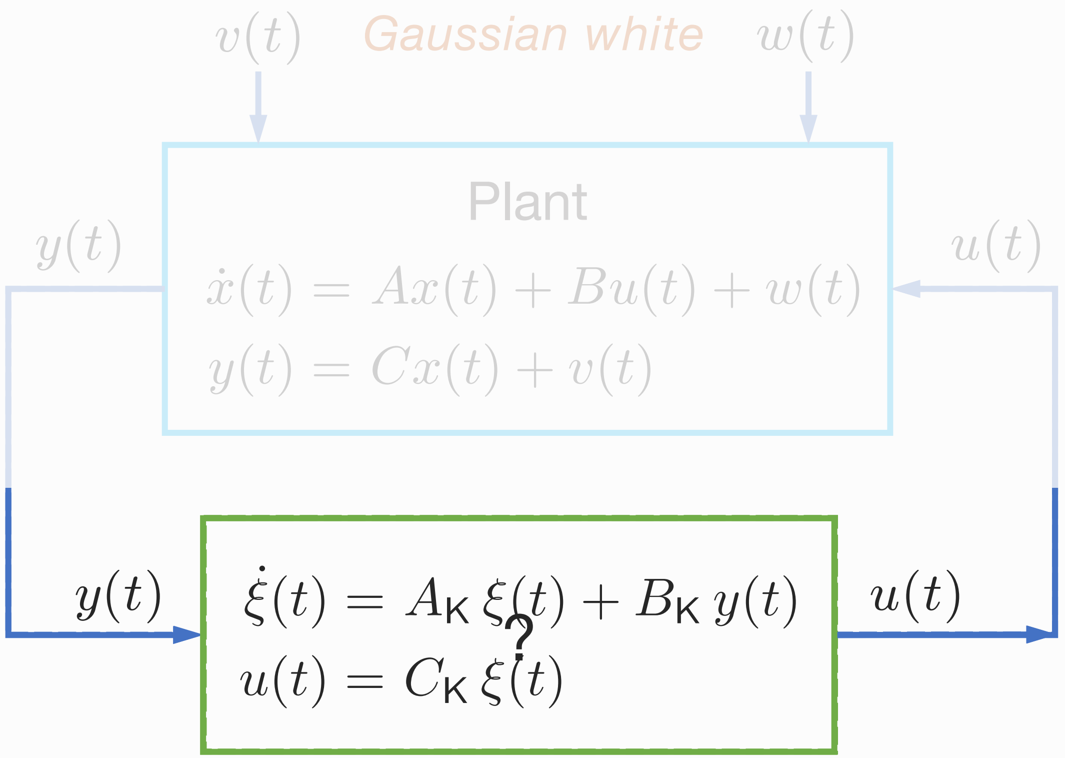

dynamic controller

$\xi(t)$ internal state of the controller

$\dim \xi(t)$ order of the controller

$\dim \xi(t) = \dim x(t)$ full-order

$\dim \xi(t) < \dim x(t)$ reduced-order

minimal controller

The input-output behavior cannot be replicated by a lower order controller.

* (A_K, B_K, C_K) controllable and observable

LQG as an Optimization Problem

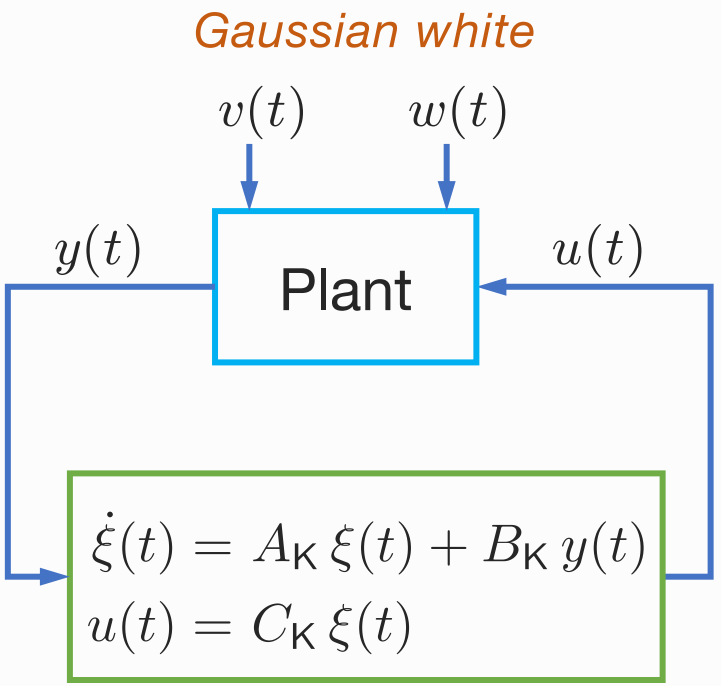

$$\begin{aligned} \min_{\mathsf{K}} \quad & J(\mathsf{K}) \\ \text{s.t.} \quad & \mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \in \mathcal{C}_{\text{full}} \end{aligned}$$

Objective: $J(\mathsf{K})$ The LQG cost

$$\lim_{T\to\infty} \frac{1}{T} \mathbb{E} \int_0^{\infty} (x^\top Q x + u^\top R u)\,dt$$

Domain: $\mathcal{C}_{\text{full}}$ The set of **full-order, stabilizing** dynamic controllers

open, unbounded and nonconvex

LQG as an Optimization Problem

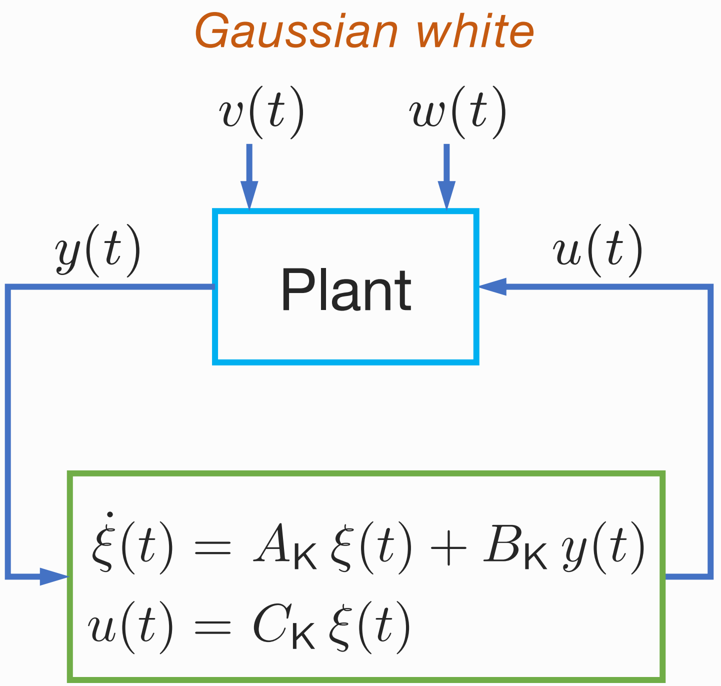

$$\min_{\mathsf{K}} \quad J(\mathsf{K})$$
$$\text{s.t.} \quad \mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \in \mathcal{C}_{\text{full}}$$

- **Connectivity of the domain $\mathcal{C}_{\text{full}}$**
  - Is it connected?
  - If not, how many connected components can it have?
- **Structure of stationary points of $J(\mathsf{K})$**
  - Are there spurious (strictly suboptimal) stationary points?
  - How to check if a stationary point is globally optimal?

Connectivity of the Domain

Theorem 1. Under some standard assumptions,

1) The set $\mathcal{C}_{\text{full}}$ can be disconnected, but has at most 2 connected components.

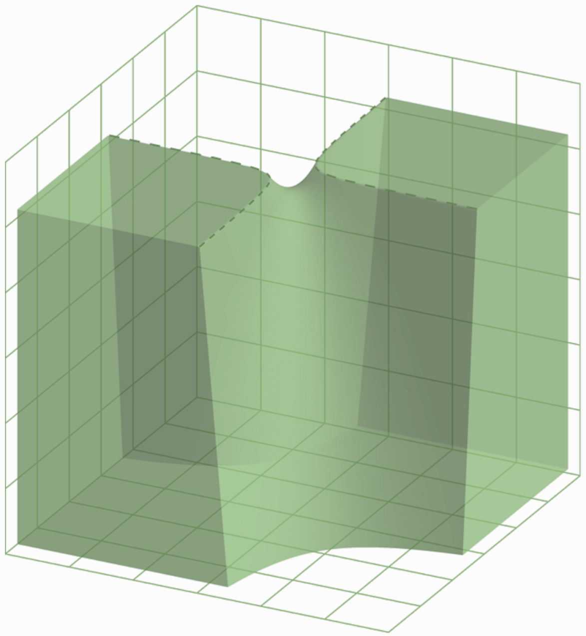

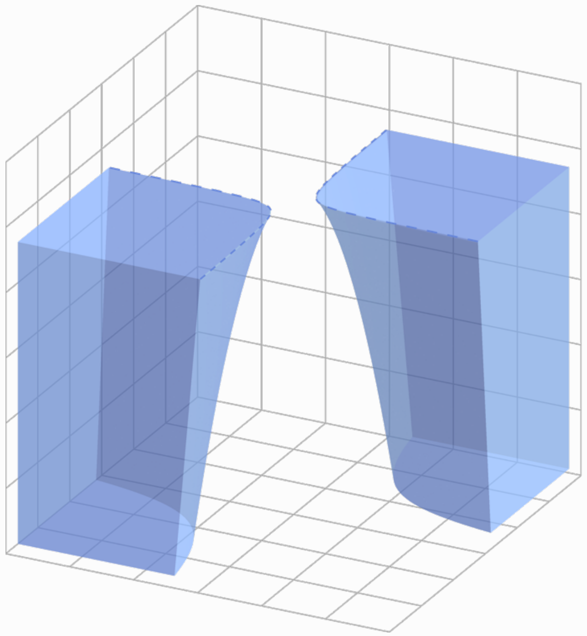

Connectivity of the Domain

Theorem 1. Under some standard assumptions,

1) The set $\mathcal{C}_{\text{full}}$ can be disconnected, but has at most 2 connected components.
2) If $\mathcal{C}_{\text{full}}$ has 2 connected components, then there is a smooth bijection T between the 2 connected components that does not change the value of $J(\mathsf{K})$.

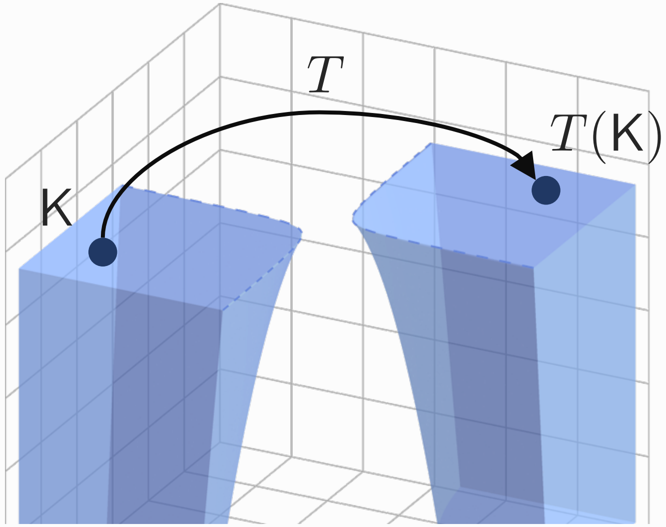

$J(\mathsf{K}) = J(T(\mathsf{K}))$

For gradient-based local search methods, it makes no difference to search over either connected component.

Connectivity of the Domain

Theorem 2. Under some standard assumptions,

1) $\mathcal{C}_{\mathrm{full}}$ is connected if the plant is open-loop stable or there exists a reduced-order stabilizing controller.

2) The sufficient condition of connectivity in 1) becomes necessary if the plant is single-input or single-output.

Example 1.
$$\dot{x}(t) = -x(t) + u(t) + w(t) \qquad x(t) \in \mathbb{R}$$
$$y(t) = x(t) + v(t)$$

- open-loop stable

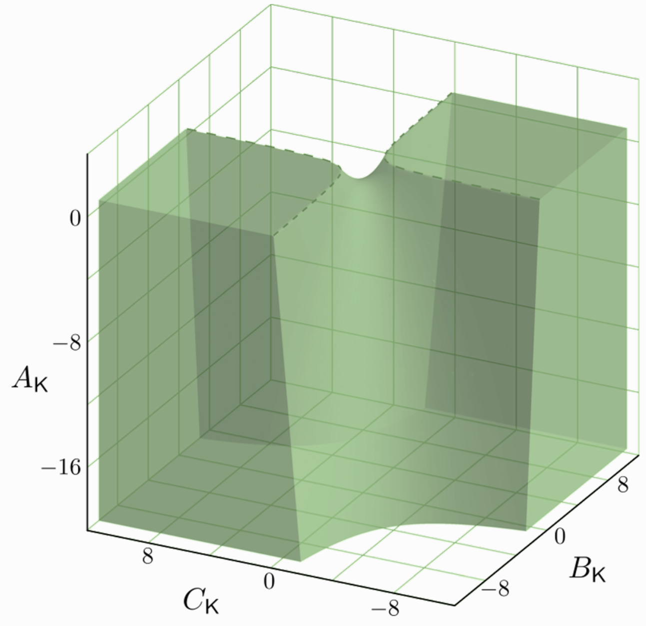

Connectivity of the Domain

Theorem 2. Under some standard assumptions,

1) $\mathcal{C}_{\text{full}}$ is connected if the plant is open-loop stable or there exists a reduced-order stabilizing controller.

2) The sufficient condition of connectivity in 1) becomes necessary if the plant is single-input or single-output.

Example 2. $$\dot{x}(t) = x(t) + u(t) + w(t) \qquad x(t) \in \mathbb{R}$$
$$y(t) = x(t) + v(t)$$

- **not** open-loop stable
- **no** reduced-order stabilizing controller
- single-input single-output

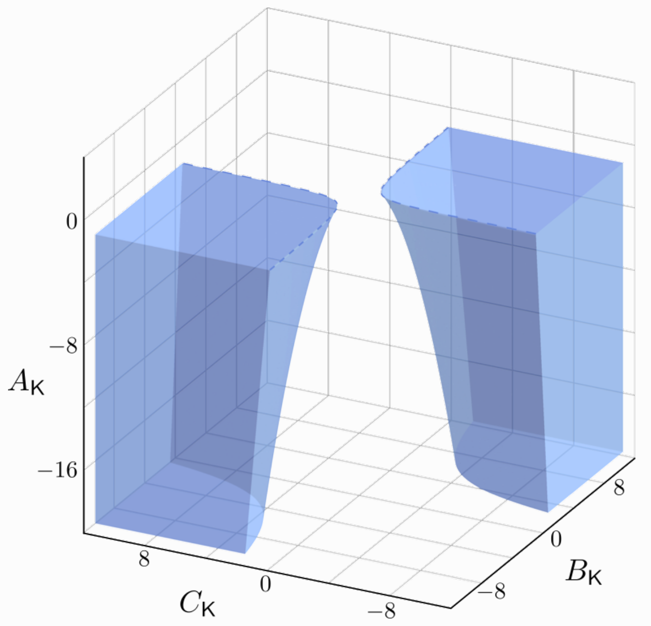

LQG as an Optimization Problem

$$\begin{aligned} \min_{\mathsf{K}} \quad & J(\mathsf{K}) \\ \text{s.t.} \quad & \mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \in \mathcal{C}_{\text{full}} \end{aligned}$$

- **Connectivity of the domain $\mathcal{C}_{\text{full}}$**
  - Is it connected? Not necessarily.
  - If not, how many connected components can it have? Two.
- **Structure of stationary points of $J(\mathsf{K})$**
  - Are there spurious (strictly suboptimal) stationary points?
  - How to check if a stationary point is globally optimal?

LQG as an Optimization Problem

$$\begin{aligned} \min_{\mathsf{K}} \quad & J(\mathsf{K}) \\ \text{s.t.} \quad & \mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \in \mathcal{C}_{\text{full}} \end{aligned}$$

- Connectivity of the domain $\mathcal{C}_{\text{full}}$
  - Is it connected? Not necessarily.
  - If not, how many connected components can it have? Two.
- **Structure of stationary points of $J(\mathsf{K})$**
  - Are there spurious (strictly suboptimal) stationary points?
  - How to check if a stationary point is globally optimal?

Structure of Stationary Points

Facts.

1) $J(\mathsf{K})$ has **non-unique** and **non-isolated** global optima
2) $J(\mathsf{K})$ may have **spurious** stationary points

Contrary to LQR

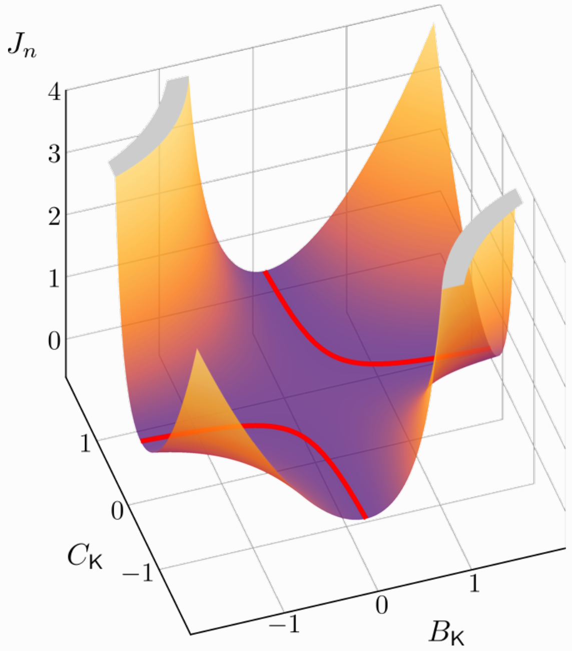

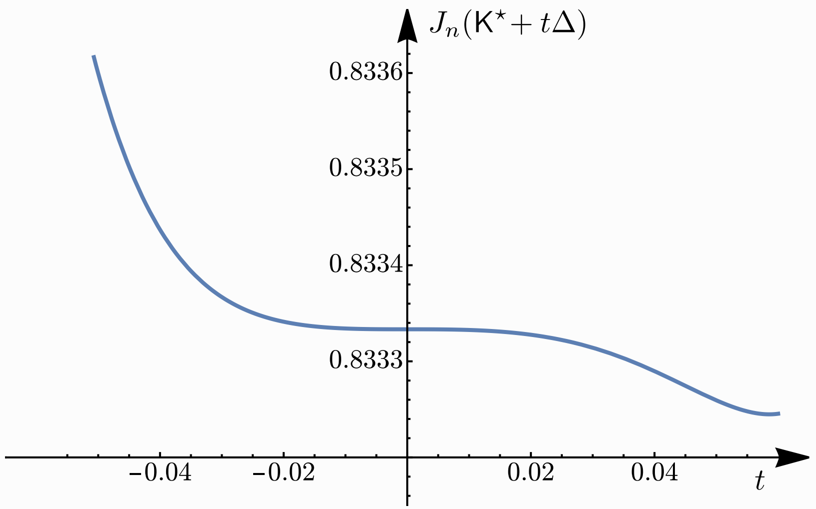

Structure of Stationary Points

Theorem 3. Suppose there exists a stationary point that is a **minimal** controller. Then

1) This stationary point is a global optimum of $J(\mathsf{K})$
2) The set of all global optima forms a manifold with 2 connected components.

Example 1

$$\dot{x}(t) = x(t) + u(t) + w(t)$$
$$y(t) = x(t) + v(t)$$

$$x(t) \in \mathbb{R}$$

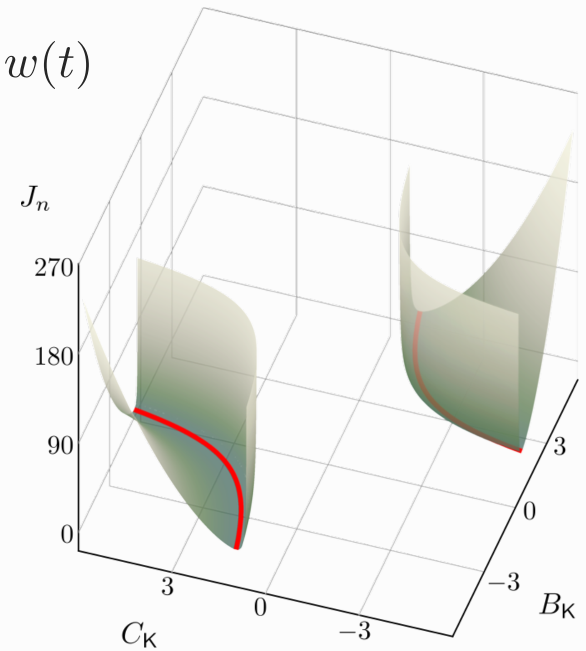

Example 2

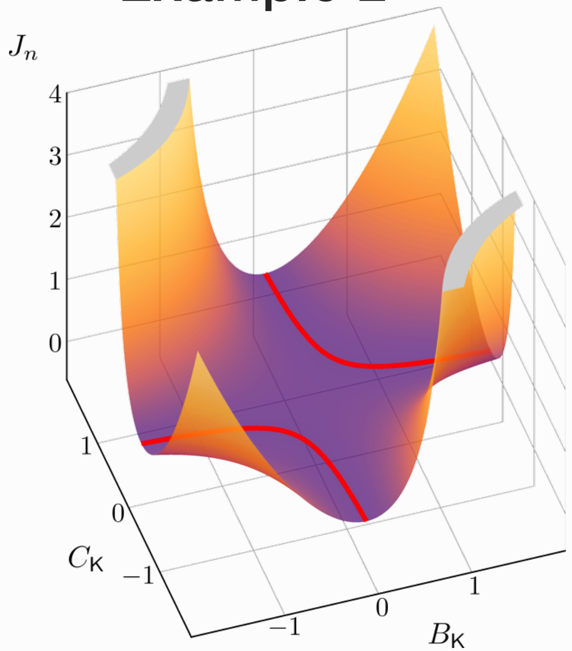

$$\dot{x}(t) = -x(t) + u(t) + w(t)$$
$$y(t) = x(t) + v(t)$$

$$x(t) \in \mathbb{R}$$

Structure of Stationary Points

Implication.

Consider gradient descent iterations

$$\mathsf{K}_{t+1} = \mathsf{K}_t - \alpha \nabla J(\mathsf{K}_t)$$

If the iterates converge to a minimal controller, then this minimal controller is a global optimum.

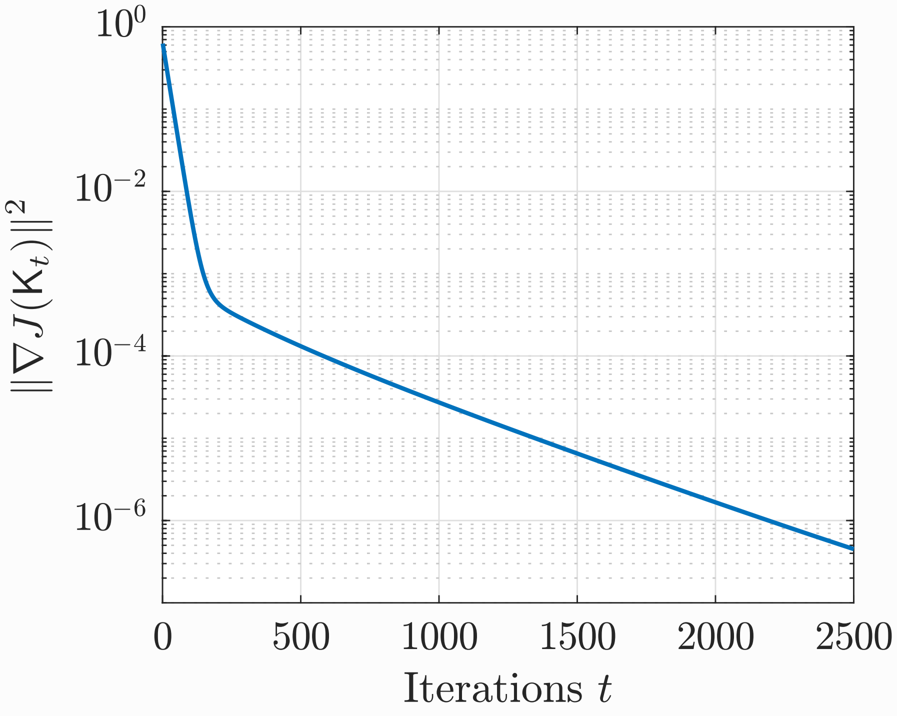

---

* How to check if a controller is minimal?

- Check its controllability and observability.

Summary

LQG as an optimization problem

$$\begin{aligned} \min_{\mathsf{K}} \quad & J(\mathsf{K}) \\ \text{s.t.} \quad & \mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \in \mathcal{C}_{\text{full}} \end{aligned}$$

Connectivity of domain

- At most **two** connected components
- The two connected components mirror each other
- Conditions for being connected

Stationary points

- **Non-unique** global optima, **spurious** stationary points
- **Minimal** stationary points are globally optimal

More results are presented in arXiv:2102.04393.

Future Directions

- A comprehensive classification of stationary points
- Conditions for existence of minimal globally optimal controllers
- Saddle points with vanishing Hessians may exist. How to deal with them?
- Alternative model-free parametrization of dynamic controllers

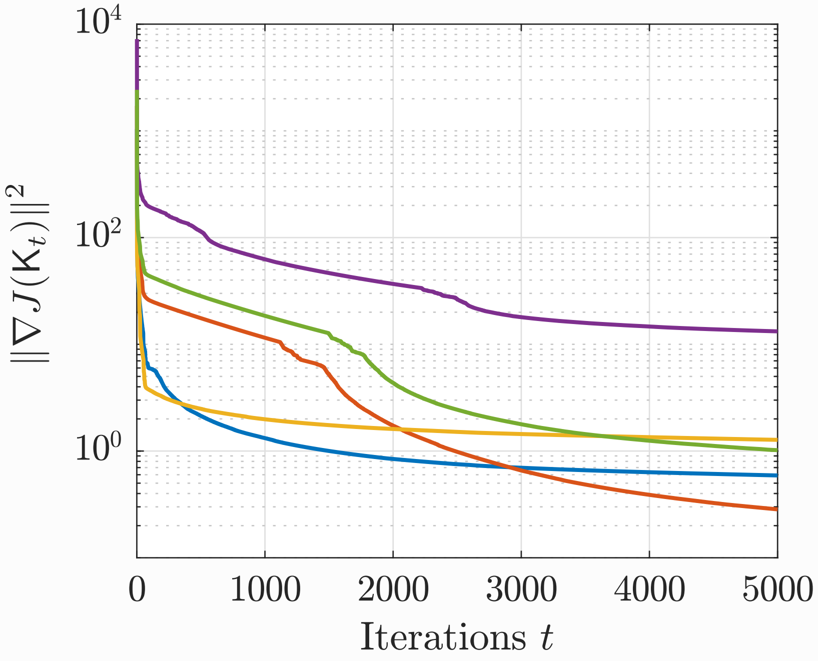

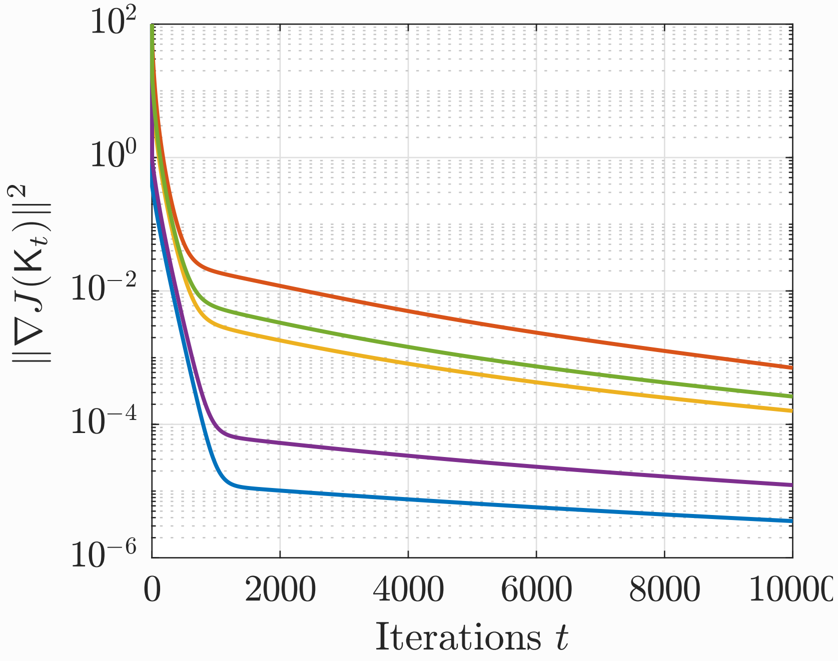

References

Full version of the paper: arXiv:2102.04393

- [Tu 2017] S. Tu, R. Boczar, A. Packard, and B. Recht. Non-asymptotic analysis of robust control from coarse-grained identification. arXiv preprint arXiv:1707.04791, 2017.
- [Boczar 2018] R. Boczar, N. Matni, and B. Recht. Finite-data performance guarantees for the output-feedback control of an unknown system. In Proceedings of the 57th IEEE Conference on Decision and Control, pages 2994–2999, 2018.
- [Simchowitz 2020] M. Simchowitz, K. Singh, and E. Hazan. Improper learning for non-stochastic control. Proceedings of 33rd Conference on Learning Theory, pages 3320-3436, 2020.
- [Zheng 2021] Y. Zheng, L. Furieri, M. Kamgarpour, and N. Li. Sample complexity of linear quadratic gaussian (LQG) control for output feedback systems. In Proceedings of the 3rd Conference on Learning for Dynamics and Control, pages 559-570, 2021.
- [Venkataraman 2019] H. K. Venkataraman and P. J. Seiler. Recovering robustness in model-free reinforcement learning. In 2019 American Control Conference, pages 4210-4216, 2019.
- [Fazel 2018] M. Fazel, R. Ge, S. Kakade, and M. Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In Proceedings of the 35th International Conference on Machine Learning, pages 1467-1476, 2018.
- [Malik 2019] D. Malik, A. Pananjady, K. Bhatia, K. Khamaru, P. Bartlett, M. Wainwright. Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. In Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics, pages 2916-2925, 2019.
- [Mohammadi 2019] H. Mohammadi, A. Zare, M. Soltanolkotabi, and M. R. Jovanović. Convergence and sample complexity of gradient methods for the model-free linear quadratic regulator problem. arXiv preprint arXiv:1912.11899, 2019.
- [Bu 2021] J. Bu, A. Mesbahi, and M. Mesbahi. On topological properties of the set of stabilizing feedback gain. IEEE Transactions on Automatic Control, vol. 66, no. 2, pages 730-744, 2021.